2022 Sagan Summer Workshop: `orvara` Hands-On Session

Timothy Brandt (UCSB)

Part I: Wednesday

In this first part, you will use `orvara` to recover a mass and orbital fit of a known companion. Example configuration files are already distributed with `orvara`, within the directory `orvara` → `tests`, and the data files you will need are within the directory `orvara` → `tests` → `data`. You can find them on `github` here:
`https://github.com/t-brandt/orvara/tree/master/orvara/tests`
`https://github.com/t-brandt/orvara/tree/master/orvara/tests/data`

There are three stars with companions that are supplied: HD 159062 (a main sequence star plus white dwarf), Gl 758 (a main sequence star plus a brown dwarf), and HD 4747 (another main sequence star plus brown dwarf). In all cases,

a) Install `orvara` or start the Jupyter notebook. Instructions and links to a Colab notebook (which does not require a local python installation) are here:
`https://nexsci.caltech.edu/workshop/2022/SSW2022_Google_Colab_Instructions.pdf`

b) Complete the tutorial Jupyter notebook supplied:
`https://github.com/t-brandt/orvara/blob/master/Tutorial.ipynb`

c) Perform an orbital fit using the configuration file of the star of your choice. You shouldn't have to modify much of anything.

d) Make plots, including a diagnostic plot. Choose a burn-in length where the chains seem to be reasonably well-converged, or if they are not converged, run a longer chain. How can you estimate convergence?

e) Make the set of plots available in `orvara` (they are saved by default in `orvara/tests/plots`), and check your answers against the published results in
`https://ui.adsabs.harvard.edu/abs/2021AJ....162..186B/abstract` (for HD 159062)
`https://ui.adsabs.harvard.edu/abs/2021AJ....162..301B/abstract` (for Gl 758, HD 4747)
I am reproducing a few of those plots here.

To the right: the astrometric orbit for HD 159062B as seen in the plane of the sky.

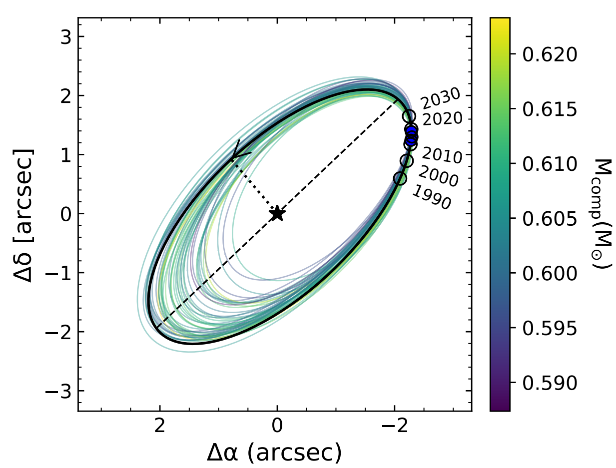

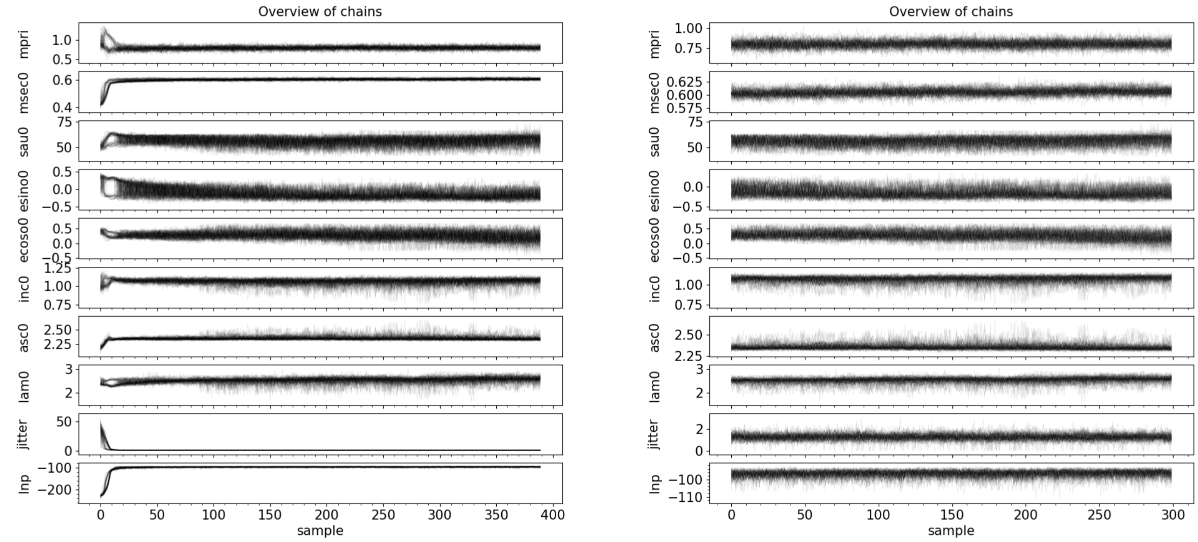

Above: diagnostic plots for HD 159062. The left-hand plot shows the early burn-in time, while the right-hand plot does not show these early points. After they have been discarded as burn-in, the chains appear to be stably clustered around the same values. Look especially at the likelihood – the bottom panel. Are all chains clustered around the same (hopefully global) best-fit orbits?

