The BONNSAI project - Documentation

Fabian R.N. Schneider

31st January 2014

Contents

1 The basic principle

1.1 In theory

The aim of the BONNSAI project is to offer a possibility to match an arbitrary set of observables $\boldsymbol{d}$ of stars such as luminosities, effective temperatures, surface abundances etc. and their uncertainties simultaneously to stellar models to find probability distribution functions of stellar model parameters $\boldsymbol{m}$ such as initial mass and stellar age. Our approach takes prior knowledge like the initial mass function or the distribution of rotational velocities into account and is based upon Bayes' theorem,

$$p(\boldsymbol{m}|\boldsymbol{d}) \propto p(\boldsymbol{d}|\boldsymbol{m})p(\boldsymbol{m}), \tag{1}$$

which connects the posterior probability, $p(\boldsymbol{m}|\boldsymbol{d})$, of the model parameters $\boldsymbol{m}$ given the observational data $\boldsymbol{d}$ to the likelihood, $p(\boldsymbol{d}|\boldsymbol{m})$, i.e. the probability of the observed data given the model, and the prior function, $p(\boldsymbol{m})$, i.e. the a priori probability of the model parameters. Bayes' theorem follows directly from the definition of conditional probabilities.

Currently we support two different kind of likelihood functions: (1) a Gaussian- and (2) a step-function. The combined likelihood function is given by the product of the individual likelihood functions for each observable $L_i = L(d_i|\boldsymbol{m})$, i.e. $p(\boldsymbol{d}|\boldsymbol{m}) = \prod_i L_i$. Doing so, we assume that all observables are independent of each other, i.e. that there are no correlations between them. This is however not necessarily always the case. The Gaussian function is used if observables including 1σ uncertainties are known and the step function if only lower and upper limits are known (e.g. if only an upper limit on a surface abundance is known).

1.2 In practice

In practice, BONNSAI does the following steps:

1. select those stellar models that are within 5σ of the observables from a database (this limits the parameter space to speed up the process)
2. compute the posterior probability given the likelihood and prior for each selected stellar model according to Bayes' theorem (Eq. 1)
3. add the posterior probabilities of all selected stellar models to create histograms and 2d posterior probability maps for stellar parameters
4. re-normalize the posterior probabilities such that $\int_{\boldsymbol{m}} p(\boldsymbol{m}|\boldsymbol{d})\,\mathrm{d}\boldsymbol{m} = 1$
5. analyse the posterior probability distributions to provide mean, median and mode including 1σ uncertainties for all requested stellar parameters

2 Submitting a job

Submitting a job is a five step process. First, the set of stellar models is chosen, second the observables including their uncertainties are provided, third the prior functions are chosen, forth the exact output is selected, fifth the advanced settings are adjusted and sixth the job is submitted. The steps are described in detail below.

2.1 Selecting a set of stellar models

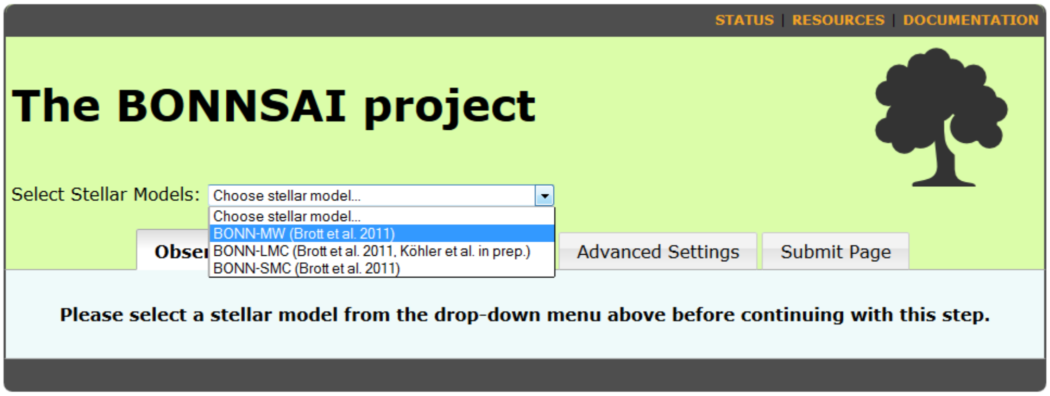

Figure 1: First step: selecting an appropriate set of stellar models for a given star.

At the very beginning, an appropriate set of stellar models has to be chosen (Fig. 1). Table 1 gives an overview of the stellar models that are currently supported by BONNSAI. The observed star has to be covered by the stellar models e.g. in mass. Otherwise, BONNSAI can fail and no output will be generated. BONNSAI automatically tests whether the provided observables are covered by the stellar grid. Note however that e.g. the luminosity and effective temperature of a Galactic main-sequence star might also be covered by stellar models of a different metallicity. After selecting the stellar models, BONNSAI assigns an internal job identification number (job ID) to each request which will be displayed at the top of the page. The job ID is needed to see the status of the job (click on "Status" and type in the job ID).

Table 1: Stellar models supported by BONNSAI.

