Astroconformer reads local and long-range patterns in stellar light curves.

MNRAS 000, 1–14 (2023) Preprint 19 January 2024 Compiled using MNRAS LATEX style file v3.0

Astroconformer: The Prospects of Analyzing Stellar Light Curves with Transformer-Based Deep Learning Models

Jia-Shu Pan (潘嘉书),[1,2]★ Yuan-Sen Ting (丁源森),[1,3,4] Jie Yu (余杰)[5,6,3]†
[1]Research School of Astronomy & Astrophysics, Australian National University, Cotter Rd., Weston, ACT 2611, Australia
[2]School of Astronomy and Space Science, Nanjing University, Nanjing 210093, China
[3]School of Computing, Australian National University, Acton, ACT 2601, Australia
[4]Department of Astronomy, The Ohio State University, Columbus, USA
[5]Max Planck Institute for Solar System Research, Justus-von-Liebig-Weg 3, 37077 Gttingen, Germany
[6]Heidelberg Institute for Theoretical Studies (HITS) gGmbH, Schloss-Wolfsbrunnenweg 35, D-69118 Heidelberg, Germany

arXiv:2309.16316v2 [astro-ph.SR] 18 Jan 2024

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

Stellar light curves contain valuable information about oscillations and granulation, offering insights into stars' internal structures and evolutionary states. Traditional asteroseismic techniques, primarily focused on power spectral analysis, often overlook the crucial phase information in these light curves. Addressing this gap, recent machine learning applications, particularly those using Convolutional Neural Networks (CNNs), have made strides in inferring stellar properties from light curves. However, CNNs are limited by their localized feature extraction capabilities. In response, we introduce *Astroconformer*, a Transformer-based deep learning framework, specifically designed to capture long-range dependencies in stellar light curves. Our empirical analysis centers on estimating surface gravity ($\log g$), using a dataset derived from single-quarter Kepler light curves with $\log g$ values ranging from 0.2 to 4.4. Astroconformer demonstrates superior performance, achieving a root-mean-square-error (RMSE) of 0.017 dex at $\log g \approx 3$ in data-rich regimes and up to 0.1 dex in sparser areas. This performance surpasses both K-nearest neighbor models and advanced CNNs. Ablation studies highlight the influence of receptive field size on model effectiveness, with larger fields correlating to improved results. *Astroconformer* also excels in extracting $\nu_{\rm max}$ with high precision. It achieves less than 2% relative median absolute error for 90-day red giant light curves. Notably, the error remains under 3% for 30-day light curves, whose oscillations are undetectable by a conventional pipeline in 30% cases. Furthermore, the attention mechanisms in *Astroconformer* align closely with the characteristics of stellar oscillations and granulation observed in light curves.

Key words: asteroseismology – methods: data analysis



1 INTRODUCTION

Stellar light curves serve as a comprehensive source for deciphering the diverse physical properties of stars, being shaped by multiple factors such as rotation, spot modulation, oscillations, and granulation (Aerts et al. 2010). Asteroseismology, the study of stellar oscillations, remains a benchmark approach for understanding essential stellar properties like mass, radius, and luminosity (Aerts et al. 2010). Particularly for evolved stars with solar-like oscillations, asteroseismology offers crucial insights into their evolutionary states (Bedding et al. 2011), internal rotation (Kawaler & Hostler 2005; Hall et al. 2021; Gehan et al. 2018), and magnetic fields (Fuller et al. 2015; Li et al. 2022; Stello et al. 2016).

The field of asteroseismology has experienced a significant boost over the past two decades, largely owing to the successful CoRoT (Auvergne et al. 2009) and *Kepler* (Koch et al. 2010) space missions. The original *Kepler* mission provided high-quality light curves for nearly 200,000 low- to intermediate-mass stars with visual magnitudes below 16, accumulating data over four years in a fixed northern sky field. Subsequent to *Kepler*, the K2 mission (Howell et al. 2014) provided an extension, focusing on various ecliptic fields and capturing valuable data for a broad spectrum of stars, albeit within a reduced observational window of approximately 90 days per field. The Transiting Exoplanet Survey Satellite (TESS) (Ricker et al. 2014), as *Kepler*'s successor, aims to monitor bright stars throughout the sky. While TESS data holds its own promise for advancing asteroseismology, it poses new challenges due to more limited continuous observation periods, ranging from 27 days to about a year.

Traditionally, asteroseismic analyses transform light curves into power spectra and analyze summary statistics within this framework (Aerts et al. 2010; Hekker & Christensen-Dalsgaard 2017; García & Ballot 2019; Aerts 2021). Two global seismic parameters, $\nu_{\rm max}$ and $\Delta\nu$, are frequently employed to estimate the parameters of solar-like oscillators. The frequency of maximal oscillation power, $\nu_{\rm max}$, is proportional to the acoustic cutoff frequency and is therefore linked to both surface gravity ($\log g$) and effective temperature ($T_{\rm eff}$) (Brown et al. 1991; Kjeldsen & Bedding 1995). The average frequency separation between consecutive radial order modes at a specific spherical degree, $\Delta\nu$, scales with mean stellar density as it correlates with the sound travel time through the star (Ulrich 1986).

★ E-mail: jspan@smail.nju.edu.cn
† E-mail: Jie.yu@anu.edu.au

© 2023 The Authors

Self-attention connects distant parts of a light curve directly.

However, focusing solely on these oscillation signals within the frequency domain has its drawbacks. For instance, oscillation amplitude is positively correlated with stellar luminosity, making the detection of oscillations in dwarf stars particularly challenging. Moreover, the short-period oscillations observed in main-sequence and subgiant stars necessitate high-cadence photometric observations, which is not widely accessible even for advanced missions like *Kepler* and TESS. Furthermore, not all giants display solar-like oscillations, leaving their correlational relationships with stellar properties yet to be fully understood (Hon et al. 2019). Even when applied to giants exhibiting oscillations, conventional pipelines require long observation duration to detect robust oscillations and ascertain global parameters precisely (Hekker et al. 2012). Furthermore, existing automated pipelines frequently require manual visual inspection—a task that becomes impractical considering the vast volume of data generated by missions like TESS.

The limitations become more pronounced when considering light curves from ground-based observatories. The upcoming US flagship Rubin Observatory (Ivezić et al. 2019) is slated to commence its sky survey in the next year, adopting a cadence of several days (Marshall et al. 2017). Similarly, the Zwicky Transient Facility follows a comparable 3-day cadence (Bellm et al. 2018). Such long-cadence observational strategies often exceed the typical timescales of stellar oscillations, except for long period variables (Yu et al. 2020), thereby rendering the detection of oscillations less viable. These considerations, along with the challenges posed by high-noise observations like those encountered in TESS data and varying observational cadences, underline the need for innovative methodologies for robust asteroseismic analysis.

Stellar light curves offer information that extends beyond oscillations and can still be harnessed within the Nyquist limit, even in long-cadence surveys like those of the Rubin Observatory. For instance, stellar granulation offers an alternative avenue for characterizing stars using only long-cadence data (Mathur et al. 2011; Kallinger et al. 2016). Emerging from the convection layer, granulation power correlates with stellar mass and provides insights into a star's evolutionary stage (Mathur et al. 2011). Previous studies have quantified the scaling relationship between granulation and global oscillation parameters (Kallinger et al. 2014). Notably, granulation consists signal on a broader range of timescales compared to oscillations. Consequently, it has been suggested as a more effective means of studying the properties of subgiants and dwarfs (Bugnet et al. 2018). In other words, light curves contain subtle features that, if fully exploited, can reveal additional aspects of stellar characteristics.

In light of this, in this study, we introduce a Transformer-based deep learning model, named *Astroconformer*, to analyze stellar light curves. Originating in natural language processing (NLP), the Transformer model utilizes a self-attention mechanism to calculate the correlations—or 'attention'—across all inputs in a sequence (Vaswani et al. 2017). We will demonstrate that Transformer-based methods can more effectively extract essential information from stellar light curves compared to other machine learning approaches. As a proof of concept, we will focus on inferring $\log g$ and $\nu_{\rm max}$ of stars from their light curves.

2 RELEVANT STUDIES AND MOTIVATION

In recent years, the field of asteroseismology has increasingly turned to machine learning techniques to glean additional physical insights from stellar light curves, which contain more than just information from stellar oscillations. For instance, Ness et al. (2018) used polynomial ridge regression on the autocorrelation function of light curves to attain a $\log g < 0.1$ dex precision for red giants with $\log g$ values between 2.0 and 3.5 dex. Building upon this work, THE SWAN (Sayeed et al. 2021) employed a k-Nearest Neighbors (k-NN) approach on the power spectra of long-cadence *Kepler* light curves.

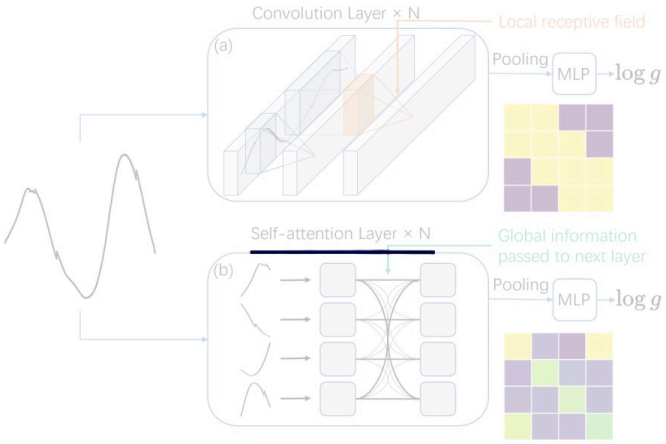


Figure 1. The figure contrasts the limited receptive field of CNNs with the global receptive field afforded by self-attention mechanisms, especially in the context of long sequences like 4-year *Kepler* light curves. (a) In CNNs, multiple convolutional layers with small receptive fields are stacked to achieve a broader receptive field in deeper layers. The heatmap illustrates the pairwise correlations between segments, showing that each segment primarily incorporates information only from its immediate neighbors. (b) Self-attention, on the other hand, calculates the correlations between all segments of the light curve and incorporates this global information into its output.

Convolutional Neural Networks (CNNs) have been provedalso been utilized in asteroseismology, particularly for their proficiency in tasks such as image recognition (Szegedy et al. 2014). For example, Hon et al. (2018) employed CNNs to scrutinize images of folded power spectra, and successfully classified the evolutionary stages of stars with high accuracy. Moreover, Blancato et al. (2022) employed a one-dimensional CNN to single-quarter *Kepler* light curves and recovered $\log g$ of red giants to 0.06 dex, using asteroseismic labels from Yu et al. (2018). Recurrent Neural Networks (RNNs), known for their sequence modelling ability, have also been implemented for *Kepler* light curves (Hinners et al. 2018).

Nevertheless, these data-driven methodologies come with their own sets of limitations in both data handling and modeling. For example, the k-NN technique used by THE SWAN relies on k nearest neighbors for inference. This approach may face challenges in terms of generalization when dealing with sparser regions of the sample space. While deep learning methods typically exhibit better generalization capabilities by adaptively learning features specific to the task at hand (Chatterjee & Zielinski 2022), the deep learning architectures so far adopted in asteroseismology, such as CNNs and RNNs, also have inherent limitations. Specifically, CNNs are limited by their local receptive fields, capturing only localized spatial information. RNNs, on the other hand, find it challenging to maintain information over long sequences. To mitigate the issue in CNNs, multiple convolutional layers are required to capture interactions between long-range information. This limitation of CNNs is illustrated in Fig. 1.

Much of the existing analysis has focused on the power spectra of light curves. This focus partially arises because CNNs excel at extracting local information, making them more suited for operating in

the frequency domain rather than capturing long-range correlations in time series data. However, relying on power spectra can lead to information loss, particularly when the underlying process is not Gaussian. Power spectra discard the complex phases of the data, which contain important frequency correlations (Toutain & Appourchaux 1994; Benomar et al. 2009; Gruberbauer et al. 2009). Since solar-like oscillations are stochastically excited and intrinsically damped by surface turbulent convection (Chaplin & Miglio 2013), they may include subtle non-Gaussian signals.

