

Understanding cirrus clouds using explainable machine learning

Journal Article**Author(s):**

Jeggle, Kai; [Neubauer, David](#) ; Camps-Valls, Gustau; [Lohmann, Ulrike](#) 

Publication date:

2023

Permanent link:

<https://doi.org/10.3929/ethz-b-000624717>

Rights / license:

[Creative Commons Attribution 4.0 International](#)

Originally published in:


Environmental Data Science 2, <https://doi.org/10.1017/eds.2023.14>

Funding acknowledgement:

860100 - innovative Machine learning to constrain Aerosol-cloud CLimate Impacts (EC)

APPLICATION PAPER  

Understanding cirrus clouds using explainable machine learning

Kai Jeggle¹ , David Neubauer¹, Gustau Camps-Valls² and Ulrike Lohmann¹

¹Institute for Atmospheric and Climate Science, ETH Zurich, Zurich, Switzerland

²Image Processing Laboratory, Universitat de València, València, Spain

Corresponding author: Kai Jeggle; Email: kai.jeggle@env.ethz.ch

Received: 26 May 2023; **Accepted:** 30 May 2023

Keywords: cirrus clouds; eXplainable AI; explainable machine learning; machine learning; SHAP; Shapley values

Abstract



Cirrus clouds are key modulators of Earth's climate. Their dependencies on meteorological and aerosol conditions are among the largest uncertainties in global climate models. This work uses 3 years of satellite and reanalysis data to study the link between cirrus drivers and cloud properties. We use a gradient-boosted machine learning model and a long short-term memory network with an attention layer to predict the ice water content and ice crystal number concentration. The models show that meteorological and aerosol conditions can predict cirrus properties with $R^2 = 0.49$. Feature attributions are calculated with SHapley Additive exPlanations to quantify the link between meteorological and aerosol conditions and cirrus properties. For instance, the minimum concentration of supermicron-sized dust particles required to cause a decrease in ice crystal number concentration predictions is $2 \times 10^{-4} \text{ mg/m}^3$. The last 15 hr before the observation predict all cirrus properties.

Impact Statement

We applied explainable machine learning to predict cirrus clouds and improve process understanding. The results can validate and improve cirrus cloud parameterizations in global climate models, essential for climate change projections and support assessing cirrus-based climate engineering methods.

1. Introduction

Cirrus clouds consist of ice crystals and occur at temperatures below -38°C in the upper troposphere (Sassen et al., 2008). Like all other clouds, they influence the Earth's radiative budget via the reflection of solar radiation and absorption and emission of terrestrial infrared radiation into space (Liou, 1986). Depending on the cloud microphysical properties (CMPs), namely ice water content IWC and ice crystal number concentration N_i , the cloud radiative effect of cirrus varies significantly and ranges from a cooling effect at lower altitudes for optically thick cirrus to a warming effect at high altitudes for optically thin cirrus (Heymsfield et al., 2017). However, due to the high spatiotemporal variability of cirrus, different formation pathways, and the nonlinear dependence on environmental conditions and aerosols,

  This research article was awarded Open Data and Open Materials badges for transparent practices. See the Data Availability Statement for details.

© The Author(s), 2023. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

constraining CMP and hence the cloud radiative effect remains an active field of study (Gasparini et al., 2018; Gryspeerdt et al., 2018; Krämer et al., 2020). An improved process understanding of cirrus CMP will help to reduce uncertainties in global climate models and climate change projections and assess climate intervention methods targeted at cirrus clouds. This study aims to quantify the impacts of individual cirrus drivers on the CMP. We approach this task with a two-stage process. First, we train a black-box machine learning (ML) model to predict IWC and N_i , given the input of meteorological variables and aerosol concentrations. Next, we calculate the individual contributions to the predictions for each input feature using post hoc feature attribution methods from the emerging field of eXplainable AI (XAI) (Belle and Papantonis, 2020). Using XAI methods to understand and discover new physical laws via ML has been discussed in recent publications and is increasingly adopted in the atmospheric science community (McGovern et al., 2019; Ebert-Uphoff and Hilburn, 2020). We train and explain a gradient-boosted regression tree on instantaneous data and a long short-term memory (LSTM) network (Hochreiter and Schmidhuber, 1997) with an attention layer on time-resolved data. The former is aimed to yield better interpretability, while the latter incorporates temporal information of the predictors. The remainder of this paper is structured as follows: A brief background about cirrus drivers is presented in Section 2, followed by a description of the data set, ML models, and feature attribution methods in Section 3. Section 4 presents the prediction performance of the ML models and the corresponding feature attributions. Finally, the results are summarized and discussed in Section 5.

2. Cirrus Driver

Previous studies have found that cirrus CMP are mainly controlled by the ice crystal nucleation mode (Kärcher et al., 2006), ice origin (Krämer et al., 2020), meteorological variables such as temperature and updraft speeds, and aerosol environment (Gryspeerdt et al., 2018). There are two possible nucleation modes—homogeneous and heterogeneous nucleation. Homogeneous nucleation is the freezing of supercooled solution droplets below -38°C and high ice supersaturations (Koop et al., 2000). It strongly depends on updraft velocity, producing higher N_i with increasing updrafts (Jensen et al., 2016). Heterogeneous nucleation occurs via crystalline or solid aerosols acting as ice-nucleating particles that can significantly lower the energy barrier needed for ice nucleation (Kanji et al., 2017). Dust is considered the most important ice-nucleating particle (Kanji et al., 2017) and is hence often the only ice-nucleating particle considered in global climate models. The availability of ice-nucleating particles can lead to fewer but larger crystals by suppressing homogeneous nucleation (Kuebbeler et al., 2014).

