

# Paying Attention to Astronomical Transients: Introducing the Time-series Transformer for Photometric Classification

Tarek Allam Jr.,<sup>1</sup>\* Jason D. McEwen<sup>1</sup>

<sup>1</sup>Mullard Space Science Laboratory, University College London, Holmbury St Mary, Dorking, Surrey RH5 6NT, UK

Accepted XXX. Received YYY; in original form ZZZ

## ABSTRACT

Future surveys such as the Legacy Survey of Space and Time (LSST) of the Vera C. Rubin Observatory will observe an order of magnitude more astrophysical transient events than any previous survey before. With this deluge of photometric data, it will be impossible for all such events to be classified by humans alone. Recent efforts have sought to leverage machine learning methods to tackle the challenge of astronomical transient classification, with ever improving success. Transformers are a recently developed deep learning architecture, first proposed for natural language processing, that have shown a great deal of recent success. In this work we develop a new transformer architecture, which uses multi-head self attention at its core, for general multi-variate time-series data. Furthermore, the proposed time-series transformer architecture supports the inclusion of an arbitrary number of additional features, while also offering interpretability. We apply the time-series transformer to the task of photometric classification, minimising the reliance of expert domain knowledge for feature selection, while achieving results comparable to state-of-the-art photometric classification methods. We achieve a logarithmic-loss of 0.507 on imbalanced data in a representative setting using data from the Photometric LSST Astronomical Time-Series Classification Challenge (PLAsTiCC). Moreover, we achieve a micro-averaged receiver operating characteristic area under curve of 0.98 and micro-averaged precision-recall area under curve of 0.87.

**Key words:** machine learning - software - data methods - time-series - transients - supernovae

## 1 INTRODUCTION

The Legacy Survey of Space and Time (LSST) of the Vera C. Rubin Observatory (Ivezić et al. 2019) will set a new precedent in astronomical surveys, expecting to produce an average of 10 million transient event alerts per night. Machine learning methods are thus essential in order to handle the sheer volume of data that will come from the LSST. Since limited resources are available for spectroscopic follow-up of observations, accurate photometric classification of astronomical transient events will be increasingly critical for subsequent scientific analyses.

There is a wide range of science that comes from analyses of astronomical transients, with one particular area of focus being the analysis of Type Ia Supernova. Type Ia Supernovae have been an important tool for cosmologists for many years, serving as a proxy for distance measure in the Universe and shedding light on the expansion rate of the Universe (Riess et al. 1998; Perlmutter et al. 1999). Observing Type Ia Supernova at ever increasing redshift helps to constrain cosmological parameters and theories of dark energy. Thus, accurate classification of Type Ia Supernova from the stream of alerts has profound consequences.

Over the last decade a plethora of photometric classification algorithms have been developed. Many of them stemming from the fruitful Supernova Photometric Classification Challenge (Kessler et al.

2010, SNPhotCC) in 2010 that focused on photometric classification of Supernovae only; and more recently the Photometric LSST Astronomical Time-Series Classification Challenge (Hložek et al. 2020, PLAsTiCC) in 2018, which included a variety of different astronomical transient events among its classes.

Several challenges arise when observing photometrically; SNPhotCC and PLAsTiCC tried to simulate such conditions in terms of photometric sampling linked to the telescope cadence, as well as the distribution of classes one expects to observe. When creating such a simulated dataset, realistic distribution of classes is of great importance as often the training data available to astronomers is not of the same distribution one would observe through a real survey. This is due to Malmquist Bias (Butkevich et al. 2005), which is caused by the inherent bias towards observing brighter and closer objects when observing the night sky. As a consequence, training datasets are skewed to have more objects that are closer in distance, lower in redshift, and brighter in luminosity. In addition, the usefulness of observations of Type Ia Supernova has induced a bias towards spectroscopic follow-up of these events, resulting in vastly imbalanced training datasets that have a large number of Type Ia Supernova samples compared to other objects. The resulting training sets are therefore typically imbalanced and non-representative of the test sets that one might observe. These issues present a major challenge when developing classifiers. Several methods have been proposed to address the problems of non-representativity and class imbalance.

\* E-mail: tarek.allam.10@ucl.ac.uk

Early attempts that applied machine learning methods to the SNPhotCC dataset can be found in [Karpenka et al. \(2013\)](#) using neural networks, in [Ishida & de Souza \(2013\)](#) using kernel PCA with nearest neighbours, as well as methods found in [Lochner et al. \(2016\)](#) which compared a variety of techniques with impressive results on representative training data. Another successful approach can be found in [Boone \(2019\)](#) which was able to specifically extend the boosted-decision-tree (BDT) method in [Lochner et al. \(2016\)](#) by achieving good performance even in the non-representative training set domain. This work used BDTs coupled with data augmentation using Gaussian processes to achieve a weighted logarithmic loss ([Malz et al. 2019](#)) of 0.68 in the PLAsTiCC competition ([Hložek et al. 2020](#)) and 0.649 in a revised model following the close of the competition. However, one drawback with many of these methods is the reliance of the *human-in-the-loop*, where well crafted feature engineering plays an important role in achieving excellent scores. With few exceptions, such as the approaches of [Lochner et al. \(2016\)](#) and [Varughese et al. \(2015\)](#) that used wavelet features, many traditional machine learning approaches for photometric classification are model dependent, relying on prior domain specific information about the light curves.

More recently, there have been attempts to apply deep learning to minimise the laborious task of feature selection and in some cases input raw time-series information only. Work by [Brunel et al. \(2019\)](#) used an Inception-V3 ([Szegedy et al. 2015](#)) inspired convolutional neural network (CNN) and earlier work by [Charnock & Moss \(2017\)](#) used a long-short-term-memory (LSTM) recurrent neural network (RNN) for Supernovae classification. [Möller & de Boissière \(2020\)](#) also achieve good results building upon the success of RNNs. Extending to the general transient case and utilising an alternative RNN architecture, gated reticular units (GRUs), work by [Muthukrishna et al. \(2019\)](#) with RAPIDS showcased the impressive results one could achieve by using the latest methods borrowed from the domain of sequence modelling and natural language processing (NLP). While these deep learning methods have been shown to yield excellent results, both RNNs and CNNs have several limitations when it comes to dealing with time-series data.

RNNs tend to struggle with maintaining context over large sequences and from the unstable gradients problem. When an input sequence becomes long, the probability of maintaining the context of one input to another decreases significantly with the distance from that input ([Madsen 2019](#)). The shorter the paths between any set of positions in the input and output sequence, the easier it is to learn long range dependencies ([Hochreiter et al. 2001](#)). Note that the *maximum path length* of an RNN is given by the length of the most direct path between the first encoder input and the last decoder output ([Grosse & Ba 2019](#)). Another problem faced by the RNN family is the inherently sequential structure, making parallelisable computation difficult as each input point needs to be processed one after the other, resulting in a computational cost of  $O(n)$ , where  $n$  is the sequence length ([Vaswani et al. 2017](#)).

CNNs overcome these problems, to some extent, with trivial parallelism across layers and, with the use of the dilated convolution, distance relations can become an  $O(\log n)$  operation, allowing for processing of larger input sequences ([Oord et al. 2016](#)). However, CNNs are known to be computationally expensive with a complexity per layer given by  $O(w \cdot n \cdot d^2)$ , where  $w$  is the kernel window size and  $d$  the representational dimensionality ([Vaswani et al. 2017](#)). For contrast, RNNs have complexity per layer  $O(n \cdot d^2)$ .

Self-attention mechanisms and the related transformer architecture, proposed by the NLP community, have been introduced to overcome the computational woes of CNNs and RNNs ([Vaswani](#)

[et al. 2017](#)). Complexity per layer is given by  $O(n^2 \cdot d)$ , with a maximum path length of  $O(1)$  and embarrassingly parallelisable operations. The self-attention mechanism has revolutionised the field of sequence modelling and is at the heart of the work presented here. We develop a new transformer architecture for the classification of general multi-variate time-series data, which uses a variant of the self-attention mechanism, and that we apply to the photometric classification of astronomical transients.

This manuscript is structured as follows. Section 2 reviews the recent breakthroughs in the domain of sequence modelling and NLP that have inspired this work, and a pedagogical overview of transformers and the attention mechanism that overcome some of the challenges faced by RNNs and CNNs. Section 3 outlines the attention-based architecture of the time-series transformer developed in this work, with the goal of photometric classification of astronomical transients in mind. Section 4 describes the implementation and performance metrics used to evaluate models. Section 5 presents the results obtained from applying the transformer architecture developed to PLAsTiCC data ([The PLAsTiCC team et al. 2018](#)). Finally, in Section 6 a summary of the work carried out and the key results is discussed.

## 2 ATTENTION IS ALL YOU NEED?


This section gives a pedagogical review of the attention mechanism, and specifically self-attention, which is the foundational element of our proposed architecture. We step through the original architecture that uses self-attention at its core and inspired this work, the transformer ([Vaswani et al. 2017](#)), and how it is generally used in the context of sequence modelling.

### 2.1 Attention Mechanisms

As humans, we tend to focus our *attention* when carrying out particular tasks or solving problems. The incorporation of this concept to problems in NLP has proven extremely successful, and in particular the development of the *attention mechanism* has been shown to have a major impact, not only in the world of sequence modelling, but also in computer vision and other areas of deep learning.

