

CROSSLAG: PREDICTING MAJOR DENGUE OUTBREAKS WITH A DOMAIN KNOWLEDGE INFORMED TRANSFORMER

Ashwin Prabu, Nhat Thanh Tran, Guofa Zhou, Jack Xin

University of California, Irvine, CA 92697, USA

ABSTRACT

A variety of models have been developed to forecast dengue cases to date. However, it remains a challenge to predict major dengue outbreaks that need timely public warnings the most. In this paper, we introduce CrossLag, an environmentally informed attention that allows for the incorporation of lagging endogenous signals behind the significant events in the exogenous data into the architecture of the transformer at low parameter counts. Outbreaks typically lag behind major changes in climate and oceanic anomalies. We use TimeXer, a recent general-purpose transformer distinguishing exogenous-endogenous inputs, as the baseline for this study. Our proposed model outperforms TimeXer by a considerable margin in detecting and predicting major outbreaks in Singapore dengue data over a 24-week prediction window.

Index Terms— Dengue, forecasting, time series, transformers, attention

1. INTRODUCTION

The forecasting of public health time series is an important scientific problem: the ability to do so can help detect dangerous anomalies, thus helping to make a quick, informed, and effective response. Dengue is one of the deadliest vector-borne diseases in a great number of tropical regions in the world. The limited availability of resources to help treat and prevent it in a large number of places gives rise to the importance of being able to predict trends, especially when there is a spike in the number of cases of infection.

Many studies [1, 2, 3, 4, 5, 6, 7] have already discussed the correlation between climate/weather data and vector-borne diseases. The climate in a region plays a crucial role in the propagation of mosquitoes and the rate of development of both vector and virus [3]. Additionally, variables based on oceanic data, such as Indian Ocean Dipole, impact the local climate and, in turn, epidemic patterns and trends [3]. The relationship between these variables and dengue outbreaks can be difficult to capture and thus difficult to predict. Smaller outbreaks may be linked to local importation and favorable local climate conditions [5], while regional outbreaks are

likely associated with large-scale climatic events [6]. Given this, a more advanced approach may be required to forecast them.

Today, many state-of-the-art models used in time series forecasting are transformers [8, 9, 10]. The attention mechanism [11] present in these models has allowed them to capture complicated properties of the data and predict future patterns with high accuracy. However, many transformers require large amounts of data to train on. This poses a challenge in using them for the dengue forecasting task, as dengue data tends to be smaller and less abundant than other time series data. For example, [12] illustrates the potential of such models to capture trends in the data, but also the difficulty for them to capture the patterns well. Thus, there is a need for a compact and efficient model that can capture the complex structure while training on a small amount of data.

TimeXer [8] is one such model that explicitly models the interactions between the endogenous feature (target feature) and the exogenous features. However, it does not directly incorporate any domain knowledge: it infers dependencies from the data directly. Prior works [9, 13, 10] have demonstrated the benefit of introducing a lag-based calculation within attention, in a variety of ways. However, the lags tend to be selected within the model by computation. Although this is logical for the general case, this does not fully utilize known lagged effects in the specific field: for example, the lagged effects of temperature and precipitation in a region on mosquito populations as observed in [7].

In this paper, we propose a new model, CrossLag, which employs a fixed lag-based attention to reduce the number of attention weights without losing meaningful information. The key feature is the cross-lag attention. Unlike TimeXer, it allows for interaction between the endogenous and the exogenous variables according to fixed lags. The model allows for a known lags to be input directly or provides a range of lag values if given a maximum lag factor. The model uses this to construct a bank of lagged keys. This ensures that the attention mechanism only lets each query attend to the specified keys. It also embeds each exogenous feature uniquely: embedding each feature at each timestep into a higher-dimensional tensor. This allows for the lag factor to act upon each feature while utilizing a high-dimensional representation of the feature at each time step.

This work was partially supported by NSF grants DMS-2151235, DMS-2219904, DMS-2309520.

The main contributions are:

- A per-feature embedding that encodes domain-relevant information such as weekly periodicity, annual drift, and local properties to enrich the embedding.
- A variant of attention that matches each query with a subset of recent keys (L out of T keys), reducing computation cost while preserving relevant information.
- A model that utilizes aforementioned embeddings and attention to allow for the interaction of the endogenous feature with each of the exogenous features of fixed lags based on domain knowledge.
- We study the model’s performance in forecasting trends in the dengue data from Singapore with the aid of statistically selected regional climate and oceanic data, demonstrating its effectiveness and potential for the task.