Wernli et al. (2016) and Krämer et al. (2020) showed that the history of the air parcel leading up to a cloud significantly influences cirrus CMP. Two main categories with respect to the ice origin were identified: Liquid-origin cirrus form by freezing cloud droplets, that is, originating from mixed-phase clouds at temperatures above -38°C and leading to higher IWC, in situ cirrus form ice directly from the gas phase or freezing of solution droplets at temperatures below -38°C and lead to thinner cirrus with lower IWC. Besides the ice origin, the air parcel's exposure to changing meteorological and aerosol conditions upstream of a cloud observation can impact the cirrus CMP. For instance, a high vertical velocity occurring prior to a cloud observation may increase IWC and N_i (Heymsfield et al., 2017). In this study, we focus on the meteorological variables and aerosol environment as cirrus drivers. Nucleation mode and ice origin are classifications that are themselves controlled by the cirrus drivers.

3. Methods and Data

3.1. Data

This study combines satellite observations of cirrus clouds and reanalysis data of meteorology and aerosols to predict IWC and N_i . We use the DARDAR-Nice (Sourdeval et al., 2018) product for satellite-based cirrus retrievals. The DARDAR-Nice product combines measurements from CloudSat's radar and CALIOP's lidar to obtain vertically resolved retrievals for IWC and N_i for ice crystals $>5\ \mu\text{m}$ in cirrus

clouds. DARDAR-Nice also provides information about each layer's distance to the cloud top. Additionally, we calculate a cloud thickness feature, which provides information about the vertical extent of a cloud. To represent cirrus cloud drivers, we use the ERA5 (Hersbach et al., 2018) and MERRA2 (Global Modeling and Assimilation Office, 2015) reanalysis data sets for meteorological and aerosol drivers, respectively. All data sources are collocated on a $0.25^\circ \times 0.25^\circ \times 300$ m spatial and hourly temporal resolution corresponding to the original ERA5 resolution. Due to the narrow swath of the satellite overpass (i.e., low spatial coverage) and long revisiting time (16 days), studies about the temporal development of cirrus clouds and their temporal dependencies on driver variables are not possible given only instantaneous DARDAR-Nice data collocated with driver variables.

CMPs are not only determined by their instantaneous environmental variables but undergo a temporal development in which they are exposed to changing atmospheric states such as shifts in temperature, updraft velocities, and aerosol environment. Thus, we add a temporal dimension to our data set by calculating 48 hr of Lagrangian backtrajectories for every cirrus cloud observation in the DARDAR-Nice data set with LAGRANTO (Sprenger and Wernli, 2015). An individual trajectory is calculated for every cirrus layer in an atmospheric column. For instance, for a 900 m thick cloud consisting of three 300 m layers, three individual trajectories are calculated and initialized simultaneously and identical latitude/longitude but at different height levels. LAGRANTO uses ERA5 wind fields for the trajectory calculation. All meteorological and aerosol variables are traced along the trajectory. These variables represent the changing meteorology and aerosol environment along the air parcel upstream of a cirrus observation in the satellite data. Figure 1 exemplarily displays the satellite overpass (gray), cirrus cloud observations (red), and the corresponding backtrajectories (blue).

To account for seasonal and regional covariations, we include categorical variables representing the different seasons (*DJF*, *MMA*, *JJA*, *SON*) and regions divided in 10° latitude bins into the data set. Additionally, the surface elevation of the ground below an air parcel is included to incorporate the influence of orographic lifting. Since an air parcel can travel through multiple 10° latitude bins in 48 hr and is exposed to varying topographies, both latitude bin and surface elevation are also traced along the trajectory. Finally, we use the land water mask provided by the DARDAR-Nice dataset. Please refer to Sourdeval et al. (2018) for more information. All used data variables are specified in Table 1. The last column in the table indicates whether the variable is available along the backtrajectories or just at the time of the satellite observation. We chose to study a domain ranging from 140°W to 40°E and 25°N to 80°N .