The attention mechanism originates from research into neural machine translation, a sub-field of sequence modelling often referred to as Seq2Seq modelling ([Sutskever et al. 2014](#)). Seq2Seq modelling attempts to build models that take in a sentence represented as a sequence of embeddings  $\mathbf{x} = [x_1, x_2, \dots, x_L]$  and tries to find a mapping to the target sequence  $\mathbf{y} = [y_1, y_2, \dots, y_L]$ <sup>1</sup>. Seq2Seq has traditionally been done by way of two bi-directional RNNs that form an encoder-decoder architecture, with the encoder taking the input sequence  $\mathbf{x}$  and transforming it into a fixed length context vector  $\mathbf{c}$ , and the decoder taking the context through transformations that lead to the final output sequence  $\mathbf{y}$ . The hope is that the context vector is a sufficiently compressed representation of the *entire* input sequence. However, trouble arises with use of RNNs due to the inherent Markov modelling property of these sequential networks, where the state is assumed to be only dependent on the previously observed state. As a consequence RNNs need to maintain memory of each input in the sequence, albeit a compressed representation, up to the desired context length. Therefore, RNNs suffer greatly with computationally maintaining memory for large sequences ([Madsen 2019](#)).

<sup>1</sup> In the domain of NLP, the inputs are word embeddings that are transfor-



**Figure 1.** A model using the attention mechanism, reading the sentence: *The FBI is chasing a criminal on the run.* Blue represents the attention for the input sequence up to the end word in red. The level of shading depicts the attention weighting for each word in the input sequence. Reproduced in full from Cheng et al. (2016).

Attention mechanisms (Bahdanau et al. 2014) were introduced to mitigate these issues and to allow for the full encoder state to be accessible to the decoder via the context vector. This context vector is built from hidden states  $\mathbf{h}$  of the encoder and decoder as well as an alignment score  $\alpha_{ti}$ , between the target  $t$  and input  $i$ . This assigns a score  $\alpha_{ti}$  to the pair  $(y_t, x_i)$ , e.g. in neural machine translation, the word at position  $i$  in the input and the word at position  $t$  in the output, according to how well they align in vector space. It is the set of weights  $\{\alpha_{ti}\}$  that define how much of each input hidden state should be considered for each output. The context vector  $\mathbf{c}_t$  is then defined as the weighted sum of the input sequence hidden states  $\mathbf{h}_i$ , and the alignment scores  $\alpha_{ti}$ . This can be expressed as

$$\mathbf{c}_t = \sum_i \alpha_{ti} \mathbf{h}_i. \quad (1)$$

A common global attention mechanism used to compute the alignments is to compare the current target hidden state  $\mathbf{h}_t$  to each input hidden state  $\mathbf{h}_i$ , as follows (e.g. Luong et al. 2015):

$$\alpha_{ti} = \text{align}(\mathbf{h}_t, \mathbf{h}_i) = \frac{\exp(\text{score}(\mathbf{h}_t, \mathbf{h}_i))}{\sum_{t'} \exp(\text{score}(\mathbf{h}_t, \mathbf{h}_{t'}))}, \quad (2)$$

where score can be any similarity function. For computational convenience this is often chosen to be the dot product of the two hidden state vectors, i.e.

$$\text{score}(\mathbf{h}_t, \mathbf{h}_i) = \mathbf{h}_t \mathbf{h}_i^\top. \quad (3)$$


See Weng (2018) for a summary table of several other popular attention mechanisms and corresponding alignment score functions.

## 2.2 Self-Attention

Self-attention is an attention mechanism that compares different positions of a single input sequence to itself in order to compute a representation of that sequence. It can make use of any similarity function, as long as the target sequence is the same as the input sequence. Prominent use of self-attention came from work in machine reading tasks where the mechanism is able to learn correlations between current words in a sentence and the words that come before (see Figure 1). This type of attention can thus be useful in determining the correlations of data at individual positions with data at other positions in a single input sequence.

Drawing from database and information retrieval literature, a common analogy of query  $q$ , key  $k$ , and value  $v$ , is used when referring to the hidden states of encoder and decoder subcomponents. The

mations of a word at a given position into a numerical vector representation for that word such as the word2vec algorithm (Mikolov et al. 2013).



**Figure 2.** Diagrammatic representation of the computation of the attention matrix  $\mathbf{A}$ . An input sequence of length  $L$  and embedding dimension  $d$  is combined with learned weights to produce query, key and value matrices  $\mathbf{Q}$ ,  $\mathbf{K}$  and  $\mathbf{V}$  respectively.

query,  $q$ , can be seen as the decoder’s hidden state,  $\mathbf{h}_t$ , and the key  $k$ , can be seen as the encoder’s hidden state,  $\mathbf{h}_i$ . The similarity between the query and key can then be used to access the encoder value  $v$ . In the case of dot-product self-attention, learnable-weights are attached to the input  $\mathbf{X} \in \mathbb{R}^{L \times d}$  for sequence length  $L$  and embedding dimension  $d$  for each of  $q$ ,  $k$  and  $v$ , which can be visualised in Figure 2. This results in a set of queries  $\mathbf{Q} = \mathbf{XW}^Q \in \mathbb{R}^{L \times d_q}$ , keys  $\mathbf{K} = \mathbf{XW}^K \in \mathbb{R}^{L \times d_k}$  and values  $\mathbf{V} = \mathbf{XW}^V \in \mathbb{R}^{L \times d_v}$  that can be calculated in parallel, where  $d_q$ ,  $d_k$ , and  $d_v$  are the respective dimensions. A self-attention matrix  $\mathbf{A} \in \mathbb{R}^{L \times d_v}$  can then be computed by

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \mathbf{A} = \text{softmax}(\mathbf{QK}^\top) \mathbf{V}. \quad (4)$$

## 2.3 The Rise of the Transformer

Seminal work by Vaswani et al. (2017) introduced an architecture dubbed the transformer, which is constructed entirely around self-attention. They showed that state-of-the-art performance in neural machine translation can be achieved without the need for any CNN or RNN components; as they put simply “attention is all you need”. Such was the impact of this work that there has since been an explosion of transformer variants as researchers strive to develop more efficient implementations and new applications (Tay et al. 2020). It is the original architecture by Vaswani et al. (2017) that inspired the architecture proposed in this article, and as such the remainder of this section focuses on describing the inner workings of this model.

As can be seen in Figure 3, the transformer consists of two sections: an *encoder* and a *decoder*. Within each encoder and decoder there exists a transformer-block, which contains the multi-head attention mechanism. In the context of neural machine translation, one could think of this set up as the encoder encoding a sentence in English, transforming the input into a certain representation, and the decoder taking this representation and performing the translation to French. To ensure the model only attends to words it has seen up to a certain point when decoding, an additional casual mask is applied to the input sentence. We focus our discussion on the transformer block without this casual mask since it is this block that is most relevant when we come to classification tasks later in this article. Notwithstanding, there is scope for further study to investigate the usefulness

of applying a casual mask to the input sequence for early light curve classification.

### 2.3.1 Multi-Headed Scaled Dot Product Self-Attention

Whilst the main building block used by [Vaswani et al. \(2017\)](#) is indeed the self-attention mechanism, they modified the typical dot-product attention by introducing a *scaled* element. This resulted in a new mechanism called the *scaled dot-product attention* which is similar to Equation 4 but with the input to the softmax scaled down by a factor of  $d_k$ . The motivation for introducing a scaling factor is to control possible vanishing gradients that may arise from large dot-products between embeddings. The new formulation for this mechanism can be expressed as

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \mathbf{A} = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d_k}}\right)\mathbf{V}. \quad (5)$$

This now scaled version of the self-attention module was extended further to also have multiple heads  $h$ , which allows for the model to be able to learn from many representation subspaces at different positions simultaneously ([Vaswani et al. 2017](#)). Similar to normal self-attention calculations show in Section 2.2, this can be pictorially understood with Figure 4 and by concatenating the attentions for each head:

$$\text{MultiHead}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \mathbf{A} = \text{Concat}[\mathbf{A}_1, \dots, \mathbf{A}_h]\mathbf{W}^O, \quad (6)$$

where  $\mathbf{A}_i = \text{Attention}(\mathbf{Q}_i, \mathbf{K}_i, \mathbf{V}_i)$ . With each  $\mathbf{A}_i \in \mathbb{R}^{L \times d_v}$ , the result of a final linear transformation of all concatenated heads,  $\text{Concat}[\mathbf{A}_1, \dots, \mathbf{A}_h] \in \mathbb{R}^{L \times h d_v}$  with learned output weights  $\mathbf{W}^O \in \mathbb{R}^{h d_v \times d}$ , produces the multi-headed attention matrix  $\mathbf{A} \in \mathbb{R}^{L \times d}$ .

### 2.3.2 Additional Transformer-Block Components

As can be seen Figure 3 inside the transformer-block, there is also a pathway that skips the multi-head attention unit and feeds directly into an *Add & Norm* layer. This skip-connection, often referred to a residual connection, allows for a flow of information to bypass potentially gradient-diminishing components. The information that flows around the multi-head attention block is combined with the output of the block and then normalised using layer normalisation ([Ba et al. 2016](#)) by

$$\mathbf{X} \leftarrow \text{LayerNorm}(\text{MultiHeadSelfAttention}(\mathbf{X})) + \mathbf{X}. \quad (7)$$

A feed-forward network follows, comprised of two dense layers with the first using ReLU activation ([Nair & Hinton 2010](#)) and the second without any activation function. A similar skip connection occurs, but instead bypasses the feed-forward network, before being combined again and layer-normalised. It should be noted that all operations inside the transformer-block are time-distributed, which is to say that each word or vector representational embedding, is applied independently at all positions. When combining these elements together, this results in a single transformer-block:

$$\begin{aligned} \mathbf{X} &\leftarrow \text{LayerNorm}(\text{MultiHeadSelfAttention}(\mathbf{X})) + \mathbf{X} \\ \mathbf{X} &\leftarrow \text{LayerNorm}(\text{FeedForward}(\mathbf{X})) + \mathbf{X}, \end{aligned} \quad (8)$$

where  $\mathbf{X}$  is the input embedding to the transformer-block.