Stellar models	Z	$M_{\rm ini}$	$v_{\rm ini}$	Age	Notes
Bonn MW[1]	0.0088	5–50 $\mathrm{M}_\odot$	0–600 $\mathrm{km\,s^{-1}}$	0–100 Myr	MS single stars
Bonn LMC[1,2]	0.0047	5–500 $\mathrm{M}_\odot$	0–600 $\mathrm{km\,s^{-1}}$	0–100 Myr	MS single stars
Bonn SMC[1]	0.0021	5–60 $\mathrm{M}_\odot$	0–600 $\mathrm{km\,s^{-1}}$	0–100 Myr	MS single stars

References: [1] Brott et al. (2011), [2] Koehler et al. (in preparation)

2.2 Providing the observables

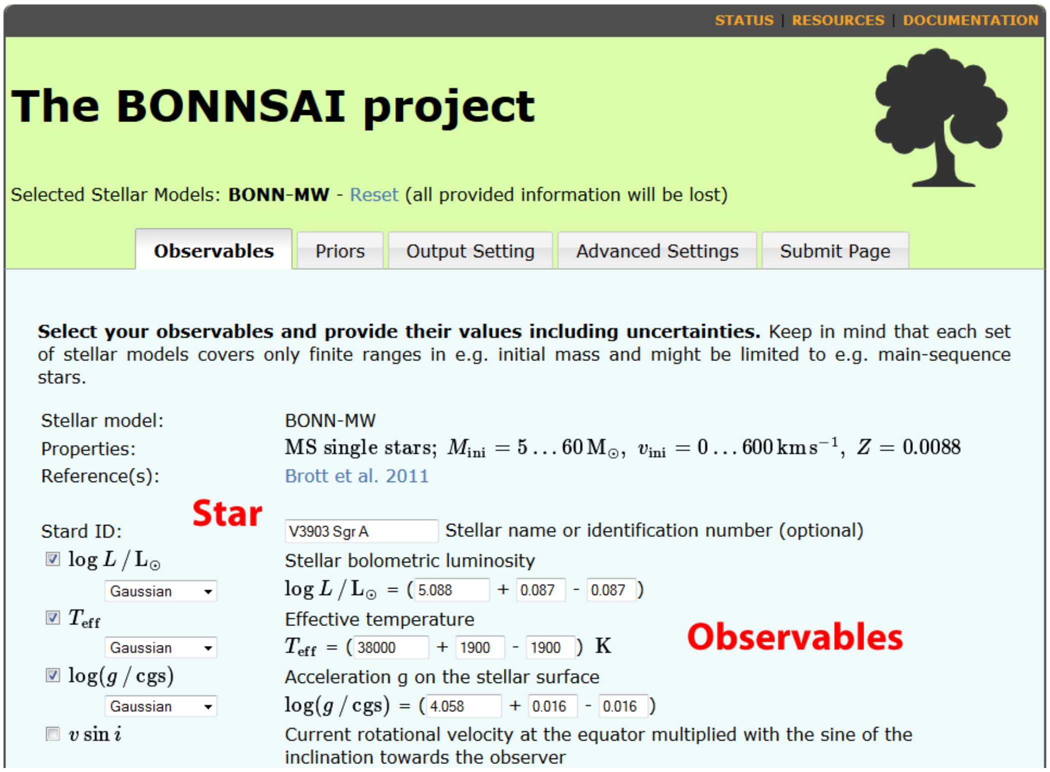

Figure 2: Second step: providing the observables including their uncertainties.

In the second step, the observables including their uncertainties have to be provided. At the top of the page, a brief summary of the chosen stellar models is given to assist the user to judge whether the chosen stellar models are the correct ones (if not, click "Reset" and chose a more appropriate set of stellar models). A stellar name or identification number should be selected; otherwise, an automatic name consisting of the current date is used.

A list of stellar quantities supported by the chosen stellar models is presented. Clicking on the stellar quantity (or checking the checkbox), makes the input fields for the values and the 1σ uncertainties of the observable visible. Furthermore the likelihood function can be changed from Gaussian to lower or upper limit. In the latter cases, no uncertainties need to be provided. Note however, that at least one observable must follow a Gaussian likelihood, i.e. providing observables with lower and upper limits only does not work at the moment.

In Fig. 2, we show an example for the primary O7V star of the Galactic binary V3903 Sgr. The observational data is from Torres et al. (2010).

2.3 Choosing the priors

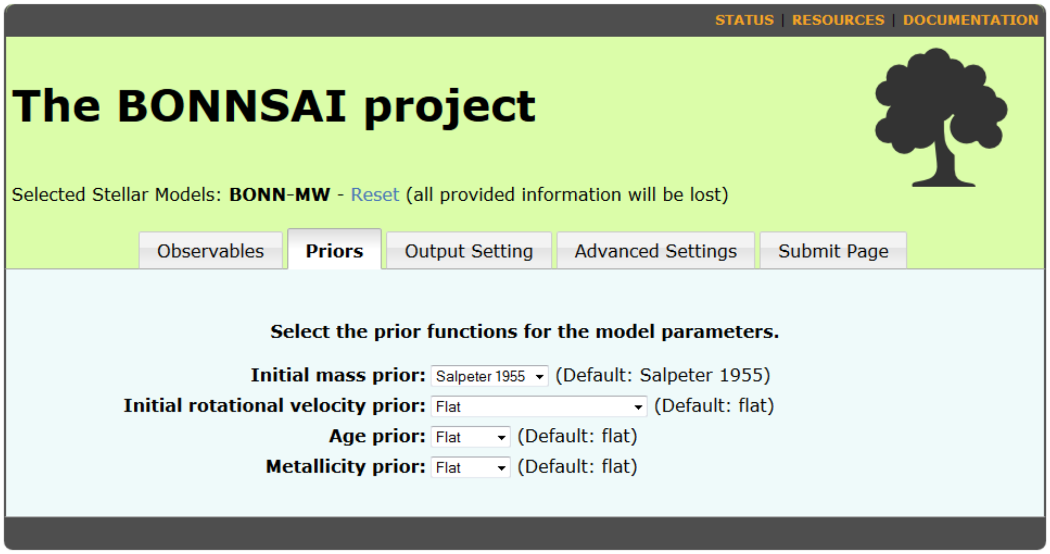

Figure 3: Third step: choosing the priors for the model parameters.

