How I learned to stop trusting my X-ray spectral best fits and love nested sampling

A comparison of X-ray spectral fitting techniques

Fabio Acero, FSLAC CNRS/IAC, CEA-Saclay
fabio.acero@cea.fr

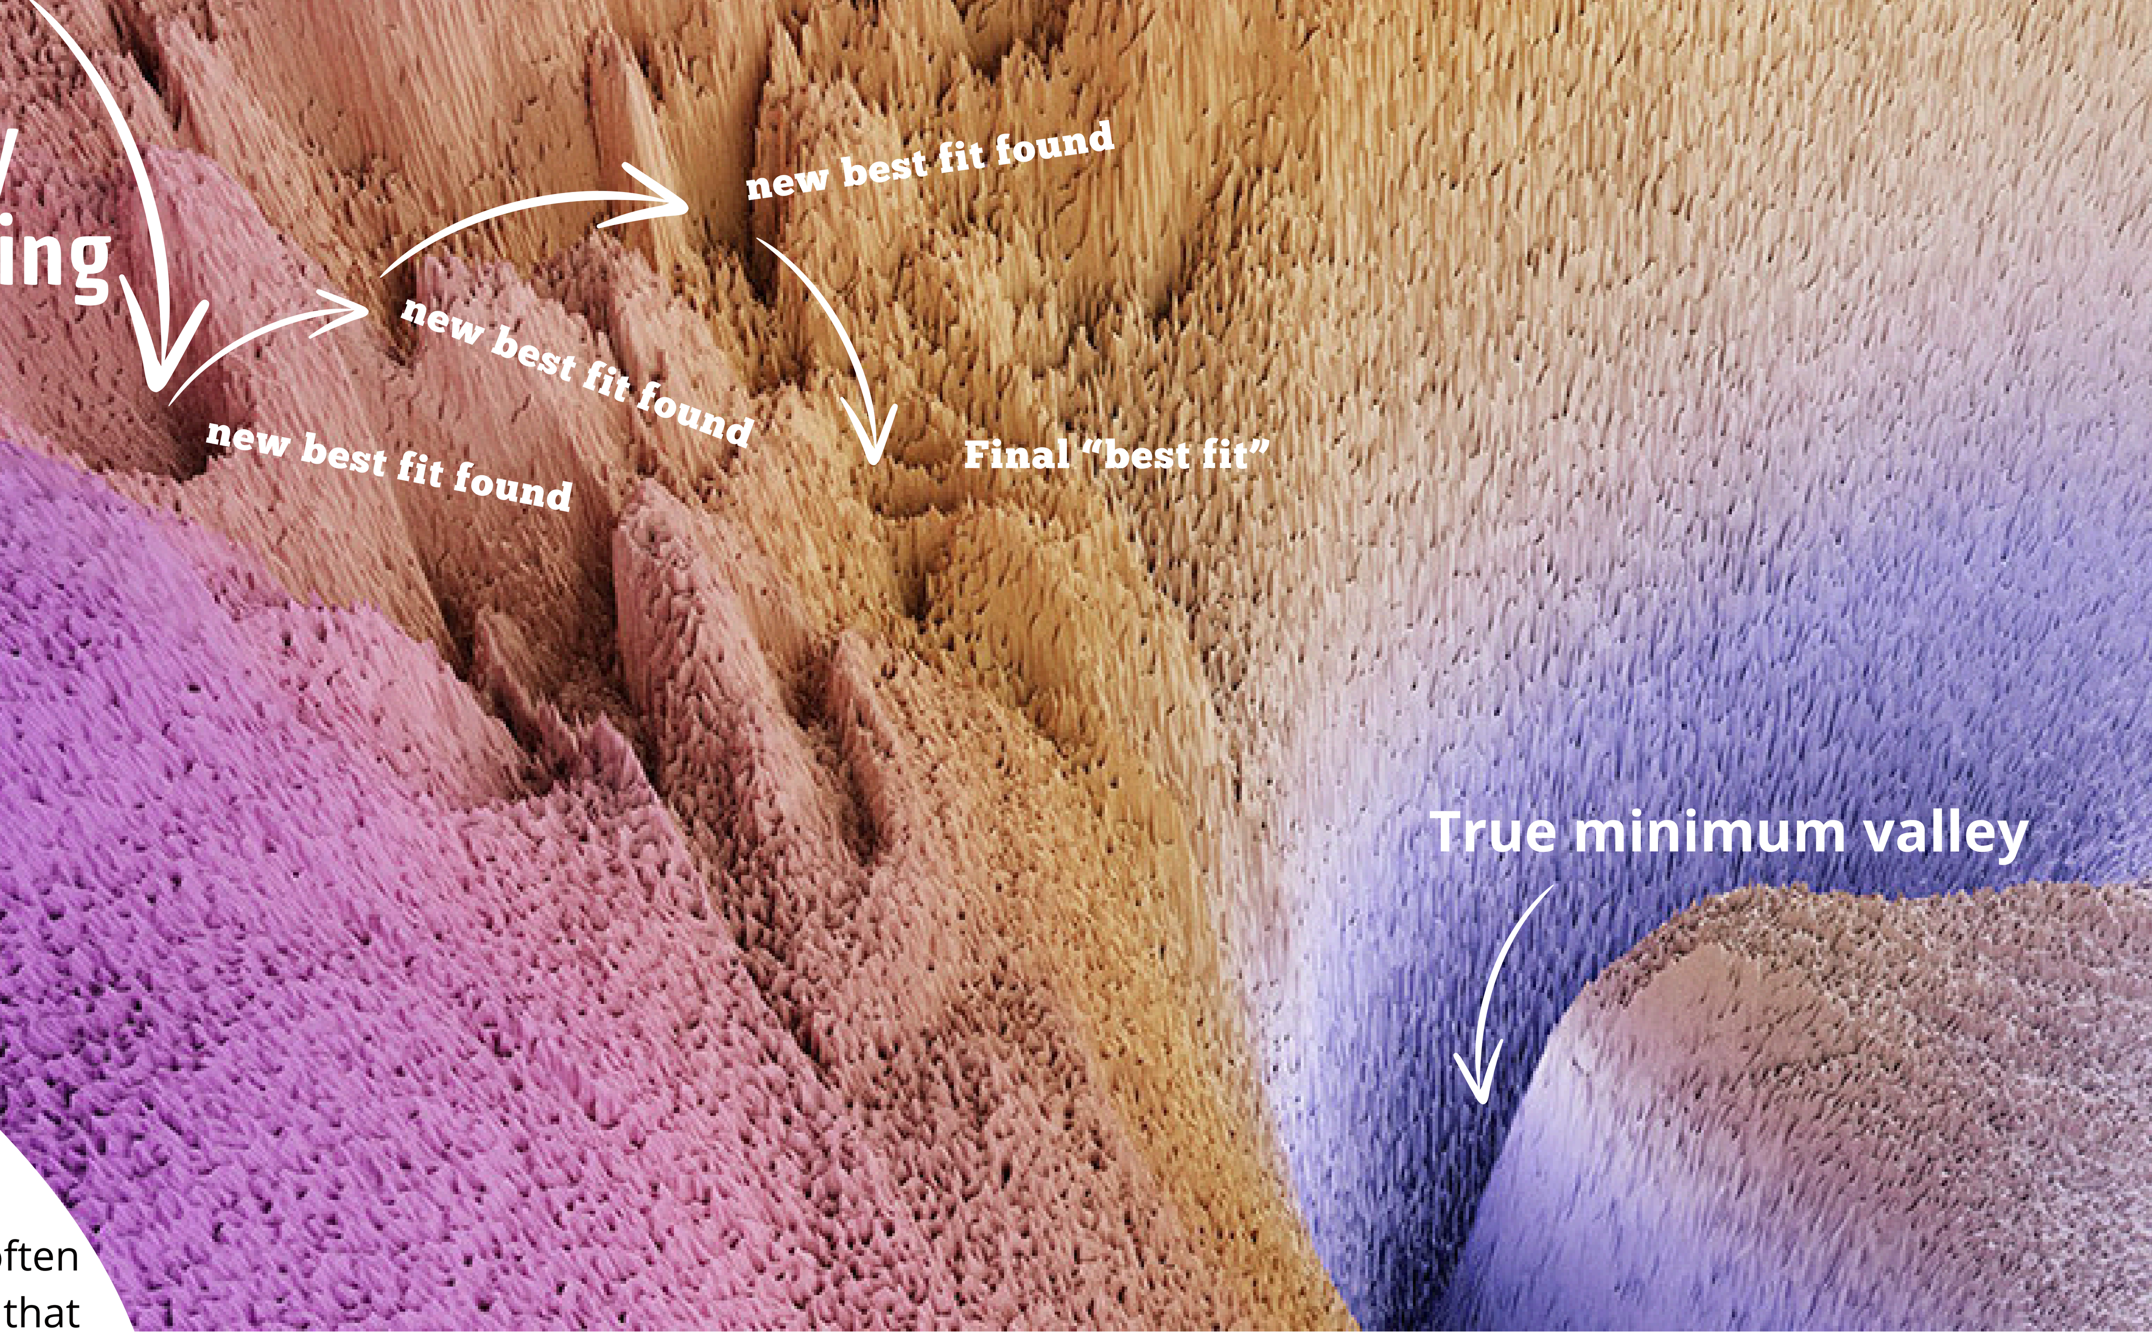

FITTING IS PRONE TO LOCAL MINIMA ...

With more than two decades of X-ray archives, the X-ray space telescopes Chandra and XMM-Newton provide highly detailed spectra of many SNRs with deep observations (> 100h) providing millions of X-ray photons.

The spectral modeling of such exquisite datasets requires complex models, often with a large number of degrees of freedom (N > 10). For such models, we show that the default fitting methods included in Xspec, Sherpa or SPEX (the most used fitting packages) are not always robust to the initialization points and are prone to stop in local minima, strongly limiting the interpretability of the results.