These limitations—namely, the challenge of directly handling time series data in asteroseismology and the absence of appropriate inductive biases in CNNs for capturing long-range correlations—have led us to investigate Transformer-based models. As we will discuss in subsequent sections, Transformers are inherently designed to capture long-range correlations, making them a more suitable choice for analyzing time series data. Utilizing Transformers allows us to work directly with observed light curves in the time domain, minimizing information loss and obviating the need for additional post-processing steps.

3 ASTROCONFORMER: A TRANSFORMER-BASED METHOD TO ANALYZE STELLAR LIGHT CURVES

In this section, we introduce *Astroconformer*, a model inspired by the Conformer architecture as presented in Gulati et al. (2020). Given that the timescale of stellar oscillations and granulation can vary from minutes for main-sequence stars to hundreds of days for most luminous red giants, depending on surface gravity, an effective inference model must capture both local and global features for accurate log g prediction from light curves. The Conformer model integrates a variety of deep learning modules, combining self-attention layers for long-range information capture with convolutional layers for local feature extraction. Our PyTorch implementation of *Astroconformer* is publicly available[1].

3.1 Self-attention mechanism

Sequence representation continues to be a vibrant area of research in deep learning, with applications spanning natural language processing (NLP), computer vision (CV) (Dosovitskiy et al. 2021), and time-series forecasting (Nie et al. 2023). Unlike images, where key information is often local, sequences frequently require the extraction of long-range correlations for a comprehensive interpretation.

Traditional CNNs are notably limited in capturing these long-range dependencies effectively, as shown in Fig. 1. The self-attention mechanism, first introduced by Vaswani et al. (2017), addresses this limitation. Thanks to its remarkable ability to model long-range correlations, self-attention has become a cornerstone in state-of-the-art NLP models like BERT (Devlin et al. 2019) and GPT-3 (Brown et al. 2020).

The crux of self-attention lies in first analyzing the correlation between timestamps (up to some linear combination) and then using this correlation to either amplify or diminish the contributions from individual timestamps. Given that the correlation matrix is calculated across all timestamps, the model inherently incorporates a strong inductive bias towards considering these correlations.

To elaborate, a schematic representation of this process is depicted in Fig. 2 (a), and can be summarized as follows:

1. Self-attention starts by duplicating the input sequence and linearly projecting each copy into one of three forms: query, key, and value. These projections are steered by learnable parameters.
2. Each projection serves a unique purpose: the query seeks relevant information within the sequence, and the key aims to align with this query. This alignment is achieved through a dot product operation, measuring the correlation between various keys and queries. Pairs with high dot products are deemed strongly correlated and are thus allocated "attention." Importantly, the dot product spans the entire sequence, allowing for attention to be distributed between any two timestamps, regardless of their separation.
3. The value vector then leverages these attention weights to decide the proportion of the original sequence's information to include in the output. As a result, the output vector, serving as the new sequence representation, is a weighted compilation based on these attention mechanisms.

To operationalize this, we consider an input sequence $\boldsymbol{X} \in \mathbb{R}^{L\times d}$, where L represents the sequence length and d denotes the dimension of each vector. For notational convenience, we'll use the subscript $L \times d$ to specify the dimensions for all tensors henceforth. Self-attention initiates the process by linearly projecting the input $\boldsymbol{X}_{L\times d}$ into three separate forms: query $\boldsymbol{Q}_{L\times d}$, key $\boldsymbol{K}_{L\times d}$, and value $\boldsymbol{V}_{L\times d}$. Mathematically, these projections are defined as:

$$\boldsymbol{Q}_{L\times d} = \boldsymbol{X}_{L\times d}\boldsymbol{W}^{\boldsymbol{Q}}_{d\times d}, \tag{1}$$

$$\boldsymbol{K}_{L\times d} = \boldsymbol{X}_{L\times d}\boldsymbol{W}^{\boldsymbol{K}}_{d\times d}, \tag{2}$$

$$\boldsymbol{V}_{L\times d} = \boldsymbol{X}_{L\times d}\boldsymbol{W}^{\boldsymbol{V}}_{d\times d}, \tag{3}$$

where $\boldsymbol{W}^{\boldsymbol{Q}}_{d\times d}, \boldsymbol{W}^{\boldsymbol{K}}_{d\times d}, \boldsymbol{W}^{\boldsymbol{V}}_{d\times d}$ are learnable matrices.

The next step involves calculating the similarity between the query and key vectors by taking an inner dot product over the d dimension, which yields an $L\times L$ similarity matrix (α_{ij}). To endow these similarities with a probabilistic interpretation, a softmax function is applied to responses belonging to each query to ensure the attention weights are both positive and normalized:

$$\alpha_{ij} \mapsto \frac{e^{\alpha_{ij}}}{\sum_{k=1}^{L} e^{\alpha_{ik}}}.$$

We call this normalized similarity matrix $\boldsymbol{A}_{L\times L} = (A_{ij})$ attention map, where each entry A_{ij} is the normalized dot product between the i-th query vector $\boldsymbol{q}^i_d$ and the j-th key vector $\boldsymbol{k}^j_d$. This key feature sets self-attention apart from convolutional layers by measuring similarity across all timestamps.

Finally, a new representation for $\boldsymbol{X}$ is generated by matrix-multiplying the attention matrix A with the value matrix $\boldsymbol{V}$. This operation effectively weights a linearly transformed version of the original input data using the computed attention. Formally, for each self-attention layer, the input $\boldsymbol{X}$ undergoes the following transformation:

$$\boldsymbol{X} \mapsto \mathrm{softmax}(\boldsymbol{Q}\boldsymbol{K}^T)\boldsymbol{V}.$$

The resulting vectors extract specialized information from the entire sequence, taking into account which pairs of timestamps should be allocated additional "attention."

In practice, it has been found (Vaswani et al. 2017) that normalizing with the attention also with the feature dimension $\sqrt{d}$ facilitates the training of the neural networks. Therefore, self-attention, including in this study, is implemented as

$$\boldsymbol{X} \mapsto \mathrm{softmax}\left(\frac{\boldsymbol{Q}\boldsymbol{K}^T}{\sqrt{d}}\right)\boldsymbol{V}. \tag{4}$$

[1] https://github.com/panjiashu/Astroconformer

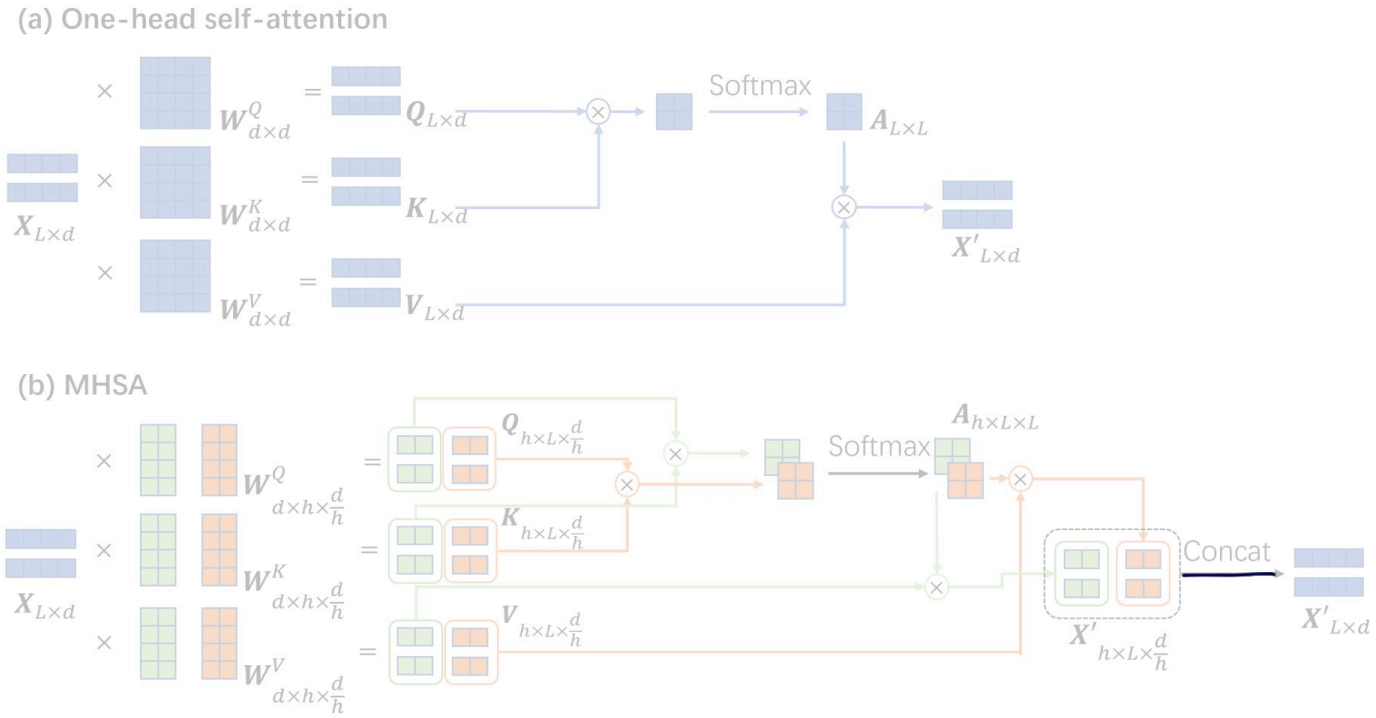


Figure 2. A schematic illustration of the self-attention mechanism within Transformer models and its capability for extracting long-range information. The top panel depicts a basic single-head self-attention mechanism. In this setup, the input is duplicated into three copies, each subjected to distinct linear transformations defined by the learnable matrices $\boldsymbol{W}^{\boldsymbol{Q}}$, $\boldsymbol{W}^{\boldsymbol{K}}$, and $\boldsymbol{W}^{\boldsymbol{V}}$. Two of these copies, $\boldsymbol{Q}$ and $\boldsymbol{K}$, are used for an inner dot product to compute similarities between different timestamps. These computed attention values are subsequently merged with the remaining $\boldsymbol{V}$ copy to produce a new representation of the input sequence. The bottom panel showcases an extension to multi-head self-attention (MHSA). While retaining the same number of learnable parameters, MHSA partitions the linear transformations into separate blocks. The cross-matching between these blocks enables the capture of diverse correlations within sequences, accommodating different types of relevancy, such as varying time scales.

This straightforward approach, which involves linear transformations and computing attention via dot products, is termed 'single-head' self-attention. While efficient, this model is somewhat limited in its expressiveness due to its reliance on a single set of weights. To enhance its versatility, the concept is extended to Multi-Head Self-Attention (MHSA), enabling the model to capture various types of correlations within the sequence.

In MHSA, h separate self-attention operations are performed, each with its own uniquely parameterized $\boldsymbol{W}^{\boldsymbol{Q}}, \boldsymbol{W}^{\boldsymbol{K}}, \boldsymbol{W}^{\boldsymbol{V}}$. Here, h denotes the number of heads in the model. To elucidate this, Fig. 2 illustrates how a two-head self-attention mechanism processes an input $\boldsymbol{X}_{2\times4}$. Rather than employing a single set of weights, the weight tensors $\boldsymbol{W}^{\boldsymbol{Q}}_{4\times4}, \boldsymbol{W}^{\boldsymbol{K}}_{4\times4}, \boldsymbol{W}^{\boldsymbol{V}}_{4\times4}$ are partitioned into $h = 2$ segments. The queries, keys, and values from these different segments are then correlated independently to construct a range of attention-weighted arrays $\boldsymbol{X}'$. When compared to single-head self-attention with an equivalent number of parameters, MHSA proves more adept at capturing diverse patterns, e.g. different time scales, within the sequence.

