



Research paper



SERT: A transformer based model for multivariate temporal sensor data with missing values for environmental monitoring

Amin Shoari Nejad ^{a,*}, Rocío Alaiz-Rodríguez ^b, Gerard D. McCarthy ^c, Brian Kelleher ^d, Anthony Grey ^d, Andrew Parnell ^a

^a Hamilton Institute, Insight Centre for Data Analytics, Maynooth University, Kildare, Ireland

^b Department of Electrical, Systems and Automation, University of León, Spain

^c ICARUS, Department of Geography, Maynooth University, Maynooth, Ireland

^d Organic Geochemical Research Laboratory, Dublin City University, Glasnevin Campus, Dublin 9, Ireland

ARTICLE INFO

Keywords:

Spatio-temporal

Deep learning

Transformers

Environmental monitoring

ABSTRACT

Environmental monitoring is crucial to our understanding of climate change, biodiversity loss and pollution. The availability of large-scale spatio-temporal data from sources such as sensors and satellites allows us to develop sophisticated models for forecasting and understanding key drivers. However, the data collected from sensors often contain missing values due to faulty equipment or maintenance issues. The missing values rarely occur simultaneously leading to data that are multivariate misaligned sparse time series. We propose two models that are capable of performing multivariate spatio-temporal forecasting while handling missing data naturally without the need for imputation. The first model is a transformer-based model, which we name **SERT** (Spatio-temporal Encoder Representations from Transformers). The second is a simpler model named **SST-ANN** (Sparse Spatio-Temporal Artificial Neural Network) which is capable of providing interpretable results. We conduct extensive experiments on two different datasets for multivariate spatio-temporal forecasting and show that our models have competitive or superior performance to those at the state-of-the-art.

1. Introduction

The importance of spatio-temporal forecasting has increased significantly in recent years due to the availability of large-scale spatio-temporal data from various sources such as sensors and satellites (Hamdi et al., 2022). Spatio-temporal forecasting involves predicting how data vary over space and time, which is critical for a wide range of applications such as water quality forecasting (Deng et al., 2022). A common approach to modelling spatio-temporal data is to use a multivariate time series structure, where each time series is associated with a variable at a specific location (Wikle, 2015). Stochastic methods such as state space models (Durbin and Koopman, 2012) and Gaussian processes (GPs) (Ren and Han, 2014) are commonly employed for this purpose. In particular, autoregressive models and its variants, such as vector autoregression (VAR) (Sims, 1980), have been widely used to predict future trends while accounting for randomness. These approaches, complemented by machine learning techniques, provide a comprehensive toolkit for handling the complexities of spatio-temporal data analysis.

Spatio-temporal data often contain missing values which is a common problem in environmental monitoring, and can be caused by

sensor failure, malfunction or communication problems (for example see Fig. 1). A common remedy for forecasting with missing data is to impute the missing values using a variety of methods such as mean substitution, interpolation, or advanced techniques like multiple imputation (Van Buuren, 2018; Stekhoven and Bühlmann, 2012). However, these methods are not always effective as they may introduce biases or fail to accurately capture the underlying patterns in the data. For instance, mean substitution can oversimplify the data's complexity, leading to a loss of variability and potentially masking important signals. Advanced techniques like multiple imputation, while more sophisticated, might not always be suitable for all types of data or missingness patterns. They may lead to models that are overfitted to the imputed values. Therefore it is important to design models that can handle missing data naturally whilst simultaneously learning the underlying patterns in the data without resorting to imputation.

We propose a new model that is capable of performing multivariate spatio-temporal forecasting named SERT (Spatio-temporal Encoder Representations from Transformers). Our model is an extension of the well-known transformer architecture that has shown remarkable success in natural language processing and also image analysis (Devlin

* Corresponding author.

E-mail address: amin.shoarinejad@gmail.com (A.S. Nejad).

<https://doi.org/10.1016/j.cageo.2024.105601>

Received 23 June 2023; Received in revised form 29 February 2024; Accepted 22 April 2024

Available online 7 May 2024

0098-3004/© 2024 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

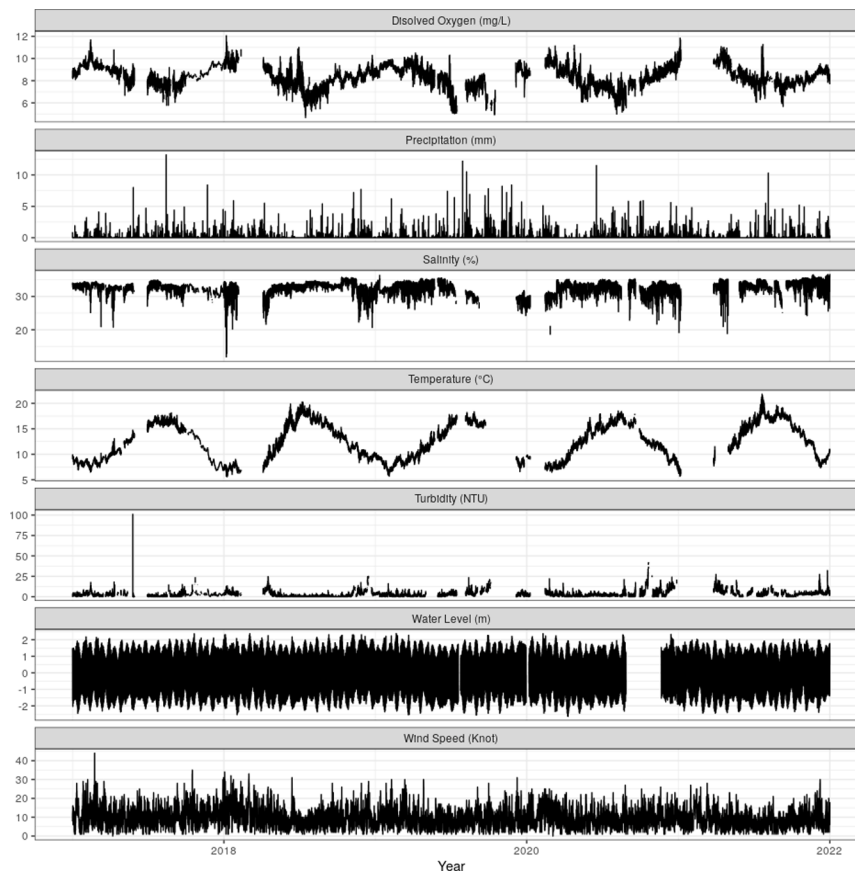


Fig. 1. An example of our environmental monitoring dataset; a multivariate environmental time-series with missing values.

et al., 2018; Dosovitskiy et al., 2020). SERT is designed to capture the complex joint temporal and spatial dependencies among the input variables. An important feature of our proposed model which differentiates it from the many other available methods is its ability to handle missing data more naturally without requiring any missing value imputation.