### 2.3.3 Input Embedding and Positional Encoding

The inputs to the transformer are word embeddings created from typical vector representation algorithms such as word2vec. Applying this transformation projects each word token into a vector representation on which computations are made. Additionally, recall that attention is computed on sets of inputs, and thus the computation itself is permutation invariant. While this gives strengths to this model in terms of parallelism, a drawback of this is the loss of temporal information that would usually be retained with RNNs. A consequence of this is the need for positional encodings to be applied to the input embeddings. In [Vaswani et al. \(2017\)](#) the positional encoding  $\mathbf{P} \in \mathbb{R}^{L \times d}$ , which is used to provide information about a specific position in a sentence ([Weng 2018](#)), is computed by a combination of sine and cosine evaluations at varying frequencies. Assume  $l$  to be a particular position location in an input sequence, with  $l = 1, \dots, L$ , and embedding index  $k = 1, \dots, d$ , then

$$\mathbf{P}_{lk} = \begin{cases} \sin(\omega_k \cdot l), & \text{if } k \text{ even} \\ \cos(\omega_k \cdot l), & \text{otherwise} \end{cases} \quad (9)$$

$$\text{where } \omega_k = \frac{1}{10000^{2k/d}}.$$

Provided the dimension of the word embedding is equal to the dimension of the positional encoding, the positional vector  $\mathbf{p}_l \in \mathbb{R}^d$  corresponding to a row of the matrix  $\mathbf{P}$  is added to the corresponding word embedding  $\mathbf{x}_l$  of the input sequence  $[\mathbf{x}_1, \dots, \mathbf{x}_n]$  ([Kazemnejad 2019](#)):

$$\mathbf{x}_l \leftarrow \mathbf{x}_l \oplus \mathbf{p}_l. \quad (10)$$

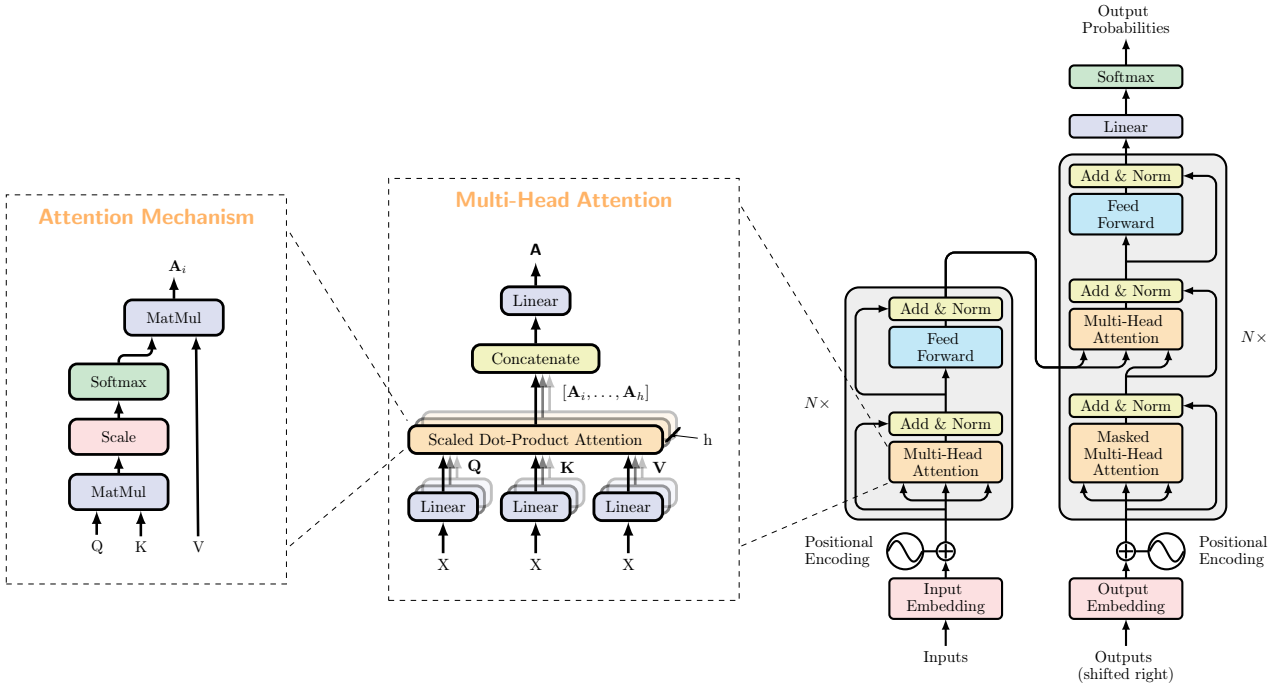
For a visual representation of the position encoding see Figure 5, which depicts the positional encoding for a 128-dimensional by 100 sequence length input embedding. Using positional encoding in this way allows for the model to have access to a unique encoding for every position in the input sequence. The motivation for using sine and cosine functions are such that the model is also able to learn relative position information since any offset,  $\mathbf{p}_{l+\text{offset}}$  can be represented as a linear function of  $\mathbf{p}_l$  ([Vaswani et al. 2017](#)).

## 3 THE TIME-SERIES TRANSFORMER: T2


In this section we present our transformer architecture for time series data, which is based on the self attention mechanism and the transformer-block. Our work is motivated by photometric classification of astronomical transients but generally applicable for classification of general time-series. The time-series transformer architecture that we propose supports the inclusion of additional features, while also offering interpretability. Furthermore, we include layers to support the irregularly sampled multivariate time-series data typical of astronomical transients.

### 3.1 Architecture

Our architecture, referred to from herein as t2, shown in Figure 6, has several key differences compared to the original transformer shown in Figure 3. The first of these differences is the removal of the decoder. As the task at hand is classification, a single transformer-block is sufficient ([Tay et al. 2020](#)). Another difference can be seen with the additional two layers prior to positional encoding unit, which are *Gaussian Process Interpolation* and *Convolutional Embedding*. In



**Figure 3.** Layout of the original transformer architecture defined in Vaswani et al. (2017). The multi-head attention unit has been zoomed-in to reveal the inner workings and key component of the scaled dot-product attention mechanism. Note the two grey boxes on the left and right of the architecture. These are both transformer blocks, with  $N$  indicating how many times each block is stacked upon itself.




**Figure 4.** Diagrammatic representation of the computation of the multi-head attention. Instead of computing attention once, the multi-head mechanism divides the input with sequence length  $L$  and embedding dimension  $d$  by number of heads  $h$  to compute the scaled dot-product attention over each subspace simultaneously. These independent attention matrices are then concatenated together and linearly transformed into an attention matrix  $\mathbf{A} \in \mathbb{R}^{L \times d}$ . The above diagrammatic representation assumes the dimensionality of keys  $d_k$  is the same as the dimensionality of the values  $d_v$ .

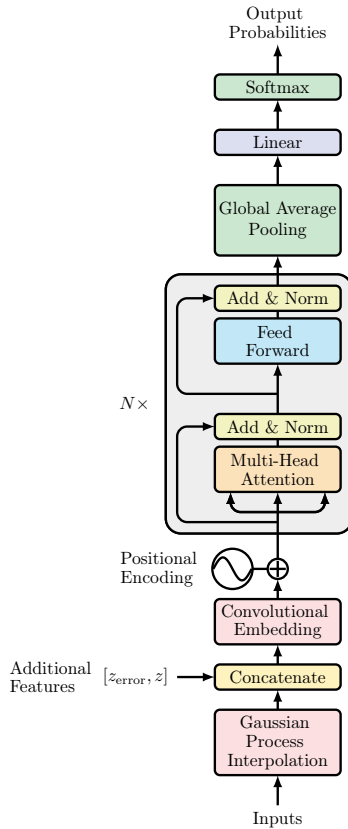
conjunction with these two layers is a *Concatenation* layer that is able to add an arbitrary number of additional features to the network. These layers process the astronomical input sequence data and pass it to a typical transformer-block. The output of the transformer-block is then passed through a new *Global Average Pooling* layer, before finally being passed through a softmax function that provides output probabilities over the classes considered.

### 3.2 Irregularly Sampled Multivariate Time-series Data

With neural machine translation the input consisted of a sequence of words that form a sentence. While this is similar for astronomical transients in the sense that one has a sequence of observations that form a light curve, there are several differences that are important to address. It will be useful to review the kind of data one is dealing



**Figure 5.** A 128-dimensional positional encoding for a sequence of length of 100. This can be understood as each row representing the encoding vector for that position in the sequence.



**Figure 6.** Schematic of the time-series transformer (t2) architecture. Raw time-series data is processed through the Gaussian process interpolation layer, followed by a concatenation layer to include any additional features. A convolutional embedding layer follows to transform the input into a vector representation, with a positional encoding applied to the embedding vector. This is passed as input into the transformer-block, where the multi-head attention block is the same as that shown in Figure 3. The output of which is then passed through a global average pooling layer and finally a linear layer with softmax to output class prediction probabilities for the objects.

with and to make some definitions (adapted from Fawaz et al. 2019) with regards to the task of astronomical transient classification. In general, the data that one observes can be viewed as an *irregular multivariate time-series* signal:

**Definition 1.** A *univariate time-series* signal  $\mathbf{x} = [x_1, x_2, \dots, x_T]$  consists of an ordered set of  $T$  real values with  $\mathbf{x} \in \mathbb{R}^T$ .

