

Explainable Weather Forecasts Through an LSTM-CBR Twin System

Craig Pirie^{1,*}, Malavika Suresh¹, Pedram Salimi¹, Chamath Palihawadana¹ and Gayani Nanayakkara¹

¹*School of Computing, The Sir Ian Wood Building, Garthdee Road, Robert Gordon University, Aberdeen, Scotland, United Kingdom, AB10 7GJ*

Abstract

In this paper, we explore two methods for explaining LSTM-based temperature forecasts using previous 14 day progressions of humidity and pressure. First, we propose and evaluate an LSTM-CBR twin system that generates nearest-neighbors that can be visualised as explanations. Second, we use feature attributions from Integrated Gradients to generate textual explanations that summarise the key progressions in the past 14 days that led to the predicted value.

Keywords

LSTM, XCBR, NLG, Integrated Gradients, Forecasting, Visualisation

1. Introduction

Inaccurate weather forecasting can have a serious impact to life, property and businesses. For instance, farmers rely on weather forecasting to plan activities such as ploughing, harvesting and cultivation [1]. In 2018, the inability of weather forecasts to detect an early monsoon indirectly trapped 12 Thai children and their football coach in a cave for 18 days – and caused the death of two of their rescuers¹. Efficiently predicting weather can help to minimise these losses in the future [2]. However, weather is stochastic and its patterns are non-linear, making it a challenge to obtain precise predictions. The problem lends itself well to neural networks for their ability to model non-linearity. Particularly, weather forecasting suits sequential methods such as the Long Short-Term Memory (LSTM) [3], to allow for the utilisation of time-series data in the prediction process [2]. By making use of the recurrent nature of the LSTM, information can be fed back through different steps along a time-series, capturing trends, seasonality and more, which all help increase predictive power.

ICCBR XCBB'22: 4th Workshop on XCBB: Case-based Reasoning for the Explanation of Intelligent Systems at ICCBR-2022, September, 2022, Nancy, France

*Corresponding author.

✉ c.pirie11@rgu.ac.uk (C. Pirie); m.suresh@rgu.ac.uk (M. Suresh); p.salimi@rgu.ac.uk (P. Salimi);

c.palihawadana1@rgu.ac.uk (C. Palihawadana); g.nanayakkara@rgu.ac.uk (G. Nanayakkara)

🌐 <https://www.linkedin.com/in/craig-pirie-aberdeen/> (C. Pirie)

🆔 0000-0002-6799-0497 (C. Pirie); 0000-0002-6572-8516 (P. Salimi); 0000-0002-9594-8683 (C. Palihawadana);

0000-0002-0017-6589 (G. Nanayakkara)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

¹Read 'The full story of Thailand's extraordinary cave rescue' here

Although weather forecasts are common in everyday life, we seldom seek for deeper insights to explain how a prediction was made. On the contrary – if we *did* ask for further explanation – we could pass better judgement on the validity of predictions. Explainable AI (XAI) has become increasingly prevalent in recent times. This is because despite deep learning methods now excelling in performance in a variety of domains, they often remain as black-boxes; we are unable to ‘look inside’ to understand why a prediction was made. Methods such as Integrated Gradients [4] have been useful for producing saliency maps in images to highlight the important pixels that weigh heavily on the outcome of a classification. They have also been applied in a time-series setting but are often unsophisticated and can be difficult to interpret – presenting a need for dedicated time-series solutions [5, 6]. Literature shows that some (albeit few) attempts have been made to explain weather forecasts in a time-series setting such as by score-maximisation and occlusion analysis visualisations [7], or applying LIME [8]. However, none of these methods propose the use of case-based reasoning (CBR) to provide explanations. Using an approach grounded in similarity, such as CBR, *a posteriori* knowledge can be drawn upon to help both prediction and explanation. Past experiences can be drawn upon to make arguments for why a decision was made. We propose an approach that takes the learned LSTM embeddings from the prediction stage and uses these to build a case-base to provide explanations that can be visualised. We will also use Integrated Gradients and NLG to provide further explanation by means of a report.

2. Methodology

2.1. Predicting Temperature with LSTM

We use the default prediction model provided for the challenge which consists of two bidirectional LSTM layers.

2.2. Building the Twin LSTM-CBR System

We use the Cloud CBR ² system as the twin. Cloud CBR is a distributed CBR system which supports high scalability. The system has a micro-service architecture which splits an application into a set of smaller and interconnected services that scale to meet varying demands [9]. The output embeddings from the final LSTM layer are extracted for each input row and are fed into the Cloud CBR system to form the case base. For a given row and LSTM prediction that we want to explain, we query the Cloud system using the nearest-neighbors approach. This gives us the top three similar rows based on the embedding similarity. To establish the quality of the twin, we compare our approach to a baseline twin system that uses the raw features instead of the embeddings.

We use Mean Absolute Error (MAE) as the evaluation metric to compare the two twin systems as it is one of the common measures of forecasting error in time-series analysis. MAE measures the absolute difference $|y_i - x_i|$ between a prediction y_i and an actual observation x_i where the individual differences share the equal weight. The 3 nearest neighbours for a prediction are retrieved from the case-base and the mean MAE is calculated between their predicted values.


²<https://github.com/rgu-computing/cloud>

Comparing the mean gives an indication of whether the cases are better aligned through the use of embeddings – or not.

$$MAE = \frac{\sum_{i=1}^n |y_i - x_i|}{n} \quad (1)$$

2.3. Generating Explanations

A template-based Natural Language Generation approach to generate explanations was taken. Based on Integrated Gradients, we extracted feature importance attributions for each day over three different time frames. These were as follows: the entire two-week period; the final week; and the final three days. The prediction, feature attributions and general information from the dataset were then slotted into the template to produce a weather forecast report (as seen in Figure 1). Furthermore, to aid interpretability, we add visualisations (emojis, graphics and colour) to our textual explanation to indicate the weather prediction.