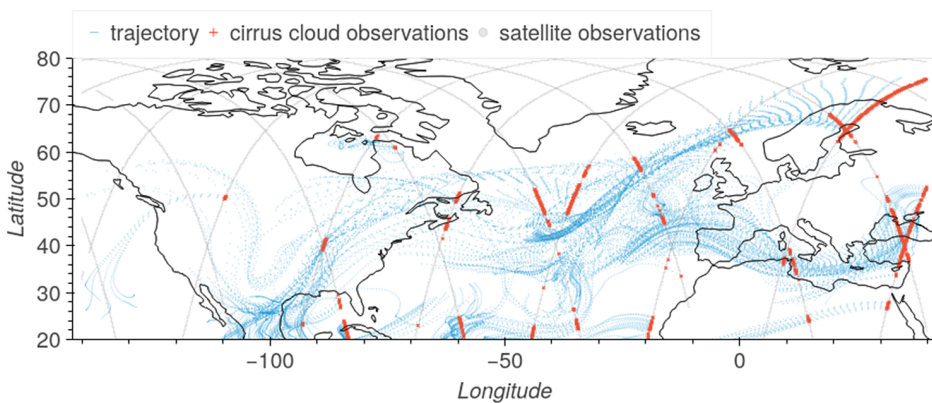


Figure 1. Illustrative example of data for 1 hr. DARDAR-Nice satellite observations along the satellite overpass are displayed as gray dots. A vertical profile of 56 300 m thick vertical layers is available for each satellite observation. To keep the diagram uncluttered, only one vertical level is shown. Red crosses mark observations containing cirrus clouds. Blue dashes represent the corresponding 48 hr of backtrajectories calculated for each cirrus observation. Along each trajectory, meteorological and aerosol variables are traced.

Table 1. Summary of data used in this study.

	Variable	Source	Backtrajectory
Target variables	Ice crystal number concentration (N_i) (mg/m ³)	DARDAR-Nice	✗
	Ice water content (IWC) (cm ⁻³)		
Predictor variables	Temperature (K)	ERA5	✓
	Vertical velocity (Pa/s)		
	Horizontal wind speed (m/s)		
	Dust aerosols ^a (DU _{sub} , DU _{sup}) (mg/kg)	MERRA2	✓
	Sulfate aerosols (SO ₄) (mg/kg)		
	Distance from cloud top (m)	DARDAR-Nice	✗
	Cloud thickness (m)		
	Land water mask		
	Surface elevation (m)		✓
	Region (10° latitude bins)	ERA5	✓
	Season		✗

Note. Lagrangian backtrajectories were calculated 48 hr back in time for each cirrus observation. Variables with ✓ in the last column are available along backtrajectories (hourly data). Variables with ✗ are only available at the satellite overpass.

^aDust concentrations are divided in particles smaller than 1 μm (DU_{sub}) and particles larger than 1 μm (DU_{sup}).

This study area covers a wide range of climate zones and topographies. In the vertical direction, we consider 56 vertical layers of 300 m thickness covering heights between 4,080 and 20,580 m. Since we only consider cirrus clouds in this study, all data above -38°C are omitted. The tropics are excluded from this study as convective regimes often drive liquid-origin cirrus clouds in this region, and convection is not well represented at this horizontal resolution. Due to gaps in the satellite data, we focus on the period between 2007 and 2009. Since sunlight can adversely affect the quality of satellite retrievals, we only consider nighttime observations.

3.2. Prediction models

We aim to predict IWC and N_i of cirrus clouds with meteorological and aerosol variables, a regression task from an ML perspective. By adding latitude bin, season, and surface elevation to the input features, we account for regional, seasonal, and topographic variability. We approach the regression task with (a) instantaneous and (b) time-resolved input data. The former has the advantage of training the ML model on a tabular data set that best suits the explainability methods described in the next section. The latter data set enables the model to incorporate the temporal dependencies of cirrus drivers on the cirrus CMP but is more difficult to interpret. Note that only the input data has a temporal dimension, namely the backtrajectories.

3.2.1. Instantaneous prediction model

We train a gradient-boosted regression tree using the XGBoost algorithm (Chen and Guestrin, 2016), which typically outperforms more complex deep-learning models on tabular data (Shwartz-Ziv and Armon, 2021). A separate XGBoost model is trained for each target variable.

3.2.2. Time-resolved prediction model

Given the temporally resolved backtrajectory data of cirrus drivers x_t with $t \in [-48, 0]$ (sequential features) and static features x_{static} , like cloud thickness and distance to cloud top, we predict the target variables IWC and N_i at $t = 0$, that is, the time of the satellite observation, using a single multi-output LSTM architecture (Hochreiter and Schmidhuber, 1997). From an ML perspective, this is a many-to-one

regression problem. LSTMs are neural networks that capitalize on temporal dependencies in the input data. The LSTM encodes the sequential features into a latent representation called hidden state h_t (1).

$$h_t = \text{LSTM}(x_t), t \in [-48, 0], \tag{1}$$

$$u_t = \tanh(W h_t + b), \tag{2}$$

$$\alpha_t = \frac{\exp(u_t u)}{\sum_t \exp(u_t u)}, \tag{3}$$

$$z = \sum_t \alpha_t h_t. \tag{4}$$

Then, the output of the LSTM is fed to a simple attention mechanism that determines the relative importance of each timestep in the sequential input data for the final prediction. The attention mechanism provides a built-in approach to interpreting the ML model’s prediction for the temporal dimension. It comprises a single fully connected layer and a context vector u . First, a latent representation u_t of h_t is created by passing h_t through the attention layer (2), then the importance of each timestep α_t is calculated by multiplying u_t with the context vector u . We normalize the importance weights α_t by applying the softmax function (3) so that all attention weights sum up to 1. Next, the weighted sum z of the importance weights α_t and the LSTM hidden state h_t is calculated (4). The attention layer, as well as the context vector u , are randomly initialized and learned during the training through backpropagation together with the weights of the LSTM and other fully connected layers. To make the final prediction $y \in \mathbb{R}_+^2$, z is concatenated with the static input features x_{static} and fed through a set of fully connected layers. Figure 2 visualizes the whole architecture consisting of LSTM, attention layer, fully connected layers, and finally, a linear layer.

