



Exoplanet Classification Through Vision Transformers with Temporal Image Analysis

Anupma Choudhary¹ , Sohith Bandari², B. S. Kushvah¹ , and C. Swastik^{3,4,5} ¹ Department of Mathematics and Computing, Indian Institute of Technology (Indian School of Mines), Dhanbad 826004, India² Department of Electronics Engineering, Indian Institute of Technology (Indian School of Mines), Dhanbad 826004, India³ Dipartimento di Fisica, Università degli Studi di Milano, Via Celoria 16, 20133 Milano, Italy; swastik.chowbay@iit.res.in⁴ Indian Institute of Astrophysics, Koramangala 2nd Block, Bangalore 560034, India⁵ Pondicherry University, R.V. Nagar, Kalapet, 605014, Puducherry, India

Received 2024 August 27; revised 2025 June 9; accepted 2025 June 17; published 2025 July 28

Abstract

The classification of exoplanets has been a longstanding challenge in astronomy, requiring significant computational and observational resources. Traditional methods demand substantial effort, time, and cost, highlighting the need for advanced machine learning techniques to enhance classification efficiency. In this study, we propose a methodology that transforms raw light curve data from NASA’s Kepler mission into Gramian angular fields (GAFs) and recurrence plots (RPs) using the Gramian angular difference field and RP techniques. These transformed images serve as inputs to the vision transformer (ViT) model, leveraging its ability to capture intricate temporal dependencies. We assess the performance of the model through recall, precision, and F1 score metrics, using a five-fold cross-validation approach to obtain a robust estimate of the model’s performance and reduce evaluation bias. Our comparative analysis reveals that RPs outperform GAFs, with the ViT model achieving an 89.46% recall and an 85.09% precision rate, demonstrating its significant ability to accurately identify exoplanetary transits. Despite using undersampling techniques to address class imbalance, data set size reduction remains a limitation. This study underscores the importance of further research into optimizing model architectures to enhance automation, performance, and generalization of the model.

Unified Astronomy Thesaurus concepts: [Convolutional neural networks \(1938\)](#); [Transit photometry \(1709\)](#); [Exoplanet astronomy \(486\)](#); [Astronomy data visualization \(1968\)](#); [Light curve classification \(1954\)](#); [Light curves \(918\)](#); [Keplerian telescopes \(2356\)](#)

1. Introduction

The groundbreaking discovery of the first exoplanet orbiting a star similar to our Sun by M. Mayor et al. (1995) sparked a new era of research in the field of astronomy. Since then, thousands of exoplanets have been discovered, ranging in size and mass from those similar to Earth to massive, Jupiter-sized planets. With the detection of over 5800 exoplanets (NASA Exoplanet Science Institute 2025), scientists are now able to study the different populations of exoplanets and their host stars in greater detail (D. A. Fischer & J. Valenti 2005; S. Udry & N. C. Santos 2007; J. A. Johnson et al. 2017; C. Swastik et al. 2022; A. Unni et al. 2022; M. L. Hill et al. 2023; T. Pyne et al. 2025). By investigating the correlations between the properties of exoplanet host stars and the planets they host, researchers can gain a deeper understanding of how these planets formed and evolved. It is also worth noting that many exoplanet systems do not resemble our solar system; for instance, the prevalence of “super-Earths” and “mini-Nep-tunes” which are not found in our solar system, and “hot Jupiters” which are extremely close to their host stars. This highlights that our solar system is somewhat atypical and that exoplanet formation and migration may have occurred differently when compared to our solar system.

Over time, various techniques have been developed for detecting exoplanets. The most direct way to detect exoplanets is direct imaging. The direct imaging technique uses high-contrast

imaging and advanced technology such as coronagraphs and adaptive optics, and is typically used to find exoplanets in wider orbits (\sim tens of astronomical unit), specifically young, self-luminous systems that are easier to detect (G. Chauvin et al. 2005; C. Swastik et al. 2021; B. B. Ren et al. 2023; Z. Wahhaj et al. 2024; T. Pyne et al. 2025). However, only a handful of planets (~ 70) have been detected by direct imaging. The majority of exoplanets detected so far are by transit, followed by radial velocity (RV) technique. The RV detection technique, which also detected the first exoplanet, 51 peg, measures the tiny wobbles in a star’s motion caused by the gravitational pull of an orbiting exoplanet, revealing the planet’s presence and mass. The transit method, on the other hand, detects exoplanets by observing the slight dimming of a star’s light when a planet passes in front of it, revealing the planet’s size and orbit (W. J. Borucki et al. 2010). These techniques have been crucial in advancing our understanding of exoplanetary systems and are continually being refined to facilitate future discoveries (J. T. Wright & B. S. Gaudi 2012; M. Perryman 2018; C. Swastik et al. 2023). However, it is important to note that different methods have different biases and limitations. For example, the RV and transit methods are currently the most successful methods for detecting exoplanets, but they are most effective for finding planets in close to moderate orbits (0.1–10 au). On the other hand, the direct imaging technique is sensitive to find exoplanets in wider orbits (\sim tens of astronomical unit), specifically young, self-luminous systems that are easier to detect (C. Swastik et al. 2024).

In a significant endeavor to expand our knowledge of exoplanets, several missions have been launched by various space agencies. The French Space Agency initiated the CoRoT mission in 2006, pioneering space-based transit observations



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

and discovering numerous exoplanets (A. Baglin et al. 2006). Building on this success, NASA launched the Kepler mission in 2009, which played a vital role in expanding our knowledge of exoplanetary systems (W. J. Borucki et al. 2010). Kepler was established to monitor over 170,000 stars in the Milky Way, especially in the Cygnus-Lyra area, with the main objective of finding Earth-sized exoplanets in the habitable zone of their host stars using the transit method (W. J. Borucki 2016). By analyzing the slight dimming of starlight caused by planetary transits, Kepler substantially advanced the discovery and classification of exoplanets. This mission revealed a rich diversity in exoplanet size, composition, and orbital dynamics (W. J. Borucki et al. 2010; N. M. Batalha 2014; W. J. Borucki 2016). Running until 2018, Kepler significantly enhanced our understanding of exoplanetary systems, confirming over 2600 exoplanets and demonstrating the widespread presence of planets in our Galaxy.

Despite the significant contributions of various space missions, the complexity and volume of Kepler data necessitated the development of advanced techniques for exoplanet classification. Early machine learning methods, such as random forests, were implemented to automate the classification of Kepler objects of interest, significantly reducing human effort and minimizing errors, thus leading to more efficient and accurate classification (S. D. McCauliff et al. 2015). Furthermore, studies such as S. Basak et al. (2021) have shown that machine learning techniques can classify new exoplanets, but capturing complex patterns inside the large data set is still a challenging task. To address these limitations, Y. LeCun et al. (1998) introduced convolutional neural networks (CNNs), a subclass of deep neural networks specifically designed for image-based classification. CNNs consist of multiple layers of convolutional filters to autonomously learn and extract spatial hierarchies of features from input images. Each successive layer of filters is designed to extract complex features, from simple edges and textures to intricate patterns (A. Krizhevsky et al. 2012). This hierarchical approach allows CNNs to effectively analyze and classify visual data with remarkable accuracy and efficiency (C. Szegedy et al. 2015).

