

MonoTools – A python package for planets of uncertain period

Hugh P. Osborn^{1, 2}

¹ NCCR/Planet S, Centre for Space and Habitability, University of Bern, Switzerland ² Kavli Institute for Space Sciences, Massachusetts Institute of Technology, Cambridge, MA, USA

DOI:

Software

- [Review](#) 
- [Repository](#) 
- [Archive](#) 

Submitted:

Published:

License

Authors of papers retain copyright and release the work under a Creative Commons Attribution 4.0 International License ([CC BY 4.0](#)).

Summary

The transit method has proved the most productive technique for detecting extrasolar planets, especially since the era of space-based photometric survey missions began with *CoRoT* (Auvergne et al., 2009) and *Kepler* (Borucki et al., 2010) in the late 2000s. This continued with *K2* (Howell et al., 2014) and *TESS* (Ricker et al., 2014), and will extend into the 2030s with *PLATO* (Rauer et al., 2014). Typically, the planets detected by these surveys show multiple consecutive transits. This means planet candidates are most often detected through algorithms which search the frequency domain [e.g.; Kovács et al. (2002); Hippke & Heller (2019)], vetted using metrics that require multiple detected transits [e.g.; Thompson et al. (2018); Shallue & Vanderburg (2018)], and modelled (and sometimes statistically validated) using the assumption that the orbital period is well-constrained and approximated by a Gaussian distribution [e.g.; Eastman et al. (2013); Morton (2012)]. However, planet candidates continue to be found that do not show multiple consecutive transits - the single (or “Mono-”) transits [e.g.; Wang et al. (2015); Osborn et al. (2016); Gill et al. (2020)]. For these transit candidates - where orbital period is not a priori known from the detection - a special approach to exoplanet detection, vetting and modelling must be taken.

In this work, we detail `MonoTools`, a python package capable of performing detection, vetting and modelling of Mono (and Duo) transit candidates. First we will describe briefly what Mono (and Duo-) transits are, and the challenges associated with them. Then in the following three sections we will outline the basis of the three parts of the code. Following that, we will validate the code using limited examples of planets with known orbital periods.

Mono- & Duo-transits

Mono-transits, which have also variously been called “single transits” or “orphan transits,” are the transits of long-period planet candidates which occur only once during photometric observations. In these cases, the orbital period is not directly evident as we do not have subsequent transits. However, the planet’s orbit can be constrained using the transit event, as we will explore later in this section.

Another special case is worth noting - that of two non-consecutive transits where intermediate transit events were not observed, therefore the orbital period is not directly constrained by the transit events. Here I class these cases as “Duotransits” in contrast to “Monotransits” and “Multitransits” which we will use as short-hand for planet candidates which show multiple (and consecutive) transit events. In these cases, the resulting planet candidate may have both a highly uncertain period ($20d < P < 750d$ in the case of two transits separated by a 2-year gap) and yet a well-constrained array of possible periods to search ($P \in (t_{tr,2} - t_{tr,1})/\{1, 2, 3, \dots, N_{max}\}$).

`MonoTools` is explicitly dealt to deal with both the monotransit and duotransit cases.

Transit shape is universally governed by the same simple geometry Seager & Mallen-Ornelas (2003). As they must be strongly detected in a single transit, the typical per-transit signal-to-noise of monotransits is often higher than for multitransits, allowing their shape to be well-constrained. This shape is important for detection, vetting and modelling of such planet candidates. Transit duration is weakly dependent on planetary period ($t_D \propto P^{1/3}$), therefore long-period planets typically have longer-duration transits. Indeed the longest-duration transit yet found belonged to a monotransit detected in K2 (Giles et al., 2018) at 54 hours.

Input Information

Detrended Lightcurve

We built a custom `MonoTools.lightcurve` package to manipulate photometric lightcurves for this package. This includes the ability to download all available Kepler, K2, CoRoT and TESS lightcurves for any target.

Kepler

For stars in or near the Kepler field, we use `astroquery` to query the Kepler input catalogue (KIC) to assess if the star was observed. The Kepler lightcurves (either 1 or 30-min cadence, depending on availability) were accessed on MAST and the PDCSAP flux (pre-datasearch conditioning simple aperture photometry) was used as the default flux. We masked points where the `quality` bits [1,2,3,4,6,7,8,9,13,15,16,17] were flagged.