3.2.3. Experimental setup

To train and evaluate the ML models, we split the data into 80% training, 10% validation, and 10% test data. Each cloud layer containing cirrus clouds is considered an individual sample, that is, a cloud consisting of 3300 m layers is considered as three different samples. To prevent overfitting due to spatiotemporally correlated samples in training and test sets, samples from the same month are put in the same split. In total, the data set consists of 6 million samples. Due to the large spread in scales, the target variables and aerosol concentrations are logarithmically transformed. The categorical features are one-hot

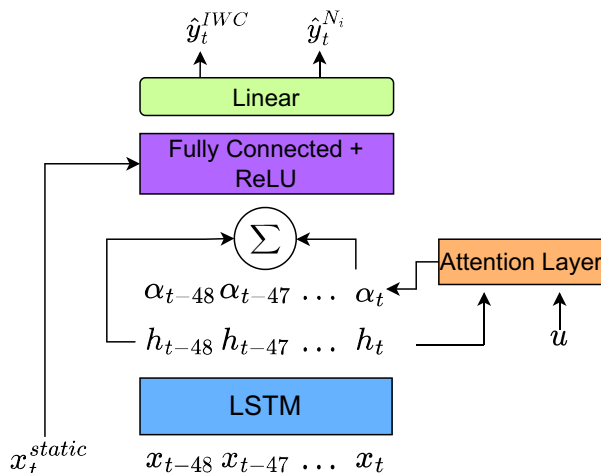


Figure 2. Architecture of the LSTM-based model with integrated attention method to predict IWC and N_i of cirrus clouds with the temporal data set. ReLU, rectified linear unit.

encoded. For the XGBoost model, no further pre-processing is required. For the LSTM model, we standardize the continuous sequential and static features by removing the mean and dividing them by the standard deviation of the training samples. The hyperparameters of the models are tuned using Bayesian optimization. Fifty hyperparameter configurations are tested per model during the optimization. The final hyperparameters are displayed in [Supplementary Appendix A](#).

3.3. Explainable AI

We apply post hoc feature attribution methods to the trained models. For each sample, the marginal contribution of each feature toward the prediction is calculated. By combining explanations for all samples, we can gain insight into the internal mechanics of the model and, eventually, an indication of the underlying physical processes. We focus on the instantaneous predictive model and only briefly interpret the temporal model by analyzing the attention weights. We use SHapley Additive exPlanations (SHAP) (Lundberg and Lee, 2017) using the TreeSHAP (Lundberg et al., 2020) implementation to calculate feature attributions. SHAP is a widely applied XAI method based on game-theoretic Shapley values. For each feature, SHAP outputs the change in expected model prediction, that is, if a feature contributes to an increase or decrease in the prediction.

To evaluate the quality of black-box ML model explanations, it is customary to calculate faithfulness and stability metrics. The notion of faithfulness states how well a given explanation represents the actual behavior of the ML model (Alvarez Melis and Jaakkola, 2018), and stability measures how robust an explanation is toward small changes in the input data (Alvarez Melis and Jaakkola, 2018; Agarwal et al., 2022). In this study, we use *Estimated Faithfulness* (Alvarez Melis and Jaakkola, 2018), which measures the correlation between the relative importance assigned to a feature by the explanation method and the effect of each feature on the performance of the prediction model. Higher importance should have a higher effect and vice versa. Explanation stability is evaluated using relative input stability (RIS) and relative output stability (ROS) (Agarwal et al., 2022). Implementation details of the XAI evaluation metrics are described in [Supplementary Appendix B](#). To further validate our explanations, we compare them with LIME (Ribeiro et al., 2016) and baseline with randomly generated attributions. LIME is another popular XAI method that calculates feature attributions by fitting linear models in the neighborhood of a single sample.

4. Results

We first present the predictive performance of the XGBoost and LSTM-based models. The feature attributions of the XGBoost model are discussed. Finally, we analyze the weights of the attention layer in the LSTM-based model to understand which timesteps are important for the prediction tasks.

4.1. Cirrus cloud prediction performance

The prediction performance of the ML models and a linear baseline regression are displayed in [Table 2](#). Both ML models outperform the baseline, and the LSTM-based model performs best overall with 0.49 (0.40) R^2 , and root mean squared error of 0.35 (0.33) for IWC (N_i). This shows that the backtrajectories contain useful information for cirrus cloud prediction. Like the LSTM-based model, the XGBoost model can capture nonlinear relationships in the data. Thus, it is fair to continue focusing on that model for model explainability.

4.2. Feature attributions with SHAP

[Figure 3](#) shows that SHAP explanations are more stable and have higher faithfulness than the other feature attribution methods and are thus best suited for our study. For each feature, the mean absolute SHAP value aggregated over all test samples is displayed in [Figure 4](#) for IWC (left) and N_i (right). It can be seen that the vertical extent of the cloud and the position of the cloud layer within the cloud are important predictors for

Table 2. Performance metrics evaluated on independent test data for the regression task of predicting cirrus CMP using meteorological and aerosol variables.

