



orbitize!: A Comprehensive Orbit-fitting Software Package for the High-contrast Imaging Community

Sarah Blunt^{1,11} , Jason J. Wang^{1,12} , Isabel Angelo² , Henry Ngo³ , Devin Cody^{1,4}, Robert J. De Rosa⁵ , James R. Graham⁶,
Lea Hirsch⁵ , Vighnesh Nagpal⁷, Eric L. Nielsen⁵ , Logan Pearce^{8,11} , Malena Rice^{9,11} , and Roberto Tejada^{5,10}

¹ Department of Astronomy, California Institute of Technology, Pasadena, CA, USA

² Department of Physics and Astronomy, University of California, Los Angeles, Los Angeles, CA, USA

³ NRC Herzberg Astronomy & Astrophysics Research Centre, Victoria, BC, Canada

⁴ Department of Electrical Engineering, California Institute of Technology, Pasadena, CA, USA

⁵ Kavli Institute for Particle Astrophysics and Cosmology, Stanford University, Stanford, CA, USA

⁶ Department of Astronomy, University of California, Berkeley, Berkeley, CA, USA

⁷ Lexington High School, Lexington, MA, USA

⁸ Steward Observatory, University of Arizona, Tucson, AZ, USA

⁹ Department of Astronomy, Yale University, New Haven, CT, USA

¹⁰ California State University, Los Angeles, Los Angeles, CA, USA

Received 2019 October 1; revised 2019 November 17; accepted 2019 November 18; published 2020 February 5

Abstract

orbitize! is an open-source, object-oriented software package for fitting the orbits of directly imaged objects. It packages the Orbits for the Impatient (OFTI) algorithm and a parallel-tempered Markov Chain Monte Carlo (MCMC) algorithm into a consistent and intuitive Python API. orbitize! makes it easy to run standard astrometric orbit fits; in less than 10 lines of code, users can read in data, perform one fit using OFTI and another using MCMC, and make two publication-ready figures. Extensive pedagogical tutorials, intended to be navigable by both orbit-fitting novices and seasoned experts, are available on our documentation page. We have designed the orbitize! API to be flexible and easy to use/modify for unique cases. orbitize! was designed by members of the exoplanet imaging community to be a central repository for algorithms, techniques, and know-how developed by this community. We intend for it to continue to expand and change as the field progresses and new techniques are developed, and call for community involvement in this process. Complete and up-to-date documentation is available at orbitize.info, and the source code is available at github.com/sblunt/orbitize.

Unified Astronomy Thesaurus concepts: [Open source software \(1866\)](#); [Orbit determination \(1175\)](#)

Supporting material: [animation](#)

1. Introduction

By repeatedly imaging exoplanets, we can directly observe them moving along their orbits. The physics behind orbits is well established and straightforward to computationally model, and can reveal much about the properties and dynamical histories of planetary systems. For example, misalignment of the orbital plane with the stellar spin axis can indicate disturbances early in the lifetime of a system before planets formed (Bate et al. 2010; Maire et al. 2019). In systems with circumstellar dust, a planet can warp a disk when the two are mutually inclined (Lagrange et al. 2010; Dawson et al. 2011) or scatter comets into the inner planetary system (Thébaud & Beust 2001; Zieba et al. 2019). Planet–planet scattering and resonant migration in a gas disk can excite observable eccentricities (Yu & Tremaine 2001; Scharf & Menou 2009), potentially implying the presence of unseen planets in planetary systems and constraining formation scenarios. Improved astrometric precision (e.g., Gravity Collaboration et al. 2019) could soon lead to the detection of unseen planets based on perturbations in the orbits we observe. In the future, accurate orbit models of exo-Earths from future space imaging missions will be critical to properly assess their climates and habitability (Williams & Pollard 2002).

While the physics of orbits is straightforward, orbit fitting is challenging, especially for the current generation of directly imaged planets. The wide angular separations needed to detect these objects translate to decades-long or longer orbital periods, which means that the parameter space of possible orbits is often large. Orbit-fitting packages such as ExoFast (Eastman et al. 2013) and PyAstrOFit (Wertz et al. 2017) explore these parameter spaces using a Markov Chain Monte Carlo (MCMC) approach, but such algorithms often converge slowly in this regime. Several specialized techniques have been developed in the past few years to address some of these difficulties (e.g., Blunt et al. 2017; Brandt et al. 2019; O’Neil et al. 2019), but it is often left as an exercise to the reader to implement, debug, and combine these techniques.

In this paper, we present orbitize!, an open-source orbit-fitting software package inspired by radvel (Fulton et al. 2018), designed to meet the needs of the high-contrast imaging community.¹³ orbitize! is designed to consolidate the algorithms, techniques, and know-how of the high-contrast orbit-fitting community in one place. It is fast and robust, but also clearly written, well documented, and easy to use. We have designed orbitize! to be flexible and easily modifiable so it can grow with the field of high-contrast orbit fitting. We believe orbitize! to be a code base with a comparatively low barrier to understanding and contribution.

¹¹ NSF Graduate Research Fellow.

¹² 51 Pegasi b Fellow.

¹³ Versions of orbitize! have been archived to Zenodo, including Version 1.7 concurrent with this article (Blunt et al. 2019).

We seek to remove obstacles to becoming an expert in direct-imaging orbit fitting, enabling the field to advance more rapidly.