K2

Like for Kepler, we check for any star near the ecliptic if the star has an EPIC (Ecliptic plane input catalogue, (Huber2014?)) ID using `astroquery` Unlike Kepler, K2 had a diversity of pipelines used to produce photometry. `MonoTools.lightcurve` has the capability to access data from Everest (luger2016?), K2SPP (Vanderburg2015?), and PDC []. Unless specified `MonoTools.lightcurve` will search in this order, which follows typical lightcurve precision, until data is found for a given EPIC.

CoRoT

The CoRoT object search API available at the NASA Exoplanet Archive is used to both search for and then download CoRoT data. Although three band photometry is available, we are typically most interested in the more precise monochrome lightcurve, so this is by default the flux parameter. As CoRoT observed continuously from its low-earth orbit, the South Atlantic Anomaly produced clear flux bumps in the data due to excess cosmic rays, which needs to be removed from the data. To do this, each flux point is compared with its 24 neighbours and any point identified to be significantly (at 3.2σ) higher than its neighbours median flux in 80% of possible bins is added to a mask. This is iterated (typically twice) to make sure an accurate standard deviation without the highest anomalies with lower but still-significant SNR to be removed.

TESS

Given an RA/Dec, we search the TIC (TESS Input Catalogue (Stassun2018?)) to find the TESS ID for a given target. As for K2, there is not one unique pipeline for TESS data, especially for those targets not observed in 2-minute TPFs but only in the FFIs. In this case, `MonoTools.lightcurve` will search MAST for a PDC (20s or 120s)

lightcurve (Jenkins?), a SPOC-TESS (10 or 30min) lightcurve, a QLP (Quick-Look Pipeline, (Huang?)), and finally an Eleanor lightcurve (Feinstein2019?).

Stellar parameters

Search

`MonoTools.search.MonoTransitSearch`

This function iteratively fits both a transit model and a polynomial to the lightcurve to detect monotransits in space telescope photometry, which we detail here.

We first create a series of reference transit models (default 5) to iterate across the lightcurve using `exoplanet` (Foreman-Mackey et al., 2021). The derived stellar parameters are used, along with a default planet-to-star radius ratio of 0.1. As input periods, logspaced values between 0.4 and 2.5 times the duration of continuous observations (in the case of lightcurves with gaps longer than 5 days, the longest individual region was used). The impact parameters were chosen such that the maximum duration transit (with $P = 2.5P_{\text{mono}}$) is given $b = 0.0$ while successively shorter durations linearly spaced up to $b = 0.85$ producing ever-shorter duration transits. 500 in-transit steps are generated for each model with exposure times fixed to that of the lightcurve, and then interpolated. This interpolated transit function forms the model which is minimized at each step in the lightcurve.

Each of the models (with differing transit durations) are then iterated over the lightcurve, where transit centres are shifted some small fraction of transit duration each iteration (default 5%). At each position, a 7-transit-duration-long window around the transit time is fitted to three different models which are minimised using `scipy.optimize`. These models are: - The interpolated transit model with varying depth (reparameterised to log depth to avoid negative depths) plus a 1D gradient in the out-of-transit flux. - A 3-degree polynomial. - A “wavelet” model with the following equation, designed to fit dips due to stellar variability where t_D is the transit duration (set, in our case, from the interpolated transit models), and a is the depth. As with the transit, a gradient was also included to account for any non-linear out-of-eclipse flux trend.

$$t' = 2\pi x / (2t_D); F = a(\exp((-t'^2)/(2\pi^2)) \sin(t' - \pi/2))$$

For each of these three models, the minimised log likelihood is used to compute a Bayesian Information Criterion. Significant detections are therefore found by choosing all transit model fits which have a log likelihood ratio with respect to non-transit models greater than some threshold (default: 4.5) as well as an SNR (calculated from the depth, transit duration, and out-of-transit RMS) greater than some SNR threshold (default: 6.0).

Multiple iterations (either in transit time or duration) may find the same significant dip. In this case the minimum DeltaBIC between transit & polynomial model is used to choose the representative detection, and all nearby detections within $0.66t_D$ of this candidate are removed to avoid double counting.

`MonoTools.search.PeriodicPlanetSearch`

Many multitransiting planets produce high-SNR individual transits that would be detected using `MonoTransitSearch`, therefore we also require a method of detecting periodic planets, as well as selecting between the monotransit and multitransit cases.