Model	Input data	IWC		N_i	
		RMSE	R^2	RMSE	R^2
Linear regression (baseline)	Instantaneous	0.44	0.24	0.45	0.19
XGBoost	Instantaneous	0.40±0.01	0.38±0.02	0.42±0.01	0.35±0.02
LSTM	Temporal	0.35±0.03	0.49±0.04	0.33±0.02	0.42±0.03

Note. Each model was trained 10 times; the values in the table represent the mean and standard deviation over the 10 model runs. Performance metrics are calculated on log-transformed values of target variables. Abbreviations: CMP, cloud microphysical property; IWC, ice water content; LSTM, long short-term memory; R^2 , coefficient of determination; RMSE, root mean squared error.

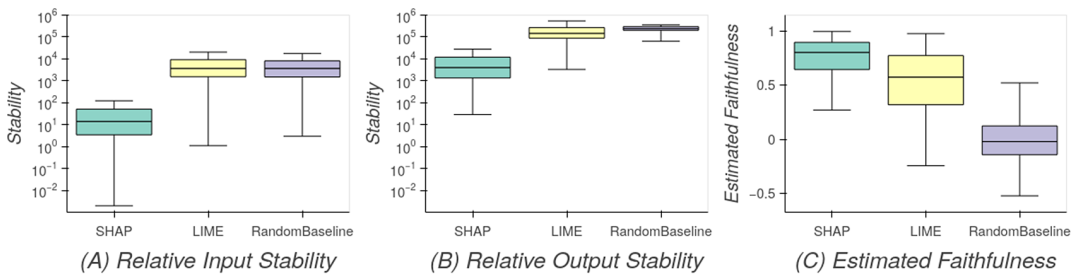


Figure 3. XAI evaluation metrics for SHAP, LIME, and a random baseline feature attribution of the XGBoost model. For the stability metrics RIS (a) and ROS (b), lower values indicate more stable explanations, that is, more robust toward small changes of the feature values. Estimated faithfulness (c) indicates whether features with high importance attributed by the feature attribution method are important for the prediction performance, where 1 denotes perfect estimated faithfulness.

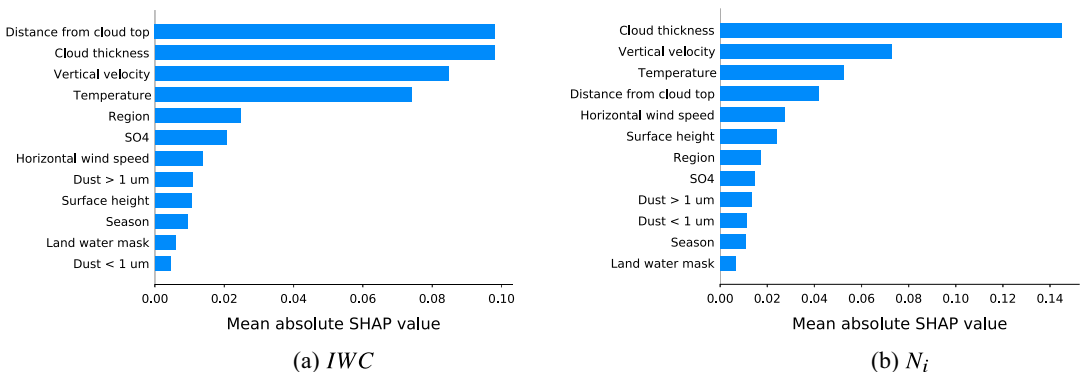


Figure 4. Mean absolute SHAP values of each feature for IWC (a) and N_i (b). The higher the value, the higher the contribution to the prediction, that is, the more important the feature.

both CMP. As expected from theory, temperature, and vertical velocity are the main drivers for the prediction. Aerosol concentrations play only a minor role in the prediction.

Figure 5 allows a more detailed analysis of the SHAP values. Here, we aggregated the SHAP values for each feature along the range of occurring feature values. Increasing updrafts (i.e., negative vertical