The third step involves choosing the appropriate priors. Priors reflect a priori knowledge on the stellar model parameters, i.e. for single stars on the stellar initial mass, the initial rotational velocity, the age and the chemical composition/metallicity. For each model parameter a prior function has to be chosen which directly goes into the computation of the posterior probabilities through Bayes' theorem (Eq. 1). By a 'flat' prior we denote a prior that attributes every value of a model parameter the same probability, i.e. every possible model parameter is equally probable. So choosing flat priors for all model parameters corresponds to doing a maximum likelihood analysis. For stars, the initial mass function tells us that not each initial mass is equally probable but that lower stellar masses are preferred over higher masses (e.g. Salpeter, 1955). Similarly, not all initial rotational velocities are equally probable. We describe the available priors in the following sections.

2.3.1 Initial mass prior

Flat All initial masses are equally probable.

Power-law A simple power-law function with user-defined slope γ which is often used to express mass functions,

$$p(M_{\text{ini}}) \propto M_{\text{ini}}^{\gamma}. \tag{2}$$

Salpeter A Salpeter initial mass function (Salpeter, 1955), i.e. a power-law (Eq. 2) with fixed slope $\gamma = -2.35$, for all initial masses. Note that this is a good approximation for stars with initial masses $\gtrsim 1\,\mathrm{M}_{\odot}$ (see e.g. Bastian et al., 2010).

2.3.2 Initial rotational velocity prior

Flat All initial rotational velocities are equally probable.

Gaussian Two-piece Gaussian for initial rotational velocities $v_{\rm ini}$ with mean μ and standard deviations σ_+ and σ_-,

$$p(v_{\rm ini};\mu,\sigma) = \frac{2}{\sqrt{2\pi}(\sigma_+ + \sigma_-)} \exp\left[-\frac{1}{2}\left(\frac{v_{\rm ini}-\mu}{\sigma}\right)^2\right], \quad \sigma = \begin{cases}\sigma_+ & \text{for } v_{\rm ini} > \mu \\ \sigma_- & \text{for } v_{\rm ini} \leq \mu\end{cases}.$$

The two-piece Gaussian results in the usual Gaussian for $\sigma_+ = \sigma_-$.

Maxwellian Maxwell-Boltzmann distribution of initial rotational velocities $v_{\rm ini}$ specified by σ (see also Deutsch, 1970),

$$p(v_{\rm ini};\sigma) = \sqrt{\frac{2}{\pi}}\frac{v_{\rm ini}^2}{\sigma^3}\exp\left[-\frac{v_{\rm ini}^2}{2\sigma^2}\right].$$

Tsallis distribution The Tsallis and Kaniadakis distributions are suggested by Carvalho et al. (2009) because they do represent the observed $v\sin i$ distributions of FG-stars. The functional form of the Tsallis distribution, a generalized Maxwellian distribution, is governed by two parameters, q and σ, and results in a Maxwellian for $q \to 1$,

$$p(v_{\rm ini};q,\sigma) = v_{\rm ini}\exp_q\left(-\frac{v_{\rm ini}^2}{\sigma^2}\right), \quad \exp_q(f) = (1+(1-q)f)^{\frac{1}{1-q}}.$$

Kaniadakis distribution The Tsallis and Kaniadakis distributions are suggested by Carvalho et al. (2009) because they do represent the observed $v\sin i$ distributions of FG-stars. The functional form of the Kaniadakis distribution, a generalized Maxwellian distribution, is governed by two parameters, κ and σ, and results in a Maxwellian for $\kappa \to 0$,

$$p(v_{\rm ini};\kappa,\sigma) = v_{\rm ini}\exp_\kappa\left(-\frac{v_{\rm ini}^2}{\sigma^2}\right), \quad \exp_\kappa(f) = \left(\sqrt{1+\kappa^2 f^2}+\kappa f\right)^{\frac{1}{\kappa}}.$$

Hunter et al. 2008 The distribution of present-day rotational velocities as determined by Hunter et al. (2008) are used as the prior for the initial rotational velocities. The observed distributions are approximated by Gaussian functions,

$$p(v_{\rm ini}) = \frac{1}{\sqrt{2\pi}\sigma_v}\exp\left(-\frac{(v_{\rm ini}-\mu_v)^2}{2\sigma_v^2}\right),$$

where the mean rotational velocity μ_v and standard deviation σ_v are different for stars in the Milky Way (MW), the Large Magellanic Cloud (LMC) and the Small Magellanic Cloud (SMC). For the MW and LMC we use $\mu_v = 100\,{\rm km\,s^{-1}}$ and $\sigma_v = 106\,{\rm km\,s^{-1}}$ and for the SMC $\mu_v = 175\,{\rm km\,s^{-1}}$ and $\sigma_v = 106\,{\rm km\,s^{-1}}$. These distributions result in low probabilities for initial non-rotating stars and stars which rotate very fast ($v_{\rm ini} \gtrsim 200$–$300\,{\rm km\,s^{-1}}$).

