

Figure 1: Avg. Runtime *vs* Avg. Log-likelihood (LL) values of MCMC-PP-Fixed with different intialization strategies and sample sizes (per walker), averaged across 1000 curves generated i.i.d. from the prior. The blue point represents LC-PFN, with the dotted line as a reference for its LL value. The shaded areas correspond to ± 1 SE.

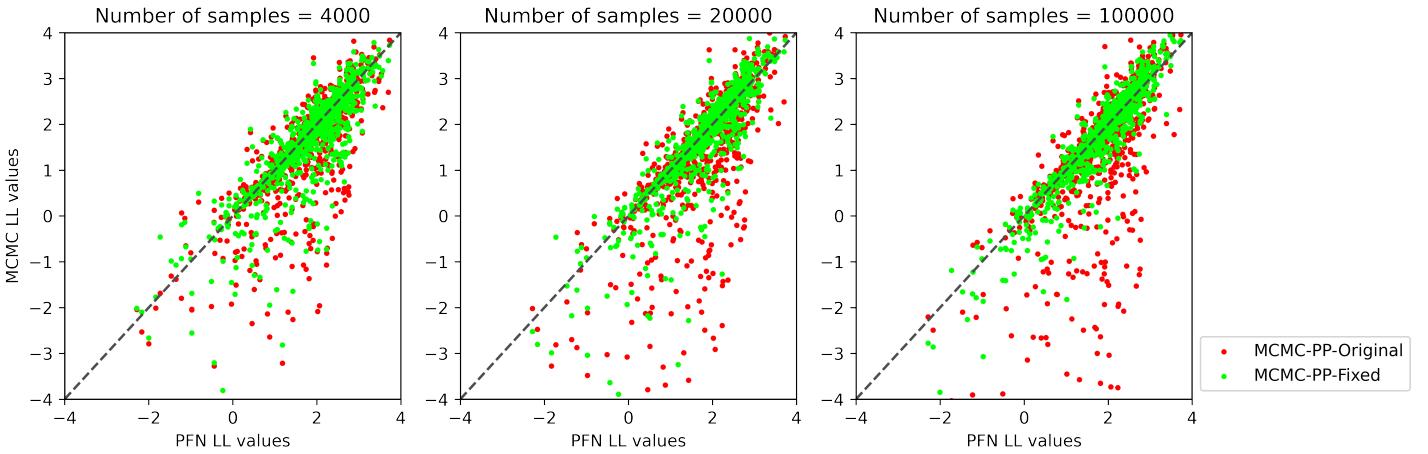


Figure 2: Pairwise comparison of per-curve Log-likelihood (LL) values between LC-PFN and MCMC with different sample sizes and initialization strategies.

Table 1: Comparison of three LC-PFN and MCMC-PP variants on prior curves in terms of log-likelihood. Values in brackets correspond to one standard error.

Method	Parameters	10%	20%	40%	80%	Avg. Runtime
MCMC	nsamples=2000, nwalkers=100, burn_in=500, thin=1	1.53 (0.01)	1.85 (0.01)	2.21 (0.01)	2.42 (0.01)	46.419409 (0.005749)
MCMC	nsamples=4000, nwalkers=100, burn_in=100, thin=1	1.54 (0.01)	1.87 (0.01)	2.21 (0.01)	2.42 (0.01)	84.596940 (0.009415)
MCMC	nsamples=4000, nwalkers=100, burn_in=100, thin=100	1.54 (0.01)	1.87 (0.01)	2.21 (0.01)	2.42 (0.01)	31.981228 (0.009055)
PFN	nb_data=10M, nlayers=3, emsize=128	1.58 (0.01)	1.99 (0.01)	2.28 (0.01)	2.43 (0.01)	0.004115 (0.000006)
PFN	nb_data=10M, nlayers=6, emsize=256	1.72 (0.01)	2.10 (0.01)	2.39 (0.01)	2.51 (0.01)	0.011155 (0.000016)
PFN	nb_data=10M, nlayers=12, emsize=512	1.76 (0.01)	2.13 (0.01)	2.40 (0.01)	2.52 (0.01)	0.050314 (0.000090)