To search for periodic transits, we first flatten long-timescale variation from the lightcurve. This is performed by fitting polynomials to sections of the lightcurve while also iteratively removing anomalies, as was adapted from (Armstrong et al., 2014). For each small step

along the lightcurve, a wide window around (but not including) each step is used to fit a polynomial. Points in this window that had already been identified as either outliers (i.e. from detrending) or within detected monotransits (from the Monotransit search), can be excluded from the polynomial fitting. A log likelihood is computed on each of ten iterated polynomial fits, and each time a new pseudo-random mask is generated by excluding points whose scaled residual to the model is greater than a randomly-generated absolute normal distribution with unit standard deviation (thereby, on average, excluding points with offset residuals). This best-fit polynomial, is then subtracted from the small central masked region. For Periodic Planet searches, a window with duration 11 times the likely maximum duration and a stepsize of 0.1 days are typically used to ensure transits do not influence the polynomial fit.

`transit least squares` [TLS; Hippke & Heller (2019)] is used to perform the periodic planet search. We iteratively run this TLS search and masked the detected transits until no more candidates are found above the SNR threshold (default:6)

During the TLS search, we necessitated a minimum of three transits. This is preferred over a limit of two for a handful reasons: - The implementation of period-epoch values in `transit least squares` means that allowing two transits also lets monotransits be detected, thereby duplicating our effort with the above search technique. - Multi-transit search is not strict about assigning only similar dips together and may connect either two monotransits, or the wrong two transits from a multi-transiting planet. Requiring three dips ensures the correct periodicity - Individual transits of the majority of good duo-transiting planet are likely to be individually detectable on their own right, as the individual transits have SNR's only $1/\sqrt{2}$ (30%) lower than the combination of both events. To make sure that at least 3 transits were detected, we excluding any candidates where one or two individual transits dominated the combined SNR (defined by computing an expected SNR from the sum of each individual transit SNRs and assuring solid detections have $\text{SNR}_i > 0.5\text{SNR}_{\text{expected}}$). If the highest periodogram peak in the TLS corresponds to a multi-transiting planet with a SNR higher than our threshold (default: 6), and

In either case, if a signal with SNR higher that the threshold is found, we mask the detected transits by replacing all points associated with the transit with flux values randomly taken from the rest of the lightcurve. The lightcurve is then re-scanned with TLS until no high-SNR candidates remain.

Vetting

Fitting

Typical Monotransit fitting approaches

We have the following information available from a monotransit:

- Epoch of transit, t_0
- Transit duration, t_D
- Ingress & Egress duration τ
- Transit depth, δ
- In-transit shape
- Stellar parameters (e.g. stellar radius and density)
- Orbital period information from the lack of additional transits in the lightcurve. At the least we have a minimum possible period below which obvious transits would be observed, and at the most we may have a complex sequence of period islands.

- Additional planets
- Complimentary observations (e.g. radial velocities)

From these observables, there are then second order parameters. These can either be derived from the observables or, more commonly, can be used directly in fitting as reparameterisations of the observed parameters:

- **Limb-darkening parameters** - These parameters due to the change in optical depth as a function of position on the stellar surface correspond to the in-transit shape and are also constrainable from the stellar parameters (as theoretical limb-darkening parameters can be calculated for a given star).
- **Radius ratio**, R_p/R_s - This is most directly linked to transit depth δ ($R_p/R_s \sim \sqrt{\delta}$), although limb-darkening and dilution can play effects here (as well as impact parameter in the case of a grazing transit/eclipse).
- **Impact parameter**, b - Impact parameter refers to the location of the transit chord between the centre and edge of the stellar disc. In the case of multitransiting planets impact parameter constraints come from both the transit shape and the known orbital distance compared with the transit duration. With monotransits we do not have this luxury and instead only the transit shape constrains b (i.e. the radius ratio, ingress duration, transit duration).

These parameters can then in turn be linked to orbital parameters. Typical transit modelling includes parameters for both transit shape (e.g. impact parameter, radius ratio, & limb-darkening parameters), semi-major axis (typically parameterised as a/R_s), and orbital period. Splitting orbital parameters into both a/R_s & P is superfluous for planets with uncertain periods.

Instead, the typical approach is to use only the transit shape parameters to constrain as few orbital parameters as possible. For example, if the impact parameter can be constrained from the shape alone, then in combination with the transit duration we can estimate the velocity of a planet across the star. In the case of a purely circular orbit, this velocity then directly produces a period. Including samples from some eccentricity and omega (argument of periasteron) distributions, these will then modify the resulting period.

