

DISS. ETH NO 28722

**Multimodal Emotion Recognition among Couples from  
Lab Settings to Daily Life using Smartwatches**

A thesis submitted to attain the degree of

DOCTOR OF SCIENCES

(Dr. sc. ETH Zurich)

presented by

GEORGE GYARTEH BOATENG

M.S. in Engineering, Dartmouth College

born on 29.03.1993

Citizen of Ghana

Accepted on the recommendation of

Examiner: Professor Dr. Elgar Fleisch  
Co-Examiners: Professor Dr. Tobias Kowatsch  
Professor David Kotz

2022

*“In God we trust. All others must bring (raw) data.”*

W. Edwards Deming

## Summary

Couples generally manage chronic diseases together and the management takes an emotional toll on both patients and their romantic partners. Consequently, recognizing the emotions of each partner in daily life could provide an insight into their emotional well-being in chronic disease management. The emotions of partners are currently inferred in the lab and daily life using self-reports which are not practical for continuous emotion assessment or observer reports which are manual, time-intensive, and costly. Currently, there exists no comprehensive overview of works on emotion recognition among couples. Furthermore, approaches for emotion recognition among couples have (1) focused on English-speaking couples in the U.S., (2) used data collected from the lab, and (3) performed recognition using observer ratings rather than partner’s self-reported / subjective emotions.

In this body of work contained in this thesis (8 papers — 5 published and 3 currently under review in various journals), we fill the current literature gap on couples’ emotion recognition, develop emotion recognition systems using 161 hours of data from a total of 1,051 individuals, and make contributions towards taking couples’ emotion recognition from the lab which is the status quo, to daily life. First, we provided a comprehensive survey of the research field of emotion recognition among couples (Paper 1). Second, we leveraged insights from psychology research and deep transfer learning approaches to develop machine learning systems to recognize each partner’s emotions using lab data from Dutch-speaking couples in Belgium (Paper 2) and German-speaking couples in Switzerland (Papers 3 and 4). We also performed emotion recognition using data from German-speaking elderly individuals (not romantic partners) in Germany (Paper 5) given the target use case for our emotion recognition system consisted of partners who were elderly and spoke German. Third, we developed ubiquitous smartwatch and smartphone systems — VADLite and DyMand — to collect relevant multimodal sensor data and self-report emotion data from the daily life interactions of German-speaking, Swiss-based couples managing type 2 diabetes (Papers 6 and 7). Finally, we developed and evaluated machine learning systems for recognizing each partner’s emotions using the collected multimodal real-world smartwatch data — heart rate, accelerometer, gyroscope, and speech (Paper 8).

This thesis contributes toward building automated emotion recognition systems that would eventually enable partners to monitor their emotions in daily life and enable the delivery of interventions to improve their emotional well-being.

## Résumé

Les couples gèrent généralement ensemble les maladies chroniques et la gestion a un impact émotionnel sur les patients et leurs partenaires amoureux. Par conséquent, la reconnaissance des émotions de chaque partenaire dans la vie quotidienne pourrait donner un aperçu de leur bien-être émotionnel dans la gestion des maladies chroniques. Les émotions des partenaires sont actuellement déduites en laboratoire et dans la vie quotidienne à l'aide d'auto-rapports qui ne sont pas pratiques pour l'évaluation continue des émotions ou de rapports d'observateurs qui sont manuels, chronophages et coûteux. Actuellement, il n'existe pas de synthèse exhaustive des travaux sur la reconnaissance des émotions chez les couples. De plus, les approches de reconnaissance des émotions chez les couples se sont (1) concentrées sur les couples anglophones aux États-Unis, (2) ont utilisé des données recueillies auprès du laboratoire et (3) ont effectué une reconnaissance en utilisant les évaluations des observateurs plutôt que les émotions autodéclarées / subjectives du partenaire. .

Dans cet ensemble de travaux contenus dans cette thèse (8 articles — 5 publiés et 3 actuellement en cours de révision dans diverses revues), nous comblons le vide bibliographique actuel sur la reconnaissance des émotions des couples, développons des systèmes de reconnaissance des émotions en utilisant 161 heures de données provenant d'un total de 1,051 personnes et contribuent à faire passer la reconnaissance des émotions des couples du laboratoire, qui est le statu quo, à la vie quotidienne. Tout d'abord, nous avons fourni une enquête complète sur le domaine de recherche de la reconnaissance des émotions chez les couples (article 1). Deuxièmement, nous avons tiré parti des connaissances de la recherche en psychologie et des approches “deep transfer learning” pour développer des systèmes d'apprentissage automatique permettant de reconnaître les émotions de chaque partenaire à l'aide de données de laboratoire de couples néerlandophones en Belgique (article 2) et de couples germanophones en Suisse (articles 3 et 4). Nous avons également effectué la reconnaissance des émotions à l'aide de données provenant de personnes âgées germanophones (pas de partenaires romantiques) en Allemagne (article 5), étant donné que le cas d'utilisation cible de notre système de reconnaissance des émotions était composé de partenaires âgés et parlant allemand. Troisièmement, nous avons développé des systèmes ubiquitaires de montres intelligentes et de smartphones — VADLite et DyMand — pour collecter des données de capteurs multimodaux pertinentes et auto-déclarer des données émotionnelles à partir des interactions quotidiennes de couples germanophones basés en Suisse qui gèrent le diabète de type 2 (articles 6 et 7) . Enfin, nous avons développé et évalué des systèmes d'apprentissage automatique pour reconnaître les émotions de chaque partenaire à l'aide des données de smartwatch multimodales du monde réel collectées - fréquence cardiaque, accéléromètre, gyroscope et parole (article 8).

Cette thèse contribue à la construction de systèmes automatisés de reconnaissance des émotions qui permettraient éventuellement aux partenaires de surveiller leurs émotions dans la vie quotidienne et de permettre la réalisation d'interventions pour améliorer leur bien-être émotionnel.

## Acknowledgement

First, I would like to express my gratitude to Prof. Dr. Elgar Fleisch for the opportunity to do my PhD in the Center for Digital Health Interventions. The ethos of doing research that is both technically rigorous and societally relevant will always stick with me. The continuous asking of what our technical results mean in the real-world has influenced my drive to keep working on answering research questions that make a difference in the lives of people. I thank Prof. Dr. Tobias Kowatch for his ever-available support, feedback and in particular, his infectious enthusiasm and optimism about the projects in our center.

I'm grateful to Prabhakaran Santhanam, my co-laborer in the DyMand project for all his assistance and support. I thank my DyMand project collaborators at the University of Zurich for the opportunity to work together: Professor Urte Scholz, Professor Guy Bodenmann, Dr. Janina Luscher and Dr. Theresa Pauly.

I thank various postdocs and professors for their feedback on my work at various stages of my PhD: Prof. Cecilia Mascolo, Prof. David Kotz, Prof. Dr. Elliott Ash, Prof. Emily Provost, Prof. Dr. Felix Wortmann, Prof. Dr. Florian Wangenheim, Prof. Dr. Peter Hilpert, Prof. Dr. Petra Schmid, Dr. Laura Sels, Prof. Dr. Stefan Feuerriegel, Prof. Temiloluwa Prioleau, Prof. Dr. Verena Tiefenbeck

I'm grateful to several master's students I've been privileged to work with who contributed to these works: Madhav Sachdeva, Malgorzata Speichert, Marine Hoche, Jacopo Biggiogera, Xiangyu Zhao. I thank the following research assistants for assisting with data collection, coding, annotating and transcribing our dataset: Kemeng Zhang, Kwabena Atobra, Elena Luzi, Denis Adamec, and Luljeta Isaki.

I'm grateful to all the individuals who gave feedback on my work in conferences, research visits, or paper reviews and in particular, my colleagues in our research group for the feedback and contributions along these years: Dr. Anselma Wörner, Caterina Bérubé, Davide Clares, Dr. Dominik Rüeegg, Eva van Weenen, Dr. Filipe Barata, Fabian Schneider, Dr. Florian Künzler, Gisbert Teepe, Dr. Iris Shih, Dr. Jan-Niklas Kramer, Kevin Koch, Dr. Klaus Fuchs, Dr. Liliane Ableitner, Mara Nägeline, Martin Maritsch, Dr. Peter Tinschert, Raphael Weibel, Robert Jakob, Simon Föll, Dr. Shu Liu, Yanick Lukic.

I thank Monica Heinz, Elisabeth Keller, Judith Holzheimer, Olivia Keller for assisting with the administrative aspects of doing PhD at ETH.

I thank my family and friends for their support throughout this journey.

Last but not least, I thank God for His grace which has been sufficient through this journey, and the strength to trust that His plans and purpose for me are always good, through it all. To Him be the glory.

## Contents

Summary	ii
Résumé	iii
Acknowledgement	iv
Contents	v
Introduction	1
References	14
Paper 1: Survey of Couples' Emotion Recognition	15
Paper 2: Peak-End Rule	38
Paper 3: BERT meets LIWC	48
Paper 4: You made me feel this way	57
Paper 5: Elderly Emotion Recognition	66
Paper 6: VADLite	75
Paper 7: DyMand	85
Paper 8: Are you ok, honey?	111

## Introduction

Romantic partners' emotions have links with the quality of their relationship, and disease management when one partner has a chronic disease. Couples heading for break-up show more negative emotions and less positive emotions than happy couples [12, 16]. And for couples in which one partner has a chronic disease such as cancer and diabetes, its joint management called dyadic coping [2, 11, 21] takes an emotional toll on both patients and spouses [26]. Consequently, recognizing the emotions of each partner could unobtrusively provide an insight into the quality of their relationship, and the emotional well-being of each partner, and enable the triggering of interventions to improve their relationship and chronic disease management.

There are mainly two models of emotions: categorical and dimensional. Categorical emotions are based on the six basic emotions proposed by Ekman: happiness, sadness, fear, anger, disgust, and surprise [15]. Dimensional approaches mainly use two dimensions: valence (pleasure) and arousal which are based on Russell's circumplex model of emotions [24]. Valence refers to how negative to positive a person feels and arousal refers to how sleepy to active a person feels. Using these two dimensions, several categorical emotions can be placed and grouped into the four quadrants: high arousal and negative valence (e.g., angry), low arousal and negative valence (e.g., depressed), low arousal and positive valence (e.g., relaxed), and high arousal and positive valence (e.g., excited) [24].

Currently, the emotions of romantic partners are inferred in the lab and in daily life using two approaches: self-report and observer reports [4]. In the lab, couples are asked to have an emotionally charged conversation that is videotaped. Afterward, each partner provides emotion ratings, for example, by adjusting a joystick while watching the videos [23] or by using a validated affect instrument such as the Multidimensional Mood questionnaire [27]. Also, people are trained to watch the video recordings and use a coding scheme (e.g., SPAFF [14]) to rate the emotional behavior of each partner. In the case of daily life, couples are periodically asked to complete self-reports [25] such as the PANAS [28] or observers can code audio data collected from couples' daily life interactions [22]. Collecting self-reports in daily life are obtrusive and impractical for continuous emotion assessment (e.g., every minute). Furthermore, self-reports could be biased (for example, if the partner desires to project a certain emotion rating rather than how they really feel) and may not reflect the partner's true emotion. The manual coding process with observer reports is costly and time-consuming as multiple coders need to be trained for this task [18] and suffers from inter-rater reliability issues [17, 19]. Additionally, observer reports do not reflect the subjective emotions of partners, but rather, an external person's assessment based on expressed behavior. Automated recognition of each partner's emotion could potentially address these limitations by leveraging sensor data (e.g., audio).

Emotion recognition among couples is the task of automatically recognizing the emotions of romantic partners based on their conversation or interaction

context [4]. In particular, it entails recognizing each partner’s emotions for every utterance/speaker turn or every few seconds — local emotion — or for the whole conversation — global emotion. These emotion ratings are the self-reports provided by the partners or observer ratings provided by external individuals. This task differs from other kinds of emotion recognition tasks mainly by the kind of stimuli that induces emotions. Some stimuli are driving [29], listening to music or watching a movie [1], and conversation between people [20]. Couples’ emotion recognition is similar to emotion recognition tasks whose stimuli are conversations since it uses a conversational context. However, its uniqueness lies in the fact that the two interacting individuals are in a romantic relationship. Consequently, various insights from psychology about couples’ interaction dynamics could be leveraged to recognize each partner’s emotions. For example, romantic partners influence each other when interacting, and that insight has been used for couples’ emotion recognition (e.g., [13, 7]). Currently, there exists no comprehensive overview of works on emotion recognition among couples. Furthermore, approaches for emotion recognition among couples have (1) focused on English-speaking couples in the U.S., (2) used data collected from the lab, and (3) performed recognition using observer ratings rather than partner’s self-reported / subjective emotions.

In this body of work contained in this thesis (8 papers), we fill the current literature gap on couples’ emotion recognition, develop emotion recognition systems using 161 hours of data from a total of 1,051 individuals, and make contributions towards taking couples’ emotion recognition from the lab which is the status quo, to daily life. First, we provided a comprehensive survey of the research field of emotion recognition among couples (Paper 1 [4]). Second, we leveraged insights from psychology research and deep transfer learning approaches to develop machine learning systems to recognize each partner’s emotions using lab data from Dutch-speaking couples in Belgium (Paper 2 [9]) and German-speaking couples in Switzerland (Paper 3 [3] and Paper 4 [7]). We also performed emotion recognition using data from German-speaking elderly individuals (not romantic partners) in Germany (Paper 5 [5]) given the target use case for our emotion recognition system consisted of partners who were elderly and spoke German. Third, we developed ubiquitous smartwatch and smartphone systems — VADLite and DyMand — to collect relevant multimodal sensor data and self-report emotion data from the daily life interactions of German-speaking, Swiss-based couples managing type 2 diabetes (Paper 6 [10] and Paper 7 [8]). Finally, we developed and evaluated machine learning systems for recognizing each partner’s emotions using the collected multimodal real-world smartwatch data — heart rate, accelerometer, gyroscope, and speech (Paper 8 [6]). Of the 8 papers included in this thesis, 5 have been published and 3 are currently under review in various journals. The rest of the thesis contains an overview of the 8 papers with the reference of each paper provided, an abstract of the papers, their scientific contributions, my specific contribution to each of the works, and the full text of all 8 papers.



## Paper 1: Survey of Couples' Emotion Recognition

### Reference

**George Boateng**, Elgar Fleisch, Tobias Kowatch. "Emotion Recognition among Couples: A Survey". *ACM Computing Surveys* (*under review*). <https://arxiv.org/abs/2202.08430>

### Abstract

Couples' relationships affect the physical health and emotional well-being of partners. Automatically recognizing each partner's emotions could give a better understanding of their individual emotional well-being, enable interventions and provide clinical benefits. In the paper, we summarize and synthesize works that have focused on developing and evaluating systems to automatically recognize the emotions of each partner based on couples' interaction or conversation contexts. We identified 28 articles from IEEE, ACM, Web of Science, and Google Scholar that were published between 2010 and 2021. We detail the datasets, features, algorithms, evaluation, and results of each work as well as present main themes. We also discuss current challenges, research gaps and propose future research directions. In summary, most works have used audio data collected from the lab with annotations done by external experts and used supervised machine learning approaches for binary classification of positive and negative affect. Performance results leave room for improvement with significant research gaps such as no recognition using data from daily life. This survey will enable new researchers to get an overview of this field and eventually enable the development of emotion recognition systems to inform interventions to improve the emotional well-being of couples.

### Contribution

Our contributions are as follows: (1) the first survey and comprehensive overview of the field of emotion recognition among couples which would enable new researchers to get started in the field easily (2) background on emotion models, elicitation, and annotation approaches (3) details of the datasets, features, modalities, algorithms, evaluation, and results of each work (4) modeling approaches that consider the unique context of couples' interactions (5) discussion of current challenges, research gaps and proposal of future research directions. For this paper, I performed the search of key terms in various databases, selected papers based on the inclusion and exclusion criteria, read and wrote summaries of the selected papers, wrote the first draft of this paper, and then the final draft after receiving feedback.

## Paper 2: Peak-End Rule

### Reference

**George Boateng**, Laura Sels, Peter Kuppens, Peter Hilpert, Urte Scholz, and Tobias Kowatsch. 2020. “Speech Emotion Recognition among Couples using the Peak-End Rule and Transfer Learning”. In *Companion Publication of the 2020 International Conference on Multimodal Interaction (ICMI '20 Companion)*, October 25–29, 2020, Virtual event, Netherlands. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3395035.3425253>

### Abstract

Extensive couples’ literature shows that how couples feel after a conflict is predicted by certain emotional aspects of that conversation. Understanding the emotions of couples leads to a better understanding of partners’ mental well-being and consequently their relationships. Hence, automatic emotion recognition among couples could potentially guide interventions to help couples improve their emotional well-being and their relationships. It has been shown that people’s global emotional judgment after an experience is strongly influenced by the emotional extremes and ending of that experience, known as the peak-end rule. In this work, we leveraged this theory and used machine learning to investigate which audio segments can be used to best predict the end-of-conversation emotions of couples. We used speech data collected from 101 Dutch-speaking couples in Belgium who engaged in 10-minute long conversations in the lab. We extracted acoustic features from (1) the audio segments with the most extreme positive and negative ratings — the peak, and (2) the ending of the audio — the end. We used transfer learning in which we extracted these acoustic features with a pre-trained convolutional neural network (YAMNet). We then used these features to train machine learning models — support vector machines — to predict the end-of-conversation valence ratings (positive vs negative) of each partner. From our results (balanced accuracy), the segments from the peak were the best for recognizing the emotions of female partners and outperformed male partners’ perception of their female partners’ emotions. The results of this work could inform how to best recognize the emotions of couples after-conversation sessions and eventually, lead to a better understanding of couples’ relationships either in therapy or in everyday life.

### Contribution

Our contributions are as follows: (1) exploration of the best way to recognize the emotions of couples after a conversation (5 - 10 minutes) through the peak-end rule lens using deep transfer learning — classification of end-of-conversation valence using acoustic features from the emotional peaks and end of the audio (2) use of a unique dataset — real-world data collected from Dutch-speaking couples with self-ratings of emotions (3) proposal and computation of a “partner perception baseline” for emotion recognition within the context of couples

interactions that leverage each partner’s perception of their partner’s emotions and gives an estimate of how well each partner could infer their partner’s emotion. For this paper, I conceptualized the key idea, preprocessed the dataset, extracted features, implemented the machine learning experiments, wrote the first draft of the paper, and then the final draft after receiving feedback.

### **Paper 3: BERT meets LIWC**

#### **Reference**

Jacopo Biggioera, **George Boateng**, Peter Hilpert, Matthew Vowels, Guy Bodenmann, Mona Neysari, Fridtjof Nussbeck, Tobias Kowatsch. “BERT meets LIWC: Exploring State-of-the-Art Language Models for Predicting Communication Behavior in Couples’ Conflict Interactions”. In *Companion Publication of the 2021 International Conference on Multimodal Interaction (ICMI ’21 Companion)*, October 18–22, 2021, Montréal, QC, Canada. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3461615.3485423>

#### **Abstract**

Many processes in psychology are complex, such as dyadic interactions between two interacting partners (e.g., patient-therapist, intimate relationship partners). Nevertheless, many basic questions about interactions are difficult to investigate because dyadic processes can be within a person and between partners, they are based on multimodal aspects of behavior and unfold rapidly. Current analyses are mainly based on the behavioral coding method, whereby human coders annotate behavior based on a coding schema. But coding is labor-intensive, expensive, slow, focuses on few modalities, and produces sparse data which has forced the field to use average behaviors across entire interactions, thereby undermining the ability to study processes on a fine-grained scale. Current approaches in psychology use LIWC for analyzing couples’ interactions. However, advances in natural language processing such as BERT could enable the development of systems to potentially automate behavioral coding, which in turn could substantially improve psychological research. In this work, we train machine learning models to automatically predict positive and negative communication behavioral codes of 368 German-speaking Swiss couples during an 8-minute conflict interaction on a fine-grained scale (10-seconds sequences) using linguistic features and paralinguistic features derived with openSMILE. Our results show that both simpler TF-IDF features as well as more complex BERT features performed better than LIWC, and that adding paralinguistic features did not improve the performance. These results suggest it might be time to consider modern alternatives to LIWC, the de facto linguistic features in psychology, for prediction tasks in couples research. This work is a further step towards the automated coding of couples’ behavior which could enhance couple research and therapy, and be utilized for other dyadic interactions as well.

## Contribution

Our contributions are as follows: (1) an evaluation of the predictive capability of BERT vis-à-vis LIWC in the context of the automatic recognition of couples' communication behavioral codes on a fine-grained time scale (every 10 seconds) (2) an investigation into how the addition of paralinguistic features affects prediction performance (3) the use of a unique dataset — spontaneous, naturalistic, speech data collected from German-speaking Swiss couples (n=368 couples, N=736 participants), and the largest ever such dataset used in the literature for automatic coding of couples' behavior. For this paper, I co-conceptualized the key idea, extracted paralinguistic and linguistic (BERT) features, provided feedback on the machine learning experiments, and co-wrote the first and final drafts of the paper.

## Paper 3: You made me feel this way

### Reference

**George Boateng**, Peter Hilpert, Guy Bodenmann, Mona Neysari, Tobias Kowatsch. ““You made me feel this way”: Investigating Partners' Influence in Predicting Emotions in Couples' Conflict Interactions using Speech Data”. In *Companion Publication of the 2021 International Conference on Multimodal Interaction (ICMI '21 Companion)*, October 18–22, 2021, Montréal, QC, Canada. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3461615.3485424>

### Abstract

How romantic partners interact with each other during a conflict influences how they feel at the end of the interaction and is predictive of whether the partners stay together in the long term. Hence understanding the emotions of each partner is important. Yet current approaches that are used include self-reports which are burdensome and hence limit the frequency of this data collection. Automatic emotion prediction could address this challenge. Insights from psychology research indicate that partners' behaviors influence each other's emotions in conflict interaction and hence, the behavior of both partners could be considered to better predict each partner's emotion. However, it is yet to be investigated how doing so compares to only using each partner's own behavior in terms of emotion prediction performance. In this work, we used BERT to extract linguistic features (i.e., what partners said) and openSMILE to extract paralinguistic features (i.e., how they said it) from a data set of 368 German-speaking Swiss couples (N = 736 individuals) who were videotaped during an 8-minutes conflict interaction in the laboratory. Based on those features, we trained machine learning models to predict if partners feel positive or negative after the conflict interaction. Our results show that including the behavior of the other partner improves the prediction performance. Furthermore, for men, considering how their female partners spoke is most important and for women

considering what their male partner said is most important in getting better prediction performance. This work is a step towards automatically recognizing each partner’s emotion based on the behavior of both, which would enable a better understanding of couples in research, therapy, and the real world.

### **Contribution**

This work builds upon Paper 2 by using (1) the same data set, (2) global emotion ratings rather than local emotion ratings, (3) self-reports rather than observer reports, and (4) BERT which was shown to perform well for linguistic feature extraction in Paper 2. Our contributions are as follows: (1) an evaluation of how well a partner’s own linguistic and paralinguistic features predict one’s own end-of-conversation emotion (2) an investigation of how the prediction performance changes when including one’s partner’s features (linguistic, paralinguistic, and both) (3) the use of a unique dataset — spontaneous, naturalistic, speech data collected from German-speaking, Swiss couples (n=368 couples, N=736 participants), which is the largest ever such dataset used in the literature for automatic recognition of partners’ end-of-conversation emotion. For this paper, I co-conceptualized the key idea, preprocessed the dataset, extracted features, implemented the machine learning experiments, and co-wrote the first and final drafts of the paper.

## **Paper 5: Elderly Emotion Recognition**

### **Reference**

**George Boateng** and Tobias Kowatsch. 2020. “Speech Emotion Recognition among Elderly Individuals using Multimodal Fusion and Transfer Learning”. In *Companion Publication of the 2020 International Conference on Multimodal Interaction (ICMI '20 Companion), October 25–29, 2020, Virtual event, Netherlands*. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3395035.3425255>

### **Abstract**

Recognizing the emotions of the elderly is important as it could give an insight into their mental health. Emotion recognition systems that work well on the elderly could be used to assess their emotions in places such as nursing homes and could inform the development of various activities and interventions to improve their mental health. However, several emotion recognition systems are developed using data from younger adults. In this work, we trained machine learning models to recognize the emotions of elderly individuals via performing a 3-class classification of valence and arousal as part of the INTERSPEECH 2020 Computational Paralinguistics Challenge (COMPARE). We used speech data from 87 participants who gave spontaneous personal narratives. We leveraged a transfer learning approach in which we used pretrained CNN and BERT models to extract acoustic and linguistic features respectively and fed them into separate

machine learning models. Also, we fused these two modalities in a multimodal approach. Our best model used a linguistic approach and outperformed the official competition of unweighted average recall (UAR) baseline for valence by 8.8% and the mean of valence and arousal by 3.2%. We also showed that feature engineering is not necessary as transfer learning without fine-tuning performs as well or better and could be leveraged for the task of recognizing the emotions of elderly individuals. This work is a step towards better recognition of the emotions of the elderly which could eventually inform the development of interventions to manage their mental health.

### Contribution

Our contribution is the evaluation of transfer learning approaches to recognize the emotions of elderly individuals using a novel dataset — speech data collected from German-speaking elderly individuals via using a pretrained CNN model (YAMNET) to extract acoustic features and a pretrained Transformer language model — Bidirectional Encoder Representations from Transformers (BERT) to extract linguistic features. For this paper, I conceptualized the key idea, preprocessed the dataset, extracted features, implemented the machine learning experiments, wrote the first draft of the paper, and then the final draft after receiving feedback.

## Paper 6: VADLite

### Reference

**George Boateng**, Prabhakaran Santhanam, Janina Lüscher, Urte Scholz, and Tobias Kowatsch. 2019. “VADLite: An Open-Source Lightweight System for Real-Time Voice Activity Detection on Smartwatches”. In *Adjunct Proceedings of the 2019 ACM International Joint Conference on Pervasive and Ubiquitous Computing and the 2019 International Symposium on Wearable Computers (UbiComp/ISWC '19 Adjunct)*, September 9–13, 2019, London, United Kingdom. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3341162.3346274>

### Abstract

Smartwatches provide a unique opportunity to collect more speech data because they are always with the user and also have a more exposed microphone compared to smartphones. Speech data could be used to infer various indicators of mental well-being such as emotions, stress, and social activity. Hence, real-time voice activity detection (VAD) on smartwatches could enable the development of applications for mental health monitoring. In this work, we present VADLite, an open-source, lightweight, system that performs real-time VAD on smartwatches. It extracts mel-frequency cepstral coefficients and classifies speech versus non-speech audio samples using a linear Support Vector Machine. The real-time implementation is done on the Wear OS Polar M600 smartwatch. An

offline and online evaluation of VADLite using real-world data showed better performance than WebRTC’s open-source VAD system. VADLite can be easily integrated into Wear OS projects that need a lightweight VAD module running on a smartwatch.

### Contribution

Our contribution is the development and evaluation of VADLite, an open-source lightweight software system <sup>1</sup> that performs real-time voice activity detection (VAD) on smartwatches and runs efficiently on constrained systems via using a linear support vector machine, consequently addressing the gap in obtaining an easy-to-use smartwatch VAD system. For this paper, I ran the user study to collect data, preprocessed the data, extracted features, implemented the machine learning experiments, designed and implemented the VADLite smartwatch app, wrote the first draft of the paper, and then the final draft after receiving feedback.

## Paper 7: DyMand

### Reference

**George Boateng**, Prabhakaran Santhanam, Elgar Fleisch, Janina Lüscher, Theresa Pauly, Urte Scholz, and Tobias Kowatsch. “Development, Deployment, and Evaluation of DyMand—An Open-Source Smartwatch and Smartphone System for Capturing Couples’ Dyadic Interactions in Chronic Disease Management in Daily Life”. *Sensors* (*under review*). <https://arxiv.org/abs/2205.07671>

### Abstract

Dyadic interactions of couples are of interest as they provide insight into relationship quality and chronic disease management. Currently, ambulatory assessment of couples’ interactions entails collecting data at random or scheduled times which could miss significant couples’ interaction/conversation moments. In this work, we developed, deployed, and evaluated DyMand, a novel open-source smartwatch and smartphone system for collecting self-report and sensor data from couples based on partners’ interaction moments. Our smartwatch-based algorithm uses the Bluetooth signal strength between two smartwatches each worn by one partner, and a voice activity detection machine-learning algorithm to infer that the partners are interacting, and then to trigger data collection. We deployed the DyMand system in a 7-day field study and collected data about social support, emotional well-being, and health behavior from 13 (N=26) Swiss-based heterosexual couples managing diabetes mellitus type 2 of one partner. Our system triggered 99.1% of the expected number of sensor and self-report data when the app was running, and 77.6% of algorithm-triggered

---

<sup>1</sup><https://https://bitbucket.org/Jojo29/vadlite/>

recordings contained partners’ conversation moments compared to 43.8% for scheduled triggers. The usability evaluation showed that DyMand was easy to use. DyMand can be used by social, clinical, or health psychology researchers to understand the social dynamics of couples in everyday life, and for developing and delivering behavioral interventions for couples who are managing chronic diseases.

### Contribution

This work incorporates VADLite from Paper 6. Our contributions are as follows: (1) designed and developed DyMand, a novel open-source smartwatch <sup>2</sup>, and smartphone <sup>3</sup> system that uses the Bluetooth signal strength between two smartwatches each worn by one partner, and a voice activity detection machine-learning algorithm to infer that the partners are interacting, and then to trigger sensor and self-report data collection (2) deployment and evaluation of DyMand in a field study with heterosexual couples in Switzerland that are managing type 2 diabetes (T2DM) of one partner. For this paper, I co-designed the DyMand system, implemented the DyMand smartwatch app, co-ran the user study to collect data, preprocessed the data, performed data analysis, wrote the first draft of the paper, and then the final draft after receiving feedback.

### Paper 8: Are you ok, honey?

#### Reference

**George Boateng**, Xiangyu Zhao, Malgorzata Speichert, Elgar Fleisch, Janina Lüscher, Theresa Pauly, Urte Scholz, Guy Bodenmann and Tobias Kowatsch. ““Are you okay, honey?”: Recognizing Emotions among Couples Managing Diabetes in Daily Life using Multimodal Smartwatch Data”. *ACM Interact. Mob. Wearable Ubiquitous Technol* (*under review*). <http://arxiv.org/abs/2208.08909>

#### Abstract

Couples generally manage chronic diseases together and the management takes an emotional toll on both patients and their romantic partners. Consequently, recognizing the emotions of each partner in daily life could provide an insight into their emotional well-being in chronic disease management. Currently, the process of assessing each partner’s emotions is manual, time-intensive, and costly. Despite the existence of works on emotion recognition among couples, none of these works have used data collected from couples’ interactions in daily life. In this work, we collected 85 hours (1,021 5-minute samples) of real-world multimodal smartwatch sensor data (speech, heart rate, accelerometer, and gyroscope) and self-reported emotion data (n=612) from 26 partners (13 couples)

---

<sup>2</sup><https://bitbucket.org/mobilecoach/dymandwatchclient/src/master/>

<sup>3</sup><https://bitbucket.org/mobilecoach/dymand-mobilecoach-client/src/master/>



managing diabetes mellitus type 2 in daily life. We extracted physiological, movement, acoustic, and linguistic features, and trained machine learning models (support vector machine and random forest) to recognize each partner’s self-reported emotions (valence and arousal). Our results from the best models — balanced accuracies of 63.8% and 78.1% for arousal and valence respectively — are better than chance and our prior work that also used data from German-speaking, Swiss-based couples, albeit, in the lab. This work contributes toward building automated emotion recognition systems that would eventually enable partners to monitor their emotions in daily life and enable the delivery of interventions to improve their emotional well-being.

### **Contributions**

This final paper builds upon all previous 7 papers. It summarizes and integrates the content of the survey paper (Paper 1), uses the data preprocessing, feature extraction and machine learning approaches detailed in Papers 2, 3, 4, and 5, and uses the smartwatch and smartphone systems from Papers 6 and 7 for data collection. This work is the first to recognize the emotions of romantic partners using data collected from everyday life. Our contributions are as follows (1) collection and use of a unique dataset — real-world, multimodal smartwatch sensor data from German-speaking, Swiss-based couples (N=13 couples, n=26 participants), which is the first such dataset used in the literature for automatic recognition of partners’ emotions (2) approaches for validating and quantifying data quality on manually coded, annotated and transcribed real-world speech data (3) development and evaluation of a machine learning system to recognize the emotions of each partner using a wide variety of sensor data — acoustic, linguistic, heart rate, accelerometer, and gyroscope (4) an investigation of the sensor modality combinations which result in the best emotion recognition performance of romantic partners. For this paper, I conceptualized the key idea, preprocessed the dataset, extracted features, implemented the machine learning experiments, wrote the first draft of the paper, and then the final draft after receiving feedback.

## References

- [1] Mojtaba Khomami Abadi et al. “DECAF: MEG-based multimodal database for decoding affective physiological responses”. In: *IEEE Transactions on Affective Computing* 6.3 (2015), pp. 209–222.
- [2] Hoda Badr and Linda K Acitelli. “Re-thinking dyadic coping in the context of chronic illness”. In: *Current Opinion in Psychology* 13 (2017), pp. 44–48.
- [3] Jacopo Biggiogera et al. “BERT Meets LIWC: Exploring State-of-the-Art Language Models for Predicting Communication Behavior in Couples’ Conflict Interactions”. In: *Companion Publication of the 2021 International Conference on Multimodal Interaction*. New York, NY, USA: Association for Computing Machinery, 2021, pp. 385–389. ISBN: 9781450384711. URL: <https://doi.org/10.1145/3461615.3485423>.
- [4] George Boateng, Elgar Fleisch, and Tobias Kowatsch. “Emotion Recognition among Couples: A Survey”. In: *arXiv preprint arXiv:2202.08430* (2022).
- [5] George Boateng and Tobias Kowatsch. “Speech Emotion Recognition among Elderly Individuals using Multimodal Fusion and Transfer Learning”. In: *Companion Publication of the 2020 International Conference on Multimodal Interaction (ICMI ’20 Companion), October 25–29, 2020, Virtual event, Netherlands*. 2020.
- [6] George Boateng et al. “Are you okay, honey?”: Recognizing Emotions among Couples Managing Diabetes in Daily Life using Multimodal Real-World Smartwatch Data. 2022. DOI: [10.48550/ARXIV.2208.08909](https://doi.org/10.48550/ARXIV.2208.08909). URL: <https://arxiv.org/abs/2208.08909>.
- [7] George Boateng et al. ““You Made Me Feel This Way”: Investigating Partners’ Influence in Predicting Emotions in Couples’ Conflict Interactions Using Speech Data”. In: *Companion Publication of the 2021 International Conference on Multimodal Interaction*. ICMI ’21 Companion. Montreal, QC, Canada: Association for Computing Machinery, 2021, pp. 390–394. ISBN: 9781450384711. DOI: [10.1145/3461615.3485424](https://doi.org/10.1145/3461615.3485424). URL: <https://doi.org/10.1145/3461615.3485424>.
- [8] George Boateng et al. “Development, Deployment, and Evaluation of DyMand—An Open-Source Smartwatch and Smartphone System for Capturing Couples’ Dyadic Interactions in Chronic Disease Management in Daily Life”. In: *arXiv preprint arXiv:2205.07671* (2022).
- [9] George Boateng et al. “Speech Emotion Recognition among Couples using the Peak-End Rule and Transfer Learning”. In: *Companion Publication of the 2020 International Conference on Multimodal Interaction (ICMI ’20 Companion), October 25–29, 2020, Virtual event, Netherlands*. 2020.

- [10] George Boateng et al. “VADLite: an open-source lightweight system for real-time voice activity detection on smartwatches”. In: *Adjunct Proceedings of the 2019 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2019 ACM International Symposium on Wearable Computers*. 2019, pp. 902–906.
- [11] Guy Bodenmann. “Dyadic coping—a systematic-transactional view of stress and coping among couples: Theory and empirical findings”. In: *European Review of Applied Psychology* 47 (1997), pp. 137–140.
- [12] Laura L Carstensen, John M Gottman, and Robert W Levenson. “Emotional behavior in long-term marriage.” In: *Psychology and aging* 10.1 (1995), p. 140.
- [13] Sandeep Nallan Chakravarthula, Brian Baucom, and Panayiotis Georgiou. “Modeling Interpersonal Influence of Verbal Behavior in Couples Therapy Dyadic Interactions”. In: *arXiv preprint arXiv:1805.09436* (2018).
- [14] James A Coan and John M Gottman. “The specific affect coding system (SPAFF)”. In: *Handbook of emotion elicitation and assessment* (2007), pp. 267–285.
- [15] Paul Ekman and Dacher Keltner. “Universal facial expressions of emotion”. In: *Seegerstrale U, P. Molnar P, eds. Nonverbal communication: Where nature meets culture* (1997), pp. 27–46.
- [16] John Mordechai Gottman. *The mathematics of marriage: Dynamic nonlinear models*. MIT Press, 2005.
- [17] Richard E Heyman. “Observation of couple conflicts: Clinical assessment applications, stubborn truths, and shaky foundations.” In: *Psychological assessment* 13.1 (2001), p. 5.
- [18] Patricia K Kerig and Donald H Baucom. *Couple observational coding systems*. Taylor & Francis, 2004.
- [19] Angeliki Metallinou and Shrikanth Narayanan. “Annotation and processing of continuous emotional attributes: Challenges and opportunities”. In: *2013 10th IEEE international conference and workshops on automatic face and gesture recognition (FG)*. IEEE. 2013, pp. 1–8.
- [20] Soujanya Poria et al. “Emotion recognition in conversation: Research challenges, datasets, and recent advances”. In: *IEEE Access* 7 (2019), pp. 100943–100953.
- [21] Tracey A Revenson and Anita DeLongis. “Couples coping with chronic illness.” In: (2011).
- [22] Megan L Robbins et al. “Cancer conversations in context: naturalistic observation of couples coping with breast cancer.” In: *Journal of Family Psychology* 28.3 (2014), p. 380.
- [23] Nicole A Roberts, Jeanne L Tsai, and James A Coan. “Emotion elicitation using dyadic interaction tasks”. In: *Handbook of emotion elicitation and assessment* (2007), pp. 106–123.

- [24] James A Russell. “A circumplex model of affect.” In: *Journal of personality and social psychology* 39.6 (1980), p. 1161.
- [25] Dominik Schoebi. “The coregulation of daily affect in marital relationships.” In: *Journal of Family Psychology* 22.4 (2008), p. 595.
- [26] Salvatore Settineri et al. “Caregiver’s burden and quality of life: Caring for physical and mental illness”. In: *International Journal of Psychological Research* 7.1 (2014), pp. 30–39.
- [27] Rolf Steyer et al. “Der Mehrdimensionale Befindlichkeitsfragebogen MDBF [Multidimensional mood questionnaire]”. In: *Göttingen, Germany: Hogrefe* (1997).
- [28] David Watson, Lee Anna Clark, and Auke Tellegen. “Development and validation of brief measures of positive and negative affect: the PANAS scales.” In: *Journal of personality and social psychology* 54.6 (1988), p. 1063.
- [29] Sebastian Zepf et al. “Driver emotion recognition for intelligent vehicles: a survey”. In: *ACM Computing Surveys (CSUR)* 53.3 (2020), pp. 1–30.

## **Paper 1: Survey of Couples' Emotion Recognition**

# Emotion Recognition among Couples: A Survey

GEORGE BOATENG, ETH Zürich, Switzerland

ELGAR FLEISCH, ETH Zürich, Switzerland and University of St. Gallen, Switzerland

TOBIAS KOWATSCH, ETH Zürich, Switzerland and University of St. Gallen, Switzerland

Couples' relationships affect the physical health and emotional well-being of partners. Automatically recognizing each partner's emotions could give a better understanding of their individual emotional well-being, enable interventions and provide clinical benefits. In the paper, we summarize and synthesize works that have focused on developing and evaluating systems to automatically recognize the emotions of each partner based on couples' interaction or conversation contexts. We identified 28 articles from IEEE, ACM, Web of Science, and Google Scholar that were published between 2010 and 2021. We detail the datasets, features, algorithms, evaluation, and results of each work as well as present main themes. We also discuss current challenges, research gaps and propose future research directions. In summary, most works have used audio data collected from the lab with annotations done by external experts and used supervised machine learning approaches for binary classification of positive and negative affect. Performance results leave room for improvement with significant research gaps such as no recognition using data from daily life. This survey will enable new researchers to get an overview of this field and eventually enable the development of emotion recognition systems to inform interventions to improve the emotional well-being of couples.

CCS Concepts: • **General and reference** → **Surveys and overviews**; • **Human-centered computing** → **Human computer interaction (HCI)**; • **Applied computing** → **Psychology**.

Additional Key Words and Phrases: Couples, emotion recognition, affective computing, literature survey

## 1 INTRODUCTION

The emotions experienced by romantic partners are linked with relationship quality and the management of chronic diseases. Couples' emotions experienced during conflicts predict if these couples stay together in the long term ([41]). For example, couples heading for break-up show more negative emotions and less positive emotions than happy couples [19, 40]. For couples with one partner having a chronic disease, the burden of the disease management is shared by both partners and it takes a toll on the emotional well-being of not just the patient but also on the supporting partner. Furthermore, social support from partners in chronic disease management has been shown to either have positive or negative effects on the emotional well-being of patients [15, 50, 76]. Because of the importance of emotions among couples, researchers are working towards understanding the emotional processes that take place in intimate relationships (e.g., [34, 90]) and the link between emotions and social support in couples' dyadic management of chronic diseases [64]. Consequently, being able to automatically recognize each partner's emotions could enable the research of social and health psychologists, and also inform the development of dyadic interventions (where partners are both involved e.g., [54]) to improve the emotional well-being, relationship quality, and chronic disease management of couples.

Emotion recognition among couples entails the recognition of the emotions of each romantic partner based on the context of their interaction. This task has a number of differences and similarities with other kinds of emotion recognition tasks. Standard emotion recognition attempts to recognize the emotion of each individual and uses various kinds of stimuli to induce an emotional reaction such as driving [104], listening to music or watching a movie [1], giving a speech [85] and engaging in a conversation with another person [75].

The stimuli used for couples' emotion recognition is a conversational context and hence, it has some similarities to emotion recognition tasks based on individuals having a conversation. The conversational context has the

---

Authors' addresses: George Boateng, gboateng@ethz.ch, ETH Zürich, Zurich, Switzerland; Elgar Fleisch, efleisch@ethz.ch, ETH Zürich, Zurich, Switzerland and University of St. Gallen, St. Gallen, Switzerland; Tobias Kowatsch, tkowatsch@ethz.ch, ETH Zürich, Zurich, Switzerland and University of St. Gallen, St. Gallen, Switzerland.

unique challenge of turn-taking dynamics which requires the system to correctly identify who is speaking for each speech segment, termed speaker diarization, consequently making this kind of emotion recognition task more challenging than others that employ stimuli such as listening to music or watching a video.

Most emotion recognition tasks using conversational contexts tend to use actors acting out hypothetical scenarios [17, 18, 66]. Actors tend to exaggerate their emotional expressions and the models trained with that type of data have been shown to perform better than models trained on real-world data [30]. It is not clear if such emotion recognition systems will perform well on data from non-actors. On the other hand, some emotion recognition tasks employ people (e.g., 2 strangers) having real conversations [43, 71]. This type of emotion recognition is the closest to couples' emotion recognition in terms of the conversational context and the dyadic and realistic nature of the interaction. However, for the couples' context, because the two individuals involved are in a romantic relationship, various insights from psychology research about couples can be leveraged to better recognize the emotion of each partner. For example, partners tend to influence each other's emotions throughout an interaction [3], and hence various interpersonal dynamics could be leveraged to adequately recognize each partner's emotions (e.g., [11, 70]).

Because of the uniqueness of couples' emotion recognition and the potential clinical utility, there is a need to synthesize emotion recognition approaches focused on the couples' context. Several works have developed systems to automatically recognize the emotions of couples. In this paper, we describe and discuss these works and give a comprehensive overview of this research field. We surveyed 28 articles published over the past 11 years (2010 - 2021) given the first set of works on this topic were published in 2010. We detail the datasets, features, algorithms, evaluation, and results of each work as well as present main themes. We also discuss current challenges, research gaps and propose future research directions. There are surveys of emotion recognition works focused on specific modalities such as visual and speech modalities ([103]), wearables ([86]), multimodality ([30, 74]) and contexts such as driving ([104]) and conversation ([75]). However, this is the first survey of works that focus on emotion recognition within the context of couples' interactions or conversations.

The rest of this paper is organized as follows. In Section 2, we describe the scope of the survey and the approach used to select the papers. In Section 3, we give an overview of the surveyed works. In Section 4, we describe emotion models, elicitation, and annotation approaches that have been used. In Section 5, we give an overview of the datasets that have been used in the surveyed works. In Section 6, we describe the modalities used, how they have been preprocessed, and the features extracted from them. In Section 7, we discuss the algorithms that have been used, modeling approaches considering the unique context of couples' interactions, evaluation approaches, and the results obtained. In Section 8, we discuss research gaps, challenges, and future directions. We conclude in Section 9.

## 2 SURVEY SCOPE AND METHODOLOGY

Our methodology is similar to those of Zepf et al. [104]. We sought to survey papers that used data to automatically recognize the emotions of each romantic partner based on the couple's interaction or conversation context. We developed a list of search terms that covered three concepts: (1) emotions/emotional behavior (emotion, affect, affective, moods, behavior), (2) recognition (recognition, prediction, classification, behavior signal processing, affective computing, machine learning, deep learning, neural networks) and (3) couples (couples, dyad, spouse, married). We entered the search terms into the following databases: IEEE, ACM, Web of Science, and Google Scholar. We also looked through the references of relevant papers.

To be included, the papers had to perform automatic recognition of each partner's emotions (e.g., positive or negative valence), emotional behavior or emotional states (e.g., positive affect/positivity, negative affect/negativity, sadness), employed statistical or machine learning approaches, and used data collected from the context of couples'

interaction or conversations. We included papers that were not peer-reviewed yet but were on archival databases such as arXiv for completeness.

We excluded papers that used data from interacting partners that are not real couples, i.e., individuals acting out dyadic interactions either using a script or engaging in spontaneous sessions such as the following datasets [17, 18, 66, 69]. We also excluded papers that recognized couple behavior which are not emotional states such as level of blame [7], conflict [92], suicidal risk [23]. In particular, we excluded one paper [2] that focused on recognizing stressful conversations from other kinds of stressful situations among couples. We also excluded papers that recognized couples' relationship state (e.g., happy vs sad couple) [97] rather than each partner's emotions. Another related paper that we excluded is [9] which proposed a research plan for emotion recognition but does not present performed analysis and results in the paper. After using these inclusion and exclusion criteria, we had 28 relevant articles published between 2010 and 2021 (Table 2).

### 3 OVERVIEW OF WORKS

Out of the 28 surveyed papers, a majority of the works ( $n=24$ ) have been done primarily by the Signal Analysis and Interpretation Laboratory (SAIL) team at the University of Southern California which published the first set of works on this topic in 2010 [6, 56]. Subsequent works extended or built upon previous works from the research group. Few works have been done by researchers outside this research lab and include the following [5, 11, 13, 26].

Together, these works have used 5 datasets that were collected from a laboratory setting (Table 1). Most of these works used emotion labels from external raters with few using self-reported data from the couples themselves. Only three modalities have been used — acoustic, lexical, and visual — with acoustic being the most used modality. Multimodal fusion has been done mostly for acoustic and lexical modalities using feature-level and decision-level fusion. Support vector machines are the most used algorithm. Various intrapersonal considerations (e.g., saliency) and interpersonal considerations (e.g., synchrony) have been leveraged (details in a future section). Evaluations have mostly been done with leave-one-couple-out cross-validation with accuracy being used as the metric.