In addition to the SERT model, we introduce an interpretable simplified version that provides insights into the underlying factors that drive the predicted values, which can assist in decision and policy-making. Our proposed simplified model is named SST-ANN (Sparse Spatio-Temporal Artificial Neural Network) and removes the transformer layers from the SERT structure. Despite being less accurate than SERT, it is capable of providing insightful results with faster computation time while similarly being able to handle missing values. Depending on the complexity of the problem, the required accuracy and the available computational resources, the user can fit both SERT and SST-ANN, using the former to provide more accurate forecasts and the latter to forecast and gain insights about how the results were obtained.

To evaluate the performance of our proposed models, we conducted extensive experiments on two different datasets for multivariate spatio-temporal forecasting. We fitted our models to a simulated dataset to assess their ability to function under different levels of sparsity. We also evaluated the performance of the models on a real-world dataset, including missing values, of environmental variables in Dublin Bay for 7 h ahead forecasting. Our experimental results show that our models are competitive with state-of-the-art models for multivariate spatio-temporal forecasting.

Our paper is organized as follows. In Section 2, we provide a brief overview of related work on models developed for analysing sequential data in general and spatio-temporal forecasting applied to environmental monitoring in particular. In Section 3, we describe the proposed SERT and SST-ANN models in detail. In Section 4, we present the experimental results and analysis. Finally, in Section 5, we conclude the paper and discuss future directions of research.

2. Related work

In this section we provide a brief overview of the recent developments in deep learning models for sequential data analysis and spatio-temporal models for environmental monitoring, and also methods for handling missing data and adding interpretability to deep learning models applied to time series data.

2.1. Deep learning models for sequential data

Deep learning models, which are fundamentally built upon neural networks, have revolutionized the way complex applications, such as multi-modal learning, are approached (Barua et al., 2023). Among these models, Recurrent Neural Networks (RNNs) have been one of the most popular deep learning models for sequential data (Li et al., 2018). However, RNNs suffer from the vanishing gradient problem which makes them unable to learn long-term dependencies in the data. To address this problem, Long Short-Term Memory (LSTM) networks (Hochreiter and Schmidhuber, 1997) and Gated Recurrent Unit (GRU) networks (Cho et al., 2014) were introduced. These models have been applied to various tasks such as machine translation (Bahdanau et al., 2014), speech recognition (Graves et al., 2013), and time series forecasting (Lim and Zohren, 2021). Although these new variations of RNNs perform better by mitigating the risk of the vanishing gradient problem, they are still somewhat limited in modelling long-range dependencies due to the recursion used in their operations.

More recently, a new type of deep learning model named transformers was introduced (Vaswani et al., 2017). Transformers rely on the self-attention mechanism, which enables them to learn dependencies between input and output sequences without recursion, significantly reducing the maximum path length of network signals to the theoretical minimum of $\mathcal{O}(1)$ (Zhou et al., 2021). This capability is a significant

advantage of transformers over recurrent methods and has contributed to their superior performance. Furthermore, the original architecture of transformers includes an encoder and a decoder network. The encoder network is responsible for learning the representation of the input sequence, while the decoder network is responsible for generating the output sequence based on the learned representation. Models developed on the transformer architecture include BERT (Devlin et al., 2018), which uses only the encoder part, and GPT (Radford et al., 2018) which uses only the decoder part. These models have been applied to various tasks in natural language processing such as question answering (Devlin et al., 2018), text classification (Sun et al., 2019), and text summarization (Liu and Lapata, 2019). Overall, transformers have proven themselves to be more effective than recurrent based models in many applications, especially in natural language processing.

An essential step in leveraging transformer-based models involves obtaining high-quality, contextualized embeddings of the input data. This process is primarily facilitated by the encoder network, which is comprised of multi-head attention mechanisms and feed-forward layers within the transformer architecture. These embeddings capture the context of each element in the input data, reflecting their meaning in relation to the entire set of input tokens. After encoding, the embeddings are often passed through a linear layer to generate the final predictions. This approach has been validated by numerous studies, demonstrating its effectiveness across a variety of tasks. For instance, in grammatical error correction, transformer-based models have significantly outperformed conventional models, showcasing their advanced language processing capabilities (Alikaniotis and Raheja, 2019). Additionally, these models have proven to be extremely versatile, finding utility in domains beyond traditional NLP tasks. An example includes polymer property prediction, where a transformer model, referred to as TransPolymer, has demonstrated superior performance in accurately predicting properties across diverse datasets (Xu et al., 2022). These examples underscore the broad applicability and effectiveness of transformer-based models in capturing and utilizing complex patterns in data.

2.2. Deep learning models for spatio-temporal data

The application of deep learning models for spatio-temporal forecasting is not new. For example, Zhang et al. (2018) used an LSTM to forecast daily land surface temperature, Xiao et al. (2019) developed a deep learning model that utilizes convolutional long short-term memory as the building block for predicting sea surface temperature fields, and Mendil et al. (2022) introduced MCxM, a deep learning approach specifically designed for emergency pollution forecasting. However, the application of transformers to spatio-temporal forecasting is relatively new and challenging because transformers have been mainly developed in the field of natural language processing (NLP). Nonetheless researchers were inspired by the success of transformers in NLP and started adapting them to spatio-temporal forecasting which can be formulated as a sequence-to-sequence problem, where the input is a sequence of historical observations of multiple variables at different locations, and the output is a sequence of future predictions of the same variables at the same locations. A common approach for sequence-to-sequence modelling is to use an encoder–decoder architecture, where an encoder network maps the input sequence into a latent representation, and a decoder network generates the output sequence from the latent representation (Sutskever et al., 2014). Grigsby et al. (2021) used this idea to develop a new model called Spacetimeformer and applied it to traffic prediction and weather forecasting. However, to the best of our knowledge, the application of transformers to environmental monitoring is limited to the recent work by Yu et al. (2023) who used a transformer-based model for hourly $PM_{2.5}$ forecasting in Los Angeles.

