
Automatically detecting anomalous exoplanet transits

Christoph J. Hönes

University of Amsterdam
christoph.hoenes@gmail.com

Benjamin Kurt Miller

University of Amsterdam
b.k.miller@uva.nl

Ana M. Heras

European Space Agency, ESTEC
ana.heras@esa.int

Bernard Foing

Leiden Observatory, VU Amsterdam,
ISU, ILEWG EMMESI
foing@strw.leidenuniv.nl

Abstract

Raw light curve data from exoplanet transits is too complex to naively apply traditional outlier detection methods. We propose an architecture which estimates a latent representation of both the main transit and residual deviations with a pair of variational autoencoders. We show, using two fabricated datasets, that our latent representations of anomalous transit residuals are significantly more amenable to outlier detection than raw data or the latent representation of a traditional variational autoencoder. We then apply our method to real exoplanet transit data. Our study is the first which automatically identifies anomalous exoplanet transit light curves. We additionally release three first-of-their-kind datasets to enable further research.

1 Introduction

Exoplanets An exoplanet obscuring the line-of-sight between a star and telescope, thereby temporarily reducing the observed brightness (light curve) of that star, is known as a transit [1]. Today, detected transits number in the thousands [2] and are normally identified using variants of the Box Least Squares algorithm [3, 4]. Deviation from a theoretical transit light curve model [5] indicates unusual or unknown astrophysical phenomena (e.g. [6, 7, 8]). Therefore, finding anomalous transit light curves is of great scientific interest. Large-scale space telescope surveys and improvements in automated detection methods will lead to unprecedentedly high detection rates of transits, demanding a fast outlier detection solution. We propose to automate the identification of anomalous transits.

Related work Deep learning has been applied to exoplanet detection [9] and confirmation, also known as automated vetting [10, 11, 12]. Feature extraction and clustering have been applied to Kepler data [13, 14] to find stellar variability outliers [15]. Similarly, [16] use an ensemble of outlier detection methods on features extracted from light curve data to look for special stellar activity.

Contribution This paper is concerned with an entirely novel topic: Automatically identifying anomalous exoplanet transit light curves. (a) Our paper represents the first systematic study of automated detection of exoplanet transit shape anomalies. (b) Our proposed architecture assigns an anomaly score to each transit light curve. On one fabricated dataset with known labels, our architecture produces an average precision 61.81% higher than using anomaly detection on raw data alone and 23.08% higher than a naive variational autoencoder-based dimension reduction technique [17, 18]. (c) We present two sets of fabricated, labeled transit data along with an unlabeled collection

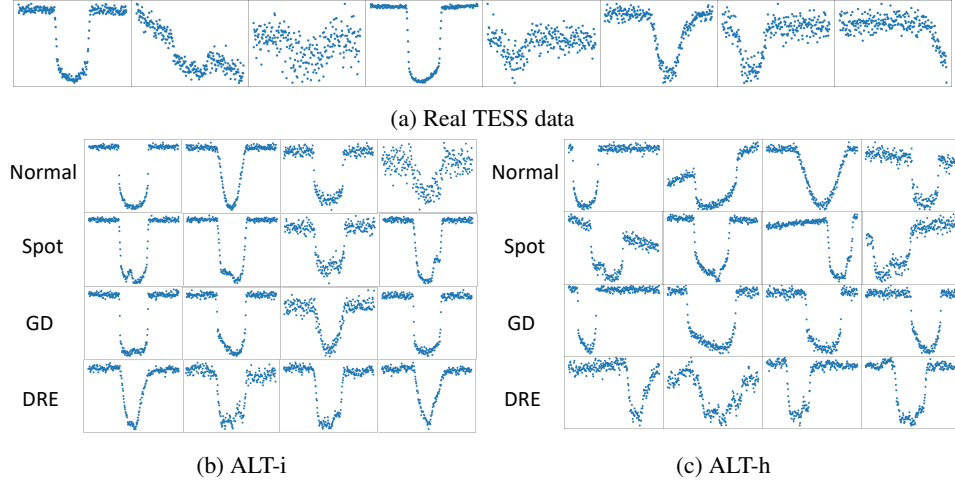


Figure 1: Examples of fabricated and genuine transit light curves. Genuine TESS data on top. Fabricated data from ALT-i and ALT-h on the bottom. The y-axes represent the normalized (unitless) flux (brightness of the star) and the x-axes represent the phase (time relative to the transit duration).

of genuine transit data.¹ Each dataset is the first-of-its-kind for deep learning methods, thereby catalyzing future deep learning research into exoplanetary science.

2 Genuine and fabricated transit light curves

Genuine transit data We curated a set of transit light curves with a total of 1668 targets from the Transiting Exoplanet Survey Satellite (TESS) mission [19], which is publicly available at MAST [20], using the *astroquery* [21] package. As the data is unlabeled, we can make no guarantees about the distribution of anomalies to regular transit light curves. Data processing closely follows the pipelines of exoplanetary scientists. The most important processing step is phase-folding where the time series is transformed into phase space (time relative to the orbital period). If the parameters used for phase-folding are accurate, the transits within a light curve perfectly overlay and have zero phase at the transit mid-point. Inaccuracies in the currently available parameter estimates of period, epoch and transit duration can lead to *horizontal shifts* of the transit shape within the extracted observation window or different *horizontal scale* (e.g. apparent transit width within the observation window). This needs to be considered in the model design. Further details can be found in Appendix A.