2. DATA

The dataset comprises weekly dengue data from Singapore, from 2000 to 2019. The data contains 1000 samples and 16 features: number of cases, average temperature, precipitation, Southern Oscillation Index (SOI), Oceanic Niño Index (ONI), ONI anomaly, Indian Ocean Dipole (IOD), IOD East, NIÑO1+2, NIÑO3, NIÑO4, and each NIÑO’s anomaly. Of these, we use only 5 of the 15 climate and oceanic features: ONI anomaly, NIÑO4 anomaly, IOD East, average temperature, and precipitation. We do this as many oceanic features are highly correlated; thus, we select only a subset of them to represent the influence of oceanic variables on dengue spread. We selected these 3 oceanic features based on their proximity to Singapore and their estimated Mutual Information (MI) [14, 15] with respect to the number of cases. We present a table with the estimated MI of each selected feature with respect to the number of cases in Tab. 1. We also present the plot, Fig. 1, of the NIÑO4 anomaly alongside the plot of the number of cases in order to help visualize the data and the reason for feature selection better.

The split ratio is 6:2:2 (train:val:test) in chronology. The data is normalized based on the training set statistics before feeding the input to the models: each feature has a mean of 0 and a standard deviation of 1. Additionally, we scale week stamps uniformly and standardize the year stamps prior to feeding these time stamps into the models. The processed data is available upon request.

3. MODEL

There is variation between the latency of effect for the different features on the number of cases. However, domain knowl-

Feature	Estimated MI
SOI	0.28
ONI	0.38
ONI Anomaly	0.38
IOD	0.29
IOD East	0.44
NIÑO1+2	0.35
NIÑO1+2 Anomaly	0.42
NIÑO3	0.30
NIÑO3 Anomaly	0.35
NIÑO4	0.38
NIÑO4 Anomaly	0.48
NIÑO3.4	0.43
NIÑO3.4 Anomaly	0.42

Table 1. Estimated MI between each feature and Cases.




Fig. 1. NIÑO4 Anomaly and number of cases in weeks.

edge allows us to recognize that certain weeks play more important roles than others. Therefore, standard attention, which looks at all keys and then learn which weeks are important, may fail to capture this, especially when there is a limited amount of data and there is a large amount of noise. Therefore, we would like to incorporate some amount of domain knowledge into our attention, allowing the model to look at a smaller and more relevant subset of time steps for the different features.

Another important part of the model is the embeddings: they must be used effectively. For example, using the traditional sinusoidal encoding would not be representative of the structure of the data, as it would not reflect the seasonality of the features and the latency of the effects.

To ensure that the model makes as much use of domain knowledge and the natural structure of the data as possible, we propose the following architecture for our model: custom embeddings for exogenous and endogenous features, a self-attention on the endogenous feature, a cross attention which incorporates different fixed lags for the keys, and finally a fully connected layer consisting of Multi Layered Perceptrons (MLPs). Additionally, we regulate the flow of data by using gates that take a weighted sum of the input and output of a layer before passing it onto the next, as seen in Fig. 2. There

are 3 gates present in the model, learnable parameters, to ensure that minimal information loss occurs in the model.




Fig. 2. Architecture of CrossLag.

3.1. Embeddings

To extract the maximum amount of information from the data before passing it on, we encode the following information before projecting to the final embedding: weekly periodicity, annual drift, local average, rate of change, and residual.

Let $w \in [-\frac{1}{2}, \frac{1}{2}]^T$ represent the normalized week vector and $y \in \mathbb{R}^T$ the normalized year vector. Let $\theta_i = \frac{2\pi w_i}{P}$ where P represents the periodicity factor, and is 50 by default. Then, $\forall i \in [1, T]$, $w_i^{enc} = \{\sin(\theta_i), \cos(\theta_i), \sin(2\theta_i), \cos(2\theta_i)\}$ is the encoded weekly periodicity, and $y_i^{enc} = cy_i + b$ is the encoded annual drift. Note, c, b are learnable scalars.