There have been numerous past efforts and theoretical works exploring fitting such transits:

- Yee & Gaudi (2008) provided a theoretical perspective on modelling such transits even before Kepler began finding them.
- Wang et al. (2015) adapted a transit model which included both circular period and semi-major axis (a/R_s) without specific priors on these quantities.
- Foreman-Mackey et al. (2016) included eccentricity and reparameterised the orbital semi-major axis & inclination into two parameters ($\sqrt{a} \sin i$ & $\sqrt{a} \cos i$), with an effective prior on the period ($P^{-2/3}$).
- Osborn et al. (2016) fitted impact parameter and a scaled velocity parameter (which encapsulates a prior equating to $P^{-5/3}$) to predict planetary periods, with the same approach being used in Giles et al. (2018).
- D. Kipping (2018) provided a purely theoretical view of the correct prior to place on such analyses, combining the geometric transit probability, a window effect prior, and the intrinsic prior to produce a value of $P^{-8/3}$.
- Sandford et al. (2019) created the `single` python package which used gaia parallaxes as a source of stellar density and allowed eccentricity to vary (with a period prior of $P^{-5/3}$)

- Becker et al. (2018) modelled the duotransit system HIP41378 using discrete period aliases and a P^{-1} prior.

As can be seen from this array, the approach and prior varies widely between study. Some directly model orbital period while others reparameterise in terms of parameters closer to the observed transit information. Some use eccentricity but most assume circular orbits. Some use information from interior multitransiting planets (e.g. Becker et al. (2018)) but most treat only the outer planet individually.

MonoTools.fit approach

The `monoModel` class of `MonoTools.fit` uses the `exoplanet` package (Foreman-Mackey et al., 2021) and `PyMC3` (Salvatier et al., 2016) to build flexible a transit model which can be easily and efficiently sampled using `PyMC3`'s Hamiltonian Monte Carlo approach.

The key development of `MonoTools` over past monotransit and duotransit tools is that it natively supports bayesian marginalisation over discontinuous period space. In the case of duotransits, this means the multiple period aliases, while in the case of monotransits, this means the multiple period gaps that can occur due to non-continuous photometric coverage.

Calculating Period Aliases & Gaps

For Duotransits, period is not a modelled quantity in `MonoTools.fit`, but is instead derived from modelling two transit centres t_0 and t_1 , with the period being part of the set $P \in (t_{tr,2} - t_{tr,1}) / \{1, 2, \dots, N\}$. Potential aliases therefore lie between $P_{max} = t_{tr,2} - t_{tr,1}$ and P_{min} , a minimum period, and are calculated by `compute_duo_period_aliases`. To calculate P_{min} , this function iterates over all potential aliases between P_{max} and 10d. For each period, the data is phase-folded and the known transits masked. Only period aliases for which there are no significant in-transit observations found elsewhere (defined as 15% of the central 90% of the transit duration) are kept in the model.

For monotransiters, a similar process is applied to find regions of the period parameter space that are rejected by photometry using `compute_period_gaps`. First, an RMS timeseries of the light curve is computed. This iterates through the flattened light curve in steps that are typically $1/7t_D$ wide, performing a weighted average & standard deviation for photometry in a $1t_D$ wide. The resulting timeseries can be converted into a theoretical transit SNR given the depth of the known transit. This timeseries can be converted to a function of period space (i.e. by phase-folded around the know transit), with regions without photometric data being given SNR values of 0.0. Period gaps can then be defined as regions in period space where the computed SNR is below some threshold value (default: 4σ).

Marginalisation

Here we have some number of discrete regions in one parameter space that we want to sample. Typically, samplers such as MCMC fail with multiple local minima, especially in the case where the gaps between regions are far wider than the regions themselves. One way to avoid this problem is to treat each region of this discontinuous parameter space as separate and therefore sample each one individually. We can then perform marginalisation over N submodels with parameters that correspond to each N period gaps. By computing the log likelihood with respect to the data and the log prior of the parameters used, their sum gives us the probability of each submodel for a given step.

$$p(\theta | y) = \sum_{k=1}^K p(\theta | y, M_{P=i}) p(M_{P=i} | y)$$

The normalised probability for each period gap or alias are then the marginalised probabilities, and the marginalised parameters are simply the average of the submodel pa-

parameters weighted by this probability. However, if all of the parameters in the model are marginalised, this can effectively require a huge number of parameters - $N_{params} \times N_{models}$. Therefore, to improve efficiency, we must choose which parameters to marginalise and which to fit only once.