Table 1. Overview of datasets

Dataset	No of Couples	Couples Context	Data	Hours of Data	Session Type	Mins per Session	Annotation Type	Annotation Scope	Emotion Model	Emotions
UCLA / UW Couples Therapy	134	English-speaking, Chronically distressed couples	Video, audio, transcripts	96	Relationship problem (1)	10	Observer	Global	Dimensional, Categorical	Positive affect, Negative affect, Sadness, Anger, Anxiety
Moffit Center Cancer Conversation	85	English-speaking, Couples managing cancer	Video, audio, transcripts	27	Neutral (1), Cancer management issue (1)	10	Observer	Local (utterance / speaker-turn)	Categorical	Hostile (Negative), Positive, Constructive (Neutral)
KU Leuven Dyadic Interaction	101	Dutch-speaking Couples	Video, audio	51	Neutral (1), Positive (1), Negative (1)	10	Self	Global, Local (continuous)	Categorical, Dimensional	Categorical: anger, sadness, anxiety, relaxation, happiness, Dimensional: Valence and Arousal
Stanford Psychotherapy	3	English-speaking Couples Therapy	Video, audio, transcripts	18	Relationship discussion (1)	60	Observer	Local (utterance)	Categorical	Anger, sadness, joy, tension, neutral
UZH Couples Interaction	368	German-speaking couples in Switzerland	Video, audio, transcripts	637	Conflict discussion (1), Mutual support discussion (2)	8	Self-Observer	Global, Local (utterance)	Dimensional	Positive, negative, happy vs sad, good mood vs bad mood, relaxed vs angry, calm vs stressed

Table 2. Studies focused on emotion recognition among couples. ACC=Accuracy, Corr=Spearman Correlation, CV=cross validation, DD=Diversity Density, DNN=Deep neural network, GMM=Gaussian Mixture Models, HMM=Hidden Markov Model, LDA=Linear discriminant analysis, LR=Logistic Regression, LNCO=Leave-n-couple-out, LOCO= Leave-one-couple-out, MAE=Mean Average Error, ML=Maximum Likelihood, MM=Markov Model, RBF=Radial Basis Function, RF=Random Forest, SPRT=Sequential Probability Ratio Test, SVM=Support vector machine, UAR=Unweighted Average Recall

Ref	Dataset	Modalities	Features	Interpersonal	Intrapersonal	Algorithms	Evaluation	Metric	Class	Main Best Results
Biggiogera et al., 2021 [5]	UZH Couples Interaction	Acoustic (A), Lexical (L)	A: Prosodic and spectral eGeMAPS features, L: Ngram + TFIDF, LIWC, Deep sentence embeddings	No	No	Linear SVM	10-fold CV couple disjoint	UAR	2	Negative vs Positive affect: 69.2%
Black et al., 2010 [6]	UCLA / UW Couples Therapy	Acoustic	Prosodic and spectral	No	No	Linear SVM, LDA	LOCO CV	ACC	2	Positive affect: 82% (female), 75% (male), Negative affect: 77% (female), 76% (male)
Black et al., 2013 [8]	UCLA / UW Couples Therapy	Acoustic	Prosodic and spectral	No	No	Linear SVM, LR	LOCO CV	ACC	2	Positive affect: 77.9% (female), 72.9 (male), Negative affect: 80% (female), 85.7% (male)
Boateng et al., 2020 [13]	KU Leuven Dyadic Interaction	Acoustic	Deep acoustic embeddings (Spectrograms + Pretrained CNN)	No	Yes	Linear SVM	LOCO CV	UAR	2	Negative vs Positive valence: 74.8% (female), 53.3% (male)
Boateng et al., 2021 [11]	UZH Couples Interaction	Acoustic (A), Lexical (L); Fusion: Feature level	A: Prosodic and spectral eGeMAPS features, L: Deep sentence embeddings	Yes (Dyadic Influence)	No	Linear SVM, RBF SVM, RF	10-fold CV couple disjoint	UAR	2	Negative vs Positive valence: 64.8% (female), 56.1% (male)
Chakravarthula et al., 2015 [21]	UCLA / UW Couples Therapy	Lexical	Unigram	No	Yes (Dy-namic modeling)	ML as baseline, HMM	LOCO CV	ACC	2	Negative affect: 88.57% (female), 83.57% (male)
Chakravarthula et al., 2018 [70]	UCLA / UW Couples Therapy	Lexical	Ngram	Yes (Dyadic Influence)	Yes (Dy-namic modeling)	ML as baseline, HMM	LOCO CV	ACC	2	Negative affect: 88.93%
Chakravarthula et al., 2019 [22]	Moffit Center Cancer Conversation	Acoustic, Lexical; Fusion: Feature level, Decision	Acoustic: Prosodic and spectral eGeMAPS features; Lexical: Deep sentence embedding	No	No	DNN	LOCO CV	UAR	3	Positive vs Negative vs Neutral: 57.42%
Chakravarthula et al., 2021 [20]	UCLA / UW Couples Therapy	Lexical	N-gram, ELMo	No	No	ML, GRU	6-fold CV couple disjoint	Corr	N/A	Positive affect: 0.5, Negative affect: 0.58, Anxiety: 0.18, Anger: 0.52, Sadness: 0.34
Crangle et al., 2019 [26]	Stanford Psychotherapy	Acoustic	Prosodic and spectral features	No	No	RF	Hold out	ACC	5	Anger, sadness, joy, tension, neutral, Couple A: 87% (male), 90% (female), Couple B: 78% (male), 84% (female), Couple C: 95% (male), 88% (female)
Georgiou et al., 2011 [37]	UCLA / UW Couples Therapy	Lexical	Unigram	No	No	ML	LOCO CV	ACC	2	Positive affect: 88.9%, Negative affect: 86.7%, Sadness: 61.6%

Table 2. Studies focused on emotion recognition among couples. ACC=Accuracy, Corr=Spearman Correlation, CV=cross validation, DD=Diversity Density, DNN=Deep neural network, GMM=Gaussian Mixture Models, HMM=Hidden Markov Model, LDA=Linear discriminant analysis, LR=Logistic Regression, LNCO=Leave-n-couple-out, LOCO= Leave-one-couple-out, MAE=Mean Average Error, ML=Maximum Likelihood, MM=Markov Model, RBF=Radial Basis Function, RF=Random Forest, SPRT=Sequential Probability Ratio Test, SVM=Support vector machine, UAR=Unweighted Average Recall

Ref	Dataset	Modalities	Features	Interpersonal	Intrapersonal	Algorithms	Evaluation	Metric	Class	Main Best Results
Gibson et al., 2011 [38]	UCLA / UW Couples Therapy	Acoustic	Spectral	No	Yes (Sailency)	RBF SVM as baseline, DD SVM	10-fold CV couple disjoint	ACC	2	Positive affect: 74.3% (female), 58.6% (male), Negative affect: 77.9% (female), 71.4% (male), Sadness: 66.4% (female), 63.6% (male)
Gibson et al., 2015 [39]	UCLA / UW Couples Therapy	Acoustic (A), Lexical (L), Visual (V); Fusion: Feature level, Decision level	A: Prosodic and spectral, L: TFIDF, V: Power spectral density (PSD) of head motion vectors	No	Yes (Sailency)	DD, Linear SVM	LOCO CV	ACC	2	Positive affect: 65.13% (all modalities), Negative affect: 70.34% (lexical)
Katsamanis et al., 2011 [52]	UCLA / UW Couples Therapy	Acoustic, Lexical; Fusion: Unclear	MFCC, TFIDF	No	Yes (Sailency)	RBF SVM as baseline, DD SVM.	10-fold CV couple disjoint	ACC	2	Positive affect: 93% (lexical), Negative affect: 95% (lexical), Sadness: 80% (lexical)
Lee et al., 2010 [56]	UCLA / UW Couples Therapy	Acoustic	Prosodic	Yes (Synchrony)	No	MM	LOCO CV	ACC	2	Positive vs negative affect: 76%
Lee et al., 2011a [58]	UCLA / UW Couples Therapy	Acoustic	Prosodic and spectral	Yes (Synchrony)	No	RBF SVM	LOCO CV	ACC	2	Positive vs negative affect: 51.79%
Lee et al., 2011b [59]	UCLA / UW Couples Therapy	Acoustic	Prosodic and spectral	Yes (Synchrony)	Yes (Sailency)	DD	LOCO CV	ACC	2	Positive vs negative affect: 53.93%
Lee et al., 2012 [60]	UCLA / UW Couples Therapy	Lexical	TFIDF	No	Yes (Sailency)	DD, SPRT	10-fold CV couple disjoint	ACC	2	Positive affect: 76.1%, Negative affect: 74.2%, Sadness: 54.2%
Lee et al., 2014 [57]	UCLA / UW Couples Therapy	Acoustic	Prosodic and spectral	Yes (Synchrony)	No	HMM, Factorial HMM	LOCO CV	ACC	2	Positive vs negative affect: 62.86%
Li et al., 2016 [61]	UCLA / UW Couples Therapy	Acoustic	Prosodic and spectral features	No	No	SVM as baseline, DNN	LOCO CV	ACC	2	Negative affect: 77.14%
Li et al., 2017 [62]	UCLA / UW Couples Therapy	Acoustic	Prosodic and spectral	No	No	DNN (autoencoder)	LOCO CV	ACC	2	Negative affect: 69.64%, Positive affect: 66.43%
Li et al., 2020 [63]	UCLA / UW Couples Therapy	Acoustic	Prosodic and spectral, Deep acoustic embeddings	No	No	CNN, GRU	LNCO (n=4)	ACC	2	Positive affect: 65.36%, Negative affect: 76.07%, Sadness: 59.29%
Tseng et al., 2016 [95]	UCLA / UW Couples Therapy	Lexical	word2vec	No	No	ML as baseline, LSTM + RBF SVR	LOCO CV	ACC	2	Negative affect: 88.93%

Table 2. Studies focused on emotion recognition among couples. ACC=Accuracy, Corr=Spearman Correlation, CV=cross validation, DD=Diversity Density, DNN=Deep neural network, GMM=Gaussian Mixture Models, HMM=Hidden Markov Model, LDA=Linear discriminant analysis, LR=Logistic Regression, LNCO=Leave-n-couple-out, LOCO= Leave-one-couple-out, MAE=Mean Average Error, ML=Maximum Likelihood, MM=Markov Model, RBF=Radial Basis Function, RF=Random Forest, SPRT=Sequential Probability Ratio Test, SVM=Support vector machine, UAR=Unweighted Average Recall

Ref	Dataset	Modalities	Features	Interpersonal	Intrapersonal	Algorithms	Evaluation	Metric	Class	Main Best Results
Tseng et al., 2017 [94]	UCLA / UW Couples Therapy	Lexical	word2vec, Deep sentence embeddings	No	No	LSTM + RBF SVR	LOCO CV	MAE	N/A	Negative affect: 1.37
Tseng et al., 2018 [96]	UCLA / UW Couples Therapy	Acoustic (A), Lexical (L); Fusion: Feature level, Decision level, Gender based, Therapy stage	A: Prosodic and spectral features, L: Deep sentence embeddings	No	No	LSTM, DNN	LOCO CV	MAE	2	Negative affect: 1.22
Tseng et al., 2019 [93]	UCLA / UW Couples Therapy	Lexical	Deep sentence embeddings	No	No	LSTM	LOCO CV	ACC	2	Positive affect: 86.8% Negative affect: 87.9%
Xiao et al., 2015 [101]	UCLA / UW Couples Therapy	Acoustic	Prosodic and spectral	No	Yes (Dynamic modeling)	SVM, LDA, Voted Perceptron as baseline, HMM + SVM, LDA, Voted Perceptron	LOCO CV	ACC	2	Positive affect: 81% (female), 78% (male), Negative affect: 79% (female), 84% (male), Sadness: 65% (female), 61% (male)
Xiao et al., 2015 [102]	UCLA / UW Couples Therapy	Visual	Line Spectral Frequencies of head motion vectors	Yes (Synchrony)	No	GMM + Linear SVM	LOCO CV	ACC	2	Positive affect: 63%, Negative affect: 57%

## 4 BACKGROUND

In this section, we describe various emotion models and approaches to eliciting and annotating emotion data from couples.

### 4.1 Emotion Models

There are mainly two models of emotions used in the literature in emotion recognition: categorical and dimensional. Categorical emotions are based on the six basic emotions proposed by Ekman: happiness, sadness, fear, anger, disgust, and surprise [31]. Over time, additional emotion categories have been included and used in literature such as anxiety, frustration, etc. Dimensional approaches mainly use two dimensions: valence (pleasure) and arousal (activation) which are based on Russell's circumplex model of emotions [82]. Valence refers to how negative to positive the person feels and arousal refers to how sleepy to active a person feels. Using these two dimensions, several categorical emotions can be placed and grouped into the four quadrants: high arousal and negative valence (e.g., stressed, angry), low arousal and negative valence (e.g., depressed), low arousal and positive valence (e.g., relaxed), and high arousal and positive valence (e.g., excited).

### 4.2 Elicitation

Approaches for eliciting emotions in couples have generally happened in the lab/controlled settings and in daily life. In the lab, couples are asked to have emotionally charged conversations that are videotaped [80]. Some of these conversations center on topics that cause distress in their relationship. These conversations elicit emotions during and after the conversations which are then annotated. In daily life, sensor data (e.g., audio) is collected from couples periodically [79] or when conversation moments are detected [12].

### 4.3 Annotation

Two approaches are used for emotion annotations by social psychologists: self-report and observer reports, and two scales: global and local (continuous, utterance-level). For self-reports, each partner provides emotion ratings right after the whole interaction/conversation (global/session ratings) with validated instruments such as the Affect Grid questionnaire [83] and Multidimensional Mood questionnaire [91] or they are asked to watch a video recording of the conversation while providing continuous (moment-by-moment) emotion ratings using a joystick (e.g., [80]). In the case of daily life, couples are periodically asked to complete self-reports such as the PANAS [100], Affect Grid questionnaire [83] and Affective Slider [4] at random time [87, 89] or after sensor data recording [12]. Additionally, the dyadic nature of couples' interactions enables the collection of partner-perceived emotions where each partner (e.g., partner A) is asked to provide a rating of their perception of the emotion of their partner (e.g., partner B) emotion after an interaction in the lab (e.g., sels2019a, sels2019b) or in daily life (e.g., sels2020). All these types of self-reports enable the collection of the subjective and perceived emotions of each partner. However, these ratings could be biased and may not reflect each partner's actual emotions about the interaction.

For observer reports, people are trained to watch the video recordings (e.g., in the case of lab data) and use a coding scheme to rate the interaction on specific emotional behaviors (e.g., SPAFF [25]) using continuous or utterance-level ratings (e.g., every 10 seconds or speaker turn) or global ratings (of the whole interaction). Such coding is also done for example, for audio data collected from couples' daily life interactions [79]. This manual coding process is costly and time-consuming as multiple coders need to be trained for this task [53] and suffers from inter-rater reliability issues [47, 67]. Furthermore, these ratings reflect the observers' perceived emotions of the partners and they do not represent the subjective emotions of the partners.

## 5 STUDIES AND DATASETS

In this section, we describe all the datasets that have been used in the surveyed emotion recognition works, how they were collected and annotated. The surveyed papers used five (5) datasets, all of which were collected in the lab. Three (3) were observer annotated, one (1) was self annotated, and one (1) had both self and observer annotations (Table 1). The distribution of papers that have used the five datasets is as follows: UCLA/UW Couples Therapy (23), Cancer Conversation (1), KU Leuven Dyadic Interaction (1), Stanford Psychotherapy (1), and UZH Couples' Interactions (2).

### 5.1 UCLA/UW Couples Therapy

Researchers conducted a longitudinal lab study at the University of California, Los Angeles, and the University of Washington in the U.S. with 134 seriously and chronically distressed heterosexual married couples [6]. Their age statistics were as follows: the range is 22 to 72 years, the median age for men was 43 years ( $SD = 8.8$ ), and the median age for women was 42 years ( $SD = 8.7$ ). They were, on average, college-educated (median level of education for both men and women was 17 years,  $SD = 3.2$ ). The sample was largely Caucasian (77%), with 8% African American, 5% Asian or Pacific Islander, 5% Latino/Latina, 1% Native American, and 4% Other. Couples were married for an average of 10.0 years ( $SD = 7.7$ ) [24].

Couples received couples therapy for 1 year. They had conversations and discussed a problem in their relationship with no therapist or research staff present. They discussed the wife's chosen topic for 10 minutes and the husband's chosen topic for 10 minutes which were considered separate sessions. The sessions were recorded at three points in time: before the therapy, 26 weeks into it, and two years after the therapy sessions ended. There were 96 hours of data across 574 sessions. The sessions were videotaped and later transcribed and annotated by 3-4 trained coders. The annotators assigned 33 session-level (global) behavioral codes for each spouse on a scale of 1 - 9 using two coding schemes. The coding schemes are the Social Support Interaction Rating System (SSIRS) which consists of 20 codes that measure the emotional component of the interaction and the topic of conversation [51] and the Couples Interaction Rating System 2 (CIRS2) which consist of 13 codes and were specifically designed for conversations involving a problem in a relationship [44]. There were no local (utterance- or speaker-turn-level) annotations.

Authors that used this dataset for emotion recognition tasks used only six codes in experiments due to low inter-evaluator agreement for the other codes. The codes that were used are level of blame, level of acceptance towards the other spouse, global positive affect, global negative affect, level of sadness, use of humor. They used the manual transcript of the data to automatically create word and speaker turn alignments which resulted in a smaller number of sessions and unique couples data used for the recognition experiments: 293 sessions. Also, they computed the mean values across raters and then selected data whose ratings were in the top 20% and bottom 20% for each of the codes. Consequently, the data used for analysis was from 60 - 85 unique husband/wife pairs. The task was cast as a binary classification for the two extremes for each code. In this survey, we consider works that used the affect-related codes: global positive affect, global negative affect, level of sadness, anxiety, and anger.

### 5.2 Moffit Center Cancer Conversation

Researchers conducted a study in which they collected data from 85 couples in the U.S. who were coping with advanced cancer (one spouse having cancer and the other being a caregiver) [22, 77, 78]. Here is the demographic data of the 82 couples whose data was eventually used: 29.3% of patients were female and 70.7% of caregivers were female, the mean age for patients was 66.8 years ( $SD = 9.2$ ), and the mean age for caregivers was 64.8 years ( $SD = 9.4$ ). On average, they had college or vocational education. The patient sample was largely Caucasian (92.7%), with 6.1% African American, 3.7% Latino/Latina, 1.2% Native American, and 0% Other. The caregiver sample was

also largely Caucasian (90.2%), with 4.9% African American, 3.7% Latino/Latina, 2.4% Native American, and 1.2% Other. Couples were married for an average of 35 years (SD = 15.8) [77].

They engaged in a 10-minute neutral discussion (daily routine) and a 10-minute stressor discussion about an issue related to cancer management in controlled settings (e.g., clinic consult rooms, participant homes) with an experimenter present without facilitating. The issue was decided based on their ratings on Cancer Inventory of Problem Situations, [46] in which a list of 20 common cancer concerns (e.g., lack of energy, finances, over-protection) are rated as being not a problem, somewhat of a problem, or a severe problem. The interactions were audio-recorded. We estimated that a total of 27 hours of data was collected.

The audio was annotated on an utterance / speaker-turn level by 2 trained coders using the Rapid Marital Interaction Coding System, 2nd Edition [48, 49] with inter-rater reliability scores of Kappas above 0.88 for 20% of all codes). Each utterance was assigned one behavioral code out of 7 codes which were then grouped into three: positive (low and high positive), hostile/negative (low and high hostile), neutral/constructive (constructive problem discussion). Additional codes were dysphoric affect and other which were not used for the recognition task. Hence, the task was framed as a 3-class classification problem. They used the manual transcripts to automatically create word, speaker turn and label alignments. This dataset has been used by [22].

### 5.3 KU Leuven Dyadic Interaction

Researchers conducted a Dyadic Interaction lab study in Leuven, Belgium with 101 heterosexual, Dutch-speaking couples [13, 87, 88]. The majority (n=96) cohabited and 7 were married. The average age was 26 years (SD=5), ranging from 18 to 53 years. The partners were together for 4.5 years (SD=2.8), ranging from 7 months to 21 years. There was no information about the ethnicity and education levels of participants.

These couples were first asked to have a neutral 10-minute conversation, then a 10-minute conversation about a negative topic (a characteristic of their partner that annoys them the most), followed by a 10-minute conversation about a positive topic (a characteristic of their partner that they value the most) [27, 87, 88].

After each conversation, each partner completed self-reports on various emotion labels such as anger, sadness, anxiety, relaxation, and happiness using a 7-point Likert scale ranging from strongly disagree (1) to strongly agree (7). Additionally, each partner watched the video recording of the conversation separately on a computer and rated his or her emotion on a moment-by-moment basis by continuously adjusting a joystick to the left (very negative) and the right (very positive), so that it closely matched their feelings, resulting in valence scores on a continuous scale from -1 to 1 [42, 81]. Additionally, each partner reported how they felt after the interaction and how they thought their partner felt, using the Affect Grid questionnaire [83] which captures the valence and arousal dimensions of Russell's circumplex model of emotions [82] resulting in values between 0 and 8 each for pleasure and arousal.

Trained research assistants (5) listened and visually inspected the audios, and annotated the exact start and end of each talking turn for each partner. Authors that used this data categorized the valence scores into two classes, negative (0-4) and positive valence (5-8) for males and females, consequently framing the task as a binary classification task. This dataset was used by [13]

### 5.4 Stanford Psychotherapy

Researchers collected audio, and video data from 3 heterosexual couples at Stanford University in the U.S. over 18-hour-long couple therapy sessions (18 sessions, 1 hour each for each couple A, B, and C) undergoing therapy over a period of 2 years [26]. We estimated that a total of 18 hours of data was collected. No demographic information about the couples was available.

The audio was first transcribed manually, including the start and end times of each word. Then trained coders watched the video and used the transcript to code the data by marking start and end times of any of the following

4 emotions: anger, sadness, joy, tension, and neutral (defined as the absence of the 4 emotions). Each label was given a rating of low, medium, or high but the levels were not used in the analysis. These codes were adapted from Gottman's 19 SPAFF affective codes [25]. To assign labels, annotators had to mark the start and end times of the occurrence of one of the emotions. Hence, no two emotion-labeled segments could overlap. Only audio data was used for recognition which was framed as a 5-class classification task. This dataset has been used by the work [26].

## 5.5 UZH Couples Interactions

Researchers collected data from 368 heterosexual German-speaking, Swiss couples (N=736 participants; age 20-80) at the University of Zurich, Switzerland over 10 years [55, 98]. The longitudinal study sought to investigate the impact of stress on the relationship development of couples and children across their lifespan. The average age was 47 (SD = 18.4) for women and 49 (SD = 18.2) for men, the mean relationship duration was 21 years (SD = 17.9) with 66% being married. For women, 6% attended the mandatory school years (9 years), 40% completed vocational training, 21% completed high school, and 32% completed college or university. For men, 3% attended the mandatory school years, 35% completed vocational training, 12% completed high school, and 49% completed an academic degree.

Couples participated in three videotaped conversations in the lab each for 8 minutes — one conflict and two mutual support conversations from years 1 to 6 (one session per year). Video-recorded data from 3 couples were not available resulting in 365 couples' data. The number of couples that took part reduced over the years with the details available at [98]. Based on the data collected over the years, our estimate of the total amount of data is 637 hours.

For the conflict interaction which was used for emotion recognition, couples had to choose one problematic topic for the conflict interaction from a list of common problems (PAQ A; [45]), and participants were then videotaped as they discussed the selected issue for 8 minutes. After each conversation, each partner provided self-report responses to the Multidimensional Mood questionnaire [91] of their emotions on four bipolar dimensions — namely “good mood versus bad mood,” “relaxed versus angry,” “happy versus sad” and “calm versus stressed” — with the scale: 1 — very much, 2 — much, 3 — a little, 4 — a little, 5 — much, 6 — very much. The authors that used the dataset preprocessed these responses by averaging the “good mood versus bad mood” and “happy versus sad” scales and then binarized the averaged values such that values greater than or equal to 3.5 were negative (0) and the rest were positive (1).

Additionally, two research assistants were trained to code communication behaviors (interobserver agreement,  $k = 0.9$ ) using an adapted version of the Specific Affect Coding System (SPAFF) [25]. The most prevalent code from the list was assigned every 10 seconds resulting in 48 sequences for each interaction. The codes were then grouped into positive and negative for emotion recognition.

The speech was manually annotated with the start and end of each speaker's turn, along with pauses and noise. The speech was manually transcribed in 15-second chunks separately for each partner. Given that Swiss German is mostly spoken with different dialects across Switzerland, the spoken words were written as the corresponding German word equivalent. The following two works — [11] and [5] — used this data to recognize global and local emotions respectively, all framed as binary classification tasks.

## 6 MODALITIES, DATA PREPROCESSING AND FEATURE EXTRACTION

In this section, we describe the modalities of the data used for emotion recognition in the surveyed works along with preprocessing approaches and features that were extracted from each modality. The surveyed papers used three distinct modalities with acoustic being the most represented modality: acoustic (19), lexical (9), and visual (2).



## 6.1 Acoustic

Given that the audio data is collected in the context of conversations, it is generally annotated manually or automatically with the segments that contain speech vs no speech (voice activity detection), and additionally, those segments are annotated to correspond to the speech of each partner, which is known as speaker diarization.

Next, various features are extracted either using feature engineering or transfer learning. Feature engineering entails using handcrafted features that have been shown to be discriminative for the recognition task. For feature engineering, standard acoustics features such as prosodic (e.g., pitch, energy, speaking rate), spectral (e.g., mel frequency cepstral coefficients) and voice quality (shimmer and jitter) are extracted over frames of short durations (e.g., 25 ms) known as low-level descriptors (LLDs) using a sliding window (e.g., 10 ms) which may or may not be overlapping [5, 6, 8, 11, 13, 22, 26, 38, 39, 52, 56–59, 61–63, 96, 101]. Various statistics called functionals (e.g., mean, median, percentiles, etc) are computed over these frames to get features for a segment (e.g., 2 seconds) or the whole audio (8-10 mins). In particular, a set of 88 features called eGeMAPS [32] have been shown to be a minimalist feature set that is effective for emotion recognition tasks and have been used in the following works [5, 11, 22]. The openSMILE toolkit [33] has been mostly used for acoustic feature extraction. Other tools such as Praat [14] have been used. Due to the likelihood of having a lot of features, various features selection methods such as forward feature selection has been used [6]. Additionally, various approaches have been used to remove the speaker, microphone, and environmental variability of the audio signal by performing mean normalizing of the LLDs for the whole session audio (e.g., [8]).

Transfer learning is an approach used to circumvent the need to develop hand-crafted features and entails using a model pre-trained on a different but related task ([35]). This process entails using the model for feature extraction or fine-tuning in which the whole model or later layers are retrained. For example, Boateng et al. took a CNN model called the YAMNET that was pretrained on an audio event classification task and used it to extract feature embeddings from spectrograms over 1-second time windows ([13]). Also, Li et al pretrained a deep learning model (CNN) on an emotion recognition task that used acted data and then used the model to extract acoustic embeddings for the recognition task [63].

## 6.2 Lexical

The audio is generally transcribed automatically or manually in order to use the content of the speech for emotion recognition. Various linguistic features ranging from simple features (bag of words and TF-IDF), dictionary-based features used in psychology such as LIWC [72], more advanced ones such as word embeddings (word2vec [68] and ELMo [73]) to deep learning models such as BERT [28] which are currently the state-of-the-art for computing linguistic features.

Here are examples of works that have used those lexical features: bag-of-words (unigram [21, 37], ngram [5, 20, 70]), TF-IDF [39, 52, 60], LIWC [5], word embeddings (word2vec [94, 95], ELMo [20]), deep sentence embeddings (seq-to-seq models [22, 93, 94, 96], BERT and Sentence-BERT [5, 11]). Transfer learning has also been used for the lexical data. For example, various sentence embeddings have been computed using pretrained models [5, 11, 22, 93, 94, 96].

## 6.3 Visual

Few works used the visual modality and in particular head movements in the videos that were recorded. The following features have been extracted: line spectral frequencies/power spectral density of head motion vectors to capture the vertical and horizontal directions of the head motion [39, 102]. Facial expressions were not used in those works because the quality of the video was not good enough to compute features from the face (e.g., varying sitting positions, camera distance/angle, and lighting conditions) [39].

## 7 DATA ANALYSIS AND EVALUATION

In this section, we describe various algorithms that have been used for emotion recognition, multimodal fusion approaches, intrapersonal and interpersonal considerations, evaluation, and results.

### 7.1 Algorithms

The surveyed works used mostly supervised learning approaches with a few using semi-supervised [21, 38, 39, 52, 59, 60, 101] and unsupervised learning [62, 93]. The algorithms used range from simple statistical algorithms and traditional machine learning to deep learning methods. Support vector machines (SVM) have been the most used algorithm.

Here are the algorithms used by various works: SVM [5, 6, 8, 11, 13, 38, 39, 52, 58, 61, 94, 95, 101, 102], linear discriminant analysis (LDA) [6, 101], markov models [21, 56, 57, 70, 101], multiple instance learning (diversity density [39, 59, 60], diversity density SVM [38, 52]), maximum likelihood [20, 21, 37, 70], sequential probability ratio test [60], logistic regression [8], perceptron [101], gaussian mixture model (GMM) [102], deep neural networks [22, 61, 62, 96], LSTM [93–96], GRU [20, 63], random forest [11, 26], CNN [63].

### 7.2 Multimodal Fusion

Modalities that have been combined include acoustic and lexical data [11, 22, 52, 96], and acoustic, lexical, and visual [39]. Various fusion methods have been used such as feature-level fusion in which the features of each modality are concatenated and decision-level fusion in which each modality is trained with a separate model and the predictions from the models are combined using various approaches such as majority vote. The following papers used feature-level fusion [11, 22, 39, 96] and decision-level fusion [22, 39, 96]. Additionally, knowledge-driven expert fusion approaches have been explored by Tseng et al. as follows: gender-based and therapy-stage fusion [96].

### 7.3 Intrapersonal Considerations

One challenge with recognizing global emotion labels — one emotion label for a long interaction duration such as 8-10 minutes — is that there is a whole range of emotions experienced and expressed throughout the interaction with different intensities. One naive approach to address this challenge is to assign every segment (e.g., 2 seconds) with the label of the whole audio and then train the model with this modified data-label pairings. This approach is used in various fields such as physical activity recognition [10] since an activity label (e.g., walking) is consistent over different segments. However, such an approach is error-prone for the context of emotion recognition as it erroneously assumes that the emotion label is the same for all segments. The standard approach used in various works is to compute statistics such as mean, median, etc., over the features that have been computed over short windows as previously seen in the approach used for extracting acoustic features.

However, the emotion recognition task may benefit from more creative modeling approaches. One such approach relates to the concept of saliency. Some interaction segments might be more salient for recognizing that one label assigned to the whole audio. Some works have leveraged some methods to identify those salient segments. One such saliency-based method is multiple instance learning which automatically identifies salient instances from a bag of instances in a semi-supervised learning fashion [29]. For example, the following works leveraged multiple instance learning to identify salient instances to use for recognition using acoustic features [38, 59], lexical features [60], both acoustic and lexical features [52], and all three modalities — acoustic, lexical, and visual [39]. Another saliency-based method used the concept of the peak-end rule which posits that how people feel after an emotional experience is predicted by the emotional extremes and the end of that experience [36]. The theory was leveraged to identify salient segments, extract features from those segments, and then

perform the recognition task [13]. Another modeling approach leveraged dynamic modeling of all segments with a Markov model using acoustic features [101] and lexical features [21, 70].

#### 7.4 Interpersonal Considerations

The dyadic nature of couples' interactions offers the opportunity to leverage various interaction dynamics to perform recognition of emotions. One major dyadic dynamic that has been used is synchrony/entrainment which refers to how similar/aligned/synchronized partners are when interacting. Various quantitative measures for synchrony have been computed for various modalities. For acoustic, some examples include prosodic entrainment measures computed with the following similarity measures (1) square of correlation coefficient, (2) mutual information, and (3) mean of spectral coherence over pitch and energy between the sequential turns of partner A and partner B when there is turn change [56]. Another approach leverages principal component analysis (PCA) to compute both prosodic and spectral entrainment while providing information about the directionality of the entrainment [57, 58]. For the visual modality, the Kullback-Leibler (KL) divergence of the features extracted from the head motion of the partners was used as the similarity measure for synchrony [102]. Another dyadic dynamic derives from the idea that partners generally influence each other while interacting. Dyadic influence has been modeled using lexical features from both partners [70] and both acoustic and lexical features from both partners [11].

#### 7.5 Evaluation

Three works performed regression [20, 94, 96] with the rest performing classification. All works that performed classification trained models to perform binary classification except [22] and [26] which performed 3-class and 5-class classification respectively. All works performed global emotion recognition except [22] and [5] which performed utterance-level recognition for every speaker turn and every 10 seconds respectively. All works have used accuracy as the evaluation metric except the following which used unweighted average recall (UAR) [5, 11, 13, 22], Spearman correlation [20] and mean absolute error (MAE) [94, 96]. It is important to note that all the works that used accuracy as the metric had balanced classes (except [26]) and hence, there should not be any concern about it not being an appropriate metric.

Most evaluations have been done with leave-one-couple-out (LOCO) cross-validation which is a robust evaluation approach as it gives a sense of how the model will perform on an unseen couple. With this approach, models are trained using data from all couples but one, and then the prediction is done on the remaining couple's data as the test set. This process is repeated till each couple has been used as a test set. Hence, if there are 300 couples, the evaluation is done 300 times. In the end, the predictions of each test couple are combined either by computing the evaluation metric (e.g., accuracy) separately for each couple and then computing the mean and standard deviation of the accuracies, or concatenating all the predictions and computing one accuracy value for all the combined predictions. LOCO is a variation of the standard leave-one-subject-out cross-validation but more robust for the context of couples data as it ensures that there is no data leakage from the same audio (as an example) being in both train and test sets. One challenge with LOCO is that the evaluation could take a long time when there are a lot of couples as it is done as many times as there are couples.

Other similarly robust evaluation approaches that have been used which also ensure that there is no data leakage but reduces the amount of time relatively are leave-N-couples-out cross-validation (LNCO) (e.g., [63] with  $n=4$ ) and K-fold cross-validation (CV) couple disjoint ( $k=10$ : [5, 11, 38, 52, 60],  $k=6$ : [20]). The "couple disjoint" refers to the fact that a couple is never in both the train and test set for the same evaluation run.

Another approach that has been used is the standard hold out (train test split) evaluation [26]. That work also performed couple-dependent evaluation [26]. That is, the authors trained and evaluated models within each couple separately. Hence, the concept of "couple disjoint" does not apply. The results from such an evaluation

could be inflated as it could leverage particularities of the data to produce high results and does not give a sense of how the model will perform on an unseen couple. Nonetheless, it gives a sense of how well the model may perform if personalized models are trained.

Furthermore, several works have performed gender-specific evaluations where the model is trained and evaluated separately for male and female partners [6, 8, 11, 13, 21, 22, 26, 38, 96, 101]. The motivation for this approach is that gender differences affect how people express their emotions [16] and in particular how they speak and hence training approaches may benefit from using models separately for each gender.

## 7.6 Results

In this section, we summarize the main results across the works and also provide some context to enable the correct interpretation of the results.

The best result for the work that performed 3-class classification (positive, negative, neutral) is 57.4% UAR [22]. The best result for the 5-class classification (anger, sadness, joy, and tension, neutral) which also used couple-dependent evaluation ranged from 78% to 95% for different couples and genders [26].

For binary classification, we provide results for 2 groups of works – (1) works that used the UCLA/UW Couples Therapy dataset in which they only considered ratings at the 2 extremes and (2) works that used other datasets without only considering the extreme ratings. We separate the two because the first task is easier than the second since only extremes are being considered rather than all ratings regardless of the intensity.

For the first group, here are the best accuracies for each emotion task and gender with the corresponding modality shown:

- Positive affect: 82% (female) [6] (acoustic), 78% (male) [101] (acoustic), 93% (combined male and female) [52] (lexical)
- Negative affect: 88.57% (female) [21] (lexical), 85.7% (male) [8] (acoustic), 95% [52] (combined male and female) (lexical)
- Sadness: 66.4% (female) [38] (acoustic), 63.6% (male) [38] (acoustic), 80% [52] (combined male and female)(lexical)
- Positive vs negative: 76% [56] (combined male and female) (lexical)

In light of these high accuracy results, it is worth noting that they are not reflective of true emotion recognition performance since the data was partitioned into two extreme ratings (top 20% and bottom 20%). Consequently, the performance would likely be much lower if all the data were used.

The best results for the second group are the following UAR for positive vs negative: 74.8% (female) [13] (acoustic), 56.1% (male) [11] (acoustic and lexical), 69.2% (combined male and female) [5] (lexical).

For the regression tasks, the best results are 1.22 MAE for negative affect [96] (acoustic and lexical) and Spearman correlation of 0.5 for positive affect, 0.58 for negative affect, 0.18 for anxiety, 0.52 for anger and 0.28 for sadness [20] (lexical).

For works that use gender-specific evaluations, performance for female partners tends to be better than for male partners. These results might suggest that it is more difficult recognizing the emotions of male partners and consistent with insights from psychology that suggest that female partners are more emotionally expressive [16]. Also, in works that consider multiple modalities, lexical modality tends to outperform other modalities including multimodal ones [5, 22, 39, 52]. Considering the different evaluation contexts (e.g., different number of classes, data subsamples, etc.), it is difficult to compare results directly. Nonetheless, excluding results from UCLA/UW Couples Therapy dataset (because of the data selection bias issue) and the Stanford Lab dataset (because of the use of couple-dependent evaluation) both of which produce inflated results, it is clear that all of the best accuracy results are below 75% with most below 70%. As a reference, the partner-perceived result reported in Boateng et al ([13]) – how well partner A could tell the emotions of their partner B – were 73.2% (male) and 74.3% (female).

Hence, there is more room for improvement to have performance results that are on par with or exceed how well, for example, husbands or wives could tell the emotions of their wives or husbands.

## 8 DISCUSSION: RESEARCH GAP, CHALLENGES AND FUTURE DIRECTION

Despite the contributions of these works, there are still significant research gaps. In this section, we discuss these research gaps, challenges, and future directions in this area of research.

### 8.1 Unexplored Modalities

Only two modalities — acoustic and lexical — have been mostly used for recognition with the visual modality explored only superficially. Several modalities such as physiological data (heart rate, heart rate variability, skin temperature, skin conductance), body and hand gestures using accelerometer, gyroscope, or even the visual modality, and facial expression are unexplored. Additionally, only standard and simple multimodal fusion approaches have been used. More complex fusion approaches such as model-level and hybrid [74] could be explored in the future.

### 8.2 Cross-lingual and Cross-cultural evaluation

None of these works have performed cross-lingual and cross-cultural evaluations. Models in these works have been developed in lingual silos (English, German, and Dutch language speakers) and cultural silos (North Americans, Western Europeans). More effort would be needed to develop and evaluate recognition systems that work across languages since multilingual language models would need to be used. Furthermore, culture affects how people experience and express emotions [65, 84]. Currently, it is not clear how well the recognition systems would work across cultural contexts. Hence, more work is needed to perform these kinds of evaluations as it is important for building systems that are easily generalizable to other contexts.

### 8.3 Intrapersonal and Interpersonal Modeling

Further intrapersonal and interpersonal modeling approaches could be explored. For example, attention mechanisms [99] could be leveraged to automatically learn the salient segments as part of the training process. Also, synchrony measures have only been computed for individual modalities. Computing and using synchrony measures multimodally is a possible future direction. Additionally, more complex dyadic influence modeling could be used such as on a turn-by-turn basis rather than only including the features of the interacting partner as was done in [11]. Other kinds of dyadic dynamics from [23] can be used such as the ratio of both partners' counts of positive and negative words and turn-taking patterns (e.g., the ratio of partners' speaker turn duration, pauses, number of words, etc).

### 8.4 Observed vs Self-Reported Emotion Data

Most of these works have used observed emotion labels from external raters with only a few works using self-reported labels from the partners themselves [11, 13]. One challenge with observed labels is that they are based on the perceptions of external individuals and consequently, do not reflect the subjective emotions of the partners. Though similar, performing recognition of these two groups is distinct and important to be mindful of depending on the downstream use case of the system. For example, if the intended use case and intervention is that partner A shows empathy to partner B based on how partner A is feeling, for example, a recognition system that only looks at emotional behavior or emotional expression will not be the best to use but rather, one that can adequately quantify partner A's subjective emotions. More work is needed to be done using self-reported labels.

## 8.5 Data from Daily Life

Currently, there is no work that has performed emotion recognition using data collected from couples' interactions in uncontrolled settings in daily life. Data from the wild tend to be noisy and could have more potential confounders such as increased heart rate arising from physical activity rather than from high emotional arousal such as stress or anger which could be the likely reason in a controlled lab conversation setting. Hence, the recognition task would be more challenging than the context of the datasets used in these works, which are couples sitting at one place, with limited mobility, and having an 8-10 minute conversation. Consequently, models developed with lab data will likely not perform well on data from daily life. Future work is needed to collect this kind of data and perform recognition with it. This work [9] is a step in that direction.

Though not unique to the context of data from daily life, it is more critical that performance evaluations go beyond the standard accuracy metrics and include detailed error analyses and assessments of conditions under which the model performs poorly. For example, the model might perform poorly when the signal-to-noise ratio is above a certain threshold (as was considered by [6]) or if the transcript has way too few words per speaker turn. The model could be preempted from performing recognition when these conditions are encountered to reduce the likelihood of the model performing poorly. Error analyses would reveal more detailed information such as these. These are key requirements for building robust systems that work using data from uncontrolled settings.

## 8.6 Real-Time Recognition Systems

Furthermore, none of the works surveyed have actually implemented systems that perform real-time recognition either in the lab or in the real world which is the holy grail of a system for couples' emotion recognition. The real test of such a machine learning system is its deployment and evaluation in the contexts in which they are to be used.

Key challenges related to the turn-taking nature of couples' interactions need to be addressed towards accomplishing this goal especially for the two most used modalities — acoustic and lexical. These include having automatic speech preprocessing pipelines — voice activity detection, speaker diarization, and speech recognition systems — that work accurately on the fly without time lags for audio data. Several of the surveyed works used manually annotated data in the preprocessing stage and hence it is a key challenge for future works to address.

Additionally, the recognition algorithm would also need to work accurately and without lag, on the systems that they are deployed on. For ubiquitous devices such as smartphones, smartwatches, or edge devices, the model would have to be small enough to fit on the device and perform computation without hoarding all the compute resources. Most of the works have used simple algorithms such as SVM which will work well for such contexts. But for the deep learning models, there might be potential challenges because of their size and compute requirements. This issue is more pertinent for the lexical modality since current state-of-the-art language models (e.g., BERT) are huge and might be impossible to fit on edge devices in their original form. Various approaches to compress, distill and quantize large models would need to be explored.

## 9 CONCLUSION

In this work, we survey 28 works that have developed and evaluated systems for emotion recognition using data collected from couples' interactions or conversations. Overall, the works in this survey have mostly used one specific data set — UCLA/UW Couples Therapy data. All works have used data collected from lab contexts. Most works used the acoustic modality and the SVM algorithm for binary classification of positive and negative affect. Various multimodal fusion and intrapersonal and interpersonal modeling approaches have been explored. Robust evaluation approaches (e.g., LOCO, LNCO, and 10-fold CV couple disjoint) and metrics (UAR and accuracy with balanced data) have been used. Performance results leave room for improvement. Substantial research gaps remain with several opportunities for future research directions such as exploring more modalities and advanced

fusion approaches, performing cross-lingual and cross-cultural evaluations, leveraging other intrapersonal and interpersonal modeling approaches, using data from daily life, and performing real-time and real-world deployment and evaluation of the recognition system. Insights from this survey would enable future research towards having better couples' emotion recognition system that would enable social and health psychology research and the development of interventions to improve the emotional well-being, relationship quality, and chronic disease management of couples.

## ACKNOWLEDGMENTS

We are grateful to Sandeep Chakravarthula for his feedback on this paper.

## REFERENCES

- [1] Mojtaba Khomami Abadi, Ramanathan Subramanian, Seyed Mostafa Kia, Paolo Avesani, Ioannis Patras, and Nicu Sebe. 2015. DECAF: MEG-based multimodal database for decoding affective physiological responses. *IEEE Transactions on Affective Computing* 6, 3 (2015), 209–222.
- [2] Rummana Bari, Md Mahbubur Rahman, Nazir Saleheen, Megan Battles Parsons, Eugene H Buder, and Santosh Kumar. 2020. Automated Detection of Stressful Conversations Using Wearable Physiological and Inertial Sensors. *Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies* 4, 4 (2020), 1–23.
- [3] Ellen Berscheid and Hilary Ammazzalorso. 2001. Emotional experience in close relationships. *Blackwell handbook of social psychology: Interpersonal processes* (2001), 308–330.
- [4] Alberto Betella and Paul FMJ Verschure. 2016. The affective slider: A digital self-assessment scale for the measurement of human emotions. *PLoS one* 11, 2 (2016), e0148037.
- [5] Jacopo Biggio, George Boateng, Peter Hilpert, Matthew Vowels, Guy Bodenmann, Mona Neysari, Fridtjof Nussbeck, and Tobias Kowatsch. 2021. *BERT Meets LIWC: Exploring State-of-the-Art Language Models for Predicting Communication Behavior in Couples' Conflict Interactions*. Association for Computing Machinery, New York, NY, USA, 385–389. <https://doi.org/10.1145/3461615.3485423>
- [6] Matthew Black, Athanasios Katsamanis, Chi-Chun Lee, Adam C Lammert, Brian R Baucom, Andrew Christensen, Panayiotis G Georgiou, and Shrikanth S Narayanan. 2010. Automatic classification of married couples' behavior using audio features. In *Eleventh annual conference of the international speech communication association*.
- [7] Matthew P Black, Panayiotis G Georgiou, Athanasios Katsamanis, Brian R Baucom, and Shrikanth Narayanan. 2011. "You made me do it": Classification of Blame in Married Couples' Interactions by Fusing Automatically Derived Speech and Language Information. In *Twelfth Annual Conference of the International Speech Communication Association*.
- [8] Matthew P Black, Athanasios Katsamanis, Brian R Baucom, Chi-Chun Lee, Adam C Lammert, Andrew Christensen, Panayiotis G Georgiou, and Shrikanth S Narayanan. 2013. Toward automating a human behavioral coding system for married couples' interactions using speech acoustic features. *Speech communication* 55, 1 (2013), 1–21.
- [9] George Boateng. 2020. Towards Real-Time Multimodal Emotion Recognition among Couples. In *Proceedings of the 2020 International Conference on Multimodal Interaction*. 748–753.
- [10] George Boateng, John A Batsis, Ryan Halter, and David Kotz. 2017. ActivityAware: an app for real-time daily activity level monitoring on the amulet wrist-worn device. In *2017 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*. IEEE, 431–435.
- [11] George Boateng, Peter Hilpert, Guy Bodenmann, Mona Neysari, and Tobias Kowatsch. 2021. "You Made Me Feel This Way": Investigating Partners' Influence in Predicting Emotions in Couples' Conflict Interactions Using Speech Data. In *Companion Publication of the 2021 International Conference on Multimodal Interaction (Montreal, QC, Canada) (ICMI '21 Companion)*. Association for Computing Machinery, New York, NY, USA, 390–394. <https://doi.org/10.1145/3461615.3485424>
- [12] George Boateng, Janina Lüscher, Urte Scholz, and Tobias Kowatsch. 2020. Emotion Capture among Real Couples in Everyday Life. In *1st Momentary Emotion Elicitation & Capture workshop (MEEC 2020), co-located with the ACM CHI Conference on Human Factors in Computing Systems*.
- [13] George Boateng, Laura Sels, Peter Kuppens, Peter Hilpert, Urte Scholz, and Tobias Kowatsch. 2020. Speech Emotion Recognition among Couples using the Peak-End Rule and Transfer Learning. In *Companion Publication of the 2020 International Conference on Multimodal Interaction (ICMI '20 Companion), October 25–29, 2020, Virtual event, Netherlands*.
- [14] Paul Boersma. 2001. Praat, a system for doing phonetics by computer. *Glottol. Int.* 5, 9 (2001), 341–345.
- [15] Niall Bolger and David Amarel. 2007. Effects of social support visibility on adjustment to stress: Experimental evidence. *Journal of personality and social psychology* 92, 3 (2007), 458.

- [16] Leslie R Brody. 1993. On understanding gender differences in the expression of emotion. *Human feelings: Explorations in affect development and meaning* (1993), 87–121.
- [17] Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeannette N Chang, Sungbok Lee, and Shrikanth S Narayanan. 2008. IEMOCAP: Interactive emotional dyadic motion capture database. *Language resources and evaluation* 42, 4 (2008), 335.
- [18] Carlos Busso, Srinivas Parthasarathy, Alec Burmania, Mohammed AbdelWahab, Najmeh Sadoughi, and Emily Mower Provost. 2016. MSP-IMPROV: An acted corpus of dyadic interactions to study emotion perception. *IEEE Transactions on Affective Computing* 8, 1 (2016), 67–80.
- [19] Laura L Carstensen, John M Gottman, and Robert W Levenson. 1995. Emotional behavior in long-term marriage. *Psychology and aging* 10, 1 (1995), 140.
- [20] Sandeep Nallan Chakravarthula, Brian RW Baucom, Shrikanth Narayanan, and Panayiotis Georgiou. 2021. An analysis of observation length requirements for machine understanding of human behaviors from spoken language. *Computer Speech & Language* 66 (2021), 101162.
- [21] Sandeep Nallan Chakravarthula, Rahul Gupta, Brian Baucom, and Panayiotis Georgiou. 2015. A language-based generative model framework for behavioral analysis of couples' therapy. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2090–2094.
- [22] Sandeep Nallan Chakravarthula, Haoqi Li, Shao-Yen Tseng, Maija Reblin, and Panayiotis Georgiou. 2019. Predicting Behavior in Cancer-Afflicted Patient and Spouse Interactions Using Speech and Language. *Proc. Interspeech 2019* (2019), 3073–3077.
- [23] Sandeep Nallan Chakravarthula, Md Nasir, Shao-Yen Tseng, Haoqi Li, Tae Jin Park, Brian Baucom, Craig J Bryan, Shrikanth Narayanan, and Panayiotis Georgiou. 2020. Automatic Prediction of Suicidal Risk in Military Couples Using Multimodal Interaction Cues from Couples Conversations. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 6539–6543.
- [24] Andrew Christensen, David C Atkins, Sara Berns, Jennifer Wheeler, Donald H Baucom, and Lorelei E Simpson. 2004. Traditional versus integrative behavioral couple therapy for significantly and chronically distressed married couples. *Journal of consulting and clinical psychology* 72, 2 (2004), 176.
- [25] James A Coan and John M Gottman. 2007. The specific affect coding system (SPAFF). *Handbook of emotion elicitation and assessment* (2007), 267–285.
- [26] Colleen E Crangle, Rui Wang, Marcos Perreau-Guimaraes, Michelle U Nguyen, Duc T Nguyen, and Patrick Suppes. 2019. Machine learning for the recognition of emotion in the speech of couples in psychotherapy using the Stanford Suppes Brain Lab Psychotherapy Dataset. *arXiv preprint arXiv:1901.04110* (2019).
- [27] Egon Dejonckheere, Merijn Mestdagh, Marlies Houben, Isa Rutten, Laura Sels, Peter Kuppens, and Francis Tuerlinckx. 2019. Complex affect dynamics add limited information to the prediction of psychological well-being. *Nature human behaviour* 3, 5 (2019), 478–491.
- [28] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [29] Thomas G Dietterich, Richard H Lathrop, and Tomás Lozano-Pérez. 1997. Solving the multiple instance problem with axis-parallel rectangles. *Artificial intelligence* 89, 1-2 (1997), 31–71.
- [30] Sidney K D'mello and Jacqueline Kory. 2015. A review and meta-analysis of multimodal affect detection systems. *ACM Computing Surveys (CSUR)* 47, 3 (2015), 1–36.
- [31] Paul Ekman and Dacher Keltner. 1997. Universal facial expressions of emotion. *Seegerstrale U, P. Molnar P, eds. Nonverbal communication: Where nature meets culture* (1997), 27–46.
- [32] Florian Eyben, Klaus R Scherer, Björn W Schuller, Johan Sundberg, Elisabeth André, Carlos Busso, Laurence Y Devillers, Julien Epps, Petri Laukka, Shrikanth S Narayanan, et al. 2015. The Geneva minimalistic acoustic parameter set (GeMAPS) for voice research and affective computing. *IEEE transactions on affective computing* 7, 2 (2015), 190–202.
- [33] Florian Eyben, Martin Wöllmer, and Björn Schuller. 2010. Opensmile: the munich versatile and fast open-source audio feature extractor. In *Proceedings of the 18th ACM international conference on Multimedia*. 1459–1462.
- [34] Allison K Farrell, Ledina Imami, Sarah CE Stanton, and Richard B Slatcher. 2018. Affective processes as mediators of links between close relationships and physical health. *Social and Personality Psychology Compass* 12, 7 (2018), e12408.
- [35] Kexin Feng and Theodora Chaspari. 2020. A Review of Generalizable Transfer Learning in Automatic Emotion Recognition. *Frontiers in Computer Science* 2 (2020), 9.
- [36] Barbara L Fredrickson. 2000. Extracting meaning from past affective experiences: The importance of peaks, ends, and specific emotions. *Cognition & Emotion* 14, 4 (2000), 577–606.
- [37] Panayiotis G Georgiou, Matthew P Black, Adam C Lammert, Brian R Baucom, and Shrikanth S Narayanan. 2011. “That’s aggravating, very aggravating”: is it possible to classify behaviors in couple interactions using automatically derived lexical features?. In *International Conference on Affective Computing and Intelligent Interaction*. Springer, 87–96.