Fabricated transit data Quantification of the precision and recall of anomaly detection requires labeled data. We generate fabricated data labeled according to its anomaly class: standard transit, occulted star spot, gravity darkened (GD), and disintegrating rocky exoplanet (DRE). All but the first option is considered an anomaly. Standard transits are simulated using the Mandel and Agol formalism [5] implemented in the *batman* [22] library, star spot anomalies using the *e11c* package [23], and the GD effects with the *PyTransit* library [24]. Details about the data generation can be found in Appendix A and physics details in Appendix B. We call our fabricated data the Artificial Labeled Transit dataset (ALT). We produced an idealized version, ALT-i, and a more difficult, or “hard,” version ALT-h. The ALT-i version is completely detrended with transits occurring at consistent horizontal shift and scale. The ALT-h version features a vast range of trends and horizontal shifts and scales. Both contain 25,680 training transits with 6,420 standard transits and 6,420 of each anomaly class. Each test sets has 2,700 regular transits and 100 of each anomaly class. The anomalous samples were hand-selected to ensure they are distinguishable from a regular transit. Data examples are shown in Figure 1.

¹The source code and the data can be found at:
<https://github.com/ChristophHoenes/AnomalousExoplanetTransits>

3 Automatic detection of anomalous transit light curves

We aim to learn a featurization of light curve data which is amenable to unsupervised anomaly detection [25] and incorporates the inductive bias that a transit can be additively decomposed into a *standard transit* component and a *potentially anomalous residual* component. Each piece is modeled by a variational autoencoder (VAE), *TransitVAE* and *ResidualVAE* respectively, such that their sum is a learned reconstruction of the input light curve. The latent representation of the ResidualVAE is therefore a clustered and dimensionality-reduced representation of relevant anomalous information.

Method The TransitVAE is trained to predict a theoretical transit model according to [5]. It is always trained on fabricated data to enable using a ground truth general transit shape label ξ . Those labels are always centered, detrended, and noise-free. To better model the observed transit x we apply a deterministic transform f with five parameters that accounts for shift, trend and scale of x . Find implementation details of f in Appendix C.

We define TransitVAE’s encoder $T_{\phi_T}^{\text{enc}}(x) = (\mu_T, \sigma_T, \hat{t})$, where x is the light curve data, μ_T and σ_T are the mean and standard deviation of a normal distribution, \hat{t} are transformation parameters, and ϕ_T are the model weights. A posterior predictive sample is produced by taking $z_T \sim \mathcal{N}(\mu_T, \sigma_T^2 \mathbf{I})$ and first evaluating the decoder $T_{\theta_T}^{\text{dec}}(z_T) = \hat{\xi}$, where θ_T are model weights and $\hat{\xi}$ is the estimated detrended and centered transit. The estimated observed transit is recovered by $f(\hat{\xi}, \hat{t}) = \hat{x}$.

ResidualVAE is a traditional VAE. Its latent distribution is defined by encoder $R_{\phi_R}^{\text{enc}}(y) = (\mu_R, \sigma_R)$, where $y = x - \hat{x}$ is the residual to reconstruct, μ_R and σ_R parametrize the latent normal distribution and ϕ_R are model weights. The reconstruction is estimated by the decoder $R_{\theta_R}^{\text{dec}}(z_R) = \hat{y}$, where $z_R \sim \mathcal{N}(\mu_R, \sigma_R^2 \mathbf{I})$ and θ_R are model weights.

TransitVAE is trained by minimizing the VAE loss where the reconstruction term compares the detrended and centered transit estimate $\hat{\xi}$ to the ground truth ξ . An additional regression loss for the predicted transformation parameters \hat{t} is applied. ResidualVAE is trained similarly but attempts to reconstruct y and has no transformation loss. Training occurs in two stages starting by optimizing ϕ_T & θ_T , then TransitVAE is fixed and ϕ_R & θ_R are optimized. An overview of the losses, the hyperparameter details and a sketch of the design can be found in Appendix C.

4 Experiments on artificial data

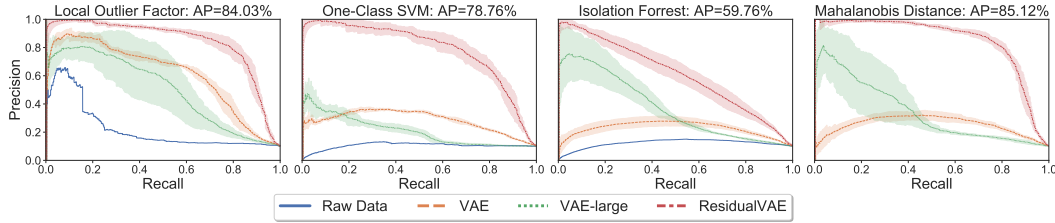


Figure 2: PR curves for various classifiers and featurizations on the ALT-i test set. Mean and 95% confidence intervals are shown across 5 runs with different random seeds. AP values of the best featurization are given in the title for each classifier. ResidualVAE always had the highest AP.

Evaluation methods We quantify the performance of each model on artificial, labeled data by plotting the precision versus the recall, i.e. a Precision-Recall (PR) curve. We summarize this with `scikit-learn`’s formulation of Average Precision (AP), essentially the area under the curve [26].

Setup We consider the following unsupervised anomaly detection classifiers: Local Outlier Factor (LOF) [27], the One-Class Support Vector Machine [28], the Isolation Forrest [29] and the Mahalanobis distance [30]. Each classifier is applied to several featurizations including the raw data x and the latent representation μ_R of ResidualVAE. We trained two more architectures that reconstruct x as a benchmark. The first, VAE, has as many parameters as ResidualVAE, and the second, VAE-large, has the same number of parameters as TransitVAE and ResidualVAE combined. Each classifier is

also applied to the mean latent representation of VAE and VAE-large. The procedure was repeated five times with different random seeds to report uncertainty bounds.