... BUT NESTED SAMPLING CAN HELP

- Nested Sampling (Skilling, 2004) is a Monte Carlo algorithm for estimating an integral over a model parameter space θ, in this case the Bayesian evidence :

$$\mathcal{Z} = \int_{\Omega_\Theta} \mathcal{L}(\Theta)\pi(\Theta)\mathrm{d}\Theta$$

- The ratio of Bayesian evidence Z1/Z2 can be used to select models given the observed data even if models are not nested (e.g. power-law vs thermal).
- The integral is approximated by evolving a collection of live points over nested shells of iso-likelihood contours (like an onion) to approximate the likelihood contributions to the integral. The posterior distributions come as a by product.

... AND OVERCOME MCMC LIMITATIONS

- Nested sampling evaluates the global likelihood landscape and refines sampling towards the best likelihood solutions (does not init walkers from a small volume).
- No need for an initialization point and no concept of burn-in period.
- Robust against multimodal problems and has well defined stopping criteria.
- Provides results for model comparison and to evalutate the posterior distribution.

BENCHMARKING WITH A TOY MODEL

To showcase this effect we have run a small data challenge by making public a fake X-ray spectrum (shown in Fig. 1) of a representative toy model of the Tycho SNR (SN 1572) as observed with Chandra.

This dataset has been circulated among X-ray colleagues, including recognized SNR experts in this field, and the best-fit results have been collected for a comparison with the simulated ground truth and the complete likelihood landscape to get a better insight of traditional fitting method limitations (Fig.2).