- [38] James Gibson, Athanasios Katsamanis, Matthew P Black, and Shrikanth Narayanan. 2011. Automatic identification of salient acoustic instances in couples' behavioral interactions using diverse density support vector machines. In *Twelfth Annual Conference of the International Speech Communication Association*.
- [39] James Gibson, Athanasios Katsamanis, Francisco Romero, Bo Xiao, Panayiotis Georgiou, and Shrikanth Narayanan. 2015. Multiple instance learning for behavioral coding. *IEEE Transactions on Affective Computing* 8, 1 (2015), 81–94.
- [40] John Mordechai Gottman. 2005. *The mathematics of marriage: Dynamic nonlinear models*. MIT Press.
- [41] John Mordechai Gottman. 2014. *What predicts divorce?: The relationship between marital processes and marital outcomes*. Psychology Press.
- [42] John M Gottman and Robert W Levenson. 1985. A valid procedure for obtaining self-report of affect in marital interaction. *Journal of consulting and clinical psychology* 53, 2 (1985), 151.
- [43] Michael Grimm, Kristian Kroschel, and Shrikanth Narayanan. 2008. The Vera am Mittag German audio-visual emotional speech database. In *2008 IEEE international conference on multimedia and expo*. IEEE, 865–868.
- [44] C Heavey, D Gill, and A Christensen. [n.d.]. Couples interaction rating system 2 (CIRS2), University of California, Los Angeles, 2002. *Unpublished manuscript*. [Google Scholar] ([n. d.]).
- [45] Christopher L Heavey, Andrew Christensen, and Neil M Malamuth. 1995. The longitudinal impact of demand and withdrawal during marital conflict. *Journal of consulting and clinical psychology* 63, 5 (1995), 797.
- [46] Richard L Heinrich, Cyndie Coscarelli Schag, and Patricia A Ganz. 1984. Living with cancer: The cancer inventory of problem situations. *Journal of clinical psychology* 40, 4 (1984), 972–980.
- [47] Richard E Heyman. 2001. Observation of couple conflicts: Clinical assessment applications, stubborn truths, and shaky foundations. *Psychological assessment* 13, 1 (2001), 5.
- [48] Richard E Heyman. 2004. Rapid marital interaction coding system (RMICS). In *Couple observational coding systems*. Routledge, 81–108.
- [49] Richard E Heyman, Robert L Weiss, and J Mark Eddy. 1995. Marital interaction coding system: Revision and empirical evaluation. *Behaviour research and therapy* 33, 6 (1995), 737–746.
- [50] Masumi Iida, Mary Ann Parris Stephens, Karen S Rook, Melissa M Franks, and James K Salem. 2010. When the going gets tough, does support get going? Determinants of spousal support provision to type 2 diabetic patients. *Personality and Social Psychology Bulletin* 36, 6 (2010), 780–791.
- [51] J Jones and A Christensen. 1998. Couples interaction study: Social support interaction rating system. *University of California, Los Angeles* 7 (1998).
- [52] Athanasios Katsamanis, James Gibson, Matthew P Black, and Shrikanth S Narayanan. 2011. Multiple instance learning for classification of human behavior observations. In *International Conference on Affective Computing and Intelligent Interaction*. Springer, 145–154.
- [53] Patricia K Kerig and Donald H Baucom. 2004. *Couple observational coding systems*. Taylor & Francis.
- [54] Dana Ketcher, Casidee Thompson, Amy K Otto, Maija Reblin, Kristin G Cloyes, Margaret F Clayton, Brian RW Baucom, and Lee Ellington. 2021. The Me in We dyadic communication intervention is feasible and acceptable among advanced cancer patients and their family caregivers. *Palliative Medicine* 35, 2 (2021), 389–396.
- [55] Monika Kuster, Katharina Bernecker, Sabine Backes, Veronika Brandstätter, Fridtjof W Nussbeck, Thomas N Bradbury, Mike Martin, Dorothee Sutter-Stickel, and Guy Bodenmann. 2015. Avoidance orientation and the escalation of negative communication in intimate relationships. *Journal of Personality and Social Psychology* 109, 2 (2015), 262.
- [56] Chi-Chun Lee, Matthew Black, Athanasios Katsamanis, Adam C Lammert, Brian R Baucom, Andrew Christensen, Panayiotis G Georgiou, and Shrikanth S Narayanan. 2010. Quantification of prosodic entrainment in affective spontaneous spoken interactions of married couples. In *Eleventh Annual Conference of the International Speech Communication Association*.
- [57] Chi-Chun Lee, Athanasios Katsamanis, Matthew P Black, Brian R Baucom, Andrew Christensen, Panayiotis G Georgiou, and Shrikanth S Narayanan. 2014. Computing vocal entrainment: A signal-derived PCA-based quantification scheme with application to affect analysis in married couple interactions. *Computer Speech & Language* 28, 2 (2014), 518–539.
- [58] Chi-Chun Lee, Athanasios Katsamanis, Matthew P Black, Brian R Baucom, Panayiotis G Georgiou, and Shrikanth Narayanan. 2011. An analysis of PCA-based vocal entrainment measures in married couples' affective spoken interactions. In *Twelfth Annual Conference of the International Speech Communication Association*.
- [59] Chi-Chun Lee, Athanasios Katsamanis, Matthew P Black, Brian R Baucom, Panayiotis G Georgiou, and Shrikanth S Narayanan. 2011. Affective state recognition in married couples' interactions using PCA-based vocal entrainment measures with multiple instance learning. In *International Conference on Affective Computing and Intelligent Interaction*. Springer, 31–41.
- [60] Chi-Chun Lee, Athanasios Katsamanis, Panayiotis G Georgiou, and Shrikanth S Narayanan. 2012. Based on isolated saliency or causal integration? toward a better understanding of human annotation process using multiple instance learning and sequential probability ratio test. In *Thirteenth Annual Conference of the International Speech Communication Association*.
- [61] Haoqi Li, Brian Baucom, and Panayiotis Georgiou. 2016. Sparsely Connected and Disjointly Trained Deep Neural Networks for Low Resource Behavioral Annotation: Acoustic Classification in Couples' Therapy. In *Proc. Interspeech 2016*. 1407–1411. <https://doi.org/10.21437/Interspeech.2016-1217>

- [62] Haoqi Li, Brian Baucom, and Panayiotis Georgiou. 2017. Unsupervised latent behavior manifold learning from acoustic features: Audio2behavior. In *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 5620–5624.
- [63] Haoqi Li, Brian Baucom, and Panayiotis Georgiou. 2020. Linking emotions to behaviors through deep transfer learning. *PeerJ Computer Science* 6 (2020), e246.
- [64] Janina Lüscher, Tobias Kowatsch, George Boateng, Prabhakaran Santhanam, Guy Bodenmann, and Urte Scholz. 2019. Social Support and Common Dyadic Coping in Couples' Dyadic Management of Type II Diabetes: Protocol for an Ambulatory Assessment Application. *JMIR research protocols* 8, 10 (2019), e13685.
- [65] David Matsumoto and Paul Ekman. 1989. American-Japanese cultural differences in intensity ratings of facial expressions of emotion. *Motivation and Emotion* 13, 2 (1989), 143–157.
- [66] Angeliki Metallinou, Chi-Chun Lee, Carlos Busso, Sharon Carnicke, Shrikanth Narayanan, et al. 2010. The USC CreativeIT database: A multimodal database of theatrical improvisation. *Multimodal Corpora: Advances in Capturing, Coding and Analyzing Multimodality* (2010), 55.
- [67] Angeliki Metallinou and Shrikanth Narayanan. 2013. Annotation and processing of continuous emotional attributes: Challenges and opportunities. In *2013 10th IEEE international conference and workshops on automatic face and gesture recognition (FG)*. IEEE, 1–8.
- [68] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*. 3111–3119.
- [69] Philipp M Müller, Sikandar Amin, Prateek Verma, Mykhaylo Andriluka, and Andreas Bulling. 2015. Emotion recognition from embedded bodily expressions and speech during dyadic interactions. In *2015 International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, 663–669.
- [70] Sandeep Nallan Chakravarthula, Brian Baucom, and Panayiotis Georgiou. 2018. Modeling Interpersonal Influence of Verbal Behavior in Couples Therapy Dyadic Interactions. In *Proc. Interspeech 2018*. 2339–2343. <https://doi.org/10.21437/Interspeech.2018-1562>
- [71] Cheul Young Park, Narae Cha, Soowon Kang, Auk Kim, Ahsan Habib Khandoker, Leontios Hadjileontiadis, Alice Oh, Yong Jeong, and Uichin Lee. 2020. K-EmoCon, a multimodal sensor dataset for continuous emotion recognition in naturalistic conversations. *Scientific Data* 7, 1 (2020), 1–16.
- [72] James W Pennebaker, Martha E Francis, and Roger J Booth. 2001. Linguistic inquiry and word count: LIWC 2001. *Mahway: Lawrence Erlbaum Associates* 71, 2001 (2001), 2001.
- [73] Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of NAACL-HLT*. 2227–2237.
- [74] Soujanya Poria, Erik Cambria, Rajiv Bajpai, and Amir Hussain. 2017. A review of affective computing: From unimodal analysis to multimodal fusion. *Information Fusion* 37 (2017), 98–125.
- [75] Soujanya Poria, Navonil Majumder, Rada Mihalcea, and Eduard Hovy. 2019. Emotion recognition in conversation: Research challenges, datasets, and recent advances. *IEEE Access* 7 (2019), 100943–100953.
- [76] Gabriele Prati and Luca Pietrantonio. 2010. The relation of perceived and received social support to mental health among first responders: a meta-analytic review. *Journal of Community Psychology* 38, 3 (2010), 403–417.
- [77] Maija Reblin, Richard E Heyman, Lee Ellington, Brian RW Baucom, Panayiotis G Georgiou, and Susan T Vadaparampil. 2018. Everyday couples' communication research: Overcoming methodological barriers with technology. *Patient education and counseling* 101, 3 (2018), 551–556.
- [78] Maija Reblin, Steven K Sutton, Susan T Vadaparampil, Richard E Heyman, and Lee Ellington. 2019. Behind closed doors: How advanced cancer couples communicate at home. *Journal of psychosocial oncology* 37, 2 (2019), 228–241.
- [79] Megan L Robbins, Ana María López, Karen L Weihs, and Matthias R Mehl. 2014. Cancer conversations in context: naturalistic observation of couples coping with breast cancer. *Journal of Family Psychology* 28, 3 (2014), 380.
- [80] Nicole A Roberts, Jeanne L Tsai, and James A Coan. 2007. Emotion elicitation using dyadic interaction tasks. *Handbook of emotion elicitation and assessment* (2007), 106–123.
- [81] Anna Marie Ruef and Robert W Levenson. 2007. Continuous measurement of emotion. *Handbook of emotion elicitation and assessment* (2007), 286–297.
- [82] James A Russell. 1980. A circumplex model of affect. *Journal of personality and social psychology* 39, 6 (1980), 1161.
- [83] James A Russell, Anna Weiss, and Gerald A Mendelsohn. 1989. Affect grid: a single-item scale of pleasure and arousal. *Journal of personality and social psychology* 57, 3 (1989), 493.
- [84] K. R Scherer, H Wallbott, D Matsumoto, and K Tsutomu. 1988. Emotional experience in cultural context: A comparison between Europe, Japan and the United States. *Faces of emotion: recent research* (1988), 98–115.
- [85] Philip Schmidt, Attila Reiss, Robert Dürichen, Claus Marberger, and Kristof Van Laerhoven. 2018. Introducing wesad, a multimodal dataset for wearable stress and affect detection. In *Proceedings of the 20th ACM international conference on multimodal interaction*. 400–408.
- [86] Philip Schmidt, Attila Reiss, Robert Dürichen, and Kristof Van Laerhoven. 2019. Wearable-Based Affect Recognition—A Review. *Sensors* 19, 19 (2019), 4079.

- [87] Laura Sels, Jed Cabrieto, Emily Butler, Harry Reis, Eva Ceulemans, and Peter Kuppens. 2019. The occurrence and correlates of emotional interdependence in romantic relationships. *Journal of personality and social psychology* (2019).
- [88] Laura Sels, Eva Ceulemans, and Peter Kuppens. 2019. All's well that ends well? A test of the peak-end rule in couples' conflict discussions. *European Journal of Social Psychology* 49, 4 (2019), 794–806.
- [89] Laura Sels, Yan Ruan, Peter Kuppens, Eva Ceulemans, and Harry Reis. 2020. Actual and perceived emotional similarity in couples' daily lives. *Social Psychological and Personality Science* 11, 2 (2020), 266–275.
- [90] Timothy W Smith and Karen Weihs. 2019. Emotion, social relationships, and physical health: concepts, methods, and evidence for an integrative perspective. *Psychosomatic medicine* 81, 8 (2019), 681–693.
- [91] Rolf Steyer, Peter Schwenkmezger, Peter Notz, and Michael Eid. 1997. Der Mehrdimensionale Befindlichkeitsfragebogen MDBF [Multidimensional mood questionnaire]. *Göttingen, Germany: Hogrefe* (1997).
- [92] Adela C Timmons, Theodora Chaspari, Sohyun C Han, Laura Perrone, Shrikanth S Narayanan, and Gayla Margolin. 2017. Using multimodal wearable technology to detect conflict among couples. *Computer* 50, 3 (2017), 50–59.
- [93] Shao-Yen Tseng, Brian Baucom, and Panayiotis Georgiou. 2019. Unsupervised online multitask learning of behavioral sentence embeddings. *PeerJ Computer Science* 5 (2019), e200.
- [94] Shao-Yen Tseng, Brian R Baucom, and Panayiotis G Georgiou. 2017. Approaching Human Performance in Behavior Estimation in Couples Therapy Using Deep Sentence Embeddings.. In *INTERSPEECH*. 3291–3295.
- [95] Shao-Yen Tseng, Sandeep Nallan Chakravarthula, Brian R Baucom, and Panayiotis G Georgiou. 2016. Couples Behavior Modeling and Annotation Using Low-Resource LSTM Language Models.. In *INTERSPEECH*. 898–902.
- [96] Shao-Yen Tseng, Haoqi Li, Brian Baucom, and Panayiotis Georgiou. 2018. "Honey, I Learned to Talk" Multimodal Fusion for Behavior Analysis. In *Proceedings of the 20th ACM International Conference on Multimodal Interaction*. 239–243.
- [97] Maximiliane Uhlich and Daniel Bojar. 2021. DeepConnection: classifying momentary relationship state from images of romantic couples. *Journal of Computational Social Science* (2021), 1–23.
- [98] UZH. University of Zurich. 2020. PASEZ Project- Impact of stress on relationship development of couples and children. <http://www.dynage.uzh.ch/en/newsevents/news/news25.html>
- [99] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*. 5998–6008.
- [100] David Watson, Lee Anna Clark, and Auke Tellegen. 1988. Development and validation of brief measures of positive and negative affect: the PANAS scales. *Journal of personality and social psychology* 54, 6 (1988), 1063.
- [101] Wei Xia, James Gibson, Bo Xiao, Brian Baucom, and Panayiotis G Georgiou. 2015. A dynamic model for behavioral analysis of couple interactions using acoustic features. In *Sixteenth Annual Conference of the International Speech Communication Association*.
- [102] Bo Xiao, Panayiotis Georgiou, Brian Baucom, and Shrikanth S Narayanan. 2015. Head motion modeling for human behavior analysis in dyadic interaction. *IEEE transactions on multimedia* 17, 7 (2015), 1107–1119.
- [103] Zhihong Zeng, Maja Pantic, Glenn I Roisman, and Thomas S Huang. 2008. A survey of affect recognition methods: Audio, visual, and spontaneous expressions. *IEEE transactions on pattern analysis and machine intelligence* 31, 1 (2008), 39–58.
- [104] Sebastian Zepf, Javier Hernandez, Alexander Schmitt, Wolfgang Minker, and Rosalind W Picard. 2020. Driver emotion recognition for intelligent vehicles: a survey. *ACM Computing Surveys (CSUR)* 53, 3 (2020), 1–30.

## Paper 2: Peak-End Rule

# Speech Emotion Recognition among Couples using the Peak-End Rule and Transfer Learning

GEORGE BOATENG, ETH Zürich, Switzerland

LAURA SELS, Ghent University, Belgium

PETER KUPPENS, KU Leuven, Belgium

PETER HILPERT, University of Surrey, United Kingdom

TOBIAS KOWATSCH, ETH Zürich, Switzerland and University of St. Gallen, Switzerland

Extensive couples' literature shows that how couples feel after a conflict is predicted by certain emotional aspects of that conversation. Understanding the emotions of couples leads to a better understanding of partners' mental well-being and consequently their relationships. Hence, automatic emotion recognition among couples could potentially guide interventions to help couples improve their emotional well-being and their relationships. It has been shown that people's global emotional judgment after an experience is strongly influenced by the emotional extremes and ending of that experience, known as the peak-end rule. In this work, we leveraged this theory and used machine learning to investigate, which audio segments can be used to best predict the end-of-conversation emotions of couples. We used speech data collected from 101 Dutch-speaking couples in Belgium who engaged in 10-minute long conversations in the lab. We extracted acoustic features from (1) the audio segments with the most extreme positive and negative ratings, and (2) the ending of the audio. We used transfer learning in which we extracted these acoustic features with a pre-trained convolutional neural network (YAMNet). We then used these features to train machine learning models – support vector machines – to predict the end-of-conversation valence ratings (positive vs negative) of each partner. From our results (balanced accuracy), the segments from the peak were the best for recognizing the emotions of female partners and outperformed male partners' perception of their female partners' emotions. The results of this work could inform how to best recognize the emotions of couples after-conversation sessions and eventually, lead to a better understanding of couples' relationships either in therapy or in everyday life.

CCS Concepts: • **Applied computing** → **Psychology**.

Additional Key Words and Phrases: Speech emotion recognition; Speech processing; Affective computing; Couples; Transfer Learning; Peak-end rule; Convolutional neural network; Support vector machine

## ACM Reference Format:

George Boateng, Laura Sels, Peter Kuppens, Peter Hilpert, and Tobias Kowatsch. 2020. Speech Emotion Recognition among Couples using the Peak-End Rule and Transfer Learning. In *Companion Publication of the 2020 International Conference on Multimodal Interaction (ICMI '20 Companion)*, October 25–29, 2020, Virtual event, Netherlands. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3395035.3425253>

## 1 INTRODUCTION

Couples' observation research has shown that the emotions that couples experience during a conflict predict if these couples stay together in the long-term (for an overview, see [19]). For instance, couples heading for break-up show more negative emotions and less positive emotions than happy couples, and are stuck in certain

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*ICMI '20 Companion, October 25–29, 2020, Virtual event, Netherlands*

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8002-7/20/10...\$15.00

<https://doi.org/10.1145/3395035.3425253>

emotional patterns [7, 18]. Although couples' observation research has delivered valuable clinical insights, it also suffers from measurement issues such as low cross-validity and interrater reliability [23] and entails some methodological challenges. One important methodological challenge is the manual coding of audio-video data, which is very costly and time-consuming [27]. Automated emotion recognition could alleviate these limitations, and therefore advance the field in important ways [36].

Several emotion recognition works on couple dyads use data that is collected from individuals acting out dyadic interactions either using a script or engaging in spontaneous sessions [5, 6, 32, 34]. A lot of emotion recognition works use such data sets [38]. The emotions are later rated by others amidst several challenges [33] and do not necessarily reflect the subjective emotions of the individuals. Additionally, these algorithms are likely to perform poorly on naturalistic data [10].

On the other hand, there are few works on detecting the emotional behavior of real couples. Some leveraged interaction dynamics among the partners (e.g., entrainment – synchrony between partners) [2, 28, 29] and salient instances [16, 17, 26] to perform recognition. These works tend to use emotion labels from external raters rather than the couples and hence do not reflect the subjective emotions of the couples.

Our aim is to build upon recent findings from fundamental psychological research to automatically recognize couples' self-reported emotions. Specifically, couples literature has shown that how couples feel after a conflict is predicted by certain emotional aspects of that conversation (e.g., [13, 14, 21, 30, 31]); and recently, it has been suggested that the emotional extremes and ending of the conversation might be particularly valuable [44]. In fact, in a variety of domains, it has been shown that judgments of emotional experiences are most impacted by the most extreme moments (peaks) and the end of that particular experience, known as the peak-end rule [12, 25]. The peak-end rule could be leveraged to develop systems to better recognize the emotions of couples.

Building upon our recommendations in [4], we investigate through a machine learning perspective which segment(s) of an audio conversation could be used to best recognize the emotions of each partner after a conversation. Our research question is as follows:

*Using features of which of the following audio segments produce the best emotion recognition result: a) segments with the most extreme positive and negative ratings, b) the ending of the audio or c) a combination of the extremes and ending?*

In this first of its kind work, our primary contribution is the exploration of the best way to recognize the emotions of couples after a conversation (5 - 10 minutes) through the peak-end rule lens using deep learning approaches. Our secondary contribution is the use of a unique dataset – real-world data collected from Dutch-speaking couples with self-ratings of emotions. Our third contribution is our proposal and computation of a "partner perception baseline" for emotion recognition within the context of couples' interactions that leverage each partner's perception of his/her partner's emotions.

We classified the end-of-conversation valence (positive vs negative) of Dutch-speaking couples using acoustic features from various segments of the audio and compared with the partner perception baseline. We used transfer learning, an approach used in deep learning to circumvent the need to develop hand-crafted features [11]. It is used to address the limitations of using small labeled datasets and has shown success in various fields including emotion recognition tasks ([35, 42]). The results of this work would inform the best way to recognize the emotions of couples' after-conversation sessions and eventually, lead to a better understanding of couples' relationships either in therapy or in everyday life.

The rest of this paper is organized as follows: In Section 2, we describe our method. In Section 3, we describe our experiments. In Section 4, we show and discuss the results. In Section 5, we present limitations of this work and future work, and we conclude in Section 6.

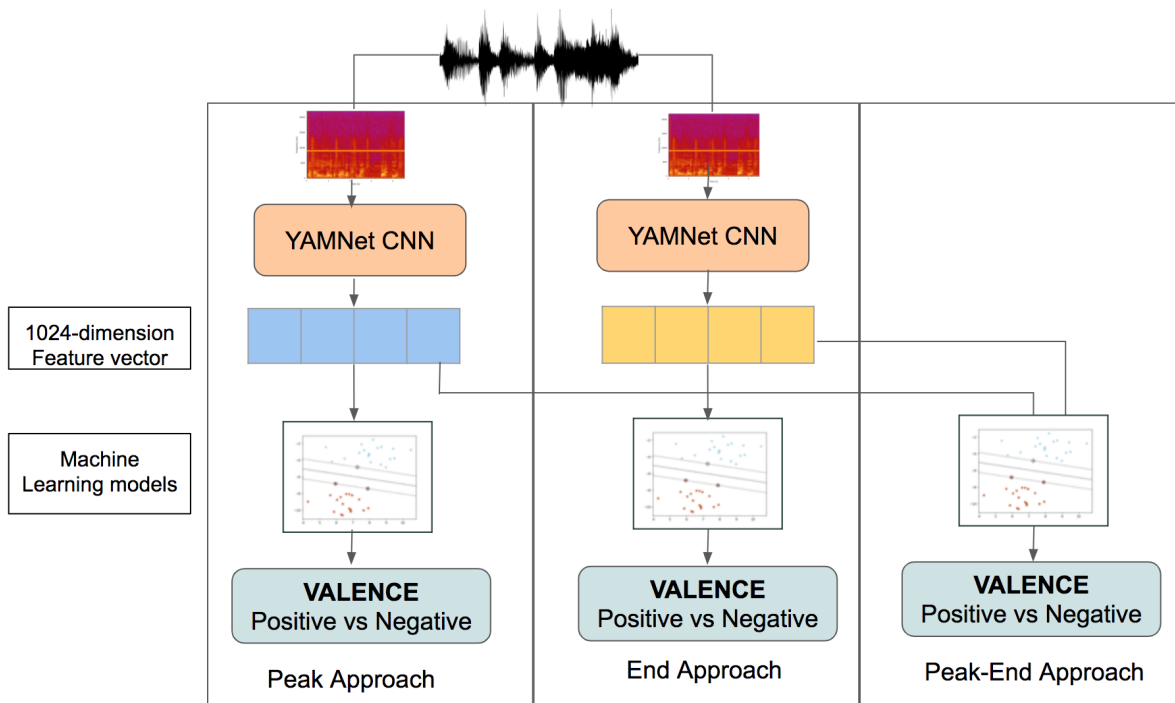


Fig. 1. Overview of Approach

## 2 METHODS

In this section, we describe the dataset and preprocessing, and the transfer learning approach (Figure 1).

### 2.1 Dataset and Preprocessing

A Dyadic Interaction lab study was conducted in Belgium with 101 Dutch-speaking, heterosexual couples. These couples were first asked to have a 10-minute conversation about a negative topic (a characteristic of their partner that annoys them the most), followed by a 10-minute conversation about a positive topic (a characteristic of their partner that they value the most) [9, 43–45]. During both conversations, couples were asked to wrap up the conversation after 8 minutes. For the negative topic, they were also asked to end on good terms. After each conversation, each partner watched the video recording of the conversation separately on a computer and rated his or her emotion on a moment-by-moment basis by continuously adjusting a joystick to the left (very negative) and the right (very positive), so that it closely matched their feelings, resulting in valence scores on a continuous scale from -1 to 1 [20, 39]. Additionally, each partner reported how they felt after the interaction and also what they thought their partner felt, using the Affect Grid questionnaire [41]. The Affect Grid captures the valence and arousal dimensions of Russell’s circumplex model of emotions [40].

Valence refers to how negative to positive the person feels and arousal refers to how sleepy to active a person feels. Using these two dimensions, categorical emotions can be placed and grouped into the four quadrants: high arousal and negative valence (e.g. stressed), low arousal and negative valence (e.g. depressed), low arousal and positive valence (e.g. relaxed) and high arousal and positive valence (e.g. excited). Subjects had to place an ‘x’ on any square on the Affect Grid corresponding to their feelings about each conversation, which translates to a

value of between 0 and 8 each for pleasure and arousal. We only used the valence dimension of the Affect Grid because the continuous rating that the end-of-conversation emotion was compared with was done only using valence. The continuous rating was restricted to valence to minimize the time spent by subjects in the lab and also because it is standard practice in such dyadic interaction designs. We categorized the valence scores into two classes, negative (0-4) and positive valence (5-8) for males and females. Also, we only used audios from the negative/conflict conversation in this work. We could use only 92 out of the 101 audios in this work as some of the data was unavailable due to several issues peculiar of real-world data collection such as missing self-ratings due to failure of the recording device, lack of speaker annotations for all couples among others. In total, for males, we had 22 negative and 70 positive ratings and for females, we had 16 negative and 76 positive ratings. This distribution shows how significantly imbalanced the dataset is which is reflective of real-world data and consistent with other couple emotion recognition works (e.g. [8]).

The audio was manually annotated showing which partner was speaking at various points of the audio. Trained research assistants (5) were instructed to listen and visually inspect the audios, and annotate the exact start and end of each talking turn for each partner. In addition, students coded pauses, cross-talk, and noise and laughter. Multiple rounds of checking were done to ensure this process was precisely done. We used the segments of the audio where the male or female spoke to extract audio segments corresponding to the peaks and ends for each partner. For the peaks, we used the continuous valence rating to find the specific second with the largest negative value (minimum) and the specific second with the largest positive value (maximum). We then used the speaker turn containing that specific second as the peak segment (each for the minimum and maximum). The average duration of the peak segments for all the couples was 3.5 seconds. For the ending, we used the last 60 seconds of the audio corresponding to 10% of the whole audio (600 seconds). There was no reference in the literature for the duration to use for the end and so we picked 60 secs (the last 10%) as we reasoned it will capture a good enough duration of each couple's interaction without being too long.

Finally, we computed a partner perception baseline for the context of emotion recognition among couples. We used the assessment of each partner's perception of his/her partner's emotion at the end of the conversation to compute the baseline. This baseline gives an estimate of how well each partner could infer his/her partner's emotion after an interaction. We argue this is a good enough human baseline with which to compare the machine learning approach since a person's partner, in theory, is the best person to know him or her albeit this perception is biased in practice [46].

## 2.2 Transfer Learning Approach

Given that the data set is small, we sought to leverage work that has been done for a related task and hence used transfer learning [35] where we used a model that is pre-trained on a similar problem. We extracted spectrograms and used a pretrained convolutional neural network (CNN) to compute embeddings as acoustic features which we used to perform classification with machine learning models. We used the YAMNet model [1] which is a CNN that was pretrained on the AudioSet dataset to predict 521 audio event classes [15, 22]. YAMNet is based on the MobileNet architecture [24]. We used the YAMNet model as a feature extractor and hence replaced the original final logistic layer which outputs 521 class with a linear support vector machine (SVM) which we trained.

We extracted a spectrogram as an input into the YAMNet model in the same way as was done for the trained model. The audio's sample rate is 16 kHz. A spectrogram is computed using magnitudes of the Short-Time Fourier Transform with a window size of 25 ms, a window hop of 10 ms, and a periodic Hann window. A mel spectrogram is computed by mapping the spectrogram to 64 mel bins covering the range 125-7500 Hz. A stabilized log mel spectrogram is computed by applying  $\log(\text{mel-spectrum} + 0.01)$  where the offset is used to avoid taking a logarithm of zero. These features are then framed into non-overlapping examples of 0.96 seconds, where each example covers 64 mel bands and 96 frames of 10 ms each [1]. This resulted in a 2D data of size 96 x 64 for each



Table 1. Results for Peak, End and Peak-End Approaches and Baseline

Approach	Balanced Accuracy (%)	
	Male	Female
Partner perception	73.2	74.3
Peak	48.8	74.8
End	50	58.6
Peak-End	53.3	54.4

second, which we used as a data point input to the YAMNet model. The output of the model is a 1024-dimensional feature vector per data point input of size 96 x 64. We then normalized the vectors to be zero mean and unit variance and then used the features vectors as inputs to a linear SVM.

### 3 EXPERIMENTS

We performed various experiments using a linear SVM and the scikit-learn library [37]. We trained models separately for males and females to perform binary classification of valence. Our main models were trained using features from the peak, end, and peak and end (peak-end). We used majority voting of the classification for all features to decide the class for the audio segment. We performed an evaluation with leave-one-couple-out cross-validation similar to [8] which is a robust evaluation approach and gives an estimate of how well the model will perform on an unseen couple. We used confusion matrices and the metric balanced accuracy for evaluation since the data is imbalanced. Balanced accuracy is the unweighted average of the recall of each class. We used different values of the hyperparameter “C” ranging from  $10^{-4}$  to  $10^1$  for separate models and present results for the hyperparameter that produced the best results. We used the “balanced” hyperparameter for all models of the SVM to account for the class imbalance while training. We compared our results to a random baseline equivalent to 50% balanced accuracy and our proposed partner perception baseline.

### 4 RESULTS AND DISCUSSION

We report the results of the best performing models in Table 1. The peak approach which used about only 1.1% of the whole 10 minute audio performed the best for the female model with 74.8%, outperforming both the random and partner perception baselines. Yet, it performed the worst for the male model. The peak-end approach performed the best for the male model with 53.3% albeit worse than the partner perception baseline and slightly better than the random baseline. Figure 2 and 3 show the confusion matrices of the best models for male and female respectively.

The peaks performing better than the end in predicting end-of-conversation affect (though for female partners only) is consistent with the results of [44], in which the peak rating was more predictive than the end. The peak approach produced the best results likely because the peak segments contained the most extreme emotional expressions (acoustically). This result was not the same for the male partners for whom the results for peak and ends were similar and worse than the baselines. This result suggests that the male partners may not have been more emotionally expressive (acoustically) at the peak segments than at the end. This reasoning is speculative and hence further investigation is needed using, for example, linguistic features before any conclusions can be drawn. These results points to the need to develop methods that can automatically identify the speaker turns with the most extreme emotional expressions, after which acoustic features can be extracted to get accurate end-of-conversation emotion predictions. This work is one step towards our goal to recognize the emotions of German-speaking couples in daily life based on 5 minutes of multimodal data from conversation moments which we are currently collecting [3].

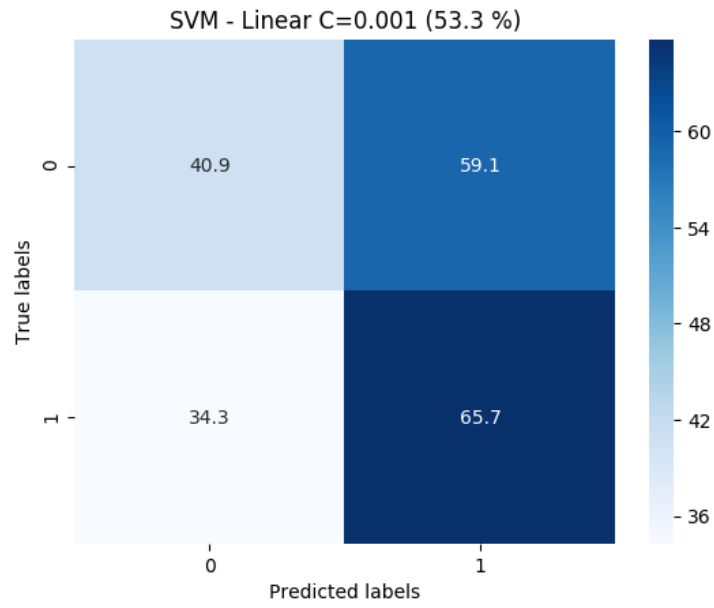


Fig. 2. Best Male Result Confusion Matrix (Peak-End Approach)

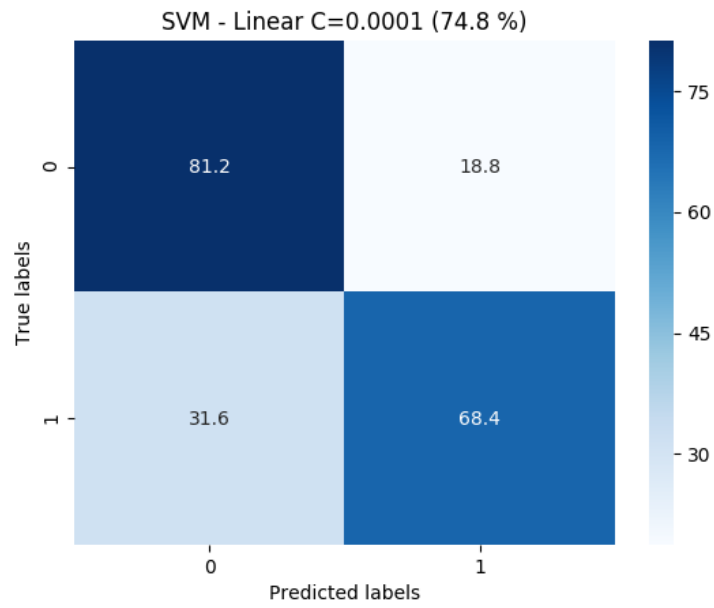


Fig. 3. Best Female Result Confusion Matrix (Peak Approach)

## 5 LIMITATIONS AND FUTURE WORK

In this work, we did not perform an evaluation with the whole audio or random segments as the focus was on the peaks and ends. Hence, we used random and partner perception baselines for comparison. Future work will use the whole audio, and random segments. Also, we focused on valence since that was the only dimension rated in the continuous rating. Future work will need to collect data with the arousal dimension and explore using the arousal dimension. Those results could be used together with this work to identify the right quadrant of the Affect grid and consequently, the kinds of emotions the person may be feeling. Additionally, we only used the negative/conflict conversation. These experiments will be repeated with the positive conversation and results will be compared to the results of this work. Furthermore, this work focused on evaluating the segments using acoustic features. We currently do not have manual transcripts of the data and automatic speech recognition systems that we tried out did not work well for this Dutch-based speech data. Hence, we plan to get manual transcript of this data and use linguistic features also. Additionally, given that the continuous ratings were done for the whole conversation including the speech of both partners, the peak rating of each partner may not always overlap with a speech segment of that partner. Hence, we first extracted the speaker turns of each partner, and then found the speaker turn with the peak rating. Consequently, the most extreme rating overall may not have used. We extracted and used features from both positive and negative peaks. Future work will evaluate using the positive and negative peaks separately and using different durations surrounding the peaks and ends. Additionally, we plan to perform a similar evaluation using self-reports other than the Affect Grid such as ratings for happy, sad, etc.

## 6 CONCLUSION

In this work, we performed an evaluation of the segments of an audio conversation that best predicts the end-of-conversation emotions of couples. We leveraged the peak-end rule, and used a transfer learning approach to extract features from (1) the audio segments with the most extreme positive and negative ratings, and (2) the ending of the audio. We used a pre-trained CNN to extract these acoustic features and a linear SVM to perform binary classification of the valence of partners. Our results showed that the segments from the peak produce the best results for recognizing the emotions of female partners and the approach was better than the partner perception baseline. This first-of-its-kind work contributes an evaluation of an approach that could be leveraged to best recognize the emotions of couples and then potentially used to improve the emotional well-being and relationship quality of couples via interventions.

## ACKNOWLEDGMENTS

We are grateful to Markus Wyss, Arthur Deschamps and Daniel Végh for their contributions to this work.

## REFERENCES

- [1] [n.d.]. YAMNet. <https://github.com/tensorflow/models/tree/master/research/audioset/yamnet>.
- [2] Matthew P Black, Athanasios Katsamanis, Brian R Baucom, Chi-Chun Lee, Adam C Lammert, Andrew Christensen, Panayiotis G Georgiou, and Shrikanth S Narayanan. 2013. Toward automating a human behavioral coding system for married couples' interactions using speech acoustic features. *Speech communication* 55, 1 (2013), 1–21.
- [3] George Boateng. 2020. Towards Real-Time Multimodal Emotion Recognition among Couples. In *Proceedings of the 2020 International Conference on Multimodal Interaction (ICMI '20), October 25–29, 2020, Virtual event, Netherlands*.
- [4] George Boateng, Laura Sels, Peter Kuppens, Janina Lüscher, Urte Scholz, and Tobias Kowatsch. 2020. Emotion Elicitation and Capture among Real Couples in the Lab. In *1st Momentary Emotion Elicitation & Capture workshop (MEEC 2020), co-located with the ACM CHI Conference on Human Factors in Computing Systems*.
- [5] Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeannette N Chang, Sungbok Lee, and Shrikanth S Narayanan. 2008. IEMOCAP: Interactive emotional dyadic motion capture database. *Language resources and evaluation* 42, 4 (2008), 335.

- [6] Carlos Busso, Srinivas Parthasarathy, Alec Burmania, Mohammed AbdelWahab, Najmeh Sadoughi, and Emily Mower Provost. 2016. MSP-IMPROV: An acted corpus of dyadic interactions to study emotion perception. *IEEE Transactions on Affective Computing* 8, 1 (2016), 67–80.
- [7] Laura L Carstensen, John M Gottman, and Robert W Levenson. 1995. Emotional behavior in long-term marriage. *Psychology and aging* 10, 1 (1995), 140.
- [8] Sandeep Nallan Chakravarthula, Haoqi Li, Shao-Yen Tseng, Maija Reblin, and Panayiotis Georgiou. 2019. Predicting Behavior in Cancer-Afflicted Patient and Spouse Interactions Using Speech and Language. *Proc. Interspeech 2019* (2019), 3073–3077.
- [9] Egon Dejonckheere, Merijn Mestdagh, Marlies Houben, Isa Rutten, Laura Sels, Peter Kuppens, and Francis Tuerlinckx. 2019. Complex affect dynamics add limited information to the prediction of psychological well-being. *Nature human behaviour* 3, 5 (2019), 478–491.
- [10] Sidney K D’mello and Jacqueline Kory. 2015. A review and meta-analysis of multimodal affect detection systems. *ACM Computing Surveys (CSUR)* 47, 3 (2015), 1–36.
- [11] Kexin Feng and Theodora Chaspari. 2020. A Review of Generalizable Transfer Learning in Automatic Emotion Recognition. *Frontiers in Computer Science* 2 (2020), 9.
- [12] Barbara L Fredrickson. 2000. Extracting meaning from past affective experiences: The importance of peaks, ends, and specific emotions. *Cognition & Emotion* 14, 4 (2000), 577–606.
- [13] Lisa Gaelick, Galen V Bodenhausen, and Robert S Wyer. 1985. Emotional communication in close relationships. *Journal of personality and social psychology* 49, 5 (1985), 1246.
- [14] Robert L Geist and David G Gilbert. 1996. Correlates of expressed and felt emotion during marital conflict: Satisfaction, personality, process, and outcome. *Personality and Individual Differences* 21, 1 (1996), 49–60.
- [15] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter. 2017. Audio set: An ontology and human-labeled dataset for audio events. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 776–780.
- [16] James Gibson, Athanasios Katsamanis, Matthew P Black, and Shrikanth Narayanan. 2011. Automatic identification of salient acoustic instances in couples’ behavioral interactions using diverse density support vector machines. In *Twelfth Annual Conference of the International Speech Communication Association*.
- [17] James Gibson, Bo Xiao, Panayiotis G Georgiou, and Shrikanth Narayanan. 2013. An audio-visual approach to learning salient behaviors in couples’ problem solving discussions. In *2013 IEEE International Conference on Multimedia and Expo Workshops (ICMEW)*. IEEE, 1–4.
- [18] John Mordechai Gottman. 2005. *The mathematics of marriage: Dynamic nonlinear models*. MIT Press.
- [19] John Mordechai Gottman. 2014. *What predicts divorce?: The relationship between marital processes and marital outcomes*. Psychology Press.
- [20] John M Gottman and Robert W Levenson. 1985. A valid procedure for obtaining self-report of affect in marital interaction. *Journal of consulting and clinical psychology* 53, 2 (1985), 151.
- [21] John M Gottman and Robert W Levenson. 1986. Assessing the role of emotion in marriage. *Behavioral Assessment* (1986).
- [22] Shawn Hershey, Sourish Chaudhuri, Daniel PW Ellis, Jort F Gemmeke, Aren Jansen, R Channing Moore, Manoj Plakal, Devin Platt, Rif A Saurous, Bryan Seybold, et al. 2017. CNN architectures for large-scale audio classification. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (icassp)*. IEEE, 131–135.
- [23] Richard E Heyman. 2001. Observation of couple conflicts: Clinical assessment applications, stubborn truths, and shaky foundations. *Psychological assessment* 13, 1 (2001), 5.
- [24] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. 2017. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861* (2017).
- [25] Daniel Kahneman. 2000. Evaluation by moments: Past and future. *Choices, values, and frames* (2000), 693–708.
- [26] Athanasios Katsamanis, James Gibson, Matthew P Black, and Shrikanth S Narayanan. 2011. Multiple instance learning for classification of human behavior observations. In *International Conference on Affective Computing and Intelligent Interaction*. Springer, 145–154.
- [27] Patricia K Kerig and Donald H Baucom. 2004. *Couple observational coding systems*. Taylor & Francis.
- [28] Chi-Chun Lee, Athanasios Katsamanis, Matthew P Black, Brian R Baucom, Panayiotis G Georgiou, and Shrikanth Narayanan. 2011. An analysis of PCA-based vocal entrainment measures in married couples’ affective spoken interactions. In *Twelfth Annual Conference of the International Speech Communication Association*.
- [29] Chi-Chun Lee, Athanasios Katsamanis, Matthew P Black, Brian R Baucom, Panayiotis G Georgiou, and Shrikanth S Narayanan. 2011. Affective state recognition in married couples’ interactions using PCA-based vocal entrainment measures with multiple instance learning. In *International Conference on Affective Computing and Intelligent Interaction*. Springer, 31–41.
- [30] Robert W Levenson, Laura L Carstensen, and John M Gottman. 1994. Influence of age and gender on affect, physiology, and their interrelations: A study of long-term marriages. *Journal of personality and social psychology* 67, 1 (1994), 56.
- [31] Robert W Levenson and John M Gottman. 1983. Marital interaction: physiological linkage and affective exchange. *Journal of personality and social psychology* 45, 3 (1983), 587.

- [32] Angeliki Metallinou, Chi-Chun Lee, Carlos Busso, Sharon Carnicke, Shrikanth Narayanan, et al. 2010. The USC CreativeIT database: A multimodal database of theatrical improvisation. *Multimodal Corpora: Advances in Capturing, Coding and Analyzing Multimodality* (2010), 55.
- [33] Angeliki Metallinou and Shrikanth Narayanan. 2013. Annotation and processing of continuous emotional attributes: Challenges and opportunities. In *2013 10th IEEE international conference and workshops on automatic face and gesture recognition (FG)*. IEEE, 1–8.
- [34] Philipp M Müller, Sikandar Amin, Prateek Verma, Mykhaylo Andriluka, and Andreas Bulling. 2015. Emotion recognition from embedded bodily expressions and speech during dyadic interactions. In *2015 International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, 663–669.
- [35] Hong-Wei Ng, Viet Dung Nguyen, Vassilios Vonikakis, and Stefan Winkler. 2015. Deep learning for emotion recognition on small datasets using transfer learning. In *Proceedings of the 2015 ACM on international conference on multimodal interaction*. 443–449.
- [36] Sally Olderbak, Andrea Hildebrandt, Thomas Pinkpank, Werner Sommer, and Oliver Wilhelm. 2014. Psychometric challenges and proposed solutions when scoring facial emotion expression codes. *Behavior Research Methods* 46, 4 (2014), 992–1006.
- [37] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in Python. *the Journal of machine Learning research* 12 (2011), 2825–2830.
- [38] Soujanya Poria, Erik Cambria, Rajiv Bajpai, and Amir Hussain. 2017. A review of affective computing: From unimodal analysis to multimodal fusion. *Information Fusion* 37 (2017), 98–125.
- [39] Anna Marie Ruef and Robert W Levenson. 2007. Continuous measurement of emotion. *Handbook of emotion elicitation and assessment* (2007), 286–297.
- [40] James A Russell. 1980. A circumplex model of affect. *Journal of personality and social psychology* 39, 6 (1980), 1161.
- [41] James A Russell, Anna Weiss, and Gerald A Mendelsohn. 1989. Affect grid: a single-item scale of pleasure and arousal. *Journal of personality and social psychology* 57, 3 (1989), 493.
- [42] Sourav Sahoo, Puneet Kumar, Balasubramanian Raman, and Partha Pratim Roy. 2019. A Segment Level Approach to Speech Emotion Recognition Using Transfer Learning. In *Asian Conference on Pattern Recognition*. Springer, 435–448.
- [43] Laura Sels, Jed Cabrieto, Emily Butler, Harry Reis, Eva Ceulemans, and Peter Kuppens. 2019. The occurrence and correlates of emotional interdependence in romantic relationships. *Journal of personality and social psychology* (2019).
- [44] Laura Sels, Eva Ceulemans, and Peter Kuppens. 2019. All’s well that ends well? A test of the peak-end rule in couples’ conflict discussions. *European Journal of Social Psychology* 49, 4 (2019), 794–806.
- [45] Laura Sels, Yan Ruan, Peter Kuppens, Eva Ceulemans, and Harry Reis. 2020. Actual and perceived emotional similarity in couples’ daily lives. *Social Psychological and Personality Science* 11, 2 (2020), 266–275.
- [46] Tessa V West and David A Kenny. 2011. The truth and bias model of judgment. *Psychological review* 118, 2 (2011), 357.

## Paper 3: BERT meets LIWC

# BERT meets LIWC: Exploring State-of-the-Art Language Models for Predicting Communication Behavior in Couples' Conflict Interactions

JACOPO BIGGIOGERA, University of Surrey, U.K.

GEORGE BOATENG, ETH Zurich, Switzerland

PETER HILPERT, University of Lausanne, Switzerland

MATTHEW VOWELS, University of Surrey, U.K.

GUY BODENMANN, University of Zurich, Switzerland

MONA NEYSARI, University of Zurich, Switzerland

FRIDTJOF NUSSBECK, University of Konstanz, Germany

TOBIAS KOWATSCH, ETH Zurich, Switzerland and University of St. Gallen, Switzerland

Many processes in psychology are complex, such as dyadic interactions between two interacting partners (e.g., patient-therapist, intimate relationship partners). Nevertheless, many basic questions about interactions are difficult to investigate because dyadic processes can be within a person and between partners, they are based on multimodal aspects of behavior and unfold rapidly. Current analyses are mainly based on the behavioral coding method, whereby human coders annotate behavior based on a coding schema. But coding is labor-intensive, expensive, slow, focuses on few modalities, and produces sparse data which has forced the field to use average behaviors across entire interactions, thereby undermining the ability to study processes on a fine-grained scale. Current approaches in psychology use LIWC for analyzing couples' interactions. However, advances in natural language processing such as BERT could enable the development of systems to potentially automate behavioral coding, which in turn could substantially improve psychological research. In this work, we train machine learning models to automatically predict positive and negative communication behavioral codes of 368 German-speaking Swiss couples during an 8-minute conflict interaction on a fine-grained scale (10-seconds sequences) using linguistic features and paralinguistic features derived with openSMILE. Our results show that both simpler TF-IDF features as well as more complex BERT features performed better than LIWC, and that adding paralinguistic features did not improve the performance. These results suggest it might be time to consider modern alternatives to LIWC, the de facto linguistic features in psychology, for prediction tasks in couples research. This work is a further step towards the automated coding of couples' behavior which could enhance couple research and therapy, and be utilized for other dyadic interactions as well.

CCS Concepts: • **Applied computing** → **Psychology**.

Additional Key Words and Phrases: Couples, predicting observer ratings, multimodal fusion, behavioral signal processing, BERT, LIWC, SVM

## ACM Reference Format:

Jacopo Biggogera, George Boateng, Peter Hilpert, Matthew Vowels, Guy Bodenmann, Mona Neysari, Fridtjof Nussbeck, and Tobias Kowatsch. 2021. BERT meets LIWC: Exploring State-of-the-Art Language Models for Predicting Communication Behavior in Couples' Conflict Interactions. In *Companion Publication of the 2021 International Conference on Multimodal*

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*ICMI '21 Companion, October 18–22, 2021, Montréal, QC, Canada*

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8471-1/21/10...\$15.00

<https://doi.org/10.1145/3461615.3485423>

*Interaction (ICMI '21 Companion), October 18–22, 2021, Montréal, QC, Canada.* ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3461615.3485423>

## 1 INTRODUCTION

There are many processes in the field of psychology that are very complex such as dyadic interactions – interactions between two people [23]. These processes are difficult to investigate because each person’s behavior is multimodal, both persons influence each other’s behavior mutually, and this process unfolds rapidly [21]. Such dynamic processes are relevant for a large number of human interactions (e.g., romantic partners, patient-therapist, student-teacher, buyer-seller).

Of the different human interactions, conflict interactions in intimate relationships have been well studied over the last decades [11]. Results indicate two principal types of communication behaviors: functional and dysfunctional. For example, contempt and criticism are a reliable predictor for later divorce and therefore seen as negative or dysfunctional, whereas providing appreciation and taking responsibility are considered to be functional and are associated with stable relationships [19, 20, 22]. It is therefore important to understand conflict interactions better as divorces are not only often emotionally and financially difficult for partners, but also have long-term negative consequences on the children involved [2].

The major reason for the disappointing progress of understanding behavioral processes during conflict interactions is the lack of methods that enable an automated approach for a fine-grained understanding of behavior. Traditionally, analyses in interaction research are mainly undertaken using data obtained from observer rating methods which are labor-intensive, expensive and time-consuming [27]. Consequently, codes are generally assigned on a global scale (e.g., one rating for 8-10 minutes sessions) rather than on a fine-grain scale (e.g., every talk turn or 10 seconds) resulting in sparse data. While observer ratings provide a means to capture global aspects of behavior (e.g., positive behavior), the analysis of such global behavioral aspects and sparse data has forced the field to focus on predictions based on average behaviors across entire interactions, thereby undermining the ability to study intra- and inter-individual processes [23].

Beyond observer rating methods, psychology has also included technology to extract linguistic (i.e., *what* was said) and paralinguistic features (i.e., *how* it was said). Various paralinguistic features have been extracted mainly using Praat [10] and openSMILE [18] which are software tools that compute various acoustic features over audio signals (e.g., pitch, fundamental frequency) over sequential time segments (e.g., 25 ms). They have been used in various works for example to show that the fundamental frequency of oscillation of the vocal folds is a valid proxy for emotional arousal [25] and a larger range in fundamental frequency is associated with more conflict interactions [4, 5]. Furthermore, a specific set of 88 features computed called eGeMAPS have been shown to be a minimalist feature set that performs well for affective recognition tasks [17].

Linguistic features have mainly been extracted through word-count-based programs like Linguistic Inquiry and Word Count (LIWC) [36] which is a software for extracting the count of words using an existing list of words and categories (e.g., positive/negative words, personal pronouns, social process). Its usage in couples research for example has shown the words partners utilize during conflict significantly affect interaction and overall marital quality. Findings indicate that greater first-person plural pronoun usage (‘we’), compared to first-person singular pronoun usage (‘I’) produces more positive resolutions to conflicts [34, 38]. Tools such as LIWC however are not without their limitations. In fact, they depend on the accuracy and comprehensiveness of the dictionary they are based upon, together with not being able to take into account both the context that words are placed in and the different meanings they might hold [3]. In a context such as conflict interactions where specific word choices and their meanings are important in affecting how the conflict unfolds [38], such limitations hold great significance for the validity and accuracy of applications that use such tools. However, recent advances in natural language processing such as Bidirectional Encoder Representations and Transformations (BERT) [16] based on the Transformer architecture [44] have been shown to set new state-of-the-art records in various natural



language understanding tasks such as natural language inference, question answering, and sentiment analysis. Some prior works have evaluated the predictive capability of BERT relative to LIWC in psychotherapy and mental health classification with BERT outperforming models based on LIWC features in populations with mental health diagnosis [24, 39]. Yet, BERT features have not yet been used in couples' interaction research for prediction tasks.

Some studies have used linguistic and or paralinguistic features specifically to predict behavioral codes for interacting romantic partners with the goal of automating behavioral coding. Most of these works have focused on session-level prediction – predicting one code for the whole 8-10 minute session [7, 8, 12, 13, 26, 30–32, 40–42, 46] with a scarcity of works focused on prediction of fine-grained behavioral codes such as at the speaker turn level or every few seconds. One such work is that of Chakravarthula et al [14] in which they trained machine learning models to predict 3 behavioral codes on a speaker turn level of 85 couples' 10-minute conversations using paralinguistic features (from openSMILE) and linguistic features (custom sentence embedding model) and achieved 57.4% unweighted average recall (balanced accuracy) for 3-class classification. Leveraging advanced sentence embedding methods such as BERT could potentially improve performance and increase the potential of automating behavioral coding. Yet, it has not been investigated in the context of couples research. Furthermore, including paralinguistic features could potentially enable better recognition.

In order to overcome current limitations, we utilized a data set collected from 368 couples (N = 736 participants) who were recorded during an 8-minute conflict interaction. Our main goal is to examine how linguistic and paralinguistic features in 10-second sequences can be used to predict how the same sequence is perceived and rated by human coders as positive or negative communication behavior. We aim to answer the following research questions (RQs):

**RQ1:** *Which linguistics features – LIWC or BERT – are better for predicting sequences-to-sequences-rated positive and negative communication behavior of partners?*

**RQ2:** *Given that the raters focused on coding linguistic aspects of behavior, how does adding openSMILE's eGeMAPS paralinguistic features affect the prediction performance?*

Our contributions are (1) an evaluation of the predictive capability of BERT vis-à-vis LIWC in the context of the automatic recognition of couples' communication behavioral codes on a fine-grained time-scale (every 10 seconds) (2) an investigation into how the addition of paralinguistic features affects prediction performance (3) the use of a unique dataset – spontaneous, real-life, speech data collected from German-speaking Swiss couples (n=368 couples, N=736 participants), and the largest ever such dataset used in the literature for automatic coding of couples' behavior. The insights from our work would enable the usage of new technologies to potentially automate the behavioral coding of couples, which could substantially improve the efficiency of couples research.