2.3. Addressing missing values in modelling

As mentioned in the introduction, a major challenge in spatio-temporal forecasting in the environmental monitoring context is dealing with missing values. A common approach for dealing with the missing values is imputation (Van Buuren, 2018) before conducting any analysis. In time series modelling care needs to be taken to avoid introducing bias; last observation carried forward is a common approach. An alternative used in the literature is that of a Bayesian framework which enables defining a prior distribution over the missing values so that they can be inferred with the other unobserved parameters when fitting the models. For example, Shoari Nejad et al. (2022) proposed a Bayesian model (called VARICH) for spatio-temporal modelling of turbidity data with many missing values. However, this approach is computationally expensive and requires a large number of samples from the posterior distribution to obtain acceptable results and thus is not suitable for large spatio-temporal datasets.

To address the missing data problem in time series, Horn et al. (2020) proposed a novel approach to encode multivariate time series using set functions and introduced a new model called SeFT for classifying time series with irregularly sampled clinical data. More recently, Tipirneni and Reddy (2022) proposed a new transformer based model called STraTS that represents each observation as a triplet of the form (time, variable name, value). As opposed to SeFT, STraTS uses a learnable positional encoding and a Continuous Value Embedding (CVE) scheme that is a one-to-many feed-forward network. STraTS was developed to perform multivariate time series forecasting during the pre-training phase and classification as the final task on irregularly sampled clinical data. It has been shown to have higher accuracy than its predecessor, the SeFT model, when applied to the classification of clinical time series. Both SeFT and STraTS can handle missing values without requiring any imputation. To the best of our knowledge, none of these novel methods have been applied to environmental monitoring challenges; we adapt STraTS to a spatio-temporal setting.

2.4. Interpretability of deep learning models

Deep learning models are often considered as black-box models because they are hard to interpret (Buhrmester et al., 2021). However, in many applications, it is important to understand the model's decision making process (Du et al., 2019). Some authors have proposed methods to help interpret the results of deep learning models. For example, SeFT uses the attention mechanism in its architecture and the authors of the work showed that the attention weights can be used to gain insights into the importance of input data, including multiple variables. Similarly, Castangia et al. (2023) used a transformer-based model to forecast floods and employed attention maps extracted from the attention layers of the transformer model for interpretation. Tipirneni and Reddy (2022), inspired by Choi et al. (2016) and Zhang et al. (2020), took a different approach and proposed an interpretable version of the STraTS model, called STraTS-I, which uses an almost identical structure to their STraTS model but instead uses encoded inputs directly to the output layer as opposed to STraTS that uses the contextualized inputs to the output layer. Their approach allows for the calculation of a contribution score for each input observation towards the prediction, achieved by multiplying the encoded input, attention weights, and output layer weights. This modification aims to compensate accuracy for interpretability while both models have similar computational complexity. We follow a similar simplification routine in the creation of our SST-ANN approach explained in Section 3.3.

3. Proposed methods

In this section we first define the problem followed by the details of the general model architectures that we use to build SERT and SST-ANN to address the problem. We then introduce a modification for encoding location information in the models' input data. Finally, we describe the masked loss function that we use for training the models.

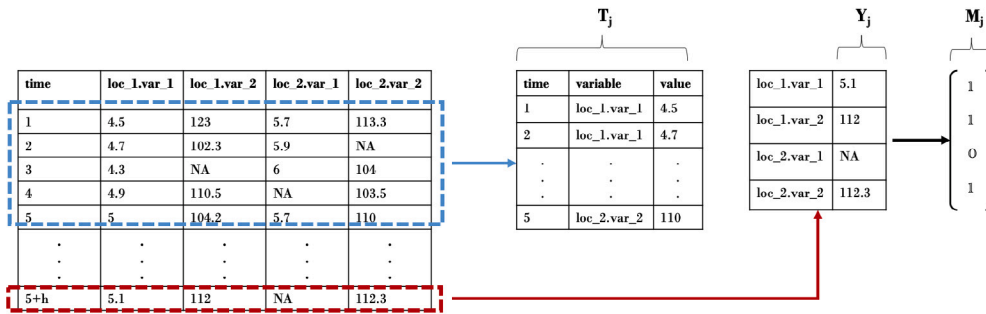


Fig. 2. An example of a sample of a spatiotemporal dataset to be used for training our proposed models where h is the desired forecast horizon.

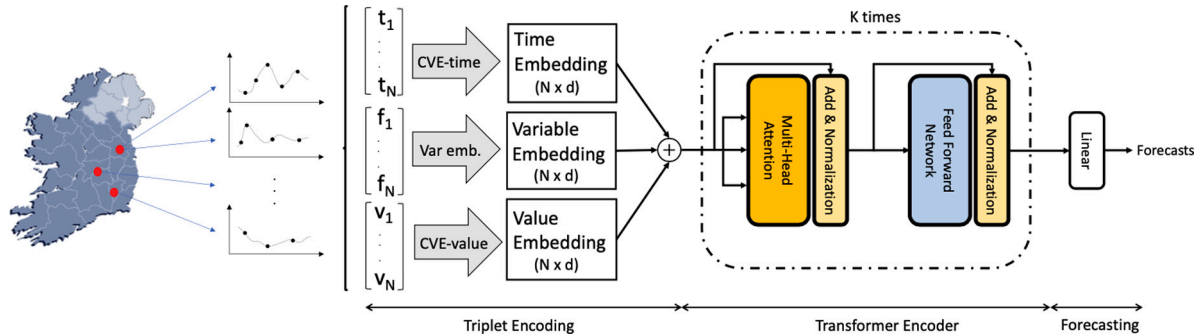


Fig. 3. Schematic diagram of the SERT model.

3.1. Problem definition

We have a dataset $D = \{(T_j, Y_j, M_j)\}_{j=1}^J$ where T_j is a multivariate time series consisting of triples (time, variable, value) written as $\{(t_i, f_i, v_i), i = 1, \dots, N\}$. Y_j is the values (and associated variables) for a future time horizon for which we want the model to forecast. M_j is a binary vector indicating whether each of element in Y_j is observed for sample j in the dataset. M_j is used in the loss function (see Section 3.5 for details) for masking the unobserved values in the forecast window. A schematic of a sample from the dataset is shown in Fig. 2. The goal is to learn a model G that maps T_j to Y_j , i.e., $G(T_j) = Y_j$, without imputing the missing values in T_j or aligning the time series.