Fig.1: simulated toy model spectrum inspired from a Tycho SNR best fit from Godinaud, 2024. The simplified model is :

tbabs*(vnei1+vnei2)

where vnei1 represents the intermediate mass elements emission lines and vnei2 the hotter ejecta with Fe-L and Fe-K emission.

With the temperatures, ionization timescales tau, redshifts, normalizations and free Mg, S, Ar, Ca, Fe abundances, the model has a total of 14 free parameters.

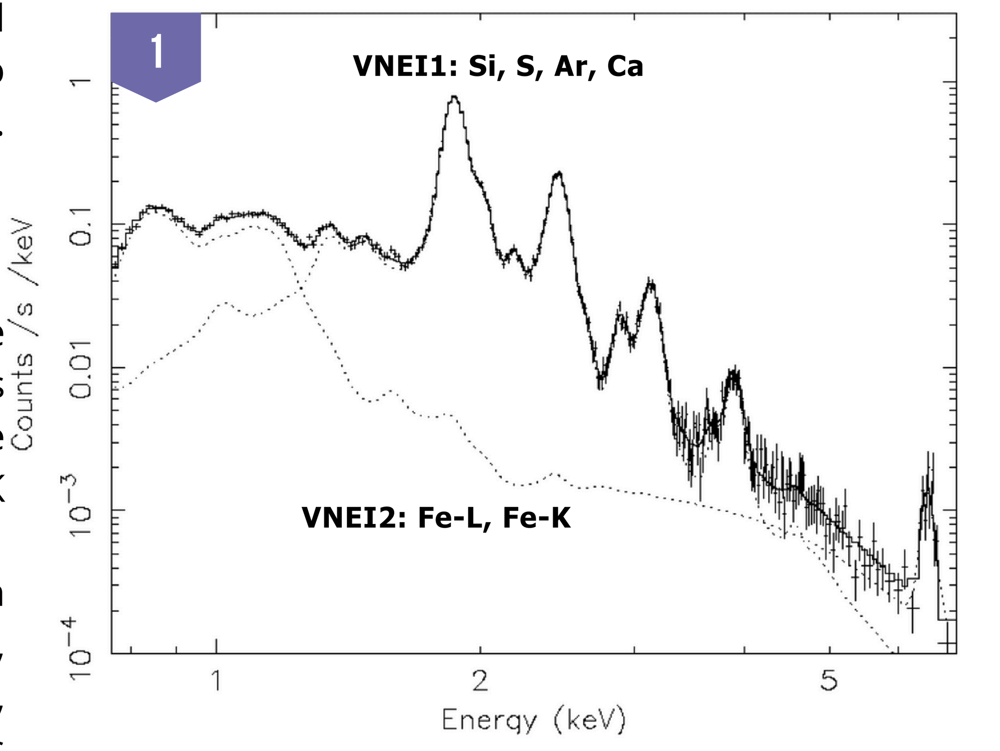

Tycho-like toy model using Chandra responses for 400 ks