Building on the limitations of CNNs in capturing long-term dependencies within data sets and the successful emergence of transformers in natural language processing (NLP) inspired researchers, including A. Dosovitskiy et al. (2020), to experiment with a standard transformer by applying it directly to images, with the fewest possible modifications. As a result, the vision transformer (ViT) model was developed, which offers a new approach to image-based classification as detailed by A. Dosovitskiy et al. (2020). The fundamental concept of the model involves a sequence of visual patches as tokens for classifying the full image. It has achieved a benchmark performance in multiple image-based recognition tasks. In addition, transformer-based architectures have been successfully applied to numerous vision problems, including semantic segmentation (S. Zheng et al. 2021), video understanding (A. Arnab et al. 2021), image processing using ViT (A. Dosovitskiy et al. 2020), and object detection (X. Dai et al. 2021). Inspired by the rapid evolution of transformers, we have also implemented this in our methodology to examine its potential in exoplanet classification.

In the field of image-based classification, J.-P. Eckmann et al. (1995) introduced recurrence plots (RPs) as a technique to visualize the periodic nature inside the time-series data. RPs

convert time-series data into a 2D image where patterns and structures in the data can be more easily identified, facilitating the detection of hidden periodicities and the analysis of dynamic behavior over time (J.-P. Eckmann et al. 1995). In contrast, Z. Wang & T. Oates (2015) introduced the concept of Gramian angular fields (GAFs), which provide a method for transforming time-series data into images by encoding angular information. GAFs represent time-series data in a polar coordinate system, capturing the temporal correlations between different points. This transformation emphasizes the angular relationships between data points, which can be particularly useful for capturing cyclical patterns and trends (Z. Wang & T. Oates 2015).

Considering the potential of the above-mentioned transformation of 1D time series to 2D image representations, researchers have used them with various machine learning models in diverse fields. For example, P. G. Freitas et al. (2023a, 2023b) highlighted the potential of RPs and GAFs by projecting 1D photoplethysmogram signals into 2D images as an input to ViT model for enhancing the overall performance and the results demonstrated a competitive prediction accuracy as compared to the current state of the art. Subsequently, these advancements have enhanced the performance in tasks such as classification accuracy, anomaly detection, and forecasting (P. G. Freitas et al. 2023a, 2023b). These techniques have also been applied in the field of astronomy, where they have been used to enhance the efficiency of solar panels (M. Mulenga et al. 2023).

In this way, recent research studies have successfully integrated image representations with advanced modeling algorithms for time-series data, showcasing their effectiveness in tasks related to image-based classification. Therefore, this study intends to build upon this foundation by integrating these representations as an input in the ViT model to evaluate its potential in exoplanet classification. Furthermore, we performed a comparative analysis between RPs and GAFs to assess their performance when integrated with the ViT model, providing a framework to explore further in the field of astronomy.

This paper is organized as follows. Section 1 offers an in-depth overview of the background of exoplanet research, detailing various missions and the machine learning techniques for exoplanet classification. In Section 2, we present our methodology, which encompasses the processes of data collection and preprocessing. Additionally, we provide details of the conversion of data into the preferred formats: RPs and GAFs. This section provides a comprehensive overview of the techniques and strategies applied to prepare the data for subsequent analysis using ViT. Section 3 gives the subtle architectural details of ViT and hyperparameters used in the model. In Sections 4 and 5, we will provide detailed information on the data used in our study and present a comprehensive analysis of our comparative study of RPs and GAFs as inputs to our model. These sections elaborate on the techniques used for categorizing exoplanets, along with an in-depth discussion of the results obtained and their implications for the field of exoplanetary science. Section 6 will culminate with a detailed discourse on the conclusions drawn from our comparative study, emphasizing that RPs have performed better on this data set.




Figure 1. This diagram illustrates the transit crossing events observed in the light curve data for (KIC) candidate with ID 6922244. The plotted points represent the variations in flux as the exoplanet transits its host star, causing periodic dips in brightness.

2. Data Preprocessing

2.1. Light Curves

NASA’s Kepler mission, launched in 2009, was specifically designed to search for exoplanets by monitoring the brightness of over 156,000 stars in a fixed field of view (W. J. Borucki et al. 2010). This mission deepens our understanding of exoplanetary systems by detecting subtle changes in stellar brightness caused by planetary transits over vast distances, typically known as light curves associated with a particular star, which could further be used to identify and classify extrasolar planets. Transiting dips are the reduction in flux values that provide valuable information about the transiting planet, whereas the width and depth of these dips provide insights about the transiting events. Specifically, the width of the transit provides information about the orbital period and distance from the star, while the depth of the dip gives the relationship between the radius of the star and the radius of the transiting planet.

The shape of the dip in a light curve reveals significant details about the transiting object and its orbital geometry. There are generally two types of dips—a U-shaped dip and a V-shaped dip. The U-shaped dip results from a central, edge-on transit, which occurs at an orbital inclination close to 90° , where the planet crosses near the center of the stellar disk. Conversely, a V-shaped dip signifies a grazing transit, where the planet partially crosses the stellar disk. While this also requires a high orbital inclination, it is slightly less edge-on compared to a central transit. Additionally, V-shaped dips are characteristic of eclipsing binary stars, where two stars pass in front of each other (C. Magliano et al. 2023).

In this study, we used the Kepler light curve data for the exoplanet classification, consisting of labeled threshold crossing events (TCEs). A TCE occurs when the Kepler pipeline (J. M. Jenkins et al. 2010) detects periodic dips in stellar brightness, indicating a potential planetary transit. By analyzing these dips (TCEs), we can distinguish between actual exoplanet transit signals and other astrophysical phenomena or instrumental noise.

The TCEs shown in Figure 1 illustrate notable reductions in the normalized flux for the Kepler ID ‘KIC 6922244’ over time. Here, KIC refers to the Kepler Input Catalog. The sequence of periodic dips (TCE) or straight-line patterns,

Table 1
Description of Key Attributes in the Exoplanet Data Set

Attribute	Meaning
rowid	Identifier for each row in the TCE table
kepid	Kepler ID
tce_plnt_num	TCE planet number
tce_period	Orbital period
tce_time0bk	Transit epoch in BJD format
tce_duration	Transit duration
av_training_set	Classification labels

Table 2
Configuration of the Data Set used in the Study, Detailing the Number of Training Examples for each Label

Data Configuration	
Labels	Number of Training Set Examples
AFP	9596
NTP	2541
PC	3600
UNK	4630

shown in Figure 1, indicate the possibility of actual exoplanetary signals. This gives insight into the role of TCEs in classifying potential exoplanet candidates. We will use these labeled TCEs for the classification. The labels, assigned through human vetting, provide a reliable basis for distinguishing between true exoplanet candidates and false positives.