Table 1: Average precision (in %) and standard deviation over 5 independent runs for various classifiers with different featurizations on the ALT-h test set. Each classifier’s best result is bolded.

	Local Outlier Factor	One-class SVM	Isolation forrest	Mahalanobis
Raw data	10.59 +/- 0	9.29 +/- 0	7.83 +/- 0	-
VAE	13.07 +/- 2.12	7.78 +/- 0.23	8.61 +/- 0.31	8.62 +/- 0.33
VAE-large	12.83 +/- 2.04	8.30 +/- 0.37	9.00 +/- 0.64	8.48 +/- 1.32
ResidualVAE	25.10 +/- 3.32	16.37 +/- 1.17	17.20 +/- 2.37	17.57 +/- 2.91

Discussion of results ResidualVAE’s featurization led to superior precision across nearly all recall levels for every classifier on the ALT-i test set, see Figure 2. AP results can be found in Appendix D.

ResidualVAE outperforms all other methods across all classifiers on the ALT-h test set as seen in Table 1 as measured by AP. ALT-h was more difficult with the highest mean AP reaching only 25%. This implies that the shifting, scaling, and trending of the data can have a significant effect on performance. The PR curves were qualitatively similar to ALT-i results, thus only AP is reported.

5 Experiments on TESS data

Evaluation methods The Classifier Two-Sample Test (C2ST) [31, 32] attempts to distinguish whether two sets of samples are drawn from the same distribution by means of a binary classifier, e.g., a Multi Layer Perceptron (MLP). Assuming the classifier has converged and is sufficiently expressive, two sets of held-out samples drawn from the same distribution will be classified at near-chance level. If the sets are drawn from sufficiently distinct distributions, they will be accurately classified. One can test statistical significance, such as with a p-value, using the accuracy on the held-out data.

Setup We apply the TransitVAE that was pre-trained on ALT-h to extract residuals from the real TESS dataset. Then we train the ResidualVAE on the residuals of the real data. We use the LOF scores based on the ResidualVAE features to rank the transits since it performed best on the artificial datasets. To increase robustness of predictions we use an ensemble of five independent runs of TransitVAEs with three independent runs of the ResidualVAE each. To assess the quality of the ranking we perform a C2ST comparing various strata to each other. Our hypothesis is that highly-ranked strata are drawn from a different distribution (outliers) than lower-ranked strata. For example, an equally sized set of held-out samples from the top 10% (90th percentile) are compared to the bottom 90%. A high C2ST with a small p-value provides some evidence that the 90th percentile contains outliers. As classifier we use a simple MLP with two hidden layers (64 and 10 neurons). As input we use the raw transit data and evaluation is done with five-fold cross validation. Our significance threshold is set to $p \leq 0.05$.

Table 2: C2ST accuracies and p-values for various strata of real TESS transits with ResidualVAE featurizations. Samples from the identified subset are compared to samples from the remainder of the data. Rank is determined by LOF score. Statistically significant results, $p \leq 0.05$, are bolded.

Percentiles	Up to 90th	Up to 50th	80th through 65th	65th through 50th
Accuracy (in %)	71.14	65.77	64.00	56.80
p-value	$1.44e^{-4}$	$2.34e^{-9}$	$1.76e^{-3}$	$9.67e^{-2}$

Discussion of results Table 2 shows the classifier accuracies and corresponding p-values comparing samples from the identified subset to the remainder of the data. Tests sampling strata from the top and middle of the ranking are shown to confirm that the highly ranked data contains more anomalies than the middle rankings. All tests resulted in a classification accuracy above chance and all but one was considered significant by our criterion.

6 Conclusion

In this paper we introduced the first automated approach for anomalous transit shape detection and evaluated it on artificial and real data. On artificial data, we showed that ResidualVAE extracted features better suited to unsupervised anomaly detection methods than using raw data or featurizations from a standard VAE. ResidualVAE identified statistically significant outliers in real data.

Limitations Our method depends on estimates of epoch, period and transit duration to extract the transits from light curve data. Current estimates provided by the space telescope’s processing pipeline may suffer from inaccuracies which leads to shifts and scale mismatches. While the deterministic transform $f(\hat{\xi}, \hat{t})$ is designed to mitigate this effect, its success is limited by the accuracy of the predictions for \hat{t} . Errors in \hat{t} cause artifacts in the residuals which can negatively affect the performance. We do not present any results on TransitVAE’s performance of predicting t because without any reference for comparison the regression loss is less meaningful.

Despite the efforts made in creating realistic fabricated data it is always possible that the real data distribution does not align with our simulated one. This might introduce unwanted biases.

Future Work A direction for future work could be improving the accuracy of predicting t , since correctly aligning the theoretical transit shape with the observed data before calculating the residuals is essential for the performance of our method. Other works on anomaly detection use the reconstruction error of an AE or the reconstruction probability of a VAE as metric for separating anomalies [33]. This would enable independence from traditional anomaly classifiers. We did not pursue this direction in our work but experimenting with such an approach and comparing the results could be done in future work.

Social Impact Our method uses existing telescope data and does not have application outside of exoplanetary science. The main positive impacts would be reducing the effort of astronomers for visual inspection of data and potentially inspiring future scientists. The main negative impact is the (relatively small) cost to the environment of training on GPUs.

Acknowledgements

This paper includes data collected with the TESS mission, obtained from the MAST data archive at the Space Telescope Science Institute (STScI). Funding for the TESS mission is provided by the NASA Explorer Program. STScI is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS 5–26555.