In the case of a transit where we want to marginalise over multiple sub-models at different orbital periods, we only need marginalise over parameters that substantially vary as a function of orbital period. Other parameters, such as transit time, limb darkening and radius ratio, can be fitted as global parameters.

In the simplest case, `MonoTools` allows some degree of flexibility in what parameters to marginalise using the `fit_params` and `marginal_params` lists as inputs to the `monoModel` class. Period is always marginalised, but so can t_D or b .

However, this implementation of marginalisation can still be slow, and suffers from drawbacks. For t_D and b one must always be globally fitted and the other marginalised. But, their connection to the orbital period means that across this marginalisation there is always going to be many aliases which do not well represent the data. For example, if a 15d planet with $b = 0.2$ fits the transit well, a 150d planet with the same transit chord is sure to produce far too long a transit duration, and therefore a very low likelihood. And, despite the fact a 150d planet might be able to well explain the data at higher impact parameters, the strong prior on period means this part of the parameter space is not explored and our 150d alias may be given an artificially low marginal probability.

Marginalising with derived in-transit velocity

The solution to this problem is to not marginalise duration or impact parameter which are both intimately connected to the observed transit shape. By keeping all the parameters required to fit transit shape global, we can remove the need to perform likelihood calculations for each of the different period parameters, greatly improving speed and sampling efficiency. Instead, we use the duration and impact parameter to derive an instantaneous velocity across the star, as was performed in Osborn et al. (2016). For each of the period aliases and the sampled stellar parameters, we can calculate a circular velocity. The derived transit velocity as a ratio of circular velocity (v/v_{circ}) for each period alias/gap then becomes the important quantity to marginalise. Of course this is incompatible with the assumption of a circular orbit - we require an eccentricity distribution for this method to work.

As we are directly modelling the transit shape the likelihood for each alias is identical (or at least negligibly different), all that is important is deriving a prior for each. The key part of this prior comes from the assumed eccentricity distribution. Observations of exoplanets show that low eccentricities are typically preferred over high ones. Two distributions are typically used to quantify this - the beta distribution of D. M. Kipping (2013) for typically single-planet RV systems, and the Rayleigh distribution of Van Eylen & Albrecht (2015) for multi-planet transiting systems.

Another observational constraint on eccentricity comes from the distribution of perihelion distances - exoplanet orbits do not typically pass within $2R_s$, as within this radius tidal circularisation occurs. In terms of semi-major axis, we include a sharp sigmoid prior at the threshold of $e = 1 - 2R_s/a$ which corresponds to this perihelion limit. We can also include another upper limit on eccentricity here - stable exoplanetary systems require that a planet's orbit does not cross the orbit of interior candidates. So in the case of transiting multi-planet systems we can use $e < 1 - R_s/a_{inner}$.

For each given v/v_{circ} we must calculate the possible eccentricity and argument of periastron. From Barnes (2007) (Eq 12) we know that a planet's azimuthal velocity can be defined as:

$$\frac{v_f}{v_{circ}} = \frac{1+e \cos f}{\sqrt{1-e^2}} \text{ where } f_{tr} = (\omega - \pi/2).$$

Rearranging to give eccentricity gives two roots, although the second root is only applicable for cases where $v/v_{\text{circ}} < 1.0$:

$$e_1 = (-v^2 \sqrt{\frac{v^2(\sin^2 \omega + v^2 - 1)}{(\sin^2 \omega + v^2)^2}} - \sin \omega^2 \sqrt{\frac{v^2(\sin^2 \omega + v^2 - 1)}{(\sin^2 \omega + v^2)^2}} - \sin \omega) / (\sin^2 \omega + v^2)$$

$$e_2 = (v^2 \sqrt{\frac{v^2(\sin^2 \omega + v^2 - 1)}{(\sin^2 \omega + v^2)^2}} + \sin \omega^2 \sqrt{\frac{v^2(\sin^2 \omega + v^2 - 1)}{(\sin^2 \omega + v^2)^2}} - \sin \omega) / (\sin^2 \omega + v^2)$$