3.2 *Astroconformer* Architecture

The self-attention mechanism outlined earlier serves as the foundation of our model, *Astroconformer*. However, to fully leverage the richness of stellar light curves, *Astroconformer* integrates both convolutional layers and MHSA layers. The convolutional layers are adept at capturing local, short-range information, while the MHSA layers excel at extracting long-range correlations. In the sections that follow, we will briefly discuss the architecture of *Astroconformer* and its essential components. Detailed explanations of the training techniques specific to *Astroconformer* are deferred to Appendix A.

Our primary focus is on individual quarters of *Kepler* data, which serve as the input for *Astroconformer*. Each quarter of *Kepler* data spans approximately 90 days and features long-cadence data collected at intervals of 29.4 minutes. After the data preprocessing, as described in Section 4, the light curves are resampled at 30-minute intervals. Consequently, a 90-day light curve consists of 4320 timestamps. During the training process, we randomly crop segments containing 4,000 timestamps from each quarter's light curve. This strategy allows for parallel computation within each batch and serves a dual purpose as a form of data augmentation. By increasing the diversity of the input data, random cropping helps mitigate challenges associated with data scarcity.

3.2.1 *Embedding*

Inspired by the Vision Transformer (Dosovitskiy et al. 2021), which divides an image into fixed-size patches and employs a Multi-Layer Perceptron (MLP) to project each into a high-dimensional space, *Astroconformer* adopts a similar approach for transforming stellar light curves into sequences of vectors. As depicted in Fig. 3, the input *Kepler* data first passes through a patch-embedding layer, where the light curve is segmented into 200 patches, each comprising 20 timestamps. These patches are then converted into 128-dimensional vectors using an MLP denoted as E. This transformation restructures the original one-dimensional input vector $\mathbb{R}^{4000}$ into a sequence of

Global attention and local convolutions work together before pooling.

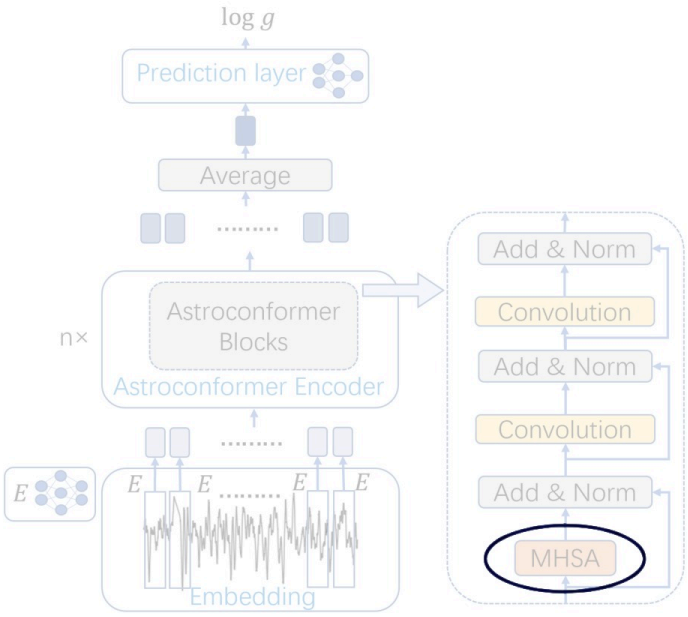


Figure 3. Architecture of *Astroconformer*. The icon labeled E represents the MLP employed in the model. Stellar light curves are partitioned into patches of size 20 (corresponding to a time span of 10 hours), each of which is then transformed into a vector via a fully connected layer denoted by E. These vectors are processed through the *Astroconformer* encoder to extract local and global features. The *Astroconformer* architecture consists of multiple such blocks, each incorporating a learnable MHSA module followed by two convolutional modules. To facilitate training, skip connections and subsequent layer normalization are applied between every pair of adjacent modules within each *Astroconformer* block. Vectors output by the *Astroconformer* encoder are aggregated through average pooling to produce the final representation of the entire light curve. This representation vector is subsequently fed into a final MLP for predicting the log g of stars.

vectors with dimensions $\mathbb{R}^{200\times128}$. Here, the first dimension of 200 can be interpreted as a coarse temporal axis, while the second 128-dimensional axis represents the localized features captured by the MLP E. When dealing with light curves that differ in length from the standard 4,000 timestamps, Astroconformer consistently follows the approach of segmenting these curves into patches of 20 timestamps. For instance, with Kepler Q1 data (33 days, 1,600 timestamps), the curve is divided into 80 patches, each individually embedded into a 128-dimensional space. No padding values at the end are introduced to the input.

This embedding serves dual objectives. Firstly, the MLP E focuses on extracting localized, short-range features, covering spans of 20 timestamps or approximately 10 hours. Secondly, by mapping the input into a 128-dimensional feature space, the model is well-positioned to exploit the power of MHSA, enabling it to cross-correlate different features across various time spans.

3.2.2 *Astroconformer Encoder*

The (200, 128)-dimensional vector sequence, produced by the embedding layer, is subsequently inputted into the *Astroconformer* encoder. This encoder is designed to simultaneously capture both global long-range and localized short-range information from the sequence of vectors.

In our architecture, the *Astroconformer* encoder is made up of multiple *Astroconformer* blocks. Each of these blocks comprises an 8-head MHSA layer and two convolutional modules. Each convolutional module includes a convolutional layer, batch normalization (Ioffe & Szegedy 2015), and the Swish activation function (Ramachandran et al. 2017). The settings for each convolutional layer are a kernel size of 3, a stride of 1, and same padding. The number of channels in each embedding vector remains constant at 128 throughout the entire forward process of the encoder. Our results indicate that the inclusion of convolutional layers following each MHSA layer enhances performance. This is possibly because MHSA modules, while capable of capturing long-range correlations, lack inherent assumptions about data locality, which often necessitates larger datasets for optimal learning compared to CNNs. In our specific scenario, the relatively limited dataset (∼14,000 stars) could have made it difficult for a model solely reliant on MHSA layers to converge satisfactorily, motivating our incorporation of convolutional layers.

Compared to the original Conformer architecture (Gulati et al. 2020), we have omitted the feed-forward modules within the conformer block, as our experiments show that they add a significant number of parameters without substantially improving performance. Focusing the primary parameters on modules that enable meaningful interactions between patches appears to be a more parameter-efficient strategy. Additionally, we refrained from employing sparse convolutions like depthwise and pointwise convolutions to ensure a fair comparison with ResNet-18, as discussed in Section 5.2.

It is worth noting that MHSA is inherently agnostic to the sequence order of its input. To address this limitation and make sequence order information available to the model, we employ Rotary Positional Encoding (RoPE) (Su et al. 2021) in each MHSA module. Briefly, RoPE compensates for the MHSA's lack of sensitivity to input order by multiplying the query and key vectors with complex phases that are proportional to their positions, $\boldsymbol{q}'_m = \boldsymbol{q}_m e^{i\omega m}$, $\boldsymbol{k}'_n = \boldsymbol{k}_n e^{i\omega n}$, where ω is a fixed frequency parameter. This ensures that their relative positions are encoded into the resulting inner products $\langle \boldsymbol{q}'_m, \boldsymbol{k}'_n \rangle = \boldsymbol{q}_m \cdot \boldsymbol{k}_n e^{i\omega(m-n)}$[2].

3.2.3 *Pooling and Prediction Layer*

The Embedding layer and the *Astroconformer* encoder together generate a final patch embedding sequence with dimensions 200×128 derived from the input light curve. We apply an average pooling to patch embeddings to the temporal dimension and obtain a final 128-dimensional embedding of the light curve (Fig. 3). Since we are dealing with individual patches of the light curves and aggregating all the information through average pooling, this enables *Astroconformer* to cope with light curves of variable lengths.

Subsequently, this 128-dimensional final embedding is processed through a two-layer MLP with a dropout probability of 0.3. This serves as the final layer tasked with predicting log g.

4 DATA

4.1 Sample Selection

In this study, we leverage two data sets, Sayeed et al. (2021), and Yu et al. (2018) to facilitate a direct comparison between *Astroconformer* and other methods (Fig. 4).

Yu et al. (2018) systematically characterized solar-like oscillations and granulation for 16,094 oscillating red giants, using end-of-mission long-cadence data. These stars feature log g values between

[2] We refer to the original paper for the details, including the choice of ω.

Preprocessing preserves amplitude ranking while making light curves easier to learn.

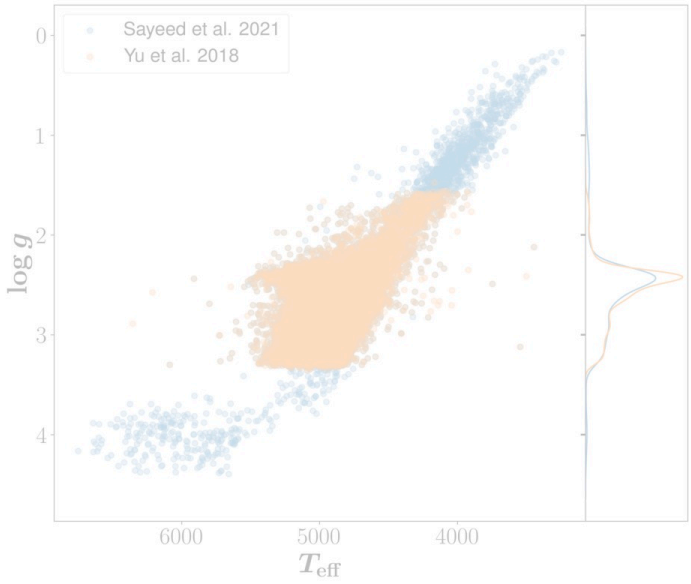


Figure 4. Distribution of surface gravity versus effective temperature for our two data sets. The data set from Yu et al. (2018) consists of 16,094 giant stars with $\log g$ values ranging from 1.5 to 3.3. Based on Yu et al. (2018), the data set from Sayeed et al. (2021) further includes turnoff stars and extends to the tip of the red giant branch (TRGB), while a data curation detailed in text is executed, leading to 14,003 stars in the data set.

Normalization keeps the relative amplitude ordering.

1.5 and 3.3. The typical uncertainties of their $\nu_{\rm max}$ and asteroseismic $\log g$ estimations are 1.6% and 0.01-0.02 dex, respectively. For the input data, we include all available quarters of light curves for each star, provided that the quarter contains at least 3,000 valid observations (quality flag = 0). This criterion results in a data set comprising 147,731 quarters of light curves for 15,874 red giants. This high-quality data is used for our training data as well as the comparison done in Section 5.3.

To further compare with the state-of-the-art machine learning method in asteroseismology, we also include an analysis with Sayeed et al. (2021). Sayeed et al. (2021) cross referenced two asteroseismic surface gravity catalogue from Mathur et al. (2017) and Yu et al. (2018). Mathur et al. (2017) compiled asteroseismic stellar properties from 11 catalogues (see references therein). Stars supplemented by Mathur et al. (2017) include high-luminosity red giants with $\log g < 2$, along with subgiants and main-sequence stars. In Section 5.1 and Section 5.2, the same asteroseismic $\log g$ as Sayeed et al. (2021) from these two work are used for training and evaluation.

Furthermore, to ensure a fair comparison with THE SWAN, we utilize an identical dataset of single-quarter long-cadence light curves as our input data. THE SWAN adopted specific data curation schemes. It excludes light curves of rotating stars, classical pulsators, eclipsing binaries, and *Kepler* exoplanet host stars. Additionally, light curves missing observations from *Kepler* Q7-11, or lacking stellar parameter estimates from Berger et al. (2020), are omitted as stellar radius data is essential for their subsequent light curve smoothing. These criteria yield a final dataset of 14,003 stars, boasting reliable 0.2–4.4 $\log g$ estimates across a range of evolutionary stages, including turnoff stars, red-giant-branch (RGB) stars, and red clump stars.