This paper is organized as follows. In Section 2, we review Bayesian orbit fitting for directly imaged astrometry and discuss the various algorithms used to perform this procedure. In Section 3, we outline the design of `orbitize!` and give examples of how to use the code. In Section 4, we discuss community involvement guidelines and provide a list of items on the `orbitize!` to-do list, and we conclude in Section 5. We direct readers to articles that use `orbitize!` to learn more about the diverse science applications of this software (Blunt et al. 2019; De Rosa et al. 2019; Ruane et al. 2019; Ruffio et al. 2019).

2. Fitting Imaged Orbits

2.1. Defining the Orbit

Orbits in `orbitize!` are parameterized using the angle conventions from Green (1985): semimajor axis (a), eccentricity (e), inclination angle (i), argument of periastron of the secondary's orbit (ω),¹⁴ longitude of ascending node (Ω), epoch of periastron passage (τ), parallax (π), and total mass (M_T). Note that we express epoch of periastron passage (τ) as a fraction of the orbital period past a specified reference date t_{ref} (default 2020 January 1),

$$\tau = \left(\frac{t_0 - t_{\text{ref}}}{P} \right) \bmod 1, \quad (1)$$

where t_0 is the time of periastron and P is the orbital period. We chose to fit in τ rather than t_0 because the prior bounds for τ are straightforward (between 0 and 1) no matter the orbital period. A screen capture from an interactive module intended to help readers visualize these parameters is shown in Figure 1, and an interactive version is available.¹⁵

In the `orbitize!` coordinate system, motion along the positive z -direction causes a redshift. The positive x -direction is in the positive R.A. direction offset from the primary star, and the positive y -direction is in the positive decl. direction. The orbital elements are defined as usual within this reference frame, with $i = 0^\circ$ corresponding to a face-on orbit.

2.2. Bayesian Orbit Fitting

For high-contrast imaging, Bayesian orbit fitting is the process of converting time-series measurements of a planet's location relative to its primary to a posterior over its orbital parameters. The inputs to this process are times at which measurements were taken and relative position measurements, most often expressed as a planet–star separation (ρ) and position angle (θ), or as offsets in R.A. (α) and decl. (δ).

To compute the posterior over orbital parameters, we use Bayes' theorem,

$$p(m|\mathcal{D}) \propto \mathcal{L}(m|\mathcal{D})p(m), \quad (2)$$

¹⁴ We caution that users must be careful comparing the outputs of radial-velocity orbit-solving packages like `radvel` with those of `orbitize!`, since `radvel` fits the orbit of the star, and `orbitize!` fits the orbit of the planet. In practice, this just means adding 180° to any argument of periastron (ω) values returned by a radial-velocity code. See Figure 1 of Fulton et al. (2018) for a visualization of this difference.

¹⁵ github.com/sblunt/orbitize/blob/master/docs/tutorials/show-me-the-orbit.ipynb




Figure 1. Screen capture of an animated, interactive Python module intended to help users visualize orbits and the `orbitize!` coordinate system. Top: R.A. vs. decl. The orbit's line of nodes is shown as a dashed line, and the primary's location as a pink star. Orbital parameters without associated sliders were fixed at the following values: $a = 10$ au, $\tau = 0.3$, $\pi = 30$ mas, and $M_T = 1.3 M_\odot$. The mass of the secondary was fixed at $0.01 M_\odot$. The dark purple arc shows the portion of the orbit in front of the sky plane, and the lighter purple arc shows the portion that is behind the sky plane. Middle: stellar radial velocity vs. time. Bottom: interactive sliders that users can drag to set the values of orbital elements in the interactive version. A narrated and captioned animation of a user varying the orbital elements by interacting with the sliders is available in the online version of this article. In addition, readers are encouraged to download and interact with this visualization at [this link](#).

(An animation of this figure is available.)

where $p(m|\mathcal{D})$ is the posterior, $\mathcal{L}(m|\mathcal{D})$ is the likelihood (the probability of the orbit model given the data), and $p(m)$ is the prior (the probability of the orbit model itself). By default, `orbitize!` uses a Gaussian likelihood

$$\log \mathcal{L}(m|\mathcal{D}) = -\frac{1}{2}\chi_\theta^2 - \frac{1}{2}\chi_\rho^2 - \frac{1}{2}\chi_\alpha^2 - \frac{1}{2}\chi_\delta^2 \quad (3)$$

$$\chi_\theta^2 = \sum_{i=1}^N \frac{\left[\arctan \left(\frac{\sin(\theta_m(t_i) - \theta_o(t_i))}{\cos(\theta_m(t_i) - \theta_o(t_i))} \right) \right]^2}{\sigma_{\theta_o(t_i)}^2} \quad (4)$$

$$\chi_\rho^2 = \sum_{i=1}^N \frac{(\rho_m(t_i) - \rho_o(t_i))^2}{\sigma_{\rho_o(t_i)}^2} \quad (5)$$

$$\chi_\alpha^2 = \sum_{i=1}^M \frac{(\alpha_m(t_i) - \alpha_o(t_i))^2}{\sigma_{\alpha_o(t_i)}^2}, \quad (6)$$

$$\chi_{\delta}^2 = \sum_{i=1}^M \frac{(\delta_m(t_i) - \delta_o(t_i))^2}{\sigma_{\delta_o(t_i)}^2} \quad (7)$$

where N is the number of observation epochs with observations measured in terms of ρ and θ , M is the number of observation epochs with observations measured in terms of α and δ , t_i is the epoch of the i th observation, $\rho_m(t_i)$ is the separation predicted by the model orbit at t_i , $\rho_o(t_i)$ is the observed separation at t_i , and $\sigma_{\rho_o(t_i)}$ is the observational uncertainty on the observed separation. The inverse tangent function in Equation (4) accounts for angle wrapping near 0 and 2π to ensure that the difference between the model and observed value is calculated correctly. $\rho_m(t_i)$, $\theta_m(t_i)$, $\alpha_m(t_i)$, and $\delta_m(t_i)$ are determined by solving Kepler’s equation, discussed further in Section 3.1.3.