Right: the corner plot for HD 159062 after discarding the burn-in (corresponding to the right-hand diagnostic plot above). If I did not discard the burn-in there would be tails toward the parameter values seen at the beginning of the chains in the left-hand plot above.

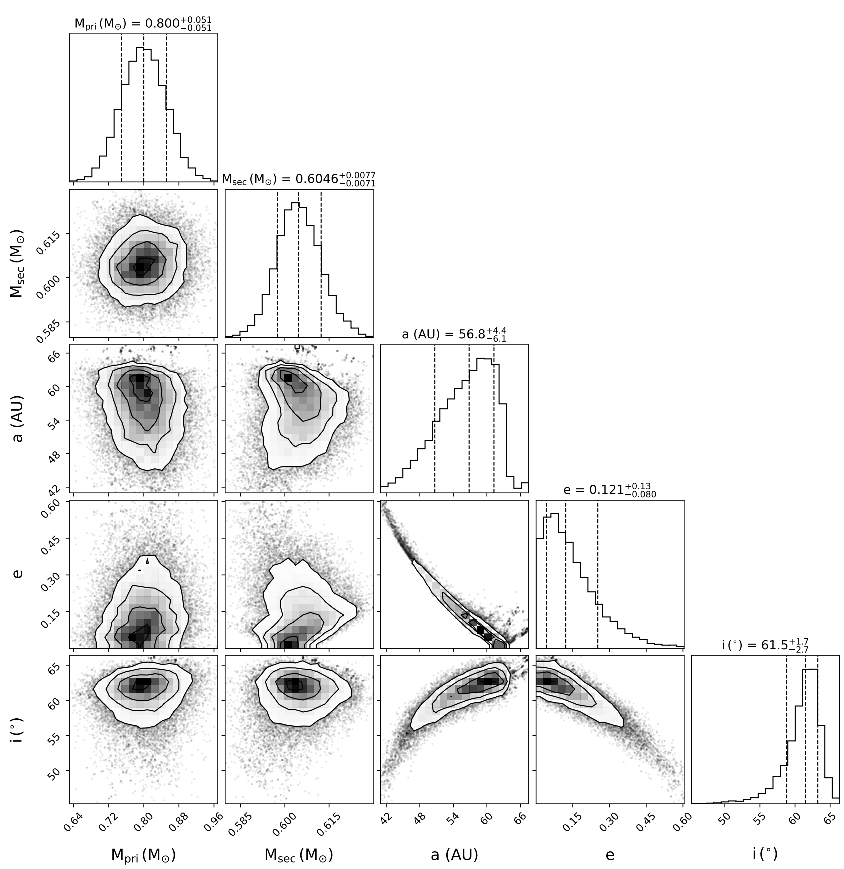

Part II: Thursday

In this second part, you will apply `orvara` to another problem, with a less straightforward result/interpretation. For this case, you have a few options.

> Note: You can use the notebooks `SSW2022_Jupyter_Orvara_Projects.ipynb` (Jupyter) or `SSW2022_Colab_Orvara_Projects.ipynb` (Google Colab) to work on these projects (you can find the links in the Python and Google Colab instructions documents, respectively). The data sets that will be fit are automatically downloaded if you use the Google Colab notebook, while links to the data sets are provided in the Jupyter notebook. Instructions and links to a Colab notebook (which does not require a local python installation) are at the same url as for Part I. A starter notebook is available here. Note that you would want to download this and then open it in Jupyter (it won't look good in a browser without Jupyter).

1. Take data from a system in the literature. My suggestion for this is HIP 113201, with radial velocities and relative astrometry taken from `https://arxiv.org/abs/2112.05457`. This time, you will have to set up the configuration file yourself, using one of the earlier files as a template. When you run and analyze your chain, keep the following questions in mind:

a) How can you assess whether your orbital fit is converged? Try to verify that your chain is long enough and/or that you used enough temperatures, and that you discarded enough data as burn-in.

b) Is your best-fit orbit a formally good fit? To assess this, you can use the χ^2 values of the best-fit orbit that are printed to a file with a name similar to `beststep_params_Gl758.dat`. The χ^2 values of a good fit should be comparable to the number of data points. For example, if there are ten relative position measurements, then a good χ^2 for relative separation would be ≈10 (classically, $N \pm \sqrt{N}$, where N is the number of data points minus the number of parameters that went into fitting this parameter. In this case, χ^2 values of 6 and 15 would both be ok, 30 would not be.)

c) If your best-fit orbit is not a formally good fit, suggest reasons why this might be the case, and what you could do to investigate the cause and run a more meaningful fit. Try out at least one of these adjustments, and see what difference it makes.