In this analysis, we accessed labeled TCEs data from the NASA Exoplanet Science Institute (2025), specifically the Autovetter Planet Candidate Catalog for Q1–Q17 DR24 (J. H. Catanzarite 2015). The data encompasses 20,367 TCEs, each with a set of attributes or columns, and we selected specific columns relevant to our analysis. Table 1 represents the seven key attributes with their usual meanings.

The most crucial attribute, *av_training_set*, consists of four distinct labels: nontransiting phenomena (NTP), astrophysical false positives (AFPs), planet candidates (PCs), and unknown (UNK). NTP includes phenomena that are not related to transits, such as stellar variability, while AFP signifies signals initially thought to be transits but later identified as noise or other astrophysical events. PC denotes probable transiting exoplanets, and UNK represents signals that cannot be clearly classified. These labels provide a comprehensive framework for categorizing various astrophysical signals. The detailed configuration is given in Table 2. For binary classification purposes, the data set is further categorized into “planet” candidates and “not-planet” candidates, the latter comprising NTP and AFP. We ignored TCEs with the UNK label. Table 2 indicates a substantial class imbalance, with 3600 “planet” candidates and the remainder classified as “not-planet” (NTP/AFP) candidates. J. H. Catanzarite (2015) reported that the inclusion criteria for these labels were developed by a precise procedure that included manual vetting and diagnostic evaluations. In this study, we assume these labels as the ground truth, considering that the potential impact of any inaccuracies on our model is negligible.

The light curves of the corresponding stars were downloaded through the Mikulski Archive for Space Telescopes (MAST), with each light curve generated through the Kepler

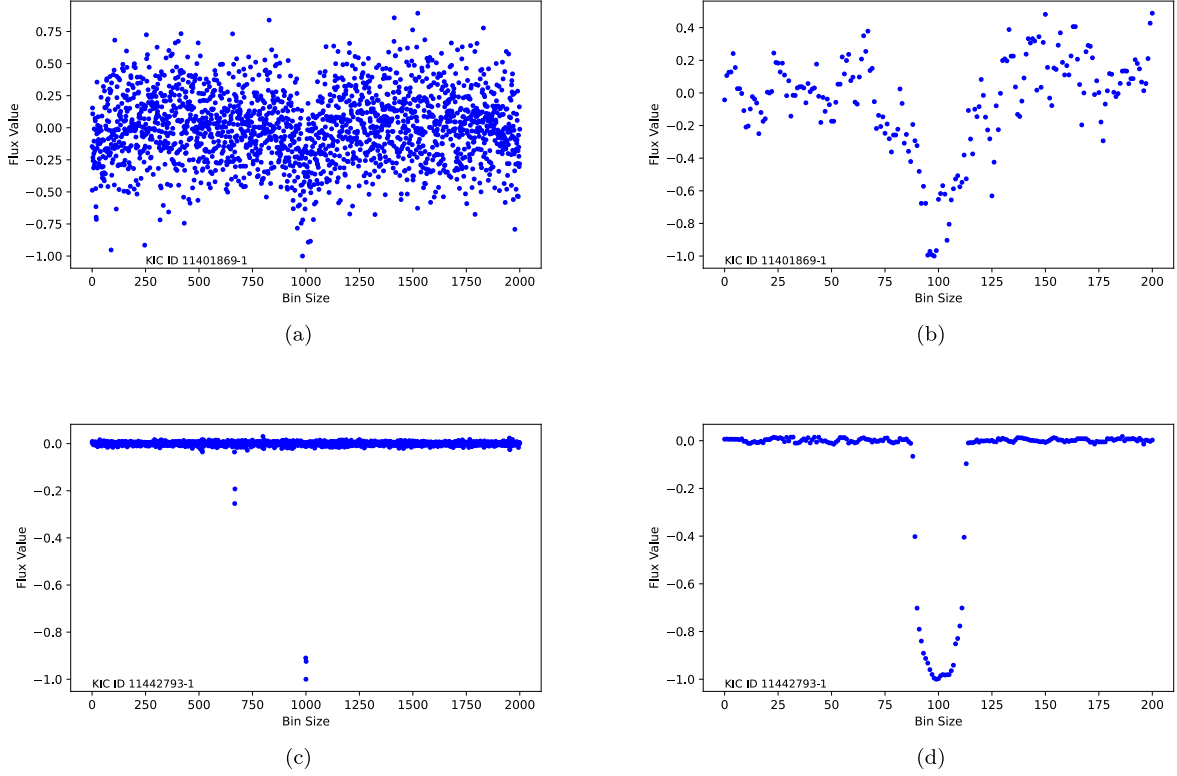


Figure 2. (a) and (c) represent the global light curves of AFP and PC candidates, respectively, consisting of full-length data with a resolution of 2001 bin size. (b) and (d) represent the local light curves for the same candidates, which are fixed-length representations focusing on the transit region with a bin size of 201. Each plot highlights the key periodic features distinguishing these astrophysical phenomena.

pipeline containing integrated flux measurements at 29.4 minutes intervals over up to four years.

Following the procedure of C. J. Shallue & A. Vanderburg (2018), we performed additional steps on generated light curves. First, we removed data points corresponding to transits of any confirmed planets within the system. Subsequently, we flattened the light curves by fitting a basis spline to the data and normalized them by dividing each light curve by the best-fit spline. Finally, the flattened light curves were phase folded on the TCE period and binned to produce a 1D vector. These binned and phase-folded light curves corresponding to each TCE provide local and global curves.

Figure 2 presents the global and local curves generated from the original light curve. The global light curve captured broad trends and changes throughout the event based on the photometric measurements, while the local light curve was concentrated to capture changes around the transit dips. Figure 2 provides the visual representation of both curves to evaluate the effectiveness of each curve in capturing critical patterns within the data. Figure 2(c) presents the global curve of the KIC ID 11442793-1, in which identifying a narrow transit was a difficult task or may provide us with a wrong prediction, whereas the local light curve (Figure 2(d)) makes it easier by providing subtle signals inside the data. Therefore, considering its effectiveness in capturing subtle patterns, we have used the local light curve as the primary approach in this study, which makes it different from the one that incorporates local and global views, such as C. J. Shallue & A. Vanderburg (2018).

Furthermore, we transformed the local curve data set into RP and GAF representations to feed them for training the model and evaluating its performance using evaluation metrics

to compare the effectiveness of each in identifying hidden patterns within the data.

2.2. Recurrence Plots

RP analysis is a technique deduced from nonlinear dynamics and is particularly appropriate for complex signals (D. F. Silva et al. 2013). It provides a strong alternative for demonstrating and studying the periodic signals. RPs are a visual matrix showing the point at which a phase space trajectory returns to a state already visited. This method provides insight into viewing and identifying hidden repetitive patterns that exhibit the characteristics of orbital movements of exoplanet data, which enables the analysis of m -dimensional trajectories within a 2D phase space. We define a recurrence matrix R , where each member is used to construct an RP (D. F. Silva et al. 2013).