Total mass and parallax are either included as free parameters and assigned priors motivated by observations or held fixed in this analysis. In `orbitize!`, users can select either option.

2.3. Algorithms for Orbit Fitting

Prior to more modern Bayesian techniques, well-constrained visual orbits were typically fit with least-square minimization techniques (e.g., Binnendijk 1960; Sozzetti et al. 1998). While such methods are computationally efficient, and they effectively find the maximum-likelihood orbit with estimates of the 1D errors and 2D covariances, they do not produce full posteriors. This is particularly problematic for poorly constrained orbits, when the posteriors are distinctly non-Gaussian.

A more computationally intensive method is the least-square Monte Carlo (LSMC), which generates input orbits by drawing randomly from a multidimensional parameter space of orbital parameters. Each of the input orbits is run through an iterative χ^2 minimizer until a stopping criterion is reached (e.g., Mugrauer et al. 2012). Several recent publications (Ginski et al. 2013; Vigan et al. 2016; Maire et al. 2019) have demonstrated that while LSMC is capable of finding families of plausible orbits, its output is often significantly different from the Bayesian posterior. Another Monte-Carlo-based method was used in Wagner et al. (2016) and Wagner et al. (2018) to fit short orbital arcs. This method used a multistep grid-search algorithm to find the maximum-likelihood orbit and Gaussian point estimates to approximate confidence intervals.

Both LSMC and grid-search algorithms have difficulty fitting systems with data covering only short orbital arcs, in part because both prioritize finding the set of orbital parameters corresponding to the maximum-likelihood solution. In a well-constrained orbit where the posterior is a 6D Gaussian, either method should find the best-fitting parameters and recover the marginalized uncertainty in each parameter. When the posteriors are decidedly non-Gaussian, however (which is much more often the case for orbit fitting in direct imaging), both techniques fail to adequately derive the shape of the posterior. Additionally, neither of these methods appears to offer a gain in computational efficiency compared to Orbits for the Impatient (OFTI; see below) or MCMC.

In the following two subsections, we describe OFTI and MCMC, the two backend algorithms available in `orbitize!`. These algorithms represent efficient Bayesian methods for producing plausible sets of orbital parameters that represent the full multidimensional posterior of orbital parameters. These posteriors are robust probability density functions with

confidence intervals that are an accurate reflection of the constraints imposed by the data.

2.3.1. OFTI

The OFTI algorithm is described in detail in Blunt et al. (2017), but we review it briefly here. Trial orbits are drawn randomly from priors, and are scaled and rotated to match the data. Scaling and rotating involves modifying the semimajor axis and position angle of nodes of the trial orbit to match the most constrained astrometric data point within the observational uncertainties, cutting down the large parameter space of possible orbits. For each scaled-and-rotated trial orbit, Kepler’s equation is solved (See Section 3.1.3) and a likelihood is computed. Finally, each orbit is either accepted or rejected by comparing the likelihood probability to a uniform random number.

The OFTI algorithm is most efficient when the orbital posteriors are similar to the priors or, in other words, when the parameter space is relatively unconstrained. This occurs most often for short orbital arcs, when the data span only a small fraction of the total orbital period.

In the case of OFTI, each individual orbit considered is uncorrelated with the rest, and so the only stopping criterion is the number of samples desired, which varies by application. To plot plausible orbit tracks, ~ 100 sets of orbital parameters are sufficient. ~ 1000 may be needed for accurate medians and 1σ confidence intervals on each parameter. Plots of 1D marginalized posteriors are well sampled with $\sim 10^4$ sets, and 2D parameters with $\sim 10^6$.

2.3.2. Affine-invariant and Parallel-tempered MCMC

MCMC algorithms are commonly used to sample the posterior of planetary orbits. `orbitize!` makes use of two such algorithms: the Affine-invariant sampler from `emcee` (Foreman-Mackey et al. 2013), and the parallel-tempered MCMC sampler from `ptemcee` (Vousden et al. 2016). Given the complex covariances and often multimodal posteriors of the orbits of directly imaged planets, the Affine-invariant sampler alone generally fails to fully sample the posterior without fine-tuned starting locations for the walkers. We offer the use of `ptemcee`, which runs multiple Affine-invariant samplers with different likelihood weights, as an alternative to overcome this difficulty.

MCMC algorithms generally have similar run times for orbits with partial phase coverage (Blunt et al. 2017). However, convergence time is cut down significantly if the orbital elements are well constrained, such that the posterior is close to a multivariate Gaussian.