**Definition 2.** An  $M$ -dimensional *multivariate time-series* signal,  $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_M]$  consists of  $M$  univariate time series with  $\mathbf{X} \in \mathbb{R}^{T \times M}$ .

**Definition 3.** An *irregular time-series* is a ordered sequence of observation time and value pairs  $(t_n, x_n)$  where the space between observation times is *not* constant.

**Definition 4.** A dataset  $\mathcal{D} = \{(X_1, Y_1), (X_2, Y_2), \dots, (X_N, Y_N)\}$  is a collection of pairs  $(X_i, Y_i)$  where  $X_i$  could either be a univariate or multivariate time series with  $Y_i$  as its corresponding one-hot label vector. For a dataset containing  $C$  classes, the one-hot label vector  $Y_i$  is a vector of length  $C$  where each element is equal to 1 for the index corresponding to the class of  $X_i$  and 0 otherwise.

The goal for general time-series classification consists of training a classifier on a dataset  $\mathcal{D}$  in order to map from the space of possible inputs to a probability distribution over the class variable labels. For photometric classification of astronomical transients, the light that is observed is collected through different filters, also known as passbands, that allow frequencies of light within certain ranges to pass through. Observations collected over a period of time for a particular celestial object form a *light-curve*. A light-curve is irregularly sampled in time and across passbands which further complicates the task.


### 3.2.1 Data Interpolation with Gaussian Processes

The raw time-series that is observed is irregularly sampled with heteroskedastic errors. A technique widely used to overcome missing data, and that can also provide uncertainty information is Gaussian process regression (Rasmussen 2004). This technique is a popular method that has been applied to Supernovae light curves for many years, e.g Lochner et al. (2016). Gaussian processes represent distributions over functions  $f$  that when evaluated at a given point  $x$  is a random variable  $f(x)$ , with mean  $\mathbb{E}[f(x)] = m(x)$  and covariance between two sampled observations  $x, x'$  as  $\text{Cov}(f(x), f(x')) = \mathbf{K}_f(x, x')$ , where  $\mathbf{K}_f(\cdot, \cdot)$  is a kernel.

An important aspect of applying Gaussian process interpolation to data is the choice of kernel. It was discovered in Boone (2019) that for general transients a 2-dimensional kernel that incorporates both wavelength (*i.e.* passband) information as well as time works well. It can be seen in Boone (2019) that by use of a 2-dimensional kernel, correlations between passbands are leveraged and predictions in passbands that do not have any observations are still possible. As such, we use the Matern kernel (Rasmussen 2004) that is parametrised by  $\nu$  which controls the smoothness of the resulting function and set to  $\nu = 3/2$ . By performing Gaussian process regression and then sampling the resulting Gaussian process at regular intervals, we transform our previously irregular multivariate time-series to a now well sampled regular multivariate signal. The Gaussian process mean is sampled at regular points in time to produce  $\mathbf{X} \in \mathbb{R}^{L \times M}$ , where  $L$  is the sampled time sequence length and  $M$  is the number of passbands. This procedure is illustrated in Figure 7 (along with a final convolutional embedding explained in the following Section 3.3).

## 3.3 Convolutional Embedding

With neural machine translation applications the inputs to the original transformer architecture take in word embeddings that had been derived from a typical vector representation algorithm such



**Figure 7.** The three-stage process of transforming an astronomical transient light curve from raw photometric time-series data to a vector representation suitable as input to the time-series transformer. First, Gaussian process regression is carried out to regularly sample the light curve. The Gaussian process mean is evaluated at  $L$  time points  $l$  for each of  $M$  passbands, and projected to dimension  $d$  via a 1-dimensional convolutional embedding (Lin et al. 2013) at each point  $l$  such that final input sequence  $\mathbf{X} \in \mathbb{R}^{L \times d}$ . The resulting embedding matrix is then shown in blue, where for example a  $d$ -length vector for a single input position highlighted in green.

as word2vec. In a similar manner, embeddings for the now interpolated time-series data are required. We adopt a simple 1-dimensional convolutional embedding, with kernel size of 1 and apply a ReLU non-linearity. Inspired by Lin et al. (2013), a 1-dimensional convolution allows for a transformation from  $k$ -dimensional space to a  $k'$ -dimensional space whilst operating over a single time window of size of 1. For our purposes, using this convolution allows for dimensionality to be scaled from  $M$  to  $d$  dimensions without affecting the spatio-temporal input. Therefore, this operation transforms the original input of  $M$ -dimensional time-series data points, *i.e.* time-series data points across  $M$  passbands, into a  $d$ -dimensional vector representation ready for input into the transformer-block. This operation is akin to a time-distributed, position-wise feed-forward neural network operating on each input position.

### 3.4 Global Average Pooling

We also introduce a layer that performs global average pooling (GAP) on the output of the transformer-block. The motivation for adding a GAP layer following the transformer-block was inspired by work in Zhou et al. (2015). The GAP layer, originally proposed by Lin et al. (2013), has become a staple in modern CNN architectures due to its usefulness in interpretable machine learning. In previous works on 2-dimensional images, GAP layers are used as a replacement to common fully connected layers to avoid overfitting since there are no parameters to optimise. Another useful advantage over the fully connected layer is the averaging in a GAP layer averages out the spatial information leaving it more robust to translations of the inputs (Lin et al. 2013). Similar to 2-dimensional inputs, using a GAP layer on the 1-dimensional time-series, proves robustness to translations in the input.

By adapting the description found in Zhou et al. (2015), one can apply a GAP layer to a time-series. Let  $f_k(l)$  represent the activation of a particular embedded dimension  $k$  at a location  $l$ , where  $k = 1, \dots, d$  and  $l = 1, \dots, L$ . Then a GAP layer can be computed by taking the average over time for each feature map  $F_k = \sum_l f_k(l)$ .

### 3.5 Class Activation Maps (CAM)

A nice feature of using a GAP layer is that one can determine the influence of  $f_k(l)$  on predictions for a given class  $c \in C$  by considering the associated score  $S_c$  that is passed into the softmax layer (Zhou et al. 2015). This is calculated from the final fully connected weights  $w_k^c$  and the feature maps  $F_k$  as  $S_c = \sum_k w_k^c \cdot F_k = \sum_l \sum_k w_k^c \cdot f_k(l)$ .

The class activation map (CAM) for a given class  $c$  is then given by

$$M_c(l) = \sum_k w_k^c f_k(l). \quad (11)$$

Since  $S_c = \sum_l M_c(l)$ , it is possible to use  $M_c(l)$  to directly gauge the importance of the activation at input location  $l$  in leading to the classification of class  $c$ .

### 3.6 Imputing Additional Information

The time-series transform,  $\tau_2$ , is designed to be malleable such that one can add further features if desired. For the purposes of the current study of photometric classification, only redshift information has been added. In many photometric classifiers, photometric redshift  $z$ , has consistently been a feature of high importance (*e.g.* Boone 2019). As one particular example of the type of additional features that can be added, photometric redshift  $z$  and the associated error  $z_{\text{error}}$  are included.

Additional features could in principle be incorporated in the time-series transformer in a variety of different manners. To leverage the power of neural networks to model complex non-linear mappings, such additional features should feed through non-linear components of the architecture. On the other hand, recall from Section 3.5 that in order to compute a CAM (class activation map), the output of the GAP layer must pass directly into the linear softmax layer. Hence, incorporating additional features at this stage of the architecture will not be effective unless a non-linear activation is introduced, which would destroy the interpretability of the model.

To preserve our ability to compute CAMs, there are several other possible locations in the architecture where one could consider including additional features. The most natural point is immediately

prior to the convolutional embedding layer (see Figure 6). Adding features at this location allows for all information to be passed throughout the entire network. Nevertheless, there are alternative ways in which additional features can be incorporated at this point.

The most obvious way to incorporate additional features is to essentially consider them as a additional channels and concatenate in the dimension of the  $M$  passbands to redefine the input as  $\mathbf{X} \in \mathbb{R}^{L \times M} \rightarrow \mathbf{X} \in \mathbb{R}^{L \times M'}$ , where  $M' = M + R$ , and  $R$  is the number of features to add. This essentially broadcasts the additional information to each input position in  $l$ .

The alternative is to concatenate in the dimension of the  $L$  time sequence samples, which transforms the input as  $\mathbf{X} \in \mathbb{R}^{L \times M} \rightarrow \mathbf{X} \in \mathbb{R}^{L' \times M}$ , where  $L' = L + R$ . There are several advantages for choosing the approach of concatenating to  $L$  rather than  $M$ . Firstly, this approach allows one to pay attention to the additional features explicitly. Secondly, it gives activation weights for the additional features, which in our case is redshift and redshift error, so the impact of the additional features can be interpreted. So, while in principle one could consider concatenating to either  $L$  or  $M$ , we advocate concatenating to  $L$ .