$R_{k,l}$ is defined by the proximity of state vectors $x(k)$ and $x(l)$ by the following formula:

$$R_{k,l} = \phi(\epsilon - \|x(k) - x(l)\|), \quad x(\cdot) \in \mathbb{R}^m, \quad k, l = 1, \dots, N. \quad (1)$$

The state vectors at indices k and l are represented by $x(k)$ and $x(l)$, respectively, and the symbol $\|\cdot\|$ specifies the application of norm, between these observations. The closeness threshold is set by the parameter ϵ , and the Heaviside function, which is represented by ϕ in Equation (2), is as follows:

$$\phi(z) = \begin{cases} 0, & \text{if } z < 0 \\ 1, & \text{otherwise} \end{cases}. \quad (2)$$

Equation (1) states that if the m -dimensional trajectory of the time series at time k lies close to time l , then value 1 is




Figure 3. RPs for exoplanet classification. The upper two frames represent RPs for “not-planet” candidate cases, showing irregular patterns indicative of false positives. The lower two frames depict RPs for “planet” candidates, displaying clear, periodic structures consistent with true exoplanets.

assigned to $R_{k,l}$; otherwise, 0 is assigned. Graphically, it is represented by an $N \times N$ matrix, with pixels corresponding to 1 being colored as black, showing recurrence behavior, and those relating to 0 being white. For generating RPs, we choose a threshold of 0.65, a time delay of 2, and a value of $m = 2$, and apply the Euclidean distance to determine recurrences between the embedded state vectors because it captures more temporal dynamics since each point is now represented as a 2D state vector. The RPs produced using these parameters have dimensions of 199×199 pixels, as shown in Figure 3. This visualization offers a powerful tool for identifying periodic patterns and randomness within time-series data, with varied degrees of structure and complexity evident across different systems.

For example, the two lower frames in Figure 3 indicate real exoplanet prospects with regular transit properties, indicative of stable exoplanetary systems. Conversely, systems characterized by high disturbance or high stellar activity may be represented by more disorganized and fractured plots (refer to the upper two frames of Figure 3), which could lead to false positive detections of exoplanets.

These RPs were generated using some essential Python libraries, such as NumPy, PyRQA, and matplotlib, where recurrence quantification analysis (RQA) offers several measures to quantify the structures within the RPs, such as determinism, recurrence rate, and entropy (T. Rawald et al. 2017). These metrics can elucidate the nature of the dynamical systems observed, providing insights into the stability and predictability of exoplanetary orbits. Due to its outstanding suitability for handling nonstationary and nonlinear data, RQA is particularly well-suited for analyzing light curves, where noise and observational gaps frequently provide substantial obstacles.

2.3. Gramian Angular Fields

Z. Wang & T. Oates (2015) proposed GAFs to convert time-series data to visual patterns. GAFs are effective in capturing




Figure 4. GAFs for exoplanet classification. The upper two frames represent GAFs for “planet” cases, characterized by distinct periodic patterns indicative of planetary transits. The lower two frames show GAFs for “not-planet” candidate cases, which lack the clear periodic structures seen in true exoplanet detections, indicating the absence of exoplanets.

long-term dependencies, sparsity, and subtle signals, so we have used the GAFs method for converting the 1D time series to 2D visual representations to provide a framework for images wherein the image is characterized by angular coordinates rather than Cartesian coordinates. The encoding process first normalizes the time-series data into the range of $[-1, 1]$. The normalized input is transformed into polar coordinates to retain the temporal information of the input. After this, the trigonometric cosine function is used for temporal correlation by comparing each point in polar coordinates with every other point to generate a Gramian matrix of dimension $n \times n$, where n is the length of time-series data.

Given a time series $Y = \{y_1, y_2, \dots, y_n\}$ of n real-valued observations, we first normalize the values into the range $[-1, 1]$ using the following equation:

$$\tilde{y}_i = \frac{2y_i - \{\max(Y) + \min(Y)\}}{\max(Y) - \min(Y)}. \quad (3)$$

After normalization, each value in the time series is encoded into a polar coordinate system. Specifically, each value \tilde{y}_i is mapped to an angular cosine value, representing the angle $\theta_i = \arccos(\tilde{y}_i)$, while the temporal index is mapped to a radial coordinate. This mapping establishes a unique angular position for each time point in the series.

We then computed the GAF, where each element of the matrix G is defined as the trigonometric sum of angles between time steps:

$$G_{i,j} = \cos(\theta_i + \theta_j). \quad (4)$$

The resulting symmetric GAF matrices encapsulate the temporal patterns and underlying relationships observed in the exoplanetary time-series data. Following the aforementioned methodological approach, a set of images with the resolution of 201×201 was generated, which is depicted in Figure 4,

where the top two plots illustrate the periodic behavior of planet candidates by capturing subtle signals and the bottom two plots represent cases where no transiting planet candidates were detected.

3. Model Architecture

The classification of exoplanets using Kepler data is inherently challenging due to the intricate patterns embedded within light curves, which often contain deviations corresponding to planetary transits. Traditional approaches, such as CNNs, have demonstrated significant efficacy in various image-based classification tasks. However, CNNs use localized convolution operations, and it may be difficult for them to capture the global patterns and long-term dependencies due to their limited scope.

To address these limitations and leverage its effectiveness in various image-based tasks, we have implemented ViT (A. Dosovitskiy et al. 2020) in the exoplanet classification. ViT can be an effective approach to image-based classification by using a transformer architecture that treats an image as a sequence of patches, similar to the way words are processed in a sentence. Each patch is embedded into a token and processed through multiple layers of self-attention mechanisms, enabling ViT to capture both local and global relationships within the data. We have used RPs and GAFs as classification input tokens, allowing ViT to analyze Kepler data and classify exoplanets effectively.

Its architecture begins by dividing an input image I into nonoverlapping square patches; our model typically uses size 8×8 . Each patch is flattened into a vector representing the pixel data, which is then projected into a higher-dimensional space via a learnable linear transformation matrix, E . These patch embeddings are analogous to word embeddings in NLP. Alongside these tokens, a special classification token $[cls]$ is introduced at the beginning of the sequence. This token will be later used to accumulate and represent the overall information of the image, which will be used for the final classification output. Moreover, positional encodings are added to each embedding to ensure the spatial relationships between patches are maintained, providing critical context to the self-attention mechanism (A. Dosovitskiy et al. 2020).