Unlike OFTI, the MCMC algorithms coded in `orbitize!` require an initial period for the walkers to fully converge before they sample the posterior in an unbiased fashion. Convergence is assessed using a combination of the walkers’ autocorrelation time as recommended by Foreman-Mackey et al. (2013) and visual inspection of the walker positions over time to determine if they are fully exploring parameter space. Unconverged posteriors typically appear “lumpy,” as different chains are still in different regions of parameter space, compared to generally smooth (though not necessarily Gaussian) posteriors for converged chains. See Figure 4 of Blunt et al. (2017) for an illustration. The multimodal posteriors of ω and Ω can also illustrate convergence. Without a measurement of radial

Table 1
Orbit Model Parameters

Parameter	Symbol	Unit	Default Prior
Semimajor axis	a	au	Log-uniform
Eccentricity	e		Uniform[0, 1]
Inclination	i	rad	Sine[0, π]
Argument of periastron	ω	rad	Uniform[0, 2π]
Longitude of ascending node	Ω	rad	Uniform[0, 2π]
Epoch of periastron passage ^a	τ		Uniform [0, 1]

Note.^a Expressed as a fraction of the orbital period past 2020 January 1.

position or velocity, values 180° apart in both ω and Ω are degenerate. Thus, the chains are more likely to have converged when the 1D marginalized posteriors on ω and Ω show symmetric equal peaks.

3. Code Design

3.1. API

`orbitize!` comprises several modules, each a separate code file. This section refers to `orbitize!` version 1.7.0, but the most up-to-date documentation can be found online.¹⁶ The primary modules are described in the following sections.

3.1.1. Input Data

The `orbitize.read_input` module is designed to read astrometric measurements as input to `orbitize!` in any file format supported by `astropy.io.ascii.read()`, including `csv`, `fixed width`, `cds`, and `LaTeX` formats. The main method in this module is `orbitize.read_input.read_file()`.

This module also contains the method `orbitize.read_input.write_orbitize_input()`, which takes a table of measurements formatted in the `orbitize!` format and writes it as an ASCII file in any file format supported by `astropy.io.ascii.read()`.

3.1.2. Priors

Priors in `orbitize!` are represented as subclasses of the abstract super class `orbitize.priors.Prior`. Each prior class can draw random values from its distribution as well as compute the probability of drawing particular values from that prior distribution. Other parts of the code use these two methods to interface with the priors. Currently, `orbitize!` implements uniform, Gaussian, log-uniform (Jeffreys), sine, and linear priors. The default priors listed in Table 1 are used for orbit fits, but users can easily replace priors by instantiating their own priors and replacing them in the `orbitize.system.System` instance.¹⁷

3.1.3. Kepler's Equation Solver

The `orbitize.kepler` module converts orbital elements to position and radial-velocity measurements by generating an orbital ellipse, placing the companion at the appropriate phase of the orbit, and rotating the orbit based on viewing geometry.

The instantaneous position of a companion along an orbital ellipse is defined by the true anomaly ν , the angle between the location of the companion and periastron, as measured from the focus of the orbital ellipse. ν cannot be analytically calculated from a given epoch and set of orbital elements. It is instead calculated from the eccentric anomaly E , the angle between the projection of the companion's location onto a circle of radius a that intersects the true orbit at periastron, which is calculated in turn from the mean anomaly M , the fraction of the period that has elapsed since periastron passage. See Murray & Correia (2010) for an excellent derivation of the relevant equations from first principles and visualizations of these quantities.

The conversion between M and E , typically called Kepler's equation, is given by

$$M = E - e \sin E. \quad (8)$$

This equation cannot be solved analytically. In `orbitize!`, we solve Kepler's equation using one of two different methods depending on the eccentricity of the orbit. For $e < 0.95$, we use Newton's method to estimate E with a default tolerance of $|E_n - E_{n-1}| < 10^{-9}$. We use $E_0 = M$ as a starting point and typically achieve the required tolerance in three or four iterations. If a solution to E is not found within a default of 50 iterations, the procedure is restarted with $E_0 = \pi$. If a solution is still not found, we instead use the numerical approximation described below to estimate E .

For highly eccentric orbits ($e \sim 1$) the number of iterations required to reach the required tolerance increases dramatically, especially as M approaches 0 or 2π , significantly reducing the speed of this solver (Figure 2, left panel). This can be especially problematic for the parallel-tempered MCMC sampler described in Section 3.1.6, where the high-temperature walkers will explore the full range of eccentricities allowed by the prior distribution. To mitigate this slowdown, we instead use the numerical approximation for Kepler's equation described by Mikkola (1987) for orbits with $e \geq 0.95$. Mikkola's algorithm invokes a cubic approximation to Kepler's equation and a numerically intensive (but single iteration) high-precision correction formula.

We verified the accuracy of Mikkola's method by comparing the value of E calculated for a range of (M, e) combinations using Newton's method (with a tolerance of 10^{-18}) and using Mikkola's method. We find a maximum absolute difference between the eccentric anomaly computed via both methods of 1.6×10^{-15} (Figure 2, middle panel), well below the nominal tolerance on the eccentric anomaly solver of $|E_n - E_{n-1}| < 10^{-9}$. While Mikkola's method is strictly more accurate than applying Newton's method with a tolerance of 10^{-9} , it is much more computationally expensive, especially for low-eccentricity orbits where it is a factor of 10 slower (Figure 2, right panel). For the vast majority of applications for `orbitize!` the default tolerance of 10^{-9} corresponds to a position angle error of $3.6 \times 10^{-7\circ}$ for a face-on circular orbit, well below the precision of current astrometric measurements.