MODEL FITTING COMPARISON : XSPEC, MCMC, BXA

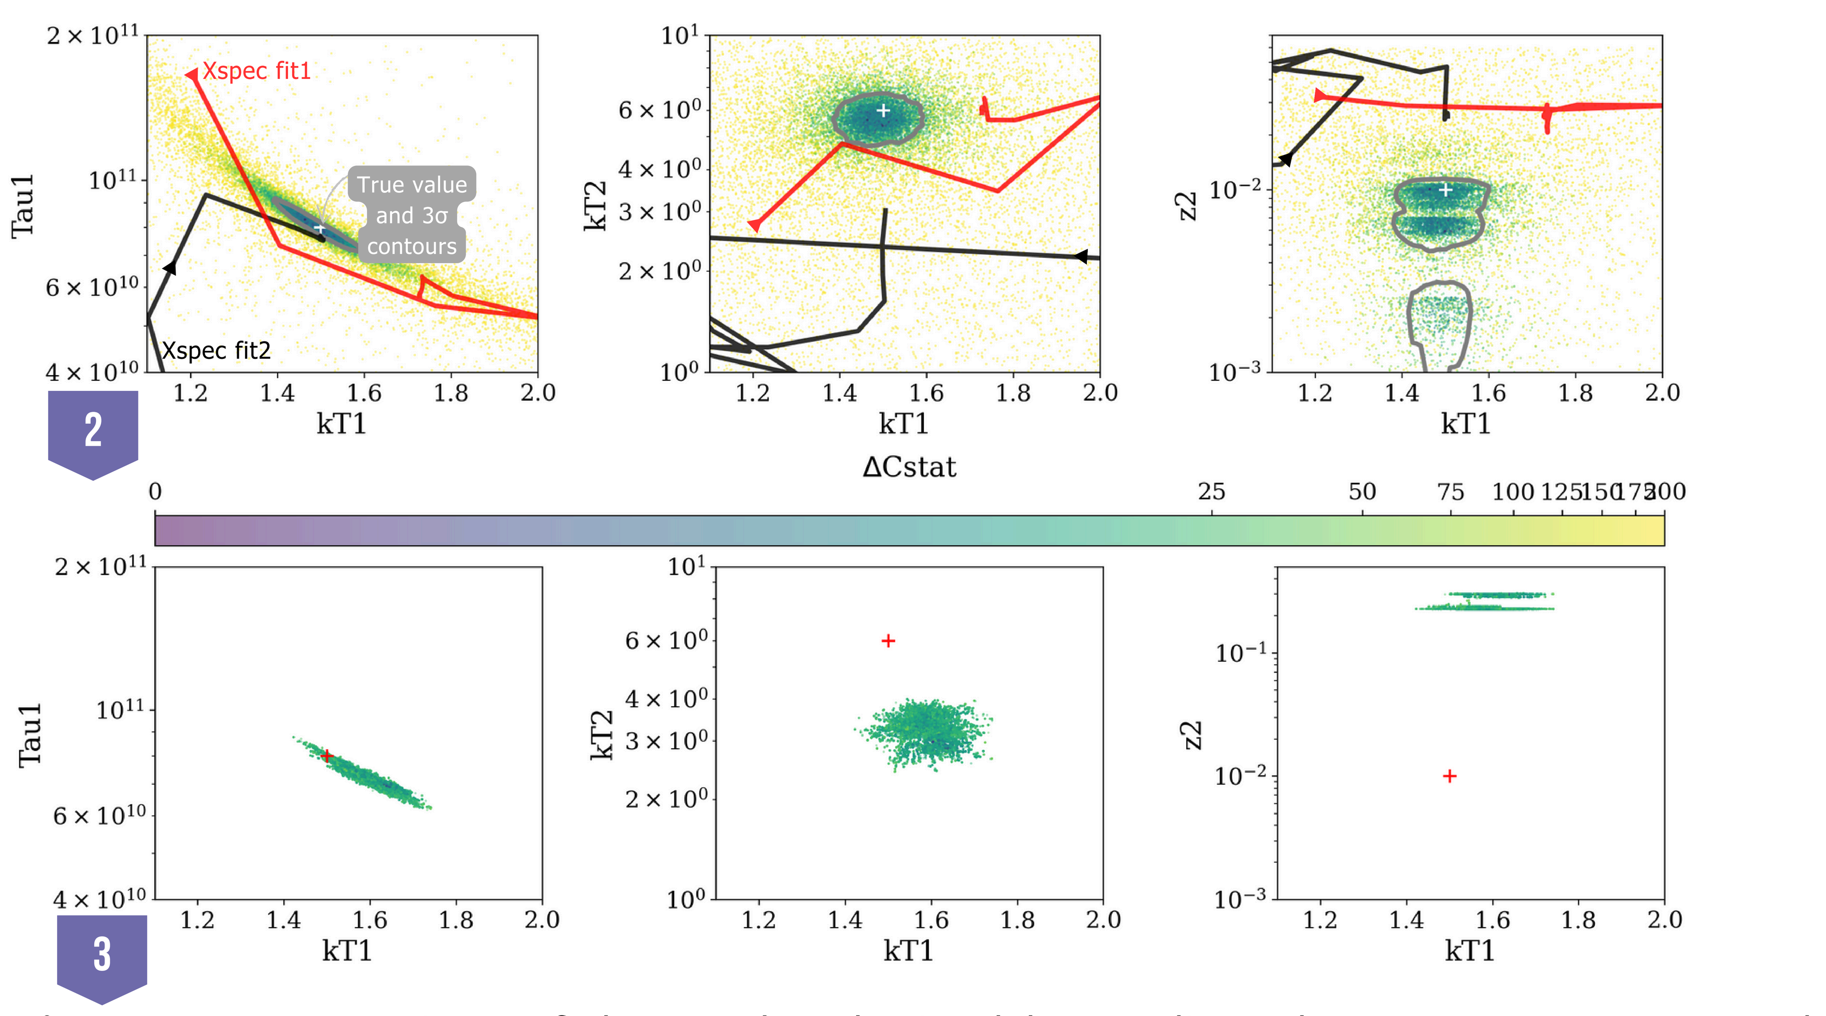

Fig. 2 & 3: comparison of the results obtained by gradient descent (Xspec) & Nested Sampling using BXA (Fig.2) and MCMC in Fig.3 (emcee+Xspec) using the same spectrum shown in Fig. 1. Both panels show some projections of the 14-dimensions likelihood landscape using the ensemble of samples drawn by both methods color coded by Cstat.

Fig. 2: includes the trace of the two Xspec fits (red & black lines) ***starting from different initializations and resulting in two different best-fits, none of them compatible with the true solution (shown by the cross).*** The BXA run took 3h to run on 16 cores.

***Fig.* 3:** MCMC approach. After a fit, the MCMC was run for 10'000 burn-in then 50'000 iterations (40h on a single core). The method was not able to recover the true parameters (show by the red cross). As a sanity check, we tested that initiating the MCMC at the true parameters value, the posterior distribution was similar to the one obtained by BXA.