Once the input sequence is ready, it is fed into the transformer encoder, a structure derived from the architecture proposed by A. Vaswani et al. (2017) and C. Chi et al. (2020). In our model, we used 10 such transformer layers. Each encoder layer has two primary components: a multihead self-attention mechanism and a feed-forward network (FFN), as shown in Figure 5. Using four heads per layer, this mechanism enables the model to capture complex interdependencies and global features from different regions of image, essential for identifying subtle planetary signals. The multihead self-attention mechanism allows every token to attend every other token in the sequence, capturing long-range dependencies and relationships between patches. It computes attention scores using the query, key, and value matrices obtained from the input embeddings. In the model architecture, the input embeddings of dimension 32 are first projected to a higher-dimensional space of dimension 128 to calculate the query, key, and value matrices across four attention heads. The concatenated output from all heads is then projected back to the original dimension 32, resulting in a higher parameter count for the multihead attention layer, as shown in the

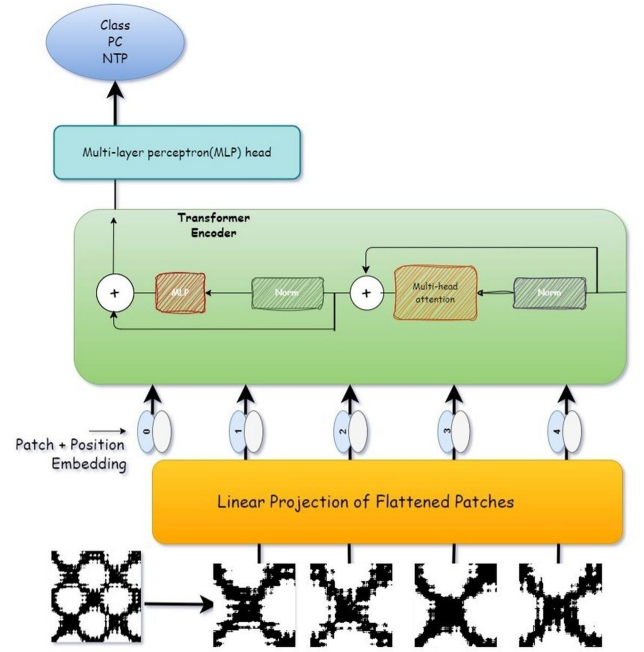


Figure 5. Architecture of the ViT model.

Appendix. By performing multiple self-attention operations in parallel, the model can understand intricate relationships across different parts of the image.

After the self-attention stage, the FFN, which consists of two fully connected layers (64 and 32 neurons) with a Gaussian error linear unit (GELU) activation function, further refines the information extracted by the attention mechanism. The GELU is a nonlinear activation function that combines the properties of the rectified linear unit (ReLU) and Gaussian distributions. Mathematically, it is defined as $\text{GELU}(x) = x \cdot \Phi(x)$, where $\Phi(x)$ is the cumulative distribution function of a Gaussian distribution. This function provides a smoother transition by gradually adjusting the activation, unlike ReLU, which immediately sets all negative values to zero. The GELU function models data more accurately by weighing the input data's transformation based on its estimated probability distribution, allowing it to smoothly activate neurons rather than switching them off completely. This characteristic makes GELU particularly effective in deep learning architectures such as ViT, where the diversity and complexity of the input data require nuanced processing. GELU helps in smoothing the learning pathway, allowing for better handling of the nonlinearities found in the data, thus contributing to more robust and higher-quality feature extraction.

At the end of the encoder sequence, the classification token, which now contains aggregated global information from all patches, passes through a multilayer perceptron (MLP) head. This MLP head, consisting of layers with 2048 and 1024 units, uses the refined global representations to determine the final classification. Each sublayer in the encoder, both the self-attention and the feed-forward layers, includes residual connections and layer normalization, which help stabilize the training process by improving gradient flow.

In this study, the ViT model was initialized with random weights and trained using the AdamW optimizer with hyperparameters $\beta_1 = 0.9$, $\beta_2 = 0.999$, a learning rate of 0.0001, and a weight decay of 0.00001 to ensure the model's

optimization and regularization. Input images were resized to 192×192 and divided into nonoverlapping patches of size 8×8 , resulting in 576 patches per image. Each patch was flattened and linearly projected into an embedding space with a dimensionality of 32.

4. Exoplanet Classification

A structured implementation of the above methodology was performed to handle complex astronomical data for exoplanet classification. To address this challenge, we conducted a comparative analysis of the ViT model based on primary inputs: RPs and GAFs. This study investigates how using different input types enhances the reliability and robustness of exoplanet classification with the model.

To train our model for exoplanet classification, we adopted a systematic approach inspired by C. J. Shallue & A. Vanderburg (2018). The data set was randomly partitioned into three distinct subsets to ensure a comprehensive and unbiased evaluation of the model. The partitioning was as follows:

1. *Training Set.* Comprising 80% of the total data set, this subset was used to train the model. In the case of RPs as input, the data set comprised 2848 “planet” candidates and 9741 “not-planet” candidates. Conversely, when using GAFs as input, the data set consisted of 2869 “planet” candidates and 9720 “not-planet” candidates. The substantial size of this set ensured that the model had sufficient data to learn and generalize from a wide range of examples.
2. *Validation Set.* This subset, comprising 10% of the data set, was used for hyperparameter tuning and model validation. In the case of RPs as input, the data set comprised 382 “planet” candidates and 1192 “not-planet” candidates. Meanwhile, for GAFs, it comprised 359 “planet” candidates and 1215 “not-planet” candidates. The validation set played a crucial role in preventing overfitting by providing a basis for tuning the model’s parameters and assessing its performance during training.
3. *Test Set.* The remaining 10% of the data set was reserved for testing the model’s performance. For RPs, this subset included 370 “planet” candidates and 1204 “not-planet” candidates. For GAFs, the test set consisted of 372 “planet” candidates and 1202 “not-planet” candidates. As detailed by C. J. Shallue & A. Vanderburg (2018), the test set provides an unbiased evaluation of the model’s accuracy and generalization capabilities.

The data for RPs and GAFs was partitioned to ensure that each split (i.e., training, validation, and testing) contained the same instances across both representations; however, slight differences in the number of planet and not-planet candidates between the two input types were due to the independent generation of RPs and GAFs images before partitioning into training, validation, and testing data sets. Additionally, the models were trained separately on each data representation. To ensure the reliability of the results, we have performed cross validation. This ensures that the performance of the model with both inputs is comparable. Given the significant class imbalance in the data set, evaluating the model’s performance based solely on accuracy would be inappropriate. Therefore, we assessed the model’s performance using recall and

precision metrics. These metrics provide a more accurate and meaningful evaluation of the model’s ability to handle an imbalanced data set. Precision and recall are critical metrics. Precision measures the model’s ability to correctly identify relevant instances (true positives), while recall assesses the model’s effectiveness in capturing all relevant instances. The F1 score, as the harmonic mean of Precision and recall, provides a balanced metric that addresses the trade-offs between these two aspects. The metrics are defined as:

Recall (sensitivity or true positive rate; TPR):

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}, \quad (5)$$

Precision (positive predictive value):

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}. \quad (6)$$

In these formulas:

1. TP (true positive) refers to the number of correctly identified planet candidates.
2. FN (false negative) refers to the number of actual planet candidates that were incorrectly classified as not-planet candidates.
3. FP (false positive) refers to the number of not-planet candidates that were incorrectly classified as planet candidates.

Upon evaluation of the test data set, the model with RPs as input exhibited significant performance as compared to GAFs by achieving a recall of 89.46% and a precision of 85.09%. These results underscore the model’s effectiveness when using RPs as input, demonstrating its ability to identify planet candidates accurately. This validates the model’s utility for classification tasks within the context of the Kepler-DR24 data set.

5. Results and Discussion

In this section, we present a comprehensive analysis of ViT performance in time-series data, which not only identifies the strengths of the proposed methodology but also points out its shortcomings, paving the way for further research and improvements in these areas.