These algorithms are implemented in C for maximum computational efficiency. A Python version of this module is also included in `orbitize!`, and the package will revert to the Python version if the C-solver fails to compile on a user's machine.

¹⁶ orbitize.info¹⁷ http://orbitize.info/en/latest/tutorials/Modifying_Priors.html

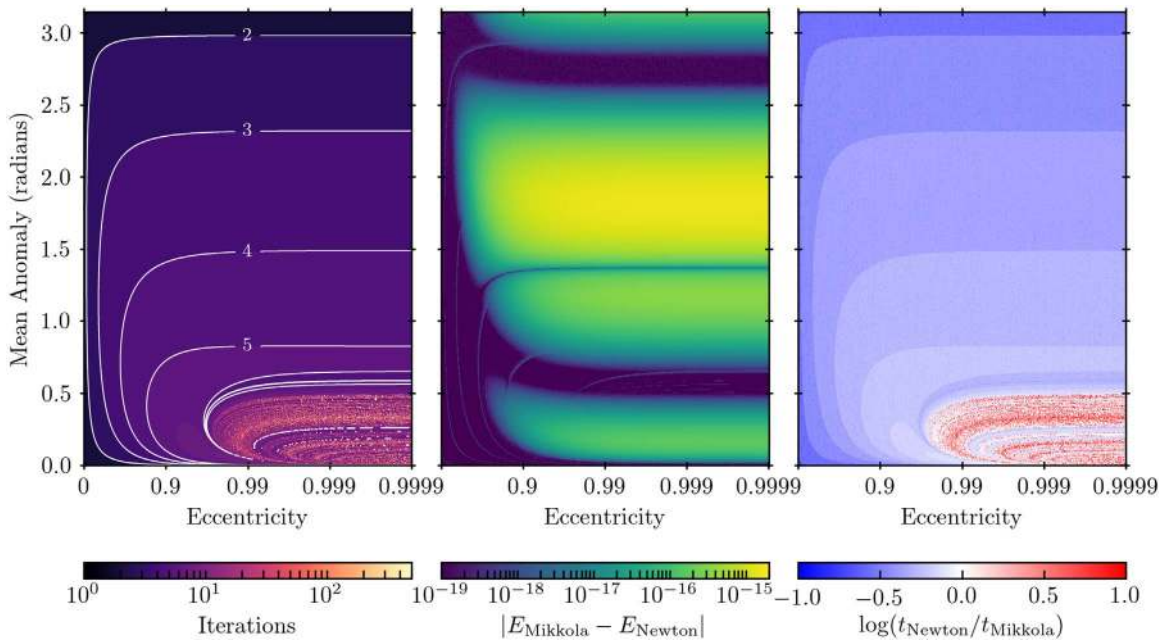


Figure 2. Performance summary of our two Keplerian solvers. Left: number of iterations by the Newton solver required to achieve a tolerance of $1e-9$ as a function of eccentricity and mean anomaly. The Mikkola solver is used for orbits with $e \geq 0.95$ to mitigate this slowdown at high eccentricities. Middle: absolute difference between the eccentric anomaly values computed by the Mikkola and Newton solvers. A tolerance of $1e-18$ was used in this experiment. The absolute difference between the two solvers is negligible compared with current astrometric precision. Right: time needed for both solvers to achieve a tolerance of $1e-18$. The Mikkola solver achieves significant performance gains over the Newton solver for high eccentricities.