Ramirez-Agudelo et al. 2013 Functional form of the observed equatorial rotational velocity v_e distribution of LMC O-stars (Ramírez-Agudelo et al., 2013). The functional form is a combination of a Gamma- and Gaussian-distribution — the Gamma distribution represents the low v_e's while the Gaussian represents the higher v_e's. The form is

$$p(v_{\rm ini}) = 0.43g(v_{\rm ini};\alpha=4.82,\beta=1/25) + 0.67N(v_{\rm ini};\mu=205\,{\rm km\,s^{-1}},\sigma=190\,{\rm km\,s^{-1}}),$$

with

$$g(x;\alpha,\beta) = \frac{\beta^\alpha}{\Gamma(\alpha)}x^{\alpha-1}\exp\left(-\beta x\right),$$

where $\Gamma(\alpha)$ is the Gamma function, and

$$N(x;\mu\sigma) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2\right].$$

Dufton et al. 2013 Observed equatorial rotational velocity distribution of LMC B-stars (Dufton et al., 2013). BONNSAI uses the empirical numbers from Table 6 in Dufton et al. (2013).

2.3.3 Age prior

The age prior might be interesting if e.g. the observed star belongs to a star cluster of known age. The age of the star cluster can then be used as prior knowledge of the age of the star. At the moment, we have no other built-in function except for a flat prior.

Flat All ages are equally probable.

Gaussian Two-piece Gaussian for stellar ages τ with mean μ and standard deviations σ_+ and σ_-,

$$p(\tau;\mu,\sigma) = \frac{2}{\sqrt{2\pi}(\sigma_+ + \sigma_-)} \exp\left[-\frac{1}{2}\left(\frac{\tau-\mu}{\sigma}\right)^2\right], \quad \sigma = \begin{cases} \sigma_+ & \text{for } \tau > \mu \\ \sigma_- & \text{for } \tau \leq \mu \end{cases}.$$

The two-piece Gaussian results in the usual Gaussian for $\sigma_+ = \sigma_-$.

2.3.4 Chemical composition/metallicity prior

At the moment, we have no other built-in function except for a flat prior. Furthermore, the currently available stellar models (Tab. 1) are for one metallicity only, i.e. the metallicity is no model parameter. Different sets of stellar models will be included in the future and therefore have this dimension.

Flat All metallicities are equally probable.

2.4 Selecting the output

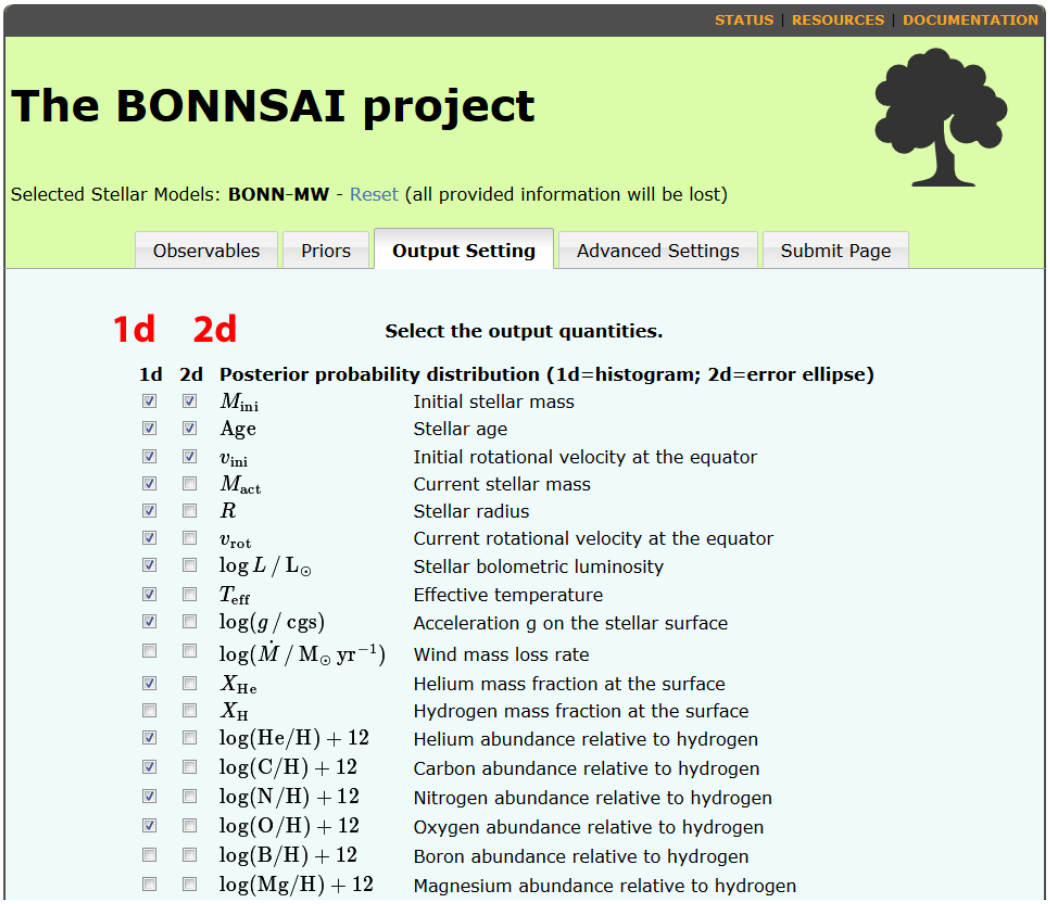

Figure 4: Forth step: selecting the output quantities.