In this study, we conducted a comparative analysis between two types of inputs, RPs and GAFs, by using a ViT model for the classification of exoplanets from time-series data. The performance of the model was evaluated over 100 epochs, and it took approximately 3–4 hr to train the model, which consists of ~ 2.39 million parameters. This computationally intensive task was executed using an NVIDIA DGX-1 8X Tesla V100 GPU, equipped with 32 GB of memory, and supports Tensor Cores, enhancing the performance of machine learning related tasks. The following sections provide a detailed analysis and interpretation of the results, supported by visual aids.

5.1. Performance Analysis of Recurrence Plots

The metrics for evaluating the model performance with RPs as input are summarized in Table 3. The recall, precision, and F1 score values are 0.8946, 0.8509, and 0.8722, respectively. These values indicate a balanced trade-off between precision and recall, highlighting the model’s proficiency in minimizing false positives, while effectively identifying true positives.

Table 3

General Model Performance for RPs and GAFs, Evaluated Using Precision, Recall, and F1 Score

Metric	RPs	GAFs
Precision	0.8509	0.7513
Recall	0.8946	0.7956
F1 Score	0.8722	0.7728

Further supporting these results, the confusion matrix for the model shown in Figure 6 visually illustrates its classification capabilities. The matrix reveals that the model accurately identified a significant number of both “not-planet” and “planet” candidates, with relatively few misclassifications. This underscores the model’s effectiveness in distinguishing between the classes.

To validate the model’s robustness, we applied cross validation, a widely used technique that assesses the model’s performance on independent data sets. Specifically, we used five-fold cross validation, where the data set is split into five equal parts (folds). In each iteration, four folds are used for training, and one fold is used for testing iteratively, ensuring that every fold is used as a test set once. Cross validation is crucial because it provides a more reliable estimate of model performance by reducing the variance associated with a single split of the data.

The cross-validation results yielded an average recall of 0.8810 and an average precision of 0.8426. To validate the consistency of the cross-validation results across all folds, we calculated the standard deviation and variance for each metric, as presented in Table 4. The low value of standard deviation indicates minimal variability across all folds. Similarly, the small variance highlights the robustness, irrespective of the specific fold partitioning. These results are consistent with the initial metrics, affirming the model’s stability and reliability across different data subsets. By using cross validation, we ensure that the model’s performance is not overly dependent on a particular subset of the data, thereby improving its generalizability to unseen data sets.

However, while cross validation confirms the model’s robustness, it also brings attention to the persistent challenge of class imbalance in the data set. In this context, the number of “not-planet” candidates substantially exceeds that of “planet” candidates. This imbalance can lead to biased model predictions, where the model is inclined to the majority class, thereby reducing its sensitivity to detecting minority class instances. To mitigate this issue, an undersampling technique was applied to the original data set by randomly reducing the number of instances in the majority class to achieve a balanced data set. This approach is designed to enhance the model’s capacity to accurately identify the minority class, thereby improving overall classification performance.

Following the application of undersampling, the model was trained again and the performance was re-evaluated following the training process outlined in Section 4. As shown in Table 5, there was a slight decrease in precision to 0.8426, recall to 0.8756, and the F1 score remained stable at 0.8587, indicating that the balance between precision and recall was effectively maintained.

Additionally, the receiver operating characteristic (ROC) curve offers a graphical representation of the model’s diagnostic ability by plotting the TPR on the y-axis and the




Figure 6. Confusion matrix for RPs data set, illustrating model performance. Class 0 represents “not-planet” candidates, while Class 1 represents “planet” candidates.

Table 4
Statistical Measures for Five-fold Cross Validation (RPs)

Metric	Mean	Standard Deviation	Variance
Precision	0.8426	0.0153	0.0002
Recall	0.8810	0.0149	0.0002
F1 Score	0.8614	0.0150	0.0002

Table 5
A Comparative Performance Analysis of RPs and GAFs as Input Data following Undersampling

Metric	RPs	GAFs
Precision	0.8426	0.7971
Recall	0.8756	0.8293
F1 Score	0.8587	0.8128

false positive rate (FPR) on the x-axis, allowing for an intuitive assessment of the model’s ability to distinguish between the positive and negative classes. A model that perfectly separates the classes would have an ROC curve that reaches the top-left corner of the plot, indicating a high TPR and a low FPR. As illustrated in Figure 7, the ROC curve showcases the model’s performance, with the diagonal line representing a random classifier, where TPR equals FPR and indicates no discriminatory power. The ROC curve for the general model reaches an ROC value of 0.9445, significantly exceeding the diagonal line and reflecting strong predictive power. Although the ROC value for the undersampled model slightly decreased to 0.9434. The model demonstrates robustness in handling class imbalance, showing strong classification performance with RPs as input, even without using undersampling techniques.

5.2. Performance Analysis of GAFs

The model’s performance using GAFs as input is also summarized in Table 3. The model achieved a precision of




Figure 7. ROC curve for the RPs-based model, comparing the performance of the general model (red) and the undersampled model (blue).




Figure 8. Confusion matrix for the GAFs data set, illustrating general model performance. Class 0 represents “not-planet” candidates, while Class 1 represents “planet” candidates.

0.7513, a recall of 0.7956, and an F1 score of 0.7728. Although these values are slightly lower than those obtained with RPs, they still reflect strong classification capabilities, demonstrating the model’s effectiveness in identifying and capturing relevant instances within the data set.

Further validation of these results is provided by the confusion matrix shown in Figure 8. The model correctly identified a substantial number of not-planet candidates (1104) and planet candidates (296), with relatively few misclassifications (98 false positives and 76 false negatives). This distribution supports the model’s significant classification

Table 6
Statistical Measures for Five-fold Cross Validation (GAFs)

Metric	Mean	Standard Deviation	Variance
Precision	0.7756	0.0470	0.0022
Recall	0.8102	0.0468	0.0021
F1 Score	0.7925	0.0468	0.0021

accuracy and effectiveness in distinguishing between positive and negative classes, corroborating the values of the metrics presented in Table 3.

To ensure the robustness of the model, we performed a cross-validation technique. The cross-validation results for the model with GAFs as input produced an average recall of 0.8102 and an average precision of 0.7756. Following the analysis performed for RPs, we also calculated the same statistical measures for GAFs as presented in Table 6 to thoroughly assess the consistency and reliability of its cross-validation results. The low standard deviations and variances confirm that the performance metrics for GAFs also exhibit minimal variability across folds. Although slightly lower than the general results, the values demonstrate the consistent performance of the model and confirm its robustness in handling various data subsets. This consistency across different data splits underscores the reliability and stability of the model when using GAFs as input.

To address the class imbalance in the data set, we applied an undersampling technique to balance the data set and reassess the model’s performance. The performance metrics for the undersampled data set are presented in Table 5. Following undersampling, the model achieved a precision of 0.7971, recall of 0.8293, and F1 score of 0.8128. This improvement suggests that the undersampling technique enhanced the model’s ability to classify exoplanets more effectively.