### 3.7 Trainable Parameters and Hyperparameters

The time-series transformer,  $\tau_2$ , model contains a set of trainable parameters that stem from the weights contained in the transformer-block as well as learned weights at the embedding layer and final fully connected layer. The first layer with trainable parameters is the convolutional embedding layer. The numbers of parameters for a general convolutional layer is given by

$$[M \times w \times d] + d,$$

where  $M$  denotes the number of input channels or passbands,  $w$  refers to the kernel window size, which in this case is 1, and  $d$  is the dimensionality of the embedding. Continuing through the model, the number of trainable parameters for the multi-head attention unit has 4 linear connections, including  $\mathbf{Q}$ ,  $\mathbf{K}$ ,  $\mathbf{V}$  and one after the concatenation layer, *i.e.*  $\mathbf{W}^Q$ ,  $\mathbf{W}^K$ ,  $\mathbf{W}^V$  and  $\mathbf{W}^O$ . Recall that for multi-head attention we set  $hd_v = d$  (see Figure 4), hence the number of parameters for  $\mathbf{W}^Q$ ,  $\mathbf{W}^K$ ,  $\mathbf{W}^V$  and  $\mathbf{W}^O$  across all of the  $h$  heads is identical. The number of layer normalisation parameters is simply the sum of weights and biases together with the feed forward neural network weights of the input multiplied by the weights of the output plus the output biases (Chai 2020). Combining all units inside the transformer block together yield

$$N \times \underbrace{(4 \times [(d \times d) + d])}_{\text{Multi-Head Attention}} + \underbrace{(2 \times 2d)}_{\text{Layer Norm}} + \underbrace{(d \times d_{\text{ff}} + d_{\text{ff}})}_{\text{Feed Forward}} + (d_{\text{ff}} \times d + d)$$

where  $N$  refers to how many times one stacks the transformer-block upon itself, and  $d_{\text{ff}}$  refers to the number of neurons in the feed forward network inside the transformer-block. Since there are no trainable parameters with the GAP layer, the final fully connected linear layer with softmax results in a remaining number of trainable parameters of

$$([d \times C] + C),$$

where  $C$  refers to the number of classes.

Of the parameters discussed above, there are some that are fixed due to the problem at hand, such as  $M$  number of passbands and  $C$  classes to classify. But there are also other parameters that are not necessarily trainable that are considered *hyperparameters*. These

include: the dimensionality of the input embedding  $d$ , the dimensionality of the feed forward network inside the transformer-block  $d_{\text{ff}}$ , the number of heads to use in conjunction with the multi-head attention unit  $h$ , the percentage of neurons to drop when in training using the dropout method (Srivastava et al. 2014) `dropout`, the number of transformer-blocks  $N$ , and the learning rate `learning_rate` (discussed further in Section 4.4).

## 4 IMPLEMENTATION, EVALUATION METRICS & TRAINING

We leverage modern machine learning frameworks to develop the time-series transformer implementation,  $\tau_2$ , in a modular manner for ease of use and future extension. We present the key evaluation metrics that we use to measure the performance of the classifier and the motivation for particular types of metrics in relation to the photometric astronomical transient classification problem that we consider in Section 5. We also discuss the loss function used for training and how hyperparameters are optimised.

### 4.1 Implementation

We use the machine learning framework of TensorFlow (Abadi et al. 2015) with the `tf.keras` API for the implementation of our  $\tau_2$  architecture. Our code is available under Apache 2.0 licence and open-sourced<sup>2</sup>. Key data processing software of pandas (Wes McKinney 2010) and numpy (Harris et al. 2020) has been used heavily for manipulation of input data, with george (Ambikasaran et al. 2015) used for fitting the Gaussian processes. Training and inference of our model has been carried out on a NVIDIA Tesla V100 GPU.

### 4.2 Performance Metrics

Choice of evaluation metrics is of high importance when considering the performance of a classifier. This is compounded when dealing with imbalanced datasets since most metrics consider the setting of an even distribution of samples among the classes. One must be careful when considering which metrics to evaluate a model's performance since relatively robust procedures can be unreliable and misleading when dealing with imbalanced data (Branco et al. 2015; Malz et al. 2019).

Typically, threshold metrics are used which consider the rate or fraction of correct or incorrect predictions. Threshold metrics are formulated by combinations of the four possible outcomes a classifier could have with regards to predicting the correct class:

- True Positive (TP): prediction of a given class and indeed it being that class.
- False Positive (FP): prediction of a given class but it does *not* belong to that class.
- True Negative (TN): prediction that an object is *not* a particular class and it is indeed *not* that class.
- False Negative (FN): prediction that an object is *not* a particular class but it is in fact that class.

From these outcomes common threshold metrics can be formulated, with perhaps the most common threshold metric being *accuracy*,

<sup>2</sup> [github.com/tallamjr/astronet](https://github.com/tallamjr/astronet)



which is the number of correctly classified samples over the total number of predictions. However, for imbalanced data results on accuracy alone can be misleading as a model can achieve high accuracy by simply classifying the majority class. More robust metrics for imbalanced data are precision and recall since their focus is on a particular class:

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}, \quad \text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}. \quad (12)$$

Precision gives the fraction of samples predicted as a particular class that indeed belong to the particular class. While recall, also known as the true positive rate, indicates how well a particular class was predicted.

#### 4.2.1 Confusion Matrix

One way to visually inspect the performance of a classifier with regards to threshold metrics is by the confusion matrix. The confusion matrix provides more insight into the performance of the model and reveals which classes are being predicted correctly or incorrectly. Often these tables are normalised across the rows to give probabilities in order to provide a more intuitive understanding. A perfect classifier across all classes would therefore be equivalent to the identity matrix with all ones along the diagonal and zero elsewhere.

#### 4.2.2 Receiver Operating Characteristic

An important point to note is that threshold metrics alone assume the class imbalance present in the training set is of the same distribution as that of the test set (He & Ma 2013). On the other hand, a set of metrics built from the same fundamental components as threshold metrics, called rank metrics, do not make any assumptions about class distributions and therefore are a useful tool for evaluating classifiers based on how effective they are at distinguishing between classes (Brownlee 2020).

Rank metrics require that a classifier predicts a probability of belonging to a certain class. From this, different thresholds can be applied to test the effectiveness of classifiers. Those models that maintain a strong probability of being a certain class across a range of thresholds will have good class separation and thus will be ranked higher.

The most common of this type of metric is the receiver-operating-characteristic (ROC) curve, which plots the true positive rate versus the false positive rate to estimate the behaviour of the model under different thresholds. The ROC curve is then used as a diagnostic tool to evaluate the model's performance, with every point on graph representing a given threshold. Interpolating between these points forms a curve, with the area under the curve (AUC) quantifying performance. A classifier is effectively random if the AUC is 0.5 and, conversely, is a perfect classifier if the AUC is equal to 1.0.

#### 4.2.3 Precision-Recall Trade-Off

An alternative diagnostic plot to the ROC curve is the precision-recall (PR) trade-off curve. This is used in a similar way to the ROC curve but instead focuses on the performance of the classifier to the minority class, and hence is more useful for imbalanced classification problems (Brownlee 2020). Much like the ROC curve, points on the curve represent different classification thresholds with a random classifier resulting in an AUC equal to 0.5 and a perfect classifier resulting in an AUC of 1.0.

### 4.3 Multi-Class Logarithmic-Loss

The underlying algorithm that governs the usefulness of neural networks is the stochastic gradient descent (SGD) optimisation algorithm that updates the weights of the network according to the backpropagation algorithm (Rumelhart et al. 1986). While performance metrics give an indicator as to how well a model is able to distinguish between classes, to be able to train and improve the model one must have a differentiable loss function. Extensive investigations by Malz et al. (2019) showed that the most suitable differentiable loss-function for the problem of transients classification is a probabilistic loss function. Probabilistic loss functions are used in cases where the uncertainty of a prediction is useful and the problem at hand is best served with quantification of the errors rather than a binary answer of correct or incorrect. The probabilistic loss function they suggest is the multi-class weighted logarithmic-loss that up-weights rarer classes and defines a perfect classifier as one that achieves a score of zero, and is given by

$$\mathcal{L} = - \left( \frac{\sum_{i=1}^C w_i \sum_{j=1}^{N_i} \frac{y_{ij}}{N_i} \ln p_{ij}}{\sum_{i=1}^C w_i} \right), \quad (13)$$

where  $C$  refers to the number of classes in the dataset and  $N_i$  the number of objects in the  $i$ -th class. The predicted probability of an observation  $i$  belonging to class  $j$  is given by  $p_{ij}$ . For our investigation we opt for a flat-weighted multi-class logarithmic-loss as described in Boone (2019) that assigns all classes in the training set the same weight of  $w_i = 1$ . To consider the original metric put forth in Malz et al. (2019) and use the weighting scheme designed for the PLAsTiCC competition, one would also need to include the additional anomaly classes (class 99) that existed in the PLAsTiCC test set. By ignoring class 99 one can better compare later analyses between the original PLAsTiCC training set and our modified dataset (described in upcoming Section 5.1).

### 4.4 Training

In order to train a model with the  $\tau_2$  architecture, we need to establish the choice of optimisation algorithm and associated parameters that will be used to update the weights of the network. We use a variant of the SGD optimisation algorithm mentioned in Section 4.3 called ADAM (Kingma & Ba 2014). An important aspect to consider when training a model using any optimisation algorithm is the learning schedule and corresponding learning rate. The initialisation value of the learning rate can be seen as a hyperparameter to be optimised for separately with hyperparameter optimisation (discussed in the next section). It is typically beneficial to introduce a learning schedule to reduce the learning rate as training progresses (Goodfellow et al. 2016). We indeed adopt a learning schedule, reducing the learning rate by 10% if it is observed that our loss value does not decrease within 5 epochs. To ensure the model does not overfit, we monitor the ratio of validation loss with the training set loss.

### 4.5 Hyperparameter Optimisation

As discussed in Section 3.7,  $\tau_2$  contains a set of fixed parameters such as  $M$  and  $C$ , and a set of tunable hyperparameters. Choosing the best set of hyperparameters can be framed as an optimisation problem expressed as

$$\theta^* = \arg \min_{\theta \in \Theta} g(\theta), \quad (14)$$

where  $g(\theta)$  is an objective score to be minimised and evaluated on a validation set, with the set of hyperparameters  $\theta$  being able to take any value defined in the domain of  $\Theta$ . The objective score for our purposes is the logarithmic-loss defined in Equation 13 and the set of hyperparameters that yield the lowest objective score is  $\theta^*$ . The goal is to find the model hyperparameters that yield the best score on the validation set metric (Koehrsen 2018).