Next, the output has to be specified. By default, several quantities such as initial stellar mass or age are selected. BONNSAI can produce two kinds of output: 1d or 2d posterior probability distributions. The 1d posterior probability distribution essentially is a histogram and the 2d distribution a 2-dimensional map of posterior probabilities. For each 1d distribution, BONNSAI computes the mean, median and mode including 1σ uncertainties. So if you want to know e.g. the stellar radius of your observed star, make sure that the radius is selected from the list of possible output quantities. For 2d posterior distributions, at least two quantities have to be chosen; 2d maps are then computed for each possible combination of the selected quantities, i.e., for the default settings, posterior probability maps of initial stellar mass vs age, initial stellar mass vs initial rotational velocity and age vs. initial rotational velocity.

2.5 Advanced settings

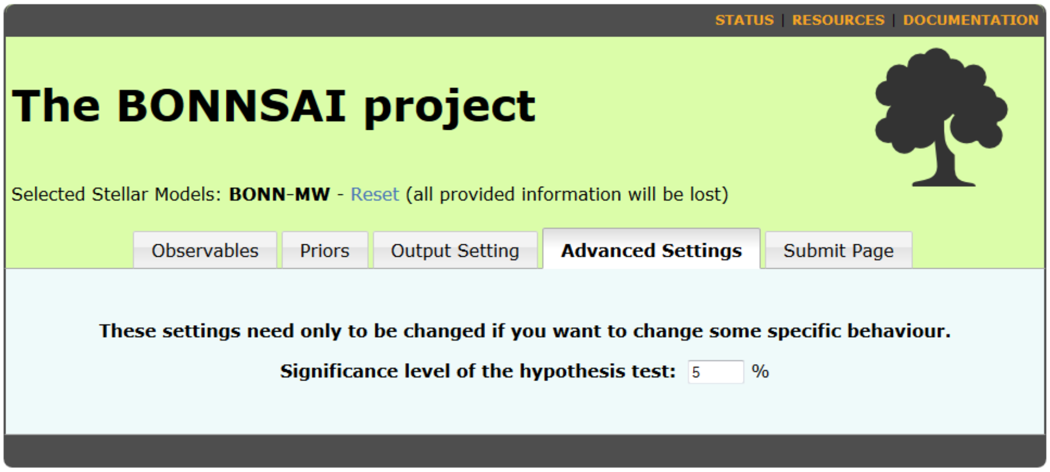

Figure 5: Fifth step: advanced settings.

Currently, there is only one advanced setting. You can set the significance level of the χ^2-hypothesis test and the posterior predictive check (see Sec. 4).

2.6 Submitting

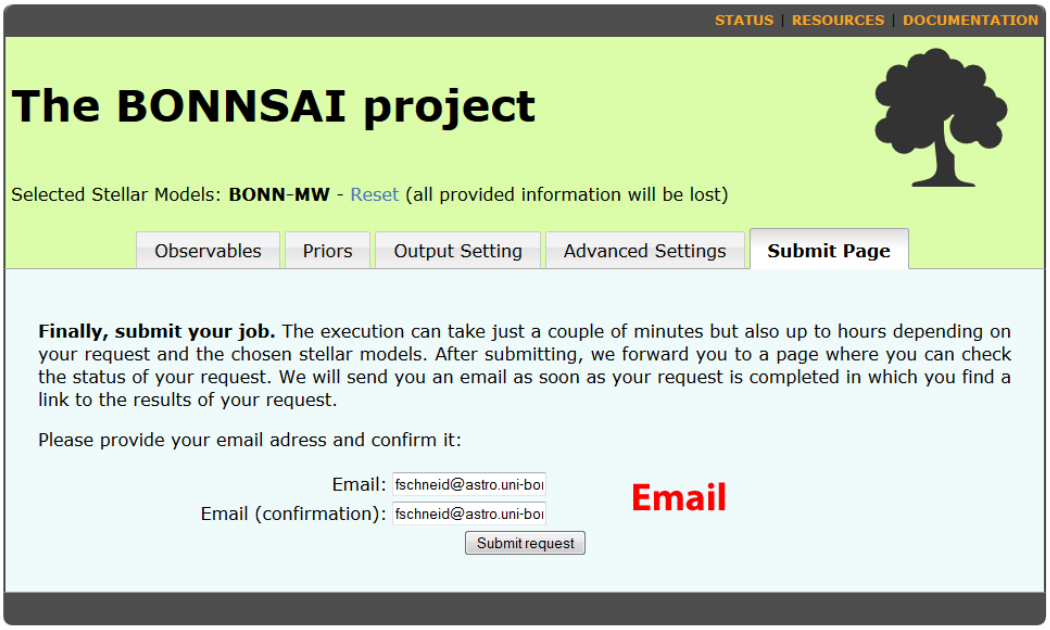

Figure 6: Sixth step: submitting the job.

Once all information are provided, the request can be submitted and will be executed as soon as there is a free computation slot in the internal queuing system of the server. You have to provide your email address because we will send you an email once your request has been finished. After submitting your request you will be forwarded to a status-page where you can see which step is currently being executed and also whether the job is still stuck in the queue.