2. As another possibility, try fitting a system with two planets: an inner one and an outer one. To get this to converge you might have to give `orvara` a starting file with a reasonably good guess. You may also want to use an informative prior on the mass of the star (`mpri` and `mpri_sig` in the configuration file). One suggestion is Gl 86 (=HIP 10138), with data taken from `https://ui.adsabs.harvard.edu/abs/2021arXiv211206394Z/abstract`. To plot your results, set `iplanet` under `[plotting]` to either zero or one (the outer white dwarf is number 1 in the files supplied for Gl 86). If you do use Gl 86 (or a similar system with an inner planet on an orbit shorter than ∼1 year), keep the following questions in mind and try to answer them:

a) Your posterior on the inclination of the inner planet should look a lot like $\sin i$. Why?

b) Assuming that you used an informative prior for the star's mass, you should have a strong covariance between mass and semimajor axis for the inner planet. Why?

c) Are your χ^2 values all reasonable for the best-fit orbit? If not, what are possible causes and what impact might they have on your results?

d) It is not possible in a system like Gl 86 to constrain the angle between the orbits of the inner and outer companions with the data we have. But if the inner companion transits, as is the case for π Mensae c (π Mensae b is an outer planet), then it is possible. Assume the inclination of the inner planet is 90°. Use the inclination values in your chain to construct a posterior probability distribution of the mutual alignment between the orbital planes of the inner companion (assumed to orbit edge-on) and the outer companion. This exercise requires a bit of geometry!

3. Fit a Gaia binary (both the primary and secondary are in Gaia) using only Gaia and Hipparcos-Gaia data. For this case, a good Gaia binary star is HIP 18267. To make the relative astrometry file for HIP 18627, please use the following approach (you can repeat this for any star with a companion seen in Gaia), modeling the format of the file after any of the other examples provided:

1. Perform a Gaia archive search (`https://gea.esac.esa.int/archive/`) for the star of interest, here HIP 18267. You should see two stars at the same parallax within 5″ of each other, one brighter (6.65 mag) and one fainter (10.43 mag). The brighter star is the primary.

2. Compute the separation in RA and in Decl. This is a little tricky. For Decl., just take the Decl. of the companion (the fainter star here), subtract that of the primary, and multiply by 3600 to convert from degrees to arcseconds. For RA you do the same thing, but you need to then multiply this difference by the cosine of the Decl. to account for the arc length in spherical coordinates. You can use the average Decl. of the two stars for the cos(Decl.) term.

3. Go from ΔRA and ΔDecl. to separation and position angle using $\sqrt{(\Delta\mathrm{RA})^2 + (\Delta\mathrm{Decl.})^2}$ and $\arctan2(\Delta\mathrm{RA}, \Delta\mathrm{Decl.})$, respectively. Make sure that your separation is in arcseconds and your position angle is in degrees!

4. Use 2016.0 as the epoch: this is the Gaia reference epoch.

5. Don't worry too much about the uncertainties (you can use 1 mas and 0.01 degrees if you want): the Gaia errors are so small that they will be irrelevant to your results. You can use zero for the covariances and for the companion ID (last column in the data file with the astrometry).

6. Add the proper motion information of the secondary to the configuration file: under `secondary_gaia`, set `companion_id` to `0` (matching your choice for companion ID

above), and enter the proper motions in RA and Decl., their uncertainties, and their correlation coefficient from the Gaia archive (for the latter, you may have to re-run your query with the `pmra_pmdec_corr` column selected). For example, for the companion to HIP 18267, you should enter 209.739 for the companion's RA proper motion.

Then run your fit! See how good the constraints are from just this one Gaia data point. If you have time,

a) Make some fake, but plausible, additional data. This could be a measurement of the relative position in 2022, for example, or some RV points between 2022 and 2024. If you add fake RV points, you will have to convert from year to BJD for the file (0:00 on Jan 1 2023 is 2459945.5, for example). How much could these additional measurements improve your constraints on the masses and orbit? Are they worth proposing for telescope time to get?

4. Find a system with only radial velocity data, or only direct imaging data. There are many systems with only radial velocities; you can find your own in one of several public catalogs. We will supply two of them: HD 221420 (=HIP 116250) and HIP 10339. We will also supply two systems with only direct imaging data: like HIP 29860 and HIP 117712. Perform an orbital fit to the data you have on one or more of these stars (or one of your own choosing). Keep the following questions in mind:

a) Is your orbital fit formally good? Note that the RV fit will usually be formally good, because RV jitter inflates the errors to make it so! But the absolute astrometry and imaging parts do not have to be formally good fits.

b) Is your fit converged? How sure are you, and how can you be more sure?

c) How important or useful would additional data be? To answer this, you have a few options. If you are using RV data, you can make a plot of the predicted relative astrometry. You can then add a couple of fake (but plausible, according to this plot) relative astrometry measurements and see how much they help. This exercise could prioritize systems for imaging follow-up. If you are using only imaging data and are missing RVs, you can do the same exercise in RVs. Look at the RV plot and see how much different orbits and masses vary in their RV predictions. Make a fake RV file with plausible data and see how much it helps: how long would you need RV data, and at what precision, to significantly help with the constraints?