## 2 METHODOLOGY

**Data Collection and Preprocessing:** This work used data from a larger dyadic interaction laboratory project conducted at the premises of the University of Zurich, Switzerland over 10 years with 368 heterosexual German-speaking, Swiss couples (N=736 participants; age 20-80) [29, 43]. The inclusion criterion was to have been in the current relationship for at least 1 year. Couples had to choose one problematic topic for the conflict interaction from a list of common problems, and participants were then videotaped as they discussed the selected issue for 8 minutes. The data used in this work had one interaction from each couple and consequently, 368 8-minute interactions.

Two research assistants were trained to code communication behaviors using an adapted version of the Specific Affect Coding System (SPAFF) [15, 29]. Both raters practiced coding for at least 60 hours on videotapes that were not part of the study, with Cohen's kappa indicating that they had achieved an acceptable interobserver agreement ( $k = 0.9$ ). Each interaction was rated by both raters, with one rater focusing on the male partner and the other rater focusing on the female partner. Ratings were produced every 10 seconds to account for the behavior

unfolding during each sequence, resulting in 48 sequences for each interaction. *Positive communication*: (1) careful listening, interest, curiosity, (2) recognition, approval, factual praise, (3) affective communication, caring, (4) constructive criticism, and *negative communication*: (1) blaming, criticism, (2) defensiveness, (3) domineering, (4) withdrawal, stonewalling, (5) formally negative interaction, (6) contempt, (7) provocation, belligerence. For each 10-second sequence, raters would thus assign the code representing the communication behavior that was most prevalent out of the ones listed above. Raters were asked to focus on the verbal aspect of the behavior in assigning the codes. Due to the vast variety of codes present, we categorized all types of positive and negative as 1 and 2 respectively, and then passed them to machine learning models in the form of a binary classification problem.

The speech was manually annotated with the start and end of each speaker’s turn, along with pauses and noise. The speech was manually transcribed in 15-second chunks separately for each partner. Given that Swiss German is mostly spoken with different dialects across Switzerland, the spoken words were written as the corresponding German word equivalent. Transcripts and audio recordings acquired from the interactions were divided along the same 10-second sequence to match the 10-second sequences used for behavioral coding. This process was done separately for each partner’s transcript and speech data. Consequently, we dropped 10-second matched transcript-audio-code sequences that contained no speech and transcribed words.

Of the original 368 Swiss heterosexual couples that took part in the study, we could only use 345 because some couples requested their data to be removed and some data were missing arising from technical problems in data collection. In addition, while the original dataset presented instances where behaviors had been coded as neutral/no communication, these were dropped from the analyses since no accurate description for what constituted neutral communication was given in the codebook, and no differentiation with instances of no communication was provided. This thus resulted in a total of 9930 10-seconds speech sequences with their matching behavioral codes. Out of that number, 6978 were instances where communication had been coded as positive, while 2952 were instances where it had been coded as negative, highlighting a significant class imbalance that is characteristic of real-world datasets and partners’ behavior as seen in other works (e.g., [14]).

**Linguistic Features:** We extracted linguistic features from each 10-second transcript sequence using the LIWC software for German [33]. It utilizes an existing list of words and categories (e.g., positive/negative words, personal pronouns, social process) to count the number of the corresponding words in the transcript sequences and categorize them across 97 different features. The internal German LIWC dictionary was used to analyze the transcript and extract the features. We normalized each transcript sequence’s feature vector by dividing the value of all the other features by the “word count” feature which represents the number of words present in each transcript sequence. We then dropped the word count feature. This procedure thus left 96 normalized features that were passed as input to the machine learning models.

Also, we extracted features from each 10-second sequence using a pretrained Sentence-BERT (SBERT) model [37]. Sentence-BERT is a modification of the BERT architecture with siamese and triplet networks to compute sentence embeddings such that semantically similar sentences are close in vector space. Sentence-BERT has been shown to outperform the mean and CLS token outputs of regular BERT models for semantic similarity and sentiment classification tasks. Given that the text is in German, we used the German BERT model [1] as SBERT’s Transformer model and the mean pooling setting. The German BERT model was pretrained using the German Wikipedia dump, the OpenLegalData dump, and German news articles. The extraction resulted in a 768-dimensional feature vector.

**Paralinguistic Features:** We extracted acoustic features from the voice recordings. For each 10-second sequence, we first used the speaker annotations to get the acoustic signal for each partner separately. Next, we used openSMILE [17] to extract the 88 eGeMAPS acoustic features which have been shown to be a minimalist set of features for affective recognition tasks [18]. The original audio was encoded with 2 channels. As a result, we extracted the features for each channel resulting in a 176-dimensional feature vector.

### 3 EXPERIMENTS AND EVALUATION

We performed multiple experiments using the support vector machine (SVM) algorithm with the radial basis function (RBF) kernel and the scikit-learn library [35]. We used RBF since it performed the best in our initial explorations in comparison to random forests, XGBoost, and linear SVM. We trained models to perform binary classification of the behavioral codes for positive and negative communication using different feature sets. Specifically, we used features from LIWC, BERT, openSMILE. Also, we explored multimodal fusion at the features level of BERT and openSMILE by concatenating features from both groups. We used TF-IDF unigram and bigram features (using the most frequent 1000 features) of the transcripts as a linguistic baseline. To train and evaluate the models, we used nested K-fold cross-validation (CV). The nested procedure consisted of utilizing an “inner” run of 3-fold CV for hyperparameter tuning, followed by an “outer” run of 5-fold CV which utilizes the best values for each hyperparameter found by the “inner” run. We prevented data from the same couple from being in both the train and test folds, thereby evaluating the model’s performance on data from unseen couples. As the data was imbalanced, we utilized the metric balanced accuracy which is the unweighted average of the recall of each class, and confusion matrices for evaluation. We used different values of the hyperparameter “C”, presenting results for the hyperparameter that produced the best results. We used the “balanced” hyperparameter for all the SVM models to mitigate the class imbalance while training. Standard errors were computed by running models repeatedly, randomizing the groups used for the K fold CV and thus gathering a set of 20 accuracy measures for each model.

### 4 RESULTS AND DISCUSSION

Table 1 presents the results of the best performing models for each of the feature modalities. The model that used only the BERT features performed the best with 69.4% accuracy, compared to the LIWC model with 65.4% accuracy. The BERT only approach also performed better than combining paralinguistic features to the BERT features. The latter performed closely with 69.2% accuracy, but still significantly worse than the BERT only model ( $p < .001$  on a Wilcoxon signed rank test). Indeed, the paralinguistic baseline approach using openSMILE features performed the worst, 61.3%, being outperformed by TF-IDF linguistic baseline (65.6%) which also slightly outperformed LIWC. The worse performance of the paralinguistic features is expected given that raters in the study were instructed to focus on the verbal aspect of the interaction rather than nonverbal behavior in assigning codes.

Our results indicate that LIWC features do not have as much discriminative potential for prediction tasks compared both with even simpler approaches such as TF-IDF, as well as more advanced methods such as BERT. One likely explanation for BERT’s superior performance is its ability to capture the semantics of text via contextualized embeddings. These results have also been shown in a similar work using emotion psychotherapy or mental health data [24, 39]. The performance gain afforded by BERT notwithstanding, the simpler approaches did perform better than expected, with the performance improvement between TF-IDF and BERT being less than 4%. It is worth noting that this BERT model is used out-of-the-box and it is out-of-domain without any customization on the couples conversational text. The results indicate that researchers in social psychology ought to consider alternatives to LIWC, such as BERT, for extracting features for prediction tasks such as automated behavioral coding and emotion recognition. Although LIWC features (and indeed, TF-IDF) have the advantage of being simpler and more easily interpretable, various approaches are being developed to make BERT features more interpretable via its multi-head attention mechanism [45] and Shapley [28]. Finally, the result of the BERT model performing better than the multimodal approach is consistent with other works that found a similar result for emotion [9] and behavioral recognition [14]. Including paralinguistic features did not seem to add any more predictive information especially considering the context of the study in which assigning codes focused on

Table 1. Results across 20 runs (with standard errors) for models using LIWC, BERT, openSMILE and multimodal input features

Input Features	Balanced Accuracy (% +/- S.E.)
openSMILE	61.28 ± .07
TF-IDF + ngrams	65.61 ± .08
LIWC	65.41 ± .05
BERT	69.39 ± .06
BERT and openSMILE	69.18 ± .06

verbal behavior. Further approaches need to be explored to better combine the openSMILE and BERT features for improved results.

## 5 LIMITATIONS AND FUTURE WORK

In this work, we used manual transcripts. To accomplish true automated behavioral coding, our approach needs to use and work for automated transcriptions. Current speech recognition systems do not work for this unique dataset given that couples speak Swiss German, which is (1) a spoken dialect and not written, and (2) varies across different parts of the German-speaking regions of Switzerland. Further work is needed to develop automatic speech recognition systems for Swiss German.

Also, we only used the BERT model as a feature extractor to make a fair comparison with the LIWC features. Fine-tuning the BERT model on this task and domain to update the weights of the model would potentially improve the prediction results. This approach will be explored in future work. Finally, BERT models have been shown to encode gender and racial bias because of the data they are trained on. This consideration needs to be factored in for the specific prediction task and context [6].

## 6 CONCLUSION

In this work, we investigated the predictive potential of BERT features for automated coding of couples' communication behavior compared to LIWC features, the de facto linguistic features in social psychology. We extracted and compared LIWC and BERT features, used openSMILE features as a paralinguistic baseline and TF-IDF with ngrams as a linguistic baseline. We trained an RBF SVM to classify positive and negative communication behavior of each romantic partner on a 10-second granularity. Our results showed that both simple TF-IDF features as well as more complex BERT features both outperform LIWC, indicating that it might be time for researchers to consider alternatives to LIWC for predictive tasks in couples interactions. Additionally, adding paralinguistic features did not perform better than the BERT-only approach. Our work is a further step towards better approaches in automating the coding of couples' behavior which could enhance couples research and assessments in couples therapy.

## ACKNOWLEDGMENTS

Funding was provided by the Swiss National Science Foundation: CR12I1\_166348/1; CRSI11\_133004/1; P3P3P1\_174466; P300P1\_164582

## REFERENCES

- [1] [n.d.]. Open Sourcing German BERT. <https://deepset.ai/german-bert>. Accessed: 2020-05-1.
- [2] Paul R Amato. 2001. Children of divorce in the 1990s: an update of the Amato and Keith (1991) meta-analysis. *Journal of family psychology* 15, 3 (2001), 355.

- [3] Erin O'Carroll Bantum and Jason E Owen. 2009. Evaluating the validity of computerized content analysis programs for identification of emotional expression in cancer narratives. *Psychological assessment* 21, 1 (2009), 79.
- [4] Brian R Baucom, David C Atkins, Kathleen Eldridge, Pamela McFarland, Mia Sevier, and Andrew Christensen. 2011. The language of demand/withdraw: verbal and vocal expression in dyadic interactions. *Journal of Family Psychology* 25, 4 (2011), 570.
- [5] Brian R Baucom, Darby E Saxbe, Michelle C Ramos, Lauren A Spies, Esti Iturralde, Sarah Duman, and Gayla Margolin. 2012. Correlates and characteristics of adolescents' encoded emotional arousal during family conflict. *Emotion* 12, 6 (2012), 1281.
- [6] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 610–623.
- [7] Matthew Black, Athanasios Katsamanis, Chi-Chun Lee, Adam C Lammert, Brian R Baucom, Andrew Christensen, Panayiotis G Georgiou, and Shrikanth S Narayanan. 2010. Automatic classification of married couples' behavior using audio features. In *Eleventh annual conference of the international speech communication association*.
- [8] Matthew P Black, Athanasios Katsamanis, Brian R Baucom, Chi-Chun Lee, Adam C Lammert, Andrew Christensen, Panayiotis G Georgiou, and Shrikanth S Narayanan. 2013. Toward automating a human behavioral coding system for married couples' interactions using speech acoustic features. *Speech communication* 55, 1 (2013), 1–21.
- [9] George Boateng and Tobias Kowatsch. 2020. Speech Emotion Recognition among Elderly Individuals using Transfer Learning and Multimodal Fusion. In *Companion Publication of the 2020 International Conference on Multimodal Interaction (ICMI '20 Companion), October 25–29, 2020, Virtual event, Netherlands*.
- [10] Paul Boersma and Vincent Van Heuven. 2001. Speak and unSpeak with PRAAT. *Glott International* 5, 9/10 (2001), 341–347.
- [11] T.N. Bradbury and B.R. Karney. 2010. *Intimate Relationships*. W.W. Norton & Company. <https://books.google.ch/books?id=YMTeHAAACAAJ>
- [12] Sandeep Nallan Chakravarthula, Brian Baucom, and Panayiotis Georgiou. 2018. Modeling Interpersonal Influence of Verbal Behavior in Couples Therapy Dyadic Interactions. *arXiv preprint arXiv:1805.09436* (2018).
- [13] Sandeep Nallan Chakravarthula, Rahul Gupta, Brian Baucom, and Panayiotis Georgiou. 2015. A language-based generative model framework for behavioral analysis of couples' therapy. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2090–2094.
- [14] Sandeep Nallan Chakravarthula, Haoqi Li, Shao-Yen Tseng, Maija Reblin, and Panayiotis Georgiou. 2019. Predicting Behavior in Cancer-Afflicted Patient and Spouse Interactions Using Speech and Language. *Proc. Interspeech 2019* (2019), 3073–3077.
- [15] James A Coan and John M Gottman. 2007. The specific affect coding system (SPAFF). *Handbook of emotion elicitation and assessment* (2007), 267–285.
- [16] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [17] Florian Eyben, Klaus R Scherer, Björn W Schuller, Johan Sundberg, Elisabeth André, Carlos Busso, Laurence Y Devillers, Julien Epps, Petri Laukka, Shrikanth S Narayanan, et al. 2015. The Geneva minimalistic acoustic parameter set (GeMAPS) for voice research and affective computing. *IEEE transactions on affective computing* 7, 2 (2015), 190–202.
- [18] Florian Eyben, Martin Wöllmer, and Björn Schuller. 2010. Opensmile: the munich versatile and fast open-source audio feature extractor. In *Proceedings of the 18th ACM international conference on Multimedia*. 1459–1462.
- [19] John Gottman. 2018. *The Seven Principles for Making Marriage Work*. Hachette UK.
- [20] John Mordechai Gottman. 1994. *What predicts divorce?: The relationship between marital processes and marital outcomes*. Lawrence Erlbaum Associates, Inc.
- [21] John Mordechai Gottman. 2005. *The mathematics of marriage: Dynamic nonlinear models*. MIT Press.
- [22] John M Gottman, James Coan, Sybil Carrere, and Catherine Swanson. 1998. Predicting marital happiness and stability from newlywed interactions. *Journal of Marriage and the Family* (1998), 5–22.
- [23] Peter Hilpert, Timothy R Brick, Christoph Flückiger, Matthew J Vowels, Eva Ceulemans, Peter Kuppens, and Laura Sels. 2020. What can be learned from couple research: Examining emotional co-regulation processes in face-to-face interactions. *Journal of Counseling Psychology* 67, 4 (2020), 475.
- [24] Zheng Ping Jiang, Sarah Ita Levitan, Jonathan Zomick, and Julia Hirschberg. 2020. Detection of Mental Health from Reddit via Deep Contextualized Representations. In *Proceedings of the 11th International Workshop on Health Text Mining and Information Analysis*. 147–156.
- [25] Patrik N Juslin and Klaus R Scherer. 2005. *Vocal expression of affect*. Oxford University Press.
- [26] Athanasios Katsamanis, James Gibson, Matthew P Black, and Shrikanth S Narayanan. 2011. Multiple instance learning for classification of human behavior observations. In *International Conference on Affective Computing and Intelligent Interaction*. Springer, 145–154.
- [27] Patricia K Kerig and Donald H Baucom. 2004. *Couple observational coding systems*. Taylor & Francis.
- [28] Enja Kokalj, Blaž Škrlič, Nada Lavrač, Senja Pollak, and Marko Robnik-Šikonja. 2021. BERT meets Shapley: Extending SHAP Explanations to Transformer-based Classifiers. In *Proceedings of the EACL Hackathon on News Media Content Analysis and Automated Report Generation*. 16–21.

- [29] Monika Kuster, Katharina Bernecker, Sabine Backes, Veronika Brandstätter, Fridtjof W Nussbeck, Thomas N Bradbury, Mike Martin, Dorothee Sutter-Stickel, and Guy Bodenmann. 2015. Avoidance orientation and the escalation of negative communication in intimate relationships. *Journal of Personality and Social Psychology* 109, 2 (2015), 262.
- [30] Chi-Chun Lee, Matthew Black, Athanasios Katsamanis, Adam C Lammert, Brian R Baucom, Andrew Christensen, Panayiotis G Georgiou, and Shrikanth S Narayanan. 2010. Quantification of prosodic entrainment in affective spontaneous spoken interactions of married couples. In *Eleventh Annual Conference of the International Speech Communication Association*.
- [31] Chi-Chun Lee, Athanasios Katsamanis, Matthew P Black, Brian R Baucom, Andrew Christensen, Panayiotis G Georgiou, and Shrikanth S Narayanan. 2014. Computing vocal entrainment: A signal-derived PCA-based quantification scheme with application to affect analysis in married couple interactions. *Computer Speech & Language* 28, 2 (2014), 518–539.
- [32] Haoqi Li, Brian Baucom, and Panayiotis Georgiou. 2016. Sparsely connected and disjointly trained deep neural networks for low resource behavioral annotation: Acoustic classification in couples' therapy. *arXiv preprint arXiv:1606.04518* (2016).
- [33] Tabea Meier, Ryan L Boyd, James W Pennebaker, Matthias R Mehl, Mike Martin, Markus Wolf, and Andrea B Horn. 2019. "LIWC auf Deutsch": The Development, Psychometrics, and Introduction of DE-LIWC2015. *PsyArXiv a* (2019).
- [34] Mona Neysari, Guy Bodenmann, Matthias R Mehl, Katharina Bernecker, Fridtjof W Nussbeck, Sabine Backes, Martina Zemp, Mike Martin, and Andrea B Horn. 2016. Monitoring pronouns in conflicts. *GeroPsych* (2016).
- [35] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in Python. *the Journal of machine Learning research* 12 (2011), 2825–2830.
- [36] James W Pennebaker, Martha E Francis, and Roger J Booth. 2001. Linguistic inquiry and word count: LIWC 2001. *Mahway: Lawrence Erlbaum Associates* 71, 2001 (2001), 2001.
- [37] Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084* (2019).
- [38] Rachel A Simmons, Peter C Gordon, and Dianne L Chambless. 2005. Pronouns in marital interaction: What do "you" and "I" say about marital health? *Psychological science* 16, 12 (2005), 932–936.
- [39] Michael J Tanana, Christina S Soma, Patty B Kuo, Nicolas M Bertagnolli, Aaron Dembe, Brian T Pace, Vivek Srikumar, David C Atkins, and Zac E Imel. 2021. How do you feel? Using natural language processing to automatically rate emotion in psychotherapy. *Behavior Research Methods* (2021), 1–14.
- [40] Shao-Yen Tseng, Brian R Baucom, and Panayiotis G Georgiou. 2017. Approaching Human Performance in Behavior Estimation in Couples Therapy Using Deep Sentence Embeddings.. In *INTERSPEECH*. 3291–3295.
- [41] Shao-Yen Tseng, Sandeep Nallan Chakravarthula, Brian R Baucom, and Panayiotis G Georgiou. 2016. Couples Behavior Modeling and Annotation Using Low-Resource LSTM Language Models.. In *INTERSPEECH*. 898–902.
- [42] Shao-Yen Tseng, Haoqi Li, Brian Baucom, and Panayiotis Georgiou. 2018. "Honey, I Learned to Talk" Multimodal Fusion for Behavior Analysis. In *Proceedings of the 20th ACM International Conference on Multimodal Interaction*. 239–243.
- [43] UZH University of Zurich. [n.d.]. PASEZ Project-Impact of stress on relationship development of couples and children. <http://www.dynage.uzh.ch/en/newsevents/news/news25.html>. Accessed: 2021-05-1.
- [44] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *arXiv preprint arXiv:1706.03762* (2017).
- [45] Jesse Vig. 2019. A multiscale visualization of attention in the transformer model. *arXiv preprint arXiv:1906.05714* (2019).
- [46] Wei Xia, James Gibson, Bo Xiao, Brian Baucom, and Panayiotis G Georgiou. 2015. A dynamic model for behavioral analysis of couple interactions using acoustic features. In *Sixteenth Annual Conference of the International Speech Communication Association*.

## Paper 4: You made me feel this way

# “You made me feel this way”: Investigating Partners’ Influence in Predicting Emotions in Couples’ Conflict Interactions using Speech Data

GEORGE BOATENG, ETH Zurich, Switzerland

PETER HILPERT, University of Lausanne, Switzerland

GUY BODENMANN, University of Zurich, Switzerland

MONA NEYSARI, University of Zurich, Switzerland

TOBIAS KOWATSCH, ETH Zürich, Switzerland and University of St. Gallen, Switzerland

How romantic partners interact with each other during a conflict influences how they feel at the end of the interaction and is predictive of whether the partners stay together in the long term. Hence understanding the emotions of each partner is important. Yet current approaches that are used include self-reports which are burdensome and hence limit the frequency of this data collection. Automatic emotion prediction could address this challenge. Insights from psychology research indicate that partners’ behaviors influence each other’s emotions in conflict interaction and hence, the behavior of both partners could be considered to better predict each partner’s emotion. However, it is yet to be investigated how doing so compares to only using each partner’s own behavior in terms of emotion prediction performance. In this work, we used BERT to extract linguistic features (i.e., what partners said) and openSMILE to extract paralinguistic features (i.e., how they said it) from a data set of 368 German-speaking Swiss couples (N = 736 individuals) who were videotaped during an 8-minutes conflict interaction in the laboratory. Based on those features, we trained machine learning models to predict if partners feel positive or negative after the conflict interaction. Our results show that including the behavior of the other partner improves the prediction performance. Furthermore, for men, considering how their female partners spoke is most important and for women considering what their male partner said is most important in getting better prediction performance. This work is a step towards automatically recognizing each partners’ emotion based on the behavior of both, which would enable a better understanding of couples in research, therapy, and the real world.

CCS Concepts: • **Applied computing** → **Psychology**.

Additional Key Words and Phrases: couples, emotion recognition, multimodal fusion, linguistic, paralinguistic, conflict

## ACM Reference Format:

George Boateng, Peter Hilpert, Guy Bodenmann, Mona Neysari, and Tobias Kowatsch. 2021. “You made me feel this way”: Investigating Partners’ Influence in Predicting Emotions in Couples’ Conflict Interactions using Speech Data. In *Companion Publication of the 2021 International Conference on Multimodal Interaction (ICMI ’21 Companion)*, October 18–22, 2021, Montréal, QC, Canada. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3461615.3485424>

## 1 INTRODUCTION

Understanding the emotions partners feel during and after conflict interactions is important because of its long-term effects on couples’ relationship quality and stability [13]. Happy couples, for example, experience more positive and less negative emotions during conflict interactions compared to unhappy couples [21]. The current

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*ICMI ’21 Companion, October 18–22, 2021, Montréal, QC, Canada*

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8471-1/21/10...\$15.00

<https://doi.org/10.1145/3461615.3485424>



study focuses on one fundamental aspect of the conflict mechanism — how the emotional experience *within* each partner is influenced by the behavioral exchange *between* partners.

A crucial aspect of conflict interaction in couples is how the behavioral exchange makes each person feel during and after the interaction [25]. But although both partners experience the same interaction, they can feel very differently about it. For example, if we assume that partner A shows contempt and criticizes partner B, we can assume that partner A might feel angry or superior whereas partner B might feel hurt or humiliated. Thus, the experience can be very different for the partner who communicates something compared to the partner who perceives it [7]. This differentiation allows us to reflect on another mechanism, namely that the emotions a person experiences are the results of two kinds of influences. Obviously, a person’s emotional experience is constantly influenced by a partner’s behavior as a kind of co-regulating force, talk turn by talk turn. In addition, however, each person has the ability to regulate one’s own emotional response (e.g., cognitive appraisal, emotion regulation) [16], which then affects one’s own subsequent behavioral response. Thus, what partners experience emotionally during and at the end of a conflict interaction is a reflection of the co-regulation and self-regulation processes [6].

To better understand emotions in couples and their impact on relationships, often self-report assessments are used in which each partner is asked to provide a rating of their own emotions right after an interaction, or partners are asked to watch a video recording of the interaction and provide continuous ratings using a joystick, for example, [17, 25]. Self-reports are burdensome to complete and may not be collected frequently. This means that the relationship between behavior and emotions cannot be studied often. Thus, an automatic emotion recognition system would allow scaling of couples research.

Various works have used linguistic features (i.e., *what* has been said) and paralinguistic features (i.e., *how* it was said) to predict the emotions of each partner in couples interactions more broadly [3, 4, 8, 9, 19, 20, 22, 28–30, 32] and in conflict interactions in particular [5, 10]. Most of these works have used observer ratings (perceived emotions) as labels rather than self-reports (one’s actual emotions). Hence, the prediction task becomes that of recognizing external individuals’ perception of each partner’s emotion rather than each partner’s emotion per their own assessment. Though similar, the latter is more challenging than the former for a number of reasons. First, the rating might be biased and may not reflect their actual emotions over the period the rating is for (e.g., the past 5 minutes). Whereas for observer ratings, coders are generally trained over several weeks, it is done by more than one person and various approaches are employed to resolve ratings that are not in agreement and ensure the validity of the labels. Second, the self-reported emotion may not be reflected in that partner’s behavior in comparison to observer ratings which are purely based on behavioral observation.

Despite these challenges, insights from psychology research could be leveraged to make the prediction task easier. Specifically, given that partners’ behaviors influence each other’s emotions in conflict interaction, the behavior of both partners could be considered to better predict each partner’s end-of-conversation emotion. However, it is yet to be investigated how doing so compares to using each partner’s own behavior only in terms of emotion prediction performance. In this work, we used a dataset collected from 368 couples who were recorded during an 8-minute conflict interaction, extracted linguistic and paralinguistic features, and used machine learning approaches to predict how each partner felt directly after the conflict interaction (self-reported emotion). We answer the following research questions (RQs)

**RQ1:** *How well can the end-of-conversation emotion of each partner be predicted by their own behavior — a combination of linguistic and paralinguistic data? (self-regulation)*

**RQ2:** *How does the prediction performance change when including the other partner’s behavior — (a) linguistic only, (b) paralinguistic only, and (c) combination of linguistic and paralinguistic data? (co-regulation)*

Our contributions are (1) an evaluation of how well a partner’s own linguistic and paralinguistic features predict one’s own end-of-conversation emotion (2) an investigation of how the prediction performance changes when including one’s partner’s features (linguistic, paralinguistic, and both) (3) the use of a unique dataset

– spontaneous, real-life, speech data collected from German-speaking, Swiss couples ( $n=368$  couples,  $N=736$  participants), which is the largest ever such dataset used in the literature for automatic recognition of partners' end-of-conversation emotion. The insights from our work would advance the use of methods to automatically recognize the emotions of each partner which could enable research and applications to better understand couples' relationships in therapy and the real world.

The rest of the paper is organized as follows: In Section 2, we describe our data collection, preprocessing and feature extraction, in Section 3, we describe our experiments and evaluation, in Section 4 we present and discuss our results, in Section 5, we present limitations and future work, and we conclude in Section 6.

## 2 METHODOLOGY

### 2.1 Data Collection and Preprocessing

This work used data from a larger dyadic interaction laboratory project conducted at the premises of the University of Zurich, Switzerland over 10 years with 368 heterosexual German-speaking, Swiss couples ( $N=736$  participants; age 20-80) [18, 31]. The inclusion criterion was to have been in the current relationship for at least 1 year. Couples had to choose one problematic topic for the conflict interaction from a list of common problems, and participants were then videotaped as they discussed the selected issue for 8 minutes. The data used in this work had one interaction from each couple and consequently, 368 8-minute interactions.

After each conversation, each partner provided self-reported responses to the Multidimensional Mood questionnaire [27] of their emotions on four bipolar dimensions – namely “good mood versus bad mood,” “relaxed versus angry,” “happy versus sad” and “calm versus stressed” – with the scale: 1 – very much, 2 – much, 3 – a little, 4 – a little, 5 – much, 6 – very much. In this work, we sort to focus on predicting emotional valence (positive or negative) based on Russell's circumplex model of emotions [26]. Hence, we used an average of the “good mood versus bad mood” and “happy versus sad” scales which enables us to get a more valid score since several dimensions that measure similar constructs are combined. We did not use the other two scales because their polarity also could represent the arousal dimension of emotion (low vs high arousal). We then binarized the averaged values similar to prior works (e.g., [3, 5]) such that values greater than or equal to 3.5 were negative (0) and the rest were positive (1). Binarization enables us to map the data into Russell's circumplex model of emotions which has 4 quadrants for emotions, further enabling its real-world utility – easily being able to tell which group of emotions are being felt by each partner using binarized valence and arousal dimensions.

The speech data were manually annotated with the start and end of each speaker's turn, along with pauses and noise. This was necessary in order to later be able to extract linguistic and paralinguistic features for each partner separately. In addition, speech content of both partners was manually transcribed for each partner separately and stored in 15-second chunks. Given that Swiss German is mostly spoken with different dialects across Switzerland, the spoken words were written as the corresponding German word equivalent.

Some couples requested their data to be removed and some data were missing due to technical problems in data collection. Of the original 368 couples that took part in the study, we could use 338 samples for females (46 negative labels) and 341 samples for males (32 negative labels). The distribution highlights a significant class imbalance that is characteristic of real-world datasets and partners' emotions as seen in other similar works (e.g., [5]).

### 2.2 Linguistic Features

We extracted linguistic features from the transcripts of the whole 8-minute interaction using a pre-trained model – Sentence-BERT (SBERT) [24]. Sentence-BERT is a modification of the BERT architecture with siamese and triplet networks to compute sentence embeddings such that semantically similar sentences are close in vector space. Sentence-BERT has been shown to outperform the mean and CLS token outputs of regular BERT models for

semantic similarity and sentiment classification tasks. Given that the text is in German, we used the German BERT model [1] as SBERT’s Transformer model and the mean pooling setting. The German BERT model was pre-trained using the German Wikipedia dump, the OpenLegalData dump, and German news articles. The extraction resulted in a 768-dimensional feature vector.

### 2.3 Paralinguistic Features

We extracted acoustic features from the voice recordings. First, we used the speaker annotations to get the acoustic signal for each gender separately. Next, we used openSMILE [11] to extract the 88 eGeMAPS acoustic features which have been shown to be a minimalist set of features for affective recognition tasks [12]. The features are extracted in 25 ms sequences and then various functions (e.g., mean, median, range, etc.) are computed over the sequences resulting in 88 features for the whole 8-minute audio. The original audio was encoded with 2 channels. As a result, we extracted the features for each channel resulting in a 176-dimensional feature vector.

### 2.4 Multimodal and Dyadic Feature Fusion

Given that emotions are reflected in what and how people say things, we performed multimodal fusion (early fusion) by concatenating the linguistic and paralinguistic features resulting in a 944-dimensional feature. We did this separately for each partner. This feature vector was used as the baseline approach to answer research question (1).

Additionally, we fused features from both partners to answer research question (2). Specifically, for partner A, we concatenated their multimodal feature vector with the features of partner B and used it to predict partner A’s emotion label. This process was done for partner B as well. In order to investigate which behavioral data of the interacting partner was most important in the prediction of the emotions, we included the features in the following order (1) linguistic only, (2) paralinguistic only, and (3) multimodal fusion of both.

Consequently, we had four feature sets: (1) Multimodal fusion (baseline — own features), (2) Multimodal + Dyadic Fusion (with partner’s linguistic only), (3) Multimodal + Dyadic Fusion (with partner’s paralinguistic features only) (4) Multimodal + Dyadic Fusion (with partner’s combined linguistic and paralinguistic only). These were passed to machine learning models to answer the two research questions.

## 3 EXPERIMENTS AND EVALUATION

We ran experiments using scikit-learn [23] the following machine learning models: support vector machine (SVM) algorithm with linear and radial basis function kernel, and random forests. We trained models to perform binary classification of each partner’s self-reported positive and negative emotion using different feature sets. We used the four feature sets described in the previous section. To train and evaluate the models, we used nested K-fold cross-validation (CV). The nested procedure consisted of utilising an “inner” run of 5-fold CV for hyperparameter tuning, followed by an “outer” run of 10-fold CV which utilizes the best values for each hyperparameter found by the “inner” run. We prevented data from the same couple from being in both the train and test folds, thereby evaluating the model’s performance on data from unseen couples. As the data was imbalanced, we used the metric balanced accuracy which is the unweighted average of the recall of each class and confusion matrices for evaluation. We used the “balanced” hyperparameter for all the models to mitigate the class imbalance while training. We compare to a random baseline of 50% balanced accuracy.

## 4 RESULTS AND DISCUSSION

Our results are shown in Table 1. The baseline approach with multimodal fusion was not better than chance in predicting men’s emotions at the end of the conflict interaction (49.8%). This result is unexpected as this indicates that men’s own behaviors during the interaction are not related to how they feel at the end of the interaction. It

Table 1. Prediction Results of the Best Models for the Fusion Approaches

Approach	Balanced Accuracy (%)	
	Male	Female
Multimodal fusion (baseline)	49.8	59
Multimodal + Dyadic Fusion (with partner's combined linguistic and paralinguistic only)	52.3	63.2
Multimodal + Dyadic Fusion (with partner's linguistic features only)	53.5	<b>64.8</b>
Multimodal + Dyadic Fusion (with partner's paralinguistic only)	<b>56.1</b>	59.9

might be due to self-regulation processes or not showing much emotions during the interaction. This finding is clearly different from women as their emotions can be predicted by their own behavior (59%). Thus, it seems that women seem to express their emotions more clearly in their behavior. These results of poorer prediction performance for men compared to women is consistent with the results of [5].

Including features from the interacting partner improved the results for both men (52.3%) and women (63.2%). These results are consistent with psychology research that the behavior of partners' have an effect on each other's emotions in conflict interaction [7, 14]. These results are a crucial finding because (i) previous research shows that the behavior of one person influences the behavior of the other person [13] and (ii) that the emotional changes of one person affect the emotions of the other [7]. However, this is the first study showing that behavioral features assessed during the conflict interaction can be used to predict one partner's emotion at the end of the conversation. In addition, the improvement in women's emotion prediction at the end of interaction is greater when including their partner's linguistic data (64.8%) whereas there is hardly any difference when including partner's paralinguistic features. This result is a surprising finding as women generally pay more attention to paralinguistic cues [15]. Notably, the results are different for men. The prediction for men's emotions slightly increases when including their partner's linguistic features but the prediction improves substantially when including women's paralinguistic features. Although we do not know which specific paralinguistic features are the main drivers for predicting the emotions, this finding is in line with prior findings – when women “nag”, men experience strong negative physiologic reactions and tend to withdraw [13, 15]. Future research is needed to investigate which aspects of one's partner's behavior exactly causes their own emotion prediction performance to decline. In addition, these results have implications for the kind of behavioral information to consider to best predict each partner's end-of-conversation emotions.

We show the confusion matrices for the best models for the men and women in Figures 1 and 2 respectively. They reveal the models' challenges at recognizing positive emotions (1), resulting in them being misclassified as negative emotions (0).

## 5 LIMITATIONS AND FUTURE WORK

Further work is needed to investigate if these results are the same for couples in a different cultural context and also explore the effect of the interacting partners' behavior at a more granular level such as on a talk-turn basis. More fine-grained emotion ratings may be needed to investigate that. Other fusion approaches like late fusion can be explored.

We used BERT as a feature extractor in this work. Generating domain-specific sentence embeddings via fine-tuning the BERT model and exploring deep transfer learning models for the paralinguistic features may

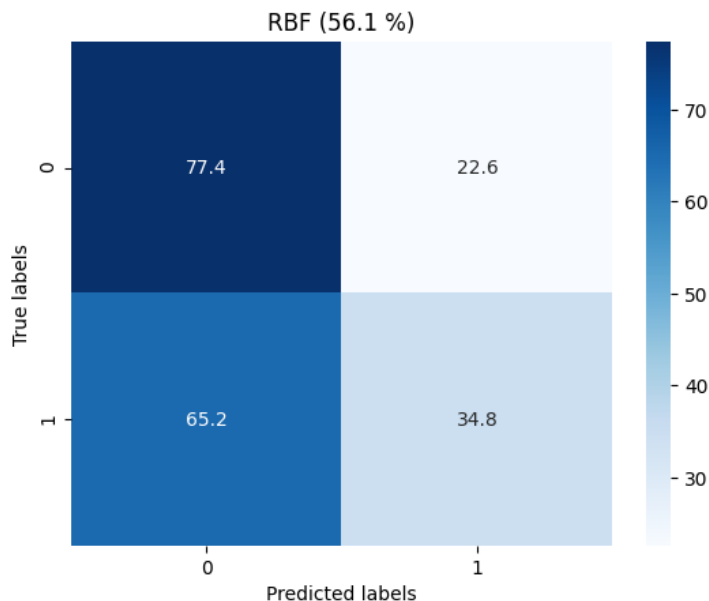


Fig. 1. Confusion matrix for best performing model for male partners (Multimodal + Dyadic Fusion with partner’s linguistic features only)

improve the results. BERT models have been shown to encode gender and racial bias because of the models they are trained on. Further investigations are needed in the future of potential biases in prediction [2].

Finally, we used manual annotations and transcripts. To accomplish true automatic emotion prediction, speaker annotations need to be done automatically and our approach needs to use and work for automated transcriptions. Current speech recognition systems do not work for this unique dataset given that couples speak Swiss German, which is (1) a spoken dialect and not written, and (2) varies across different parts of the German-speaking regions of Switzerland. Further work is needed to develop automatic speech recognition systems for Swiss German.

## 6 CONCLUSION

In this work, we investigated one’s own and the partner’s behavior in predicting end-of-conversation emotions in the context of conflict interactions in German-speaking Swiss couples. We extracted linguistic features using BERT and paralinguistic features using openSMILE. We fused both features in a multimodal approach for each partner. We also fused the features of both partners to predict the emotions of each partner. Our results show that including the behavior of the other partner improves the prediction performance. Furthermore, for men, considering *how* their female partners spoke is most important, and for women considering *what* their male partners said is most important in getting better prediction performance. These insights have implications for the behavioral information to (not) include to better predict each partner’s end-of-conversation emotions which will enable a better understanding of couples relations in research, therapy, and the real world.

## ACKNOWLEDGMENTS

Funding was provided by the Swiss National Science Foundation: CR12I1\_166348/1; CRSI11\_133004/1; P3P3P1\_174466; P300P1\_164582

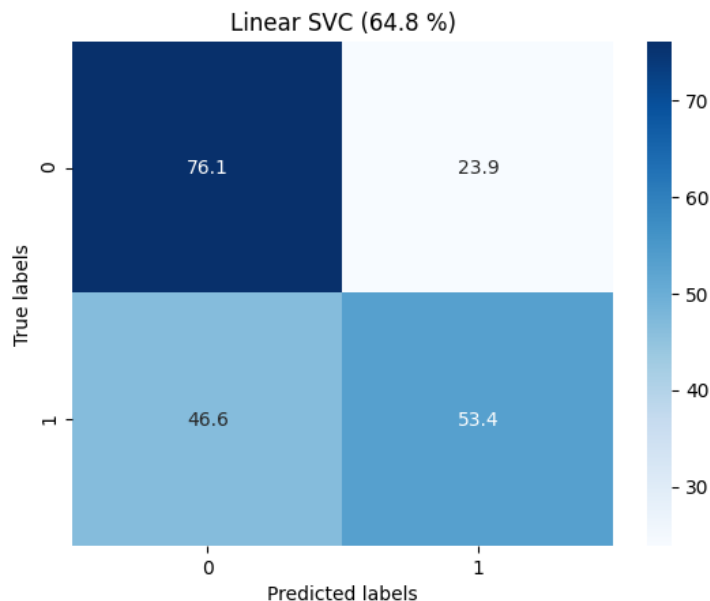


Fig. 2. Confusion matrix for best performing model for female partners (Multimodal + Dyadic Fusion with partner’s paralinguistic features only)

## REFERENCES

- [1] [n.d.]. Open Sourcing German BERT. <https://deepset.ai/german-bert>. Accessed: 2020-05-1.
- [2] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 610–623.
- [3] Matthew Black, Athanasios Katsamanis, Chi-Chun Lee, Adam C Lammert, Brian R Baucom, Andrew Christensen, Panayiotis G Georgiou, and Shrikanth S Narayanan. 2010. Automatic classification of married couples’ behavior using audio features. In *Eleventh annual conference of the international speech communication association*.
- [4] Matthew P Black, Athanasios Katsamanis, Brian R Baucom, Chi-Chun Lee, Adam C Lammert, Andrew Christensen, Panayiotis G Georgiou, and Shrikanth S Narayanan. 2013. Toward automating a human behavioral coding system for married couples’ interactions using speech acoustic features. *Speech communication* 55, 1 (2013), 1–21.
- [5] George Boateng, Laura Sels, Peter Kuppens, Peter Hilpert, and Tobias Kowatsch. 2020. Speech Emotion Recognition among Couples using the Peak-End Rule and Transfer Learning. In *Companion Publication of the 2020 International Conference on Multimodal Interaction (ICMI '20 Companion)*, October 25–29, 2020, Virtual event, Netherlands.
- [6] Steven M Boker and Jean-Philippe Laurenceau. 2006. Dynamical systems modeling: An application to the regulation of intimacy and disclosure in marriage. *Models for intensive longitudinal data* 63 (2006), 195–218.
- [7] Emily A Butler. 2011. Temporal interpersonal emotion systems: The “TIES” that form relationships. *Personality and Social Psychology Review* 15, 4 (2011), 367–393.
- [8] Sandeep Nallan Chakravarthula, Brian Baucom, and Panayiotis Georgiou. 2018. Modeling Interpersonal Influence of Verbal Behavior in Couples Therapy Dyadic Interactions. *arXiv preprint arXiv:1805.09436* (2018).
- [9] Sandeep Nallan Chakravarthula, Rahul Gupta, Brian Baucom, and Panayiotis Georgiou. 2015. A language-based generative model framework for behavioral analysis of couples’ therapy. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2090–2094.
- [10] Sandeep Nallan Chakravarthula, Haoqi Li, Shao-Yen Tseng, Maija Reblin, and Panayiotis Georgiou. 2019. Predicting Behavior in Cancer-Afflicted Patient and Spouse Interactions Using Speech and Language. *Proc. Interspeech 2019* (2019), 3073–3077.
- [11] Florian Eyben, Klaus R Scherer, Björn W Schuller, Johan Sundberg, Elisabeth André, Carlos Busso, Laurence Y Devillers, Julien Epps, Petri Laukka, Shrikanth S Narayanan, et al. 2015. The Geneva minimalistic acoustic parameter set (GeMAPS) for voice research and

- affective computing. *IEEE transactions on affective computing* 7, 2 (2015), 190–202.
- [12] Florian Eyben, Martin Wöllmer, and Björn Schuller. 2010. Opensmile: the munich versatile and fast open-source audio feature extractor. In *Proceedings of the 18th ACM international conference on Multimedia*. 1459–1462.
- [13] John Mordechai Gottman. 1994. *What predicts divorce?: The relationship between marital processes and marital outcomes*. Lawrence Erlbaum Associates, Inc.
- [14] John Mordechai Gottman. 2014. *What predicts divorce?: The relationship between marital processes and marital outcomes*. Psychology Press.
- [15] John M Gottman and Robert W Levenson. 1992. Marital processes predictive of later dissolution: behavior, physiology, and health. *Journal of personality and social psychology* 63, 2 (1992), 221.
- [16] James J Gross. 2014. Emotion regulation: Conceptual and empirical foundations. (2014).
- [17] Peter Hilpert, Timothy R Brick, Christoph Flückiger, Matthew J Vowels, Eva Ceulemans, Peter Kuppens, and Laura Sels. 2020. What can be learned from couple research: Examining emotional co-regulation processes in face-to-face interactions. *Journal of Counseling Psychology* 67, 4 (2020), 475.
- [18] Monika Kuster, Katharina Bernecker, Sabine Backes, Veronika Brandstätter, Fridtjof W Nussbeck, Thomas N Bradbury, Mike Martin, Dorothee Sutter-Stickel, and Guy Bodenmann. 2015. Avoidance orientation and the escalation of negative communication in intimate relationships. *Journal of Personality and Social Psychology* 109, 2 (2015), 262.
- [19] Chi-Chun Lee, Matthew Black, Athanasios Katsamanis, Adam C Lammert, Brian R Baucom, Andrew Christensen, Panayiotis G Georgiou, and Shrikanth S Narayanan. 2010. Quantification of prosodic entrainment in affective spontaneous spoken interactions of married couples. In *Eleventh Annual Conference of the International Speech Communication Association*.
- [20] Chi-Chun Lee, Athanasios Katsamanis, Matthew P Black, Brian R Baucom, Andrew Christensen, Panayiotis G Georgiou, and Shrikanth S Narayanan. 2014. Computing vocal entrainment: A signal-derived PCA-based quantification scheme with application to affect analysis in married couple interactions. *Computer Speech & Language* 28, 2 (2014), 518–539.
- [21] Robert W Levenson, Laura L Carstensen, and John M Gottman. 1994. Influence of age and gender on affect, physiology, and their interrelations: A study of long-term marriages. *Journal of personality and social psychology* 67, 1 (1994), 56.
- [22] Haoqi Li, Brian Baucom, and Panayiotis Georgiou. 2016. Sparsely connected and disjointly trained deep neural networks for low resource behavioral annotation: Acoustic classification in couples' therapy. *arXiv preprint arXiv:1606.04518* (2016).
- [23] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in Python. *the Journal of machine Learning research* 12 (2011), 2825–2830.
- [24] Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084* (2019).
- [25] Anna Marie Ruef and Robert W Levenson. 2007. Continuous measurement of emotion. *Handbook of emotion elicitation and assessment* (2007), 286–297.
- [26] James A Russell. 1980. A circumplex model of affect. *Journal of personality and social psychology* 39, 6 (1980), 1161.
- [27] Rolf Steyer, Peter Schwenkmezger, Peter Notz, and Michael Eid. 1997. Der Mehrdimensionale Befindlichkeitsfragebogen MDBF [Multidimensional mood questionnaire]. *Göttingen, Germany: Hogrefe* (1997).
- [28] Shao-Yen Tseng, Brian R Baucom, and Panayiotis G Georgiou. 2017. Approaching Human Performance in Behavior Estimation in Couples Therapy Using Deep Sentence Embeddings.. In *INTERSPEECH*. 3291–3295.
- [29] Shao-Yen Tseng, Sandeep Nallan Chakravarthula, Brian R Baucom, and Panayiotis G Georgiou. 2016. Couples Behavior Modeling and Annotation Using Low-Resource LSTM Language Models.. In *INTERSPEECH*. 898–902.
- [30] Shao-Yen Tseng, Haoqi Li, Brian Baucom, and Panayiotis Georgiou. 2018. "Honey, I Learned to Talk" Multimodal Fusion for Behavior Analysis. In *Proceedings of the 20th ACM International Conference on Multimodal Interaction*. 239–243.
- [31] UZH University of Zurich. [n.d.]. PASEZ Project-Impact of stress on relationship development of couples and children. <http://www.dynage.uzh.ch/en/newsevents/news/news25.html>. Accessed: 2021-05-1.
- [32] Wei Xia, James Gibson, Bo Xiao, Brian Baucom, and Panayiotis G Georgiou. 2015. A dynamic model for behavioral analysis of couple interactions using acoustic features. In *Sixteenth Annual Conference of the International Speech Communication Association*.

## Paper 5: Elderly Emotion Recognition



# Speech Emotion Recognition among Elderly Individuals using Multimodal Fusion and Transfer Learning

GEORGE BOATENG, ETH Zürich, Switzerland

TOBIAS KOWATSCH, ETH Zürich, Switzerland and University of St. Gallen, Switzerland

Recognizing the emotions of the elderly is important as it could give an insight into their mental health. Emotion recognition systems that work well on the elderly could be used to assess their emotions in places such as nursing homes and could inform the development of various activities and interventions to improve their mental health. However, several emotion recognition systems are developed using data from younger adults. In this work, we trained machine learning models to recognize the emotions of elderly individuals via performing a 3-class classification of valence and arousal as part of the INTERSPEECH 2020 Computational Paralinguistics Challenge (COMPARE). We used speech data from 87 participants who gave spontaneous personal narratives. We leveraged a transfer learning approach in which we used pretrained CNN and BERT models to extract acoustic and linguistic features respectively and fed them into separate machine learning models. Also, we fused these two modalities in a multimodal approach. Our best model used a linguistic approach and outperformed the official competition of unweighted average recall (UAR) baseline for valence by 8.8% and the mean of valence and arousal by 3.2%. We also showed that feature engineering is not necessary as transfer learning without fine-tuning performs as well or better and could be leveraged for the task of recognizing the emotions of elderly individuals. This work is a step towards better recognition of the emotions of the elderly which could eventually inform the development of interventions to manage their mental health.

CCS Concepts: • **Applied computing** → **Psychology**.

Additional Key Words and Phrases: Speech emotion recognition; Affective computing; Transfer learning; Computational paralinguistics; Elderly individuals; Multimodal fusion; Deep learning; CNN; LSTM; BERT; SBERT; Support vector machine

## ACM Reference Format:

George Boateng and Tobias Kowatsch. 2020. Speech Emotion Recognition among Elderly Individuals using Multimodal Fusion and Transfer Learning. In *Companion Publication of the 2020 International Conference on Multimodal Interaction (ICMI '20 Companion)*, October 25–29, 2020, Virtual event, Netherlands. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3395035.3425255>

## 1 INTRODUCTION

Digital technologies are needed to aid in managing the physical and emotional well-being of elderly individuals [24]. Awareness of the emotions of the elderly could give an insight into their mental health. Emotion recognition systems that work well on the elderly could be used to assess their emotions in places such as nursing homes and could inform the development of various activities and interventions to improve their mental health. However, several emotion recognition works use data collected from actors and younger adults for their development and evaluation (e.g. IEMOCAP dataset [3]). In this work, we develop and evaluate emotion recognition models using the first public speech data collected from elderly individuals in the real world for emotion recognition as part of the INTERSPEECH 2020 Computational Paralinguistics Challenge (COMPARE) [26]. The task was to perform

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*ICMI '20 Companion, October 25–29, 2020, Virtual event, Netherlands*

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8002-7/20/10...\$15.00

<https://doi.org/10.1145/3395035.3425255>

a 3-class classification of the arousal and valence dimensions of emotions based on speech data from elderly individuals.

Deep learning has been used for speech emotion recognition involving various approaches such as convolutional neural networks (CNN), Recurrent Neural Networks (RNN) such as Long Short-Term Memory (LSTM) – with and without attention – bidirectional LSTM (BLSTM), mostly together with handcrafted features ([17]). Other approaches have used the raw signal in an end-to-end approach leveraging 1D CNNs and LSTMs [29]. Transfer learning is another approach used in deep learning to circumvent the need to develop hand-crafted features and also deals with the challenge of small labeled datasets. Transfer learning entails pretraining a model on a different but related task and using it for feature extraction or fine-tuning in which the whole model or later layers are retrained ([7]). Transfer learning has shown success in various fields such as computer vision ([13, 16]), speech processing ([15]), and natural language processing ([12, 23]). Transfer learning has also been used in emotion recognition tasks ([7, 14, 25]).

Our contribution is the evaluation of transfer learning approaches to recognize the emotions of elderly individuals using a novel dataset – speech data collected from German-speaking elderly individuals. Specifically, we used a pretrained CNN model to extract acoustic features and a pretrained Transformer language model – Bidirectional Encoder Representations from Transformers (BERT) [6] – to extract linguistic features. We trained and evaluated separate models for acoustic and linguistic modalities. Also, we used a multimodal approach in which we fused the features (early fusion) and trained models using the combined features [20].