Let $x \in \mathbb{R}^T$ be any feature. We compute the rate of change, $x' \in \mathbb{R}^T$ by $x'_t = x_t - x_{t-1} \forall t \in [2, T]$ and $x'_1 = 0$.

We compute the local average using a 1d convolution for the previous N values (with padding), set to 4 by default. We let $\beta \in \mathbb{R}^N$ be the convolution weights. Thus, our local average is

$$x_t^{(m)} = \sum_{i=1}^N \beta_i x_{t-i+1} \quad \forall t \in [1, T].$$

Using this information, we also compute the residual $x^{(r)} = x - x^{(m)}$. For the endogenous feature, $x \in \mathbb{R}^T$,

$$\begin{aligned} X^{enc} &= (x \| y^{enc} \| w^{enc} \| x^{(m)} \| x^{(r)} \| x') \\ X^{endo} &= \text{Dropout}(\text{Linear}(X^{enc})) \in \mathbb{R}^{T \times d_m} \end{aligned} \quad (1)$$

is the endogenous embedding ($\|$ represents concatenation). For each exogenous feature $Z_{:,f}$ where $f \in [1, F]$, we have have

$$\begin{aligned} Z_{:,f}^{enc} &= (Z_{:,f} \| y^{enc} \| w^{enc} \| Z_{:,f}^{(m)} \| Z_{:,f}^{(r)} \| Z_{:,f}') \\ Z_{:,f}^{exo} &= \text{Dropout}(\text{Linear}(Z_{:,f}^{enc})) \in \mathbb{R}^{T \times d_m}. \end{aligned} \quad (2)$$

We compute the residual for the endogenous feature only. We do this to give the endogenous feature, our target feature, more context, as the model is expected to learn more from it. We perform the exogenous embedding per feature, thus our result is a 3d tensor $Z^{exo} \in \mathbb{R}^{T \times F \times d_m}$ instead of the usual 2d matrix. We do so to preserve the lag effect of each feature while also being able to enrich the data with the embeddings.

3.2. Cross Lag Attention

The proposed attention mechanism introduces the idea of using a bank of tensors to collect different lagged inputs before projecting them as keys. We use a lag vector $\ell \in \mathbb{Z}_{>0}^L$, with L lag values. There are two options offered for the choice of lag vector: an input vector by the user, or a *max_lag* hyperparameter, which then selects the lag vector to be $\ell \in \mathbb{Z}_{>0}^L$: $\ell_i = 2 \cdot (i - 1) \forall 1 \leq i \leq \lfloor \text{max_lag}/2 \rfloor + 1$. By default, the model uses *max_lag* = 8 to generate the lag vector, which is at the responsive end of the lagged effects of climate on dengue spread, keeping it epidemiologically meaningful [3].

Let $L = |\ell|$. For each feature $1 \leq f \leq F$ and each lag $1 \leq \delta \leq L$, we create a tensor $Z^{lag} \in \mathbb{R}^{T \times L \cdot F \times d_m}$ where

$$Z_{t,f+(F \cdot (\delta-1)),:}^{lag} = \begin{cases} Z_{t-\ell_\delta,f,:}^{exo} & t - \ell_\delta \geq 1, \\ 0 & t - \ell_\delta < 1. \end{cases} \quad (3)$$

We also define our mask by our valid key set $M(t) = \{f + F \cdot (\delta - 1) : t - \ell_\delta \geq 1\} \subset [1, F \cdot L]$.

We have our queries $Q = \text{Linear}(X^{endo}) \in \mathbb{R}^{T \times d}$, keys and values $K = V = \text{Linear}(Z^{lag}) \in \mathbb{R}^{T \times F \cdot L \times d}$. Thus, our attention mechanism follows the following:

$$S_{t,j} = \frac{\langle Q_{t,:}, K_{t,j,:} \rangle}{\sqrt{d}} \quad (4)$$

$$\hat{S}_{t,j} = \begin{cases} S_{t,j} & j \in M(t) \\ -\infty & j \notin M(t) \end{cases} \quad (5)$$

$$A_{t,j} = \begin{cases} \frac{\exp(\hat{S}_{t,j})}{\sum_{i \in M(t)} \exp(\hat{S}_{t,i})} & |M(t)| > 0 \\ 0 & |M(t)| = 0 \end{cases} \quad (6)$$

$$O_{t,:} = \text{Dropout}(A)_{t,:} V_{t,:} \quad (7)$$

The difference between our attention and standard attention is that we fix $t \in [1, T]$ for all Q, K, V instead of allowing the attention to compute for different time indices. This allows the attention to check only certain indices for each query. This reduces the calculation complexity of the attention scores $O(F \cdot T)$ to $O(F \cdot L)$ per query.