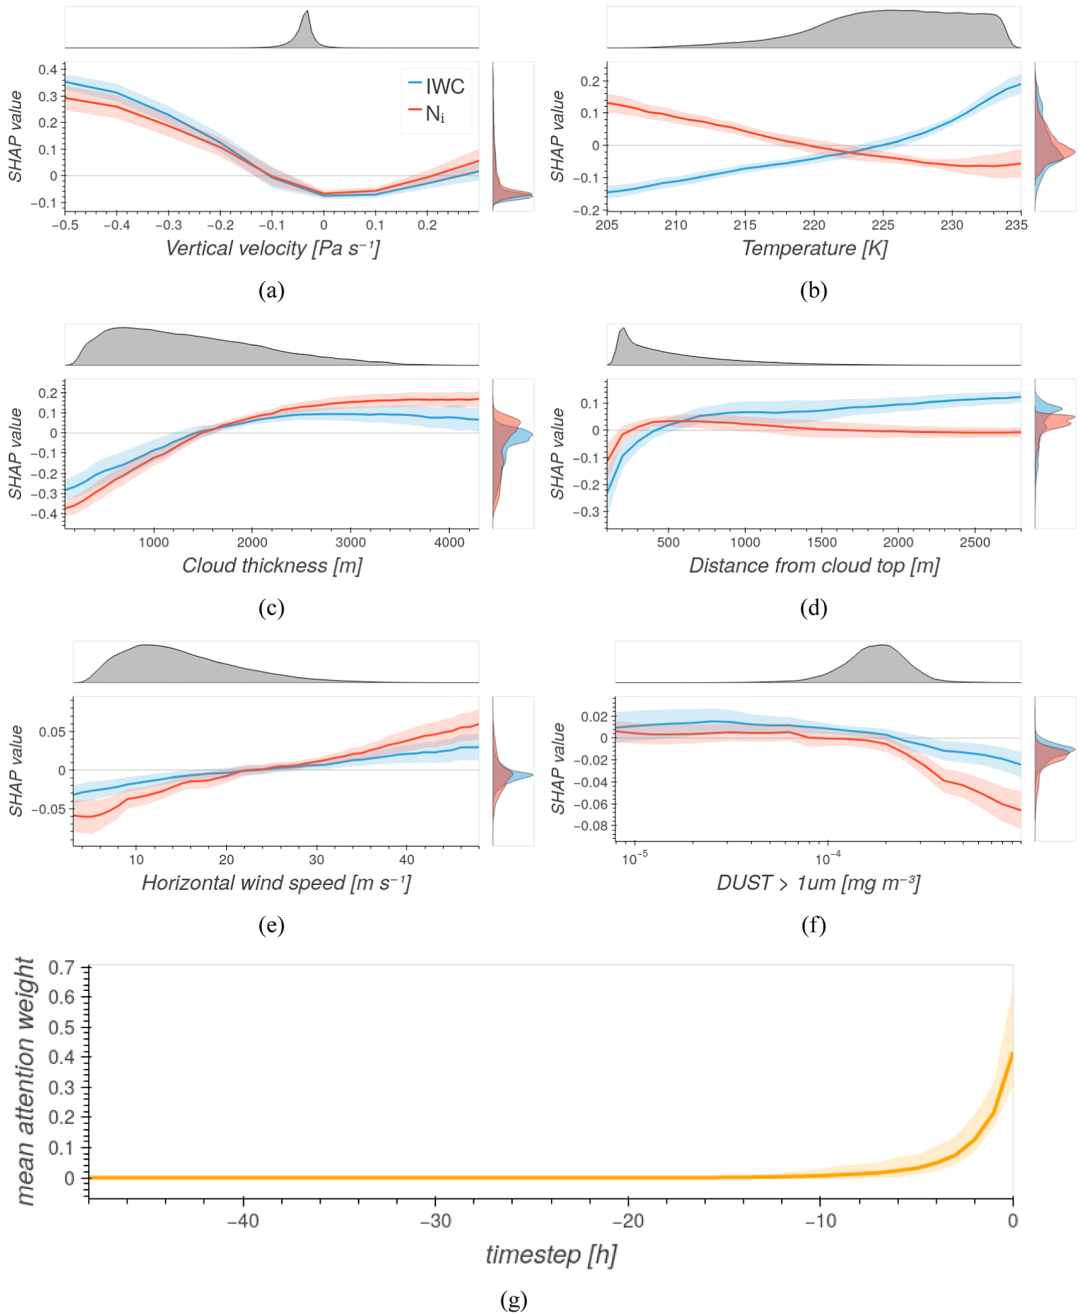


Figure 5. (a–f) Partial SHAP dependence plots for representative features showing the mean SHAP value per feature value for IWC in blue color and N_i in red color as solid lines. The shaded area represents the standard deviation. Additionally, the marginal distribution of the feature is displayed in gray above each plot, and the marginal SHAP value distributions are displayed on the right of each plot. Note that the absolute SHAP values of IWC and N_i are not directly comparable as the SHAP value indicates the contribution of a feature to the prediction value, and the two variables have different yet similar distributions. The plots show only feature values with at least 5,000 occurrences in the test data set. (g) Mean of attention weight per timestep. Attention weights are learned during the training process of the LSTM model and represent the relative importance of the input data per timestep.

velocities (a)) are increasing the model predictions for IWC and N_i . Horizontal wind speed (e) has a near-linear contribution to the predictions with a positive impact for wind speeds exceeding 25/m/s. Horizontal wind speed is not considered much in cirrus driver literature but could be a good indicator for dynamical systems such as warm conveyor belts. For temperature (b), a positive dependency is found for IWC and a negative dependency for N_i . At temperatures between 230 and 235 K, there is still a strong positive dependency for IWC, while the curve for N_i is flattening out. This could indicate that in this temperature range, keeping all other conditions the same, the main processes are growing or colliding ice crystals instead of the formation of new ice crystals. Looking at cloud thickness, it can be observed that vertically thicker clouds increase CMP regardless of where the cloud layer is positioned in the cloud. A strong decrease in CMP is visible for cloud layers at the top of the cloud, with a more pronounced effect for IWC (d). This suggests that the entrainment of dry air at the cloud top is a dominant process for cirrus, causing the ice crystals to shrink and partially sublimate. Dust $>1 \mu\text{m}$ slightly decreases the predictions when surpassing a concentration of $2 \times 10^{-4} \text{ mg/m}^3$, with a stronger effect for N_i . The physical effect acting here could be the suppression of homogeneous nucleation.