The rest of our paper is organized as follows. In Section 2, we describe our methodology. In Section 3, we describe our experiments. In Section 4, we show the results, discuss them and present future work. We conclude in Section 5.

## 2 METHODS

In this section, we describe the dataset, the competition baseline approaches, and our acoustic, linguistic, and multimodal approaches as shown in Figure 1.

### 2.1 Dataset

We used the Ulm State of mind elderly (USOMS-e) database collected from German-speaking elderly individuals [26]. The dataset contains speech data of 87 participants (55 f, 32 m, age 60–95 years, mean 71.01 years, std. dev. 9.14 years), each of whom told two negative and one positive personal narrative. Participants' emotions were assessed post every narrative by the subject and later by 4 experts on a scale of 0 (very sleepy and very bad) to 10 (very excited and very good) for the "arousal" and "valence" dimensions respectively. The audio data was converted to 16 KHz mono and was segmented into 5-sec chunks. The audio was also transcribed manually and automatically. The mean values of each dimension were used to create 3 classes: low (0-6), medium (7-8), and high (9-10).

### 2.2 Competition Baseline Approach

The organizers of the competition used various approaches to generate the baseline results for the competition such as feature engineering, transfer learning, unsupervised learning and end-to-end learning [26]. For feature engineering, they used the openSMILE toolkit to extract 6373 static features (functionals), and the OPENXBOW toolkit to extract Bag of Audio Words (BoAW) features. For transfer learning, they used the DEEP SPECTRUM toolkit which used a pretrained CNN (ResNet50) to extract embeddings from the spectrograms of the audio. They also used the Linguistic Feature Extractor (LIFE) toolkit to extract linguistic embeddings which used a BERT model that was pretrained on German text followed by Global Maximum pooling or bidirectional LSTM with attention. For unsupervised learning, they used the AUDEEP toolkit which used recurrent sequence-to-sequence

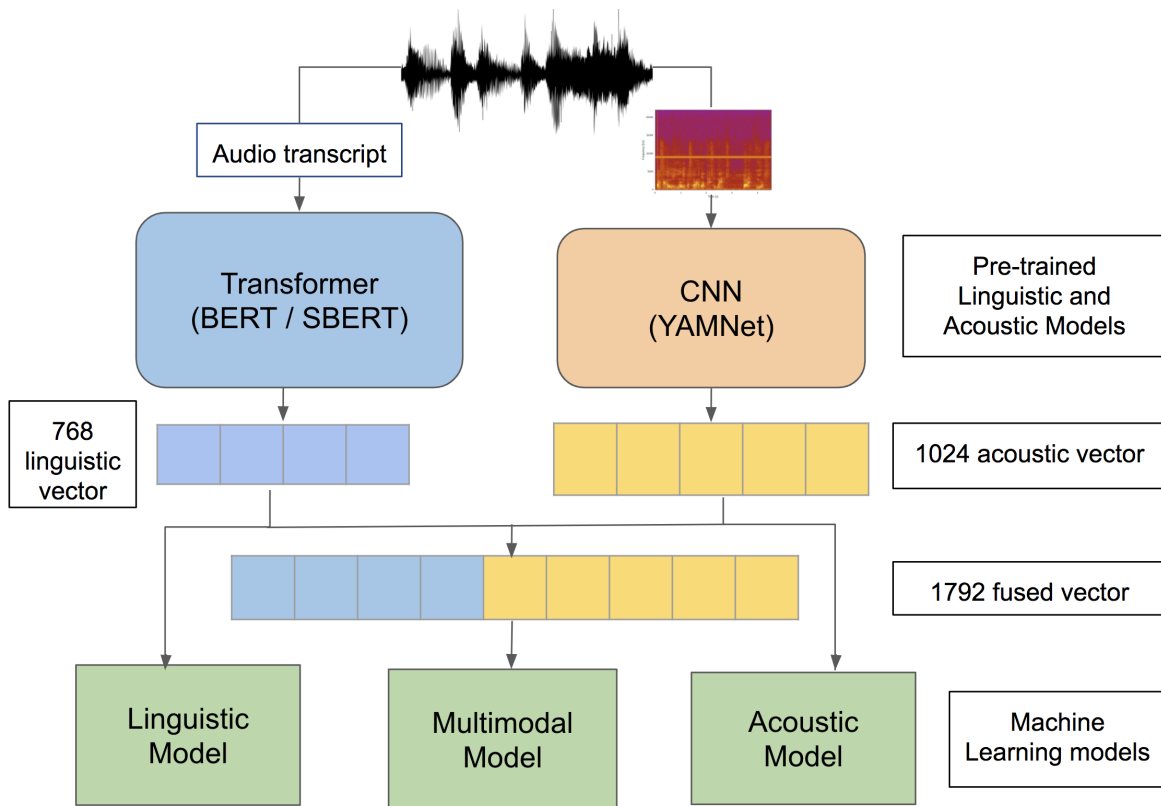


Fig. 1. Overview of Acoustic, Linguistic and Multimodal Approaches

autoencoders to learn representations of mel-spectrograms of the audio in an unsupervised manner. These different feature sets were then fed into separate linear support vector machines with different hyperparameters.

### 2.3 Acoustic Approach

We used the acoustic characteristics of the audio to perform classification. We extracted spectrograms and used a pretrained CNN to compute embeddings which we used as acoustic features to perform classification with various machine learning models (Figure 1). We used the YAMNet model which is a CNN that was pretrained on the AudioSet dataset to predict 521 audio event classes [8, 9]. YAMNet is based on the MobileNet architecture [11]. We used the YAMNet model as a feature extractor and hence replaced the original final logistic layer which outputs 521 class with various machine learning algorithms.

We extracted a spectrogram as an input into the YAMNet model in the same way as was done for the trained model. The audio is downsampled from 44.1Kz to 16 kHz mono. A spectrogram is computed using magnitudes of the Short-Time Fourier Transform with a window size of 25 ms, a window hop of 10 ms, and a periodic Hann window. A mel spectrogram is computed by mapping the spectrogram to 64 mel bins covering the range 125-7500 Hz. A stabilized log mel spectrogram is computed by applying  $\log(\text{mel-spectrum} + 0.01)$  where the offset is used to avoid taking a logarithm of zero. These features are then framed into non-overlapping examples of 0.96 seconds, where each example covers 64 mel bands and 96 frames of 10 ms each. This resulted in a 2D data of size 96 x

64 for each second, which we used as a data point input to the YAMNet model. The output of the model is a 1024-dimensional feature vector per data point input of size 96 x 64. We then normalized the feature vectors to be zero mean and unit variance and then used them as inputs to various machine learning models.

## 2.4 Linguistic Approach

We used the content of the speech – the manual transcript – to perform classification. Specifically, we used pretrained Transformer language models to extract linguistic features and then performed classification with various models (Figure 1). We used a pretrained BERT model to extract a 768-dimensional embedding vector for each narrative [6]. BERT is a deep learning model that has achieved state-of-the-art results for several natural language tasks. The BERT model we used is a case sensitive German BERT that was trained using a German Wikipedia dump, the OpenLegalData dump, and news articles [2]. We preprocessed each story’s transcript by first tokenizing each word and ensuring that the total number of tokens was less than or equal to the 512 maximum that the BERT model takes. Hence, we ignored subsequent words in each story which was over 512 length. We added special tokens for sentence classification (such as [CLS] at the first position). After passing each story into the model, we took the 768-dimensional embedding vector of the first token [CLS] of the last hidden layer and used that as the embedding for the whole story. We then normalized the vectors to be zero mean and unit variance and then used the features vectors as inputs to various machine learning models.

We also used Sentence BERT (SBERT), a modification of the BERT architecture with siamese and triplet network structures for generating sentence embeddings such that semantically similar sentences are close in vector space [21]. The SBERT network was shown to outperform state-of-the-art sentence embedding methods such as BERT and Universal Sentence Encoder for semantic similarity and sentence classification tasks such as sentiment detection. We used the multilingual version of the SBERT model [22]. The network, like the original BERT outputs a 768-dimensional embedding for each story. We normalized the vectors to be zero mean and unit variance and then used the feature vectors as inputs to various machine learning models.

## 2.5 Multimodal Approach

We also explored using a multimodal approach in which we fused aspects of the acoustic and linguistic modalities because multimodal approaches have been shown to outperform unimodal approaches in emotion recognition tasks [20]. Specifically, we fused the feature vectors from the acoustic and linguistic approaches producing a 1792-dimensional feature vector for each story (Figure 1). Since there were several acoustic feature vectors for each story, we performed a weighted sum of the acoustic feature vectors for each story. We then normalized the vectors to be zero mean and unit variance and then used these fused vectors as inputs to various machine learning models.

## 3 EXPERIMENTS

We performed various experiments using the following libraries scikit-learn [19], keras [5], and PyTorch [18]. We trained models separately for valence and arousal, and used a hyperparameter search to get models that produced the best results. We used a linear support vector machine (SVM), and a 2-layer LSTM [10] with 16 and 8 units, and 50% dropout [27] after each layer. We used the LSTM model for the acoustic approach to take advantage of the sequential nature of the acoustic embeddings. Also, for the acoustic approach, we used majority voting of the classification of the 5-sec audio chunks to decide the class for each story. For evaluation, we used the metric unweighted average recall (UAR) which is used for unbalanced data and confusion matrices.

Given that the data is imbalanced, we upsampled the minority classes so the data was balanced using the SMOTE algorithm [28] and imblearn library [1]. We used the train and development data sets provided by the competition organizers for developing the model. The organizers had a held-out test whose labels were not

Table 1. Results for Competition Baseline Approaches and our Acoustic, Linguistic, and Multimodal Approaches

Model	Dev (UAR %)		Test (UAR %)	
	Val	Arous	Val	Arous
<b>Competition Baseline Approach</b>				
Functionals + SVM	33.3	39.1	33.3	47.9
BoAW + SVM	33.3	40.5	31.5	49.1
Autoencoder + SVM	36.7	34.9	33.8	44.3
ResNet50 + SVM	31.6	35.0	40.3	<b>50.4</b>
BERT + LSTM + SVM	49.2	40.6	<b>49.0</b>	44.0
<b>Acoustic Approach</b>				
YAMNet + SVM	44.3	43.9	34.7	43.9
YAMNet + LSTM	37	40.2	—	47.9
<b>Linguistic Approach</b>				
BERT + SVM	51.1	45.7	56.3	<b>48</b>
SBERT + SVM	57.42	30.33	<b>57.8</b>	—
<b>Multimodal Approach</b>				
Fusion + SVM	49	43.8	52.3	47.4

made available to researchers. We had to submit our predictions on the held-out test which was evaluated by the organizers, and the prediction result sent to us. Also, we had a constraint of five submissions on the held-out test set and hence we used only our best models for those submissions. The official competition baseline was based on the performance on the held-out test set.

#### 4 RESULTS, DISCUSSION AND FUTURE WORK

We present the results for the competition baseline, and our acoustic, linguistic, and multimodal approaches in Table 1 where a "—" means that the model was not used for the held-out test. The best results for the competition baseline and our approaches in the valence and arousal columns are highlighted in **bold**. Also, we show the confusion matrices of the best models in Figure 2 (valence) and 3 (arousal).

The competition organizers' best methods which produced the results that were used as the official competition baseline results were the DEEP SPECTRUM (ResNet50 + SVM) for acoustic (50.4%) and LIFE (BERT + LSTM + SVM) for valence (49.0%) and an average of valence and arousal of 49.7%.

Among our approaches, the linguistic models performed the best for both valence and arousal, with the multimodal model being the second best for valence and the acoustic model being the second best for arousal. Our best model for valence was SBERT + SVM with a UAR of 57.8% and the best model for arousal is BERT + SVM with UAR of 48% and an overall mean UAR of valence and arousal being 52.9%. Our best models outperformed the official baseline (using the held-out test set) for valence by 8.8% and the mean of valence and arousal by 3.2%. Our best arousal model is however below the official arousal baseline by 2.4%. Our acoustic models not performing better than the baseline suggests that using the pretrained YAMNet model as feature extractor is not adequate. Hence, fine-tuning the model additionally or pretraining the model on a related emotion recognition task might be necessary for good performance.

The linguistic model performing better than the acoustic model is consistent with the results of other works such as an emotion recognition task among real-world couples whose best recognition result for a 3-class classification of valence was 57.42% (UAR) [4]. A possible explanation is that we used the manual transcript which is a perfect representation of the narratives which the linguistic model used as compared to the acoustic

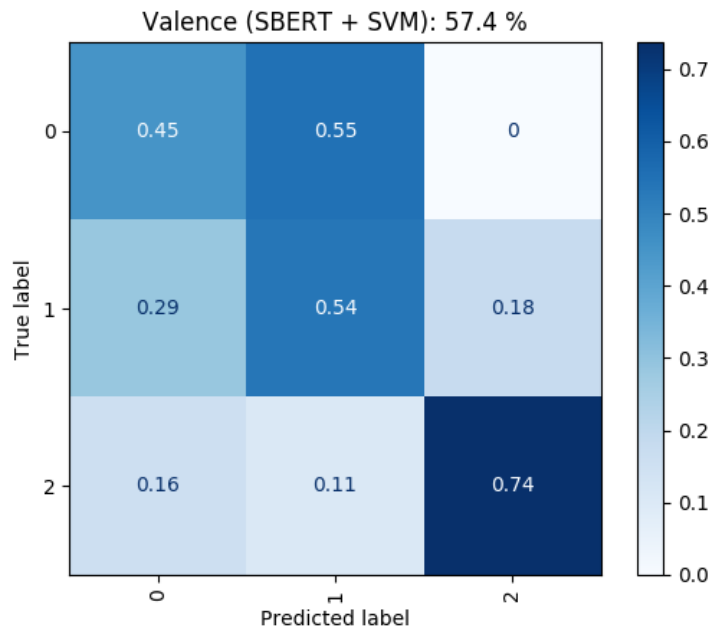


Fig. 2. Confusion matrix for development set evaluation for the best model for valence — SBERT + SVM: 57.4%

models which worked on raw, noisy, audio data. The model might have performed poorly with the automatic transcript but we did not evaluate that as we used only the best model for evaluation. Also, the SBERT model performed better than the regular BERT model for valence. This result is consistent with [21] which showed that SBERT extracts better sentence embeddings than BERT for sentiment detection tasks.

The multimodal model surprising did not perform the best considering multimodal approaches have been shown to perform better than unimodal approaches. This performance is however consistent with the result of [4]. It is possible that the limitations of the acoustic features affected the multimodal results since we performed feature-level fusion. Exploring other forms of fusion like decision-level and hybrid may improve the results of the multimodal approach.

Our transfer learning approaches performed as well or better than the competition baseline approaches that used feature engineering (static features and BoAW). These results show that feature engineering is not necessary to get good emotion classification results for real-world speech data from older adults. This work focused on using pretrained models as feature extractors. Hence, we did not fine-tune the pretrained YAMNet and BERT models on this data. Doing so in the future could improve the recognition results.

Finally, this work is a key step towards recognizing the emotions of elderly individuals in daily life. We have collected speech and video data with self-reported emotion labels from German-speaking elderly individuals in their daily life after they underwent inpatient cardiovascular rehabilitation. Our future work will build upon this work and explore emotion recognition within that unique context.

## 5 CONCLUSIONS

In this work, we used a transfer learning approach to classify low, medium, and high emotion labels of the valence and arousal dimension of audio data collected from German-speaking elderly individuals. We used pretrained

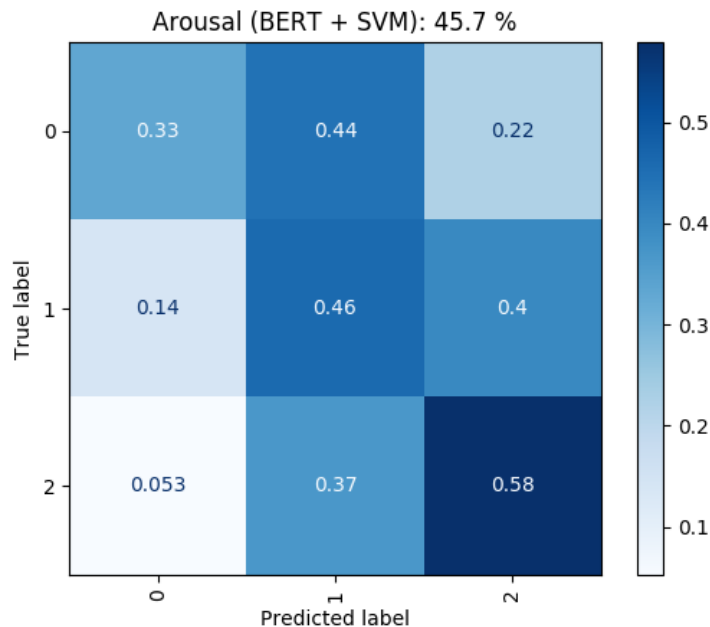


Fig. 3. Confusion matrix for development set evaluation for the best model for arousal: BERT + SVM: 45.7%

CNN and BERT models to extract acoustic and linguistic features respectively and fed them into separate machine learning models. Additionally, we fused the features in a multimodal approach and fed them to machine learning models. Our models using a linguistic approach performed better than the official competition baseline for the valence recognition task by 8.8%. Also, our results showed that feature engineering is not necessary and transfer learning can be leveraged to produce decent performance for the task of recognizing the emotions of elderly individuals. This work is a step towards better recognition of the emotions of the elderly which could eventually inform the development of interventions to manage their mental health.

## REFERENCES

- [1] [n.d.]. imbalanced-learn. <https://imbalanced-learn.readthedocs.io/en/stable/index.html>. Accessed: 2020-05-1.
- [2] [n.d.]. Open Sourcing German BERT. <https://deepset.ai/german-bert>. Accessed: 2020-05-1.
- [3] Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeannette N Chang, Sungbok Lee, and Shrikanth S Narayanan. 2008. IEMOCAP: Interactive emotional dyadic motion capture database. *Language resources and evaluation* 42, 4 (2008), 335.
- [4] Sandeep Nallan Chakravarthula, Haoqi Li, Shao-Yen Tseng, Maija Reblin, and Panayiotis Georgiou. 2019. Predicting Behavior in Cancer-Afflicted Patient and Spouse Interactions using Speech and Language. (2019).
- [5] François Chollet et al. 2018. Keras: The python deep learning library. *Astrophysics Source Code Library* (2018).
- [6] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [7] Kexin Feng and Theodora Chaspari. 2020. A Review of Generalizable Transfer Learning in Automatic Emotion Recognition. *Frontiers in Computer Science* 2 (2020), 9.
- [8] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter. 2017. Audio set: An ontology and human-labeled dataset for audio events. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 776–780.

- [9] Shawn Hershey, Sourish Chaudhuri, Daniel PW Ellis, Jort F Gemmeke, Aren Jansen, R Channing Moore, Manoj Plakal, Devin Platt, Rif A Saurous, Bryan Seybold, et al. 2017. CNN architectures for large-scale audio classification. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 131–135.
- [10] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [11] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. 2017. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861* (2017).
- [12] Jeremy Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. *arXiv preprint arXiv:1801.06146* (2018).
- [13] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. 2014. Large-scale video classification with convolutional neural networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*. 1725–1732.
- [14] Hong-Wei Ng, Viet Dung Nguyen, Vassilios Vonikakis, and Stefan Winkler. 2015. Deep learning for emotion recognition on small datasets using transfer learning. In *Proceedings of the 2015 ACM on international conference on multimodal interaction*. 443–449.
- [15] Charles C. Onu, Jonathan Lebensold, William L. Hamilton, and Doina Precup. 2019. Neural Transfer Learning for Cry-Based Diagnosis of Perinatal Asphyxia. In *Interspeech 2019, 20th Annual Conference of the International Speech Communication Association, Graz, Austria, 15-19 September 2019*, Gernot Kubin and Zdravko Kacic (Eds.). ISCA, 3053–3057. <https://doi.org/10.21437/Interspeech.2019-2340>
- [16] Maxime Oquab, Leon Bottou, Ivan Laptev, and Josef Sivic. 2014. Learning and transferring mid-level image representations using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1717–1724.
- [17] Sandeep Kumar Pandey, HS Shekhawat, and SRM Prasanna. 2019. Deep Learning Techniques for Speech Emotion Recognition: A Review. In *2019 29th International Conference Radioelektronika (RADIOELEKTRONIKA)*. IEEE, 1–6.
- [18] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*. 8024–8035.
- [19] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in Python. *the Journal of machine Learning research* 12 (2011), 2825–2830.
- [20] Soujanya Poria, Erik Cambria, Rajiv Bajpai, and Amir Hussain. 2017. A review of affective computing: From unimodal analysis to multimodal fusion. *Information Fusion* 37 (2017), 98–125.
- [21] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 3973–3983.
- [22] Nils Reimers and Iryna Gurevych. 2020. Making Monolingual Sentence Embeddings Multilingual using Knowledge Distillation. *arXiv preprint arXiv:2004.09813* (04 2020). <http://arxiv.org/abs/2004.09813>
- [23] Julian Risch, Anke Stoll, Marc Ziegele, and Ralf Krestel. 2019. hpiDEDIS at GermEval 2019: Offensive Language Identification using a German BERT model. In *Preliminary proceedings of the 15th Conference on Natural Language Processing (KONVENS 2019)*. Erlangen, Germany: German Society for Computational Linguistics & Language Technology. 403–408.
- [24] Wendy A Rogers and Tracy L Mitzner. 2017. Envisioning the future for older adults: Autonomy, health, well-being, and social connectedness with technology support. *Futures* 87 (2017), 133–139.
- [25] Sourav Sahoo, Puneet Kumar, Balasubramanian Raman, and Partha Pratim Roy. 2019. A Segment Level Approach to Speech Emotion Recognition Using Transfer Learning. In *Asian Conference on Pattern Recognition*. Springer, 435–448.
- [26] Bjorn W. Schuller, Anton Batliner, Christian Bergler, Eva-Maria Messner, Antonia Hamilton, Shahin Amiriparian, Alice Baird, Georgios Rizos, Maximilian Schmitt, Lukas Stappen, Harald Baumeister, Alexis Deighton MacIntyre, and Simone Hantke. 2020. The INTERSPEECH 2020 Computational Paralinguistics Challenge: Elderly Emotion, Breathing & Masks. In *Proceedings of Interspeech*. Shanghai, China, 5 pages. to appear.
- [27] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* 15, 1 (2014), 1929–1958.
- [28] Juanjuan Wang, Mantao Xu, Hui Wang, and Jiwu Zhang. 2006. Classification of imbalanced data by using the SMOTE algorithm and locally linear embedding. In *2006 8th international Conference on Signal Processing*, Vol. 3. IEEE.
- [29] Zixiaofan Yang and Julia Hirschberg. 2018. Predicting Arousal and Valence from Waveforms and Spectrograms Using Deep Neural Networks.. In *Interspeech*. 3092–3096.



## Paper 6: VADLite

# VADLite: An Open-Source Lightweight System for Real-Time Voice Activity Detection on Smartwatches

GEORGE BOATENG, ETH Zürich  
PRABHAKARAN SANTHANAM, ETH Zürich  
JANINA LÜSCHER, University of Zürich  
URTE SCHOLZ, University of Zürich  
TOBIAS KOWATSCH, ETH Zürich, University of St. Gallen

Smartwatches provide a unique opportunity to collect more speech data because they are always with the user and also have a more exposed microphone compared to smartphones. Speech data could be used to infer various indicators of mental well being such as emotions, stress and social activity. Hence, real-time voice activity detection (VAD) on smartwatches could enable the development of applications for mental health monitoring. In this work, we present VADLite, an open-source, lightweight, system that performs real-time VAD on smartwatches. It extracts mel-frequency cepstral coefficients and classifies speech versus non-speech audio samples using a linear Support Vector Machine. The real-time implementation is done on the Wear OS Polar M600 smartwatch. An offline and online evaluation of VADLite using real-world data showed better performance than WebRTC's open-source VAD system. VADLite can be easily integrated into Wear OS projects that need a lightweight VAD module running on a smartwatch.

CCS Concepts: • **Human-centered computing** → **Ubiquitous and mobile computing systems and tools**; • **Applied computing** → *Psychology*.

Additional Key Words and Phrases: Voice Activity Detection, Support Vector Machine, Wearable Computing, Smartwatch

## ACM Reference Format:

George Boateng, Prabhakaran Santhanam, Janina Lüscher, Urte Scholz, and Tobias Kowatsch. 2019. VADLite: An Open-Source Lightweight System for Real-Time Voice Activity Detection on Smartwatches. In *Adjunct Proceedings of the 2019 ACM International Joint Conference on Pervasive and Ubiquitous Computing and the 2019 International Symposium on Wearable Computers (UbiComp/ISWC '19 Adjunct)*, September 9–13, 2019, London, United Kingdom. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3341162.3346274>

## 1 INTRODUCTION

Smartwatches as a platform provide a unique opportunity to assesses the mental health of individuals because of their sensors which have high proximity to the human body. Multimodal sensing of gestures (from the gyroscope and accelerometer), heart rate, physical activity, ambient light, Bluetooth signal strength, among others could be used to infer the mental state of people [1].

On the specific topic of audio data, smartwatches provide a unique opportunity to collect more speech data in the everyday lives of individuals since they are more likely to always be with the user given they are worn on the wrist. Additionally, the microphone is also prone to be more exposed as compared to a smartphone, which might be in a bag or a pocket.

---

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

*UbiComp/ISWC '19 Adjunct*, September 9–13, 2019, London, United Kingdom

© 2019 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6869-8/19/09.

<https://doi.org/10.1145/3341162.3346274>

Speech data could be used to infer various indicators of mental well being such as emotions [27], stress [18] and social activity [29]. Hence, real-time voice activity detection (VAD) on smartwatches could enable the development of applications for mental health monitoring. Researchers and developers that need real-time VAD on smartwatches, which use the Wear OS operating system have to build their own custom module since an API is not provided.

Prior work have developed VAD systems but they have not focused on real-time implementation of the developed algorithms [10, 24, 30, 36]. Important aspects such as computational efficiency, latency and accuracy in a naturalistic context were not addressed. Hence, it is not clear how well they will perform if they are implemented to run in real-time on smartwatches.

On the other hand, there are VAD systems that have been implemented to run in real-time as a smartphone app [19, 34, 35]. Unfortunately, the machine learning models that were used are not easily available for others who want to simply use those pre-trained models in their work. It is also not clear how well the models will work when they run on smartwatches with reference to computational efficiency and latency.

There was another VAD system developed to run on a smartwatch as a component of a context recognizer [12]. The authors use 13 mel-frequency cepstral coefficients (MFCC) features and a convolutional neural network. Using three seconds of data, it takes 1.9 seconds to give a classification. The performance evaluation (e.g., accuracy) is not provided and most importantly, the software component is not available for others to easily use in their smartwatch-based sensing work.

Then there are also open-source VAD systems such as WebRTC's VAD [13]. It is, however, a computer-based system, which does not have a module for smartwatches. Additionally, it has been reported that WebRTC's VAD performs poorly on real-world data collected with smartwatches [17].

Given the gap in obtaining an easy-to-use smartwatch-based VAD system, we developed VADLite, an open-source lightweight software system that performs real-time voice activity detection on smartwatches. VADLite extracts MFCC as features and classifies speech versus non-speech audio samples using a linear Support Vector Machine (SVM). We designed VADLite to meet our specific requirements; it is lightweight (i.e., runs efficiently on a constrained system such as a smartwatch), and also performs well in a real-world context in which we will be deploying the devices. In this work, we describe the process of developing and evaluating VADLite, and comparing its performance to WebRTC's VAD system since it is a widely used open-source VAD system.

The real-time implementation is done on a Wear OS smartwatch and our project files with the source code are available for use by others <sup>1</sup>. VADLite can be used by including the Java source code files in a Wear OS project. Also, the parameters of the trained model are in the Java files and hence, users can use our model parameters to build their own VAD pipeline for their Wear OS projects.

The rest of this paper is organized as follows. First, we give the motivation and use case of VADLite. Next, we give an overview and describe the development of VADLite. Then, we describe the real-time implementation of VADLite. After, we detail our experiments and results including a comparison of VADLite with WebRTC's VAD system. We address ethical implications and privacy concerns of our work. Finally, we summarize and conclude.

## 2 MOTIVATION

Our primary motivation for developing VADLite was our previous work DyMand, an open-source mobile and wearable system for assessment of couples' dyadic management of chronic diseases in everyday life [7, 8]. The DyMand system collects self-report data about health behavior, and emotions, and sensor data about couples' dyadic management of chronic diseases. DyMand first determines if the partners are close using the Bluetooth signal strength between their watches. Then, VADLite is used as the next optimizing step to trigger collection of multimodal sensor data (heart rate, accelerometer, gyroscope, and ambient light) ensuring that data is collected

<sup>1</sup><https://bitbucket.org/Jojo29/vadlite/>

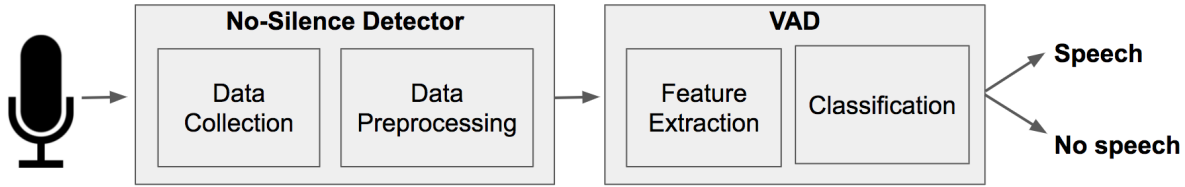


Fig. 1. Overview of VADLite System

only when the partners are speaking. This two-step process in which VADLite plays a key role ensures that DyMand collects more relevant speech data which is an improvement over how social psychologists currently collect ambulatory audio data for analysis: triggering data collection at random times of the day. [21, 22, 31, 32, 37].

Our secondary motivation for developing VADLite is for it to be used in the development of a real-time smartwatch-based app for recognizing emotions of couples using speech prosody and the semantics of speech [27]. Emotion recognition from speech will then be used in combination with other sensor modalities from the smartwatch to perform real-time multimodal emotion recognition among couples [6]. Prior work has shown that social support among couples results in better health behavior when one partner has diabetes and it also affects the emotions of the couple [15, 23, 28]. Real-time emotion recognition among couples would give an assessment of a key outcome of social support which could be used to develop just-in-time adaptive interventions [25] to enable couples better manage chronic diseases. In order to accomplish that goal, speech episodes in everyday life need to be recognized accurately and efficiently. VADLite fills that gap.

Beyond these specific use cases, real-time VAD in combination with other sensors could be used to infer social isolation or a lack of social activity, which is a predictor of mental health issues such as depression or suicidal ideation [29, 38]. An accurate measure of speech data could enable better prediction of social isolation. Using VADLite which runs on a smartwatch will adequately enable the accomplishment of this goal.

### 3 OVERVIEW AND DEVELOPMENT OF VADLITE

VADLite is a 2-stage system consisting of a no-silence detector as the first part, and a voice activity detector as the second part (Figure 1). In developing VADLite, we used the pipeline of data collection and preprocessing, feature extraction and classification. We used a linear SVM. An SVM is a classifier that constructs a high-dimensional hyperplane to separate data of different classes [14]. SVM selects a hyperplane that maximizes the distance to the nearest data points on either side of the hyperplane in the case of binary classification. Previous work have used SVM for VAD successfully [2, 9, 10, 16, 30].

We used a linear SVM because it is memory and computationally efficient when incoming data is classified. For example, in comparison with a linear SVM, a radial basis function (RBF) SVM though only slightly outperformed a linear SVM for VAD took twice as much time for classification [16]. Prior work have used an implementation of linear SVM for real-time prediction on smartwatches for stress detection [5] and activity detection [3, 4].

#### 3.1 Data Collection

We collected real-world data using a protocol that was approved by the ethics commission of ETH Zurich. We collected data using a Polar M600 smartwatch where subjects (1) in the lab read a written text as a smartwatch recorded audio data for approximately 1-2 mins and (2) in the everyday life wore a smartwatch during waking hours as it continuously collected audio data. We used 16-PCM mono audio data and a sampling frequency of 8KHz. The data was annotated as speech or non-speech data. The speech data contained mostly conversations

among several people (with at least 10 distinct speakers) at varying distances from the smartwatch's microphone. The non-speech data contained sounds from cars, trams, buses, wind, and music. The overall duration of the recorded sound data was 3.5 hours.

### 3.2 Data Preprocessing

We processed the data by first removing silence portions of the data using a one-second time window. Liaquat et al. found out that the real-world audio data they collected contained about 61.7% of silence and hence they implemented a silence detection algorithm to remove the silence part of their data [17]. Given that we used real-world audio data, we also removed silence segments of the data. We computed the root mean square (RMS) of each one-second time window of the whole data. We then checked if the RMS value is below a certain threshold, in which case we marked that segment as silence and then removed it. To determine the threshold, we created a scatter plot of the RMS values of silence, speech and noise signals and then chose the value that separates silence from both speech and noise.

### 3.3 Feature Extraction

We extracted 13 MFCC features and use 12 of them (excluding the 1st coefficient, which is the DC component) over a time window of 25 ms. MFCC features have been widely used for VAD [12]. The parameters we used are as follows: 8KHz sample rate, window length of 25 ms, window step, 12 coefficients, 26 filters in the filterbank, FFT size of 512, 0 Hz as lowest band edge of mel filters, 4KHz as highest band edge of mel filter of (i.e., half the sampling rate), 22 lifters to apply to final cepstral coefficients and a Hamming windowing function. We used a Java implementation for the feature extraction.

### 3.4 Classification

Using the feature sets, we trained a linear SVM to classify speech or non-speech. We also performed grid search to pick the most optimal hyper-parameters of the linear SVM. We first normalized the features by subtracting the mean and dividing by the variance. This normalization is important for various algorithms such as SVM whose optimization assume that the features have a normal distribution [14]. We used the following metrics for evaluation: accuracy, speech hit rate (SHR), and false alarm rate (FAR). The SHR is the ratio of correctly detected speech frames to the total number of speech frames. By contrast, FAR is one minus the noise hit rate, where noise hit rate is the ratio of correctly detected noise frames to the total number of noise frames.

## 4 REAL-TIME IMPLEMENTATION

We coded VADLite in Java for smartwatches that use the Wear OS operating system (formerly Android Wear). We used the Android Studio Integrated Development Environment (IDE) and the Android Software Development Environment (SDK). VADLite can potentially work on every Wear OS device. We used a Polar M600 smartwatch running Wear OS version 2.1 for testing which has the following specifications: Dual-Core 1.2GHz processor based on ARM Cortex-A7, 512MB RAM, 4GB flash storage, 500 mAh Battery.

Our implementation of VADLite is a Wear OS app, which collects 16-PCM mono audio data every second at a frequency of 8KHz (see Figure 2). We check if the one-second data is a non-silence segment by using an implementation of the non-silence detector from the previous section. We then process the data if it is non-silence. The one-second non-silence signal is then segmented into 25 ms frames. We then extract 12 MFCC features for each frame, which is then fed to a linear SVM for classification. We used the settings described in the previous section for the feature extraction. We used the Java implementation from the offline evaluation for extracting the MFCC features online.

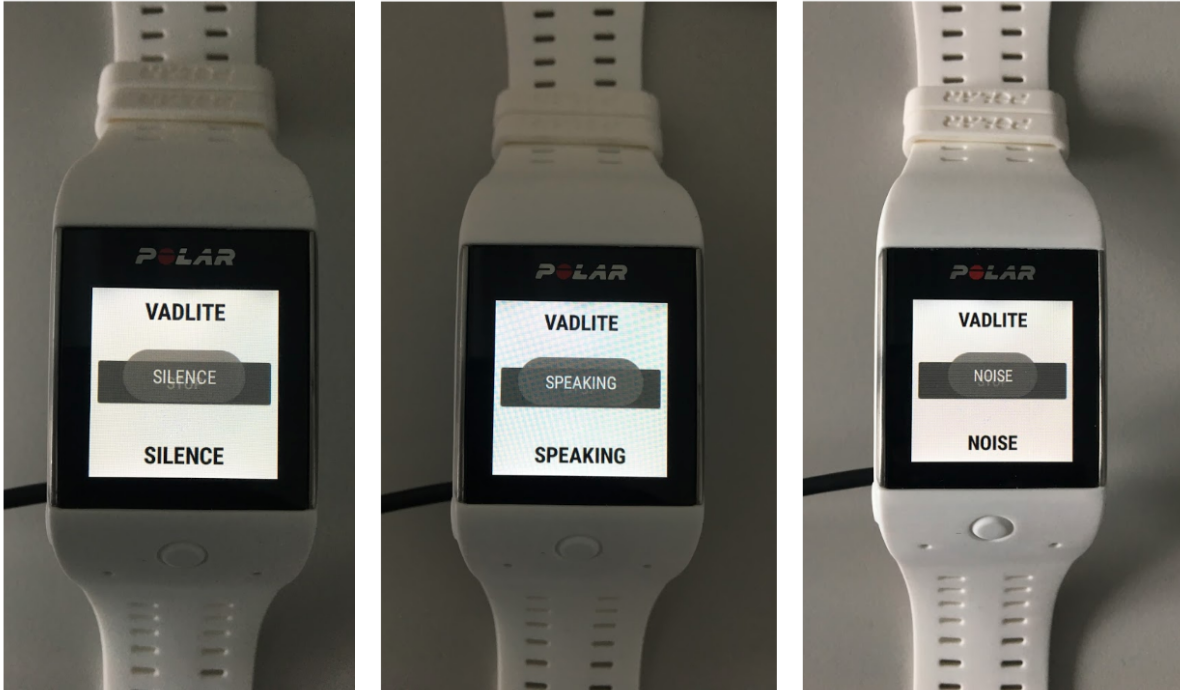


Fig. 2. Real-time running of VADLite

We normalized the features using the stored normalization vectors before performing classification with a linear SVM. Our implementation of the linear SVM is a dot product of the stored coefficients with the features:

$$y = wx + b \quad (1)$$

where  $y$  is the result of the evaluation,  $w$  is the coefficient vector of length 12, and  $b$  is the intercept. We then assigned  $y$  to be 1 (speech) if it is greater than zero, otherwise we assign it to be 0 (non-speech). We obtained  $w$  and  $b$  from the previously trained linear SVM. We output speech or non-speech classification for the whole one-second data. To accomplish this, we used majority voting of all the classified 25 ms samples within the one-second data.

VADLite had an average processing time of 2 ms for each 25 ms frame and 76 ms for the total one-second duration. As a result, throughput was met since the frame processing time was less than the 25 ms segment duration. Likewise, the processing time for whole duration was less than one second.

## 5 EXPERIMENTS AND RESULTS

We evaluated VADLite offline and also online. Additionally, we compared the classification performance of VADLite with a popular open-source VAD system, WebRTC's VAD [13]. WebRTC's VAD uses frequency band features and a pre-trained Gaussian Mixture Model (GMM) classifier [33]. We used a Python implementation of the system [39]. It gives the option to set an aggressiveness mode using an integer from zero to three with zero being the least aggressive about filtering non-speech audio. It only accepts 16-bit PCM mono audio sampled at

Table 1. Evaluation results from the offline evaluation of VADLite and WebRTC’s VAD

<i>Model</i>	<i>Accuracy</i>	<i>SHR(%)</i>	<i>FAR(%)</i>
WebRTC(0)	66.7	91.6	60.7
WebRTC(1)	71.7	89.0	47.4
WebRTC(2)	71.4	79.5	37.5
WebRTC(3)	75.9	55.3	2.1
VADLite	83.7	83.4	16.0

8KHz, 16KHz, 32KHz or 48KHz. It also processes data in frames with a duration of either 10, 20 or 30 ms. Our implementation used 8 KHz sampling rate and 20 ms time window to match the settings of VADLite.

### 5.1 Offline Evaluation

We split the data into train and test using about 70%-30% split. The speech and noise train data were 73.9 and 71.8 minutes long respectively. The speech and noise test were 24.6 and 22.4 minutes long respectively. We performed 10-fold stratified cross-validation on the train data using VADLite’s linear SVM model. We used the scikit-learn library for our experiments [26]. The model achieved 82.6% accuracy, 80.2% SHR and 14.9% FAR. We then trained the VADLite model on the whole train dataset and then we used the test data to evaluate both the VADLite’s model and those of WebRTC’s VAD. The results of the evaluation are shown in Table 1.

VADLite’s model outperforms WebRTC’s VAD when its aggressiveness mode is two and three. WebRTC’s VAD with settings zero and one though have very high SHR, their FAR are high, which will result in a lot of noise being classified as speech, which is not acceptable. VADLite’s model provides a good enough tradeoff between SHR and FAR. These results indicate that VADLite is better than WebRTC’s VAD. These results support those by Liaqat et al. who found that WebRTC’s VAD performed poorly on real-world smartwatch-based audio [17].

### 5.2 Online Evaluation

To evaluate the real-time performance of VADLite with real-world data, we recorded audio data from a naturalistic context. We then played the recorded audio through a loudspeaker as the VADLite app performed real-time classification of the audio just like was done by Feng et al [11]. The audio had a duration of 15 minutes each for speech and noise. We stored the classification and compared it with real labels of the audio. We report the classification results below. We also ran the audio data through WebRTC’s VAD. VADLite had SHR and FAR of 91.6% and 5.5% respectively. WebRTC’s VAD’s best performing mode had SHR and FAR of 73% and 18% respectively. Consistent with the results from the offline evaluation, VADLite outperforms WebRTC’s VAD. Once again, these results supports those by Liaqat et al. who found that WebRTC’s VAD performed poorly on real-world smartwatch-based audio [17].

## 6 ETHICAL IMPLICATIONS AND PRIVACY CONCERNS

This work has ethical implications as the system could be used in a manner that violates the privacy of others. We envision that this system can be used in two main ways.

The first is that it could be used to collect raw speech data from subjects, which will be stored for processing later. For this approach, it is especially important that the study protocol is subjected to review and approval from the ethics committee of the overseeing institution, as is standard practice. And additional steps need to be

taken such as giving subjects the option to listen to the recorded audio and to delete any as they wish without any explanation. This approach has been used in our studies [20] and those of others [31, 32]. Depending on the use case, the app may need to give subjects the option to completely disable audio recording as needed. Also, to protect the privacy of subjects not taking part in the study, it might be necessary to have subjects wear a tag indicating to others around them that they may be recorded.

The other way we envision this system being used is to derive important summary statistics. In this case, no raw audio will be stored. Rather, various inference such as conversation frequency and duration (total duration of speech per day, times of the day with most speech etc.) will be computed [29]. This use case is less invasive but again as usual, ethical approval needs to be obtained for such a study since summary data could be considered personal and private for some subjects.

## 7 FUTURE WORK

VADLite could be extended so that it additionally makes various inference such as conversation frequency and duration. These additions would make VADLite more useful in smartwatch-based applications that seek to improve the mental well-being of people.

## 8 CONCLUSION

In this work, we developed VADLite an open-source lightweight software system for real-time VAD on smartwatches. VADLite uses MFCC as features and classifies speech versus non-speech audio samples using linear SVM with a real-time implementation on a Wear OS smartwatch. Our evaluation of VADLite showed SHR and FAR of 83.4% and 16.0% respectively for offline, and 91.6% and 5.5% respectively for real-time classification. Benchmarking of our system against WebRTC's VAD showed better performance. Our open-source system, VADLite can be easily integrated into Wear OS projects that need a lightweight voice activity module running on a smartwatch. VADLite can be integrated into the development of various well-being specific apps.

## ACKNOWLEDGMENTS

We are grateful to Kemeng Zhang who helped collect the audio data. We thank Filipe Barata and Chen-Hsuan Shih for their feedback during this work. This project was funded by a grant from the Swiss National Science Foundation with funding number CR12I1\_166348/1.

## REFERENCES

- [1] Saeed Abdullah and Tanzeem Choudhury. 2018. Sensing technologies for monitoring serious mental illnesses. *IEEE MultiMedia* 25, 1 (2018), 61–75.
- [2] M Baig, S Masud, and M Awais. 2006. Support vector machine based voice activity detection. In *2006 International Symposium on Intelligent Signal Processing and Communications*. IEEE, 319–322.
- [3] George Boateng, John A Batsis, Ryan Halter, and David Kotz. 2017. ActivityAware: an app for real-time daily activity level monitoring on the amulet wrist-worn device. In *2017 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*. IEEE, 431–435.
- [4] George Boateng, John A Batsis, Patrick Proctor, Ryan Halter, and David Kotz. 2018. GeriActive: Wearable app for monitoring and encouraging physical activity among older adults. In *2018 IEEE 15th International Conference on Wearable and Implantable Body Sensor Networks (BSN)*. IEEE, 46–49.
- [5] George Boateng and David Kotz. 2016. StressAware: An app for real-time stress monitoring on the amulet wearable platform. In *2016 IEEE MIT Undergraduate Research Technology Conference (URTC)*. IEEE, 1–4.
- [6] George Boateng, Prabhakaran Santhanam, Janina Lüscher, Urte Scholz, and Tobias Kowatsch. 2018. Multimodal Affect Detection among Couples for Diabetes Management. In *Poster: 2nd Black in AI Workshop at the 32nd Annual Conference on Neural Information Processing Systems (NeurIPS)*.



- [7] George Boateng, Prabhakaran Santhanam, Janina Lüscher, Urte Scholz, and Tobias Kowatsch. 2019. DyMand: An Open-Source Mobile and Wearable System for Assessing Couples' Dyadic Management of Chronic Diseases. In *14th International Conference on Design Science Research in Information System and Technology (DESIRIST)*.
- [8] George Boateng, Prabhakaran Santhanam, Janina Lüscher, Urte Scholz, and Tobias Kowatsch. 2019. Poster: DyMand – An Open-Source Mobile and Wearable System for Assessing Couples' Dyadic Management of Chronic Diseases. In *The 25th Annual International Conference on Mobile Computing and Networking (MobiCom)*.
- [9] Cheng Dai, Linkai Luo, Hong Peng, and Qingyun Sun. 2018. A Method Based on Support Vector Machine for Voice Activity Detection on Isolated Words. In *2018 13th International Conference on Computer Science & Education (ICCSE)*. IEEE, 1–4.
- [10] Dong Enqing, Liu Guizhong, Zhou Yatong, and Zhang Xiaodi. 2002. Applying support vector machines to voice activity detection. In *6th International Conference on Signal Processing, 2002.*, Vol. 2. IEEE, 1124–1127.
- [11] Tiantian Feng, Amrutha Nadarajan, Colin Vaz, Brandon Booth, and Shrikanth Narayanan. 2018. TILES audio recorder: an unobtrusive wearable solution to track audio activity. In *Proceedings of the 4th ACM Workshop on Wearable Systems and Applications*. ACM, 33–38.
- [12] Claudio Forlivesi, Utku Günay Acer, Marc van den Broeck, and Fahim Kawsar. 2018. Mindful interruptions: a lightweight system for managing interruptibility on wearables. In *Proceedings of the 4th ACM Workshop on Wearable Systems and Applications*. ACM, 27–32.
- [13] Google. 2011. WebRTC. Retrieved April, 2019 from <https://webrtc.org/>
- [14] Chih-Wei Hsu, Chih-Chung Chang, Chih-Jen Lin, et al. 2003. A practical guide to support vector classification. (2003).
- [15] Masumi Iida, Mary Ann Parris Stephens, Karen S Rook, Melissa M Franks, and James K Salem. 2010. When the going gets tough, does support get going? Determinants of spousal support provision to type 2 diabetic patients. *Personality and Social Psychology Bulletin* 36, 6 (2010), 780–791.
- [16] Tomi Kinnunen, Evgenia Chernenko, Marko Tuononen, Pasi Fränti, and Haizhou Li. 2007. Voice activity detection using MFCC features and support vector machine. In *Int. Conf. on Speech and Computer (SPECOM07), Moscow, Russia*, Vol. 2. 556–561.
- [17] Daniyal Liaqat, Robert Wu, Andrea Gershon, Hisham Alshaer, Frank Rudzicz, and Eyal de Lara. 2018. Challenges with real-world smartwatch based audio monitoring. In *Proceedings of the 4th ACM Workshop on Wearable Systems and Applications*. ACM, 54–59.
- [18] Hong Lu, Denise Frauendorfer, Mashfiqui Rabbi, Marianne Schmid Mast, Gokul T Chittaranjan, Andrew T Campbell, Daniel Gatica-Perez, and Tanzeem Choudhury. 2012. Stresssense: Detecting stress in unconstrained acoustic environments using smartphones. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*. ACM, 351–360.
- [19] Hong Lu, Wei Pan, Nicholas D Lane, Tanzeem Choudhury, and Andrew T Campbell. 2009. SoundSense: scalable sound sensing for people-centric applications on mobile phones. In *Proceedings of the 7th international conference on Mobile systems, applications, and services*. ACM, 165–178.
- [20] Janina Lüscher, Tobias Kowatsch, George Boateng, Prabhakaran Santhanam, and Urte Scholz. 2019. Social Support and Common Dyadic Coping in Couple's Dyadic Management of Type II Diabetes: Study Protocol for an Ambulatory Assessment Application. In *JMIR Protocols*. JMIR.
- [21] Matthias R Mehl, James W Pennebaker, D Michael Crow, James Dabbs, and John H Price. 2001. The Electronically Activated Recorder (EAR): A device for sampling naturalistic daily activities and conversations. *Behavior Research Methods, Instruments, & Computers* 33, 4 (2001), 517–523.
- [22] Matthias R Mehl, Megan L Robbins, and Fenne große Deters. 2012. Naturalistic observation of health-relevant social processes: The Electronically Activated Recorder (EAR) methodology in psychosomatics. *Psychosomatic Medicine* 74, 4 (2012), 410.
- [23] Daisy Miller and J Lynne Brown. 2005. Marital interactions in the process of dietary change for type 2 diabetes. *Journal of Nutrition Education and Behavior* 37, 5 (2005), 226–234.
- [24] Mohammad H Moattar, Mohammad M Homayounpour, and Nima Khademi Kalantari. 2010. A new approach for robust realtime voice activity detection using spectral pattern. In *2010 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 4478–4481.
- [25] Inbal Nahum-Shani, Shawna N Smith, Bonnie J Spring, Linda M Collins, Katie Witkiewitz, Ambuj Tewari, and Susan A Murphy. 2017. Just-in-time adaptive interventions (JITAs) in mobile health: key components and design principles for ongoing health behavior support. *Annals of Behavioral Medicine* 52, 6 (2017), 446–462.
- [26] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in Python. *Journal of machine learning research* 12, Oct (2011), 2825–2830.
- [27] Soujanya Poria, Erik Cambria, Rajiv Bajpai, and Amir Hussain. 2017. A review of affective computing: From unimodal analysis to multimodal fusion. *Information Fusion* 37 (2017), 98–125.
- [28] Gabriele Prati and Luca Pietrantonio. 2010. The relation of perceived and received social support to mental health among first responders: a meta-analytic review. *Journal of Community Psychology* 38, 3 (2010), 403–417.
- [29] Mashfiqui Rabbi, Shahid Ali, Tanzeem Choudhury, and Ethan Berke. 2011. Passive and in-situ assessment of mental and physical well-being using mobile sensors. In *Proceedings of the 13th international conference on Ubiquitous computing*. ACM, 385–394.