4.2 Light Curve Pre-Processing

For the preprocessing of these light curves, we adhere to the following steps:

1. Only flux values with good quality flags (flags = 0) are retained. We apply sigma-clipping within a rolling window of 50 timestamps, clipping 3σ outliers.
2. We employ a Savitzky-Golay filter to eliminate any low-frequency variations (Garcí a et al. 2011). The filter size, denoted as d, is chosen to vary with the stellar radius according to the formula $d[\mu\text{Hz}] = 0.92 + 23.03e^{-0.27\text{R}/\text{R}_\odot}$ (Sayeed et al. 2021). This ensures that the filter size increases with $\nu_{\rm max}$, allowing for the adaptive removal of long-period variations.
3. To address any missing values that may have been introduced in Step 1, we use linear interpolation to fill in the gaps. The interpolated segments are shorter than 2.1 days for 90% stars, with the largest gap being 13.6 days.
4. To further aid in model training, which benefits from inputs having similar orders of magnitude in terms of standard deviation, we normalize the light curve using the following equation:

$$\frac{x}{\text{std(x)}} \times \frac{1}{\ln(\text{maxstd}/\text{std(x)}) + 1} \mapsto x, \tag{5}$$

where maxstd represents the largest standard deviation observed across all light curves. This normalization ensures that after preprocessing, standard deviations of light curves are restricted between 0.1-1.

It is important to note that the normalization process serves another purpose: to preserve the relative amplitude ranking of stellar oscillations. The amplitude is a critical attribute of these oscillations, and this normalization method ensures that light curves with inherently larger oscillations amplitudes maintain their amplitude rankings even after processing.

5 RESULTS

In this section, we discuss the advantages of *Astroconformer* in comparison with other approaches. Section 5.1 and Section 5.2 contrast *Astroconformer* with other proposed machine learning methods like THE SWAN and CNNs, underscoring the utility of Transformer models in optimally leveraging information from stellar light curves. Following that, Section 5.3 explores a comparative analysis between *Astroconformer* and two asteroseismic pipelines. Our results suggest that *Astroconformer* could offer a precise and reliable method for predicting stellar parameters, particularly when data is constrained in timespan—an issue of critical importance for data from missions like K2, TESS.

5.1 Compared with k-NN Based Methods

We first compare the performance of *Astroconformer* with that of THE SWAN. We focus on the same dataset as used in THE SWAN, which comprises 14,003 stars with asteroseismic $\log g$ values ranging from 0.2 to 4.4 dex. Our analysis is also confined to *Kepler* data with a 90-day duration to enable a direct comparison. For *Astroconformer*, the 14,003 light curves are divided into training, validation, and test sets in proportions of 72%, 8%, and 20%, respectively. A 5-layer *Astroconformer* is trained on the training set, and the best-performing model is selected based on its performance on the validation set. This model is then used to generate predictions on the test set. We emphasise that our training labels are from the Mathur et al. (2017)

Gravity errors are smallest near the well-sampled part of the training set.

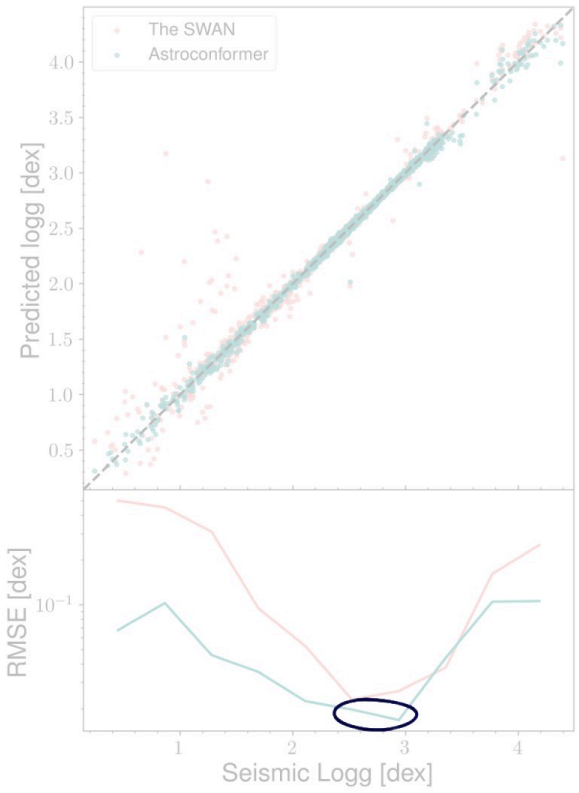


Figure 5. A comparative study of *Astroconformer* and the k-nearest neighbor-based method employed in THE SWAN. *Astroconformer* demonstrates enhanced generalization and yields fewer outlier inferences, particularly in the upper giant branch. It also consistently outperforms THE SWAN in terms of RMSE across the full $\log g$ range. The top and middle panels show the inferred $\log g$ values obtained from *Astroconformer* and THE SWAN, respectively, plotted against the asteroseismic $\log g$. The bottom panel illustrates the running mean in RMSE (with a window size of 0.4 dex) for both algorithms.

The 0.017 dex minimum is local, not the error across the whole sample.

and Yu et al. (2018), and we are only evaluating the test predictions by THE SWAN.

To facilitate a direct comparison, we also evaluate THE SWAN's test prediction on the same test set. Note that *Astroconformer* operates on a smaller training set compared to THE SWAN. In its approach to inferring surface gravity, THE SWAN utilizes a linear regression method based on a star's nearest neighbors, effectively implementing leave-one-out cross-validation. Conversely, *Astroconformer* sets aside 20% of the data as a test set.

As for performance metrics, we opt for the RMSE between the predictions and the asteroseismic $\log g$ values instead of the median absolute deviation $\sigma_{\rm mad}$ employed by THE SWAN. This choice is motivated by our observation that $\sigma_{\rm mad}$—which is less sensitive to outliers due to the use of the median—may not fully capture the nuances of the dataset.

Fig. 5 illustrates the performance comparison between *Astroconformer* and THE SWAN. THE SWAN's RMSE hovers around 0.03 dex at $\log g = 3$, and then increases to between 0.3 and 0.5 dex at both the subgiant phase and the upper part of the red giant branch. In contrast, *Astroconformer* consistently outperforms THE SWAN across the entire $\log g$ spectrum, achieving a root-mean-square error (RMSE) of at most 0.1 dex in the 0.2–4.4 $\log g$ range. It performs particularly well around $\log g = 3$, where the training sample is denser, with a minimum RMSE of 0.017 dex.

The result is perhaps unsurprising, k-NN algorithms often falter in high-dimensional feature spaces and are highly sensitive to sampling density. This is evident at the upper giant branch, where the training sample is sparse, leading THE SWAN to incur more outliers and reduced robustness. In contrast, deep learning methods like *Astroconformer* aim to generalize by learning underlying features, making them more robust even when training samples are sparse.

It is noteworthy that *Astroconformer* successfully determines $\log g$ to an RMSE < 0.1 dex for dwarf stars with $\log g > 4$ and high-luminosity red giants with $\log g < 1$. Considering that oscillation signals in dwarf stars are too brief for detection in Kepler long cadence data, this achievement indirectly demonstrates *Astroconformer*'s capability to utilize granulation information within the dataset. Similarly, for giants near the tip of the red giant branch, their extended oscillation periods > 10 days can not be resolved well by the typical 90-day span. *Astroconformer*'s ability to accurately infer $\log g$ for these stars further indicates its proficiency in leveraging information beyond just the oscillation peak.

Finally, we note that a certain degree of scatter may arise from uncertainties in the training labels. Given that the asteroseismic $\log g$ values have an uncertainty of 0.01–0.02 dex, this factor could contribute to the error. However, considering that this uncertainty is at most half of the prediction errors observed in both algorithms, we anticipate its impact on our overall analysis to be minimal. For a discussion on addressing label uncertainty, please refer to Section 6.2.

5.2 Transformer Models versus Convolutional Neural Networks

We now turn our attention to comparing *Astroconformer* with various time-domain deep learning models. As mentioned in the introduction, this study is motivated by the hypothesis that subtle yet crucial information in stellar light curves may be concealed at longer time scales.

In our ablation study, we study the performance of five distinct models, each characterized by varying extents of receptive fields: (a) *Astroconformer*; (b) a variant of *Astroconformer* employing only MHSA; (c) *Astroconformer* devoid of MHSA; (d) the ResNet-18 model; and (e) a modified ResNet architecture, which we designate as 'Short-Sighted ResNet' (Fig. 6).

To ensure a rigorous comparison, each model under consideration adopts an architecture similar to Astroconformer. This architecture incorporates identical components such as MHSA, convolutional modules, an average pooling layer, and a prediction layer, as delineated in Section 3.2. For each module, we either substitute the MHSA with convolutional modules or vice versa. Where appropriate, we also adjust the number of blocks to maintain roughly equivalent counts of learnable parameters across the neural networks. An exception to this is ResNet-18, which remains consistent with its established architecture as cited in He et al. (2015). The specifics are itemized in Table 1.

As illustrated in Table 1, the 'without convolution' variant (b) replaces all convolutional modules with MHSA modules while retaining the same global receptive field inherent to *Astroconformer*. The 'without MHSA' variant (c) substitutes all MHSA modules with convolutional modules, yielding a receptive field of 15 days. Variant (e) 'Short-Sighted ResNet,' serves as an extended modification of the without MHSA variant (c), but omits the embedding layer. By accepting raw timestamps instead of patch embeddings as inputs, its receptive field is truncated to 18 hours. Lastly, for benchmarking purposes, we include ResNet-18, known for its proficiency in time series (Ismail Fawaz et al. 2019), featuring a 4.5-day receptive field.

We subject all models to 10-fold cross-validation, using the same dataset outlined in Section 5.1, which includes 14,003 high-quality

The ablation favors global attention plus a local convolutional bias.

Table 1. The five models conducted in the ablation study, and their constituent submodules within each block, the total number of blocks, and the utilization of an embedding layer.

Model	Block	# Blocks	Embedding
Astroconformer	MHSA, Conv, Conv	5	✓
without Conv	MHSA, MHSA, MHSA	5	✓
without MHSA	Conv, Conv, Conv	6	✓
Short-sighted ResNet	Conv, Conv, Conv	6	✗
ResNet-18	-	-	✗

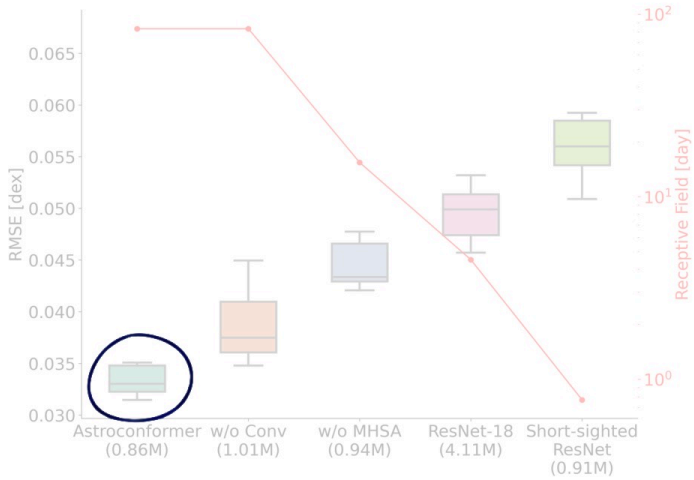


Figure 6. Ablation study illustrating the RMSE for $\log g$ inference across various models with different receptive fields. The respective numbers of parameters for each model are annotated alongside their names. The boxplot depicts the distribution of RMSE values on the test set, as assessed by 10-fold cross-validation for each model. We omit outliers more than 1.5 times the interquartile range (IQR) above the third quartile to investigate the general trend. Receptive fields for the models are indicated by red lines on a logarithmic scale. The figure reveals that as the receptive field size diminishes, the median RMSE for each model correspondingly increases. Moreover, Transformer-based models outperform a conventional CNN, ResNet-18, despite the latter having four times more parameters. Furthermore, the first two models on the plot indicate that while MHSA is crucial for performance, the incorporation of convolutional modules further enhances *Astroconformer*'s effectiveness.