2.7 The status page

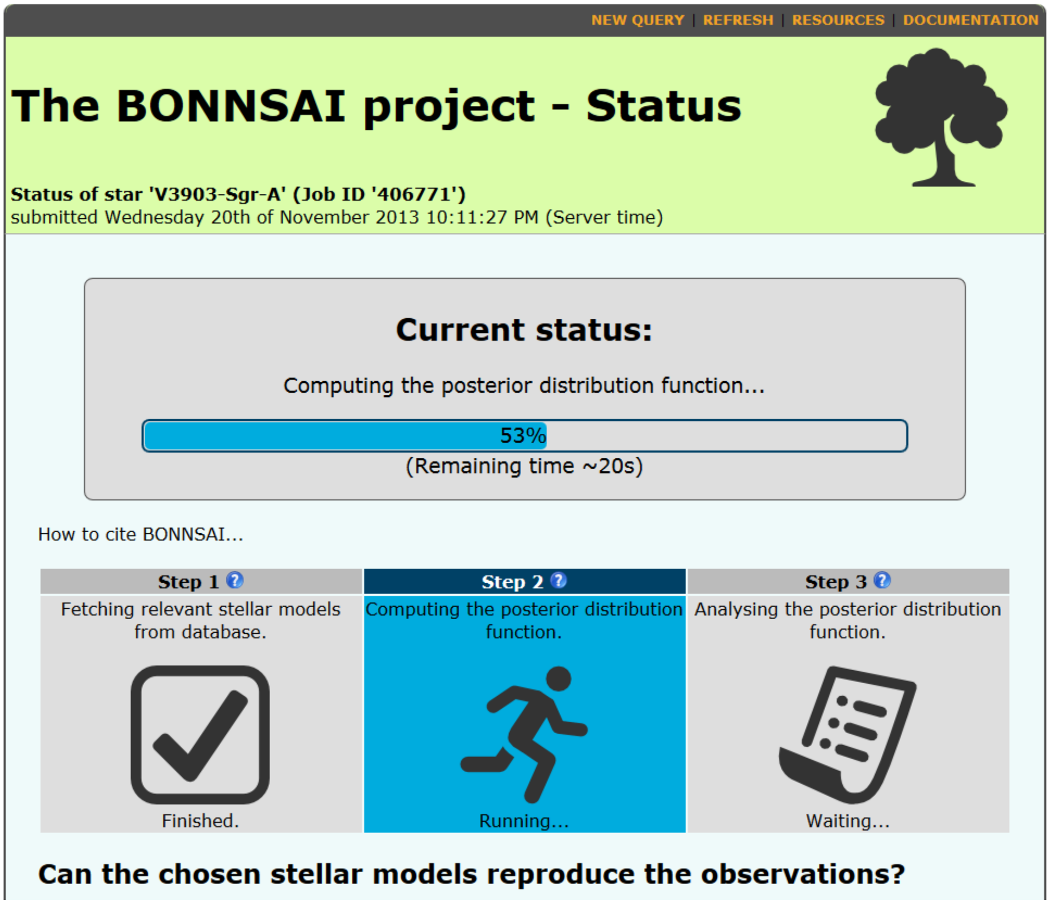

Figure 7: The status page.

Once the job is submitted, the status page of the job is loaded. There are three steps to finish the job (Fig. 7):

1. all stellar models within 5σ of the observables are selected from the database
2. the posterior probabilities are computed
3. the posterior probability distributions are analysed and the output is created

A progress bar indicates the approximate remaining time. The status page can be closed at any time and re-accessed under bonnsai.astro.uni-bonn.de/status.php by typing in the job ID which is written at the top of the page (Fig. 7 and 8). The status page is reloaded every 30 s. When the job is finished, a brief summary of the results is presented. A download link is provided at the top of the page (Fig. 8) if the stellar models can reproduce the observations and if the resolution test is passed (Sec. 4). The density of stellar models around the observation can be displayed by clicking the link. It might happen that the uncertainties of a stellar parameter could not be determined reliably. In such a case, a warning will be displayed and the corresponding posterior probability distribution has to be inspected to judge how trustworthy the mode value and the uncertainties are. At the bottom of the status page, the standard output of BONNSAI is displayed. This is useful for advanced users because it contains all details of a request.