3.2. SERT

We first describe the data encoding scheme and then the model architecture including an encoder network and a linear layer. The schematic diagram of the model is shown in Fig. 3. We will slightly modify this structure to show an alternative approach for encoding the location information in Section 3.4.

Data Encoding Scheme The input data to our model is a multivariate time series dataset T that consists of N time series, each of which is associated with a time series variable f which is a sequence of observation values v 's. Accordingly, an individual data point i is represented as a triplet (t_i, f_i, v_i) where t_i is time, f_i is the variable name and v_i is the measured value of the data point. To use the triplet in the model, we encode each component into an embedding and then add the embeddings together. Let $e_i^f \in R^d$ be the embedding of the variable name f_i which can be encoded similar to words using a lookup table, $e_i^t \in R^d$ be the embedding of the time index t which can be encoded using a continuous value embedding (CVE) scheme which is a one-to-many feed forward neural network (Tipirneni and Reddy, 2022) and $e_i^v \in R^d$ be the embedding of the value v_i which can also be encoded using the CVE. The embedding of the triplet i is then defined as $e_i = e_i^t + e_i^f + e_i^v$. The size of the embedding vector d is a hyperparameter of the model.

Encoder Network Similar to the well-known BERT model (Devlin et al., 2018), the main component of our model is the encoder part of the transformer model introduced by Vaswani et al. (2017). Since transformers have become very common, we omit the details of the architecture and refer the reader to Vaswani et al. (2017) for the full description. Intuitively, we can think of the encoder network as layers that take the triplet embeddings of the input data and transform them into contextualized embeddings that capture the long-range dependencies within a time series as well as cross dependencies between different time series.

Linear Layer After obtaining the contextualized embeddings of the input data using the encoder network, we then flatten the embeddings and apply a linear layer to them to generate the predictions. The linear layer is a feed-forward network with a single hidden layer and a ReLU activation function.

3.3. SST-ANN

The SST-ANN model is a simplified version of the SERT model that consists of only the triplet encoding and a linear layer to the output with no transformer structure in between. SST-ANN first encodes the input data using the triplet encoding scheme and then uses the embeddings as the input to single layer feed forward network to generate the predictions. Since there is no transformer structure in between, the SST-ANN model is much faster than the SERT model and, using the embeddings and the weights of the linear layer, we can compute a contribution score for each observation to the final prediction. This is useful for interpretability and variable importance analysis. More formally the output of the model can be expressed as follows:

$$\hat{y}_k = \sum_{i=1}^N c_i + b, \text{ with } c_i = W_{ik}^T \cdot e_i, \tag{1}$$

where \hat{y}_k is the prediction of variable k , c_i is the contribution of the triplet (t_i, f_i, v_i) to the prediction, N is the number of observations in the input sample, b is the bias term, e_i are the embeddings of the triplet

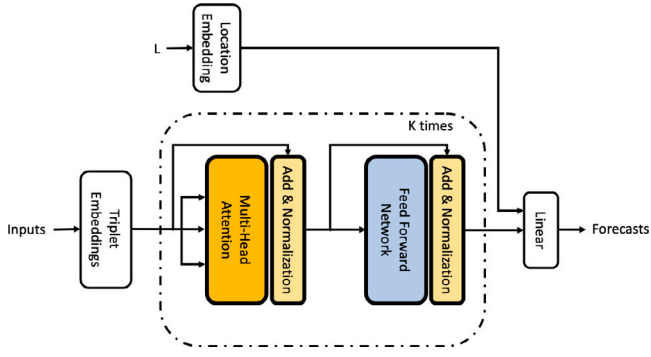


Fig. 4. Schematic diagram of the SERT model with a separated location embedding layer.

and W_{ik} is the vector of output weights associated with the embedding e_i and the target variable k .

Using the contribution scores, we can define a variable importance index. We first calculate the average contribution value of all observations belonging to the same variable as the average contribution of that variable. This calculation can be performed for a single sample to gain insights into the importance of the predictor variable for a specific target prediction, or for multiple samples used in multiple predictions to obtain an overall understanding of the predictor variable's importance in general. Next, we compute the importance of each variable by normalizing the absolute value of the average contribution values for the variables. More formally, we can express this as:

$$I_k = \frac{|\bar{c}_k|}{\sum_{k=1}^K |\bar{c}_k|} \times 100 \quad (2)$$

where I_k is the importance (in percentage) and \bar{c}_k is the average contribution value of the variable k .

3.4. Location encoding

We consider two different approaches to encode the location information in the input data. The first approach is to encode the location information together with the variable name f_i in the triplet encoding scheme. For example, our naming scheme for the variables can be $f_i = \{location\}_i \cdot \{variable\}_i$ where $\{location\}_i$ is the location of the time series i and $\{variable\}_i$ is the variable name of the time series i (e.g. Tolka.Turbidity). This approach uses the exact same architecture explained in Section 3.2. The second approach is to encode the location information separately from the variable name. This way the input time series are all assumed to arise from the same location and the location embedding is concatenated to the contextualized embeddings before the linear layer is used for prediction. In this approach, we need to structure the dataset such that all the time series from the same location are grouped together. Formally, we can define the grouped dataset as $D' = \{(T_j^L, Y_j^L, M_j^L)\}_{j=1}^S$ where $L \in S$ is the location and S is the set of all locations. This approach needs a minimal modification to the previously described architecture and its schematic diagram is shown in Fig. 4.

The first approach is similar to how the Spacetimeformer model (Grigsby et al., 2021) encodes the location information while the second approach is similar to how the STraTS model (Tipirneni and Reddy, 2022) encodes the non-temporal information (patient demographics in their work).

3.5. Masked loss function

We use a masked Mean Squared Error (MSE) loss function to train our models. Masking in the loss function is used to handle the missing

values in the output data. The masked MSE loss function is defined as follows:

$$\mathcal{L} = \frac{1}{J} \sum_{j=1}^J \sum_{k=1}^K m_k^j (\hat{y}_k^j - y_k^j)^2 \quad (3)$$

where y_k^j is the ground truth, \hat{y}_k^j is the predicted value, and m_k^j is the mask value of the target variable k in sample j .