- [30] J Ramírez, P Yélamos, JM Górriz, and JC Segura. 2006. SVM-based speech endpoint detection using contextual speech features. *Electronics letters* 42, 7 (2006), 426–428.
- [31] Megan L Robbins, Elizabeth S Focella, Shelley Kasle, Ana María López, Karen L Weihs, and Matthias R Mehl. 2011. Naturalistically observed swearing, emotional support, and depressive symptoms in women coping with illness. *Health Psychology* 30, 6 (2011), 789.
- [32] Megan L Robbins, Ana María López, Karen L Weihs, and Matthias R Mehl. 2014. Cancer conversations in context: Naturalistic observation of couples coping with breast cancer. *Journal of Family Psychology* 28, 3 (2014), 380.
- [33] Sergey Salishev, Andrey Barabanov, Daniil Kocharov, Pavel Skrelin, and Mikhail Moiseev. 2016. Voice Activity Detector (VAD) Based on Long-Term Mel Frequency Band Features. In *International Conference on Text, Speech, and Dialogue*. Springer, 352–358.
- [34] Abhishek Sehgal and Nasser Kehtarnavaz. 2018. A convolutional neural network smartphone app for real-time voice activity detection. *IEEE Access* 6 (2018), 9017–9026.
- [35] Abhishek Sehgal, Fatemeh Saki, and Nasser Kehtarnavaz. 2017. Real-time implementation of voice activity detector on ARM embedded processor of smartphones. In *2017 IEEE 26th International Symposium on Industrial Electronics (ISIE)*. IEEE, 1285–1290.
- [36] Jongseo Sohn, Nam Soo Kim, and Wonyong Sung. 1999. A statistical model-based voice activity detection. *IEEE signal processing letters* 6, 1 (1999), 1–3.
- [37] Adela C Timmons, Theodora Chaspari, Sohyun C Han, Laura Perrone, Shrikanth S Narayanan, and Gayla Margolin. 2017. Using multimodal wearable technology to detect conflict among couples. *Computer* 50, 3 (2017), 50–59.
- [38] Fabian Wahle, Tobias Kowatsch, Elgar Fleisch, Michael Rufer, and Steffi Weidt. 2016. Mobile sensing and support for people with depression: a pilot trial in the wild. *JMIR mHealth and uHealth* 4, 3 (2016), e111.
- [39] John Wiseman. 2016. Python interface to the WebRTC Voice Activity Detector. Retrieved April, 2019 from <https://github.com/wiseman/py-webrtcvad>

## Paper 7: DyMand

# Development, Deployment, and Evaluation of DyMand – An Open-Source Smartwatch and Smartphone System for Capturing Couples’ Dyadic Interactions in Chronic Disease Management in Daily Life

GEORGE BOATENG, ETH Zürich, Switzerland

PRABHAKARAN SANTHANAM, ETH Zürich, Switzerland

ELGAR FLEISCH, ETH Zürich, Switzerland and University of St. Gallen, Switzerland

JANINA LÜSCHER, University of Zürich, Switzerland

THERESA PAULY, University of Zürich, Switzerland

URTE SCHOLZ, University of Zürich, Switzerland

GUY BODENMANN, University of Zurich, Switzerland

TOBIAS KOWATSCH, ETH Zürich, Switzerland and University of St. Gallen, Switzerland

Dyadic interactions of couples are of interest as they provide insight into relationship quality and chronic disease management. Currently, ambulatory assessment of couples’ interactions entails collecting data at random or scheduled times which could miss significant couples’ interaction/conversation moments. In this work, we developed, deployed and evaluated DyMand, a novel open-source smartwatch and smartphone system for collecting self-report and sensor data from couples based on partners’ interaction moments. Our smartwatch-based algorithm uses the Bluetooth signal strength between two smartwatches each worn by one partner, and a voice activity detection machine-learning algorithm to infer that the partners are interacting, and then to trigger data collection. We deployed the DyMand system in a 7-day field study and collected data about social support, emotional well-being, and health behavior from 13 (N=26) Swiss-based heterosexual couples managing diabetes mellitus type 2 of one partner. Our system triggered 99.1% of the expected number of sensor and self-report data when the app was running, and 77.6% of algorithm-triggered recordings contained partners’ conversation moments compared to 43.8% for scheduled triggers. The usability evaluation showed that DyMand was easy to use. DyMand can be used by social, clinical, or health psychology researchers to understand the social dynamics of couples in everyday life, and for developing and delivering behavioral interventions for couples who are managing chronic diseases.

CCS Concepts: • **Human-centered computing** → **Ubiquitous and mobile computing**; • **Applied computing** → *Consumer health; Psychology*.

Additional Key Words and Phrases: Multimodal Sensor Data; Couples; Smartwatches; Smartphones; Mobile Computing; Wearable Computing; Machine Learning; Speech Processing; Chronic Disease Management; Social Support

## 1 INTRODUCTION

Romantic relationships have powerful effects on people’s mental and physical health (see [51] for an overview). For instance, conflicts and negative qualities of one’s intimate relationship are associated prospectively with morbidity and mortality [41]. Romantic or social relationships play an important role in illness management if partners share the responsibility and consider the disease to be their joint problem instead of being only the problem of the afflicted partner [48, 55] and it can involve social support and common dyadic coping (CDC) [14]. Social support entails providing resources to help a receiver cope in a time of need and can be emotional (e.g.,

---

Authors’ addresses: George Boateng, gboateng@ethz.ch, ETH Zürich, Zurich, Switzerland; Prabhakaran Santhanam, psanthanam@ethz.ch, ETH Zürich, Zurich, Switzerland; Elgar Fleisch, efleisch@ethz.ch, ETH Zürich, Zurich, Switzerland and University of St. Gallen, St. Gallen, Switzerland; Janina Lüscher, janina.luescher@psychologie.uzh.ch, University of Zürich, Zurich, Switzerland; Theresa Pauly, theresa.pauly@psychologie.uzh.ch, University of Zürich, Zurich, Switzerland; Urte Scholz, urte.scholz@psychologie.uzh.ch, University of Zürich, Zurich, Switzerland; Guy Bodenmann, guy.bodenmann@psychologie.uzh.ch, University of Zurich, Zurich, Switzerland; Tobias Kowatsch, tkowatsch@ethz.ch, ETH Zürich, Zurich, Switzerland and University of St. Gallen, St. Gallen, Switzerland.

providing comfort or encouragement) or instrumental (e.g., help with practical problems and tasks; [24, 43, 54]). CDC is the "we approach" to dealing with stressors in a couple's relationship [22] which can be assessed objectively by counting first-person plural pronouns [52], by questionnaire or behavioral observation. Social support among couples and CDC in chronic disease management have been shown to have mostly positive effects on emotional well-being [23, 32, 53, 60], and result in healthier eating habits among diabetes patients [46]. Consequently, it is of interest to better understand couples' dyadic interactions in-situ, for example, in couples' management of diabetes in daily life [35, 42] as they could enable the development and delivery of behavioral interventions to, for example, improve physical activity, diet, and medication adherence.

Ubiquitous devices such as smartphones and smartwatches provide a good opportunity to collect relevant data such as sensor and self-report data from couples in daily life. Smartwatches in particular could be leveraged to collect data on couples' dyadic interactions and chronic disease management. Several features of smartwatches make them uniquely positioned for this task. Firstly, they are mostly with the wearer since they are worn on the wrist in comparison with a smartphone which could be in various places like the pocket, or bag, and just not in proximity to the user, or devices like Amazon Echo or Google Home which can only be in one place and not always around the owners. Additionally, commercial smartwatches could be used to collect a wide variety of sensor data such as audio and heart rate (for stress detection, emotion recognition), Bluetooth (for proximity detection), accelerometer, and gyroscope (for gestures and physical activity), and ambient light (to detect context). Finally and importantly, smartwatches could be leveraged in novel ways to capture dyadic interactions of partners (e.g., triggering data collection when partners are close and speaking) as we do in this work.

In this work, we describe the development, deployment, and evaluation of DyMand, a novel open-source smartwatch<sup>1</sup>, and smartphone<sup>2</sup> system for ambulatory assessment of couples' **Dyadic Management** of chronic diseases in daily life. The DyMand system collects self-report and sensor data based on partners' interaction moments. In particular, we developed a smartwatch-based algorithm that uses the Bluetooth signal strength between two smartwatches each worn by one partner, and a voice activity detection machine-learning algorithm to infer that the partners are interacting, and then to trigger data collection. We deployed the system in a field study with heterosexual couples in Switzerland that are managing type 2 diabetes (T2DM) of one partner, a common chronic disease affecting 6% of the Swiss population [31]. The specific use case in this work was to collect data to understand the association between multimodal sensor data and self-report data of social support, and CDC in the context of diabetes management. This understanding will provide a sound basis for theory- and evidence-based development of dyadic interventions in the context of couples' dyadic illness management. The DyMand system can be used by, for example, social psychologists to understand the social dynamics of couples in everyday life and their impact on relationship quality, and also by clinical or health psychologists for developing and delivering behavioral interventions for couples who are managing chronic diseases. This work builds upon a study protocol published in 2019 [42] and it is an extension of our prior work [16, 20] and it includes a more detailed description of our DyMand system and its real-world deployment and evaluation.

This paper is organized as follows. Next, we discuss related work (Section 2). We then describe the system design (Section 3) and its implementation (Section 4). In Section 5, we describe the deployment of the system in a user study and, in Section 6, we evaluate its technical performance and usability. Finally, we describe limitations and future work in Section 7 and conclude this work in Section 8 with a summary.

## 2 RELATED WORK

Various smartphone applications have been developed for ambulatory data collection by social and health psychologists. For example, the Electronic Activated Recorder (EAR) has been used in several studies [44, 49],

<sup>1</sup><https://bitbucket.org/mobilecoach/dymandwatchclient/src/master/>

<sup>2</sup><https://bitbucket.org/mobilecoach/dymand-mobilecoach-client/src/master/>

especially for the collection of audio data in various couples' interactions such as couples managing breast cancer [34, 50]. The EAR triggers data collection at random or scheduled times in the day and collects snippets of ambient sound (e.g., 50 seconds every 9 minutes), which are later transcribed and coded. The EAR does not collect self-report data. On the other hand, a mobile and wearable system was used to collect sensor and self-report data for conflict detection among couples [57, 58]. Similar to the EAR, audio data were collected at random or scheduled times in the day (3 minutes of audio every 12 minutes). Another work [47] used a digital recorder for a whole-day recording of couples managing cancer.

Despite these advances in the ambulatory assessment of couples' interactions, there are still gaps. Firstly, any random or scheduled triggering of data collection does not take advantage of the dyadic nature of couples interactions (e.g., by inferring if partners are actually interacting) and could miss key conversations/interaction moments. The EAR collects only audio and does not leverage potentially more information from other sensor data or self-reports. The all-day recording of [47] has significant privacy concerns. Consequently, there is currently no ubiquitous system that leverages the dyadic nature of couples' interaction for the collection of sensor and self-report data that are relevant for social interactions and chronic disease management in everyday life.

### 3 DEVELOPMENT: SYSTEM DESIGN

In developing DyMand, experts from the field of computer science, information systems, and health psychology used justificatory knowledge from prior work [19, 32, 42, 46, 58] about social support, CDC, health behavior, and emotional well-being to derive a list of design specifications that are important for collecting corresponding data in-situ, in the context of chronic disease management. We describe the specifications.

#### 3.1 Physical Closeness Monitoring

The system should track the physical closeness of the partners during waking hours. This information can be used to infer how much time romantic partners spend in various forms of daily interactions (engaging in a shared activity, talking, and arguing), which could be used to predict the couple's relationship outcomes [28]. Furthermore, physical closeness tracking can be used in combination with other kinds of data to capture moments when partners are interacting, which would enable the collection of data relevant to chronic disease management such as emotional well-being, social support, and CDC based on partners' interaction context.

#### 3.2 Multimodal Sensor Data Collection

The system should collect relevant 5-minutes worth of multimodal sensor data (in particular audio) once each hour during waking hours (set by couples), ideally when partners are interacting. We restricted the data collection requirement so that we collect only one 5-minute sample per hour (audio data) for privacy reasons to ensure that we do not collect too much audio of the couples' daily life. Furthermore, the requirement for collecting this data when partners are interacting ensures there is a high likelihood of capturing conversations between partners as compared to collecting data at a random or scheduled time. Sensor data such as audio, heart rate, gestures, physical activity, and step count could be used to manually and automatically infer behavioral information such as social support, CDC, emotional well-being, and health behavior, which are relevant for chronic disease management [42]. Audio for example can be used to code constructs such as emotions, social support, and CDC. Also, audio together with other data such as heart rate and movements data could be used to automatically detect the emotions of each partner [13].

#### 3.3 Self-report Data Collection

The system should collect self-report data immediately after sensor data collection and at the end of the day. This requirement ensures that various validated self-reported instruments also collect relevant data about social

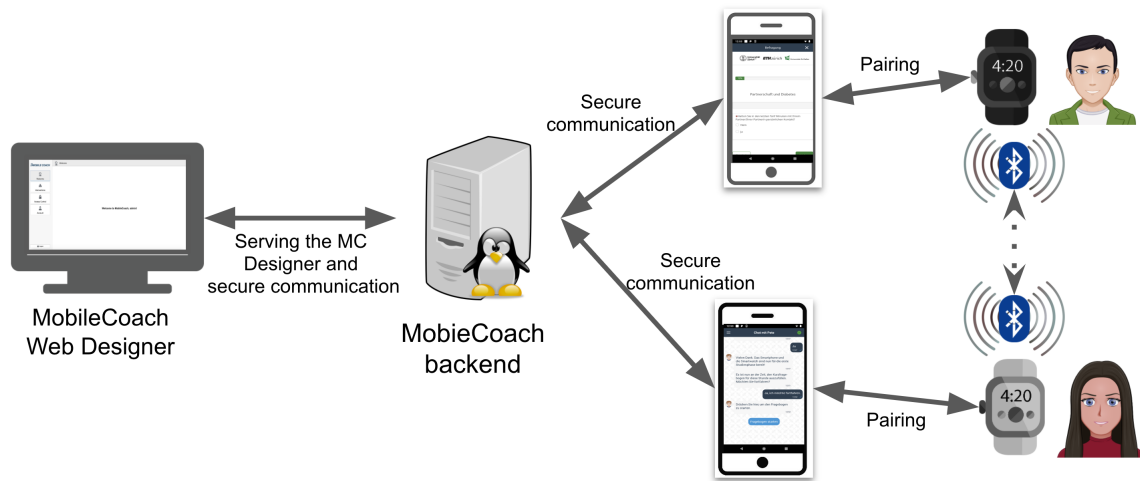


Fig. 1. Overview of the DyMand system

support, CDC, emotional well-being, and health behavior. Collecting this data right after sensor data collection enables matching the responses to the inferences made from the sensor data based on the partners' interactions. Additionally, collecting this data at the end of the day enables relevant daily summative data such as whether the patient took their medication for the day.

#### 4 DEVELOPMENT: SYSTEM IMPLEMENTATION

In this section, we describe our implementation of the system requirements to develop the DyMand system consisting of an overview of the DyMand system, devices used, the MobileCoach platform, smartphone app, and smartwatch app.

##### 4.1 Overview of the DyMand system

The DyMand system (Figure 1) consists of a smartwatch app (section 4.7), and a smartphone app (section 4.6) built on top of the MobileCoach platform (section 4.3) [26, 38] that consists of a web-based intervention designer and backend. Each partner is given a smartwatch and a smartphone (both paired) running the DyMand apps.

The smartwatch app collects 5-mins of sensor data (section 4.7.4) based on the detection of the partners' interactions via Bluetooth Low Energy (BLE) and Voice Activity Detection (VAD) (Figure 2). After data collection, it gives a vibration alert on the smartwatch, and then sends an intent to the smartphone app to also give an alert (push notification) and trigger the self-report for each partner to complete separately. The smartphone app then sends a signal to the MobileCoach backend (section 4.5) to trigger the showing of the self-report on the smartphone. The smartphone app is customized with two digital coaches PIA (interacting with the partner with diabetes) and PETE (interacting with the partner without diabetes) that send various messages to each partner (e.g., "it is time to complete the self-report"). MobileCoach is also used for collecting data from the partners during the setup phase of the devices and also triggers end-of-the-day diary questionnaires. MobileCoach also sends reminders and escalation messages when the expected number of self-reports is not completed.

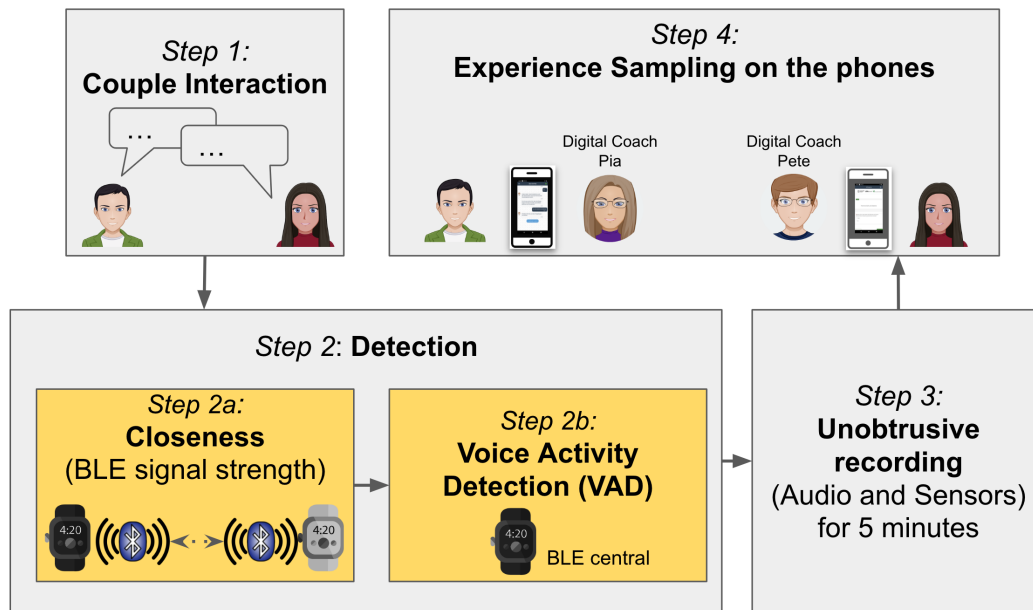


Fig. 2. Use case of DyMand system

#### 4.2 Devices: Polar M600 Smartwatch and Nokia 6.1 Smartphone

For implementing the data collection from partners during their conversation moments, we decided to use two smartwatches (one for each partner) as they are more likely to be with wearers often compared to a smartphone, and could better capture physical closeness and interactions. Furthermore, a smartwatch has a better chance of capturing conversations if audio is recorded on it than a smartphone that may not be in the proximity of the partners. Additionally, relevant data of interest such as heart rate and gestures can be captured on a smartwatch and not a smartphone.

We started working with Apple's WatchOS platform but to the best of our knowledge at the time of development, we could not find a solution to start an audio recording as a background process. Hence we moved on to use smartwatches with the Wear OS (based on Android) platform because it provided a lot of flexibility for data collection. We chose a device from the company, Polar as their devices are regularly used in research [11] and we found the Polar M600 smartwatch to be practically priced and the battery seemed to last well with continuous sensing and with a daily charge cycle. Additionally, the heart rate sensor had good performance in a comparison study in which it was shown to be accurate during periods of steady-state activities like cycling, walking, jogging, and running, but less accurate during some exercise intensity changes [29].

We decided to use a smartphone to collect self-report data as the screen size of the watch was too small for filling out self-reports, and to serve as an intermediary between the MobileCoach server and the smartwatch app. We used two Nokia 6.1 smartphones running Android 9.0 – one per partner – which we outfitted with SIM



cards for Internet access. We found Nokia 6.1 to be practically priced, with a moderate screen size (5.5”) and with good haptics.

### 4.3 MobileCoach Platform

We implemented the DyMand system on top of MobileCoach, an open-source software platform for the design of behavioral interventions and ecological momentary assessments [26, 38]. MobileCoach provides a server setup and a mobile client app. The server offers a web-based intervention designer, called MobileCoach Designer (section 4.5). The intervention author can use this interface to design dialogues that can be sent to the mobile app. The author can also put together rules in the designer and connect them to these dialogues so that the messages can be sent when the conditions in these rules are satisfied. The rules could be based on intent from the user (e.g., button press in the mobile app), time range (e.g., between 2 and 3 pm), and other variables which can be created and used in the designer. When the mobile app is opened for the first time, it can show information to the user about the app and let the user pick a conversational agent, i.e. a computer program that imitates the conversation with a human being [37]. By default, the app provides two conversational agents and after one is chosen, the agent can communicate with the user. The agent converses using the dialogues designed in the MobileCoach Designer. With some customizations and a smartwatch app connected with the MobileCoach framework, it fitted the DyMand system’s use case.

### 4.4 MobileCoach Backend

The MobileCoach server (Figure 3) of the DyMand system is hosted in an Ubuntu 16.04 virtual machine in the ETH Zurich network. The framework provides a Docker-based setup. The Docker platform [45] performs OS-level virtualization and helps in the portability of software. The server setup of MobileCoach comes with containers running Tomcat [1], DeepStream [3], MongoDB [7] and Nginx [8]. Tomcat provides a Java web server environment and it serves the MobileCoach Designer (section 4.5) which is built using the Vaadin framework [12]. The intervention engine of MobileCoach processes the rules and schedules the messages. It connects with the user interface (UI) of the MobileCoach Designer and communicates (i.e., sends/receives the messages) with the mobile app using the real-time server DeepStream. DeepStream allows clients and backend services to synchronize data. The database is managed by MongoDB, which is a NoSQL database program. The communication between the server and the mobile app happens through secure SSL [61] encrypted connections. We used the certificate authority Let’s Encrypt [5] for setting up the free SSL certificates. We performed regular security upgrades and monitored the server. In addition to sending messages via chat, the MobileCoach framework was configured to send SMS or emails. For emails, we used the SMTP server provided by ETH Zurich. For SMS, we used ASPSMS [2] service. Additionally, for the questionnaires, we hosted a LimeSurvey setup in the DyMand server. LimeSurvey [40] is a free and open-source online survey tool we used to create our questionnaires. We picked LimeSurvey because it provides free community editions and detailed setup instructions, it is easily customizable, and we could host it on our DyMand server so that the data can reside within the ETH Zurich network.

### 4.5 MobileCoach Intervention Designer

The MobileCoach intervention designer has user interface (UI) elements (Figure 4) where dialogs can be created with messages using multiple answer option types. The conversational agent (PIA or PETE) serves these dialogues to the partners. There are three main modes of rule execution in MobileCoach designer as follows: (1) daily execution (once every day at midnight), (2) almost-continuously execution (every 5 minutes) and (3) event-triggered execution (e.g., events can be triggered passively by sensor inputs or manually by user interactions in the mobile application). Each dialogue is connected to either of these rules. When conditions of rules are satisfied, then corresponding dialogues are sent. The conditions are designed based on system and user-defined variables

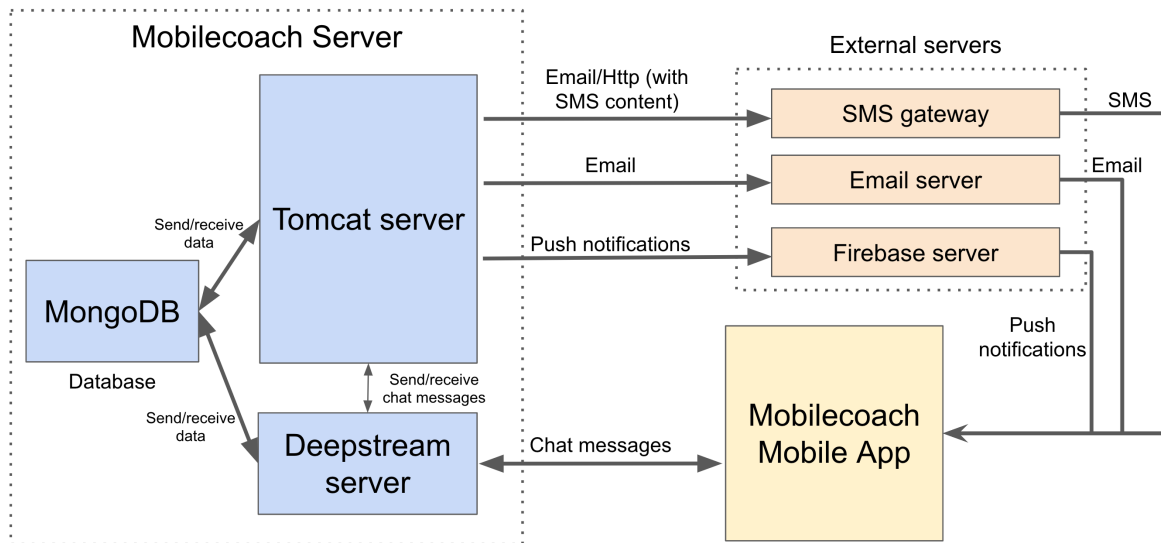


Fig. 3. DyMand MobileCoach Server and its functions

available in the MobileCoach Designer. Each participant has a copy of these variables, which defines the state of the participant. The MobileCoach designer also provides an interface through which study participants and their states (based on the variables) can be monitored. This interface also allows the export of all messages sent and received with timestamps, and the most-recent variable values for each participant for data analyses. For the DyMand system, we designed the following dialogues:

- (1) Onboarding dialogue is sent when the app is opened for the first time during the setup of the study apps. It is used to get the personal information from each partner (e.g., to approach them by their nicknames), to complete the baseline questionnaire before the start of the study, to deliver relevant study information, and help with onboarding the participants for the study.
- (2) Reminder messages in the form of text messages (SMS) were sent to the personal phones of the couples on Sunday evening and again on Monday morning reminding them to take their study smartphones and wear the study smartwatches directly after getting up to prepare them for the 7-day data collection.
- (3) Self-report dialogue requests participants to fill a self-report questionnaire after the sensor data collection on the smartwatch. The self-report was a LimeSurvey questionnaire that asked questions about social support, CDC, health behavior, and emotions (e.g., short form of the PANAS self-report [59]). This dialogue used the event-triggered rule execution mode which was initiated when the sensor data collection was successful. In particular, the MobileCoach designer provides various commands to be sent to the mobile app. These commands can have different purposes in the app. For example, a command “show-web <https://abc.de> <Button-name>”, will show a button with the name “Button-name” which when pressed would open the link “https://abc.de” in a new screen within the app. These links in the DyMand system (Figure 2) were our pre-configured questionnaire from LimeSurvey which additionally provided choices to enter JavaScript code. We added codes in each of our surveys that enabled the app to close the survey screen and continue the dialog with the conversational agent. Using JavaScript code, we added timers

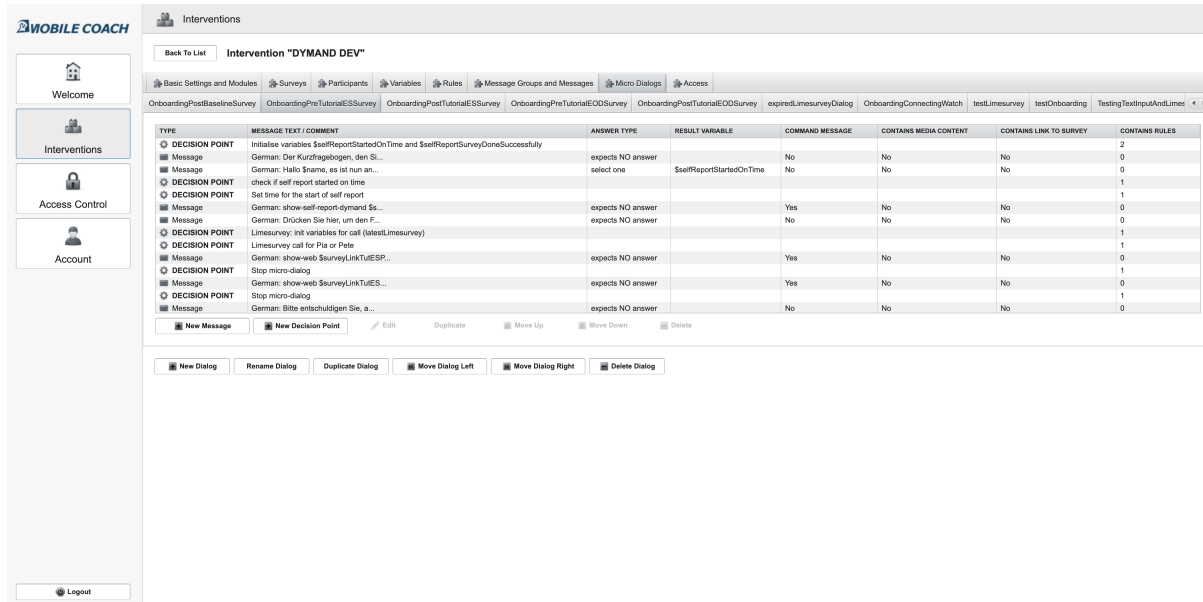


Fig. 4. MobileCoach Intervention Designer

to the surveys. They were set to expire after 4 minutes to ensure that partners started filling out the self-report within 4 minutes of getting the alert. Given that the self-report questions were about the previous 5 minutes of the couple's interaction when the sensor data was collected, the time limit helped to ensure the survey data matches the sensor data.

- (4) End-of-day diary questionnaire dialogue requests the partners to fill a self-report survey at the end of the day with more comprehensive questions on social support and CDC, healthy eating, medication adherence, and emotional well-being [42]. Similar to the implementation for the self-report dialogue, they were set to expire after 45 minutes.
- (5) Follow-up dialogue was sent when the 7-day field study was finished. This dialogue asked partners to fill out a follow-up survey about their study experience [42].
- (6) Escalation messages were designed as both text messages and emails that were sent to the partners and study supervisors when the partners were less adherent to completing self-reports in the study. As shown in Figure 5, there were two checks done by the system automatically every day during the 7-day study period — one at 2 pm, and another one at the end of the day, after the end-of-day diary is sent. As shown in Figure 5, when the partners do not fill at least 60% percent of the self-reports in the morning or the evening period, a reminder SMS is sent to them. Additionally, if the end-of-day diary is not completed or if the total number of the self-reports done for the whole day was less than 30%, then the participant received an SMS and the study supervisor was sent an email additionally to follow up with the couple via phone call.

#### 4.6 MobileCoach Smartphone App

MobileCoach provides a skeleton app in React Native, which is a framework for building cross-platform applications for Android and iOS devices. We customized it for the DyMand system and used only the Android app.

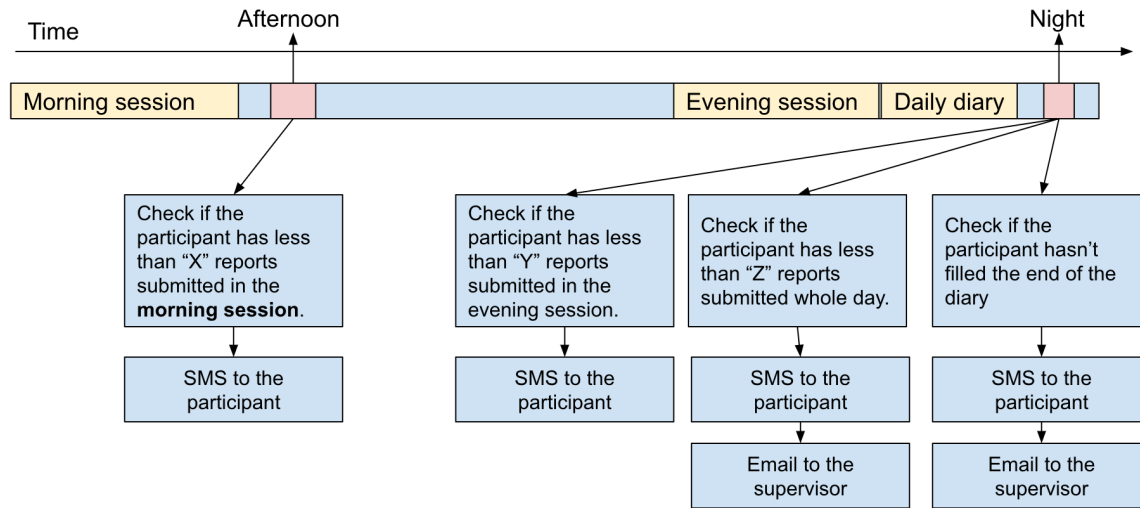


Fig. 5. DyMand Escalation mechanism: The values X, Y, and Z are calculated every day based on the number of hours the participants specify that they are available for the study in the morning or the evening period and the percentages of self-reports that we set as thresholds for escalation.

We customized it to have the two conversational agents PIA (interacting with the partner with diabetes) and PETE (interacting with the partner without diabetes). The smartphone app acts as an intermediary between the smartwatch and the server. It relays the “user intent” messages that indicate that “a recording in the smartwatch has been completed” to the server and also informs the smartwatch as soon as the user has finished completing the self-report triggered by the server. The self-report (Figure 6) contains a questionnaire in LimeSurvey [40] and the Affective Slider, a digital emotion measuring instrument that assesses emotions along the dimensions of valence and arousal [15]. Furthermore, the app collects video, audio, and ambient light for three seconds when each partner is completing the Affective Slider on their smartphone. When a user needs to fill a questionnaire prepared in LimeSurvey, the user sees a button that says “Fragebogen ausfüllen” (Start the survey). When the button is pressed, the link to the survey is opened. The link also includes relevant metadata such as participant code that uniquely identifies the participant, which is needed to link survey answers to the specific partner in LimeSurvey.

*4.6.1 Continuously Running Smartphone Android App.* For the DyMand system to function properly, the mobile app is required to run continuously (in the background) on the smartphone. As the smartwatch app was not expected to connect to the Internet on its own, if the connection to the smartphone app was lost, we could not relay the “recording done” message to the server. We implemented a Foreground service [4] in the mobile app which ran continuously and kept the app in a ready state throughout the study. In addition to the Foreground service, we changed the settings on the study phones that we gave to the couples such that the Android OS did not optimize the battery for our DyMand app. Even then, we found that in Nokia 6.1, one of the system apps (com.evenwell.powersaving.g3) still shut down our app after a few hours. We disabled this system app by using the Android Debug Bridge (ADB) and logging into the shell of the device.

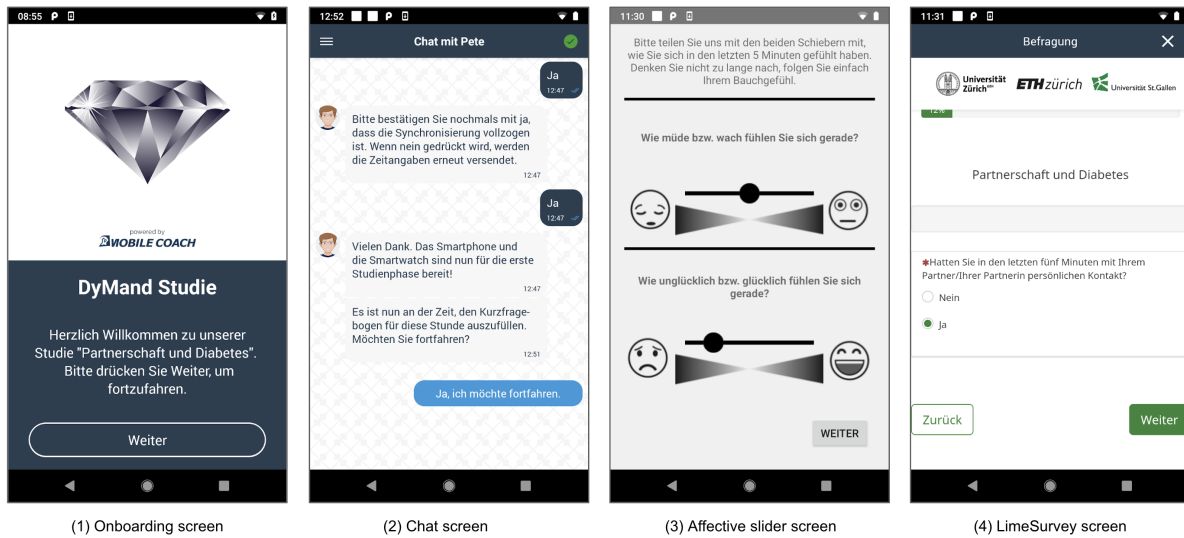


Fig. 6. Screenshots from the DyMand smartphone app. (1) Onboarding screen: These screens welcome the participants to the study. (2) Chat screen: The digital coach Pia or Pete chats with the participant on this screen. (3) Affective slider screen: The participants can choose their pleasure and arousal levels in the slider and submit them. (4) LimeSurvey screen: The self-reports, the end of the day diary, baseline, and follow-up surveys are filled by the couples on this screen.

**4.6.2 Smartphone Data Collection.** When the Affective Slider is shown on the smartphone, a 3-second sensor data recording is made on the smartphone. The recording includes video from the front camera and continuous data from the ambient light sensor. For the video recording from the front camera, we used Android's MediaRecorder API [6]. For the ambient light sensor, we used the SensorManager API [10]. We registered this sensor with the parameter "SENSOR\_DELAY\_FASTEST" to get the sensor data as fast as possible. We stored the data locally on the phone which was retrieved when the couples returned their devices after they finished the study.

**4.6.3 Smartphone and Smartwatch Communication.** For the communication between the DyMand apps in the smartphone and the paired smartwatch, we used the Wearable Data Layer API [9]. The following messages were sent between the smartphone and smartwatch applications.

- (1) The weekday and weekend hours during which the couple is available for data collection [9] (chosen by the couple in the smartphone app) are sent from the smartphone app to the smartwatch app during the setup phase (section 4.8).
- (2) Text indicating that the smartwatch has finished collecting sensor data for 5 minutes is sent from the smartwatch app to the smartphone app throughout the study.
- (3) Text indicating that the self-report on the smartphone has been completed is sent from the smartphone app to the smartwatch app throughout the study.
- (4) Other messages for acknowledging received messages and logging are sent between the apps on the smartwatch and smartphone.

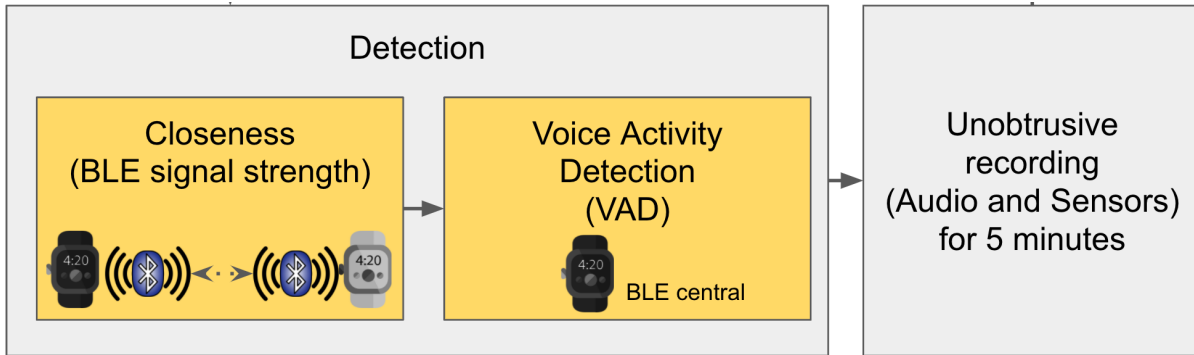


Fig. 7. Smartwatch App System Overview

#### 4.7 Smartwatch App

Similar to the smartphone app, we implemented a Foreground service [4] in the smartwatch app (Figure 7) that collected sensor data. The smartwatch app collected five minutes of the following sensor data once per hour within the morning and evening hours set by the couples: audio, heart rate, accelerometer, gyroscope, and ambient light. We collected a maximum of 5 minutes of data per hour for privacy reasons. Hence, to optimize the quality of data collected within that hour and to ensure that we recorded the most relevant 5 minutes of data (when partners are interacting), rather than triggering data collection at random or scheduled times, the app collected data when 1) the partners were physically close and 2) when there was speech. Our algorithm uses a two-step process. First, the app determines physical closeness using the BLE signal strength between the smartwatches (section 4.7.1) and checks if the signal strength is within a certain threshold, which corresponds to a distance estimate (section 4.7.2). Second, the app determines if the partners are speaking by using a voice activity detection (VAD) machine-learning algorithm, which is implemented on the smartwatch (section 4.7.3). In the case in which this condition of physical closeness and speaking is not met in the hour, the app triggers a backup recording in the last 15 minutes of the hour.

After the 5-minute recording ends, the watch vibrates and sends a trigger to the smartphone app via the Wearable Data Layer API [9] to bring up the self-report for that partner to complete. If the smartwatch does not receive a message from the smartphone app within 2 mins indicating that the self-report has been started, it gives another vibration alert. If once more, within the next 2 minutes, there is still no response about the start or completion of the self-report, it implies the self-report was not completed. Consequently, the app deletes the audio and attempts to trigger another sensor data collection and self-report for the rest of the hour. Doing this ensured that we collected data with matching sensor and self-report samples. The app also ensured that there were at least 20 minutes between subsequent data collection to reduce the burden of the partners completing the self-reports.

*4.7.1 Physical Closeness Estimation.* We used the BLE signal strength between the two smartwatches – one acting as the central (does BLE scanning) and the other acting as the peripheral (does BLE advertising) – to estimate the physical closeness of the partners. We conducted an experiment to measure the signal strength based on the distance between the two watches. In a lab setting, we placed two smartwatches at the same level and without any barrier between them. With one BLE peripheral watch fixed at a position, we placed the BLE central watch at a given distance from the former and we measured the signal strength. We varied the distances between the watches and repeated the experiment 10 times. We averaged the values and plotted them as shown in

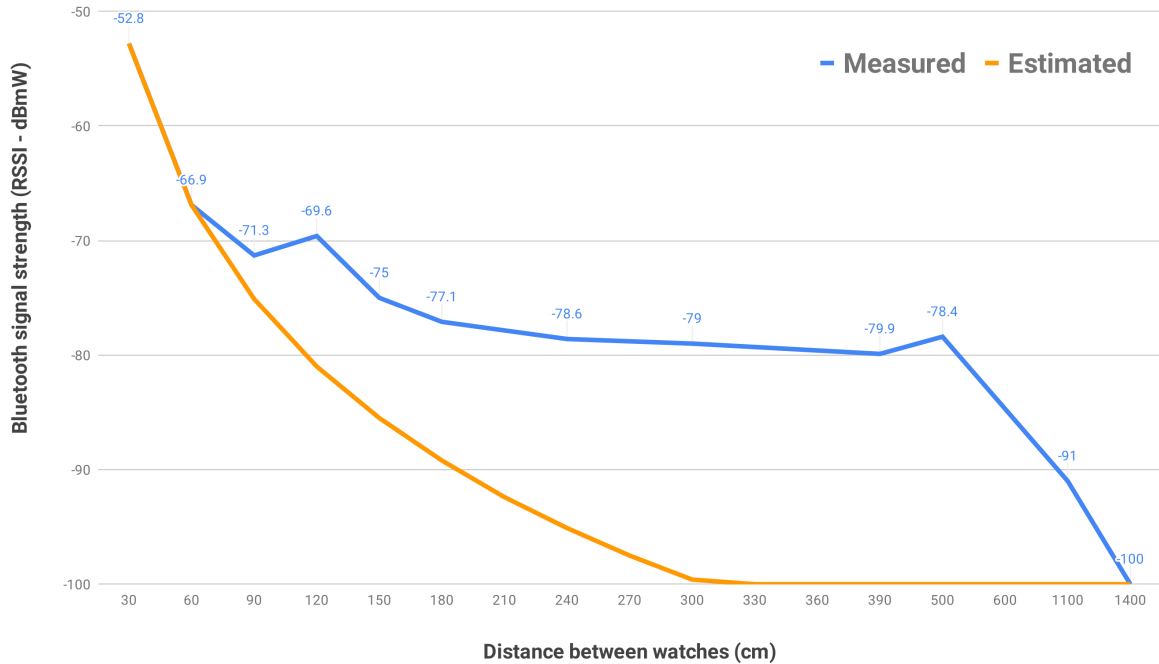


Fig. 8. Plot of RSSI between two smartwatches versus distance for real-world experiment and theoretical expectation (that the signal strength increases exponentially when the devices get closer and closer [33])

Figure 8 along with the theoretical expectation that the signal strength increases exponentially when the devices get closer and closer [33]. Our measurement showed that the signal strength is proportional to the closeness of the smartwatches but not perfect as in any BLE measurement. We chose a threshold of  $-80\text{dB}$  in the app which covers a range typically less than 5 meters which we assumed should adequately capture the distance between partners when they are interacting. We acknowledge that in the field, the presence of objects such as walls and furniture will affect the signal strength which we did not factor into the experiment. Nonetheless, our goal was not to have a very precise distance versus signal strength mapping but an approximate value to use for closeness between partners.

**4.7.2 Physical Closeness Detection.** For the DyMand smartwatch app, we implemented a BLE service in the smartwatch app for real-time physical closeness detection. At the start of each of the hours assigned for data collection, the peripheral smartwatch continuously advertises its universally unique identifier (UUID). The UUID is created based on the couple ID which is set uniquely for a couple during the onboarding of the study (e.g., P001 for the supporting partner and Z001 for the patient) (section 4.8). The central smartwatch scans with the UUID and when it finds the corresponding peripheral smartwatch, it checks the signal strength. If the signal strength is greater than  $-80\text{ dB}$ , then it tries to connect to the peripheral smartwatch. After a successful connection, both BLE services acknowledge it by sharing a message through the Bluetooth channel. After this, the central smartwatch does voice activity detection (section 4.7.3). If the signal strength is less than  $-80\text{ dB}$ , the BLE scanning waits until the strength breaches the threshold. If there are any problems when the devices connect, the BLE scanning is

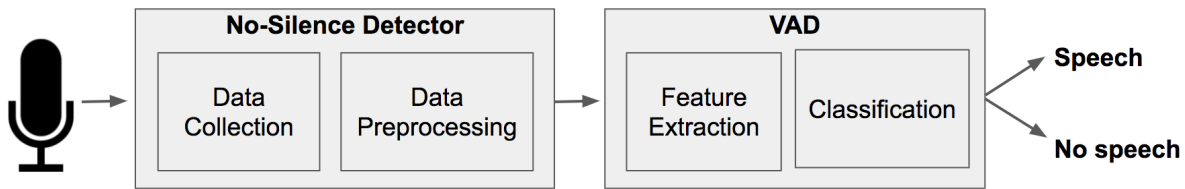


Fig. 9. System Architecture of VADLite

reset and started from the beginning. Both the central and peripheral BLE services are attached to the Foreground service so that they can run in the background continuously.

**4.7.3 Voice Activity Detection.** We developed VADLite, a lightweight, open-source voice activity detection (VAD) system that runs in real-time on the smartwatch (see [21] for details of the system) (Figure 9). An offline and online evaluation of VADLite using real-world data showed better performance than WebRTC’s VAD, a popular open-source VAD system. VADLite is a 2-stage system consisting of a no-silence detector and a voice activity detector.

The no-silence detector computes the root mean square (RMS) of segments of the audio signal and marks them as non-silence if they are above a certain threshold. The voice activity detector consists of a feature extractor, and a machine-learning algorithm, which we trained to classify speech versus non-speech [21]. In particular, it extracts mel-frequency cepstral coefficients and classifies speech versus non-speech audio samples using a linear Support Vector Machine (SVM). An SVM is a classifier that constructs a high-dimensional hyperplane to separate data of different classes [30]. SVM selects a hyperplane that maximizes the distance to the nearest data points on either side of the hyperplane in the case of binary classification. We used a linear SVM because it is memory and computationally efficient when doing predictions. Prior work has used an implementation of linear SVM for real-time prediction on smartwatches for stress detection [19] and activity detection [17, 18].

To train VADLite, we ran a study that was approved by the ethics commission of ETH Zurich. We collected lab and field audio data (16-PCM mono, 8KHz, 3.5 hours total) using the Polar M600 smartwatch from several people (at least 10 distinct individuals) at varying distances from the smartwatch. We annotated the audio samples as speech or non-speech. We preprocessed the data to remove silence segments by computing the RMS of each one-second time window and checking if the RMS value is below a certain threshold determined empirically. We then extracted frequency-based features – 13 MFCC features which have been widely used for VAD [27] and used 12 of them (excluding the 1st coefficient, which is the DC component) – over a time window of 25 ms non-overlapping time window. We normalized the features and used them to train a linear SVM to classify speech or non-speech. We then implemented the system to run in real-time on the smartwatch (Figure 10). The real-time implementation continuously collects audio and processes them in 1-second segments for no-silence detection and then voice activity detection.

We performed offline and online evaluations using the following metrics for evaluation: accuracy, speech hit rate (SHR), and false alarm rate (FAR). The SHR is the ratio of correctly detected speech frames to the total number of speech frames. By contrast, FAR is one minus the noise hit rate, where noise hit rate is the ratio of correctly detected noise frames to the total number of noise frames. For offline evaluation, we split the data into train and test using about 70%-30% stratified split and performed 10-fold stratified cross-validation with hyperparameter tuning on the train data and evaluated on the test set. The model achieved 82.6% accuracy, 80.2% SHR, and 14.9% FAR. For online evaluation, we played 15-minute audio collected from a naturalistic context through a loudspeaker as the VADLite app performed real-time classification of the audio just like was done



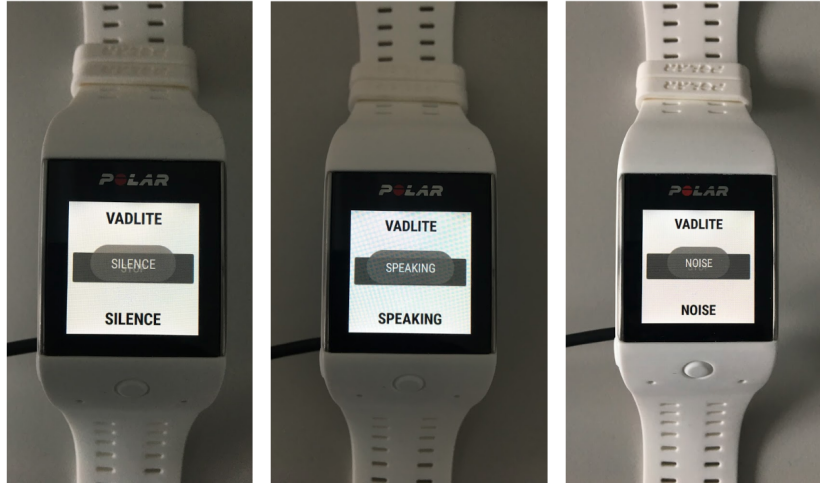


Fig. 10. Real-time testing of VADLite showing detection states of silence (left), speech (middle) and noise (right)

by Feng et al [25]. VADLite had SHR and FAR of 91.6% and 5.5% respectively. These results were better than WebRTC’s VAD: offline accuracy, SHR and FAR of 71.4%, 79.5%, and 37.5% respectively, and online SHR and FAR of 73% and 18% respectively. Furthermore, VADLite had an average processing time of 2 ms for each 25 ms frame and 76 ms for the total one-second duration. As a result, throughput was met since the frame processing time was less than the 25 ms segment duration. Likewise, the processing time for the whole duration was less than one second.

For the DyMand smartwatch app, the VAD component is triggered after the physical closeness detection. It records and samples audio at 8KHz and processes them in 5-second chunks by first performing no-silence detection and then speech detection. If it classifies the segment as speech, the VAD audio recording is stopped, and then sensor data collection is started on the central device (since it runs the VAD component) and a signal is sent to the peripheral device to immediately start sensor data collection.

**4.7.4 Smartwatch Data Collection.** The data collection component of the smartwatch app collected audio, heart rate, accelerometer, and gyroscope for 5 minutes. For the audio data, we used the MediaRecorder API [6] available in Android. We collected 16-PCM mono at 44.1 kHz. We set the output format of the audio file as ".wav" which is a lossless file format. All the other sensor data were collected using the SensorManager API [10]. We registered these sensors with the parameter "SENSOR\_DELAY\_FASTEST" to get the sensor data as fast as possible. The data collected was stored locally on the smartwatches and retrieved after the devices were returned.

**4.7.5 Exception Handling.** We added try-catch statements in various parts of the code where they could be exceptions (e.g., writing text or data to a file, BLE scanning, etc.). Additionally, similar to [36], we included the "DefaultUncaughtExceptionHandler" to catch all uncaught exceptions. Our implementation of the Exception class spawns a thread that counts the number of exceptions in that hour and saves it (for later logging), saves the exception message (for logging), schedules a restart of the app with an AlarmManager for the next second, stops the current Foreground service, and shuts down the app.

**4.7.6 Logging.** To ensure that we understood how well the system was performing, we included various logs in the smartwatch app that were saved in files on the smartwatch. In particular, we had 1) configuration logs at

the time of setup, 2) hourly logs between the setup time of the devices and the start of the study, 3) hourly logs during the 7-day study period, 4) continuous log of the BLE signal strength between the two smartwatches, 5) hourly app function logs and 6) error logs as and when they happened.