Kepler samples with $\log g$ values ranging from 0.2 to 4.4. Instead of investigating running RMSE as in Section 5.1, here we use overall RMSE as the performance measure of different models. This is motivated by the fact that long-range correlation such as granulation is important for $\log g$ estimation for all stars, and a smaller overall RMSE is expected if a model can extract long-range correlation better. The overall RMSE values from the 10 cross-validation tests are presented in Fig. 6, while we omit outliers more than 1.5 times the interquartile range (IQR) above the third quartile to investigate the general trend. In this figure, the boxplot represents the distribution of RMSE values across the different cross-validation runs for each model, while the red solid line indicates the corresponding receptive field for each model. We report the median RMSE of the five models to be 0.033 dex, 0.037 dex, 0.043 dex, 0.050 dex, and 0.056 dex, respectively.

This figure illustrates the combined influence of both receptive field size and MHSA on the accuracy of $\log g$ extraction from stellar light curves. A clear performance decline is observed as the receptive field size decreases across all five models. Importantly, this decline is not affected by the number of parameters, as demonstrated by ResNet-18. Despite having four to five times more parameters than the other models, ResNet-18 still adheres to the same trend. Additional evidence of the importance of receptive field size can be seen when comparing '*Astroconformer* without MHSA' and 'Short-sighted ResNet.' The key difference between these models lies in their input: the former processes patches of the light curve, thereby having a receptive field 20 times larger than the latter, which deals directly with light curve timestamps. This expanded receptive field contributes to a significant RMSE reduction from 0.056 dex to 0.043 dex .

Similarly, the only distinction between *Astroconformer* and '*Astroconformer* without MHSA' is the incorporation of the MHSA module in the former, which enhances its ability to assimilate global information. This enhancement results in a significant decrease in RMSE from 0.043 dex to 0.033 dex for *Astroconformer*. These observations underscore the pivotal role played by receptive fields. In addition to local information like amplitudes that can be extracted by every model, long-range correlations resulting from granulation and oscillations is also important for precise $\log g$ extraction from light curves.

Finally, a comparison between the default *Astroconformer* model and the '*Astroconformer* without Conv' model demonstrates the advantages of combining convolutional modules with self-attention mechanisms. Although the underlying Transformer architecture of *Astroconformer* is theoretically capable of learning the convolution operation, the scarcity of labeled data makes the inclusion of convolutional modules particularly beneficial. Light curve data likely encompass both long-range and short-range correlations. By utilizing both convolutional modules and MHSA, *Astroconformer* is better equipped to leverage these local and global correlations for more accurate predictions. This serves to underscore the significance of receptive field size, as well as the synergistic benefits of incorporating self-attention with convolution, in the modeling of long sequences of stellar light curves.

5.3 Comparison with Asteroseismic Pipelines

In addition to comparing *Astroconformer* with other deep learning methods, this section evaluates its performance against traditional asteroseismic pipelines. Asteroseismic pipelines such as SYD typically focus on identifying $\nu_{\rm max}$ by fitting and removing granulation and the white noise background in the Fourier spectrum and searching for the peak in the oscillations power excess. To facilitate a more rigorous comparison, we will study $\nu_{\rm max}$ as the focal stellar parameter. This is driven by the understanding that asteroseismic $\log g$ are derived from the scaling laws, which also take into account external information such as $T_{\rm eff}$, which is not assumed in *Astroconformer*.

The challenge of benchmarking *Astroconformer* against asteroseismic pipelines arises from the absence of a priori ground truth labels for $\nu_{\rm max}$. To overcome this hurdle, we utilize high-quality $\nu_{\rm max}$ labels (with an uncertainty of approximately 1.6%) provided by Yu et al. (2018), based on end-of-mission (~4 years) *Kepler* data.

5.3.1 On 90-day Light Curves

An 20-layer *Astroconformer* is trained from scratch on the identical dataset in Yu et al. (2018) (refer to Appendix A for training details). We then apply both *Astroconformer* and the SYD pipeline to a test set containing 9,321 single-quarter light curves from *Kepler* Quarters 2-

Single-quarter predictions agree with long-baseline labels and vary less between quarters.

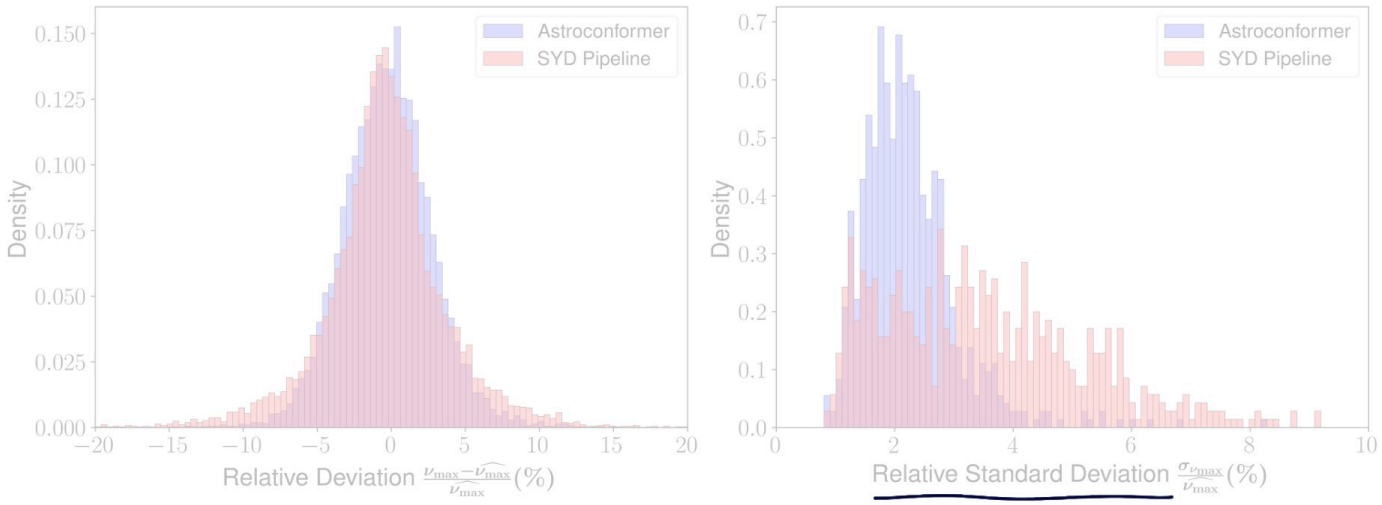


Figure 7. We compare the predictive accuracy and consistency of *Astroconformer* and the SYD pipeline using single-quarter *Kepler* light curves. *Left Panel:* The distribution of relative deviation for $\nu_{\max}$ predictions from *Astroconformer* and the SYD pipeline is shown, each based on single-quarter light curves and compared to the 'ground truth' $\widehat{\nu_{\max}}$ provided by Yu et al. (2018) using full four-year *Kepler* data. *Astroconformer* achieves a relative median absolute error of 1.93%, compared to the SYD pipeline's result of 2.10%. *Right Panel:* The relative standard deviations in predictions for each star across various quarters are displayed. *Astroconformer* has a median relative standard deviation of 2.10%, compared to the SYD pipeline's median of 3.37%.

The reference labels come from four years of Kepler data.

16 except for Q8, 12[3] (13 quarters in total) for 717 red giants[4]. We utilize two key metrics for evaluation: relative deviation and relative standard deviation, aimed at assessing prediction accuracy and internal consistency, respectively.

For relative deviation, we compute the fractional difference $(\nu_{\max} - \widehat{\nu_{\max}})/\widehat{\nu_{\max}}$, comparing the predicted $\nu_{\max}$ values obtained from single-quarter light curves against the 'ground truth' $\widehat{\nu_{\max}}$ values from Yu et al. (2018). In terms of relative standard deviation, we calculate the sample standard deviation $\sigma_{\nu_{\max}}$ across all 13 predictions, each derived from different quarters' light curves for the same star, and then divide this by $\widehat{\nu_{\max}}$. Essentially, the relative deviation evaluates how closely the predicted $\nu_{\max}$ from single-quarter light curves approximates the 'ground truth' from the long time-baseline derivation, while the relative standard deviation assesses the internal consistency of predictions across different quarters.

In Fig. 7, we present the relative deviation on the left-hand side and the relative standard deviation on the right-hand side. While the relative deviations from both *Astroconformer* and the SYD pipeline are broadly comparable, the SYD pipeline exhibits a longer tail in its distribution. Quantitatively, *Astroconformer* achieves a relative median absolute error of 1.93%, in comparison to the 2.10% observed for the SYD pipeline. Moreover, *Astroconformer*'s predictions are more consistent. The right panel of Fig. 7 demonstrates that *Astroconformer*'s relative standard deviation is statistically much smaller than that of the SYD pipeline. Specifically, the relative median standard deviation for *Astroconformer* is 2.13%, in comparison to the 3.37% observed for the SYD pipeline.

It is noteworthy that although *Astroconformer* relies on training labels for $\nu_{\max}$ obtained from the SYD pipeline's analysis of end-of-mission *Kepler* light curves (~4 years), it provides predictions that are closer to 'ground truth' $\widehat{\nu_{\max}}$ and more internally consistent than that of the SYD pipeline when applied to single-quarter light curves (~90 days). *Astroconformer* leverages information beyond just the oscillations signal. As a result, even when working with single-quarter data, the model benefits from multiple constraints across information from the entire light curves, its corresponding power spectra but also the subtle correlation between the phases, enabling it to make precise and consistent predictions.

It is a plausible consideration that oscillation modes frequently exhibit stochastic variations over time, as referenced in Dupret et al. (2009). Such variability could explain the SYD pipeline's deviations from global time-averaged quantities. Conversely, *Astroconformer* is specifically trained to predict time-averaged $\nu_{\max}$ as outlined in Yu et al. (2018). This training approach accounts for its consistent performance across various quarters, indicating its capacity to harness additional information from the time series, extending beyond the mere oscillation modes.

5.3.2 *On 30-day Light Curves*

To more effectively illustrate the value of *Astroconformer* as a supplement to existing pipelines, we delve into a scenario where light curves only have a 33-day timespan. We compare *Astroconformer* with asteroseismic results in Hekker et al. (2011) based on the same 33-day Q1 light curves. Hekker et al. (2011) conducted an asteroseismic characterization for 16,511 red giants using their Q1 data, with oscillations detected in 10,956 (71%) of these stars.

We fine-tune the 20-layer Astroconformer, as discussed in the previous section, on randomly cropped 30-day segments from the Q2-16 dataset. We then perform $\nu_{\max}$ inference on Q1 light curves for 14,997 red giants, as defined in Hekker et al. (2011). Notably, all these samples possess more reliable $\nu_{\max}$ labels from Yu et al. (2018), derived from end-of-mission light curves that extend beyond these 30-day segments. Within this group, the pipeline from Hekker et al. (2011) successfully detected oscillations in approximately 70% of the stars (10,457 stars), marked as "Detection". The remaining

[3] We omit Q8, 12due to an insufficient number of valid data points with a flag of 0, as discussed in Section 4.

[4] The seemingly limited test set size is because only 717 red giants have all 13 quarters of light curves out of the initial 1000 test stars.

Short light curves still constrain stars without detected oscillations.