4. Experiments

In this section, we will describe the experiments we conducted to evaluate the performance of our proposed models. We evaluated our models using both a simulated dataset and a real-world dataset. The primary objective of the simulation experiment is to investigate how the models perform under different sparsity levels. The real-world experiment aimed to serve as a proof of concept for the models' ability to conduct spatiotemporal multi-step ahead forecasting in real-world scenarios.

To assess the effectiveness of our models, we compared them with a baseline Naive forecaster model, which simply uses the present hour's observation as the next hour forecast. We also compare our models against a VAR model which is common for multivariate time series modelling, the LSTM model and the STraTS model. Since the Naive forecaster and LSTM model cannot handle missing values, we first imputed these values using a forward filling method (Van Buuren, 2018, p. 16). The source code for our experiments is available at <https://github.com/Aminsn/SERT2023>.

4.1. Sparsity analysis

We first simulate a dataset that consists of 16 time series each with 40,000 observations. Then, we use this dataset to compare the predictive performance of the previously mentioned models. The total length of the time series matters for the models' performance in terms of the number of observations, which could affect both deep learning and classical models differently. Deep learning models are believed to require more observations to perform well. Here, we aim to replicate a scenario typical of multiple sensor data, which usually encompasses large, high-frequency datasets.

We denote $Y_t \in R^{16}$ as the vector of observations at time t generated from the following process:

$$Y_t = 2 + 0.4Y_{t-1} + X_t + s_t, \quad s_t \sim MVN(0, \Sigma) \quad (4)$$

where s_t is spatial random effect with mean zero and variance-covariance matrix Σ that is generated with:

$$\Sigma = U \cdot U^T, \quad U \in R^{16 \times 16}, \quad U_{i,j} \sim Uniform(-1, 1), \quad (5)$$

and $X_t \in R^{16}$ is the vector of temporal effects generated as:

$$\begin{aligned} X_t^T = & [10 \sin(p_1 t), \cos(p_2 t), p_3 t, -p_3 t \\ & + 10 \sin(p_1 t), 5 \sin(p_2 t), 12 \cos(p_2 t), 7 \sin(p_2 t), 8 \cos(p_2 t), \\ & 2 \sin(p_2 t), 3 \cos(p_2 t), 12 \sin(p_2 t), 18 \cos(p_2 t), 4 \sin(p_2 t), \\ & 15 \cos(p_2 t), 11 \sin(p_2 t), 10 \cos(p_2 t)] \end{aligned} \quad (6)$$

where $p_1 = 0.005$, $p_2 = 0.0005$ and $p_3 = 0.002$. The simulated data includes multiple time series with various characteristics including seasonality, trend, a combination of seasonality and trend, and different levels of noise. Fig. 8 in the appendix demonstrates these time series.

We use the first 37,000 time steps to train the models and the remaining 3000 time steps to evaluate the performance of them. We used the model structure shown in Fig. 3 to train our proposed models. We trained all models for 1 step ahead forecasting using the previous 10 h of observations. The input length acts as a hyperparameter of the model, akin to the order in classical autoregressive models. In this example, we have chosen an input length of 10 for convenience across

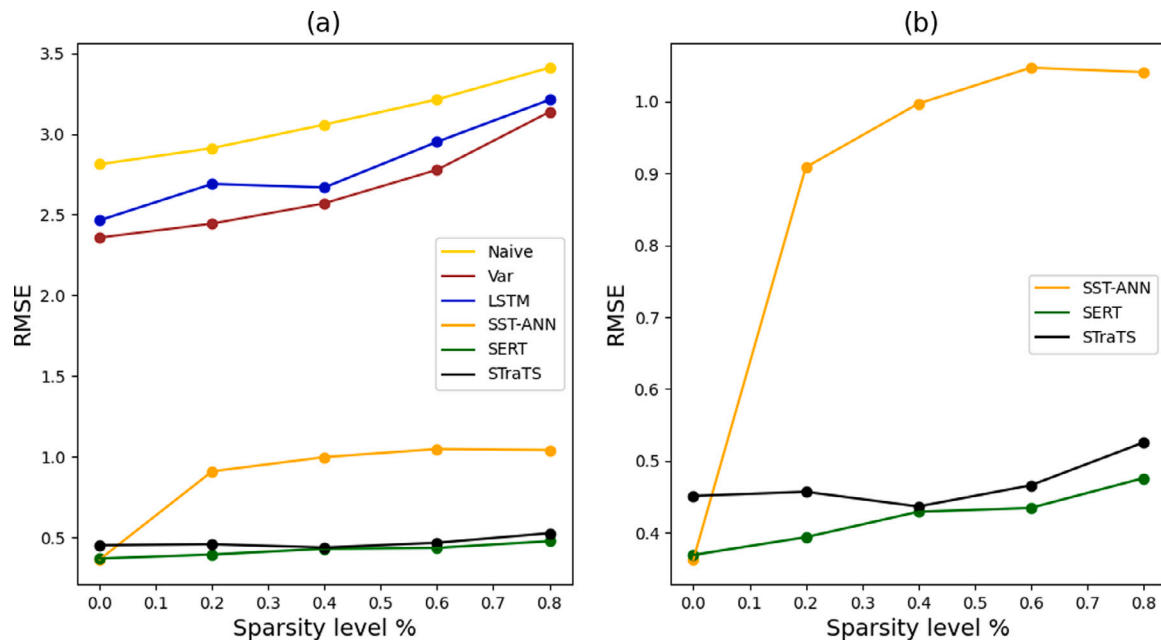


Fig. 5. (a) Performance of all the models used on the simulated dataset. (b) Comparison of SST-ANN, STraTS, and SERT zoomed in.

Table 1

RMSE of the models for 7 h ahead forecasting of the 7 environmental variables in Dublin bay.

| Model | Dissolved oxygen | Precipitation | Salinity | Temperature | Turbidity | Water level | Wind speed |
|----------------|------------------|---------------|-------------|-------------|-------------|-------------|-------------|
| Naive | 0.64 | 1.19 | 1.22 | 0.21 | 0.93 | 1.87 | 0.8 |
| LSTM | 0.65 | 0.89 | 0.8 | 0.41 | 0.81 | 0.5 | 0.82 |
| STraTS | 0.53 | 0.88 | 0.7 | 0.18 | 0.72 | 0.38 | 0.76 |
| SERT (ours) | 0.49 | 0.88 | 0.67 | 0.18 | 0.72 | 0.4 | 0.73 |
| SST-ANN (ours) | 0.51 | 0.88 | 0.73 | 0.45 | 0.88 | 0.63 | 0.79 |

all models. The appropriate length, or number of lags, for prediction is more a characteristic of the data than of the models themselves. Naturally, a model that utilizes all relevant information is anticipated to perform better. For a fair comparison, we ensure that all models are fed the same quantity of input data. We consider five different sparsity levels to fit the models. Accordingly, we remove $n\%$ of the observations randomly for $n = \{0\%, 20\%, 40\%, 60\%, 80\%\}$. We use the root mean squared error (RMSE) as the evaluation metric. The results are presented in Fig. 5.