These two roots make it impractical to solve for the probability of v analytically, so we instead compute this numerically. Ultimately, we must derive the probability of each velocity (v/v_{circ} , or v hereafter) given that a transit occurs by marginalising over all compatible eccentricities and arguments of periastron:

$$p(v | \text{Tr}, e_{\text{max}}) = \frac{\int_0^{2\pi} \int_0^{e_{\text{max}}} p(e, \omega | v) p(\text{Tr} | e, \omega, v) de d\omega}{\int_{v_{\text{min}}}^{v_{\text{max}}} \int_0^{2\pi} \int_0^{e_{\text{max}}} p(e, \omega | v) p(\text{Tr} | e, \omega, v) de d\omega dv}$$

Using the equations for e , we can feasibly generate eccentricity for each v/v_{circ} & ω sample. As the geometric probability of transit is a function of the distance in-transit, and eccentricity & argument of periastron directly affect this quantity, we also calculate a geometric correction (i.e. the distance at transit compared to semi major axis):

$$\frac{d_{\text{Tr}}}{a} = \frac{1 + e \sin \omega}{1 - e^2}$$

Therefore the probability of each grid position is then determined by the probability derived from selected prior distribution (i.e. 'kipping', 'vaneylen' or 'uniform') multiplied by the geometric correction. In the case that the derived eccentricity is above e_{max} , a log prior of -500 is added.

As all velocities here are normalised to circular velocities and the joint argument of periastron – eccentricity distributions remain constant with period, these calculations should remain constant for any period across all model samples. However, the maximum permitted eccentricity (e_{max}) can vary for each sample due to e.g. the sampled stellar radius and parameters for the orbits of interior planets. Therefore, we need a way to compute on-the-fly a prior probability for a particular velocity and e_{max} , as well as a marginal eccentricity and argument of periastron. We choose to generate a 2D interpolation function for each eccentricity prior distribution.

Effectively the equation required to produce the marginalised probability distribution for v (given some maximum eccentricity and the fact that a transit occurs) is:

$$p(v | \text{Tr}, e_{\text{max}}) = \int_0^{2\pi} \int_0^{e_{\text{max}}} p(e, \omega | v, e_{\text{max}}) p(\text{Tr} | e, \omega, v) de d\omega$$

Where, for example in the case of the D. M. Kipping (2013) β distribution where $\alpha = 0.867$ and $\beta = 3.03$, the probability on $p(e | v, e_{\text{max}})$ (and, therefore, $p(e, \omega | v, e_{\text{max}})$ as ω is uniform) is:

$$p(e | v, e_{\text{max}}) = \begin{cases} 0 & \text{if } e > e_{\text{max}} \\ \frac{\exp(\alpha-1)(1-e)^{\beta-1}}{B(\alpha, \beta)} & \text{otherwise.} \end{cases}$$

By generating a grid of v/v_{circ} (16000-steps flat in \log_{10} between 0.07 and 14), ω (8000 steps flat between 0 & 2π) and e_{max} (96 steps sampled using $e_{\text{max}} \in 1 - 10^{\{-3, \dots, -0.05\}}$), we can derive eccentricities for each point in the $\omega - v$ plane and therefore compute marginalised probabilities for each point on the $v - e_{\text{max}}$ plane. For each of the e_{max} steps, the sum of probabilities for v must sum to 1.0, therefore we must renormalise the above

equation using the integral over all possible velocities using the following normalisation factor:

$$\int_{v_{\min}}^{v_{\max}} \int_0^{2\pi} \int_{1 \times 10^{-4}}^{e_{\max}} p(e, \omega | v) p(\text{Tr} | e, \omega, \log v) de d\omega dv$$

The resulting $v - e_{\max}$ distributions can be seen in Figure XX.

Choice of transit shape parameters

Transit duration and impact parameter can both

Treatment of limb darkening parameters

Limb-darkening parameters define the relative variation in surface brightness of a star from centre to limb, and therefore govern the shape of the transit between ingress and egress. For high-SNR transits where other parameters including orbital period are well-constrained, it is typically preferred to fit for limb-darkening parameters directly from the transit. However, for analyses where we wish to use the transit shape to constrain other parameters, it instead makes sense to constrain the limb-darkening using priors derived from theoretical predictions. For this reason, in the default case, `MonoTools.fit` constrains the limb darkening parameters using the derived stellar parameters (although it is also possible to fit unbiased).