Traditionally hyperparameter optimisation has been performed with either random search or a grid search over the set of parameters in  $\Theta$ , which can be time consuming and inefficient. Instead a Bayesian optimisation approach is used that attempts to form a probabilistic model mapping hyperparameters to a probability distribution for a given score.

To choose the best performing hyperparameters we use the Tree-structured Parzen Estimator (TPE) algorithm (Bergstra et al. 2011) that is implemented in the optuna package (Akiba et al. 2019) with 5 fold cross-validation.

## 5 RESULTS

We apply our time-series transformer architecture to the problem of photometric classification of astronomical transients. As noted in Section 1 typical astronomical data that are available for training a photometric classifier are highly imbalanced, with a large number of spectroscopically confirmed Supernova Type Ia compared to other classes, and non-representative, since observations are biased towards lower redshift objects. Consequently, the training data are non-representative of the test data. For robust and accurate classification, training datasets should be representative of the test data. Works by Boone (2019) and Revsbech et al. (2018) present techniques that help address this problem of non-representativity, transforming the training data to be more representative of the true test data through data augmentation. This process is involved and can be decoupled from the design of architecture of the classifier. Therefore in this current article, as a first step we consider training data that is representative in redshift but imbalanced. In future work we will consider the combination of  $t_2$  with augmentation techniques to address the representativity problem.

### 5.1 Astronomical Transients Dataset

To be able to evaluate our architecture in a representative setting, but also to test the models resilience to class imbalance, we utilise the PLAsTiCC dataset (The PLAsTiCC team et al. 2018). The complete dataset contains synthetic light curves of approximately 3.5 million transient objects from a variety of classes simulated to be observed in 6 passbands using a cadence defined in Kessler et al. (2019).

The majority of events that exist in the dataset were simulated to be observed with the Wide-Fast-Deep (WFD) mode, which compared to the Deep-Drilling-Fields (DDF) observing mode, is more sparsely sampled in time and has larger errors. Originally crafted for a machine learning competition<sup>3</sup>, the entire PLAsTiCC dataset was divided into two parts, with < 1% initially being given to participants in the competition that was highly non-representative of the other part. Following the close of the competition all data are now publicly available<sup>4</sup>. For our purposes, we use the complement to what was initially released and construct a new training and test set from the remaining 99% of the data (without anomaly class 99). By doing so,

**Table 1.** Number of samples of the PLAsTiCC data used for evaluation of the  $t_2$  model. Note the largely imbalanced dataset distribution of SNIa objects compared to other classes.

Class	Number of Samples (%)
$\mu$ – Lens-Single	1,303 (0.037%)
TDE	13,552 (0.389%)
EB	96,560 (2.775%)
SNII	1,000,033 (28.741%)
SNIax	63,660 (1.830%)
Mira	1,453 (0.042%)
SNIbc	175,083 (5.032%)
KN	132 (0.004%)
M-dwarf	93,480 (2.686%)
SNIa-91bg	40,192 (1.155%)
AGN	101,412 (2.915%)
SNIa	1,659,684 (47.700%)
RRL	197,131 (5.666%)
SLSN-I	35,780 (1.028%)
Total	3,479,456 (100%)

**Table 2.** The time-series transformer,  $t_2$ , contains 6 hyperparameters to be optimised. The set of parameters and learning rate that scored the lowest objective score using 5-fold cross-validation and the TPE Bayesian optimisation method is shown here. To be concise we only show `learning_rate` to 3 decimal places and advise the reader to refer to the code for full details.

Parameter	Value
$d$	32
$h$	16
$d_{ff}$	128
$N$	1
<code>droprate</code>	0.1
<code>learning_rate</code>	0.017

the dataset is now representative in terms of redshift, but remains highly imbalanced in terms of the classes. The number of samples per class used to evaluate our architecture can be found in Table 1.

### 5.2 Classification Performance

Of the model parameters in the time-series transformer,  $t_2$ , there are a subset of hyperparameters that are tunable and can be optimised for (see Section 4.5). Through application of the TPE Bayesian optimisation method on a validation set constructed from 10% of the training set, using 5-fold cross-validation we obtained the parameters which gave the lowest objective score. The results of which can be found in Table 2.

When we build our time-series transformer with the hyperparameters shown in Table 2, and train a model using the training data set described in 5.1 we are able to achieve a logarithmic-loss of 0.507. The confusion matrix depicted in Figure 8 shows good performance across all classes. Both receiver operating characteristic (ROC) and precision-recall (PR) plots, Figure 9 and Figure 10 respectively, show reasonable multi-class classification accuracy, with the exception being towards the Kilonovae and SNIax classes. We suspect this is purely down to the scarcity of sample for Kilonovae and light curve similarity to SNIa in the case of SNIax as mentioned in 5. The performance of our model expectedly degrades when auxiliary information of redshift and redshift error is not included. However,

<sup>3</sup> [kaggle.com/c/PLAsTiCC-2018](https://kaggle.com/c/PLAsTiCC-2018)

<sup>4</sup> [zenodo.org/record/2539456#.YIIVA5NKjlz](https://zenodo.org/record/2539456#.YIIVA5NKjlz)

we find it promising that our model with raw time-series information only can still achieve a logarithmic-loss of 0.873.

It is expected that if a full hyperparameter search can be performed on the full training set by leveraging greater computational resources, it is likely better parameters could be discovered leading to improved performance. While a direct comparison with other methods presented in Hložek et al. (2020) cannot be made since they have been trained with non-representative datasets, the time-series transformer is able to achieve excellent classification performance with minimal feature selection and few trainable parameters by deep learning standards.

It is often the case with machine learning models that, as remarked upon in Hložek et al. (2020) and Lochner et al. (2016), in order to overcome a classification bias towards particular classes, an equal distribution of samples among the classes is often necessary for accurate classification. However, the  $\tau_2$  architecture is able to handle class imbalance very well, and as such our model did not require any data augmentation in order to achieve a good score, unlike other methods. It is uncertain at this time whether this is an inherent property of transformers or the attention mechanism, or perhaps the architecture is simply able to find sufficient discriminative features with far fewer training samples than was previously thought is required for deep learning approaches such as CNNs and RNNs. As discussed already, we are yet to consider the case of data that is not representative in redshift, where augmentation techniques will certainly be necessary, which will be the focus of future work.

### 5.3 Interpretable Machine Learning

Work by Zhou et al. (2015) lead the way forward with major improvements for model interpretability. Their use of the GAP (global average pooling) layer for the localisation of feature importance helped researchers discover methods of visually inspecting a classifier’s performance. In a similar regard, a GAP layer is included in the  $\tau_2$  architecture to allow for model interpretability through the visualisation of the various feature maps as a function of sequence length. As discussed in Section 3.5, one can compute a CAM (class activation map) which can help determine how the features at each input position have influenced the final prediction. Also recall from Section 3.6 that  $\tau_2$  allows for concatenation of arbitrary additional features; in this work we consider the addition of redshift information.

Of the two options for concatenation, either in time or passband, we adopt the approach of concatenating to  $L$  in time to give  $L' = L + R$ , where  $R = 2$  with redshift and redshift error added as additional features. This has the advantage that we explicitly pay attention to redshift information and also get interpretability with respect to redshift information (see Section 3.6). For completeness, we also re-run the photometric classification analysis discussed previously by concatenation to  $M$ , but we do not observe as good a performance as concatenating to  $L$ . As we suspected, this may well be because we are explicitly paying attention to redshift in the multi-head attention mechanism, whereas by concatenating to  $M$  we do not get this benefit.

The CAM can then be computed by Equation 11, where  $M_c(l')$  indicates the influence each position of the input sequence has on classification, which also includes redshift information, *i.e.*  $l' = 1 \dots, L + R$ . We apply a min-max scaling and normalise the CAM such that  $\sum_1^{L+R} M_c(l') = 1$ , so that the relative activation weights can be interpreted as a percentage.

We show in Figure 11 illustrative CAMs for two Supernova classes, over-plotted with the lightcurves themselves. In each panel CAM probabilities for each light curve time point are shown, in addition

to the CAM probabilities for the additional features of redshift and redshift error. Notice that for both examples the activation weight is generally low before the initial rise of the light curve, larger at the rise, with the strongest weight around the peak. Moderate weights are observed in the tail, presumably to detect any secondary peak, with the weights becoming insignificant once the light curve falls back to zero.

As our architecture is able to include additional features, these can also be inspected and visualised to gain further understanding as to how much importance the model is paying towards them. In our case, with the addition of redshift and redshift error information, we also include bar plots in Figure 11 that depict the activation weight for redshift and redshift error. We inspect the distribution of the activation weights for redshift and redshift error for all classes combined, which can be seen in Figure 12. The majority of activation weighting relating to redshift and redshift error falls around 1%. We also explored this distribution on an individual class by class basis but did not find a significant difference across classes. Therefore, there does not seem to be a particular class that benefits from redshift information over another. The distributions indicate that for most objects redshift informations accounts for a relative small proportion of the total activation weights, with a mean of  $\sim 1.92\%$ . However, it should be noted that this is related to the  $L = 100$  regularly sampled points on the light curve, many of which are highly informative. Furthermore, we recall that redshift on the whole is indeed important for accurate classification where we achieve a logarithmic-loss of 0.507 when including redshift information and 0.873 when it is not included (Section 5.2).