4.3. Interpretation of attention weights

Figure 5g displays the mean attention weight per timestep of the attention layer of the LSTM model. Compared to the post hoc feature attribution with SHAP, the attention weights are learned during the training process in the attention layer. The weights indicate how important the model considers a given timestep for the prediction. Results suggest that all necessary information for cirrus prediction is encapsulated in the last 15 hr before the observations, with more information the closer the timestep to the observation.

5. Conclusion

This work shows that ML is well suited to capture the nonlinear dependencies between cirrus drivers and CMP. We applied a twofold approach by training an XGBoost model on instantaneous cirrus driver data and an LSTM model with an attention layer on temporally resolved cirrus driver data. The LSTM-based model yielded the best predictive performance, suggesting that the history of an air parcel leading up to a cirrus cloud observation contains useful information. By analyzing the weights of the attention layer, we can show that only the last 15 hr before the observation is of interest. Furthermore, we calculate post hoc feature attributions for the XGBoost model using SHAP values. The analysis of SHAP values enables us to quantify the impact of driver variables on the prediction of CMP. While the physical interdependencies are mostly already known and discussed in the literature, our study can convert qualitative estimates of the dependencies to quantitative ones, such as the slope of the temperature dependency or the dust aerosol concentration needed to see an impact on the CMP. We also demonstrate the importance of vertical cloud extent and cloud layer position to the CMP, apart from meteorological and aerosol variables. Our study is limited by the imperfect prediction of cirrus CMP, reducing the accuracy of the feature attributions. The authors assume that subgrid-scale processes in reanalysis data, like small-scale updrafts and satellite retrieval uncertainties, reduce predictive performance. In conclusion, ML algorithms can predict cirrus properties from reanalysis and macrophysical observational data. Model explainability methods, namely SHAP and attention mechanisms, are helpful tools in improving process understanding.

Acknowledgments. We are grateful for the funding from the European Union's Horizon 2020 program. The authors would like to thank Heini Wemli, Hanin Binder, and Michael Sprenger for providing the backtrajectory data, Odran Sourdeval for helpful discussions on the DARDAR-Nice data set, and Miguel-Ángel Fernández-Torres for discussion on the ML implementation.

Author contribution. Conceptualization: K.J., U.L., D.N.; Data curation: K.J.; Data visualization: K.J.; Methodology: K.J, U.L., G.C.-V., D.N.; Writing—original draft: K.J. All authors approved the final submitted draft.

Competing interest. The authors declare no competing interests exist.

Data availability statement. The collocated data set and source code are available at <https://doi.org/10.5281/zenodo.7590909> and <https://doi.org/10.5281/zenodo.7988844>, respectively. The original data sets are freely available online: DARDAR-Nice (<https://doi.org/10.25326/09>), ERA5 (Hersbach et al., 2018), and MERRA2 (Global Modeling and Assimilation Office, 2015).

Ethics statement. The research meets all ethical guidelines, including adherence to the legal requirements of the study country.

Funding statement. This research was supported by grants from the European Union's Horizon 2020 research and innovation program iMIRACLI under Marie Skłodowska-Curie grant agreement No. 860100 and Swiss National Supercomputing Centre (Centro Svizzero di Calcolo Scientifico, CSCS; project ID s1144).

Supplementary material. The supplementary material for this article can be found at <https://doi.org/10.1017/eds.2023.14>.

Provenance statement. This article is part of the Climate Informatics 2023 proceedings and was accepted in *Environmental Data Science* on the basis of the Climate Informatics peer review process.