4. EXPERIMENT SETUP

We compare our proposed model, CrossLag, with TimeXer on the Singapore data. We use the default parameters of both

models in the experiments, adjusting only the encoder input size (6), decoder input size (1), and model output size (1) in TimeXer. TimeXer has $d_m = 512$, $d_{ff} = 2048$, $n_{heads} = 8$ as default, whereas CrossLag has $d_m = 128$, $d_{ff} = 16$, $n_{heads} = 1$ as default. TimeXer uses 2 encoder layers, whereas ours uses only 1. Both were trained for 30 epochs, with a patience of 10 epochs. We use the default optimizer (Adam), initial learning rate ($lr=10^{-4}$), and scheduler for TimeXer; we use AdamW with its default PyTorch parameters and a custom scheduler with an exponential decay (factor = 0.9) with floor ($lr = 5 \cdot 10^{-5}$) for CrossLag. We use these defaults for each model to ensure that each can perform to the best of its abilities. We feed the same features and timestamps to both models as explained in Sec. 2. Finally, we use an input sequence length of 16, a label length of 12, and a prediction length of 24. All implementations use the PyTorch framework and are trained on CPU.

5. EXPERIMENTAL RESULTS

For this study, we look at the models’ ability to predict spikes in dengue cases. Our main goal is to predict the largest spike in the test set. All plots have the number of weeks since the initial input for the x-axis and the number of dengue cases as the y-axis. The black solid line is the input history available to the models, the blue solid line is the ground truth, and the orange dashed line is the prediction of either model.

The largest spike in the test set is a jump from approximately 200 cases to 600 cases reported per week, which can be found at index 30 (hence the label “Truth 30”). Again, due to the limited availability of data to train on, there are only a few such large spikes for the models to train on, as seen in Fig. 1. Therefore, this serves as a test for the models to observe their versatility and how well they can perform when working with limited data.




Fig. 3. Predictions in default parameter settings.

Fig. 3 shows a sample of both models’ predictions. CrossLag (left) can recognize the spike and match the overall pattern with high accuracy. TimeXer (right) can not recognize that there is a spike in cases. This is explainable as CrossLag is a model that has a greater sensitivity to changes in data by design: the different information encoded in the embeddings

provides the model with more information about changes in the data, and the lag selection ensures the model pays attention only to the possible relevant information. Fig. 4 shows the robustness of CrossLag to recognize the large spike. Although predictions may vary in quality, CrossLag foresees a rise in case numbers as d_m and d_{ff} are halved, respectively. It also suggests optimality of the default parameters.




Fig. 4. Robustness of CrossLag as hyperparameters vary.

Over 10 independent runs, the default CrossLag produces a lower MSE and MAE for predictions just before and during the spike episode (Tab. 2). However, CrossLag has a higher variance, which reflects its sensitive nature. CrossLag is still able to predict that there is a spike in the rare large spike case, which is of great importance as that indicates a more serious period that warrants civilian warning, logistic, medical, and administrative preparations. We note that its predictions do not always accurately portray how the spike pattern will be. This weakness does not understate the fact that CrossLag is effective in providing a trend (not in detail) prediction on noisy time series data when certain correlated exogenous factors are identified via domain knowledge.

Metric	TimeXer	CrossLag
MSE (mean $\pm$ std)	31802 $\pm$ 1216	22268 $\pm$ 4773
MAE (mean $\pm$ std)	131.72 $\pm$ 4.26	106.21 $\pm$ 13.54

Table 2. Model performances over 10 runs.

6. CONCLUSION

We proposed CrossLag, a transformer that takes a domain knowledge-informed approach to data embedding and cross attention. We tested it against a strong baseline model, TimeXer, and compared their performance for forecasting trends in dengue cases in Singapore weeks in advance. CrossLag shows a greater ability to detect dengue outbreaks compared to TimeXer. Despite its limitations in predicting details, CrossLag offers a new direction for modeling and forecasting vector-borne diseases, utilizing both static domain knowledge as well as dynamic contextual information present in the data.