While we have shown CAMs to be useful for a first attempt to bring interpretability to light curve classification, we acknowledge more recent saliency mapping techniques that address some of the shortcomings of CAMs. We commented earlier that in order to compute CAMs we require the GAP layer. Although we have provided separate motivation for using a GAP layer (see Section 3.4) it may be the case that in the pursuit of better interpretability, requiring a GAP layer unnecessarily restricts the flexibility of our model for possible model extensions. Therefore it would be preferable to have an explainable methodology that does not impose certain characteristics on the architecture itself, and that can ideally probe a model in a black-box fashion. Follow-up work by Selvaraju et al. (2017) presented Grad-CAM that did away with the need for a GAP layer to feed directly into the softmax and was agnostic to the downstream task, but still required access to the internals of the model with gradients. An interpretability method proposed by Petsiuk et al. (2018) introduced randomised input sampling for explanation of black-box models (RISE) to better estimate how salient aspects of the input are for a model’s prediction, without the need for access to model internals nor re-implementation of existing models.

It is expected that future studies for interpretability of photometric classification architectures use techniques similar to RISE that can treat the model as a black-box and yet provide more refined saliency maps.

## 6 CONCLUSIONS

We have constructed a new deep learning architecture designed for photometric classification of astronomical transients that we call the time-series transformer or  $\tau_2$ . The architecture is designed in such a way to pay attention not only to light curves but also to any additional features considered (*e.g.* redshift information) and to also provide interpretability, again not only to light curves but also to additional




Figure 8. Confusion matrix resulting from application of the time-series transformer,  $\tau_2$ , to the PLAsTiCC dataset in a representative setting with imbalanced classes, achieving a logarithmic-loss of 0.507.





Figure 9. Receiver operating characteristic (ROC) curve, under the same setting as those described in Figure 8. Micro- and macro-averaged AUC scores of 0.98 are achieved across the classes.



**Figure 10.** Precision-recall trade-off curve, under the same setting as those described in Figure 8. A micro-averaged AUC score of 0.87 is achieved across the classes. The model understandably struggles with precision for Kilonovae (KN), which only constitutes 0.004% of the training sample.

features. While we are motivated by the problem of astronomical transient classification, the architecture is suitable for general multi-variate time-series data.

The time-series transformer,  $\tau_2$ , is able to achieve results comparable to the state-of-the-art in photometric classification and does so on extremely imbalanced datasets. Our architecture is able to achieve a logarithmic-loss of 0.507 on the PLAsTiCC dataset defined in Section 5.1 and Table 1. A direct comparison to other latest methods laid out in Hložek et al. (2020) and Gabruseva et al. (2020) is understandably not possible since each classifier has been evaluated on different data under different conditions, nonetheless,  $\tau_2$  is able to achieve the lowest logarithmic-loss on such imbalanced data, without the need for augmentation. Having such an imbalanced dataset, one would expect that there would be bias towards the most common classes, but  $\tau_2$  is robust enough to handle this. As noted in Lochner et al. (2016), accurate photometric classification requires a representative training dataset, but as discussed in Section 1 the data that will be observed with upcoming surveys will be non-representative of the training datasets that are currently available. While this work focuses on the representative setting, the architecture lends itself well to be able to be used in conjunction with latest augmentation techniques, particularly Boone (2019) and Alves et al. (2022) with use of Gaussian processes, that should help to alleviate non-representative training dataset issues, and as such this will be considered in detail in future work.


While  $\tau_2$  is already able to compete with state-of-the-art methods, improvements could be made in future work to modify the components of the architecture, while keeping the broad structure, e.g. by replacing self-attention layers with alternative mixing transforms such as Fourier transforms, which in recent work by Lee-Thorp et al.

(2021) have been shown to greatly improve efficiency and yet achieve comparable or, in certain scenarios, superior performance.


The relatively few parameters involved, and hence faster training times, compared to other deep learning methods makes  $\tau_2$  an attractive architecture for potentially combining with active learning methods or even off-line retraining should new data become available. With the small model size,  $\tau_2$  should also appeal to upcoming brokering systems such as FINK (Möller et al. 2021), ANTARES (Matheson et al. 2021) etc. that benefit from low latency and fast inference times when put into production. As we touched on in Section 2.3, the current architecture forgoes the additional decoder found in the original transformer architecture (Vaswani et al. 2017) that applies a casual mask to the input. However, the inclusion of such a mask would provide a natural mechanism within the time-series transformer architecture for early light curve classification, which provides another avenue of future work.

The time-series transformer,  $\tau_2$ , minimises the reliance of expert feature selection. Moving away from feature engineering allows the model the freedom to discover patterns that are missed by humans but yet provide powerful discriminative information for classification. The architecture, by virtue of CAMs (class activation maps), offers up a helpful tool for interpretability by inspecting the importance of both light curves and any additional features that are included<sup>5</sup>. It is hoped that with the introduction of the attention mechanism to the field of astronomical photometric classification, further studies will build on this work to improve our ability to attend to the night sky.

<sup>5</sup> The reader is reminded of alternative interpretability techniques that may provide better explainability such as RISE (Petsiuk et al. 2018).




(a) Class activation map for a Supernova Type Ia drawn from the test set.



(b) Class activation map for a Supernova Type II drawn from the test set.

**Figure 11.** Class activation maps (CAM) for two types of Supernova drawn from the test set, with lightcurves for bands *giruyz* over-plotted. For visualisation purposes, a min-max scaling is applied to the class activations as well as a normalisation to each CAM such that  $\sum_l M_c(l) = 1$ , such that the relative activation weights can be interpreted as a percentage. The left hand side depicts the percentage of activation weight attributed to each position in the sequence, while on the right hand side we show the percentage activation weights associated with any additional features that have been added; in our case redshift and redshift error. Notice that for both examples the activation weight is generally low before the initial rise of the light curve, larger at the rise, with the strongest weight around the peak. Moderate weights are observed in the tail, presumably to detect any secondary peak, with the weights becoming insignificant once the light curve falls back to zero. The influence of redshift and information can be seen on the right hand side, with  $\sim 0.6\%$  and  $\sim 1.4\%$  of the total activation weight being attributed these additional features for each object, respectively.



**Figure 12.** Distribution of activation weights for redshift and redshift error for all classes combined. This plot is constructed for all classes combined (minimal variability was observed across classes when plotted separately). The mean redshift and redshift error activation weights are both 0.96. In the centre of the plotted distribution we plot letter-value plots (Hofmann et al. 2011) that are better suited to large datasets such as this. The middle box contains 50% of the data, with the median indicated by a line at the midpoint. The next smaller boxes combined contain 25% of the data, with each successive level outward continuing in this fashion containing half of the remaining data.

**ACKNOWLEDGEMENTS**

TA would like to thank Catarina S. Alves for the helpful discussions. This work was partially enabled by funding from the UCL Cosmoparticle Initiative for use of Hypatia compute facilities. The authors acknowledge the use of the UCL Myriad High Performance Computing Facility (Myriad@UCL), and associated support services, in the completion of this work. The work was also supported by the Science and Technology Facilities Council (STFC) Centre for Doctoral Training in Data Intensive Science at UCL.

**DATA AVAILABILITY**

All data referenced here can be found freely available online. For the original PLAsTiCC dataset, the reader is advised to explore details laid out in The PLAsTiCC team et al. (2018) and Team & Modelers (2019).