4.2. Real dataset; environmental monitoring in Dublin bay, Ireland

The dataset includes hourly measurements from 2017-01-01 to 2021-12-31. An example of the data is shown in Fig. 1. The locations of the data are shown in Fig. 6. We use the data from the first four years to train the models and the last year to evaluate their performance. We train the models using the previous 10 h of observations as the input and forecasting seven hours ahead (forecasting only 1 or 2 h ahead is a relatively trivial task, and forecasting > 12 h ahead reduces the performance of models that cannot account for non-seasonality). We use the same evaluation metric as in the simulated dataset. We tried both location encoding approaches (explained in Section 3.4) with our proposed models on the real-world dataset and found that the second approach, depicted in Fig. 4, performed better and here we only report the results of the superior approach. The results are presented in Table 1.

We utilized the same computational resources (a single NVIDIA P100 GPU) to train the models. The specifications and speed performance details of the models are reported in the Table 2.

Table 2

Computational specifications of the fitted models.

| Model | Specifications | Sec/epoch |
|---------|-------------------------------|-----------|
| Naive | - | - |
| LSTM | units = 60 | 2 |
| STraTS | d = 60, # of heads = 6, k = 6 | 230 |
| SERT | d = 60, # of heads = 6, k = 6 | 40 |
| SST-ANN | d = 60 | 9 |

4.3. Interpretability of the SST-ANN model

The results of the variable importance for the real-world dataset experiment are presented in Fig. 7. The sign of the average contribution score \bar{c}_j (explained in Section 3.3) is multiplied by the importance index I_j to give an insight into the direction of the average contribution of the predictor variable to the prediction of the target variable. In this example, we only consider the contributions of water level, temperature, wind speed and precipitation to predictions of turbidity, dissolved oxygen and salinity, since we know that the former variables could affect the latter variables but not vice versa. According to the results, temperature followed by precipitation are the most important variables in predicting the target variables. However, one should take caution and not interpret these numbers as they would in a causal model. Instead, the application of variable importance in SST-ANN should be more closely compared to the concept of variable importance as used in a random forest model (for example, see Archer and Kimes (2008)).

5. Conclusion

In this paper, we proposed two novel models for spatio-temporal forecasting called SERT and SST-ANN. SERT is a transformer-based



Fig. 6. Buoys measuring various environmental variables in Dublin Bay. The 4 main environmental variables (Turbidity, Salinity, Dissolved Oxygen, Temperature) are available at the number buoys. The weather measurements (Rainfall, Wind Speed) come from Dublin Airport. The water level variable comes from Howth Harbour.




Fig. 7. Variable importance of the selected variables in predicting the target variables.

model, meaning it utilizes the encoder component of the original transformer architecture, albeit with certain modifications. Specifically, it substitutes the conventional positional encoding with a triplet encoding schema. This adjustment enables SERT to efficiently process sparse spatio-temporal data. SST-ANN, on the other hand, is a simple ANN model combined with triplet encoding of the input data.

Furthermore, we showed that STraTS, a model originally developed for sparse and irregularly sampled clinical time series classification, can be used for spatio-temporal forecasting, especially when missing values are present in the data. SERT’s architecture closely resembles that of STraTS, but there is a key difference: STraTS employs a fusion self-attention layer prior to its prediction layer, while in SERT, this has been

substituted with a simpler flattening layer. Although their prediction accuracies are similar, with SERT being slightly more accurate, this adjustment makes SERT faster and better suited for environmental forecasting.

Our proposed approaches do not require aggregation or missing value imputation techniques, and avoid the problems introduced by such methods. We evaluated the performance of the proposed models on a simulated dataset with varying levels of sparsity and showed that in general increasing sparsity has a negative effect in the performance of all the models, but SERT followed by STraTS and SST-ANN are more robust to the increase in sparsity. We also evaluated the performance of the proposed models on a real-world dataset of environmental variables in Dublin Bay, Ireland. The results indicate that SERT outperformed the other models in 7-hour ahead forecasting for 6 out of the 7 variables, with 3 of them being on par with STraTS while demonstrating significantly faster performance. We demonstrated that our proposed models are capable of forecasting various environmental variables using sparse data. Given the prevalence of issues such as sensor malfunctioning leading to missing values, our models remain applicable for utilizing incomplete information to inform decision-makers about future events. For instance, forecasting a surge in water levels could be crucial for implementing preventive measures. We then showed how SST-ANN can be used to interpret the predictions of the model by calculating and using the contribution score of the input data to develop an importance index using the average contribution scores. This approach can improve our understanding of the model’s decision-making process, similar to the utilization of other interpretable machine learning approaches in geoscience and environmental analysis (see, for example, Reichstein et al. (2019)). We demonstrated the performance of our proposed models on long multivariate time series data that are rich in time but poor in location. However, these models can also be effectively applied to other sampling scenarios, such as short time series from many locations or long time series from several locations. It is important to note, however, that an increase in the number of inputs for analysis leads to higher computational costs. Specifically, the self-attention mechanism in transformers exhibits a time and space complexity that is quadratic with respect to the length of the input sequence (Kees et al., 2023). Despite this, computations can be efficiently parallelized on modern