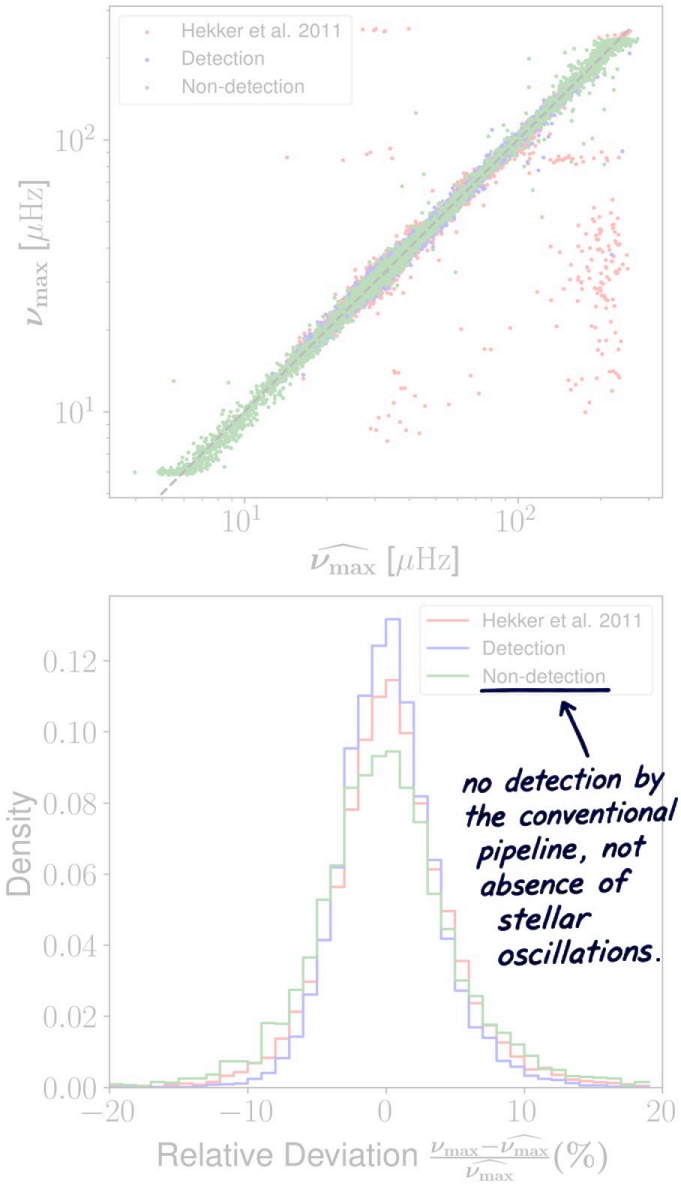


Figure 8. A comparative study of *Astroconformer* and Hekker et al. (2011) applied to Kepler Q1 light curves with an only 33-day time span. The top panel plot $\nu_{\rm max}$ estimations from both approaches against 'ground truth' $\widehat{\nu_{\rm max}}$ from Yu et al. (2018). Results of *Astroconformer* are categorized into two subgroups, 'Detection' and 'Non-detection', depending on their detectability in Hekker et al. (2011). The bottom panel illustrates the relative deviation from $\widehat{\nu_{\rm max}}$. Astroconformer's relative median absolute error for the 'Detection' group stands at 2.14%, smaller than the 2.52% of the conventional pipeline from Hekker et al. (2011). Even in the 'Non-detection' subgroup where Hekker et al. (2011) fails to identify oscillation modes, *Astroconformer* maintains a relative median absolute error of 2.91%.

30% (4,540 stars) did not exhibit detectable oscillations and are consequently labeled as "Non-detection".

We assess the $\nu_{\rm max}$ estimations made by Hekker et al. (2011) and *Astroconformer*, and juxtapose these with the 'ground truth' $\nu_{\rm max}$ from Yu et al. (2018), denoted as $\widehat{\nu_{\rm max}}$. In the top panel of Figure 8, we display the $\nu_{\rm max}$ estimations by Hekker et al. (2011) and *Astroconformer*, aligned against $\widehat{\nu_{\rm max}}$. The *Astroconformer* estimates are categorized into two subgroups, 'Detection' and 'Non-detection', based on the availability of $\nu_{\rm max}$ estimation in Hekker et al. (2011). Notably, the $\nu_{\rm max}$ estimations from Hekker et al. (2011) apply exclusively to the 'Detection' group. In the bottom panel, we present a histogram of the relative deviation from $\widehat{\nu_{\rm max}}$ as per Yu et al. (2018). Astroconformer's relative median absolute error for the 'Detection' group stands at 2.14%, compared with the 2.52% of the conventional pipeline from Hekker et al. (2011). Remarkably, even in the 'Non-detection' subgroup where Hekker et al. (2011) fails to identify oscillation modes, *Astroconformer* maintains a relative median absolute error of 2.91%.

Additionally, *Astroconformer* exhibits exceptional performance for high-luminosity red giants with $\nu_{\rm max} < 10\mu$Hz, an achievement that is noteworthy given the challenges the automated pipeline faces with these stars. The primary challenge lies in the short timespan of light curves, which reduces the frequency resolution in power spectra, thereby complicating the detection and characterization of oscillations. This effect is significant even when stellar oscillation are detectable, resulting in numerous outliers in $\nu_{\rm max}$ estimates by the automated pipeline, especially for faint red giants with $\nu_{\rm max} > 100\mu$Hz whose oscillation amplitudes are lower. In contrast, *Astroconformer* demonstrates remarkable resilience against such intricate cases, further evidencing its robustness.

6 DISCUSSION

In this study, we demonstrated that *Astroconformer* is able to provide precise and consistent results when analyzing single-quarter light curves. We also confirm its potential to provide a robust oscillation parameter estimation based on 30-day Kepler light curves. This performance advantage is likely attributable to *Astroconformer*'s ability to ingest the entire time-series data, thereby capturing not only stellar oscillation patterns but also granulation, which are typically discarded in traditional pipelines. Moreover, *Astroconformer*'s direct learning from light curves allows it to account for non-Gaussian phase information that is absent in power spectra. This capability is hypothesized to stem from *Astroconformer*'s unique facility for capturing both local and global correlations within the data.

Supporting this hypothesis, our ablation study indicates that the *Astroconformer* block employs an attention mechanism to calculate correlations between timestamp patches. To delve deeper into *Astroconformer*'s inner workings, we intend to visualize these learned attention patterns below.

6.1 Understanding Astroconformer's Learning Process

One of the distinct benefits of the self-attention mechanism is its capacity for global analysis, which contrasts with the localized receptive fields in CNNs. Furthermore, the self-attention framework provides an intuitive framework for interpreting what the neural network has learned (Caron et al. 2021). As outlined in Section 3.1, the key operation in self-attention involves taking the dot product of key and query vectors, which are linear transformations of the original input. The softmax of this product yields an attention score, quantifying the similarity between two patches in the input sequence. This attention matrix offers a valuable perspective into the specific aspects of the light curve that *Astroconformer* prioritizes.

In particular, the input of *Astroconformer* is 200 patches, each contains 20 timestamps (10 hours), of the light curve. In self-attention mechanism, each patch is converted to a query patch and a key patch. The key patch at position j is paired with the query patch at position i, yielding their correlation, denoted as $A_{ij} = softmax(\boldsymbol{q}^i \cdot \boldsymbol{k}^j)$.

Attention patterns reveal useful timescales, but they are not causal proof.

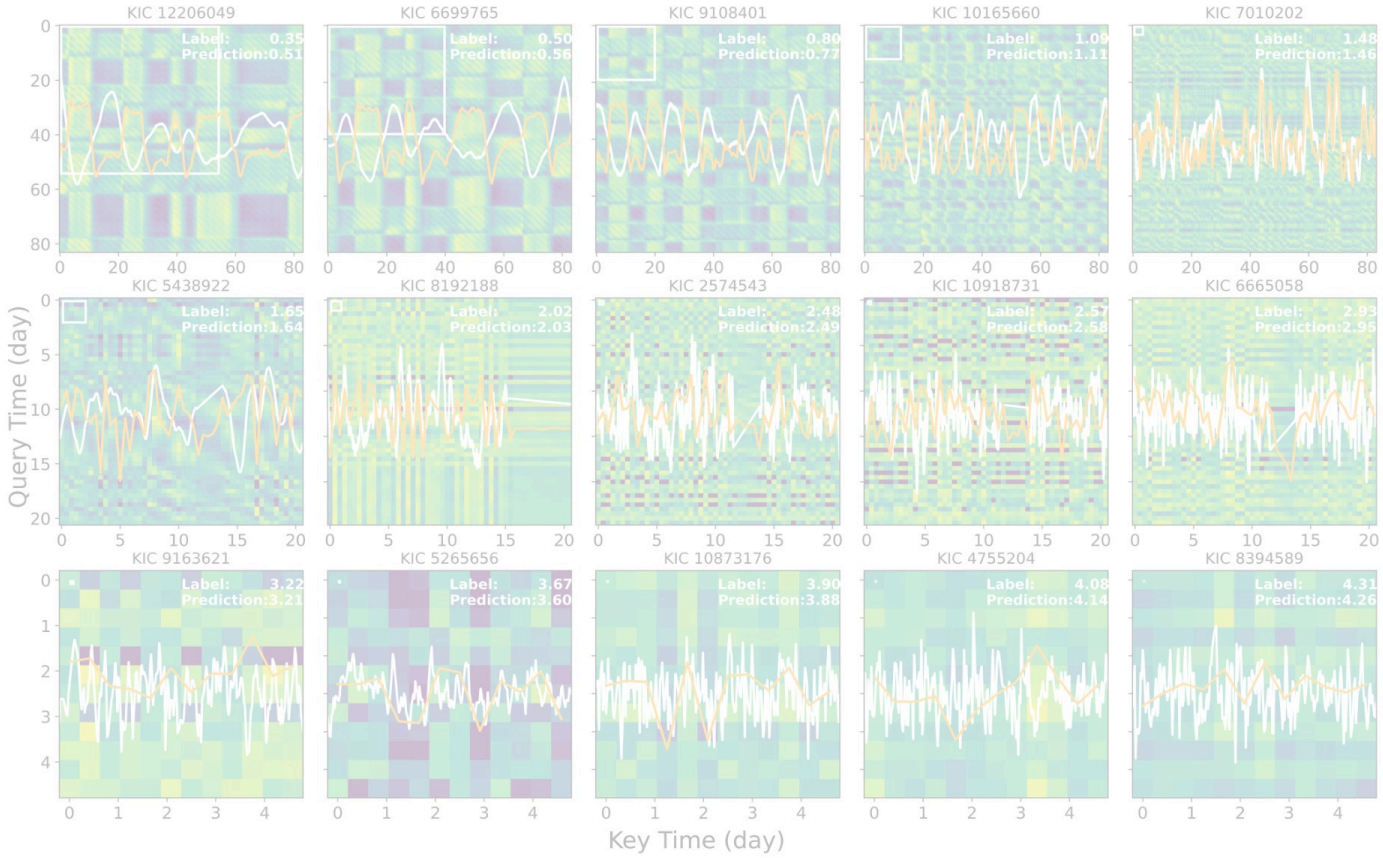


Figure 9. Attention maps for a selected set of *Kepler* stars that span a broad range of $\log g$ values and encompass various evolutionary phases, from turnoff stars to high-luminosity red giants. The brighter color indicate a larger attention weight. The star's normalized light curve is included for contextual reference. The row-averaged attention map are in orange to investigate the general pattern. Shown on top of each panel is the KIC number. A white square in the top-left corner of each panel indicates the star's characteristic oscillation period, as derived from Yu et al. (2018). For better details, panels in the second row and third row are zoomed in by 4 and 16, respectively. For stars with lower $\log g$ values ($\log g < 2$), the periodic patterns in the attention maps closely match the stars' light curve waveforms, highlighting *Astroconformer*'s proficiency in leveraging oscillations and granulation signals to accurately estimate $\log g$. Interestingly, even for stars with higher $\log g$ values—where one would expect the oscillation period to be unresolved in the attention map—long-range patterns remain discernible. These patterns have periods significantly longer than those of the stellar oscillations, further confirming that *Astroconformer* is capable of capturing features associated with larger-period granulation, particularly in stars with higher $\log g$ values.

Diagonal stripes can come from positional encoding.

Subsequently, these correlations serve as weights of how the patch i combines information from all 200 patches.