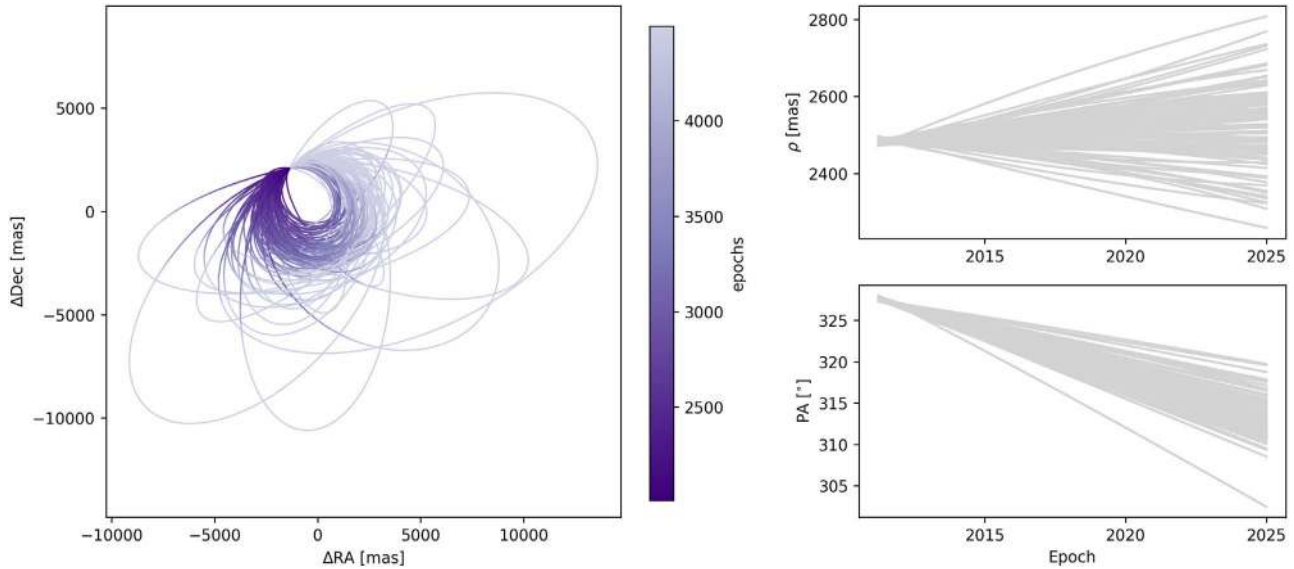


Figure 3. Orbit figure produced from the code snippet. Left: 100 orbits drawn from the posteriors computed for GJ 504 b, projected on the sky and color coded by year. Upper right: separation vs. time for the same 100 orbits. Lower right: position angle vs. time for the same 100 orbits.

3.1.4. Planetary System

The main component of the `orbitize.system` module is the `orbitize.system.System` class, which stores observational data, priors, and posteriors for a star-planet system. To initialize an instance of `orbitize.system.System`, a user inputs a data table, created using the `orbitize.read_input` module, a total system mass, and a parallax. The `__init__()` method automatically initializes a list of `orbitize.priors.Prior` objects in a standard order, detailed in the online documentation. `orbitize.system.System` also has a `compute_model()` method that takes in

an array of potential orbital parameters and computes model predictions to be compared against the data. This is the method called the Keplerian orbit-solver (Section 3.1.3). Note that this abstraction makes the sampler agnostic of physics; a user could use any model computation code as a drop-in replacement. Finally, this class has an attribute `results`, a list of `orbitize.results.Results` objects (Section 3.1.5). In version 1.7.0, `orbitize!` is limited to single-planet systems. Users can create `orbitize.system.System` objects for each planet in a multiplanet system, but orbit fits using this framework currently do not take into account the dynamical effects of other planets on the orbit fit.

3.1.5. Results

The `orbitize.results` module contains code for the `orbitize.results.Results` class. Each instance of this class represents a posterior (whether calculated using OFTI or MCMC). The `orbitize.sampler.Sampler` objects add their output to an instance of this class. This module also controls saving, loading, and visualizing results.

3.1.6. Samplers

`orbitize!` implements the OFTI algorithm (Section 2.3.1) in the `orbitize.sampler.OFTI` class and two MCMC algorithms (Section 2.3.2) in the `orbitize.sampler.MCMC` class.

The `run_sampler()` method of the `orbitize.sampler.OFTI` class generates a posterior of permissible orbits. This method iteratively runs a series of methods that define the process of generating and statistically vetting potential orbits until a desired number of orbits is accepted. Accepted orbits are then added to an instance of the `orbitize.Results` class. This process is arbitrarily parallelizable, and users can easily set the number of CPU cores available for fits.

`orbitize.sampler.MCMC` uses MCMC algorithms to generate a Markov Chain representing the posterior. Setting the attribute `num_temps` to 1 invokes the affine-invariant sampler, and setting it to a number greater than 1 invokes the parallel-tempered sampler. Users also choose the number of walkers, the number of threads, and whether they would like to fix certain parameters so that they are not sampled by the MCMC algorithm. A convenient API for inspecting chains to assess convergence is available, and is explained in the online tutorials and documentation.

3.1.7. Driver

Finally, the `orbitize.driver` module automates the creation of the data table and `orbitize.system.System` objects. It is a convenient shortcut for standard orbit fits, and allows new users to begin using the code relatively quickly. While using this object makes it simple to run `orbitize!` in a standard way, users are encouraged to learn about the underlying API in order to learn how to, e.g., modify priors, customize plots, and set parameters specific to OFTI or MCMC.

3.2. Example Usage

To illustrate the ease with which a user can run an orbit fit with `orbitize!`, we provide a minimal code example below. The first code block is a Python program a user could run to produce orbit posteriors and figures, and the second code block shows the `.csv` file used as input. One of the strengths of `orbitize!` is its customizability, however, and this example is merely intended to show how easy it is to run a standard orbit fit, not to give a sense of the full scale of the code’s capabilities. We encourage users to peruse the online tutorials¹⁸ for more in-depth examples of how to use and modify `orbitize!`.

In the example below, an `orbitize.driver.Driver` object is initialized, and the OFTI algorithm runs until 10,000 orbits have been accepted. To use the MCMC instead of the OFTI algorithm, the user just needs to switch out the keyword in line 6. In-line comments are provided to aid understanding.