References

- Agarwal C, Johnson N, Pawelczyk M, Krishna S, Saxena E, Zitnik M and Lakkaraju H (2022) *Rethinking Stability for Attribution-based Explanations*.
- Alvarez Melis D and Jaakkola T (2018) Towards robust interpretability with self-explaining neural networks. In *Advances in Neural Information Processing Systems*. Red Hook, NY: Curran Associates.
- Belle V and Papantonis I (2020) Principles and practice of explainable machine learning. arXiv:2009.11698 [cs, stat]. arXiv: 2009.11698.
- Chen T and Guestrin C (2016) XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. New York, NY: ACM.
- Ebert-Uphoff I and Hilburn KA (2020) Evaluation, tuning and interpretation of neural networks for meteorological applications. arXiv:2005.03126 [physics]. arXiv: 2005.03126.
- Gasparini B, Meyer A, Neubauer D, Münch S and Lohmann U (2018) Cirrus cloud properties as seen by the CALIPSO satellite and ECHAM-HAM global climate model. *Journal of Climate* 31(5), 1983–2003.
- Global Modeling and Assimilation Office (2015) MERRA-2 inst3_3d_aer_nv: 3d,3-Hourly,Instantaneous,Model-Level,Assimilation,Aerosol Mixing Ratio V5.12.4.
- Gryspeerd E, Sourdeval O, Quaas J, Delanö J, Krämer M and Kühne P (2018) Ice crystal number concentration estimates from lidar–radar satellite remote sensing – Part 2: Controls on the ice crystal number concentration. *Atmospheric Chemistry and Physics* 18(19), 14351–14370.
- Hersbach H, Bell B, Berrisford P, Horányi A, Muñoz-Sabater J, Nicolas J, Peubey C, Radu R, Schepers D, Simmons A, Soci C, Biavati G, Dee D and Thépaut J (2018) ERA5 hourly data on pressure levels from 1959 to present. *Copernicus Climate Change Service (C3S) Climate Data Store (CDS)*.
- Heysfield AJ, Krämer M, Luebke A, Brown P, Cziczo DJ, Franklin C, Lawson P, Lohmann U, McFarquhar G, Ulanowski Z and Van Tricht K (2017) Cirrus clouds. *Meteorological Monographs* 58, 2.1–2.26.
- Hochreiter S and Schmidhuber J (1997) Long short-term memory. *Neural Computation* 9(8), 1735–1780.
- Jensen EJ, Ueyama R, Pfister L, Bui TV, Lawson RP, Woods S, Thornberry T, Rollins AW, Diskin GS, DiGangi JP and Avery MA (2016) On the susceptibility of cold tropical cirrus to ice nuclei abundance. *Journal of the Atmospheric Sciences* 73(6), 2445–2464.
- Kanji ZA, Ladino LA, Wex H, Boose Y, Burkert-Kohn M, Cziczo DJ and Krämer M (2017) Overview of ice nucleating particles. *Meteorological Monographs* 58, 1.1–1.33.
- Kärcher B, Hendricks J and Lohmann U (2006) Physically based parameterization of cirrus cloud formation for use in global atmospheric models. *Journal of Geophysical Research: Atmospheres* 111(D1), D01205.
- Koop T, Luo B, Tsias A and Peter T (2000) Water activity as the determinant for homogeneous ice nucleation in aqueous solutions. *Nature* 406(6796), 611–614.
- Krämer M, Rolf C, Spelten N, Afchine A, Fahey D, Jensen E, Khaykin S, Kuhn T, Lawson P, Lykov A, Pan LL, Riese M, Rollins A, Stroh F, Thornberry T, Wolf V, Woods S, Spichtinger P, Quaas J and Sourdeval O (2020) A microphysics guide to cirrus – Part 2: Climatologies of clouds and humidity from observations. *Atmospheric Chemistry and Physics* 20(21), 12569–12608.
- Kuebbeler M, Lohmann U, Hendricks J and Kärcher B (2014) Dust ice nuclei effects on cirrus clouds. *Atmospheric Chemistry and Physics* 14(6), 3027–3046.
- Liou K-N (1986) Influence of cirrus clouds on weather and climate processes: A global perspective. *Monthly Weather Review* 114(6), 1167–1199.
- Lundberg SM, Erion G, Chen H, DeGrave A, Prutkin JM, Nair B, Katz R, Himmelfarb J, Bansal N and Lee S-I (2020) From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence* 2(1), 56–67.
- Lundberg SM and Lee S-I (2017) A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, vol. 30. Red Hook, NY: Curran Associates, Inc.

- McGovern A, Lagerquist R, Gagne DJ, Jergensen GE, Elmore KL, Homeyer CR and Smith T** (2019) Making the black box more transparent: Understanding the physical implications of machine learning. *The Bulletin of the American Meteorological Society* 100(11), 2175–2199.
- Ribeiro MT, Singh S and Guestrin C** (2016) “Why should i trust you?” Explaining the predictions of any classifier. [arXiv:1602.04938](https://arxiv.org/abs/1602.04938) [cs, stat]. [arXiv: 1602.04938](https://arxiv.org/abs/1602.04938).
- Sassen K, Wang Z and Liu D** (2008) Global distribution of cirrus clouds from CloudSat/cloud-aerosol Lidar and infrared pathfinder satellite observations (CALIPSO) measurements. *Journal of Geophysical Research – Atmospheres* 113(D8), 1–3.
- Shwartz-Ziv R and Armon A** (2021) Tabular Data: Deep Learning is Not All You Need. [arXiv:2106.03253](https://arxiv.org/abs/2106.03253) [cs]. [arXiv: 2106.03253](https://arxiv.org/abs/2106.03253).
- Sourdeval O, Gryspeerdt E, Krämer M, Goren T, Delanoë J, Afchine A, Hemmer F and Quaas J** (2018) Ice crystal number concentration estimates from lidar–radar satellite remote sensing – Part 1: Method and evaluation. *Atmospheric Chemistry and Physics* 18(19), 14327–14350.
- Sprenger M and Wernli H** (2015) The LAGRANTO Lagrangian analysis tool – Version 2.0. *Geoscientific Model Development* 8(8), 2569–2586.
- Wernli H, Boettcher M, Joos H, Miltenberger AK and Spichtinger P** (2016) A trajectory-based classification of ERAInterim ice clouds in the region of the North Atlantic storm track. *Geophysical Research Letters* 43(12), 6657–6664.