With this in mind, in Figure 9, we display attention maps for a diverse set of *Kepler* stars, spanning a range of $\log g$ values and stellar types, from turnoff stars to high-luminosity red giants. For the sake of illustration, we examine the attention map generated by the first head of the MHSA module within *Astroconformer*'s third layer. The background image of each map shows the attention values for individual patches, indicating how correlated signals are at different moments in time. Overlaying each map in the top-left corner is a white square, which signifies the star's characteristic oscillation period. Additionally, the star's normalized light curve is included for contextual reference. For the readability of general attention trend, a row-averaged attention map is plotted against each light curve. Lastly, we note that diagonal stripes in the attention maps are artifacts of the positional encoding process, since it scales the correlation of position i, j by a complex phase proportional to their relative distance $i - j$.

First, as seen in the figure, *Astroconformer*'s attention mechanism operates across varied time intervals to enable a more comprehensive analysis of light curves directly in the time domain. This is in contrast with CNNs, which are limited by their local receptive fields, making them less adaptable for analyzing extended time series data like light curves. Consequently, CNN approaches in asteroseismology are often restricted to the scrutiny of power spectra (e.g., Hon et al. 2018). For stars with $\log g < 2$ that the general attention patterns generated by *Astroconformer* are consistent with the light curve stellar oscillations. But on top of the characteristic oscillation periods, the patterns also include low frequency components, potentially from granulation.

Equally illuminating are the attention maps for stars with $\log g > 3$. These maps highlight frequency below those typically associated with oscillations. Given that these maps examine 10-hour patches—exceeding the oscillation periods of stars with $\log g > 3$—one would expect any oscillation patterns to remain unresolved. If stellar oscillations were the only feature captured, the attention maps would offer limited insights. Yet, as Figure 9 illustrates, averaged attention maps show consistency with trends of light curves. These extended patterns, substantially longer than the oscillation periods, support the idea that *Astroconformer* is tapping into additional data layers, possibly related to granulation and the non-gaussian nature of oscillations, to accurately infer the $\log g$ of stars. We also

Uniform sampling and careful preprocessing remain important limitations.

further validate the correlation between attention maps and granulation timescales for red giants in Appendix B.

Finally, another significant observation relates to the interpolated segments within the light curve. In these areas, the attention map typically assigns lower weights, thereby minimizing the influence of interpolated information. For instance, in the case of KIC 6665058 ($\log g \approx 3$), an interpolated segment around the 15th day appears notably dimmer in the attention map compared to surrounding areas. This suggests that *Astroconformer* demonstrates resilience against contaminated signals, effectively distinguishing spurious data from genuine light curve features. This capability may also account for *Astroconformer*'s robust performance.

6.2 Limitations and Future Work

While we have demonstrated that *Astroconformer* holds considerable potential, there remain limitations that necessitate additional research. A fundamental constraint is our reliance on uniformly sampled light curves. While this is a reasonable assumption for data from missions like *Kepler* or TESS, it could present challenges when dealing with non-uniform datasets, such as those generated by the Rubin Observatory. Despite the model's ability to leverage long-range temporal information, its current implementation may not be directly applicable to non-uniform data sources.

This assumption of uniform sampling is critical for two primary reasons. First, directly inputting individual time stamps without patching into the *Astroconformer* block would significantly increase computational costs. For example, this approach would necessitate performing 4000 x 4000 attention calculations instead of a more manageable 200 x 200 grid. Second, the assumption aids the integration of convolutional modules within the *Astroconformer* architecture. As discussed in Section 5.2, the limited availability of training samples makes a stronger inductive bias via convolutional modules advantageous. While switching to a pure Transformer architecture could bypass the need for uniform sampling, it may also result in compromised model performance. Future work should investigate alternative methods for embedding or patching to accommodate non-uniformly sampled light curves.

Additionally, the ability of *Astroconformer* to process all available information including instrumental effects simultaneously presents both advantages and disadvantages. Although the model can effectively filter out brief instances of spurious data, as evidenced in our empirical study in Section 6.1, it is susceptible to retaining misleading trends that persist throughout the entire light curve. Our empirical findings further substantiate this observation. For instance, we examine outliers in Section 5.1 and identify anomalies in their light curves, which are attributed to factors such as contamination, rotation modulation, or incomplete removal of extreme flux variations. Moreover, when applied to `Kepseismic` light curves[5], where only signals with periods longer than 80 days are filtered and significant residual trends remain, *Astroconformer*'s performance deteriorates, yielding an error of approximately 0.14 dex. This underscores the importance of meticulous data preprocessing, as outlined in Section 4. We consider this a limitation that could potentially be mitigated through deep learning techniques, such as constrastive learning (Chen et al. 2020), that allow for better data representation and spurious trend removal.

In this research, we have utilized training labels derived from the end-of-mission Kepler light curves in Yu et al. (2018). These labels are of relatively low uncertainty, especially when compared to the inherent systematics present in inference models, including our own *Astroconformer*. Therefore, they are not the dominant factor in this study. However, it is important to acknowledge that these labels are not free from noise. In future studies, exploring a more sophisticated probabilistic deep learning approach that accounts for label noise could be highly beneficial. Such an approach would aim to refine the inference process by effectively 'deconvolving' the label noise, thereby enhancing the accuracy and reliability of the results.

[5] `Kepseismic` data are available at MAST via http://dx.doi.org/10.17909/t9-mrpw-gc07

7 CONCLUSIONS

In this study, we introduce *Astroconformer*, a novel method based on the Transformer neural network architecture, renowned for its capacity to adaptively extract long-range contextual information from sequential data. This is particularly suited for time-series data such as stellar light curves, which often exhibit long-range dependencies. We specifically employ *Astroconformer* to infer the surface gravities of stars from their light curves. The primary findings of our research can be summarized as follows:

1. *Astroconformer* excels in the precise inference of $\log g$ values. When benchmarked against asteroseismic $\log g$ measurements, spanning a $\log g$ range of 0.2–4.4 dex, *Astroconformer* outperforms the k-nearest neighbor-based method known as THE SWAN. *Astroconformer* achieves a minimal RMSE of 0.017 dex at $\log g \sim 3$ and presents an increase in RMSE, up to 0.1 dex, towards the boundaries of the investigated $\log g$ range.
2. The efficacy of a deep learning model in extracting $\log g$ values from light curves is highly dependent on the model's receptive field. When self-attention modules are replaced with convolutional modules, there is a marked decline in performance, underscoring the superiority of Transformer models over convolutional neural networks for handling stellar light curves.
3. When predicting $\nu_{\rm max}$ from single-quarter light curves from *Kepler*, *Astroconformer* achieves a relative median absolute error of 1.93% and a relative median standard deviation of 2.10%, compared with the SYD pipeline's relative median absolute error of 2.10% and relative median standard deviation of 3.37%.
4. For shorter 33-day light curves from Kepler, our analysis reveals that *Astroconformer* attains a relative median absolute error of less than 3%. It consistently delivers reliable $\nu_{\rm max}$ estimations in nearly all instances. Conversely, the conventional pipeline fails to detect solar-like oscillations in approximately 30% of these samples, highlighting *Astroconformer*'s robust capabilities to infer $\nu_{\rm max}$, especially for short time series.
5. *Astroconformer* also provides valuable interpretability via its attention maps. The attention maps show that the model has learned to identify not only stellar oscillations but also longer-period granulation patterns, which are crucial for making accurate $\log g$ estimations, especially for less evolved stars.

While Transformer-based models have achieved considerable advancements in areas such as computer vision and natural language processing, their application in astronomy is still in its early stages. This study zeros in on applying Transformer models to stellar light curves, demonstrating their potential to better extract long-range information—information that has often been overlooked by CNNs. This capability could be particularly crucial in upcoming surveys like those from the Rubin Observatory and the Roman Space Telescope, as well as in extracting valuable stellar properties from data with more complex noise characteristics like TESS.

Looking at the broader picture, the utility of Transformer models undoubtedly extends well beyond these specific applications. Beyond the realm of asteroseismology, Transformer architectures hold immense potential for a wide array of other astronomical investigations. Given their proficiency in processing long-range information—a ubiquitous trait in astronomical data—these models open up avenues for more synergistic progress between the domains of machine learning and astronomy.

ACKNOWLEDGEMENTS

We extend our gratitude to Maryum Sayeed and Saskia Hekker for generously sharing the data used in their work and for providing meticulous guidance on their utilization. YST acknowledges financial support received from the Australian Research Council via the DECRA Fellowship, grant number DE220101520. We also recognize the invaluable contribution of public data from the *Kepler* mission, supplied by NASA's Ames Research Center. The *Kepler* mission is funded by NASA's Science Mission Directorate. Our special thanks go to the *Kepler* team for making their data openly accessible, thereby facilitating this research.

DATA AVAILABILITY

The data used in this study are publicly available and can be accessed through the Kepler Public Archive hosted by the Mikulski Archive for Space Telescopes (MAST). Our PyTorch implementation of *Astroconformer* is publicly available at `https://github.com/panjiashu/Astroconformer`. Any additional data or materials related to this paper may be requested from the corresponding author.