References

- [1] Michael Perryman. *Transits*, page 153–328. Cambridge University Press, 2 edition, 2018.
- [2] NASA. Exoplanet archive. [Online; accessed September 09, 2021].
- [3] G. Kovács, S. Zucker, and T. Mazeh. A box-fitting algorithm in the search for periodic transits. *Astronomy & Astrophysics*, 391:369–377, August 2002.
- [4] Michael Hippke and René Heller. Optimized transit detection algorithm to search for periodic transits of small planets. *Astronomy & Astrophysics*, 623:A39, Feb 2019.
- [5] Kaisey Mandel and Eric Agol. Analytic light curves for planetary transit searches. *The Astrophysical Journal Letters*, 580(2):L171, 2002.
- [6] Silva-Valio, A., Lanza, A. F., Alonso, R., and Barge, P. Properties of starspots on corot-2 *. *A&A*, 510:A25, 2010.
- [7] John P. Ahlers, Ethan Kruse, Knicole D. Colón, Patrick Dorval, Geert Jan Talens, Ignas Snellen, Simon Albrecht, Gilles Otten, George Ricker, Roland Vanderspek, and et al. Gravity-darkening analysis of the misaligned hot jupiter mascara-4 b. *The Astrophysical Journal*, 888(2):63, Jan 2020.
- [8] Rik van Lieshout and Saul A. Rappaport. Disintegrating rocky exoplanets. *Handbook of Exoplanets*, page 1527–1544, 2018.

- [9] Kyle A. Pearson, Leon Palafox, and Caitlin A. Griffith. Searching for exoplanets using artificial intelligence. *Monthly Notices of the Royal Astronomical Society*, 474(1):478–491, February 2018.
- [10] Sean D. McCauliff, Jon M. Jenkins, Joseph Catanzarite, Christopher J. Burke, Jeffrey L. Coughlin, Joseph D. Twicken, Peter Tenenbaum, Shawn Seader, Jie Li, and Miles Cote. Automatic classification of keplerplanetary transit candidates. *The Astrophysical Journal*, 806(1):6, Jun 2015.
- [11] Susan E. Thompson, Fergal Mullally, Jeff Coughlin, Jessie L. Christiansen, Christopher E. Henze, Michael R. Haas, and Christopher J. Burke. A machine learning technique to identify transit shaped signals. *The Astrophysical Journal*, 812(1):46, Oct 2015.
- [12] Christopher J. Shallue and Andrew Vanderburg. Identifying exoplanets with deep learning: A five-planet resonant chain around kepler-80 and an eighth planet around kepler-90. *The Astronomical Journal*, 155(2):94, Jan 2018.
- [13] Susan E Thompson, Dorothy Fraquelli, Jeffrey E van Cleve, and Douglas A Caldwell. Kepler: A search for terrestrial planets. *Kepler Archive Manual*, 2016.
- [14] William J Borucki. Kepler mission: development and overview. *Reports on progress in physics. Physical Society (Great Britain)*, 79(3):036901, March 2016.
- [15] Daniel Giles and Lucianne Walkowicz. Systematic serendipity: a test of unsupervised machine learning as a method for anomaly detection. *Monthly Notices of the Royal Astronomical Society*, 484(1):834–849, Dec 2018.
- [16] Isadora Nun, Pavlos Protopapas, Brandon Sim, and Wesley Chen. Ensemble Learning Method for Outlier Detection and its Application to Astronomical Light Curves. *The Astronomical Journal*, 152(3):71, September 2016.
- [17] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [18] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *International conference on machine learning*, pages 1278–1286. PMLR, 2014.
- [19] George R. Ricker, Joshua N. Winn, Roland Vanderspek, David W. Latham, Gaspar A. Bakos, Jacob L. Bean, Zachory K. Berta-Thompson, Timothy M. Brown, Lars Buchhave, Nathaniel R. Butler, R. Paul Butler, William J. Chaplin, David B. Charbonneau, Jørgen Christensen-Dalsgaard, Mark Clampin, Drake Deming, John P. Doty, Nathan De Lee, Courtney Dressing, Edward W. Dunham, Michael Endl, François Fressin, Jian Ge, Thomas Henning, Matthew J. Holman, Andrew W. Howard, Shigeru Ida, Jon M. Jenkins, Garrett Jernigan, John Asher Johnson, Lisa Kaltenegger, Nobuyuki Kawai, Hans Kjeldsen, Gregory Laughlin, Alan M. Levine, Douglas Lin, Jack J. Lissauer, Phillip MacQueen, Geoffrey Marcy, Peter R. McCullough, Timothy D. Morton, Norio Narita, Martin Paegert, Enric Palle, Francesco Pepe, Joshua Pepper, Andreas Quirrenbach, Stephen A. Rinehart, Dimitar Sasselov, Bun’ei Sato, Sara Seager, Alessandro Sozzetti, Keivan G. Stassun, Peter Sullivan, Andrew Szentgyorgyi, Guillermo Torres, Stephane Udry, and Joel Villaseñor. Transiting Exoplanet Survey Satellite. *Journal of Astronomical Telescopes, Instruments, and Systems*, 1(1):1 – 10, 2014.
- [20] NASA. Mikulski archive for space telescopes (mast). [Online; accessed June 27, 2021].
- [21] Adam Ginsburg, Brigitta M Sipőcz, CE Brasseur, Philip S Cowperthwaite, Matthew W Craig, Christoph Deil, Austen M Groener, James Guillochon, Giannina Guzman, Simon Liedtke, et al. Astroquery: an astronomical web-querying package in python. *The Astronomical Journal*, 157(3):98, 2019.
- [22] Laura Kreidberg. Batman: basic transit model calculation in python. *Publications of the Astronomical Society of the Pacific*, 127(957):1161, 2015.
- [23] PFL Maxted. ellc: A fast, flexible light curve model for detached eclipsing binary stars and transiting exoplanets. *Astronomy & Astrophysics*, 591:A111, 2016.
- [24] Hannu Parviainen. PYTRANSIT: fast and easy exoplanet transit modelling in PYTHON. *MNRAS*, 450(April):3233–3238, 2015.
- [25] Rong Yao, Chongdang Liu, Linxuan Zhang, and Peng Peng. Unsupervised anomaly detection using variational auto-encoder based feature extraction. In *2019 IEEE International Conference on Prognostics and Health Management (ICPHM)*, pages 1–7, 2019.