The temperature for Saturday the 29th of January 2000 is predicted to be 25.4 degrees 🌡️.
 This is because in the two weeks prior (from Saturday the 1st of January 2000 to Friday the 28th of January 2000) humidity was high and pressure was low.
 Furthermore, our system suggests that:

- in the last 14 days the most important feature was the average humidity on the 7th day (2000-01-21);
- that in the last week the most important feature was the average humidity on the 8th day (2000-01-22);
- and finally in the last three days the most important feature was the minimum humidity on the 14th day (2000-01-28).

Figure 1: An example weather report showing the prediction, an explanation of the most important features over two-week, one-week and 3-day periods. A thermometer with a high reading, a hot-faced emoji and red-text are included to help visualise the results.

3. Results

3.1. LSTM-CBR Twin Evaluation

The results in Table 1 prove that the twin LSTM-CBR system based on embeddings performs better than a twin modelled around raw features. The embedded approach achieved an MAE of 0.066 degrees whereas the raw features approach achieved an MAE of 0.133 degrees. This means that using embeddings, the twin system is able to find better nearest-neighbours, therefore can produce better explanations.

3.2. Nearest Neighbors Explanations Using LSTM-CBR Twin

In Figures 2 and 3, we visualize the progression of the average humidity and pressure, respectively, over the previous 14 days for a sample query row, top 3 nearest neighbors and 3 randomly selected rows. As can be seen, the query and the nearest neighbors have a similar progression of the predictor variables as well as close predictions for the subsequent day temperature.

Embedded Features	Raw Features
0.066	0.133

Table 1
LSTM-CBR MAEs (Embedded & Raw) compared with the original LSTM predictions.




Figure 2: Nearest Neighbors Explanation (Average Humidity): Relation=1 depicts the query, Relation=2 depicts the 3 nearest neighbors from the lstm-cbr twin, Relation=3 depicts 3 randomly chosen data points (best viewed in colour)




Figure 3: Nearest Neighbors Explanation (Average Pressure): Relation=1 depicts the query, Relation=2 depicts the 3 nearest neighbors from the lstm-cbr twin, Relation=3 depicts 3 randomly chosen data points (best viewed in colour)

4. Conclusion

In this paper we explore two methods for explaining LSTM predictions of weather forecast. We first propose an LSTM-CBR twin system that explains a prediction using nearest-neighbors and then visualises the explanations through parallel coordinate plots. We show that using the LSTM embeddings can result in better twinning compared to raw features. As a second explanation method, we make use of feature attributions from Integrated Gradients in text

templates to generate text explanations. We note that a current limitation of the twin system is that it considers all days as equally important when retrieving the nearest-neighbors. Therefore, future work could explore different weighting strategies (such as weights derived from feature attributions) for strengthening the twin system. A second limitation lies with the need to provide a template for the NLG-based explanation, which would differ between domains. Further work could explore the application of more advanced NLG techniques that allow the inference of templates in a new domain.

Acknowledgments

We would like to thank Professor Nirmalie Wiratunga for motivating us to take part in this challenge and for her advise throughout the competition.

References

- [1] P. Hewage, A. Behera, M. Trovati, E. Pereira, M. Ghahremani, F. Palmieri, Y. Liu, Temporal convolutional neural (tcn) network for an effective weather forecasting using time-series data from the local weather station, *Soft Computing* 24 (2020) 16453–16482.
- [2] D. N. Fente, D. Kumar Singh, Weather forecasting using artificial neural network, in: 2018 Second International Conference on Inventive Communication and Computational Technologies (ICICCT), 2018, pp. 1757–1761. doi:10.1109/ICICCT.2018.8473167.
- [3] S. Hochreiter, J. Schmidhuber, Long short-term memory, *Neural computation* 9 (1997) 1735–1780.
- [4] M. Sundararajan, A. Taly, Q. Yan, Axiomatic attribution for deep networks, *CoRR* abs/1703.01365 (2017). URL: <http://arxiv.org/abs/1703.01365>. arXiv:1703.01365.
- [5] T. Rojat, R. Puget, D. Filliat, J. Del Ser, R. Gelin, N. Díaz-Rodríguez, Explainable artificial intelligence (xai) on timeseries data: A survey, arXiv preprint arXiv:2104.00950 (2021).
- [6] U. Schlegel, H. Arnout, M. El-Assady, D. Oelke, D. A. Keim, Towards a rigorous evaluation of xai methods on time series, in: 2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW), 2019, pp. 4197–4201. doi:10.1109/ICCVW.2019.00516.
- [7] I. A. Abdellaoui, S. Mehrkanon, Deep multi-stations weather forecasting: explainable recurrent convolutional neural networks, 2020. URL: <https://arxiv.org/abs/2009.11239>. doi:10.48550/ARXIV.2009.11239.
- [8] M. Zdravković, I. Ćirić, M. Ignjatović, Explainable heat demand forecasting for the novel control strategies of district heating systems, *Annual Reviews in Control* (2022).
- [9] I. Nkisi-Orji, N. Wiratunga, C. Palihawadana, J. A. Recio-García, D. Corsar, Cloud cbr: Towards microservices oriented case-based reasoning, in: I. Watson, R. Weber (Eds.), *Case-Based Reasoning Research and Development*, Springer International Publishing, Cham, 2020, pp. 129–143.