- (1) The configuration log contained the dates and hours of data collection set by the couples which we used for post-study analysis.
- (2) The before-study logs contained the log's timestamp, battery level, the number of days, hours, minutes, and seconds until the start of the study, and the number of exceptions that happened in the previous hour. We could use this data to infer various things such as whether the battery died before the study started and the partners forgot to charge it, or whether an error shut down the app.
- (3) The during-study hourly logs contained several important fields. These included the following from the previous hour: timestamp of the log, battery level, date and number of times the BLE started advertising (peripheral device) or scanning (for the central device), date and number of times the device met the closeness condition, date and number of times of no-silence detection, date, and number of times of voice activity detection, date and number of times the two watches connected, date and number of times sensor data was collected, date and number of times self-report was triggered (1st and 2nd alert), started, and completed, whether the recording was a backup recording, if the audio was discarded (because self-report was not completed), number and dates of errors, number of times and dates the app restarted, whether internet was available on the smartphone, and amount of space remaining on the smartwatch. These data allowed us to assess the performance of the system (see Section 6).
- (4) The hourly app's function logs contained various "print" statements we had in various functions in our code as well as the Wear OS system logs which our app logged and saved in a file. We did this to debug and better understand which code block may have caused any errors.
- (5) The BLE log saved the signal strength between the devices throughout the continuous scanning during the hours of data collection. This data can be used to infer how much time partners spent together.
- (6) The error logs were the exceptions (described in the previous section) that we logged. We used those to debug the app.

*4.7.7 DyMand App Checker.* In early deployments, we realized that the smartwatch app could sometimes be shut down by the Wear OS system after an error, and our implementation for the app to restart itself was not always reliable. Consequently, we developed a smartwatch app – DyMand App Checker – to continuously check each hour if the DyMand smartwatch app is running and then start it if it is not running.

#### 4.8 Setup Components of Smartwatch and Smartphone Apps

We implemented setup components of the smartphone and smartwatch apps for the partners to set up their devices, to ensure all aspects of the apps were working correctly before they take the devices away, and also for them to experience the process of sensor and self-report data collection. The smartphone app collects personal data and also the hours that the partners indicate for data collection. This information is transferred to the smartwatch app during the setup of the smartwatch after the "send" button is pressed on the smartphone. If the watch receives it, it shows a screen saying that the configuration process is successful. Otherwise, it needs to be sent from the smartphone again through the same button. The communication between a paired smartwatch and smartphone is explained in detail in Section 4.6.3. The smartwatch app collects information such as the couple ID (e.g., 001, 002) for logging purposes and creates a unique UUID (section 4.7.2) that can be used by the BLE services for advertising and scanning, the color of their smartwatch for the app to know if the watch belongs to the patient or supporting partner (by default, our smartwatch app chooses white to be the BLE peripheral and the black to be a BLE central device), receives the data collection hours from the smartphone, and collects a voice sample for 1 minute from the partner by asking them to read some text on a paper. Next, the BLE central

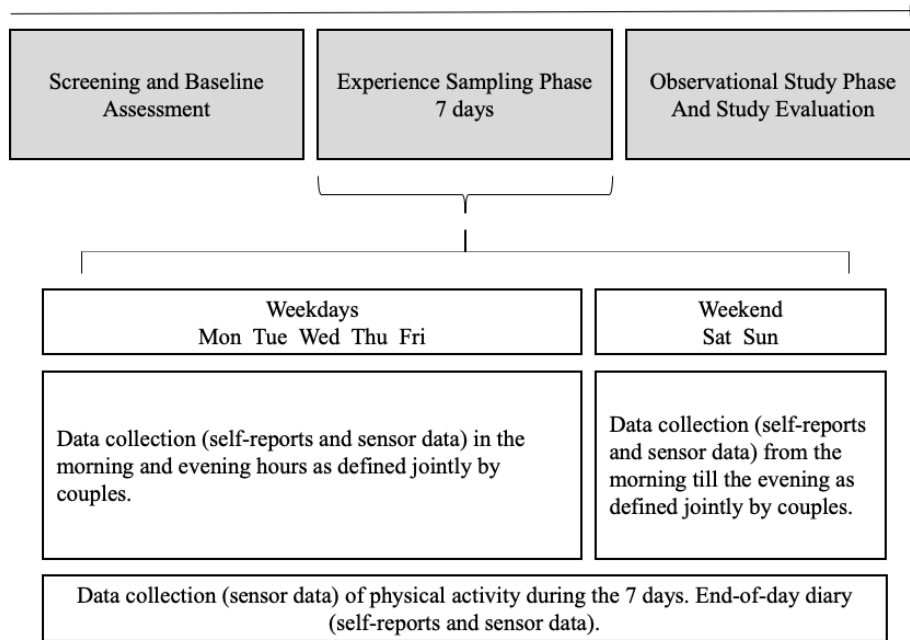


Fig. 11. Overview of DyMand study [42]

smartwatch starts scanning and the BLE peripheral smartwatch starts advertising. After successful connection, a message is shared via this Bluetooth connection between the smartwatches, and the recording of the sensor data starts in the watches, completes after 5-minutes, triggers self-report on the smartphone, and then triggers the end-of-day diary.

## 5 DEPLOYMENT: FIELD STUDY

After various internal pilot tests of the DyMand system, we deployed it in a field study with couples. We ran the DyMand study between 2019 and 2021 with heterosexual romantic couples from the German-speaking part of Switzerland in which one partner had T2DM (Figure 11) [42]. In total, we collected data from 13 couples aged 47 to 81 years, with a mean age of 68 (SD = 9) resulting in a total of 1,019 5-minute samples of sensor data (85 hours) and 598 corresponding completed self-report data.

The study was advertised in various places including hospitals, magazines, local newspapers, and the diabetes association in Switzerland. Interested couples completed a web-based questionnaire to screen them for the inclusion and exclusion criteria, and collect socio-demographic information. Those who met the eligibility criteria were able to pick a date for a baseline assessment at the Applied Social and Health Psychology laboratory of the University of Zurich. During this session, both partners received comprehensive information about the study, signed the informed consent form, and completed a web-based questionnaire that captured constructs of interest at baseline that were not assessed daily.

They also received instructions on the study and then trained research assistants helped them to set up their devices and pair the corresponding smartphone and smartwatch (see Section 4.8 for details of the setup process). Each partner was given a smartwatch and smartphone running the DyMand apps and they were instructed to

have all devices with them every day for 7 days from getting up until going to bed. To prevent mistakes from one partner accidentally using the other partner's watch and phone, one set of phones and watches had black covers and the other set had white covers. The patient was given the white set and the supporting partner was given the black set. The partners picked the hours during which we could collect data from them. They could choose any period from 4 am to 11 am for the morning hours and from 4 pm to 11 pm for the evening hours. During the weekend, only the early morning hours and late evening hours were set (e.g., from 6 am to 10 pm). With this procedure, privacy aspects were addressed by primarily focusing on situations, in which the couples spent time together and thus reducing the number of audio recordings during the day of weekdays when chances are higher that subjects are working or moving around in public places.

We collected data from their daily life for 7 consecutive days starting the next Monday after their visit until the following Sunday night. The DyMand system collected sensor and self-report data as described in previous sections. During the study, we had a process for monitoring the study to ensure that the system was working correctly and to enable us to intervene if needed. We had a spreadsheet with cells corresponding to the hours during which we should have received data. Research assistants checked the DyMand MobileCoach server daily to see if the self-reports were triggered and completed, with each cell receiving one of the following values: 0 – no self-report trigger received, 1 – self-report trigger received, but no survey completed 2 – trigger received and survey completed. Additionally, we had a sheet that we used to track the details of issues and complaints that the couples communicated to the research team either via email or calls. We collected the following relevant information to understand and solve any problems that came up: subject id, device id, does the issue involve the watch, does the issue involve the phone, detailed description of the issues and pictures if available, date and time of issue, where they were when the issue happened, was the watch switched on, was the phone switched on, were the two watches physically close together, were the phone and watch physically close together, was there a connection to the digital coach indicated by the top-right icon on the smartphone app being green, were the phone and watch connected, does the phone have internet, and are both apps on the watch running.

There are significant ethical and privacy concerns of such a system and study since we collect audio which is sensitive data, and more so in the context of couples' interactions with the likelihood of speech about private topics. We took several measures as follows. First, our study received ethical clearance from the cantonal ethics committee of the Canton of Zurich, Switzerland (Req-2017\_00430). Second, we ensured that we collected a maximum of 5 minutes of audio per hour in order not to record a significant percentage of the couples' everyday life. Consequently, even if the system triggers multiple recordings in the hour, the app always deletes all but the last one before the end of the hour. Third, to protect the privacy of subjects not taking part in the study, we asked subjects to wear a tag that we give them to indicate to others around that recording may be happening and that they may be recorded. Finally, after subjects returned their devices, we gave them the option to listen to and request the deletion of any audio samples without any explanation before the study team could listen to the audio files. Similar measures have been used in previous studies [44, 50] and have proven adequate to safeguard the privacy of study subjects and others not taking part in the study.

## 6 EVALUATION

We evaluated the DyMand system regarding its technical performance and usability. For technical performance, the DyMand apps on the smartwatch and smartphone performed hourly logs (section 4.7.6) of relevant system performance metrics such as whether data collection started and completed, and whether self-report data was triggered, started, and completed. Furthermore, we annotated the audio data with relevant information such as whether it contained speech, and whether there was a conversation between partners for the algorithm triggered recordings and the backup recordings.

Table 1. Expected and actual number of sensor and self-report data that were collected

Data	# of Samples
Total expected # of sensor and self-report	1392
Expected # of samples with the app running	1028
# of sensor data collected	1019
# of self-reports triggered	1019
# of self-reports started	618
# of self-reports completed	598

Table 2. Percentages of the expected number of sensor and self-report data that were collected

Data	Percentage (%)
% of total expected sensor data that was collected	73.2
% of expected sensor data with the app running that was collected	99.1
% of total expected self-report triggers that happened	73.2
% of total expected self-report triggers with the app running that happened	99.1
% of triggered self-report that were started	60.6
% of triggered self-report that were completed	58.7

### 6.1 Self-report and Sensor Data Collection

We investigated the percentage of the total expected number of sensor data and corresponding self-report data that was collected. Given that couples indicated the hours during which we could collect data, we estimated the total expected number for each couple and then summed them for all couples. Furthermore, given that the DyMand smartwatch app may not always be running due to circumstances such as the device being off because the couples did not charge them or the app crashing because of various errors such as in the BLE stack, a better metric would be the percentage of the expected number of sensor data and self-report data that was collected for hours when the app was running. For example, our logs showed that one couple did not turn on the smartwatch for the duration of the study. Hence, the app did not run and could not collect any data. Using the app's hourly logs, we estimated the number of hours that the app was running. We then used the following log status data: Was the sensor data collected? Was the self-report triggered? Was the self-report started? and Was the self-report completed. We computed sums of these log events for each partner, and then for all couples (Table 1). Additionally, we computed relevant percentages as shown in Table 2.

Our DyMand system collected 73.2% of the total expected sensor data and triggered 73.2% of the total expected self-reports. Considering only the case where the app was running, these percentages become 99.1% which shows that the system adequately triggered data collection as expected. Additionally, partners started 60.6% and completed 58.7% of the triggered self-reports. This percentage is not very high yet understandable for the context of this study. Partners had a maximum of 4 minutes to start the self-report after the trigger. The self-report was dismissed if it was started after this time. There are several reasons partners may not have started the self-report such as they did see or hear the alert because they were not wearing the watch, or they did not have the phone close by them. For those who saw the alert and attempted to start the self-report, some partners complained that there was a delay in the self-report loading on the smartphone due to internet connection issues (see Section 6.5) and may have affected the start of the self-report in time.

Table 3. Number of audio samples with speech and conversation between partners broken down by triggered and backup recordings

Data Field	# of Samples
Total audios	1014
Audios with speech	791
# of triggered recordings	277
Triggered recordings with speech	256
# of backup recordings	737
Backup recordings with speech	535
Conversation between partners for all audios	538
Conversation between partners for triggered recording	215
Conversation between partners for backup recording	323

Table 4. Percentage of audio samples with speech and conversation between partners broken down by triggered and backup recordings

Data	Percentage (%)
% of total audios with speech	78
% of triggered recordings with speech	92.4
% of backup recording with speech	72.6
% of total audios with conversation between partners	53.1
% of triggered recordings with conversation between partners	77.6
% of triggered recordings where one partner spoke	88.1
% of backup recordings with conversation between partners	43.8

## 6.2 Capturing Partners' Conversation Moments

Given the key novelty of our system is its touted capability of capturing partners' conversation moments using physical closeness and voice activity detection, we investigated how well the DyMand system captured speech and conversation between partners. We define a conversation as the presence of speech from both male and female partners in that audio. We had 3 trained research assistants annotate all valid audios (i.e., were not corrupted and hence playable,  $N=1014$ ) by providing 'yes' or 'no' to relevant information such as: does it contain speech, did the male partner speak, did the female partner speak, and was there a conversation between partners. We automatically extracted the information about whether each audio was a triggered or a backup recording by looking at the timestamp. Audios collected anytime between the 44th minute and the end of the hour (e.g., between 6:44 am and 7 am) are backup recordings. All audios collected before then were triggered by our algorithm. We show the sum of audio status information for all partners in Table 3 and calculate relevant percentages in Table 4.

For triggered audios, 92.4% contained speech, which shows VADLite has good performance in capturing speech in the real world considering prior work has shown that open-source VAD tools perform poorly on real-world speech data collected with smartwatches [21, 39]. Furthermore, 77.6% of triggered recordings had a conversation between partners in comparison to 43.8% of backup recordings which had a conversation between partners. This result shows that our novel approach of physical closeness plus speech detection was better at capturing couples conversation/interaction moments (77.6%) in comparison to using random or scheduled times per hour in our study (43.8%). It is difficult to have a direct comparison to other works. Nonetheless, in the following work [56], random triggering (10 times a day for one week) of self-report data collection among couples resulted in data

for which partners were together (a proxy for interaction) in 39% of the triggers. Also, another work that used the EAR triggered data collection among couples at a scheduled time sequence (every 9 minutes during waking hours) resulting in data for which patients and spouses talked (a proxy for partners' speech), on average, 47.9% and 45.0% of their waking hours [50]. These percentages are lower than the 77.6% from the DyMand system.

There are some likely reasons for the absence of conversation between partners in triggered recordings — the outstanding 22.4%. Firstly, our definition and evaluation metric of conversation between partners is strict — both male and female partners spoke in that 5-minute audio. In actuality, the partners may have been having a conversation but one partner may not have spoken in the specific 5-minute period in which we collected data, resulting in a “no” for conversation between partners. When we relax that definition to be either partner spoke, the percentage increases to 88.1%. Secondly, our physical closeness detection approach assumes that the partners are always wearing the smartwatch which was not always the case. For example, if the two smartwatches are left together on a table with the radio or TV in the background, a recording was triggered but it did not contain a conversation between partners. One way to address this issue is to include a pre-step to estimate if the device is being worn using accelerometer or heart rate data as an example. Thirdly, another edge case that happened was partners sitting together and watching TV but not having a conversation. In this case, the recording was triggered because of the physical closeness and speech from the TV but the audio did not contain speech from the partners. One way to address this issue is to use an extra step for speaker identification to check if the speech is from either partner. Doing this will entail collecting a voice sample from both partners at setup, and then automatically creating and storing an acoustic fingerprint of each partner on the watch, and checking if there is a match after the voice activity detection.

### 6.3 Usability Results

At the end of the 7-day field study, partners completed a self-report on their experience with the DyMand apps. They responded to the statement “The study app was easy to use” on a 7-point Likert scale ranging from strongly disagree (1) to strongly agree (7). As shown in Figure 12, most ratings were high with a mean of 5.8 (std=0.98, N=24), which shows the DyMand system was easy to use. Also, partners wrote open-ended responses about potential areas of improvement for the DyMand system. One key recurring theme was that there was a delay in the self-report showing up as mentioned before. Another common suggestion was to reduce the number of questions.

### 6.4 Errors

On our server, when there were too many user intent messages (messages that indicate a recording is done in the watch), DeepStream could not maintain the connection between the server and the client — a socket hang-up error. Though we could not find a technical solution to the problem, we managed to fix this by restarting the Docker container of DeepStream. As this happened only a couple of times during the study and as we found out the problem both times quickly during our periodic monitoring of the system, this error had a negligible effect on the data obtained.

There were BLE connection issues, too. In the smartwatch, we found that after a few successful BLE connections, the smartwatches could not connect anymore with each other. We fixed this by periodically destroying and recreating the whole Bluetooth stack and not just the BLE services (scanning and advertising).

### 6.5 Challenges

Unavailability of Internet access sometimes resulted in self-reports not being shown on the smartphone in time as described before. This issue is a major limitation of the MobileCoach platform as the mobile app has to send a request to the server to then send a request back to the phone to then show the self-report on the screen

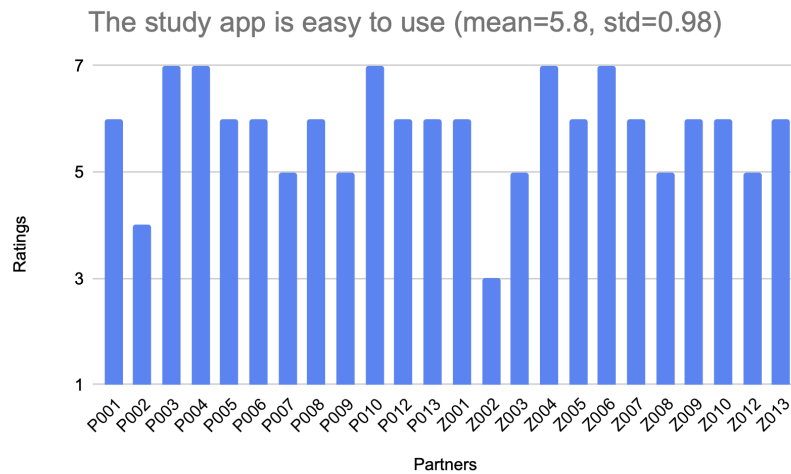


Fig. 12. Usability result showing partners' responses to the statement "The study app was easy to use" on a 7-point Likert scale ranging from strongly disagree (1) to strongly agree (7)

introducing several points of failure. When this happens, the corresponding audio recording gets deleted as it is not valid for us to analyze without a corresponding self-report. The app then tries to record another sample and trigger another self-report, increasing the burden of self-report completion. A solution during the study would have been an offline trigger that syncs when the internet connection is restored. We did not anticipate this as we gave our participants phones with SIM cards for Internet access. Due to resource constraints, we could not change this later during the study. A future extension of the MobileCoach platform with an in-app offline trigger to show the self-report will circumvent these multiple pathways and will prevent this issue.

MobileCoach is designed such that the DeepStream server (section 4.4) syncs the messages designed in the intervention designer to the app only when the app is open and used by the user. After many "recording done" user-intent messages to the server, the server triggers many self-reports and when the app is opened after an inactive period, it takes some time for the messages to sync with the DeepStream server. This issue caused some delay for the most recent message to appear when the app was opened. As the solution would have required major customizations in the MobileCoach framework, we did not address this issue.

Making the smartwatch app run continuously (section 4.6.1) was quite challenging. In addition to the Foreground service and changing the battery optimization settings, one of the system apps in the Nokia 6.1 kept shutting down our app after a few hours. We spent a lot of time trying to figure out the root cause and fix it by disabling the respective system app (`com.evenwell.powersaving.g3`). This issue is important to note for future studies that plan to implement long-running background services on the Android platform.

## 7 LIMITATIONS AND FUTURE WORK

When issues happened during the study, we sometimes did not have enough information to infer the cause until the devices were returned. Even though we had various key logs that would have helped, we did not implement a way to transfer the logs from the watch to the phone, and then to our server for monitoring. The DyMand system was already quite complex and we had time constraints, and hence, we decided not to invest the time and effort into implementing this feature. One key implementation challenge was the fact that the watch did



not have a direct internet connection and hence, any such data transfer would have had to happen through the phone and then to the server, further introducing potential points of failure. To enhance better error analysis in the field, in the future, the DyMand system should be extended to transfer the logs from the watch to the phone and then to an external server.

Our algorithm for capturing couples' interaction moments though performed well, it did not address some edge cases as described in the evaluation section (section 6.2). Future work will update the algorithm to check if the partners are wearing the smartwatch as part of the closeness and speech detection check by leveraging other sensor data such as acceleration. Furthermore, we will implement a speaker identification method to check if the speech is from either of the partners and avoid being triggered by other people or speech from the radio or TV. This implementation will entail training a speaker identification model to extract a speech embedding from each partner during the setup and use it for comparison in real-time on the smartwatch.

Our DyMand system currently only collects relevant sensor and self-report data but it does not perform real-time recognition of constructs that are relevant for understanding couples' chronic disease management such as emotional well-being, social support, and CDC. Future work will augment the DyMand system with such capabilities.

The DyMand system is generic in that it is suitable not only for couples' diabetes management but also for studies in the context of related diseases such as hypertension, or mental health disorders in which sensor and self-reports about emotional well-being and health behavior are relevant for disease management and health intervention designs. Given it is open source, it can be extended or adapted, and our novel method for capturing partners' conversations/interactions can be used by other researchers to develop similar apps to collect data to understand various constructs among other dyadic constellations such as friendships, and sibling and parent-child dyads. Also, another potential research use case is to better understand communication patterns in-situ and performance measures of teams in organizations.

## 8 CONCLUSION

In this work, we developed, deployed, and evaluated the DyMand smartwatch and smartphone system that captures couples' dyadic interactions in daily life in the context of chronic disease management. It consists of a smartwatch app, a smartphone app, built on top of the MobileCoach platform that collects sensor and self-report data that are relevant for chronic disease management on couples' interaction/conversation moments. We deployed DyMand in a 7-day field study and collected 85 hours of data from 13 heterosexual romantic couples from the German-speaking part of Switzerland. Key challenges affected the system's performance and usability such as software errors, poor Internet connectivity, and long self-report questionnaires. Nonetheless, our evaluations showed that the system had good performance in triggering the collection of the expected number of sensor and self-report data, and capturing couples' conversation moments, and it was easy to use. The DyMand system would enable social, health, and clinical psychologists to understand the social dynamics of couples in everyday life and for developing and delivering behavioral interventions for couples who are managing chronic diseases. Our system could be customized and extended to be used in other contexts besides chronic disease management such as couples' daily dynamics more broadly, workplace interactions, and other dyad constellations such as parent-child and sibling-sibling or roommate dyads.

## ACKNOWLEDGMENTS

Funding was provided by the Swiss National Science Foundation (CR12I1\_166348/1).

## REFERENCES

- [1] 2022. Apache Tomcat. Retrieved March 28, 2022 from <https://tomcat.apache.org/>
- [2] 2022. ASPSMS. Retrieved March 28, 2022 from <https://www.aspsms.com/en/>

- [3] 2022. Deepstream. Retrieved March 28, 2022 from <https://deepstream.io/>
- [4] 2022. Foreground services. Retrieved March 28, 2022 from <https://developer.android.com/guide/components/foreground-services>
- [5] 2022. Letsencrypt. Retrieved March 28, 2022 from <https://letsencrypt.org/>
- [6] 2022. MediaRecorder. Retrieved March 28, 2022 from <https://developer.android.com/reference/android/media/MediaRecorder>
- [7] 2022. MongoDB. Retrieved March 28, 2022 from <https://www.mongodb.com/>
- [8] 2022. NGINX. Retrieved March 28, 2022 from <https://www.nginx.com/>
- [9] 2022. Send and sync data on Wear OS. Retrieved March 28, 2022 from <https://developer.android.com/training/wearables/data/data-layer>
- [10] 2022. SensorManager. Retrieved March 28, 2022 from <https://developer.android.com/reference/android/hardware/SensorManager>
- [11] 2022. Using polar products in research. Retrieved March 28, 2022 from <https://www.polar.com/en/science/research-tools>
- [12] 2022. Vaadin Framework. Retrieved March 28, 2022 from <https://vaadin.com/>
- [13] Tuka AlHanai and Mohammad Ghassemi. 2017. Predicting latent narrative mood using audio and physiologic data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 31.
- [14] Isabella C Bertschi, Fabienne Meier, and Guy Bodenmann. 2021. Disability as an interpersonal experience: a systematic review on dyadic challenges and dyadic coping when one partner has a chronic physical or sensory impairment. *Frontiers in psychology* 12 (2021).
- [15] Alberto Betella and Paul FMJ Verschure. 2016. The affective slider: A digital self-assessment scale for the measurement of human emotions. *PloS one* 11, 2 (2016), e0148037.
- [16] George Boateng. 2020. Towards a wearable system for assessing couples' dyadic interactions in daily life. In *Adjunct Proceedings of the 2020 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2020 ACM International Symposium on Wearable Computers*. 208–211.
- [17] George Boateng, John A Batsis, Ryan Halter, and David Kotz. 2017. ActivityAware: an app for real-time daily activity level monitoring on the amulet wrist-worn device. In *2017 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*. IEEE, 431–435.
- [18] George Boateng, John A Batsis, Patrick Proctor, Ryan Halter, and David Kotz. 2018. GeriActive: Wearable app for monitoring and encouraging physical activity among older adults. In *2018 IEEE 15th International Conference on Wearable and Implantable Body Sensor Networks (BSN)*. IEEE, 46–49.
- [19] George Boateng and David Kotz. 2016. StressAware: An app for real-time stress monitoring on the amulet wearable platform. In *2016 IEEE MIT Undergraduate Research Technology Conference (URTC)*. IEEE, 1–4.
- [20] George Boateng, Prabhakaran Santhanam, Janina Lüscher, Urte Scholz, and Tobias Kowatsch. 2019. Poster: DyMand—An Open-Source Mobile and Wearable System for Assessing Couples' Dyadic Management of Chronic Diseases. In *The 25th Annual International Conference on Mobile Computing and Networking*. 1–3.
- [21] George Boateng, Prabhakaran Santhanam, Janina Lüscher, Urte Scholz, and Tobias Kowatsch. 2019. VADLite: an open-source lightweight system for real-time voice activity detection on smartwatches. In *Adjunct Proceedings of the 2019 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2019 ACM International Symposium on Wearable Computers*. 902–906.
- [22] Guy Bodenmann. 2005. Dyadic coping and its significance for marital functioning. (2005).
- [23] Niall Bolger and David Amarel. 2007. Effects of social support visibility on adjustment to stress: Experimental evidence. *Journal of personality and social psychology* 92, 3 (2007), 458.
- [24] Sheldon Cohen, Lynn G Underwood, and Benjamin H Gottlieb. 2000. Social support measurement and intervention: A guide for health and social scientists. (2000).
- [25] Tiantian Feng, Amrutha Nadarajan, Colin Vaz, Brandon Booth, and Shrikanth Narayanan. 2018. TILES audio recorder: an unobtrusive wearable solution to track audio activity. In *Proceedings of the 4th ACM Workshop on Wearable Systems and Applications*. ACM, 33–38.
- [26] Andreas Filler, Tobias Kowatsch, Severin Haug, Fabian Wahle, Thorsten Staake, and Elgar Fleisch. 2015. MobileCoach: A novel open source platform for the design of evidence-based, scalable and low-cost behavioral health interventions: overview and preliminary evaluation in the public health context. In *2015 Wireless Telecommunications Symposium (WTS)*. IEEE, 1–6.
- [27] Claudio Forlivesi, Utku Günay Acer, Marc van den Broeck, and Fahim Kawsar. 2018. Mindful interruptions: a lightweight system for managing interruptibility on wearables. In *Proceedings of the 4th ACM Workshop on Wearable Systems and Applications*. ACM, 27–32.
- [28] Jasara N Hogan, Alexander O Crenshaw, Katherine JW Baucom, and Brian RW Baucom. 2021. Time spent together in intimate relationships: Implications for relationship functioning. *Contemporary Family Therapy* 43, 3 (2021), 226–233.
- [29] John F Horton, Pro Stergiou, Tak S Fung, and Larry Katz. 2017. Comparison of Polar M600 optical heart rate and ECG heart rate during exercise. *Med Sci Sports Exerc* 49, 12 (2017), 2600–2607.
- [30] Chih-Wei Hsu, Chih-Chung Chang, Chih-Jen Lin, et al. 2003. A practical guide to support vector classification. (2003).
- [31] International Diabetes Federation Europe (IDF). 2022. IDF member(s) in Switzerland. Retrieved 7th May, 2022 from <https://idf.org/our-network/regions-members/europe/members/161-switzerland.html>
- [32] Masumi Iida, Mary Ann Parris Stephens, Karen S Rook, Melissa M Franks, and James K Salem. 2010. When the going gets tough, does support get going? Determinants of spousal support provision to type 2 diabetic patients. *Personality and Social Psychology Bulletin* 36, 6 (2010), 780–791.

- [33] Zhu Jianyong, Luo Haiyong, Chen Zili, and Li Zhaohui. 2014. RSSI based Bluetooth low energy indoor positioning. In *2014 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE, 526–533.
- [34] Alexander Karan, Robert C Wright, and Megan L Robbins. 2017. Everyday emotion word and personal pronoun use reflects dyadic adjustment among couples coping with breast cancer. *Personal Relationships* 24, 1 (2017), 36–48.
- [35] Cynthia M Khan, Mary Ann Parris Stephens, Melissa M Franks, Karen S Rook, and James K Salem. 2013. Influences of spousal support and control on diabetes management through physical activity. *Health Psychology* 32, 7 (2013), 739.
- [36] Noah Klugman, Veronica Jacome, Meghan Clark, Matthew Podolsky, Pat Pannuto, Neal Jackson, Aley Soud Nassor, Catherine Wolfram, Duncan Callaway, Jay Taneja, et al. 2018. Experience: Android Resists Liberation from Its Primary Use Case. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*. 545–556.
- [37] Tobias Kowatsch, Theresa Schachner, Samira Harperink, Filipe Barata, Ullrich Dittler, Grace Xiao, Catherine Stanger, Florian v Wangenheim, Elgar Fleisch, Helmut Oswald, et al. 2021. Conversational agents as mediating social actors in chronic disease management involving health care professionals, patients, and family members: multisite single-arm feasibility study. *Journal of medical Internet research* 23, 2 (2021), e25060.
- [38] Tobias Kowatsch, Dirk Volland, Iris Shih, Dominik Rügger, Florian Künzler, Filipe Barata, Andreas Filler, Dirk Büchter, Björn Brogle, Katrin Heldt, et al. 2017. Design and evaluation of a mobile chat app for the open source behavioral health intervention platform MobileCoach. In *International Conference on Design Science Research in Information System and Technology*. Springer, 485–489.
- [39] Daniyal Liaqat, Robert Wu, Andrea Gershon, Hisham Alshaer, Frank Rudzicz, and Eyal de Lara. 2018. Challenges with real-world smartwatch based audio monitoring. In *Proceedings of the 4th ACM Workshop on Wearable Systems and Applications*. 54–59.
- [40] LimeSurvey Project Team / Carsten Schmitz. 2012. *LimeSurvey: An Open Source survey tool*. LimeSurvey Project, Hamburg, Germany. <http://www.limesurvey.org>
- [41] Timothy J Loving and Richard B Slatcher. 2013. Romantic relationships and health. *The Oxford handbook of close relationships* (2013), 617–637.
- [42] Janina Lüscher, Tobias Kowatsch, George Boateng, Prabhakaran Santhanam, Guy Bodenmann, and Urte Scholz. 2019. Social Support and Common Dyadic Coping in Couples' Dyadic Management of Type II Diabetes: Protocol for an Ambulatory Assessment Application. *JMIR research protocols* 8, 10 (2019), e13685.
- [43] Janina Lüscher and Urte Scholz. 2018. Soziale unterstützung. *Psychologie in der Gesundheitsförderung. Bern, Germany: Hogrefe* (2018).
- [44] Matthias R Mehl, Megan L Robbins, and Fenne große Deters. 2012. Naturalistic observation of health-relevant social processes: The Electronically Activated Recorder (EAR) methodology in psychosomatics. *Psychosomatic Medicine* 74, 4 (2012), 410.
- [45] Dirk Merkel. 2014. Docker: lightweight linux containers for consistent development and deployment. *Linux journal* 2014, 239 (2014), 2.
- [46] Daisy Miller and J Lynne Brown. 2005. Marital interactions in the process of dietary change for type 2 diabetes. *Journal of Nutrition Education and Behavior* 37, 5 (2005), 226–234.
- [47] Maija Reblin, Richard E Heyman, Lee Ellington, Brian RW Baucom, Panayiotis G Georgiou, and Susan T Vadaparampil. 2018. Everyday couples' communication research: Overcoming methodological barriers with technology. *Patient education and counseling* 101, 3 (2018), 551–556.
- [48] Tuula-Maria Rintala, Pia Jaatinen, Eija Paavilainen, and Päivi Åstedt-Kurki. 2013. Interrelation between adult persons with diabetes and their family: a systematic review of the literature. *Journal of family nursing* 19, 1 (2013), 3–28.
- [49] Megan L Robbins, Elizabeth S Focella, Shelley Kasle, Ana María López, Karen L Weihs, and Matthias R Mehl. 2011. Naturalistically observed swearing, emotional support, and depressive symptoms in women coping with illness. *Health Psychology* 30, 6 (2011), 789.
- [50] Megan L Robbins, Ana María López, Karen L Weihs, and Matthias R Mehl. 2014. Cancer conversations in context: naturalistic observation of couples coping with breast cancer. *Journal of Family Psychology* 28, 3 (2014), 380.
- [51] Theodore F Robles, Richard B Slatcher, Joseph M Trombello, and Meghan M McGinn. 2014. Marital quality and health: A meta-analytic review. *Psychological bulletin* 140, 1 (2014), 140.
- [52] Michael J Rohrbaugh, Matthias R Mehl, Varda Shoham, Elizabeth S Reilly, and Gordon A Ewy. 2008. Prognostic significance of spouse we talk in couples coping with heart failure. *Journal of consulting and clinical psychology* 76, 5 (2008), 781.
- [53] Nina Rottmann, Dorte Gilså Hansen, Pia Veldt Larsen, Anne Nicolaisen, Henrik Flyger, Christoffer Johansen, and Mariët Hagedoorn. 2015. Dyadic coping within couples dealing with breast cancer: A longitudinal, population-based study. *Health Psychology* 34, 5 (2015), 486.
- [54] Ralf Schwarzer and Nina Knoll. 2010. Social support. *Health psychology, 2nd ed.* (2010), 283–293.
- [55] Amber J Seidel, Melissa M Franks, Mary Ann Parris Stephens, and Karen S Rook. 2012. Spouse control and type 2 diabetes management: moderating effects of dyadic expectations for spouse involvement. *Family relations* 61, 4 (2012), 698–709.
- [56] Laura Sels, Jed Cabrieto, Emily Butler, Harry Reis, Eva Ceulemans, and Peter Kuppens. 2019. The occurrence and correlates of emotional interdependence in romantic relationships. *Journal of personality and social psychology* (2019).
- [57] Adela C Timmons, Brian R Baucom, Sohyun C Han, Laura Perrone, Theodora Chaspari, Shrikanth S Narayanan, and Gayla Margolin. 2017. New frontiers in ambulatory assessment: Big data methods for capturing couples' emotions, vocalizations, and physiology in daily life. *Social Psychological and Personality Science* 8, 5 (2017), 552–563.

- [58] Adela C Timmons, Theodora Chaspari, Sohyun C Han, Laura Perrone, Shrikanth S Narayanan, and Gayla Margolin. 2017. Using multimodal wearable technology to detect conflict among couples. *Computer* 50, 3 (2017), 50–59.
- [59] David Watson, Lee Anna Clark, and Auke Tellegen. 1988. Development and validation of brief measures of positive and negative affect: the PANAS scales. *Journal of personality and social psychology* 54, 6 (1988), 1063.
- [60] Katharina Weitkamp, Fabienne Feger, Selina A Landolt, Michelle Roth, and Guy Bodenmann. 2021. Dyadic Coping in Couples Facing Chronic Physical Illness: A Systematic Review. *Frontiers in psychology* 12 (2021).
- [61] Wikipedia contributors. 2022. Transport Layer Security — Wikipedia, The Free Encyclopedia. [https://en.wikipedia.org/w/index.php?title=Transport\\_Layer\\_Security&oldid=1079821520](https://en.wikipedia.org/w/index.php?title=Transport_Layer_Security&oldid=1079821520) [Online; accessed 28-March-2022].

## Paper 8: Are you ok, honey?

# “Are you okay, honey?”: Recognizing Emotions among Couples Managing Diabetes in Daily Life using Multimodal Real-World Smartwatch Data

GEORGE BOATENG, ETH Zürich, Switzerland

XIANGYU ZHAO, TU München, Germany

MALGORZATA SPEICHERT, ETH Zürich, Switzerland

ELGAR FLEISCH, ETH Zürich, Switzerland and University of St. Gallen, Switzerland

JANINA LÜSCHER, University of Zürich, Switzerland

THERESA PAULY, University of Zürich, Switzerland

URTE SCHOLZ, University of Zürich, Switzerland

GUY BODENMANN, University of Zurich, Switzerland

TOBIAS KOWATSCH, ETH Zürich, Switzerland and University of St. Gallen, Switzerland

Couples generally manage chronic diseases together and the management takes an emotional toll on both patients and their romantic partners. Consequently, recognizing the emotions of each partner in daily life could provide an insight into their emotional well-being in chronic disease management. Currently, the process of assessing each partner’s emotions is manual, time-intensive, and costly. Despite the existence of works on emotion recognition among couples, none of these works have used data collected from couples’ interactions in daily life. In this work, we collected 85 hours (1,021 5-minute samples) of real-world multimodal smartwatch sensor data (speech, heart rate, accelerometer, and gyroscope) and self-reported emotion data (n=612) from 26 partners (13 couples) managing diabetes mellitus type 2 in daily life. We extracted physiological, movement, acoustic, and linguistic features, and trained machine learning models (support vector machine and random forest) to recognize each partner’s self-reported emotions (valence and arousal). Our results from the best models – balanced accuracies of 63.8% and 78.1% for arousal and valence respectively – are better than chance and our prior work that also used data from German-speaking, Swiss-based couples, albeit, in the lab. This work contributes toward building automated emotion recognition systems that would eventually enable partners to monitor their emotions in daily life and enable the delivery of interventions to improve their emotional well-being.

CCS Concepts: • **Human-centered computing** → **Ubiquitous and mobile computing**; • **Applied computing** → **Consumer health**; **Psychology**.

Additional Key Words and Phrases: Affective Computing; Emotion Recognition; Multimodal Sensor Data; Couples; Smartwatches; Wearable Computing; Speech Processing; Natural Language Processing; Machine Learning; Deep Learning; Transfer Learning; BERT; Chronic Disease Management

## 1 INTRODUCTION

For couples in which one partner has a chronic disease such as cancer and diabetes, their relationship plays a key role in the disease management if partners share the responsibility of its management [44, 51]. Such joint disease management, also called dyadic coping [5, 16, 43] takes an emotional toll on both patients and spouses [52]. Consequently, understanding each partner’s emotion within the context of their interactions and disease

---

Authors’ addresses: George Boateng, gboateng@ethz.ch, ETH Zürich, Zurich, Switzerland; Xiangyu Zhao, xiangyu.zhao@tum.de, TU München, Munich, Germany; Malgorzata Speichert, mspeichert@student.ethz.ch, ETH Zürich, Zurich, Switzerland; Elgar Fleisch, efleisch@ethz.ch, ETH Zürich, Zurich, Switzerland, University of St. Gallen, St. Gallen, Switzerland; Janina Lüscher, janina.luescher@psychologie.uzh.ch, University of Zürich, Zurich, Switzerland; Theresa Pauly, theresa.pauly@psychologie.uzh.ch, University of Zürich, Zurich, Switzerland; Urte Scholz, urte.scholz@psychologie.uzh.ch, University of Zürich, Zurich, Switzerland; Guy Bodenmann, guy.bodenmann@psychologie.uzh.ch, University of Zurich, Zurich, Switzerland; Tobias Kowatsch, tkowatsch@ethz.ch, ETH Zürich, Zurich, Switzerland, University of St. Gallen, St. Gallen, Switzerland.

management in daily life could enable the triggering of various dyadic interventions (where partners are both involved e.g. [33]) to improve their emotional well-being and chronic disease management.

However, assessing emotions among couples is challenging. Two approaches are used for emotion assessment in the lab and in daily life: self-report and observer reports. For self-reports, couples can be asked to have an emotionally charged conversation that is videotaped (e.g., in the lab), and then afterward, each partner provides emotion ratings, for example, while watching the videos [46] or by using a validated affect instrument such as the PANAS [55]. In the case of daily life, couples are periodically asked to complete self-reports [49] such as the PANAS which can be obtrusive and impractical for continuous emotion assessment. These ratings could be biased (for example, if the partner desires to project a certain emotion rating rather than how they really feel) and may not reflect the partner's actual emotion. For observers' reports, people are trained to watch the video recordings (e.g., in the case of lab data) and use a coding scheme to rate the emotional behavior of each partner (e.g., SPAFF [24]). Such coding is also done for example, for audio data collected from couples' daily life interactions [45]. This manual coding process is costly and time-consuming as multiple coders need to be trained for this task [32] and suffers from inter-rater reliability issues [29, 37]. Automated emotion recognition of each partner's emotion could potentially address these limitations. Current approaches for automatic emotion recognition among couples have all used data collected from the lab [11]. There exists no system that automatically recognizes the emotions of romantic partners using real-world data from couples' interactions in daily life. One potential reason for this gap is that collecting and processing such data is non-trivial, time-intensive, and costly [11].

Smartwatches have been used for mood recognition of individuals [18] and they could be leveraged for recognizing each partner's emotions based on the couple's interactions in daily life. Several features of smartwatches make them well suited for this task. They are mostly with the wearer as opposed to a smartphone which could be in various places like the pocket, or bag, and just not in proximity with the user. Also, consumer smartwatches could be used to collect a wide variety of sensor data that have been used for emotion recognition in the past: audio [50], heart rate, accelerometer, and gyroscope (for gestures e.g., [48]), and ambient light (to detect the context of couples). Multimodal fusion of these sensor data could produce better recognition results [25, 39]. Furthermore, smartwatches could be leveraged in novel ways (using Bluetooth signal strength and voice activity detection) to specifically capture partners' interaction or conversation moments in daily life [13] for use in emotion recognition.

In this work, we collected 85 hours (1,021 5-minute samples) of real-world multimodal smartwatch sensor data (speech, heart rate, accelerometer, and gyroscope) and trained machine learning models to recognize each partner's emotions. Specifically, we trained models to recognize each partner's emotional valence (negative vs positive) and emotional arousal (high vs low) during the conversation using sensor and self-report data from German-speaking, Swiss-based couples managing type 2 diabetes in daily life. We addressed the following research questions:

**RQ1:** *How well can romantic partners' emotions be recognized using multimodal real-world sensor data from daily life?*

**RQ2:** *Which modality and multimodal combinations produce the best emotion recognition results?*

This work is the first to recognize the emotions of romantic partners using data collected from everyday life. Our contributions are as follows (1) collection and use of a unique dataset — real-world, multimodal smartwatch sensor data from German-speaking, Swiss-based couples (N=13 couples, n=26 participants), which is the first such dataset used in the literature for automatic recognition of partners' emotions (2) approaches for validating and quantifying data quality on manually coded, annotated and transcribed real-world speech data (3) development and evaluation of a machine learning system to recognize the emotions of each partner using a wide variety of sensor data — acoustic, linguistic, heart rate, accelerometer, and gyroscope (4) an investigation of the sensor modality combinations which result in the best emotion recognition performance of romantic partners. This

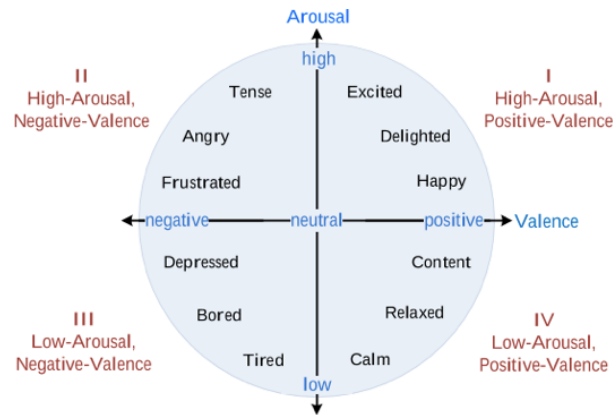


Fig. 1. Russell's circumplex model of emotions[57]

work is an extension of the work [10] and implements the research plan that was proposed in that work, along with machine learning experiments and reported results.

In the rest of this paper, we discuss background and related work in Section 2, methodology in Section 3, experiments and evaluation in Section 4 and results and discussion in Section 5, limitations and future work in Section 6, and we conclude in Section 7.

## 2 BACKGROUND AND RELATED WORK

In this section, we describe various emotion models, multimodal emotion recognition, and works that have been done to recognize emotions among couples.

### 2.1 Emotion Models

There are mainly two models of emotions used in the literature in emotion recognition: categorical and dimensional. Categorical emotions are based on the six basic emotions proposed by Ekman: happiness, sadness, fear, anger, disgust, and surprise [26]. Dimensional approaches mainly use two dimensions: valence (pleasure) and arousal which are based on Russell's circumplex model of emotions [47]. Valence refers to how negative to positive a person feels and arousal refers to how sleepy to active a person feels. Using these two dimensions, several categorical emotions can be placed and grouped into the four quadrants: high arousal and negative valence (e.g., angry), low arousal and negative valence (e.g., depressed), low arousal and positive valence (e.g., relaxed), and high arousal and positive valence (e.g., excited) [47].

### 2.2 Multimodal Emotion Recognition

Multimodal fusion entails combining data collected from various modalities and leverages the idea that data contained in different modalities could provide a better understanding of a certain context. Various works have employed multimodal fusion approaches for emotion recognition and they have been shown to give better results than unimodal approaches [25, 39]. There are two main fusion approaches – fusion at the feature level (early fusion) and at the decision level (late fusion). Early fusion entails combining features from different data modalities, for example, through concatenation and feeding them into the same machine learning algorithm. For late fusion, a separate algorithm is used for each data modality and then the predictions of the individual algorithms are combined using, for example, majority voting. Additional approaches include some hybrid of



early and late fusion [56] and model-level fusion which leverages interactions between different modalities at the model level e.g [30].

### 2.3 Emotion Recognition among Couples

Emotion recognition among couples is the task of recognizing the emotion of each romantic partner based on the context of their interaction /conversation [11]. Specifically, it entails recognizing each partner's emotions for every utterance/speaker turn, every few seconds, or for the whole conversation. It differs from other kinds of emotion recognition tasks mainly by the kind of stimuli that induces emotions. Some stimuli are driving [58], listening to music or watching a movie [2], and conversation between people [40]. Couples' emotion recognition is similar to emotion recognition tasks whose stimuli are conversations since it uses a conversational context. However, its uniqueness lies in the fact that the two interacting individuals are in a romantic relationship. Consequently, various insights from psychology about couples' interaction dynamics could be leveraged to recognize each partner's emotions. For example, romantic partners influence each other when interacting, and that insight has been used for couples' emotion recognition (e.g., [12, 21]).

There are several works that have developed machine learning systems to recognize the emotions among couples (see [11] for a detailed overview of the research field). Most of these works have been done by the Signal Analysis and Interpretation Laboratory (SAIL) team at the University of Southern California [11]. The works have mainly used emotion labels from external raters, support vector machines as the algorithm, the following three modalities — acoustic, lexical, and visual — with acoustic being the most used modality, and feature-level fusion of acoustic and lexical modalities [11]. There are other related work focused on recognizing behaviors among couples other than emotions such as level of blame [8, 9], conflict [53], and suicidal risk [23].

Most of these works have used observer ratings (perceived emotions) rather than self-reports (one's actual emotions) as labels. Consequently, the emotion recognition task essentially becomes recognizing external individuals' perception of each partner's emotion rather than each partner's emotion per their own assessment. Though similar, the latter is more challenging. For observer ratings, coders are generally trained over several weeks, and various approaches are used to resolve ratings that are not in agreement and ensure the validity of the labels. Also, the self-reported emotion may not be reflected in that partner's behavior in comparison to observer ratings which are purely based on behavioral observation.

Also, several of these works have used data from English-speaking couples in the U.S. with a few using data from German-Speaking couples in Switzerland [7, 12] and Dutch-speaking couples in Belgium [15]. Additionally, several modalities such as physiological data, hand gestures, and body movement have not been explored. More importantly, none of these works have used data collected from couples' interactions in daily life. Our work fills the current research gap by performing emotion recognition using multimodal real-world smartwatch data — speech, accelerometer, gyroscope, heart rate — and self-reported emotion data collected from German-speaking, Swiss-based couples.

### 2.4 Emotion Recognition using Smartwatch Data

There are a number of works that have performed emotion recognition using smartwatch data. AlHanai et al. [3] trained neural network models to recognize emotions using smartwatch and smartphone data collected from 10 subjects who told 31 personal stories (15 happy) in a lab. They used an iPhone to collect audio which they transcribed. They also collected physiological and movement data with the smartwatch. They extracted 386 acoustic (functionals over low-level descriptors), linguistic (average positive and negative sentiment of words), physiological, and movement features (mean, median, variance of electrocardiogram, photoplethysmogram, accelerometer, gyroscope, bioimpedance, electric tissue impedance, galvanic skin response, and skin temperature) and selected 10 features for use using sequential forward features selection. They classified the whole narration

as happy or sad (indicated by the subject) and 5-sec segments as positive, negative, or neutral (annotated by a research assistant with balanced distribution). Though they used naturalistic data (personal narratives), this work did not use data collected from an uncontrolled, real-world context such as in daily life.

Budner et al. [18] trained Random Forest models to recognize moods using smartwatch data collected from 60 subjects in daily life. They classified 9 mood states (angry, sad, tired, excited, happy, quiet, elated, very happy, relaxed) and 3 levels of pleasure and activation. They extracted the following features related to body sensor data: vector magnitude counts (a measure of the total amount of movement), heart rate, and external influences: light level and GPS coordinates (variance), weather, (humidity, temperature, cloudiness, windiness, air pressure), the hour of the day, whether it was the weekend and day of the week. Arano et al. [4] built upon the work by Budner et al. [18] and proposed the use of the smartwatch-based system to measure emotions in a real-world scenario: classroom. They were able to collect data from 30 subjects related to body sensor data: accelerometer, light, audio, heart rate (from a smartwatch), GPS data (from a smartphone), and environmental variables (e.g., weather, longitude, latitude, altitude, room temperature, humidity, pressure, wind level, clouds level, noise level). Subjects indicated their level of Activation, Tiredness, Pleasance, Quality (of lecturer's presentation), and Understanding (of lecturer's presentation) on a scale from 0-2. They extracted statistical features and used 9 models: K-Nearest Neighbor, Decision Trees, Support Vector Machines, Multilayer perceptron, logistic regression, Gradient Boost, XGBoost, and LSTM.

Kanjo et al [31] developed models to recognize emotions using body sensor and environmental data collected from the wild. They collected data from 40 females walking around Nottingham city center, UK for 45 mins: body movement, activity, heart rate, Electrodermal activities and body temperature and, environmental data including noise level (Env-noise), air pressure and ambient light levels, and GPS data. User emotions labels are collected using self-report input, based on a scale for valence (1-5). They used the Microsoft Band and Android phones (to collect noise, GPS, and self-report). They extracted 87 features: mean, median, max, min, range, and standard deviation and quartiles and selected 21 after feature selection. They trained ensemble models (stacking) to perform classification of the 5 levels of valence. They had a base model for each modality and a stacking model which fused the results of both models. They used the following models: Support Vector Machine, Random Forest, and K Nearest Neighbour as the base models, and Naive Bayes as the stacking model Learner which fused the base models' predictions.

Quiroz et al. [41] developed a smartwatch-based method to recognize emotions based on movement data. They collected data from 50 subjects: (43 females; mean age 23.18 [SD 4.87] years), North-West, UK. They collected emotion data with the PANAS before and after emotion elicitation; happy, sad, and neutral. They used audiovisual movie clips and audio music clips to elicit emotions. They asked the subject to walk for 250 meters while wearing a smartwatch and heart rate monitor strap on the chest. 18 were assigned to the audiovisual condition and watched the movie before walking. Out of the 32 assigned to audio, half of them listened while walking. It took 20 mins for each subject. They extracted 107 features over 1-second sliding windows with 50% overlap over filtered signal (accelerometer and gyroscope, and heart rate). They trained personalized models from 44 subjects to classify happy vs sad and happy vs sad vs neutral. The data was balanced. They used 10-fold stratified cross-validation with logistic regression and random forest.