**REFERENCES**

Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker,

- P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., & Zheng, X., 2015. TensorFlow: Large-scale machine learning on heterogeneous systems, Software available from tensorflow.org.
- Akiba, T., Sano, S., Yanase, T., Ohta, T., & Koyama, M., 2019. Optuna: A next-generation hyperparameter optimization framework, in *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 2623–2631.
- Alves, C. S., Peiris, H. V., Lochner, M., McEwen, J. D., Allam, T., Biswas, R., Collaboration, L. D. E. S., et al., 2022. Considerations for optimizing the photometric classification of supernovae from the rubin observatory, *The Astrophysical Journal Supplement Series*, **258**(2), 23.
- Ambikasaran, S., Foreman-Mackey, D., Greengard, L., Hogg, D. W., & O’Neil, M., 2015. Fast direct methods for gaussian processes, *IEEE transactions on pattern analysis and machine intelligence*, **38**(2), 252–265.
- Ba, J. L., Kiros, J. R., & Hinton, G. E., 2016. Layer normalization, *arXiv preprint arXiv:1607.06450*.
- Bahdanau, D., Cho, K., & Bengio, Y., 2014. Neural machine translation by jointly learning to align and translate, *arXiv preprint arXiv:1409.0473*.
- Bergstra, J., Bardenet, R., Bengio, Y., & Kégl, B., 2011. Algorithms for hyperparameter optimization, in *25th annual conference on neural information processing systems (NIPS 2011)*, vol. 24, Neural Information Processing Systems Foundation.
- Boone, K., 2019. Avocado: Photometric classification of astronomical transients with gaussian process augmentation, *The Astronomical Journal*, **158**(6), 257.
- Branco, P., Torgo, L., & Ribeiro, R., 2015. A survey of predictive modelling under imbalanced distributions, *arXiv preprint arXiv:1505.01658*.
- Brownlee, J., 2020. Tour of evaluation metrics for imbalanced classification.
- Brunel, A., Pasquet, J., PASQUET, J., Rodriguez, N., Comby, F., Fouchez, D., & Chaumont, M., 2019. A cnn adapted to time series for the classification of supernovae, *Electronic imaging*, **2019**(14), 90–1.
- Butkevich, A., Berdyugin, A., & Teerikorpi, P., 2005. Statistical biases in stellar astronomy: the malquist bias revisited, *Monthly Notices of the Royal Astronomical Society*, **362**(1), 321–330.
- Chai, Y., 2020. Counting the number of parameters in deep learning.
- Charnock, T. & Moss, A., 2017. Deep recurrent neural networks for supernovae classification, *The Astrophysical Journal Letters*, **837**(2), L28.
- Cheng, J., Dong, L., & Lapata, M., 2016. Long short-term memory-networks for machine reading, *arXiv preprint arXiv:1601.06733*.
- Fawaz, H. I., Forestier, G., Weber, J., Idoumghar, L., & Muller, P.-A., 2019. Deep learning for time series classification: a review, *Data Mining and Knowledge Discovery*, **33**(4), 917–963.
- Gabruseva, T., Zlobin, S., & Wang, P., 2020. Photometric light curves classification with machine learning, *Journal of Astronomical Instrumentation*, **9**(01), 2050005.
- Goodfellow, I., Bengio, Y., & Courville, A., 2016. *Deep Learning*, MIT Press, <http://www.deeplearningbook.org>.
- Grosse, R. & Ba, J., 2019. Neural networks and deep learning.
- Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., del Río, J. F., Wiebe, M., Peterson, P., Gérard-Marchant, P., Sheppard, K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C., & Oliphant, T. E., 2020. Array programming with NumPy, *Nature*, **585**(7825), 357–362.
- He, H. & Ma, Y., 2013. *Imbalanced learning: foundations, algorithms, and applications*, John Wiley & Sons.
- Hložek, R., Ponder, K., Malz, A., Dai, M., Narayan, G., Ishida, E., Allam Jr, T., Bahmanyar, A., Biswas, R., Galbany, L., et al., 2020. Results of the photometric lsst astronomical time-series classification challenge (plasticc), *arXiv preprint arXiv:2012.12392*.
- Hochreiter, S., Bengio, Y., Frasconi, P., Schmidhuber, J., et al., 2001. Gradient flow in recurrent nets: the difficulty of learning long-term dependencies.
- Hofmann, H., Kafadar, K., & Wickham, H., 2011. Letter-value plots: Boxplots for large data, *The American Statistician*.
- Ishida, E. E. & de Souza, R. S., 2013. Kernel pca for type ia supernovae photometric classification, *Monthly Notices of the Royal Astronomical Society*, **430**(1), 509–532.
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., Abel, B., Acosta, E., Allsman, R., Alonso, D., AlSayyad, Y., Anderson, S. F., Andrew, J., et al., 2019. Lsst: from science drivers to reference design and anticipated data products, *The Astrophysical Journal*, **873**(2), 111.
- Karpenka, N. V., Feroz, F., & Hobson, M., 2013. A simple and robust method for automated photometric classification of supernovae using neural networks, *Monthly Notices of the Royal Astronomical Society*, **429**(2), 1278–1285.
- Kazemnejad, A., 2019. Transformer architecture: The positional encoding.
- Kessler, R., Bassett, B., Belov, P., Bhatnagar, V., Campbell, H., Conley, A., Frieman, J. A., Glazov, A., González-Gaitán, S., Hložek, R., et al., 2010. Results from the supernova photometric classification challenge, *Publications of the Astronomical Society of the Pacific*, **122**(898), 1415.
- Kessler, R., Narayan, G., Avelino, A., Bachelet, E., Biswas, R., Brown, P., Chernoff, D., Connolly, A., Dai, M., Daniel, S., et al., 2019. Models and simulations for the photometric lsst astronomical time series classification challenge (plasticc), *Publications of the Astronomical Society of the Pacific*, **131**(1003), 094501.
- Kingma, D. P. & Ba, J., 2014. Adam: A method for stochastic optimization, *arXiv preprint arXiv:1412.6980*.
- Koehrsen, W., 2018. A conceptual explanation of bayesian hyperparameter optimization for machine learning.
- Lee-Thorp, J., Ainslie, J., Eckstein, I., & Ontanon, S., 2021. Fnet: Mixing tokens with fourier transforms, *arXiv preprint arXiv:2105.03824*.
- Lin, M., Chen, Q., & Yan, S., 2013. Network in network, *arXiv preprint arXiv:1312.4400*.
- Lochner, M., McEwen, J. D., Peiris, H. V., Lahav, O., & Winter, M. K., 2016. Photometric supernova classification with machine learning, *The Astrophysical Journal Supplement Series*, **225**(2), 31.
- Luong, M.-T., Pham, H., & Manning, C. D., 2015. Effective approaches to attention-based neural machine translation, *arXiv preprint arXiv:1508.04025*.
- Madsen, A., 2019. Visualizing memorization in rnns, *Distill*, <https://distill.pub/2019/memorization-in-rnns>.
- Malz, A., Hložek, R., Allam, T., Bahmanyar, A., Biswas, R., Dai, M., Galbany, L., Ishida, E., Jha, S., Jones, D., et al., 2019. The photometric lsst astronomical time-series classification challenge plasticc: Selection of a performance metric for classification probabilities balancing diverse science goals, *The Astronomical Journal*, **158**(5), 171.
- Matheson, T., Stubens, C., Wolf, N., Lee, C.-H., Narayan, G., Saha, A., Scott, A., Soraisam, M., Bolton, A. S., Hauger, B., et al., 2021. The antares astronomical time-domain event broker, *The Astronomical Journal*, **161**(3), 107.
- Mikolov, T., Chen, K., Corrado, G., & Dean, J., 2013. Efficient estimation of word representations in vector space, *arXiv preprint arXiv:1301.3781*.
- Möller, A. & de Boissière, T., 2020. Supernnova: an open-source framework for bayesian, neural network-based supernova classification, *Monthly Notices of the Royal Astronomical Society*, **491**(3), 4277–4293.
- Möller, A., Peloton, J., Ishida, E. E., Arnault, C., Bachelet, E., Blaineau, T., Boutigny, D., Chauhan, A., Gangler, E., Hernandez, F., et al., 2021. Fink, a new generation of broker for the lsst community, *Monthly Notices of the Royal Astronomical Society*, **501**(3), 3272–3288.
- Muthukrishna, D., Narayan, G., Mandel, K. S., Biswas, R., & Hložek, R., 2019. Rapid: early classification of explosive transients using deep learning, *Publications of the Astronomical Society of the Pacific*, **131**(1005), 118002.
- Nair, V. & Hinton, G. E., 2010. Rectified linear units improve restricted boltzmann machines, in *Icml*.
- Oord, A. v. d., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A., & Kavukcuoglu, K., 2016. Wavenet: A generative model for raw audio, *arXiv preprint arXiv:1609.03499*.
- Perlmutter, S., Aldering, G., Goldhaber, G., Knop, R., Nugent, P., Castro, P. G., Deustua, S., Fabbro, S., Goobar, A., Groom, D. E., et al., 1999. Measurements of  $\omega$  and  $\lambda$  from 42 high-redshift supernovae, *The Astrophysical Journal*, **517**(2), 565.
- Petsiuk, V., Das, A., & Saenko, K., 2018. Rise: Randomized input sampling for explanation of black-box models, *arXiv preprint arXiv:1806.07421*.

- Rasmussen, C. E., 2004. Gaussian processes in machine learning, in *Advanced lectures on machine learning*, pp. 63–71, Springer.
- Revsbech, E. A., Trotta, R., & van Dyk, D. A., 2018. Staccato: a novel solution to supernova photometric classification with biased training sets, *Monthly Notices of the Royal Astronomical Society*, **473**(3), 3969–3986.
- Riess, A. G., Filippenko, A. V., Challis, P., Clocchiatti, A., Diercks, A., Garnavich, P. M., Gilliland, R. L., Hogan, C. J., Jha, S., Kirshner, R. P., et al., 1998. Observational evidence from supernovae for an accelerating universe and a cosmological constant, *The Astronomical Journal*, **116**(3), 1009.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J., 1986. Learning representations by back-propagating errors, *nature*, **323**(6088), 533–536.
- Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D., 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization, in *Proceedings of the IEEE international conference on computer vision*, pp. 618–626.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R., 2014. Dropout: a simple way to prevent neural networks from overfitting, *The journal of machine learning research*, **15**(1), 1929–1958.
- Sutskever, I., Vinyals, O., & Le, Q. V., 2014. Sequence to sequence learning with neural networks, *arXiv preprint arXiv:1409.3215*.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., & Wojna, Z., 2015. Rethinking the inception architecture for computer vision. corr abs/1512.00567 (2015).
- Tay, Y., Dehghani, M., Bahri, D., & Metzler, D., 2020. Efficient transformers: A survey, *arXiv preprint arXiv:2009.06732*.
- Team, P. & Modelers, P., 2019. Unblinded data for plasticc classification challenge, doi: 10.5281/zenodo.2539456.
- The PLAsTiCC team, Allam Jr, T., Bahmanyar, A., Biswas, R., Dai, M., Galbany, L., Hložek, R., Ishida, E. E., Jha, S. W., Jones, D. O., Kessler, R., et al., 2018. The photometric lsst astronomical time-series classification challenge (plasticc): Data set, *arXiv preprint arXiv:1810.00001*.
- Varughese, M. M., von Sachs, R., Stephanou, M., & Bassett, B. A., 2015. Non-parametric transient classification using adaptive wavelets, *Monthly Notices of the Royal Astronomical Society*, **453**(3), 2848–2861.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I., 2017. Attention is all you need, *arXiv preprint arXiv:1706.03762*.
- Weng, L., 2018. Attention? attention!, *lilianweng.github.io/lil-log*.
- Wes McKinney, 2010. Data Structures for Statistical Computing in Python, in *Proceedings of the 9th Python in Science Conference*, pp. 56 – 61.
- Zhou, B., Khosla, A., A., L., Oliva, A., & Torralba, A., 2015. Learning Deep Features for Discriminative Localization., *CVPR*.

This paper has been typeset from a  $\text{\TeX}/\text{\LaTeX}$  file prepared by the author.