UNDERSTANDING LIKELIHOOD LANDSCAPE

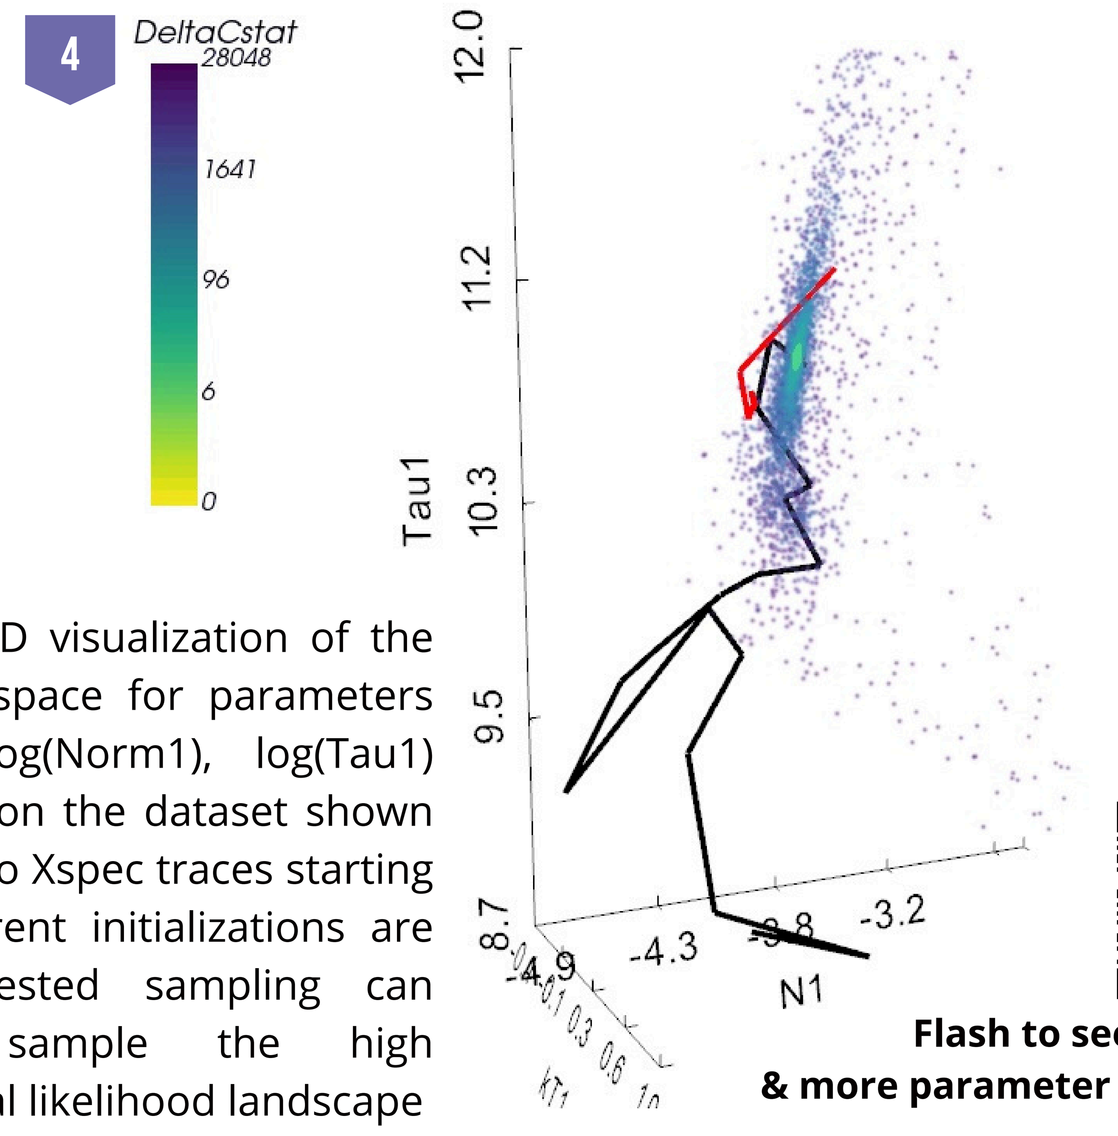

Fig. 4: A 3D visualization of the likelihood space for parameters log(kT1), log(Norm1), log(Tau1) using BXA on the dataset shown in Fig.1. Two Xspec traces starting from different initializations are shown. Nested sampling can sparsely sample the high dimensional likelihood landscape to get a full view of the problem. This is out of reach of MCMC methods which are only local. It provides a unique way to assess the fit difficulties and visualize the parameter degeneracies. Complex landscape (bananas, sharp edges, narrow valleys) are hard to "descend" in Xspec and hard to explore for MCMC.

Flash to see the movie & more parameter projections

WHEN SHOULD I USE NESTED SAMPLING ?

- For complex models with N> 10 parameters and high-statistics spectra.
- To estimate posterior distributions with smaller or similar run time than MCMC with guarantee of robustness and not sensitive to the init point.
- Or simply always as a new method in the tool box to check robustness of the fit, compute errors, perform a model comparison with Bayes factors.
- *In order to have an efficient sampling, it is important to understand some NS parameters. Get in touch via email if you want to discuss about it !*

NESTED SAMPLING REVIEWS: ASHTON 2022, BUCHNER 2023
SNR STUDIES USING BXA: ELLIEN 2023, VINK 2024, GODINAUD 2024, ... AND YOUR NEXT PAPER !