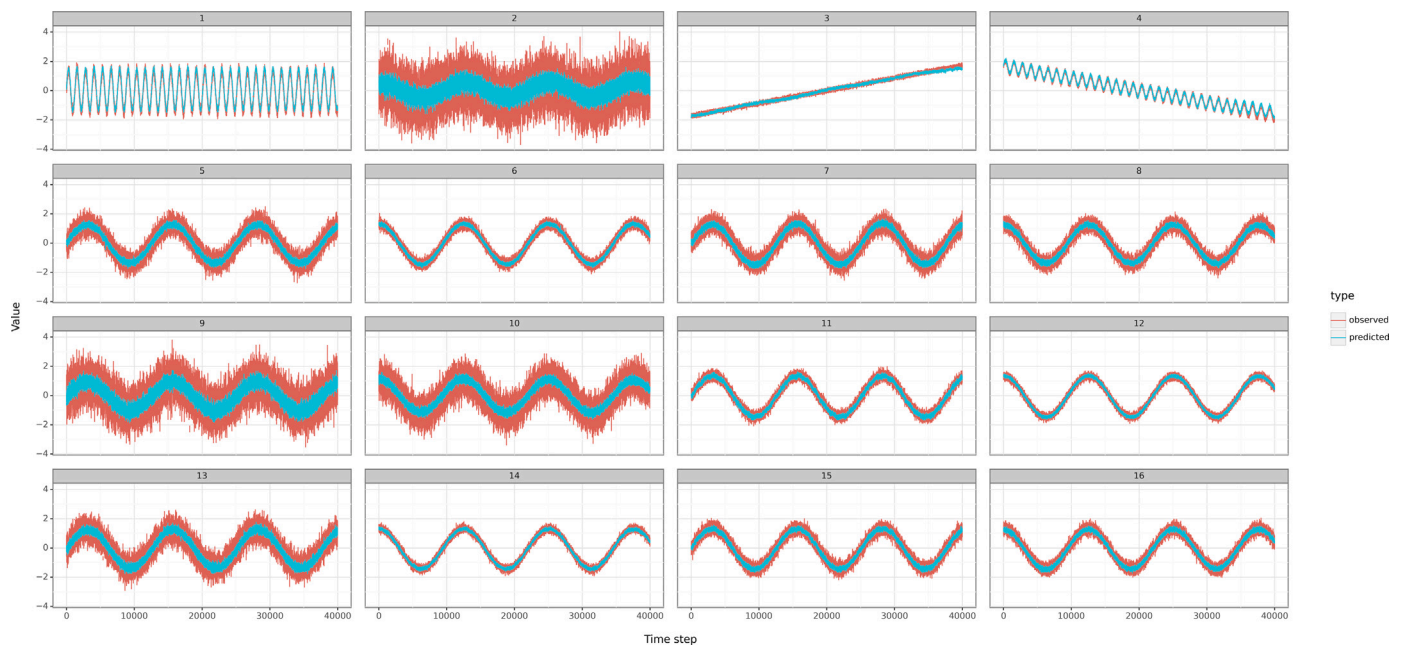


Fig. 8. Multivariate simulated time series data along with the predicted values from SERT for one step ahead forecasting, covering both training and test periods.

GPUs, offering a significant advantage over recurrent methods, which require sequential processing.

We introduced two different methods to encode the location information in our proposed models, including encoding the time series variable name with the location name simultaneously and encoding the location name separately. The first location encoding scheme accounts for the interactions among variables from different locations, thereby attempting to learn a more complex relationship in the data. In contrast, the second approach assumes that relationships among variables do not vary across locations. Depending on the true nature of the data-generating process, as well as the data quality and quantity, one approach may perform better than the other. For example, the first approach is typically expected to require more data because it aims to learn a more complicated relationship. Meanwhile, the second approach might perform poorly in scenarios where significant interaction exists between variables from different locations. However, neither of these methods takes into account the distance between the locations, which is a limitation of our work. We believe future research should focus on incorporating this information into the models, as it has the potential to improve forecasting performance and be utilized for spatiotemporal interpolation tasks.

Moreover, future work could consider incorporating uncertainty estimation into the proposed models. Enhancing SERT and STT-ANN to offer prediction with confidence intervals could present significant advantages. This enhancement can be achieved through model-agnostic methods like conformal prediction (Fontana et al., 2023) or by modifying the model's structure and loss function to predict data distribution parameters rather than mere point predictions. Such adjustments can significantly improve the sustainability of geoscientific explorations under conditions of uncertainty. These advancements allow for a more nuanced understanding of the reliability of predictions and associated risks, thereby fostering more informed decision-making processes. This approach enhances the scope of field-scale applicability of our proposed models across geographic spectra by providing a framework for handling the inherent uncertainties in geoscientific data. This leads to more sustainable exploration strategies that can adapt to diverse environmental conditions and data variability.

Computer code availability

Software name: SERT

Year first available: 2023

Programming language: Python

Software required: Python 3.7 or higher.

The source codes are available for downloading at the link: <https://github.com/Aminsn/SERT2023>

CRediT authorship contribution statement

Amin Shoari Nejad: Conceptualization, Formal analysis, Investigation, Methodology, Resources, Software, Validation, Visualization, Writing – original draft. **Rocío Alaiz-Rodríguez:** Supervision. **Gerard D. McCarthy:** Supervision. **Brian Kelleher:** Data curation, Funding acquisition. **Anthony Grey:** Data curation, Writing – review & editing. **Andrew Parnell:** Supervision.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Amin Shoari Nejad reports financial support was provided by Maynooth University.

Data availability

I have shared the link to the data and code used in the paper.

Acknowledgements

We would like to express our sincere gratitude to Dublin Port Company for providing us with the real dataset. This work was supported by an SFI Investigator award (16/IA/4520).

Appendix

See Fig. 8.

