Reinforcement Learning for Channel Coding

International Zurich Seminar on Information and Communication (IZS), February 26 – 28, 2020

Reinforcement Learning for Channel Coding

(Extended Abstract)

Mengke Lian*, Fabrizio Carpi†, Christian Häger‡, and Henry D. Pfister*
*Department of Electrical and Computer Engineering, Duke University, Durham, North Carolina
†Department of Electrical and Computer Engineering, New York University, Brooklyn, New York, USA
‡Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden

***Abstract*—We discuss the application of reinforcement learning (RL) to problems associated with decoding binary linear codes. In particular, we consider two different iterative decoding algorithms that involve sequential decisions and apply RL to optimize those decisions. For the first example, we consider bit-flipping (BF) decoders and observe that learned BF decoders can offer a range of performance–complexity trade-offs and achieve near-optimal performance in some cases. For the second example, we consider protograph low-density parity-check (LDPC) codes and use RL to optimize the decoding schedule. Our results show that, the comparison with the flooding schedule, a fixed error rate can be achieved with fewer update operations.**

I. Overview

The decoding of error-correcting codes can be seen as a classification problem and solved using methods introduced for supervised machine learning. The general idea is to treat the decoder as a parameterized function (e.g., a neural network) and learn good parameter configurations with data-driven optimization [1]–[4]. Without further restrictions on the code and decoder, these methods do not work well if the codes have moderate length or if they are unstructured. For linear codes, the problem simplifies considerably because one has to learn only a single decision region instead of one region per codeword. One can take advantage of linearity by using message-passing [2] or syndromes [3], [4]. Still, the problem remains challenging because good codes typically have complicated decision regions due to the large number of neighboring codewords.

This talk focuses on applications of reinforcement learning (RL) [5] to problems in channel coding. Despite impressive results in other fields, RL has yet to received significant attention in this area. In fact, with the exception [10] and recent work by the authors [6], no references were found that discuss RL for channel coding. For a general survey of RL in the general context of communications, see [11].

The unifying idea behind this work is that iterative decoding algorithms can adjust their behavior based on their current state. Thus, they can be modeled as sequential decision processes and RL can be used to optimize their behavior.

The bit-flipping (BF) decoder was introduced in [7], [8] and has been studied extensively in the literature. In [6], a subset of the authors apply RL to optimize a sequential BF decoder where, based on the syndrome, one bit is flipped in each step [6]. Rather than learning a direct mapping from observations to estimated codewords (or bits) in a supervised fashion, decoding is done in steps and the problem is mapped to a Markov decision process (MDP). Then, RL is applied to optimize the choice of which bit to flip. Following [3], [4], this approach is syndrome-based and the state space of the MDP is formed by all possible binary syndromes. This also decouples the decoding problem from the transmitted codeword. We also consider a parallel BF algorithm where the decision to flip a bit is based on the local neighborhood of that bit.

For the second example, we consider the belief-propagation decoding of protograph low-density parity-check (LDPC) codes. For these codes, the standard decoding schedule can be inefficient. In particular, windowed decoding is known to improve the performance–complexity trade-off of spatially-coupled codes [9]. Thus, we employ RL to optimize the decoding schedule and show that the optimized schedule can achieve good performance with significantly fewer operations.

In summary, we believe that RL is a promising technique for optimizing sequential decisions in decoding algorithms.

References

[1] T. Gruber, S. Cammerer, J. Hoydis, and S. ten Brink, "On deep learning-based channel decoding," in *Proc. Annual Conf. Information Sciences and Systems (CISS)*, Baltimore, MD, 2017.
[2] E. Nachmani, E. Marciano, L. Lugosch, W. J. Gross, D. Burshtein, and Y. Be'ery, "Deep learning methods for improved decoding of linear codes," *IEEE J. Select. Topics in Signal Processing*, vol. 12, no. 1, pp. 119–131, 2018.
[3] L. G. Tallini and P. Cull, "Neural nets for decoding error-correcting codes," in *Proc. IEEE Technical Applications Conf. and Workshops*, Portland, USA, 1995.
[4] A. Bennatan, Y. Choukroun, and P. Kisilev, "Deep learning for decoding of linear codes - a syndrome-based approach," in *Proc. IEEE Int. Symp. Information Theory (ISIT)*, Vail, CO, 2018.
[5] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. A Bradford Book, 1998.
[6] F. Carpi, C. Häger, M. Martalò, R. Raheli, and H. D. Pfister, "Reinforcement learning for channel coding: Learned bit-flipping decoding," in *Proc. Annual Allerton Conf. on Commun., Control, and Comp.*, 2019.
[7] M. Bossert and F. Hergert, "Hard- and soft-decision decoding beyond the half minimum distance—an algorithm for linear codes (corresp.)," *IEEE Trans. Inf. Theory*, vol. 32, no. 5, pp. 709–714, Sept. 1986.
[8] Y. Kou, S. Lin, and M. Fossorier, "Low-density parity-check codes based on finite geometries: a rediscovery and new results," *IEEE Trans. Inf. Theory*, vol. 47, no. 7, pp. 2711–2736, Nov. 2001.
[9] A. R. Iyengar, M. Papaleo, P. H. Siegel, J. K. Wolf, A. Vanelli-Coralli, and G. E. Corazza, "Windowed decoding of protograph-based LDPC convolutional codes over erasure channels," *IEEE Trans. Inform. Theory*, vol. 58, no. 4, pp. 2303–2320, 2011.
[10] X. Wang, H. Zhang, R. Li, L. Huang, S. Dai, Y. Huangfu, and J. Wang, "Learning to flip successive cancellation decoding of polar codes with LSTM networks," *arXiv:1902.08394*, Feb. 2019.
[11] N. C. Luong, D. T. Hoang, S. Gong, D. Niyato, P. Wang, Y.-C. Liang, and D. I. Kim, "Applications of deep reinforcement learning in communications and networking: A survey," *arxiv:1810.07862*, 2018.