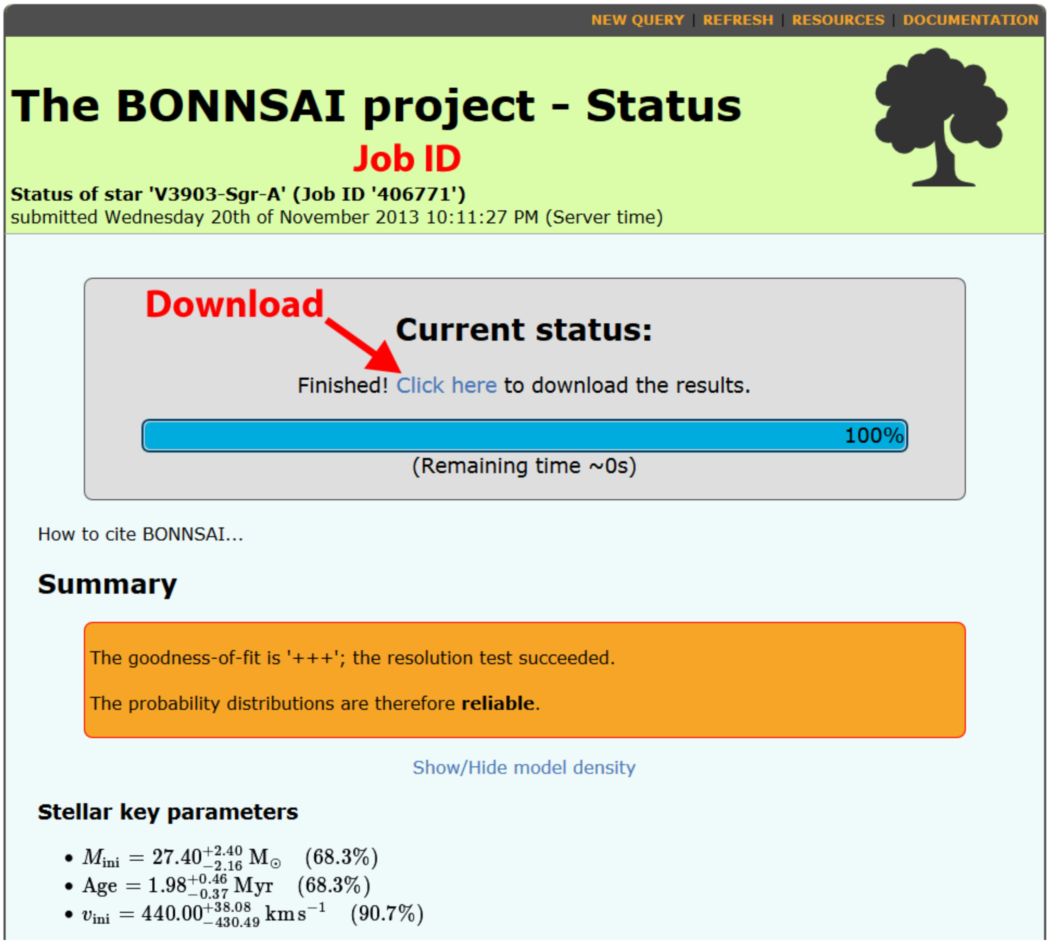

Figure 8: The status page.

3 The output

Once the job is finished and the observed star could be reproduced by the set of stellar models, a link to a zip-file with the output is provided. The zip-file contains three directories (D) and six files (F). An overview of these files and folders is given in Tab. 2 and details follow below.

3.1 Raw data — 'error-ellipses', 'histograms' and 'posterior-predictive-checks'

The raw data in the 'histograms' and 'error-ellipses' directories contain the binned and marginalised posterior probabilities. If you prefer to plot probability densities, you have to divide the provided probabilities by the widths of the bins. The binwidths are written in the header of the files. The data in the 'posterior-predictive-checks' folder contains the probability distribution functions of the differences between the predicted and original observables used as a goodness-of-fit test. A graphical representation of these differences is in the 'posterior-predictive-checks.eps' file.

Table 2: Overview of the BONNSAI output files. Directories are indicated by a 'D' and files by a 'F'.

	Filename	Note
D	error-ellipses	Raw 2d posterior probability distributions
D	histograms	Raw 1d posterior probability distributions
D	plots	Graphical representations of the 1d and 2d posterior probability distributions
F	model_density.dat	Raw data of the stellar model density around the observations
F	model-density.eps	Graphical representation of the stellar model density
F	parameter-space-coverage.png	Graphical representation of the coverage of the observable parameter space with stellar models
F	query.dat	Summary of the query send to and std-output of BONNSAI
F	re-analyse.sh	Shell script to redo all plots
F	results.dat	Summary of the determined parameters

3.2 Graphical representation — 'plots'

For each 1d and 2d posterior probability distribution, a plot is automatically generated. Shown are again the probabilities and not probability densities. Furthermore, the mean, median and mode including their 1σ uncertainties are written to the histogram-plots. The probabilities contained within the uncertainties are provided in parenthesis. Sometimes it is not possible to define 1σ uncertainties in which case the probability in parenthesis is different from the 1σ value of 68.3%. Additionally, the 1, 2 and 3σ ranges are indicated by colours in the 1d representations. Contour lines for the 1, 2 and 3σ ranges can be added to the 2d maps but are disabled by default because of computational costs. But you can enable them when using our plotting routines that can be downloaded from the BONNSAI homepage (see resources page).

3.3 Model density — 'model_density.dat' and 'model-density.eps'

The model density and the cumulative number of models around the observables are provided in raw format in the 'model_density.dat' text-file and plotted in 'model-density.eps'. The model density serves to test whether a observed star can be reproduced by the stellar models and the cumulative number of stellar models served to test the resolution of the stellar models around the observables (for more information see Sec. 4).

3.4 Coverage of parameter space with stellar models — 'parameter-space-coverage.png'

All selected stellar models, i.e. all stellar models within 5σ of the observations, are plotted as little dots in this figure together with the observables. If more than two observables are matched to the stellar models, a projection of all observables into two-dimensional planes is shown. The idea of this plot is to check visually how the stellar models are distributed around the observables to check how good the stellar models can reproduce the observations. Note that these plots can be misleading because they are projections — e.g. if the stellar models are aligned on the surface of a sphere while the observation is located right in the middle of the sphere, the projections would give a wrong impression of the coverage of the parameter space with stellar models. The model density helps out in such a situation.

3.5 Query — 'query.dat'

The full request and the standard output of BONNSAI are captured in the 'query.dat' file. This file can be used to immediately re-run a job.