References

- Alikaniotis, D., Raheja, V., 2019. The unreasonable effectiveness of transformer language models in grammatical error correction. arXiv preprint arXiv:1906.01733.
- Archer, K.J., Kimes, R.V., 2008. Empirical characterization of random forest variable importance measures. *Comput. Statist. Data Anal.* 52 (4), 2249–2260.
- Bahdanau, D., Cho, K., Bengio, Y., 2014. Neural machine translation by jointly learning to align and translate. arXiv preprint arXiv:1409.0473.
- Barua, A., Ahmed, M.U., Begum, S., 2023. A systematic literature review on multimodal machine learning: Applications, challenges, gaps and future directions. *IEEE Access*.
- Buhmester, V., Münch, D., Arens, M., 2021. Analysis of explainers of black box deep neural networks for computer vision: A survey. *Mach. Learn. Knowl. Extr.* 3 (4), 966–989.
- Castangia, M., Grajales, L.M.M., Aliberti, A., Rossi, C., Macii, A., Macii, E., Patti, E., 2023. Transformer neural networks for interpretable flood forecasting. *Environ. Model. Softw.* 160, 105581.
- Cho, K., Van Merriënboer, B., Bahdanau, D., Bengio, Y., 2014. On the properties of neural machine translation: Encoder-decoder approaches. arXiv preprint arXiv:1409.1259.
- Choi, E., Bahadori, M.T., Sun, J., Kulas, J., Schuetz, A., Stewart, W., 2016. Retain: An interpretable predictive model for healthcare using reverse time attention mechanism. *Adv. Neural Inf. Process. Syst.* 29.
- Deng, L., Chen, K., Liu, Z., Wu, B., Chen, Z., He, S., 2022. Spatiotemporal variation evaluation of water quality in middle and lower Han River, China. *Sci. Rep.* 12 (1), 1–14.
- Devlin, J., Chang, M.W., Lee, K., Toutanova, K., 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al., 2020. An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929.
- Du, M., Liu, N., Hu, X., 2019. Techniques for interpretable machine learning. *Commun. ACM* 63 (1), 68–77.
- Durbin, J., Koopman, S.J., 2012. *Time series analysis by state space methods*, vol. 38. OUP Oxford.
- Fontana, M., Zeni, G., Vantini, S., 2023. Conformal prediction: a unified review of theory and new challenges. *Bernoulli* 29 (1), 1–23.
- Graves, A., Mohamed, A.R., Hinton, G., 2013. Speech recognition with deep recurrent neural networks. In: 2013 IEEE International Conference on Acoustics, Speech and Signal Processing. Ieee, pp. 6645–6649.
- Grigsby, J., Wang, Z., Qi, Y., 2021. Long-range transformers for dynamic spatiotemporal forecasting. arXiv preprint arXiv:2109.12218.
- Hamdi, A., Shaban, K., Erradi, A., Mohamed, A., Rumi, S.K., Salim, F.D., 2022. Spatiotemporal data mining: a survey on challenges and open problems. *Artif. Intell. Rev.* 1–48.
- Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. *Neural Comput.* 9 (8), 1735–1780.
- Horn, M., Moor, M., Bock, C., Rieck, B., Borgwardt, K., 2020. Set functions for time series. In: *International Conference on Machine Learning*. PMLR, pp. 4353–4363.
- Keles, F.D., Wijewardena, P.M., Hegde, C., 2023. On the computational complexity of self-attention. In: *International Conference on Algorithmic Learning Theory*. PMLR, pp. 597–619.
- Li, S., Li, W., Cook, C., Zhu, C., Gao, Y., 2018. Independently recurrent neural network (indrnn): Building a longer and deeper rnn. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 5457–5466.
- Lim, B., Zohren, S., 2021. Time-series forecasting with deep learning: a survey. *Phil. Trans. R. Soc. A* 379 (2194), 20200209.
- Liu, Y., Lapata, M., 2019. Text summarization with pretrained encoders. arXiv preprint arXiv:1908.08345.
- Mendil, M., Leirens, S., Armand, P., Duchenne, C., 2022. Hazardous atmospheric dispersion in urban areas: A deep learning approach for emergency pollution forecast. *Environ. Model. Softw.* 152, 105387.
- Radford, A., Narasimhan, K., Salimans, T., Sutskever, I., et al., 2018. Improving language understanding by generative pre-training.
- Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., Carvalhais, N., Prabhat, F., 2019. Deep learning and process understanding for data-driven earth system science. *Nature* 566 (7743), 195–204.
- Ren, W., Han, M., 2014. Modeling of multivariate time series using variable selection and Gaussian process. In: *Proceedings of the 33rd Chinese Control Conference*. IEEE, pp. 5071–5074.
- Shoari Nejad, A., McCarthy, G.D., Kelleher, B., Grey, A., Parnell, A., 2022. Vector time series modelling of turbidity in Dublin Bay. arXiv preprint arXiv:2209.06880.
- Sims, C.A., 1980. Macroeconomics and reality. *Econometrica* 1–48.
- Stekhoven, D.J., Bühlmann, P., 2012. MissForest—non-parametric missing value imputation for mixed-type data. *Bioinformatics* 28 (1), 112–118.
- Sun, C., Qiu, X., Xu, Y., Huang, X., 2019. How to fine-tune bert for text classification? In: *Chinese Computational Linguistics: 18th China National Conference, CCL 2019, Kunming, China, October 18–20, 2019, Proceedings 18*. Springer, pp. 194–206.
- Sutskever, I., Vinyals, O., Le, Q.V., 2014. Sequence to sequence learning with neural networks. *Adv. Neural Inf. Process. Syst.* 27.
- Tipirneni, S., Reddy, C.K., 2022. Self-supervised transformer for sparse and irregularly sampled multivariate clinical time-series. *TKDD, ACM Trans. Knowl. Discov. Data TKDD*, 16 (6).1–17.
- Van Buuren, S., 2018. *Flexible imputation of missing data*. CRC Press.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I., 2017. Attention is all you need. *Adv. Neural Inf. Process. Syst.* 30.
- Wikle, C.K., 2015. Modern perspectives on statistics for spatio-temporal data. *Wiley Interdiscip. Rev. Comput. Stat.* 7 (1), 86–98.
- Xiao, C., Chen, N., Hu, C., Wang, K., Xu, Z., Cai, Y., Xu, L., Chen, Z., Gong, J., 2019. A spatiotemporal deep learning model for sea surface temperature field prediction using time-series satellite data. *Environ. Model. Softw.* 120, 104502.
- Xu, C., Wang, Y., Farimani, A., 2022. TransPolymer: a transformer-based language model for polymer property predictions. arXiv.
- Yu, M., Masrur, A., Blaszczyk-Boxe, C., 2023. Predicting hourly PM_{2.5} concentrations in wildfire-prone areas using a SpatioTemporal transformer model. *Sci. Total Environ.* 860, 160446.
- Zhang, X., Qian, B., Cao, S., Li, Y., Chen, H., Zheng, Y., Davidson, I., 2020. INPREM: An interpretable and trustworthy predictive model for healthcare. In: *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. pp. 450–460.
- Zhang, X., Zhang, Q., Zhang, G., Nie, Z., Gui, Z., Que, H., 2018. A novel hybrid data-driven model for daily land surface temperature forecasting using long short-term memory neural network based on ensemble empirical mode decomposition. *Int. J. Environ. Res. Public Health* 15 (5), 1032.
- Zhou, H., Zhang, S., Peng, J., Zhang, S., Li, J., Xiong, H., Zhang, W., 2021. Informer: Beyond efficient transformer for long sequence time-series forecasting. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. 35, (12), pp. 11106–11115.