Schmidt et al, 2019 [48] trained a convolutional neural network (CNN) model to predict emotions (arousal, valence), anxiety, and stress from real-world smartwatch-based physiological and motion data. They used an Empatica E4 to collect 1,400 hours of data; accelerometer, photoplethysmogram (PPG), EDA, and skin-temperature data from 12 subjects (7 male). Subjects received EMA prompts every 2-2.5 hours or triggered manually: 1) self-assessment mannequins assessing valence and arousal 2) State-Trait Anxiety Inventory (STAI) on 6 levels and 3) Stress level scored on a four-point Likert scale. The data was skewed for all the labels. They preprocessed the data resulting in 1083 valid windows/questionnaires. The data was split between 3 levels for all the labels except stress which was binarized. They extracted 62 features (e.g., mean, standard deviation, heart rate, heart

rate variability). They used leave-one-subject-out cross-validation (LOSO) and leave-target-questionnaires-out (LTQO). As baseline models, they used different tree-based classifiers (decision-tree (DT), randomized decision trees (ET), and random forest (RF)). They used a single-task and multi-task CNN which takes the raw sensor data as the main model with late fusion.

Park et al., 2020 [38] developed WellBeat, a smartwatch-based system for assessing the emotional well-being of individuals. They used a Samsung smartwatch to collect PPG and heart rate data from 12 subjects (3 female) continuously throughout the day: 1121 hours of data from a 3-week study (about 445 hrs eliminated) and 1032 self-report labels related to happiness, awakeness, and relaxedness levels (1-5). Subjects were asked to complete the self-report 3 times a day at random times during waking hours. They performed data preprocessing by removing samples where the watch was not worn, partitioned data into consecutive 5-min slices, filtered out signals without heart rate signals, extracted heart rate and RR intervals, and HRV parameters such as RMSSD, and estimated their validity. The label was matched to the day -10 mins to +10 mins around the label timestamp similar to Schmidt et al (2019). They performed classification with logistic regression and 10-fold cross-validation.

Our work builds upon some of these works by using similar preprocessing approaches and features (for heart rate, accelerometer data, speech data, gyroscope), algorithms and evaluation approach. One modality that is missing in most of these works is the linguistic modality. We leveraged recent advances in deep learning and natural language processing to extract linguistic features from speech. Also, we systematically evaluated the performance of individual modalities and various modality combinations. The key way our work differs from these works though is our use of the context of couples' interactions in daily life to recognize each partner's emotion.

### 3 METHODOLOGY

In this section, we describe how we collected and preprocessed the data and the features that we extracted.

#### 3.1 Data Collection

We developed DyMand, an open-source smartwatch and smartphone system which we used to collect data from couples in daily life in a user study (see [13] for a detailed description). The DyMand system (Figure 2) consists of a smartwatch app, and a smartphone app built on top of the MobileCoach platform [28, 34] that consists of a web-based intervention designer and backend.

We ran the DyMand study between 2019 and 2021 with heterosexual romantic couples from the German-speaking part of Switzerland in which one partner had type 2 diabetes [35]. In total, we collected 85 hours of sensor and self-report data from 13 couples aged 47 to 81 years, with a mean age of 68 (SD = 9).

The study was advertised in various places including hospitals, magazines, local newspapers, and the diabetes association in Switzerland. Interested couples completed a web-based questionnaire to screen them for the inclusion and exclusion criteria, and collect socio-demographic information. Those who met the eligibility criteria were able to pick a date for a baseline assessment. During this session, both partners received comprehensive information about the study, signed the informed consent form, and completed a web-based questionnaire that captured constructs of interest at baseline that were not assessed daily.

Each partner was given a smartwatch (Polar M600 running Wear OS) and a smartphone (Nokia 6.1 running Android 9.0), both paired and running the DyMand apps. They also received instructions on the study and then trained research assistants helped them to set up their devices and pair the corresponding smartphone and smartwatch. They were instructed to have all devices with them every day for 7 days from getting up until going to bed. To prevent mistakes from one partner accidentally using the other partner's watch and phone, one set of phones and watches had black covers and the other set had white covers. The patient was given the white set and the supporting partner was given the black set. The partners picked the hours during which we could collect

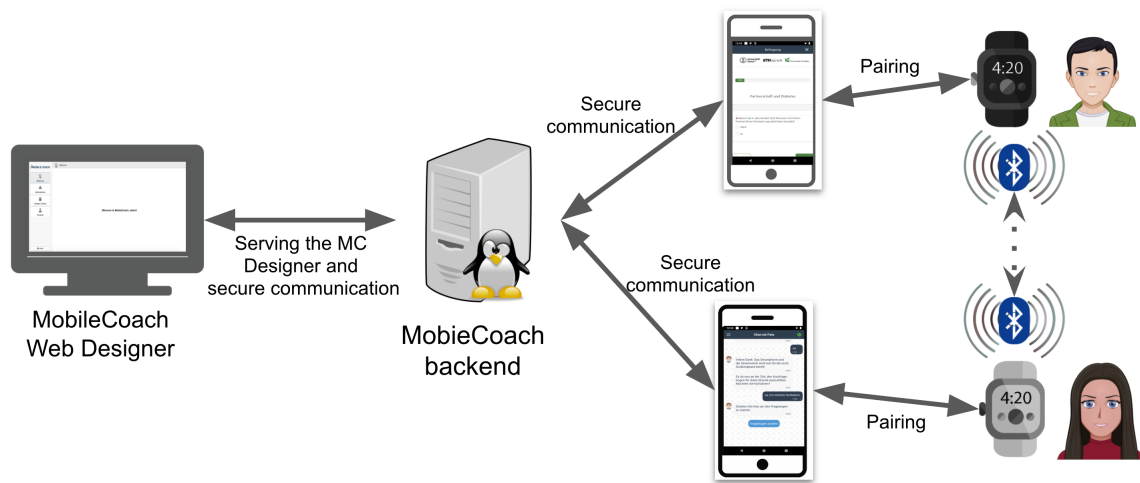


Fig. 2. Overview of the DyMand system

data from them. During the week, they could choose a period for the morning hours (any time between 4 am to 11 am, at least 2 hrs) and a period for the evening hours (any time between 4 pm to 11 pm, at least 2 hrs). During the weekend, data was collected all day and couples chose a start time in the early morning hours and an end time in the late evening hours (e.g., from 6 am to 10 pm). With this procedure, privacy aspects were addressed by primarily focusing on situations, in which the couples spent time together and thus reducing the number of audio recordings during the day of weekdays when chances are higher that subjects are working or moving around in public places.

We collected data from their daily life for 7 consecutive days starting the next Monday after their visit until the following Sunday night. The DyMand system triggered the collection of sensor and self-report data for 5 minutes each hour during the hours that partners pick. We collected the following sensor data from the smartwatch: audio, heart rate, accelerometer, gyroscope, Bluetooth low energy (BLE) signal strength between watches, and ambient light.

We collected a maximum of 5 minutes of data per hour for privacy reasons. Hence, to optimize the quality of data collected within that hour and to ensure that we recorded the most relevant 5 minutes of data (when partners are interacting), rather than triggering data collection at random or scheduled times which is the norm ([36, 45]), the app on each of the two smartwatches collected data when 1) the partners were physically close and 2) when there was speech (see [13] for the full details).

Our algorithm used a two-step process. First, the app determines physical closeness using the BLE signal strength between the two smartwatches with one acting as the central and the other acting as the peripheral. The central smartwatch scans for the peripheral device, and checks if the signal strength between them is within a certain threshold, which corresponds to a distance estimate. If this condition is met, the app on the central device determines if the partners are speaking by using a voice activity detection (VAD) machine-learning algorithm, which is implemented on the smartwatch [14]. If these two conditions are met, the central device connects with the peripheral smartwatch, starts recording, and also sends a signal to the peripheral watch to also start recording. Consequently, each smartwatch records the same interaction, albeit, sometimes with a start delay of a few seconds



Fig. 3. Emotion rating with Affective Slider

on the peripheral smartwatch. It is important to note that depending on the proximity of the partners, and the presence of other individuals, parts of ongoing conversations were captured to different degrees separately on each of the smartwatches even for the same recording time period. Hence it cannot be assumed that the two recordings at the same hour and minute are exact duplicates. For example, there was a case where both partners were together with two friends with all four being in proximity, having conversations. Yet, the male partner was talking directly with the male friend, and the female partner was talking directly with the female friend, and though the two smartwatches were recorded at the same time, they captured different conversations.

In the case in which the condition of physical closeness and speaking is not met in the hour, the app triggers a backup recording in the last 15 minutes of the hour. Our evaluation showed this approach for triggering data collection to capture conversation moments between partners performed better than the backup recording [13]. The app also ensured that there were at least 20 minutes between subsequent data collection to reduce the burden of the partners completing the self-reports.

After the 5-minute sensor data collection, the smartwatch vibrates and triggers a self-report on the smartphone for that partner to complete. The self-report asks about emotions over the last 5 minutes using the Affective Slider, a digital affect measuring tool that assesses the valence and arousal dimensions of their emotions [6]. In particular, they responded to “how unhappy vs. happy did you feel in the last 5 minutes?” and “how tired vs. awake did you feel in the last 5 minutes” by moving a slider from 0 to 100 on a visual scale — the Affective Slider (Figure 3). If the smartwatch does not receive a message from the smartphone app within 2 mins indicating that the self-report has been started, it gives another vibration alert. If once more, within the next 2 minutes, there is still no response about the start or completion of the self-report, it implies the self-report was not completed. The self-report is then dismissed. Doing this ensured that we collected data with matching sensor and self-report samples. For privacy reasons, the app deletes that audio sample if the self-report is not completed and attempts to trigger another sensor data collection and self-report later in the hour, optimizing for the case detection partner’s

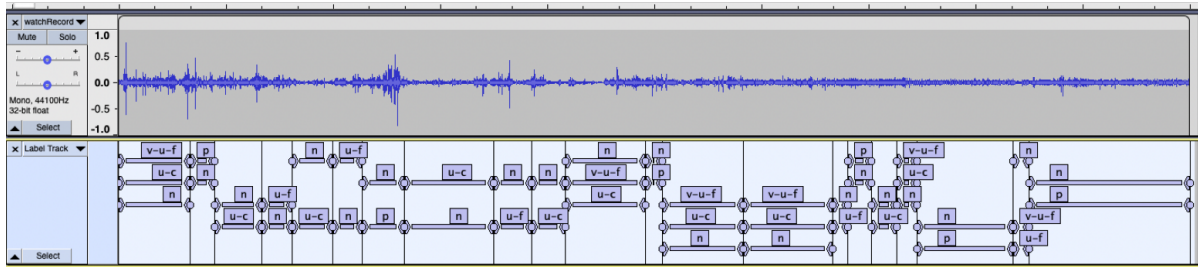


Fig. 4. Screenshot of the annotation process of the audio

interactions. Other sensor samples are still kept which could result in several sensor samples per hour without audio. If a backup recording is done, which implies that it was the last recording in that hour, the audio is not deleted even if the self-report is not completed. Doing this ensures we have at least one audio recording per hour. This approach resulted in a significant number of sensor recordings without labels. At the end of the day, the system triggered the Affective Slider, and also a short form of the PANAS self-report [55] for the couples to report their emotions over the whole day.

There are significant ethical and privacy concerns of such a system and study since we collect audio which is sensitive data, and more so in the context of couples' interactions with the likelihood of speech about private topics. We took several measures as follows. First, our study received ethical clearance from the cantonal ethics committee of the Canton of Zurich, Switzerland (Req-2017\_00430). Second, we ensured that we collected a maximum of 5 minutes of audio per hour in order not to record a significant percentage of the couples' everyday life. Consequently, even if the system triggered multiple recordings in the hour, the app always deleted all but the last one before the end of the hour. Third, to protect the privacy of subjects not taking part in the study, we asked subjects to wear a tag that we give them to indicate to others around that recording may be happening and that they may be recorded. Finally, after subjects returned their devices, we gave them the option to listen to and request the deletion of any audio samples without any explanation before the study team could listen to the audio files. Similar measures have been used in previous studies [36, 45] and have proven adequate to safeguard the privacy of study subjects and others not taking part in the study.

### 3.2 Data Annotation, Transcription and Coding

Four trained research assistants (RA) annotated, transcribed, and coded the audios. Using the software Audacity, each 5-minute audio was annotated with the start and end times of the speaker turns of each partner (m, f), unknown speakers (u), cross-talk between partners (c), vocalizations such as laughs, sighs (v) and the context (e.g., TV, radio), silence with no one speaking (p), noise such as music, movements of the watch, vehicles, etc. (n) and speech from radio or tv (u-tv/radio) (Figure 4).

The speech of both partners within each audio was transcribed in separate documents. In particular, each Microsoft Word document was used for each partner and RAs wrote the transcript in 15 secs chunks with “//” to separate the chunks. RAs wrote the German equivalent of any Swiss German words that were used since Swiss German is not a written language and there are different dialects of Swiss German spoken in Switzerland. Words that were not intelligible were written as “XY” in the document.

RAs coded the context of each audio in a spreadsheet as they listened to the audio using a protocol based on Mehl et al. [36]. They indicated if the audio contained speech, each partner spoke, and there was a conversation and a conversation between both partners. They also provided information about the conversational context

Context (what is going in this audio)	General Information about Participant	Current location	Activity	Interaction partner(s)	Emotional expression	Conversation type
	1-completely alone 2-alone but others are nearby (eg in café) 3-with one person 4-with group	1-indoors 2-home 3-at work 4-outdoors 5-in transit 6-healthcare center	1-listening to radio/music 2-watching TV 3-eating 4-talking on phone/computer 5-social interaction 6-cooking/baking 7-walking 8-healthcare related activity 9-housework 10-sleeping	1-alone 2-friend(s) 3-romantic partner 4-other family 5-stranger(s) 6-pet(s) 7-professional caregiver other (write the new option here)	1-laughter 2-crying 3-sighing 4-funny/goofy 5-singing 6-neutral 7-concerned other (write the new option here)	1-practical 2-small talk 3-substantive 4-gossip 5-emotional disclosure

Fig. 5. Screenshot of the coding context options

(what was going on in the audio), location, interaction partners, conversation type, activity, and emotional expression (Figure 5)

The real-world nature of the data posed challenges for our RAs. There were cases where the partners were having conversations with friends which made it difficult to distinguish the voices. For other cases, one partner was far away from the smartwatch while speaking, making it difficult to hear their voice. Given the challenging nature of the annotation, transcription, and coding tasks using real-world data (85 hours of audio) and the susceptibility to error, we implemented several manual and automatic approaches to perform sanity checks. We reviewed the codes to make sure the entries for different fields were consistent. For example, we cross-checked that if it is indicated that both partners spoke, then the field “interaction partner” should be “romantic partner”. Also, for each audio file, we automatically checked if there existed a non-empty transcription file and an annotation file with ‘m’ or ‘f’ if there was a “yes” for “male spoke” or “female spoke”. We also verified the accuracy of the annotations by automatically checking that for each 15-second chunk in the transcript file that contained text for the male or female partner, there existed an ‘m’ or ‘f’ in the corresponding 15-sec time period in the annotation file. We computed a percentage overlap for ‘m’ or ‘f’ with the corresponding transcript as a proxy for the quality of the annotation of that audio. Furthermore, we computed the percentage of “XY”s – inaudible words – for each audio that had speech which was a proxy for the quality of the audio and difficulty of the transcription task for that audio. RAs were given a list of files that failed these checks to then fix.

### 3.3 Data Preprocessing

We had a total of 612 self-report samples consisting of arousal and valence ratings (0 - 100) of each partner collected using the Affective Slider after the sensor data collection. We had a total of 1021 5-minute samples of sensor data (85 hours) consisting of audio, heart rate, accelerometer, gyroscope, and ambient light collected from each partner’s smartwatch. Some of the 5-minute samples were without audio data since the data collection protocol resulted in the deletion of audio samples without completed self-reports due to privacy reasons. Furthermore, because of software errors, a few of the sensor data were collected outside the data collection window and some audios were corrupted and hence could not be played or processed. We inferred these audios by eliminating audios whose size was smaller than the expected file size for 5-minute audio. Consequently, we automatically selected 5-minute samples that met the following conditions: 1) had both audio (non-corrupted) and other sensor data, and 2) were within the data collection hours specified by the partners. In total, that resulted in 1014 5-minute samples. Given our task is a supervised learning task, we also filtered and selected sensor samples that had a corresponding completed self-report.

We binarized the arousal and valence data into high (above 50) and low (less than or equal to 50) for arousal and negative (less than or equal to 50) and positive (above 50) for valence similar to the approach by previous

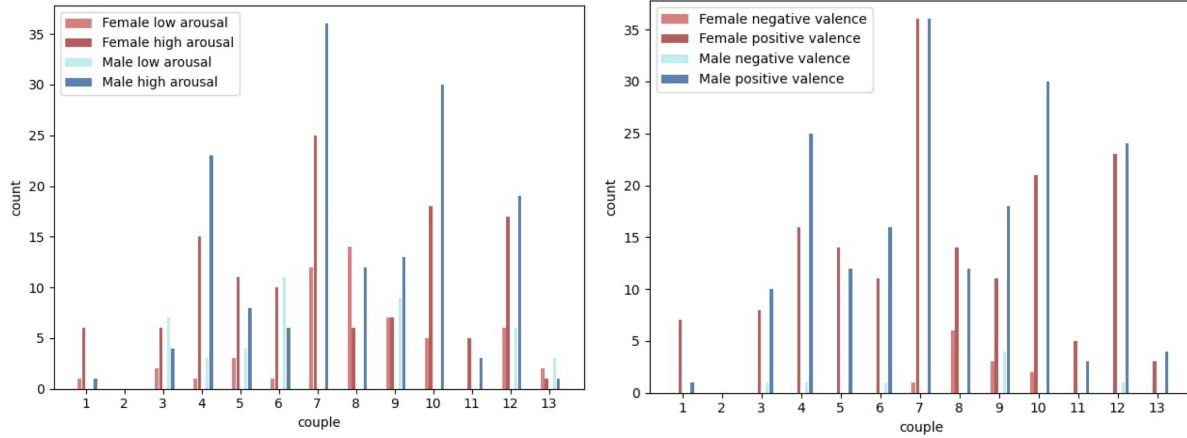


Fig. 6. Distributions (bar chart) of data for arousal (left) and valence (right) per couple per gender

Table 1. Number of negative valence and positive valence, and high arousal and low arousal samples per gender after applying selection criteria to 5-minute samples 1) had both audio (non-corrupted) and other sensor data, and 2) were within the data collection hours specified by the partners.

	Valence		Arousal	
	Negative	Positive	Low	High
Male	8	191	43	156
Female	12	169	54	127
<b>Total</b>	<b>20</b>	<b>360</b>	<b>97</b>	<b>283</b>

works [7, 12, 15]. With binarization, the binarized arousal and valence labels can be mapped to one of the four quadrants of Russell’s circumplex model of emotions, enabling its usefulness in the real world since we can tell which group of emotions are being felt by each partner. We split at 50 because with the design of the Affective Slider, 50 was understood to be the midpoint for the labels while partners responded to it. Furthermore, taking a per subject median rating as the midpoint would be problematic if there is not a good distribution of ratings for that partner (e.g., if there are just 3 ratings which are all 80, 90, and 100, it will not be correct to assume that 80 implies negative emotions for them). Next, we filtered for samples for which we had ‘yes’ for ‘male spoke’ and ‘female spoke’ as a proxy for the context of a conversation between both partners. This filtering resulted in 380 sensor-self-report samples: 20 negative valence and 360 positive valence, 97 low arousal, and 283 high arousal. The data is highly imbalanced which is typical of real-world emotion data. Table 1 shows the sensor-self-report samples per gender. Figures 6, and 7 show the distributions of low and high arousal and negative and positive valence per couple per gender. We observe the skewness of the labels per couple with some couples’ data not containing any negative valence samples (1, 2, 5, 11, 13) and any low arousal samples (2, 11).

We filtered the audio using a low-pass filter with a cut-off frequency of 4 kHz given it was collected as raw audio at a sampling rate of 44.1KHz and human speech is less than 4 kHz. For each 5-minute data for all the modalities, we removed outliers (data points that were more than 2 standard deviations from the mean). We resampled the data points at 50 Hz for accelerometer and gyroscope and 1 Hz for heart rate given that the Wear OS platform did not sample the signal at our specified frequency in the app. We additionally preprocessed the



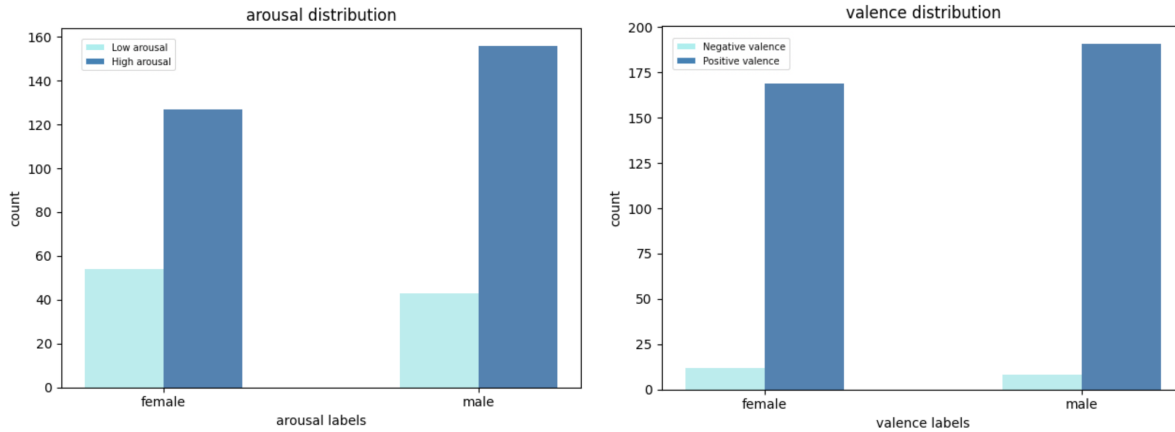


Fig. 7. Total distribution (histogram/bar chart) of data for all couples per gender for arousal (left) and valence (right)

heart rate data to remove outlier samples such as heart rate values outside the normal range (30 to 200 beats per minute). We automatically inferred samples for which the watch was not worn. For each physiological and movement sample, we logged a value (between 0 and 3) that provides a confidence estimate of the wear state of the watch. If 50% of the data points within the 5-minute sample had a value of 0 for the wear state, we marked the same as non-worn.

### 3.4 Features Extraction

We describe the extraction of physiological, movement, context, linguistic and acoustic features.

**3.4.1 Physiological.** Similar to prior work [31], we extracted the following statistical features from the heart rate data: mean, median, max, min, 25th percentile, 75th percentile, standard deviation, range, skewness, and kurtosis. The extraction resulted in a 10-dimensional feature vector.

**3.4.2 Movement.** Similar to prior work [41, 48], we extracted the following statistical features from the accelerometer and gyroscope: mean, median, max, min, 25th percentile, 75th percentile, standard deviation, range, skewness, and kurtosis. For the accelerometer and gyroscope data, we computed the magnitude of the x, y, and z axes before using them to compute the features. We did this so that the orientation of the device does not affect the results. The extraction resulted in a 10-dimensional feature vector.

**3.4.3 Acoustic Features.** For speech, we used openSMILE to extract the 88 eGeMAPS features using the annotations corresponding to the sections of the audio where each partner spoke similar to prior work on couples' emotion recognition [7, 12, 22]. These features have been shown to be a minimalist set of features adequate for emotion recognition [27]. Researchers use the extraction of acoustic parameters from the speech signal as a method to understand the patterning of the vocal expression of different emotions and other affective dispositions and processes. They used a number of acoustic parameters, including parameters: time-domain (e.g., speech rate), frequency domain (e.g., fundamental frequency or formant frequencies), amplitude domain (e.g., intensity or energy), distribution domain (e.g., relative energy in different frequency bands).

The use of machine learning led to the increase in the variety and quantity of acoustic features employed: basic (low-level ones) and derived (functionals). Therefore, finding relevant acoustic parameters is crucial in order to understand the mechanism of production and perception of emotions. Minimalistic standard parameters set for

acoustic analysis of speech and other vocal sounds might lead to better generalization in real-world scenarios. There are three criteria that guided the choice of parameters: the potential of an acoustic parameter to index physiological changes in voice production during affective processes; the frequency and success with which the parameter has been used in the past literature; its theoretical significance

The minimalistic acoustic parameter set contains 18 Low-level descriptors (LLD), which are grouped into frequency-related parameters, energy/amplitude-related parameters, and spectral (balance) parameters. They are smoothed over time with symmetric moving average filters 3 frames long (for pitch, jitter, and shimmering, only performed within voiced regions). The following functionals are applied:

- arithmetic mean and coefficient of variation (standard deviation normalized by arithmetic mean) to all 18 LDD
- 20th, 50th, and 80th percentiles, range of 20th to 80th percentile, and mean and std of the slope of rising/falling signal parts are added to loudness and pitch
- arithmetic mean of Alpha Ratio, Hammarberg Index, spectral slopes from 0-500Hz and 500-1500Hz over all unvoiced segments
- rate of loudness peaks, mean length & standard deviation of continuously voiced regions, mean length and std of unvoiced regions, number of continuous voiced regions per second

The above functionals yield 62 parameters in the Geneva Minimalistic Standard Parameter Set. For extended Geneva Minimalistic Standard Parameter Set (eGeMAPS) the following are added so the final set contains 88 parameters: arithmetic means and coefficients of variation are applied to 7 additional LLD to all segments; arithmetic mean of the spectral flux in unvoiced regions, arithmetic mean and coefficient of variation of the spectral flux and MFCC 1-4 in voices regions + equivalent sound level. In evaluations, eGeMAPS was shown to be superior or equal to the GeMAPS [27]. Hence, we extracted the eGeMAPS features resulting in an 88-dimensional feature vector.

*3.4.4 Linguistic Features.* We extracted linguistic features from the transcripts of the whole 5-minute interaction using a pre-trained model – Sentence-BERT (SBERT) [42] Similar to prior work [7, 12]. Sentence-BERT is a modification of the BERT architecture with siamese and triplet networks to compute sentence embeddings such that semantically similar sentences are close in vector space. Sentence-BERT has been shown to outperform the mean and CLS token outputs of regular BERT models for semantic similarity and sentiment classification tasks. Given that the transcripts are in German, we used the German BERT model [1] as SBERT’s Transformer model and the mean pooling setting. The German BERT model was pre-trained using the German Wikipedia dump, the OpenLegalData dump, and German news articles. The extraction resulted in a 768-dimensional feature vector.

### 3.5 Unimodal and Multimodal Fusion

We used the features of each modality separately as input for our machine learning experiments: physiological (heart rate), movement (accelerometer and gyroscope), acoustic and linguistic. We also used a multimodal approach with feature-level fusion [12, 22, 54]. We compared performance for individual modalities and various modality combinations as follows to answer our two research questions: physiological and movement, acoustic and linguistic, and physiological, movement, acoustic and linguistic.

## 4 EXPERIMENTS AND EVALUATION

We trained models for each gender to perform binary classification for arousal and valence. We trained separate models for each gender since gender differences affect how people express their emotions [17]. Hence, building gender-specific models [12, 15] may benefit the emotion recognition task. Similar to prior work [11], we performed couple disjoint cross-validation in which data from the same couples are never in both the train and test sets. This evaluation approach is a specific form of the subject independent evaluation but more robust as it accounts

Table 2. Evaluation results (balanced accuracy) for unimodal and multimodal models for arousal and valence for each gender

Modalities	Arousal (%)		Valence (%)	
	Male	Female	Male	Female
<b>Unimodal</b>				
Physiological	51.7	54.2	62.9	49
Movement	<b>59.6</b>	62.9	58.2	50
Linguistic	58.2	<b>63.2</b>	59.9	<b>64.8</b>
Acoustic	54.8	56.8	<b>78.1</b>	61.5
<b>Multimodal</b>				
Physiological and Movement	54	<b>62.3</b>	62.3	48
Linguistic and Acoustic	<b>63.8</b>	55.9	<b>62.6</b>	<b>64.9</b>
Physiological, Movement, Linguistic and Acoustic	59.7	59.1	59.6	<b>64.9</b>

for the situation in which data from one partner (e.g., speech) may be contained in the data of the other partner [11]. We did not perform leave-one-couple-out cross-validation which is the most used evaluation approach in couples’ emotion recognition tasks [11]. Given that most couples did not have negative samples, using this evaluation approach could lead to inflated results since the model could just predict all positive results without any learning. Rather, we performed 3-fold couple disjoint stratified cross-validation. In this setup, we trained on two folds, perform prediction on the third fold as a test set, and repeated this process with each fold serving as the test fold. The stratification aspect ensures that the same ratio for classes is maintained in the train and test splits, guaranteeing that each test fold will have some negative samples. The predicted labels of each test fold are combined and the evaluation metric is computed. We used the metric balanced accuracy / unweighted average recall (UAR) due to data imbalance and confusion matrices to perform an evaluation of the predictions. We also performed hyperparameter tuning within the train split using 2-fold stratified cross-validation. We used the following machine learning models: random forest (RF), support vector machines – linear, and radial basis function (RBF) with the ‘weight’ hyperparameter set to ‘balanced’ to account for the class imbalance. We used a random baseline of 50% for comparison.

## 5 RESULTS AND DISCUSSION

The results of the best models are shown in Table 2 to answer the research questions. Among the unimodal models, for arousal, movement, and linguistic modalities performed the best for male partners (59.6%) and female partners (63.2%) respectively, and for valence, acoustic and linguistic performed the best for male partners (78.1%) and female partners (64.8%) respectively. Among the multimodal models, for arousal, “Linguistic and Acoustic” and “Physiological and Movement” performed the best for male partners (63.8%) and female partners (62.3%) respectively, and for valence, “Linguistic and Acoustic” performed the best for valence for both male partners (62.6%) and female partners (64.9%), with “Physiological, Movement, Linguistic and Acoustic” also performing the same for female partners. Figure 8 shows the confusion matrices for valence’s and arousal’s best models. The linguistic and acoustic modalities produced most of the best results alone or in combination, particularly for valence, which indicates that what partners say and how they speak during their conversations are the most informative for recognizing how negative or positive they feel. This result is in line with the use of these two modalities in several couples’ emotion recognition works [11]. Furthermore, the movement modality alone or in combination with physiological modality performed the best for arousal. This result is consistent with the intuition that the greater body and hand movement are expected the more active a person feels – the arousal dimension of emotion.

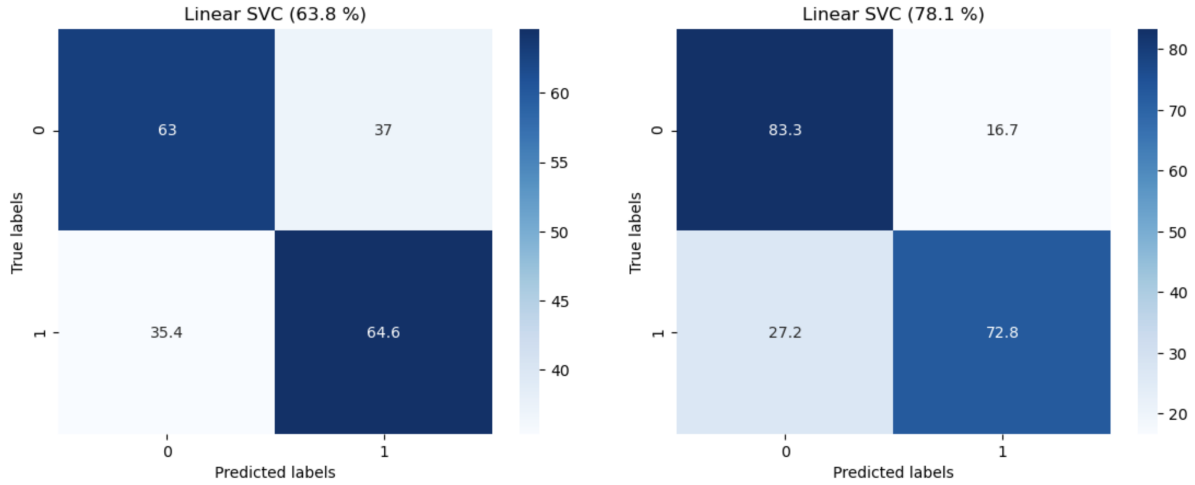


Fig. 8. Confusion matrix for the best model models for arousal (left) and valence (right). Linear SVC = Linear Support Vector Machine

We compare our results to the best results of our prior work that similarly performed global emotion recognition of positive vs negative valence from German-speaking, Swiss-based couples, albeit, with lab data. The best results (UAR) were 64.8% (female) and 56.1% (male) using fusion of acoustic and linguistic modality [12]. Our best valence results of 64.9% (female) and 78.1% (male) outperform that work, albeit, only slightly for female partners. Also, as a reference, the partner-perceived emotion results reported in Boateng et al. [15] that indicate how well partner A could tell the emotions of their partner B were 73.2% (for male partners) and 74.3% (for female partners). We did not collect such perception data from partners in this work. Hence, a direct comparison is not possible. Nonetheless, it is worth noting that our results for male partners slightly outperform female partners' perceptions of their male partners' emotions from that work.

There are significant privacy concerns with a system that recognizes the emotion of romantic couples especially considering such interactions have a high tendency to entail sensitive information. We argue that a smartwatch-based emotion recognition system such as ours has a better potential to be privacy-preserving. Compared to a facial-based emotion recognition system which can infer people's emotions without their consent (e.g., via CCTV camera), the smartwatch system can be designed to work only with the consent of the partner. For example, the device would need to be worn to be able to collect the relevant physiological and movement data needed for emotion recognition. The speech processing component would need to work only for the speech of the partner and hence would require a speech sample from that partner to work. Furthermore, the system using partner A's smartwatch could be designed to require a confirmation from partner B to allow their watch to share data (e.g., BLE signal or physiological data) which is needed to be able to recognize partner B's emotion from partner A's smartwatch. Hence, partner A could be prevented from recognizing the emotion of partner B without their consent. Furthermore, all processing of signals could be implemented to run on the device further restricting potentially sensitive data such as audio from ever leaving the smartwatch.

## 6 LIMITATION AND FUTURE WORK

The biggest limitation of this work is how highly skewed the data is especially for valence. The count for negative valence is 20 vs positive valence which is 360 with 5 out of the 13 couples having no negative valence labels.

There is a self-selection bias for this study which may have resulted in having couples that are less likely to have negatively rated interactions. Future studies could target couples in therapy who may have more negative conversation moments. Furthermore, collecting data for longer than 7 days may potentially capture conversation moments with negative emotion ratings.

Additionally, data was collected from only 13 couples (26 partners). Though small, for reference, two of the most popular public emotion datasets used in emotion recognition works – IEMOCAP [19] and MSP Improv [20] – contain data from 10 individuals (12 hours) and 12 individuals (9 hours) respectively, all collected from actors in the lab. Hence, our dataset has a greater variety with reference to subjects.

Future work would explore other fusion approaches such as decision-level fusion or some hybrid approach, multitask learning since the prediction entails two target variables – valence and arousal –, and pretraining on a related dataset and then fine-tuning on this dataset. Also, further work is needed to understand the conditions under which the model performs poorly e.g., indoors vs outdoors, when the partners are together alone or with other individuals. Such analysis could provide insight into potential changes that could improve the results (e.g., additional preprocessing of the dataset). It is also critical to better understand the conditions that could degrade performance before deploying for use in the real world.

Given that we had several audio samples without labels, we had three research assistants code all the audios with emotion labels so we could have more labeled data. Unfortunately, the inter-rater agreements were poor with an intraclass correlation coefficient of 0.21 – average for all the emotions coded. The poor agreement further demonstrates the difficulty even for humans in recognizing the emotions of romantic partners. Hence, we did not use those labels for our emotion recognition experiments. This agreement could potentially become better by improving the quality of the annotation instructions and having several rounds of annotation to ensure consistency in the annotation.

Our emotion recognition system used manual speaker annotations and transcription data. Hence, there are several steps needed in the future for this system to be usable in the real world such as implementing an automatic speaker diarization (detecting when each person spoke) and a speech recognition system. In particular, current speech recognition systems do not work for this unique dataset given that the couples spoke Swiss German, which is (1) a spoken dialect and not written, and (2) varies across different parts of the German-speaking regions of Switzerland. Hence, further work is needed to develop automatic speech recognition systems for Swiss German. Also, the machine learning system needs to be implemented on the smartwatch and evaluated in real-time in the real world. The pipeline of preprocessing, feature extraction, and machine learning classification would have to be implemented using libraries and frameworks that run on smartwatch platforms such as Google’s Wear OS and or Apple’s Watch OS. Then, the system would need to be validated in a field study to evaluate the algorithm in a new, unseen context.

## 7 CONCLUSION

In this work, we trained machine learning models to predict the emotions of romantic partners using multimodal smartwatch data collected from daily life. We used the following sensor data: heart rate, accelerometer, gyroscope, and ambient light. We performed binary classification of valence and arousal using linear SVM, RBF SVM, and random forest. We used individual modalities and explored various combinations of modalities using feature-level fusion. Our results from the best models – balanced accuracies of 63.8% and 78.1% for arousal and valence respectively – are better than chance and our prior work that also used data from German-speaking, Swiss-based couples, albeit, in the lab. This work contributes toward building automated emotion recognition systems that would eventually enable couples to monitor their emotions in daily life and enable the delivery of interventions to improve their emotional well-being. This approach could also be useful in couple therapy.

## ACKNOWLEDGMENTS

We are grateful to Prabhakaran Santhanam for assisting with the study's data collection and the following research assistants for coding, annotating and transcribing our dataset: Kwabena Atobra, Elena Luzi, Denis Adamec, and Luljeta Isaki. Funding was provided by the Swiss National Science Foundation (CR12I1\_166348/1).

## REFERENCES

- [1] [n.d.]. Open Sourcing German BERT. <https://deepset.ai/german-bert>. Accessed: 2020-05-1.
- [2] Mojtaba Khomami Abadi, Ramanathan Subramanian, Seyed Mostafa Kia, Paolo Avesani, Ioannis Patras, and Nicu Sebe. 2015. DECAF: MEG-based multimodal database for decoding affective physiological responses. *IEEE Transactions on Affective Computing* 6, 3 (2015), 209–222.
- [3] Tuka AlHanai and Mohammad Ghassemi. 2017. Predicting latent narrative mood using audio and physiologic data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 31.
- [4] Keith April Arano, Peter Gloor, Carlotta Orsenigo, and Carlo Vercellis. 2020. "Emotions are the Great Captains of our Lives": Measuring Moods through the Power of Physiological and Environmental Sensing. *IEEE Transactions on Affective Computing* (2020).
- [5] Hoda Badr and Linda K Acitelli. 2017. Re-thinking dyadic coping in the context of chronic illness. *Current Opinion in Psychology* 13 (2017), 44–48.
- [6] Alberto Betella and Paul FMJ Verschure. 2016. The affective slider: A digital self-assessment scale for the measurement of human emotions. *PLoS one* 11, 2 (2016), e0148037.
- [7] Jacopo Biggiogera, George Boateng, Peter Hilpert, Matthew Vowels, Guy Bodenmann, Mona Neysari, Fridtjof Nussbeck, and Tobias Kowatsch. 2021. *BERT Meets LIWC: Exploring State-of-the-Art Language Models for Predicting Communication Behavior in Couples' Conflict Interactions*. Association for Computing Machinery, New York, NY, USA, 385–389. <https://doi.org/10.1145/3461615.3485423>
- [8] Matthew Black, Athanasios Katsamanis, Chi-Chun Lee, Adam C Lammert, Brian R Baucom, Andrew Christensen, Panayiotis G Georgiou, and Shrikanth S Narayanan. 2010. Automatic classification of married couples' behavior using audio features. In *Eleventh annual conference of the international speech communication association*.
- [9] Matthew P Black, Panayiotis G Georgiou, Athanasios Katsamanis, Brian R Baucom, and Shrikanth Narayanan. 2011. "You made me do it": Classification of Blame in Married Couples' Interactions by Fusing Automatically Derived Speech and Language Information. In *Twelfth Annual Conference of the International Speech Communication Association*.
- [10] George Boateng. 2020. Towards Real-Time Multimodal Emotion Recognition among Couples. In *Proceedings of the 2020 International Conference on Multimodal Interaction*. 748–753.
- [11] George Boateng, Elgar Fleisch, and Tobias Kowatsch. 2022. Emotion Recognition among Couples: A Survey. *arXiv preprint arXiv:2202.08430* (2022).
- [12] George Boateng, Peter Hilpert, Guy Bodenmann, Mona Neysari, and Tobias Kowatsch. 2021. "You Made Me Feel This Way": Investigating Partners' Influence in Predicting Emotions in Couples' Conflict Interactions Using Speech Data. In *Companion Publication of the 2021 International Conference on Multimodal Interaction* (Montreal, QC, Canada) (*ICMI '21 Companion*). Association for Computing Machinery, New York, NY, USA, 390–394. <https://doi.org/10.1145/3461615.3485424>
- [13] George Boateng, Prabhakaran Santhanam, Elgar Fleisch, Janina Lüscher, Theresa Pauly, Urte Scholz, and Tobias Kowatsch. 2022. Development, Deployment, and Evaluation of DyMand—An Open-Source Smartwatch and Smartphone System for Capturing Couples' Dyadic Interactions in Chronic Disease Management in Daily Life. *arXiv preprint arXiv:2205.07671* (2022).
- [14] George Boateng, Prabhakaran Santhanam, Janina Lüscher, Urte Scholz, and Tobias Kowatsch. 2019. VADLite: an open-source lightweight system for real-time voice activity detection on smartwatches. In *Adjunct Proceedings of the 2019 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2019 ACM International Symposium on Wearable Computers*. 902–906.
- [15] George Boateng, Laura Sels, Peter Kuppens, Peter Hilpert, Urte Scholz, and Tobias Kowatsch. 2020. Speech Emotion Recognition among Couples using the Peak-End Rule and Transfer Learning. In *Companion Publication of the 2020 International Conference on Multimodal Interaction (ICMI '20 Companion)*, October 25–29, 2020, Virtual event, Netherlands.
- [16] Guy Bodenmann. 1997. Dyadic coping—a systematic-transactional view of stress and coping among couples: Theory and empirical findings. *European Review of Applied Psychology* 47 (1997), 137–140.
- [17] Leslie R Brody. 1993. On understanding gender differences in the expression of emotion. *Human feelings: Explorations in affect development and meaning* (1993), 87–121.
- [18] Pascal Budner, Joscha Eirich, and Peter A Gloor. 2017. "Making you happy makes me happy"—Measuring Individual Mood with Smartwatches. *arXiv preprint arXiv:1711.06134* (2017).
- [19] Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeannette N Chang, Sungbok Lee, and Shrikanth S Narayanan. 2008. IEMOCAP: Interactive emotional dyadic motion capture database. *Language resources and evaluation* 42, 4 (2008), 335.

- [20] Carlos Busso, Srinivas Parthasarathy, Alec Burmania, Mohammed AbdelWahab, Najmeh Sadoughi, and Emily Mower Provost. 2016. MSP-IMPROV: An acted corpus of dyadic interactions to study emotion perception. *IEEE Transactions on Affective Computing* 8, 1 (2016), 67–80.
- [21] Sandeep Nallan Chakravarthula, Brian Baucom, and Panayiotis Georgiou. 2018. Modeling Interpersonal Influence of Verbal Behavior in Couples Therapy Dyadic Interactions. *arXiv preprint arXiv:1805.09436* (2018).
- [22] Sandeep Nallan Chakravarthula, Haoqi Li, Shao-Yen Tseng, Maija Reblin, and Panayiotis Georgiou. 2019. Predicting Behavior in Cancer-Afflicted Patient and Spouse Interactions Using Speech and Language. *Proc. Interspeech 2019* (2019), 3073–3077.
- [23] Sandeep Nallan Chakravarthula, Md Nasir, Shao-Yen Tseng, Haoqi Li, Tae Jin Park, Brian Baucom, Craig J Bryan, Shrikanth Narayanan, and Panayiotis Georgiou. 2020. Automatic Prediction of Suicidal Risk in Military Couples Using Multimodal Interaction Cues from Couples Conversations. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 6539–6543.
- [24] James A Coan and John M Gottman. 2007. The specific affect coding system (SPAFF). *Handbook of emotion elicitation and assessment* (2007), 267–285.
- [25] Sidney K D’mello and Jacqueline Kory. 2015. A review and meta-analysis of multimodal affect detection systems. *ACM Computing Surveys (CSUR)* 47, 3 (2015), 1–36.
- [26] Paul Ekman and Dacher Keltner. 1997. Universal facial expressions of emotion. *Seegerstrale U, P. Molnar P, eds. Nonverbal communication: Where nature meets culture* (1997), 27–46.
- [27] Florian Eyben, Klaus R Scherer, Björn W Schuller, Johan Sundberg, Elisabeth André, Carlos Busso, Laurence Y Devillers, Julien Epps, Petri Laukka, Shrikanth S Narayanan, et al. 2015. The Geneva minimalistic acoustic parameter set (GeMAPS) for voice research and affective computing. *IEEE transactions on affective computing* 7, 2 (2015), 190–202.
- [28] Andreas Filler, Tobias Kowatsch, Severin Haug, Fabian Wahle, Thorsten Staake, and Elgar Fleisch. 2015. MobileCoach: A novel open source platform for the design of evidence-based, scalable and low-cost behavioral health interventions: overview and preliminary evaluation in the public health context. In *2015 Wireless Telecommunications Symposium (WTS)*. IEEE, 1–6.
- [29] Richard E Heyman. 2001. Observation of couple conflicts: Clinical assessment applications, stubborn truths, and shaky foundations. *Psychological assessment* 13, 1 (2001), 5.
- [30] Jian Huang, Jianhua Tao, Bin Liu, Zheng Lian, and Mingyue Niu. 2020. Multimodal Transformer Fusion for Continuous Emotion Recognition. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 3507–3511.
- [31] Eiman Kanjo, Eman MG Younis, and Nasser Sherkat. 2018. Towards unravelling the relationship between on-body, environmental and emotion data using sensor information fusion approach. *Information Fusion* 40 (2018), 18–31.
- [32] Patricia K Kerig and Donald H Baucom. 2004. *Couple observational coding systems*. Taylor & Francis.
- [33] Dana Ketcher, Casidee Thompson, Amy K Otto, Maija Reblin, Kristin G Cloyes, Margaret F Clayton, Brian RW Baucom, and Lee Ellington. 2021. The Me in We dyadic communication intervention is feasible and acceptable among advanced cancer patients and their family caregivers. *Palliative Medicine* 35, 2 (2021), 389–396.
- [34] Tobias Kowatsch, Dirk Volland, Iris Shih, Dominik Rüeegger, Florian Künzler, Filipe Barata, Andreas Filler, Dirk Büchter, Björn Brogde, Katrin Heldt, et al. 2017. Design and evaluation of a mobile chat app for the open source behavioral health intervention platform MobileCoach. In *International Conference on Design Science Research in Information System and Technology*. Springer, 485–489.
- [35] Janina Lüscher, Tobias Kowatsch, George Boateng, Prabhakaran Santhanam, Guy Bodenmann, and Urte Scholz. 2019. Social Support and Common Dyadic Coping in Couples’ Dyadic Management of Type II Diabetes: Protocol for an Ambulatory Assessment Application. *JMIR research protocols* 8, 10 (2019), e13685.
- [36] Matthias R Mehl, Megan L Robbins, and Fenne große Deters. 2012. Naturalistic observation of health-relevant social processes: The Electronically Activated Recorder (EAR) methodology in psychosomatics. *Psychosomatic Medicine* 74, 4 (2012), 410.
- [37] Angeliki Metallinou and Shrikanth Narayanan. 2013. Annotation and processing of continuous emotional attributes: Challenges and opportunities. In *2013 10th IEEE international conference and workshops on automatic face and gesture recognition (FG)*. IEEE, 1–8.
- [38] Sungkyu Park, Marios Constantinides, Luca Maria Aiello, Daniele Quercia, and Paul Van Gent. 2020. Wellbeat: A framework for tracking daily well-being using smartwatches. *IEEE Internet Computing* 24, 5 (2020), 10–17.
- [39] Soujanya Poria, Erik Cambria, Rajiv Bajpai, and Amir Hussain. 2017. A review of affective computing: From unimodal analysis to multimodal fusion. *Information Fusion* 37 (2017), 98–125.
- [40] Soujanya Poria, Navonil Majumder, Rada Mihalcea, and Eduard Hovy. 2019. Emotion recognition in conversation: Research challenges, datasets, and recent advances. *IEEE Access* 7 (2019), 100943–100953.
- [41] Juan Carlos Quiroz, Elena Geangu, and Min Hooi Yong. 2018. Emotion recognition using smart watch sensor data: Mixed-design study. *JMIR mental health* 5, 3 (2018), e10153.
- [42] Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084* (2019).
- [43] Tracey A Revenson and Anita DeLongis. 2011. Couples coping with chronic illness. (2011).

- [44] Tuula-Maria Rintala, Pia Jaatinen, Eija Paavilainen, and Päivi Åstedt-Kurki. 2013. Interrelation between adult persons with diabetes and their family: a systematic review of the literature. *Journal of family nursing* 19, 1 (2013), 3–28.
- [45] Megan L Robbins, Ana María López, Karen L Weihs, and Matthias R Mehl. 2014. Cancer conversations in context: naturalistic observation of couples coping with breast cancer. *Journal of Family Psychology* 28, 3 (2014), 380.
- [46] Nicole A Roberts, Jeanne L Tsai, and James A Coan. 2007. Emotion elicitation using dyadic interaction tasks. *Handbook of emotion elicitation and assessment* (2007), 106–123.
- [47] James A Russell. 1980. A circumplex model of affect. *Journal of personality and social psychology* 39, 6 (1980), 1161.
- [48] Philip Schmidt, Attila Reiss, Robert Dürichen, and Kristof Van Laerhoven. 2019. Wearable-Based Affect Recognition—A Review. *Sensors* 19, 19 (2019), 4079.
- [49] Dominik Schoebi. 2008. The coregulation of daily affect in marital relationships. *Journal of Family Psychology* 22, 4 (2008), 595.
- [50] Björn W Schuller. 2018. Speech emotion recognition: Two decades in a nutshell, benchmarks, and ongoing trends. *Commun. ACM* 61, 5 (2018), 90–99.
- [51] Amber J Seidel, Melissa M Franks, Mary Ann Parris Stephens, and Karen S Rook. 2012. Spouse control and type 2 diabetes management: moderating effects of dyadic expectations for spouse involvement. *Family relations* 61, 4 (2012), 698–709.
- [52] Salvatore Settineri, Amelia Rizzo, Marco Liotta, and Carmela Mento. 2014. Caregiver’s burden and quality of life: Caring for physical and mental illness. *International Journal of Psychological Research* 7, 1 (2014), 30–39.
- [53] Adela C Timmons, Theodora Chaspari, Sohyun C Han, Laura Perrone, Shrikanth S Narayanan, and Gayla Margolin. 2017. Using multimodal wearable technology to detect conflict among couples. *Computer* 50, 3 (2017), 50–59.
- [54] Shao-Yen Tseng, Haoqi Li, Brian Baucom, and Panayiotis Georgiou. 2018. "Honey, I Learned to Talk" Multimodal Fusion for Behavior Analysis. In *Proceedings of the 20th ACM International Conference on Multimodal Interaction*. 239–243.
- [55] David Watson, Lee Anna Clark, and Auke Tellegen. 1988. Development and validation of brief measures of positive and negative affect: the PANAS scales. *Journal of personality and social psychology* 54, 6 (1988), 1063.
- [56] Martin Wöllmer, Felix Weninger, Tobias Knaup, Björn Schuller, Congkai Sun, Kenji Sagae, and Louis-Philippe Morency. 2013. Youtube movie reviews: Sentiment analysis in an audio-visual context. *IEEE Intelligent Systems* 28, 3 (2013), 46–53.
- [57] Liang-Chih Yu, Lung-Hao Lee, Shuai Hao, Jin Wang, Yunchao He, Jun Hu, K Robert Lai, and Xuejie Zhang. 2016. Building Chinese affective resources in valence-arousal dimensions. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 540–545.
- [58] Sebastian Zepf, Javier Hernandez, Alexander Schmitt, Wolfgang Minker, and Rosalind W Picard. 2020. Driver emotion recognition for intelligent vehicles: a survey. *ACM Computing Surveys (CSUR)* 53, 3 (2020), 1–30.