7. REFERENCES

- [1] G. Zhou, N. Minakawa, A.K. Githeko, and G. Yan, “Association between climate variability and malaria epidemics in the east african highlands,” *Proceedings of the National Academy of Sciences*, vol. 101, no. 8, pp. 2375–2380, February 2004.
- [2] C.W. Morin, A.C. Comrie, and K. Ernst, “Climate and dengue transmission: Evidence and implications,” *Environmental Health Perspectives*, vol. 121, no. 11-12, pp. 1264–1272, September 2013.
- [3] T.-W. Chuang, L.F. Chaves, and P.-J. Chen, “Effects of local and regional climatic fluctuations on dengue outbreaks in southern taiwan,” *PLOS ONE*, vol. 12, no. 6, pp. 1–20, June 2017, e0178698.
- [4] P. Liyanage, Y. Tozan, H.J. Overgaard, H.A. Tissera, and J. Rocklöv, “Effect of el niño–southern oscillation and local weather on aedes vector activity from 2010 to 2018 in kalutara district, sri lanka: a two-stage hierarchical analysis,” *The Lancet Planetary Health*, vol. 6, no. 7, pp. e577–e585, July 2022.
- [5] C. Stephenson, E. Coker, S. Wisely, S. Liang, R.R. Dinglasan, and J.A. Lednicky, “Imported dengue case numbers and local climatic patterns are associated with dengue virus transmission in florida, usa,” *Insects*, vol. 13, no. 2, 2022.
- [6] S. Mokhtar, D.C.P. Ratterree, A.F. Britt, R. Fisher, and M.L. Ndeffo-Mbah, “Global risk of dengue outbreaks and the impact of el niño events,” *Environmental Research*, vol. 262, pp. 119830, 2024.
- [7] M.H. Myer, C.M. Fizer, K.R. McPherson, A.C. Neale, A.N. Pilant, A. Rodriguez, P.-Y. Whung, and J.M. Johnston, “Mapping aedes aegypti (diptera: Culicidae) and aedes albopictus vector mosquito distribution in brownsville, tx,” *Journal of Medical Entomology*, vol. 57, no. 1, pp. 231–240, August 2020.
- [8] Y. Wang, H. Wu, J. Dong, G. Qin, H. Zhang, Y. Liu, Y. Qiu, J. Wang, and M. Long, “Timexer: Empowering transformers for time series forecasting with exogenous variables,” in *Advances in Neural Information Processing Systems*, A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, Eds. 2024, vol. 37, pp. 469–498, Curran Associates, Inc.
- [9] H. Wu, J. Xu, J. Wang, and M. Long, “Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting,” in *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, Eds. 2021, vol. 34, pp. 22419–22430, Curran Associates, Inc.
- [10] Y. Yang, Q. Zhu, and J. Chen, “Vcformer: Variable correlation transformer with inherent lagged correlation for multivariate time series forecasting,” in *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*, Kate Larson, Ed. August 2024, pp. 5335–5343, International Joint Conferences on Artificial Intelligence Organization, Main Track.
- [11] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Advances in Neural Information Processing Systems*, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds. 2017, vol. 30, Curran Associates, Inc.
- [12] N.T. Tran, J. Xin, and G. Zhou, “Fwin transformer for dengue prediction under climate and ocean influence,” in *Machine Learning, Optimization, and Data Science (LOD 2024, ACAIN 2024)*, G. Nicosia, V. Ojha, S. Gies-selbach, M.P. Pardalos, and R. Umeton, Eds., Cham, 2025, vol. 15509 of *Lecture Notes in Computer Science*, pp. 160–175, Springer Nature Switzerland.
- [13] Q.M. Nguyen, L.M. Nguyen, and S. Das, “Correlated attention in transformers for multivariate time series,” in *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2025, pp. 1–5.
- [14] A. Kraskov, H. Stögbauer, and P. Grassberger, “Estimating mutual information,” *Phys. Rev. E*, vol. 69, pp. 066138, June 2004.
- [15] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and É. Duches-nay, “Scikit-learn: Machine learning in python,” *Journal of Machine Learning Research*, vol. 12, no. 85, pp. 2825–2830, 2011.