Similar to RPs, we also assessed the model’s performance using the ROC curve. The ROC curve, shown in Figure 9, provides a detailed illustration of the model’s ability to




Figure 9. ROC curve for the GAFs-based model, comparing the performance of the general model (red) and the undersampled model (blue).

distinguish between positive and negative classes. The model achieved an ROC value of 0.7934, indicating strong predictive power. After applying the undersampling technique, the ROC value improved to 0.9115, highlighting a significant enhancement in the model’s ability to discriminate between classes. This improvement suggests that the model benefits significantly from the undersampling technique, enhancing its ability to classify exoplanets more effectively.

We analyzed the results on balanced and imbalanced data sets to evaluate the model’s performance on both input representations. An undersampling technique was applied to address the class imbalance and the performance was observed on each representation under varying data conditions. The results indicate that RPs demonstrate consistent performance, highlighting their effectiveness in managing class imbalance, whereas GAFs exhibit reasonable performance on balanced data sets, suggesting their suitability with evenly distributed classes but limited adaptability to imbalanced data.

We also compared the performance of our results with other state-of-the-art models based on deep learning, as shown in Table 7, and found that our model achieved comparable results (a recall of 0.89 with a precision of 0.85). Although C. J. Shallue & A. Vanderburg (2018) achieved higher recall and precision by using deep learning models with local light curves, our proposed model is better suited for learning long-term dependencies with respect to previous approaches. Similarly, A. Malik et al. (2022) proposed a planet detection method based on classical machine learning on the Kepler and TESS data sets; although the precision of their model is lower than that of the current study, the recall is higher, which indicates that without human supervision, the model is prone to make mistakes on unseen data. In contrast, our study has also provided a balanced trade-off between precision and recall that was shown by the variation in the mean of cross validation, highlighting the effectiveness of ViT in exoplanet classification. Unlike previous studies, which primarily relied on deep network architectures, our approach used ViT, i.e., a

Table 7
Comparative Analysis of Exoplanet Classification Studies

Type	Recall	Precision
C. J. Shallue & A. Vanderburg (2018)	0.95	0.93
A. Malik et al. (2022)	0.96	0.82
Our method	0.89	0.85

deep network with an additional multihead self-attention layer for effective feature extraction. We have also applied dropout regularization in the model to prevent overfitting by randomly dropping some of the output neurons. So, future research should explore this in more detail to fully use the potential of ViT in this field.

Our study highlights the use of RPs and GAFs as input representations using the ViT model, which has not been explored much in the field of exoplanet classification. The results demonstrate the efficacy of RPs in tasks related to classification by highlighting their ability to deliver reliable results even under class imbalance. Moreover, our investigation highlights the potential of RPs and ViT in capturing complex relationships within diverse data distribution compared to GAFs.

Several limitations were identified in this study. While undersampling addresses class imbalance with GAFs effectively, its impact on RPs was limited because the samples were randomly dropped, indicating the need for more advanced data balancing techniques. Additionally, RPs as input to the ViT model have not demonstrated significant improvements in performance as compared to the earlier studies shown in Table 7. However, their performance highlights the potential for further enhancement and modifications in the architecture of the ViT model. One possible approach to improve the model’s performance is through a hybrid architecture of CNN and ViT because by combining the strengths of both models, the overall performance of the model may be improved. Another key consideration is that using local light curves as the basic input for generating image representations may not fully exploit ViT’s capabilities, so future exploration with global light curves as input could yield better results by providing a more comprehensive view of the data. Furthermore, implementing a ViT model increases computational cost and model complexity. Therefore, future research should focus on optimizing and fine-tuning the model’s architecture to balance performance with computational efficiency.

6. Summary and Conclusion

In conclusion, we implemented a ViT model for exoplanet classification, focusing on a comparative analysis of image-based inputs. We evaluated the model’s performance using two types of transformed inputs: RPs and GAFs. The results consistently revealed that images derived from RPs achieved higher classification results as compared to those obtained from GAFs. Additionally, the model that uses RPs is significant for handling imbalanced data sets effectively. This characteristic makes RPs a more suitable choice for scenarios involving data imbalance, in contrast to GAFs. For this application, RPs emerge as a more effective method for enhancing classification performance. Here, we summarize the key points of our paper.

1. This study presents a new approach to classify exoplanets. It uses ViT to analyze image representations of Kepler light curve data. The data is transformed into RPs and GAFs, allowing the model to effectively capture the intricate temporal relationships within the data.
2. The ViT model showed reasonable performance, especially when using RPs as input. It achieved an 89.46% recall rate and an 85.09% precision rate. These results emphasize the model's ability to accurately identify exoplanetary transits, demonstrating it as a potentially powerful tool for automated exoplanet detection.
3. We found that RPs outperformed GAFs in classifying exoplanets. By using RPs, a higher performance in classification was achieved, indicating that they are more successful in capturing the essential characteristics required to differentiate exoplanets from nonexoplanetary occurrences.
4. We also addressed the issue of class imbalance in the data set using undersampling techniques and the model still performed well, but the smaller data set size was seen as a limitation.
5. The evaluation metrics such as precision, recall, and F1 score provided a thorough assessment of the model's performance. The robustness of the model was also validated through ROC analysis. The RPs-based model achieved a high ROC value, indicating its reliability in accurately identifying true exoplanets.
6. Although this study yielded promising results, it is important to acknowledge some limitations. Specifically, the training of ViT was computationally intensive and the improvements obtained through undersampling were modest. To address these issues, it is recommended that future research prioritize the development of more advanced techniques that can effectively capture temporal dependencies and manage class imbalance. Furthermore, optimizing the ViT architecture to reduce training time and computational resource requirements will enhance the scalability and applicability of the model in various astrophysical scenarios.

This research significantly contributes to the classification of exoplanets by introducing a significant approach using ViT with image-based time-series representations. The findings suggest that this methodology has the potential to improve the accuracy and reliability of exoplanet detection, which is vital for future space missions and automated data analysis in astronomy. Moving forward, it is important to develop more sophisticated techniques that can effectively capture the temporal complexities in RPs. Additionally, optimizing the ViT architecture to reduce computational demands and training time is crucial to expand the model's usability in different astrophysical contexts. By addressing these challenges, future research can build upon this work to create even more powerful and versatile tools for classifying exoplanets.

Acknowledgments

In this work, we have used data from (a) the NASA Exoplanet Archive, run by the California Institute of Technology under an Exoplanet Exploration Program contract with NASA, and (b) the exoplanet.eu database, which compiles and organizes data from various sources, including ground-based and space-based telescopes. We are thankful to

the Inter-University Centre for Astronomy and Astrophysics (IUCAA) in Pune, India for providing access of campus research resources under the IUCAA Visiting Associateship program. S.C. is funded by the European Union (ERC, UNVEIL, 101076613). Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. S.C. acknowledges financial contribution from PRIN-MUR 2022YP5ACE.

Software: Numpy (C. R. Harris et al. 2020), Astropy (T. P. Robitaille et al. 2013), Scikit-learn (F. Pedregosa et al. 2011), Matplotlib (J. D. Hunter 2007), Scipy (P. Virtanen et al. 2020), PyRQA (T. Rawald et al. 2017), Pandas (The Pandas Development Team 2024), Tensorflow (J. Montalt Tordera 2022).