Tables of theoretical limb darkening parameters typically produce grids of each parameter with respect to stellar effective temperature, $\log g$, metallicity and micro-turbulence. We select the nearest unique metallicity value (default of $[\text{Fe}/\text{H}]=0.0$) and micro-turbulence (default 1km/s), and perform a 2D interpolation of the Teff-logg values for each of the two quadratic limb darkening parameters. We then generate a sample of stellar Teff & logg values from the stellar parameters using a minimum uncertainty of 100K and 0.1dex to allow for potential systematic errors in stellar parameters. The interpolation functions then produce samples for each of the two quadratic parameters, from which we construct a normally-distributed prior.

In both cases, the Kipping reparameterisation of limb darkening ([kipping2013?](#)) is used to allow for efficient & physically-motivated sampling.

Treatment of period gaps & aliases

Eccentricity distribution marginalisation

Other photometric fitting parameters

Gaussian processes

Jitter

Dilution may be important, especially when unresolved stellar companions are detected through high-resolution imaging. To allow for this, we include the ability to include either constrained or unconstrained dilution from a stellar companion. In the case of unconstrained dilution, we allow the magnitude difference of the companion to vary from -10 to 10, while in the constrained option the user can define mean and standard deviations for each observed band. This is converted from Δ_{mag} to a correction factor for the undiluted flux to the total flux, $F_{\text{targ}}/F_{\text{total}} = 2.511^{-\Delta_{\text{mag}}}$.

Radial Velocity modelling

Derived semi-amplitude

Validation

Installation

Acknowledgements

References

- Armstrong, D., Osborn, H., Brown, D., Faedi, F., Gómez Maqueo Chew, Y., Martin, D., Pollacco, D., & Udry, S. (2014). On the abundance of circumbinary planets. *Monthly Notices of the Royal Astronomical Society*, *444*(2), 1873–1883.
- Auvergne, M., Bodin, P., Boissard, L., Buey, J.-T., Chaintreuil, S., Epstein, G., Jouret, M., Lam-Trong, T., Levacher, P., Magnan, A., & others. (2009). The CoRoT satellite in flight: Description and performance. *Astronomy & Astrophysics*, *506*(1), 411–424.
- Barnes, J. W. (2007). Effects of orbital eccentricity on extrasolar planet transit detectability and light curves. *Publications of the Astronomical Society of the Pacific*, *119*(859), 986.
- Becker, J. C., Vanderburg, A., Rodriguez, J. E., Omohundro, M., Adams, F. C., Stassun, K. G., Yao, X., Hartman, J., Pepper, J., Bakos, G., & others. (2018). A discrete set of possible transit ephemerides for two long-period gas giants orbiting hip 41378. *The Astronomical Journal*, *157*(1), 19.
- Borucki, W. J., Koch, D., Basri, G., Batalha, N., Brown, T., Caldwell, D., Caldwell, J., Christensen-Dalsgaard, J., Cochran, W. D., DeVore, E., & others. (2010). Kepler planet-detection mission: Introduction and first results. *Science*, *327*(5968), 977–980.
- Eastman, J., Gaudi, B. S., & Agol, E. (2013). EXOFAST: A fast exoplanetary fitting suite in IDL. *Publications of the Astronomical Society of the Pacific*, *125*(923), 83.
- Foreman-Mackey, D., Luger, R., Agol, E., Barclay, T., Bouma, L. G., Brandt, T. D., Czekala, I., David, T. J., Dong, J., Gilbert, E. A., & others. (2021). Exoplanet: Gradient-based probabilistic inference for exoplanet data & other astronomical time series. *arXiv Preprint arXiv:2105.01994*.
- Foreman-Mackey, D., Morton, T. D., Hogg, D. W., Agol, E., & Schölkopf, B. (2016). The population of long-period transiting exoplanets. *The Astronomical Journal*, *152*(6), 206.
- Giles, H., Osborn, H., Blanco-Cuaresma, S., Lovis, C., Bayliss, D., Eggenberger, P., Cameron, A. C., Kristiansen, M., Turner, O., Bouchy, F., & others. (2018). Transiting planet candidate from K2 with the longest period. *Astronomy & Astrophysics*, *615*, L13.
- Gill, S., Wheatley, P. J., Cooke, B. F., Jordán, A., Nielsen, L. D., Bayliss, D., Anderson, D. R., Vines, J. I., Lendl, M., Acton, J. S., & others. (2020). NGTS-11 b (TOI-1847 b): A transiting warm saturn recovered from a TESS single-transit event. *The Astrophysical Journal Letters*, *898*(1), L11.
- Hippke, M., & Heller, R. (2019). Optimized transit detection algorithm to search for periodic transits of small planets. *Astronomy & Astrophysics*, *623*, A39.