- [26] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [27] Markus M. Breunig, Hans-Peter Kriegel, Raymond T. Ng, and Jörg Sander. Lof: Identifying density-based local outliers. *SIGMOD Rec.*, 29(2):93–104, May 2000.
- [28] Bernhard Schölkopf, Robert Williamson, Alex Smola, John Shawe-Taylor, and John Platt. Support vector method for novelty detection. In *Proceedings of the 12th International Conference on Neural Information Processing Systems*, NIPS’99, page 582–588, Cambridge, MA, USA, 1999. MIT Press.
- [29] Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. Isolation forest. In *2008 Eighth IEEE International Conference on Data Mining*, pages 413–422, 2008.
- [30] Peter Rousseeuw and Katrien Driessen. A fast algorithm for the minimum covariance determinant estimator. *Technometrics*, 41:212–223, 08 1999.
- [31] David Lopez-Paz and Maxime Oquab. Revisiting classifier two-sample tests, 2018.
- [32] Jerome H Friedman. On multivariate goodness-of-fit and two-sample testing. *STATISTICAL PROBLEMS IN PARTICLE PHYSICS, ASTROPHYSICS AND COSMOLOGY*, page 311, 2003.
- [33] Jinwon An and Sungzoon Cho. Variational autoencoder based anomaly detection using reconstruction probability. *Special Lecture on IE*, 2(1):1–18, 2015.
- [34] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- [35] Claret, A. Limb and gravity-darkening coefficients for the tess satellite at several metallicities, surface gravities, and microturbulent velocities. *A&A*, 600:A30, 2017.
- [36] Andrea Asperti and Matteo Trentin. Balancing reconstruction error and kullback-leibler divergence in variational autoencoders, 2020.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? **[Yes]** Main contributions are listed in section 1 (Introduction) paragraph ‘Contributions’. Our proposed method performs better than the mentioned baselines in terms of AP (Tables 1 and 4) and PR curves (Figure 2)
 - (b) Did you describe the limitations of your work? **[Yes]** Limitations discussed in section 6 (Conclusion)
 - (c) Did you discuss any potential negative societal impacts of your work? **[Yes]** see section 6 (Conclusion)
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? **[Yes]** We have read the ethics review guidelines and made sure that our paper certainly conforms. Since this is a scientific work on previously collected telescope data, the ethical externalities are limited.
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? **[N/A]**
 - (b) Did you include complete proofs of all theoretical results? **[N/A]**
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **[Yes]** Relevant source code is available in a GitHub repository (see link in footnote 1 in section 1 paragraph ‘Contributions’). Binaries of our datasets are provided in the same repository.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **[Yes]** Train/test sets are described in section 2; Architecture and hyperparameter details are provided in Appendix C; Further details on data processing and parameters for artificial transit synthesis are described in Appendix A.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **[Yes]** 95% confidence intervals or standard deviation over 5 runs with random seeds are used in Figure 2 and Table 1 respectively.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **[Yes]** This is discussed in Appendix C.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? **[Yes]** TESS data [20, 21] cited in section 2; The software packages used for this work have been properly attributed [21, 22, 23, 24, 26, 34]
 - (b) Did you mention the license of the assets? **[Yes]** Publicly available open data (see section 2; Genuine Transit Data)
 - (c) Did you include any new assets either in the supplemental material or as a URL? **[Yes]** Link to source code and datasets is available (see 3a).
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? **[N/A]** Not necessary since open data and properly attributed via citation of origin.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **[N/A]** Not necessary since no personal information or offensive content is present in the data
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? **[N/A]**
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? **[N/A]**
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **[N/A]**

A Data details

A.1 Genuine data

We use the targets from the TESS Objects of Interest (TOI) catalogue. It consists of known exoplanets identified by previous missions and planetary candidates detected by TESS. At the time this work was conducted the TOI catalogue contained about 2600 targets. Using the *astroquery* package [21] we were able to retrieve data for 1668 of those targets by querying their TESS Input Catalogue (TIC) IDs in MAST [20]. We use the PDCSAP flux channel generated by the TESS pipeline which contains some data correction. Not a Number (NaN) values and statistical outliers (values higher than five standard deviations from the mean flux) are removed from the light curve. We phase-fold the light curve according to the estimated orbital period and extract the cadences within a time window of three transit durations centered around the transit midpoint. This is because some anomalous features might occur shortly before or after the transit and out-of-transit data can give some context about possible trends in the data. The parameter estimates for the period, epoch and transit duration are obtained from the respective data validation file if available or directly from the TOI table in NASA’s exoplanet archive [2] otherwise. The folded light curve segments are binned to 256 equally spaced cadences and data gaps are filled with linear interpolation if present. The values of each transit shape are scaled to be in a range between -1 and 1.