The two figures produced by this code snippet are shown in Figures 3 and 4.

```

1  from orbitize import driver
2  from orbitize import DATADIR
3
4  myDriver = driver.Driver(
5      '{}/GJ504.csv'.format(DATADIR), # data file
6      'OFTI', # choose from: ["OFTI", "MCMC"]
7      1, # number of planets in system
8      1.22, # total system mass [M_Sun]
9      56.95, # system parallax [mas]
10     mass_err = 0.08, # mass error [M_Sun]
11     plx_err = 0.26 # parallax error [mas]
12 )
13 orbits = myDriver.sampler.run_sampler(100000)
14
15 # plot the results
16 myResults = myDriver.sampler.results
17 orbit_figure = myResults.plot_orbits(
18     # minimum MJD for colorbar (choose first data epoch)
19     start_mjd = myDriver.sampler.epochs[0]
20 )
21 corner_figure = myResults.plot_corner()
```

epoch	object	sep	sep_err	pa	pa_err
55645.95	1,2479,16	327.94	0.39		
55702.89	1,2483,8	327.45	0.19		
55785.015	1,2481,33	326.84	0.94		
55787.935	1,2448,24	325.82	0.66		
55985.19400184	1,2483,15	326.46	0.36		
56029.11400323	1,2487,8	326.54	0.18		
56072.30200459	1,2499,26	326.14	0.61		

3.3. Comparison to Published Results

As an example of the reliability of `orbitize!`, we calculated the orbital posterior for GJ 504 b and compared it with the results presented in Blunt et al. (2017). The orbital elements in that reference were determined using an independent implementation of OFTI. We obtained 1,000,000 sample orbits using `orbitize!` and the published astrometry. Our `orbitize!` orbital parameter posteriors agree with the published distributions with 95% confidence, as shown in Table 2. We carried out an identical comparison to the orbit fit presented in Pearce et al. (2019) and find that the results also agree for this system.

4. Future Prospects

4.1. Community Involvement Guidelines

In order to realize our goals of consolidating direct-imaging orbit-fitting best practices in one place and continually adapting `orbitize!` to serve the direct-imaging community, we require substantial community involvement. We strongly encourage anyone interested in using `orbitize!` to contribute code. We maintain a document with specific contributor guidelines¹⁹ on our GitHub page.

¹⁸ <http://orbitize.info/en/latest/tutorials.html>

¹⁹ https://github.com/sblunt/orbitize/blob/master/contributor_guidelines.md




Figure 4. Corner plot produced from the code snippet. The diagonals show 1D marginalized posteriors for each free parameter in the GJ 504 b orbit fit (semimajor axis, eccentricity, inclination, argument of periastron, position angle of nodes, epoch of periastron passage, parallax, and total mass). The off-diagonals show 2D covariances between these parameters. This plot was produced using `corner` (Foreman-Mackey 2016).

In addition to directly contributing code, we encourage community members to request features, report bugs, and provide feedback through raising issues on GitHub. This is the most efficient way to reach the entire development team.

4.2. Future Work

`orbitize!` has a long and exciting list of planned updates. A version 2.0 release, intended to implement many of these, is planned for 2020. Key upgrades in version 2.0 will be enabling

dynamical mass measurements by jointly fitting radial velocity and stellar astrometry data sets, as well as fitting multiplanet systems. To summarize, the major upcoming features are:

1. jointly fitting radial velocities;
2. incorporating independent radial-velocity fits as priors;
3. jointly fitting stellar astrometry;
4. fitting multiplanet and hierarchical systems;
5. fitting using other orbital element parameterizations (e.g., $\sqrt{e} \cos \omega$ and $\sqrt{e} \sin \omega$ rather than e and ω)

Table 2
Orbital Parameter Posteriors for GJ 504 b Obtained Using `orbitize!` Compared to Those Published in Blunt et al. (2017)

Element	<code>orbitize!</code>			Blunt et al. (2017)		
	Median	68% CI	95% CI	Median	68% CI	95% CI
a (au)	47	(40, 65)	(32, 111)	48	(39, 69)	(31, 129)
e^a	0.16	(0.04, 0.36)	(0.00, 0.56)	0.19	(0.05, 0.40)	(0.01, 0.62)
i ($^\circ$)	140	(126, 160)	(116, 172)	140	(125, 157)	(111, 171)
ω^b ($^\circ$)	103	(36, 154)	(10, 176)	95	(31, 151)	(4, 176)
Ω^b ($^\circ$)	106	(52, 153)	(11, 175)	97	(46, 146)	(8, 173)
T_0^c	2159	(2073, 2293)	(2004, 2459)	2145	(2068, 2310)	(2005, 2825)

Notes.

^a A linearly descending eccentricity prior was applied, as in Blunt et al. (2017).

^b ω and Ω have been wrapped by 180° , as in Blunt et al. (2017).

^c Epoch of periastron passage is reported as T_0 here (as a decimal year, contrary to the `orbitize!` convention) in order to facilitate easy comparison with Blunt et al. (2017).

- incorporating an N -body integrator to replace the standard *Kepler* solver where three-body interactions are non-negligible;
- adding instrument-specific jitter and offset terms to the parameterization to account for potential uncorrected instrumental calibration systematics;
- adding a Hamiltonian MCMC algorithm backend;
- and incorporating observation-driven priors (e.g., O’Neil et al. 2019).

5. Conclusion

In this paper, we have presented `orbitize!`, an open-source Python package for fitting the orbits of directly imaged planets. `orbitize!` uses OFTI and MCMC, two efficient Bayesian methods for computing posteriors. We aim to encourage community contributions, remove barriers to becoming an expert in orbit fitting, and provide an open-source development environment in order to meet the orbit fitting needs of the high-contrast exoplanet imaging community.

The authors thank those at academic and telescope facilities whose labor maintains spaces for scientific inquiry, particularly those whose communities are excluded from the academic system.