- Howell, S. B., Sobeck, C., Haas, M., Still, M., Barclay, T., Mullally, F., Troeltzsch, J., Aigrain, S., Bryson, S. T., Caldwell, D., & others. (2014). The K2 mission: Characterization and early results. *Publications of the Astronomical Society of the Pacific*, 126(938), 398.
- Kipping, D. (2018). The orbital period prior for single transits. *Research Notes of the AAS*, 2(4), 223.
- Kipping, D. M. (2013). Parametrizing the exoplanet eccentricity distribution with the beta distribution. *Monthly Notices of the Royal Astronomical Society: Letters*, 434(1), L51–L55.
- Kovács, G., Zucker, S., & Mazeh, T. (2002). A box-fitting algorithm in the search for periodic transits. *Astronomy & Astrophysics*, 391(1), 369–377.
- Mandel, K., & Agol, E. (2002). Analytic light curves for planetary transit searches. *The Astrophysical Journal Letters*, 580(2), L171.
- Morton, T. D. (2012). An efficient automated validation procedure for exoplanet transit candidates. *The Astrophysical Journal*, 761(1), 6.
- Osborn, H., Armstrong, D., Brown, D., McCormac, J., Doyle, A., Loudon, T., Kirk, J., Spake, J., Lam, K., Walker, S., & others. (2016). Single transit candidates from K2: Detection and period estimation. *Monthly Notices of the Royal Astronomical Society*, 457(3), 2273–2286.
- Rauer, H., Catala, C., Aerts, C., Appourchaux, T., Benz, W., Brandeker, A., Christensen-Dalsgaard, J., Deleuil, M., Gizon, L., Goupil, M.-J., & others. (2014). The PLATO 2.0 mission. *Experimental Astronomy*, 38(1-2), 249–330.
- Ricker, G. R., Winn, J. N., Vanderspek, R., Latham, D. W., Bakos, G. Á., Bean, J. L., Berta-Thompson, Z. K., Brown, T. M., Buchhave, L., Butler, N. R., & others. (2014). Transiting exoplanet survey satellite. *Journal of Astronomical Telescopes, Instruments, and Systems*, 1(1), 014003.
- Salvatier, J., Wiecki, T. V., & Fonnesbeck, C. (2016). Probabilistic programming in python using PyMC3. *PeerJ Computer Science*, 2, e55.
- Sandford, E., Espinoza, N., Brahm, R., & Jordán, A. (2019). Estimation of singly transiting K2 planet periods with gaia parallaxes. *Monthly Notices of the Royal Astronomical Society*, 489(3), 3149–3161.
- Seager, S., & Mallen-Ornelas, G. (2003). A unique solution of planet and star parameters from an extrasolar planet transit light curve. *The Astrophysical Journal*, 585(2), 1038.
- Shallue, C. J., & Vanderburg, A. (2018). 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.
- Thompson, S. E., Coughlin, J. L., Hoffman, K., Mullally, F., Christiansen, J. L., Burke, C. J., Bryson, S., Batalha, N., Haas, M. R., Catanzarite, J., & others. (2018). Planetary candidates observed by kepler. VIII. A fully automated catalog with measured completeness and reliability based on data release 25. *The Astrophysical Journal Supplement Series*, 235(2), 38.
- Van Eylen, V., & Albrecht, S. (2015). Eccentricity from transit photometry: Small planets in kepler multi-planet systems have low eccentricities. *The Astrophysical Journal*, 808(2), 126.
- Wang, J., Fischer, D. A., Barclay, T., Picard, A., Ma, B., Bowler, B. P., Schmitt, J. R., Boyajian, T. S., Jek, K. J., LaCourse, D., & others. (2015). Planet hunters. VIII.

Characterization of 41 long-period exoplanet candidates from kepler archival data.
The Astrophysical Journal, 815(2), 127.

Yee, J. C., & Gaudi, B. S. (2008). Characterizing long-period transiting planets observed by kepler. *The Astrophysical Journal*, 688(1), 616.