A.2 Fabricated data

Realistic transit data can be synthesized via theoretical transit models based on mathematical formalisms such as the one of Mandel and Agol [5]. Fabricated transits have the same format as the real data (vector size of 256 and range of values between -1 and 1).

Regular transits We use the *batman* library by Kreidberg [22] which requires the following parameters: planet-star radii ratio r , semi-major axis in stellar radii a , orbital inclination angle i , orbital eccentricity e , longitude of periastron $\bar{\omega}$ and the limb darkening coefficients. For our simulations we use the quadratic limb darkening law which has two coefficients c_1 and c_2 . Limb darkening coefficients are obtained via a lookup table for TESS photometry provided by [35]. This requires knowledge about the stellar surface temperature T , surface gravity $\log g$ and metallicity Z . A uniform distribution over those parameters might not reflect the general transit shapes observed in real data well. Hence, we sample constellations of confirmed exoplanets from a table with parameter estimates available at NASA’s exoplanet archive [2]. For some systems certain parameter estimates are missing. In these cases we uniformly sample them from a plausible range of values. An overview of parameter ranges for missing values is presented in Table 3. The radius ratio is obtained by

Table 3: Parameter ranges for missing planetary system parameters during artificial transit synthesis.

	R^*	R_p	a	b	e	$\bar{\omega}$
Minimum	0.3	0.3	2.3	0.0	0.0	0.0
Maximum	4.0	7.0	30.0	1.0	1.0	90.0
Unit	R_\odot	R_e	R^*	R^*	-	degrees

$r = \frac{R_p}{R^*}$, where R_p is the planetary radius and R^* is the stellar radius. The inclination is obtained by $i = \arccos\left(\frac{b}{a}\right) \frac{180}{\pi}$, where b is the impact parameter. T , $\log g$ and Z are determined sequentially in this order. The range of values for each of those parameters is constrained by the combination of values available in [35].

The theoretical transit models do not take into account photon noise or other sources of disturbances. To make the data more similar to real data we add Gaussian-distributed noise to every data point. For each modeled transit we uniformly sample a signal-to-noise ratio (SNR) in the range between 1 and 30. With SNR we refer to the ratio between the maximal transit depth of the signal divided by the standard deviation of the added Gaussian noise. The range of SNRs is inspired by the values observed in TESS data. Note that in TESS data the distribution over different SNRs is not uniform. There can occur trends in the data that are due to stellar variability (e.g due to rotational modulation through star spots). To simulate this effect we add a linear trend to half of the generated signals. The

slope of the trend can be positive or negative and its magnitude is maximally two times the transit depth within the observation window. The generation of anomalous transits is analogous, just using appropriate software designed for the anomaly class.

Transits with (occulted) star spots The software package `e11c` [23] supports the modelling of circular star spots. It requires further parameters for each spot: the longitude λ_s and latitude ϕ_s on the stellar surface, the radius R_s of the spot, the brightness factor b_s of the spot relative to the stellar surface and the rotation period p^* of the star.

Our transit models with star spot features contain up to four star spots with a higher probability for a lower number of star spots. Configurations with more than one spot are always half as likely as with one spot less. ϕ_s is biased to lie close to the latitude of the planet’s transit which is parametrized by the impact parameter b . Their latitude has an offset from the line of transit no larger than five degrees in either direction. This ensures that the transit models contain examples of occulted star spots but there will be also configurations where no spot is occulted and only the modulation by the stellar rotation is visible. In the test set, only examples that exhibit occulted spot features are selected.

λ_s is uniformly sampled in a range between -60 and 60 degrees from the center of the stellar disk. This is because further at the limb of the star the effect of occulted star spots can become extremely small due to interaction with the limb darkening effect. We uniformly sample b_s in the range between 0.7 and 1.3 modeling both dark and bright spots with a realistic magnitude. The sizes are sampled dependent on each other to reduce the probability of overlapping star spots and avoid too much of the surface being covered by star spots. The maximal spot radius size for the first spot is 20 degrees, which is reduced by five degrees for every further spot. Each spot radius is sampled uniformly in a range between two degrees and the individual maximal radius. p^* is taken from the list of parameters if available or uniformly sampled from a range between 10 and 40 days otherwise.

GD transits For modelling the gravity darkening effect on the transit shape we use `PyTransit` which offers an oblate star model. Additional parameters for this model are the gravity darkening coefficient β , the temperature at the poles T_{pole} , the stellar inclination i^* and the sky-projected obliquity λ . We obtain β from [35] analogously to the limb darkening coefficients. We uniformly sample the pole temperature from a range between 5700 and 12 000 Kelvin. The parameters i^* and λ are uniformly sampled from a range between zero and 90 degrees which covers all possible orientations that would result in a different effect (the remaining constellations are redundant due to symmetry).

Gravity darkening effects reported in the literature have usually magnitudes between $1e-3$ and $1e-4$ in a normalized light curve. For each configuration of i^* , λ and β we test up to ten different values for T_{pole} . If a configuration shows an effect larger than $1e-3$ or smaller than $1e-4$ it is skipped. After ten unsuccessful configurations the current stellar parameters are discarded and the next parameter set is tested.

DRE transits Accurate theoretical modeling of DREs is not as well understood as the other phenomena. This is why we decided to use real world examples to model DREs. Specifically, we use light curves of Kepler-1520 b, a known DRE, to create residual profiles from its transits and the best-fit theoretical model from the Kepler pipeline. Those residual profiles are added to regular theoretical transit models and the scale of the SNR is sampled in the same way as for regular transits. To increase the variety in DRE samples we add a small Gaussian noise with standard deviation $\sigma < 0.5res$, where res is the average magnitude of the residual profiles. This ensures diverse examples while not overshadowing the internal structure of the residual profile.