Appendix Model Architecture

The model comprises ~ 2.39 million parameters, and the configuration is in Table 8.

Table 8
Total Number of Trainable Parameters in the ViT Model

Sr. No.	Component	Parameters
1	Patch projection	2080
2	Positional embedding	18,432
3	Transformer encoder layers ($\times 10$)	
	Layer normalization (before)	64
	Multihead attention layer	16,800
	Layer normalization (after)	64
	Dense 1	2112
	Dense 2	2080
4	Layer normalization (final)	64
5	MLP head	
	Dense 1	67,584
	Dense 2	2,098,176
6	Final classification layer	2050
	Total	2,399,586

ORCID iDs

Anupma Choudhary  <https://orcid.org/0009-0004-6657-2260>

B. S. Kushvah  <https://orcid.org/0000-0001-9569-2297>

C. Swastik  <https://orcid.org/0000-0003-1371-8890>

References

- Arnab, A., Dehghani, M., Heigold, G., et al. 2021, in 2021 IEEE/CVF Int. Conf. on Computer Vision (ICCV) (New York: IEEE), 6816
- Baglin, A., Auvergne, M., Boisdard, L., et al. 2006, in 36th COSPAR Scientific Assembly, 36 (Beijing: COSPAR), 3749
- Basak, S., Mathur, A., Theophilus, A. J., Deshpande, G., & Murthy, J. 2021, *EPJST*, 230, 2221
- Batalha, N. M. 2014, *PNAS*, 111, 12647
- Borucki, W. J. 2016, *RPPh*, 79, 036901
- Borucki, W. J., Koch, D., Basri, G., et al. 2010, *Sci*, 327, 977

- Catanzarite, J. H. 2015, NASA Ames Research Center
- Chauvin, G., Lagrange, A. M., Zuckerman, B., et al. 2005, *A&A*, **438**, L29
- Chi, C., Wei, F., & Hu, H. 2020, *Adv. Neural Inf. Process Syst.*, **33**, 13564
- Dai, X., Chen, Y., Yang, J., et al. 2021, in *2021 IEEE/CVF Int. Conf. on Computer Vision (ICCV)* (New York: IEEE), 2968
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., et al. 2020, arXiv:2010.11929
- Eckmann, J.-P., Kamphorst, S. O., Ruelle, D., et al. 1995, in *World Scientific Series on Nonlinear Science Series A*, ed. D. Ruelle, **16** (Singapore: World Scientific), 441
- Fischer, D. A., & Valenti, J. 2005, *ApJ*, **622**, 1102
- Freitas, P. G., De Lima, R. G., Lucafo, G. D., & Penatti, O. A. 2023a, in *2023 15th Int. Conf. on Quality of Multimedia Experience (QoMEX)* (Piscataway, NJ: IEEE), 165
- Freitas, P. G., De Lima, R. G., Lucafo, G. D., & Penatti, O. A. B. 2023b, in *31st European Signal Processing Conf. (EUSIPCO)* (New York: IEEE), 1035
- Harris, C. R., Millman, K. J., Van Der Walt, S. J., et al. 2020, *Natur*, **585**, 357
- Hill, M. L., Bott, K., Dalba, P. A., et al. 2023, *AJ*, **165**, 34
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Jenkins, J. M., Caldwell, D. A., Chandrasekaran, H., et al. 2010, *ApJL*, **713**, L87
- Johnson, J. A., Petigura, E. A., Fulton, B. J., et al. 2017, *AJ*, **154**, 108
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. 2012, *Adv. Neural Inf. Process Syst.*, **1097**, 1
- LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. 1998, *Proc. IEEE*, **86**, 2278
- Magliano, C., Kostov, V., Cacciapuoti, L., et al. 2023, *MNRAS*, **521**, 3749
- Malik, A., Moster, B. P., & Obermeier, C. 2022, *MNRAS*, **513**, 5505
- Mayor, M., Queloz, D., Marcy, G., et al. 1995, *IAUC*, **6251**, 1
- McCaulliff, S. D., Jenkins, J. M., Catanzarite, J., et al. 2015, *ApJ*, **806**, 6
- Montalt Tordera, J. 2022, TensorFlow MRI, v0.22.0, Zenodo, doi:10.5281/zenodo.7120930
- Mulenga, M., Phiri, M., Simukonda, L., & Alaba, F. A. 2023, *IEEEA*, **11**, 129449
- NASA Exoplanet Science Institute 2025, Planetary Systems Table, IPAC, doi:10.26133/NEA12
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *JMLR*, **12**, 2825
- Perryman, M. 2018, *The Exoplanet Handbook* (Cambridge: Cambridge Univ. Press)
- Pyne, T., Banyal, R. K., Swastik, C., & De, A. 2025, *AJ*, **169**, 13
- Rawald, T., Sips, M., & Marwan, N. 2017, *CG*, **104**, 101
- Ren, B. B., Benisty, M., Ginski, C., et al. 2023, *A&A*, **680**, A114
- Robitaille, T. P., Tollerud, E. J., Greenfield, P., et al. 2013, *A&A*, **558**, A33
- Shallue, C. J., & Vanderburg, A. 2018, *AJ*, **155**, 94
- Silva, D. F., Souza, V. M. D., & Batista, G. E. 2013, in *IEEE 13th Int. Conf. on Data Mining (New York: IEEE)*, 687
- Swastik, C., Banyal, R. K., Narang, M., Unni, A., & Sivarani, T. 2024, *AJ*, **167**, 270
- Swastik, C., Banyal, R. K., Narang, M., et al. 2021, *AJ*, **161**, 114
- Swastik, C., Banyal, R. K., Narang, M., et al. 2022, *AJ*, **164**, 60
- Swastik, C., Banyal, R. K., Narang, M., et al. 2023, *AJ*, **166**, 91
- Szegedy, C., Liu, W., Jia, Y., et al. 2015, in *Proc. IEEE Conf. on Computer Vision and Pattern Recognition* (New York: IEEE), 1
- The Pandas Development Team 2024, pandas-dev/pandas: Pandas, Zenodo, doi:10.5281/zenodo.3509134
- Udry, S., & Santos, N. C. 2007, *ARA&A*, **45**, 397
- Unni, A., Narang, M., Sivarani, T., et al. 2022, *AJ*, **164**, 181
- Vaswani, A., Shazeer, N., Parmar, N., et al. 2017, *Adv. Neural Inf. Process Syst.*, **30**, 1
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, **17**, 261
- Wahhaj, Z., Benisty, M., Ginski, C., et al. 2024, *A&A*, **687**, A257
- Wang, Z., & Oates, T. 2015, in *Workshops at the Twenty-ninth AAAI Conf. on Artificial Intelligence* (Washington, DC: AAAI), <https://aaai.org/proceeding/ws15-14/>
- Wright, J. T., & Gaudi, B. S. 2012, in *Planets, Stars and Stellar Systems*, ed. T. D. Oswalt, L. M. French, & P. Kalas (Berlin: Springer), 489
- Zheng, S., Lu, J., Zhao, H., et al. 2021, in *IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)* (New York: IEEE), 6877