REFERENCES

Aerts C., 2021, Reviews of Modern Physics, 93
Aerts C., Christensen-Dalsgaard J., Kurtz D., 2010, Asteroseismology. Springer, Netherlands, doi:10.1007/978-1-4020-5803-5
Auvergne M., et al., 2009, Astronomy & Astrophysics, 506, 411
Bedding T. R., et al., 2011, Nature, 471, 608
Bellm E. C., et al., 2018, Publications of the Astronomical Society of the Pacific, 131, 018002
Benomar O., Appourchaux T., Baudin F., 2009, http://dx.doi.org/10.1051/0004-6361/200911657, 506
Berger T. A., Huber D., van Saders J. L., Gaidos E., Tayar J., Kraus A. L., 2020, The Gaia-Kepler Stellar Properties Catalog. I. Homogeneous Fundamental Properties for 186,301 Kepler Stars (arXiv:2001.07737)
Blancato K., Ness M. K., Huber D., Lu Y., Angus R., 2022, ApJ, 933, 241
Brown T. M., Gilliland R. L., Noyes R. W., Ramsey L. W., 1991, ApJ, 368, 599
Brown T. B., et al., 2020, Language Models are Few-Shot Learners (arXiv:2005.14165)
Bugnet L., Garcí a R. A., Davies G. R., Mathur S., Corsaro E., Hall O. J., Rendle B. M., 2018, A&A, 620, A38
Caron M., Touvron H., Misra I., J é gou H., Mairal J., Bojanowski P., Joulin A., 2021, Emerging Properties in Self-Supervised Vision Transformers (arXiv:2104.14294)
Chaplin W. J., Miglio A., 2013, ARA&A, 51, 353
Chatterjee S., Zielinski P., 2022, On the Generalization Mystery in Deep Learning (arXiv:2203.10036)
Chen T., Kornblith S., Norouzi M., Hinton G., 2020, A Simple Framework for Contrastive Learning of Visual Representations (arXiv:2002.05709)
Devlin J., Chang M.-W., Lee K., Toutanova K., 2019, BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding (arXiv:1810.04805)
Dosovitskiy A., et al., 2021, An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale (arXiv:2010.11929)
Dupret M. A., et al., 2009, A&A, 506, 57
Erickson N., Mueller J., Shirkov A., Zhang H., Larroy P., Li M., Smola A., 2020, arXiv preprint arXiv:2003.06505
Fuller J., Cantiello M., Stello D., Garcia R. A., Bildsten L., 2015, Science, 350, 423
García R. A., Ballot J., 2019, Living Reviews in Solar Physics, 16
Garcí a R. A., et al., 2011, Monthly Notices of the Royal Astronomical Society: Letters, 414, L6
Gehan C., Mosser B., Michel E., Samadi R., Kallinger T., 2018, A&A, 616, A24
Gruberbauer M., Kallinger T., Weiss W. W., Guenther D. B., 2009, A&A, 506, 1043
Gulati A., et al., 2020, Conformer: Convolution-augmented Transformer for Speech Recognition (arXiv:2005.08100)
Hall O. J., et al., 2021, Nature Astronomy, 5, 707
He K., Zhang X., Ren S., Sun J., 2015, Deep Residual Learning for Image Recognition (arXiv:1512.03385)
Hekker S., Christensen-Dalsgaard J., 2017, The Astronomy and Astrophysics Review, 25
Hekker S., et al., 2011, Monthly Notices of the Royal Astronomical Society, 414, 2594
Hekker S., et al., 2012, A&A, 544, A90
Hinners T. A., Tat K., Thorp R., 2018, AJ, 156, 7
Hon M., Stello D., Yu J., 2018, Monthly Notices of the Royal Astronomical Society, 476, 32333244
Hon M., Stello D., Garcí a R. A., Mathur S., Sharma S., Colman I. L., Bugnet L., 2019, Monthly Notices of the Royal Astronomical Society, 485, 5616
Howell S. B., et al., 2014, Publications of the Astronomical Society of the Pacific, 126, 398
Huber D., Stello D., Bedding T. R., Chaplin W. J., Arentoft T., Quirion P.-O., Kjeldsen H., 2009, Automated extraction of oscillation parameters for Kepler observations of solar-type stars (arXiv:0910.2764)
Ioffe S., Szegedy C., 2015, Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift (arXiv:1502.03167)
Ismail Fawaz H., Forestier G., Weber J., Idoumghar L., Muller P.-A., 2019, Data Mining and Knowledge Discovery, 33, 917963
Ivezić Ž., et al., 2019, ApJ, 873, 111
Kallinger T., et al., 2014, A&A, 570, A41
Kallinger T., Hekker S., Garcia R. A., Huber D., Matthews J. M., 2016, Science Advances, 2, 1500654
Kawaler S. D., Hostler S. R., 2005, The Astrophysical Journal, 621, 432
Kjeldsen H., Bedding T. R., 1995, A&A, 293, 87
Koch D. G., et al., 2010, ApJ, 713, L79
Li G., Deheuvels S., Ballot J., Lignières F., 2022, Nature, 610, 43
Loshchilov I., Hutter F., 2019, Decoupled Weight Decay Regularization (arXiv:1711.05101)
Marshall P., et al., 2017, LSST Science Collaborations Observing Strategy White Paper: "Science-driven Optimization of the LSST Observing Strategy", doi:10.5281/ZENODO.842713, https://zenodo.org/record/842713
Mathur S., et al., 2011, ApJ, 741, 119
Mathur S., et al., 2017, The Astrophysical Journal Supplement Series, 229, 30
Ness M. K., Aguirre V. S., Lund M. N., Cantiello M., Foreman-Mackey D., Hogg D. W., Angus R., 2018, The Astrophysical Journal, 866, 15
Nie Y., Nguyen N. H., Sinthong P., Kalagnanam J., 2023, A Time Series is Worth 64 Words: Long-term Forecasting with Transformers (arXiv:2211.14730)
Ramachandran P., Zoph B., Le Q. V., 2017, Searching for Activation Functions (arXiv:1710.05941)
Ricker G. R., et al., 2014, in Oschmann Jacobus M. J., Clampin M., Fazio G. G., MacEwen H. A., eds, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 9143, Space Telescopes and

Instrumentation 2014: Optical, Infrared, and Millimeter Wave. p. 914320 (arXiv:1406.0151), doi:10.1117/12.2063489
Sayeed M., Huber D., Wheeler A., Ness M. K., 2021, The Astronomical Journal, 161, 170
Smith L. N., 2017, Cyclical Learning Rates for Training Neural Networks (arXiv:1506.01186)
Stello D., Cantiello M., Fuller J., Huber D., García R. A., Bedding T. R., Bildsten L., Aguirre V. S., 2016, Nature, 529, 364
Su J., Lu Y., Pan S., Wen B., Liu Y., 2021, RoFormer: Enhanced Transformer with Rotary Position Embedding (arXiv:2104.09864)
Szegedy C., et al., 2014, Going Deeper with Convolutions (arXiv:1409.4842)
Toutain T., Appourchaux T., 1994, A&A, 289, 649
Ulrich R. K., 1986, ApJ, 306, L37
Vaswani A., Shazeer N., Parmar N., Uszkoreit J., Jones L., Gomez A. N., Kaiser L., Polosukhin I., 2017, Attention Is All You Need (arXiv:1706.03762)
Wang H., Ma S., Dong L., Huang S., Zhang D., Wei F., 2022, DeepNet: Scaling Transformers to 1,000 Layers (arXiv:2203.00555)
Yu J., Huber D., Bedding T. R., Stello D., Hon M., Murphy S. J., Khanna S., 2018, The Astrophysical Journal Supplement Series, 236, 42
Yu J., Bedding T. R., Stello D., Huber D., Compton D. L., Gizon L., Hekker S., 2020, Monthly Notices of the Royal Astronomical Society, 493, 1388

APPENDIX A: TRAINING

In Section 5.1 and Section 5.2, the 5-layer *Astroconformer* and other models are trained with the AdamW optimizer (Loshchilov & Hutter 2019) and used a cyclic learning rate scheduler (Smith 2017). This scheduler features three linear warm-up and cool-down stages (warm-up ratio=0.1), with each subsequent stage having a peak learning rate that is half of its predecessor. The lower bound of learning rate is set to be 0.1% of the first peak value. We set the first peak learning rate to 0.003 such that the model is converged while fully trained given the other setting, except for 'Without MHSA', which adopts a 0.001 learning rate due to training instability. Regarding the batch size selection, we conducted tests with batch sizes of 128, 256, and 512. Our findings indicated that batch size does not significantly impact the final result, provided it is greater than 256. Consequently, we opted for a batch size of 256 for our experiments. Training was carried out over 20,000 steps, with validation loss results indicating that this step count is adequate; we observed a plateau in the validation loss beyond this point. The typical training duration for *Astroconformer* on an Nvidia A100 40GB GPU is approximately 3.5 hours.

In Section 5.3.1, the training of a 20-layer *Astroconformer* involves weight initialization scheme proposed by Wang et al. (2022), which sets the shortcut weight to 2 and weight initialization gain to 0.5. We then implement gradient clipping by 1.0. The model was trained 100,000 training steps, with a batch size of 48, a warm-up stage lasting 10,000 steps, a peak learning rate of 0.0003, and a lower bound learning rate, 10% of the peak value. The training takes about 50 hours. In Section 5.3.2, the finetuning of the 20-layer *Astroconformer* contains a warm-up and cool-down stage lasting 10,000 steps, with a peak learning rate of 0.0001.

APPENDIX B: RELATION BETWEEN ATTENTION MAPS AND GRANULATION

We hypothesize that *Astroconformer*'s robust performance compared to conventional pipelines is due to its ability to extract information from granulation within the light curves, in addition to focusing on oscillation modes. This is supported by observations in Figure 9, where the representation of high $\log g$ values seems to incorporate more than just oscillation modes, suggesting that additional data, including granulation, contributes to the inference of $\log g$ for dwarf stars. However, it's important to acknowledge that while attention maps offer an intuitive insight into the workings of the attention mechanism, they do not directly link to granulation information. To clarify this connection, we aim to demonstrate a relationship between the patterns observed in Figure 9 and granulation timescales measured by the SYD pipeline. Specifically, the SYD pipeline fits the power spectrum background with the form

$$P(\nu) = P_n + \sum_{i=1}^{k} \frac{4\sigma_i^2 \tau_i}{1 + (2\pi\nu\tau_i)^2 + (2\pi\nu\tau_i)^4}, \tag{B1}$$

where P_n is the white noise component, k is the number of power laws used, and σ and τ are the root-mean-square intensity and timescale of granulation, respectively (Huber et al. 2009). The timescale of the first power law component, τ_1, is used to characterize the granulation timescale.

To explore the general response trends in attention maps, we randomly selected 3,487 single-quarter light curves from our dataset, as discussed in Section 5.3. For each of these light curves, we calculated the attention map using the first head in the third self-attention layer and then computed the power spectrum for each row of the attention map. An averaged power spectrum over all rows is utilized as our input variable. Meanwhile, we run the SYD pipeline on these light curves, and the derived τ_1 serves as our target variable. If the attention maps are indeed capturing granulation information, it should be feasible to deduce granulation timescales from the power spectra of these attention maps. However, this non-linear regression task presents a wide array of methodological options.

To streamline this analysis, we employ Autogluon (Erickson et al. 2020), an auto-ML package that systematically applies various machine learning models and combines their results. This approach enables us to effectively learn the correlation between the averaged power spectra and the granulation timescales. The 3,487 samples are split into subgroups of 2487, 500, 500 for training, validation, and test set, respectively.

We adopt the default `Tabularpredictor` in Autugluon to fit our training data, which automatically search through the methods including K-Nearest Neighbor, Light Gradient Boosting Machine (LightGBM), Random Forest, Categorical Boosting , Extra Trees, Multi Layer Perceptron, and eXtreme Gradient Boost (XGBoost), as well as optimising their parameters. Each model is evaluated by its RMSE on the validation set. After training individual models, an ensemble model is derived by weighting predictions from individual models by their RMSE on the validation set. In our experiment, the LightGBM model with an RMSE of 0.0098 day is the most precise individual model, while the ensemble model achieves a better RMSE of 0.0094 day.

We use the ensemble model to predict τ_1 based on averaged power spectra of attention maps in test set. The outcomes of our analysis are presented in Fig B1, with an RMSE of 0.0104 day. While this does not conclusively establish causality, this initial investigation into the power spectrum of the attention map lends further credence to the idea that the attention mechanism, as applied to light curves, indeed captures information reflective of the granulation patterns in stars. This evidence supports our assertion that *Astroconformer* excels at extracting insights beyond just oscillation modes.

This paper has been typeset from a TEX/LATEX file prepared by the author.

Attention features predict granulation timescales with scatter, not certainty.

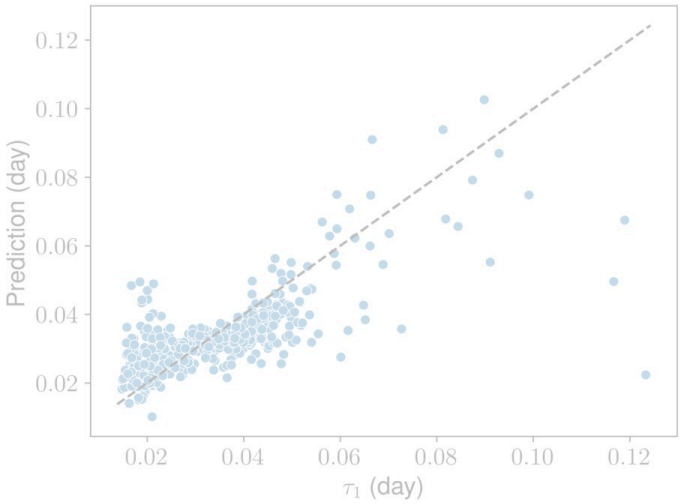
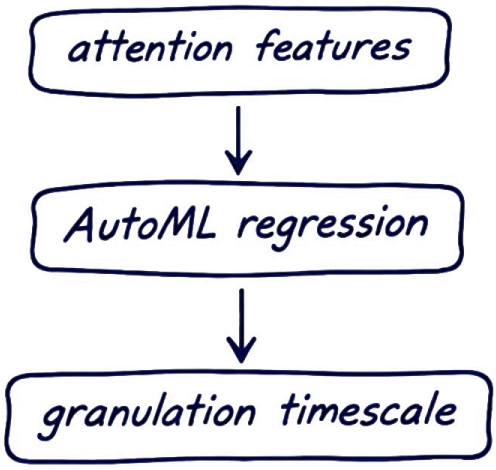


Figure B1. Averaged power spectra of attention maps encode granulation information. The horizontal axis is characteristic timescale of the first power law component of granulation, τ_1, derived by the SYD pipeline, and the vertical axis is τ_1 predicted by AutoML using averaged power spectra of attention maps.



association, not causation