B Physics of anomalous transits

To give readers interested in the astronomical part of this work a better understanding of the anomaly classes defined for our artificial datasets, we shortly elaborate on the astrophysical source of the anomalous signal and possible indications for astronomy research. For regular transits the largest variability in shape comes from the *limb-darkening* effect and the *impact parameter* $b = \frac{a \cos(i)}{R_*}$, where a is the semi-major axis of the system, i is the orbital inclination angle with respect to our point-of-view and R_* is the stellar radius. The limb-darkening effect causes the star to appear brighter

at the center and dimmer at the edges (or limb). An impact parameter of $b = 0$ refers to a planet that transits along the equator of the star producing a U-shaped transit while a high impact parameter corresponds to a grazing planet that transits at the edge of the stellar disk and produces a shallower more V-shaped transit. The difference in shape is caused by the different ratios between ingress/egress (partial occlusion of planet at beginning/end of transit) duration and full occlusion duration. In the case of occulted star spots dark or bright areas on the stellar surface can temporarily change the amount of light that is blocked by the transiting planet resulting in "bumps" in the transit shape. Accurate modeling of occulted star spots can for example increase the precision of radius estimates for the planet. Gravity darkening happens when a star is rapidly rotating around its own axis. The centrifugal forces drag the poles closer to the center and push away mass at the equator making the star more ellipsoid shaped. This causes hotter, brighter poles and a cooler, dimmer "belt" at the equator. Depending on the angle of the planetary orbit with respect to the rotation axis of the star from our point of view (projected obliquity) this can cause asymmetries in the transit shape. Based on those asymmetries the projected obliquity can be determined. DREs drag a gas/dust tail behind them that blocks additional light. An exponential decay of the light blocked by the tail leads to a characteristic asymmetry of ingress and egress as well as a bump shortly before ingress that is produced by light reflected from particles. Follow-up studies of DREs can answer questions about their internal composition.

C Architecture details and hyperparameters

C.1 Architecture, hyperparameters and compute

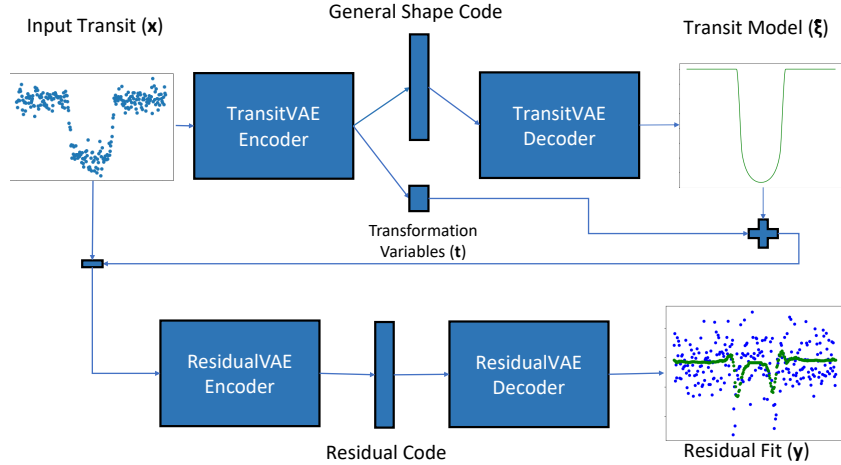


Figure 3: Step-wise feature extraction by TransitVAE and ResidualVAE. "+" represents the application of the deterministic transform f (see eq. 1) and "-" represents the calculation of the residuals.

The illustration of our approach is summarized in Figure 3. The detailed neural architecture of the VAE is summarized in Figure 4.

This architecture has about 527k trainable parameters. It is used for the TransitVAE, ResidualVAE and standard VAE. Note that the Transit VAE has a few more trainable parameters due to the additional output of the five transformation parameters, which is negligible. For the VAE-large the number of filters is changed from 16, 32 and 64 to 23, 46 and 92 respectively to yield a total of about 1.09M trainable parameters which is roughly the number of parameters of the TransitVAE and ResidualVAE combined.

The models are implemented in PyTorch [34] (version 1.7.1). We use the Adam optimizer with an initial learning rate of $5e^{-4}$ and PyTorch's default parameters for the weight updates. Training was executed on a single GPU (2GB RAM) with CUDA 10.0 of a modern notebook. This took about 40 minutes for 100 epochs of one of the ALT training sets. VAE-large was trained for 200 epochs which took about 80 minutes. We make use of a dynamical weighting between the reconstruction loss and the Kullback-Leibler divergence loss as described in [36].