This package was born and developed at the winter 2018 and 2019 AAS Hack Days, and the authors wish to thank the organizers of these events. S.B. would like to thank Logan Cody for superb video editing. The authors would also like to thank Dillon Dong, Jasmine Garani, Melisa Tallis, and Daniel Yahalomi for their time and initial work on `orbitize!`, and Junellie Gonzalez, Kelly Kosmo O’Neil, Ryan Rubenzahl, and Jean-Baptiste Ruffio for participating in our hackathons and for their anticipated future contributions to `orbitize!`. S.B. would like to thank Andrew Howard, B.J. Fulton, and Erik Petigura for helpful conversations. Finally, we thank the anonymous individual who unknowingly named `orbitize!` at the 2018 AAS Hack Day.

Yes, you need to use the exclamation point!

ORCID iDs

Sarah Blunt  <https://orcid.org/0000-0002-3199-2888>
Jason J. Wang  <https://orcid.org/0000-0003-0774-6502>

Isabel Angelo  <https://orcid.org/0000-0002-9751-2664>
Henry Ngo  <https://orcid.org/0000-0001-5172-4859>
Robert J. De Rosa  <https://orcid.org/0000-0002-4918-0247>
Lea Hirsch  <https://orcid.org/0000-0001-8058-7443>
Eric L. Nielsen  <https://orcid.org/0000-0001-6975-9056>
Logan Pearce  <https://orcid.org/0000-0003-3904-7378>
Malena Rice  <https://orcid.org/0000-0002-7670-670X>

References

- Bate, M. R., Lodato, G., & Pringle, J. E. 2010, *MNRAS*, **401**, 1505
Binnendijk, L. 1960, Properties of Double Stars; A Survey of Parallaxes and Orbits (Philadelphia, PA: Univ. Pennsylvania Press)
Blunt, S., Endl, M., Weiss, L. M., et al. 2019, *AJ*, **158**, 181
Blunt, S., Ngo, H., Wang, J., et al. 2019, `sblunt/orbitize`, Version 1.7.0, Zenodo, doi:10.5281/zenodo.3534546
Blunt, S., Nielsen, E. L., De Rosa, R. J., et al. 2017, *AJ*, **153**, 229
Brandt, T. D., Dupuy, T. J., & Bowler, B. P. 2019, *AJ*, **158**, 140
Dawson, R. I., Murray-Clay, R. A., & Fabrycky, D. C. 2011, *ApJL*, **743**, L17
De Rosa, R. J., Esposito, T. M., Hirsch, L. A., et al. 2019, *AJ*, **158**, 225
Eastman, J., Gaudi, B. S., & Agol, E. 2013, *PASP*, **125**, 83
Foreman-Mackey, D. 2016, *JOSS*, **1**, 24
Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, **125**, 306
Fulton, B. J., Petigura, E. A., Blunt, S., & Sinukoff, E. 2018, *PASP*, **130**, 044504
Ginski, C., Neuhäuser, R., Mugrauer, M., Schmidt, T. O. B., & Adam, C. 2013, *MNRAS*, **434**, 671
Gravity Collaboration, Lacour, S., Nowak, M., et al. 2019, *A&A*, **623**, L11
Green, R. M. 1985, Spherical Astronomy (Cambridge: Cambridge Univ. Press)
Lagrange, A. M., Bonnefoy, M., Chauvin, G., et al. 2010, *Sci*, **329**, 57
Maire, A. L., Rodet, L., Cantalloube, F., et al. 2019, *A&A*, **624**, A118
Mikkola, S. 1987, *CeMec*, **40**, 329
Mugrauer, M., Röhl, T., Ginski, C., et al. 2012, *MNRAS*, **424**, 1714
Murray, C. D., & Correia, A. C. M. 2010, in Exoplanets, ed. S. Seager (Tucson, AZ: Univ. Arizona Press), 15
O’Neil, K. K., Martinez, G. D., Hees, A., et al. 2019, *AJ*, **158**, 4
Pearce, L. A., Kraus, A. L., Dupuy, T. J., et al. 2019, *AJ*, **157**, 71
Ruane, G., Ngo, H., Mawet, D., et al. 2019, *AJ*, **157**, 118
Ruffio, J.-B., Macintosh, B., Konopacky, Q. M., et al. 2019, *AJ*, **158**, 200
Scharf, C., & Menou, K. 2009, *ApJL*, **693**, L113
Sozzetti, A., Spagna, A., & Lattanzi, M. G. 1998, *EM&P*, **81**, 103
Thébaud, P., & Beust, H. 2001, *A&A*, **376**, 621
Vigan, A., Bonnefoy, M., Ginski, C., et al. 2016, *A&A*, **587**, A55
Vousden, W. D., Farr, W. M., & Mandel, I. 2016, *MNRAS*, **455**, 1919
Wagner, K., Apai, D., Kasper, M., et al. 2016, *Sci*, **353**, 673
Wagner, K., Dong, R., Sheehan, P., et al. 2018, *ApJ*, **854**, 130
Wertz, O., Absil, O., Gómez González, C. A., et al. 2017, *A&A*, **598**, A83
Williams, D. M., & Pollard, D. 2002, *IJAsB*, **1**, 61
Yu, Q., & Tremaine, S. 2001, *AJ*, **121**, 1736
Zieba, S., Zwintz, K., Kenworthy, M. A., & Kennedy, G. M. 2019, *A&A*, **625**, L13