3.6 Redoing the analysis and plots — 're-analyse.sh'

This shell script will re-start the analysis process of the posterior probabilities and will re-compute confidence intervals and re-plot the data. You need to download the plotting and analysing routines from the resources page to use this script.

3.7 Summary of results — 'results.dat'

Probably one of the most important output files. It contains all results from BONNSAI in a text-file such that they can be further processed. The first block of data contains the result of the resolution test and a summary of the goodness-of-fit (i.e. whether the stellar models reproduce the observations). The next blocks contain the results for each requested output parameter. This data contains a 'flag' which indicates whether there were problems to compute the uncertainties, followed by the dispersion of mean, mode and median. The dispersion is useful to quickly find out how broad the overall distribution of posterior probabilities of a certain stellar parameter is. The next three numbers give the mode of the posterior distribution and the 1σ uncertainties. The last number is again the total probability contained within the uncertainties around the mode. If it is not about 68.3%, the given uncertainties are not 1σ confidence intervals!

The same data as described above, is printed in one row following the data blocks to make it easy to further use the information. The numbers in front of the data in the blocks give the column number of a quantity.

4 Do the stellar models reproduce the observables?

It is a crucial aspect of our method and it is of utmost importance to check the goodness of the fit, i.e. whether a set of stellar models can actually reproduce the observations. There are several reasons why this might not be possible:

- The stellar models do not cover the star because they do not cover the mass of the observed star or the evolutionary stage (e.g. post main-sequence)
- Missing physics in the stellar models like rotation or magnetic fields
- The observed star is actually a post-interaction binary star and is therefore not covered by a set of single stellar models (note however that post binary interaction stars can look like normal single stars — maybe except for chemical peculiarities)
- The observed stellar parameters have problems (e.g. unseen binary companion), were difficult to measure or allow for degenerate solutions
- ...

BONNSAI automatically conducts two tests, a χ^2-test and a posterior predictive check, to ensure that the stellar models do reproduce the observations for a given significance level. We say that the stellar models cannot reproduce the observables if the resolution test is passed but the χ^2-test and/or the posterior predictive check failed for the given significance level.

4.1 χ^2-hypothesis test

In a classical goodness-of-fit test, the χ^2 of the best fitting model is compared to the χ^2-distribution to judge upon the goodness of the best fitting model. We exactly conduct this test and say that the best-fitting model cannot reproduce the observations if the p-value of the χ^2-test is smaller than the given significance level, i.e. if the best-fitting model deviates significantly from the observations.

4.2 Posterior predictive check

In a Bayesian approach, there is more known about the estimated model parameters than just the parameters of the best-fitting model. To make use of this advantage, we conduct a so-called posterior predictive check. The idea is to compare the observables including their uncertainties to what is predicted by the stellar model for the estimated model parameters. If the deviation is "too large", we say that the models cannot reproduce the observations.

In Bayesian statistics, the model predictions are called replicated observables, $\boldsymbol{d}_{\rm rep}$, and are computed from the full posterior probability distribution in the same fashion as we compute the histograms or 2d probability maps of stellar parameters,

$$p(\boldsymbol{d}_{\rm rep}|\boldsymbol{d}) = \int_{\boldsymbol{m}} p(\boldsymbol{d}_{\rm rep}|\boldsymbol{m}) p(\boldsymbol{m}|\boldsymbol{d})\,{\rm d}\boldsymbol{m}. \tag{3}$$

From the likelihood of the observables and the posterior (predictive) probability distributions of the replicated observables (Eq. 3), we compute the probability distribution of the difference between the replicated and the original observables. We say that the stellar models cannot reproduce the observations, if the probability distribution of the difference of replicated and original observables is not compatible with being zero for the given significance level α, i.e. if the probability that the difference is larger/smaller than zero is larger than $1-\alpha$.

4.3 Resolution test

Furthermore we have to test the resolution of the model grid to judge upon the reliability of the determined probability distribution functions. To that end, we select those 10 stellar models that are closest to the best fitting model and compare their average spacing to the observed 1σ uncertainties. We require that the average spacing in each dimension of the parameter space of the observables is smaller than one fifth of the corresponding 1σ uncertainty. If the resolution of the stellar models is too sparse, we do not provide any output. But you might want to consult the BONNSAI team if your observables are extremely accurately known because it can simply happen that the grid of stellar models is too sparse. We have grids with higher resolution and can potentially help you out.

References

Bastian, N., Covey, K. R., & Meyer, M. R. 2010, ARA&A, 48, 339

Brott, I., de Mink, S. E., Cantiello, M., et al. 2011, A&A, 530, A115

Carvalho, J. C., do Nascimento, J. D., Silva, R., & DeMedeiros, J. R. 2009, ApJ, 696, L48

Deutsch, A. J. 1970, in IAU Colloq. 4: Stellar Rotation, ed. A. Slettebak, 207

Dufton, P. L., Langer, N., Dunstall, P. R., et al. 2013, A&A, 550, A109

Hunter, I., Lennon, D. J., Dufton, P. L., et al. 2008, A&A, 479, 541

Ramírez-Agudelo, O. H., Simón-Díaz, S., Sana, H., et al. 2013, A&A, 560, A29

Salpeter, E. E. 1955, ApJ, 121, 161

Torres, G., Andersen, J., & Giménez, A. 2010, A&A Rev., 18, 67