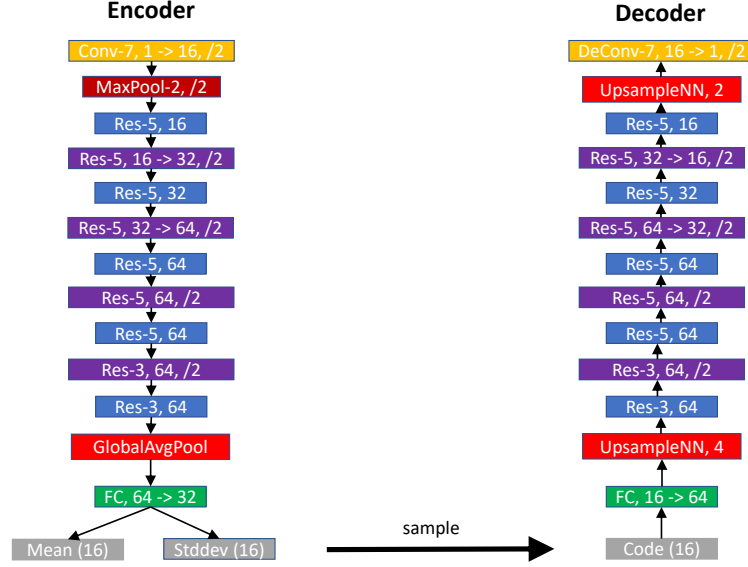


Figure 4: VAE Architecture. Conv refers to a 1D Concolutional layer, DeConv to a transposed 1D convolution. The first integer describes the kernel size and the second the number of filters. 1->16 refers to one input and 16 output channels, /2 refers to a stride of 2. All layers are followed by a batch normalization layer and a ReLU activation function. Res describes a Residual block with two Conv layers in the encoder and DeConv layers in the decoder. FC ia a fully connected layer.

C.2 Details on deterministic transform f

The deterministic transform is defined as:

$$f(\xi, t) = \text{trend}(\text{scale}_v(\text{shift}_h(\text{scale}_h(\xi, s_h), l_h), s_v), t_s, t_e) \text{ with } (s_h, l_h, s_v, t_s, t_e)^T = t. \quad (1)$$

Recall ξ is the predicted general transit (noise-free, centered and detrended) and t are the transformation parameters consisting of the horizontal scale s_h , the horizontal shift (or lag) l_h , the vertical scale s_v , the start value of the trend t_s and end of value of the trend t_e . First the horizontal scale s_h is applied according to $\text{scale}_h(\xi, s_h) \equiv \text{cut}(\text{resize}(\xi, s_h n))$, where resize corresponds to resizing with linear interpolation, n is the size of ξ and cut applies either center-cropping to size n if $s_h \geq 1$ or applies edge padding equally to each side up to size n if $s_h < 1$. The horizontal shift is applied via cyclic permutation $\text{shift}_h(\xi, l_h) \equiv$ for each $\xi_i \in \xi : i = \text{mod}(i + l_h, n)$, where $\text{mod}()$ is the modulo operation, i is an index of ξ and n is its size. The vertical scale is applied by $\text{scale}_v(\xi, s_v) \equiv s_v(\xi - \max(\xi)) + \max(\xi_h)$. A linear trend is applied by $\text{trend}(\xi, t_s, t_e) \equiv \xi + \text{interpl}(t_s, t_e, n)$, where interpl is a linear interpolation between the first two arguments of size n , which is the size of ξ . Since the scale_h and shift_h operations are non-differentiable s_h and l_h are learned via regression loss on their ground truth values (only available for fabricated data). Parameters s_v , t_s and t_e are learned via mean-squared error loss between ξ and input data x .

C.3 Overview of Losses

Transformation Loss:

$$L_{\text{transform}} = 2(MSE(\hat{l}_h, l_h) + MSE(\hat{s}_h, s_h)) + MSE(\hat{x}, x)$$

TransitVAE Loss:

$$\frac{MSE(\hat{\xi}, \xi)}{2\gamma^2} + \beta KL(\mu_T, \sigma_T) + \lambda L_{\text{transform}}$$

ResidualVAE Loss:

$$\frac{MSE(\hat{y}, y)}{2\gamma^2} + \beta KL(\mu_R, \sigma_R)$$

For an explanation of the variables see section 3 and the preceeding section about the deterministic transform f . MSE is the mean squared error and $KL(\cdot)$ refers to the analytical form of the Kullback-Leibler Divergence between a Gaussian with diagonal covariance matrix and a standard normal distribution. The dynamical recunstruction term weight γ is a monotonically decreasing estimate of the current reconstruction loss as defined in [36]. β is a fixed weight of the KL term that was tuned for each model. For the VAE-large model we used $\beta = 1$, for all others $\beta = 0.1$. λ is a fixed loss weight for the transformation loss, which should be high in order to achieve quick learning of the semantics of shift, scale and trend. The performance of the models is not very sensitive towards this parameter, we chose $\lambda = 64$.

C.4 Anomaly detection details

For the unsupervised anomaly detection methods we use `scikit-learn`'s [26] implementation of the classifiers with default settings. For the LOF we follow the recommendation from the original paper to determine the final score of each example by taking the maximum over several evaluations of the LOF with different values for k , where k defines the number of nearest neighbours that are used to estimate the score. We evaluate the LOF for $k \in \{15, 20, 25, 35, 50\}$ which are samples within the range recommended by the authors.

D Additional results

Table 4 shows the full AP results on the ALT-i dataset for each classifier on different featurizations. Reported values correspond to the mean and standard deviation over five independent runs with different random seeds. The ResidualVAE features perform the best across all classifiers. The Mahalanobis distance with ResidualVAE features achieves the overall best score closely followed by the LOF with ResidualVAE features.

Table 4: Average precision (in %) and standard deviation over 5 independent runs for various classifiers with different featurizations on the ALT-i test set. Each classifier's best result is bolded.

	Local Outlier Factor	One-class SVM	Isolation forrest	Mahalanobis distance
Raw data	22.22 +/- 0	10.40 +/- 0	12.01 +/- 0	-
VAE	60.95 +/- 1.29	28.35 +/- 0.60	22.19 +/- 2.41	25.18 +/- 2.43
VAE-large	50.04 +/- 10.85	20.73 +/- 1.86	37.88 +/- 3.87	34.15 +/- 8.44
ResidualVAE	84.03 +/- 1.97	78.76 +/- 4.29	59.76 +/- 4.23	85.12 +/- 1.41