spux

SPUX report for "randomwalk" inference

Jonas Šukys

May 24, 2019

1 Setup

The configurations of all SPUX components and the attached executors are provided below.

Component	Class	Options
Model	Randomwalk	stepsize=1
Likelihood	PF	particles=[4, 128], adaptive=True, accuracy=0.1, margin=0.05, threshold=-5, factor=2, log=1, noresample=0
Sampler	EMCEE	chains=8, a=2.0, attempts=100, reset=10

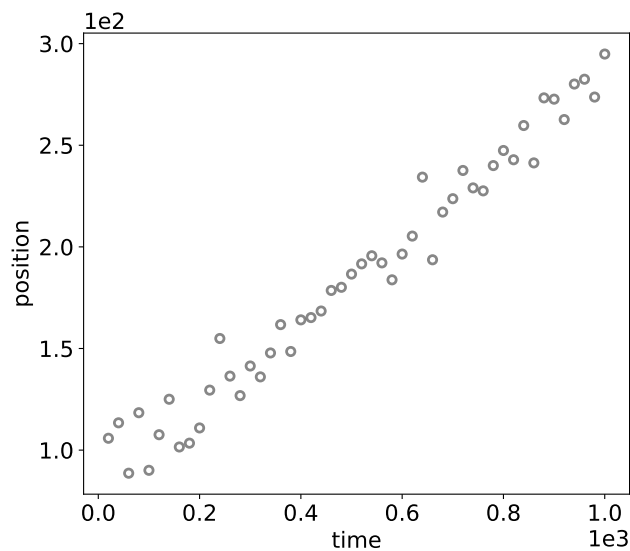Table 1: SPUX components configuration.
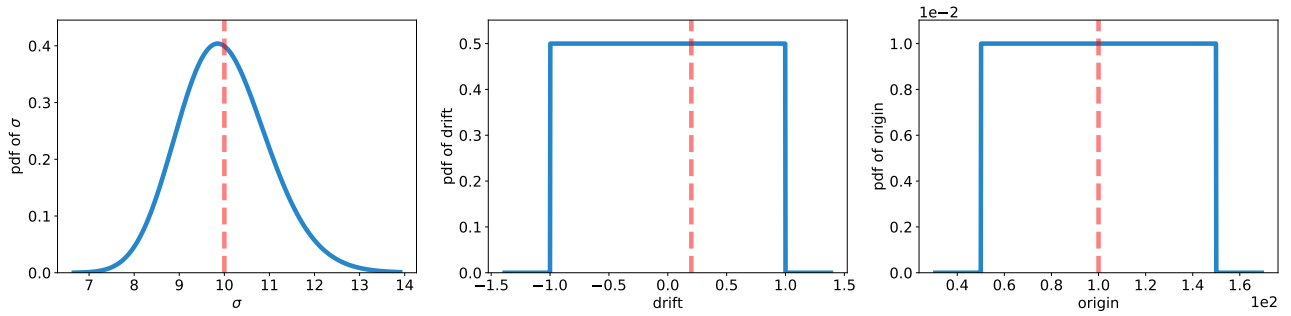

Figure 1: Observational dataset.
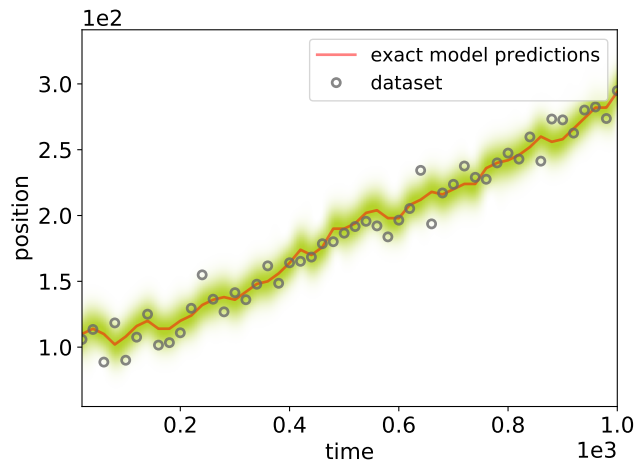
Figure 2: Marginal distributions (prior).

Figure 3: Observational dataset and the associated error model, evaluated using exact model predictions and exact model parameters. The circles (or thick dots) indicate the dataset values, the thick solid line indicates the model predictions used in the error model. The shaded green regions indicate the density of the error model distribution.

2 Results

The results of the inference with SPUX are available below, with the configurations of the attached executors.

Descriptor	Value
Timestamp	2019-05-23 00:58:01
Version	0.3.0
GIT branch (plain)	dev_jonas
GIT revision	f98e24ae81841a98a39d7d9e3b655ae6cdc97d4f

Table 2: Computational environment.

Component	Class	Task	Executor	manager	workers	resources	cumulative
Model	Randomwalk	-	-	0	1	1	1
Likelihood	PF	Randomwalk	Serial	0	1	1	1
Sampler	EMCEE	PF	Serial	0	1	1	1

Table 3: Required computational resources.

seed	thin	lock (batch)	lock (sample)
[8]	None	100	800

Table 4: Setup argument list.

Component	Class	tasks	sizes	cumulative
Model	Randomwalk	1	1	1
Likelihood	PF	128	1	128
Sampler	EMCEE	1K	128	128K

Table 5: Number of model evaluations.

Metric	Value
Maximum A Posteriori (MAP) estimate	batch:119, chain:2, sample:954, log-posterior:-2.01e+02
Multivariate Effective Sample Size (mESS)	not implemented
Multivariate thin period	not implemented
Univariate Effective Sample Size (ESS)	1 - 7 (across chains), with average 2 and sum 22
Univariate thin period	7 - 35 (across chains), with mean 22

Table 6: Metrics for the inference efficiency..

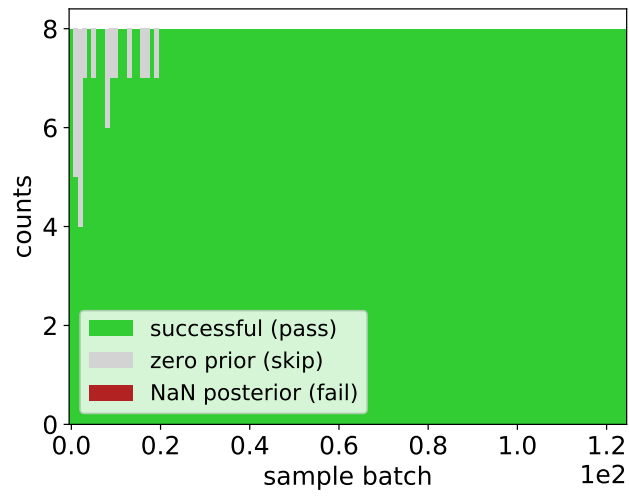| σ	drift	origin
$1.01e+01$	$2.11e-01$	$9.93e+01$

Table 7: Maximum A Posteriori (MAP) estimate parameters.

Figure 4: Diagnostics of the posterior sampler, indicating the successes and failures of the likelihood estmation procedures. Legend: green - successfully passed, gray - estimation skipped due to (numerically) zero prior, red - estimation failure due to failed model simulatios and/or failed PF filtering.
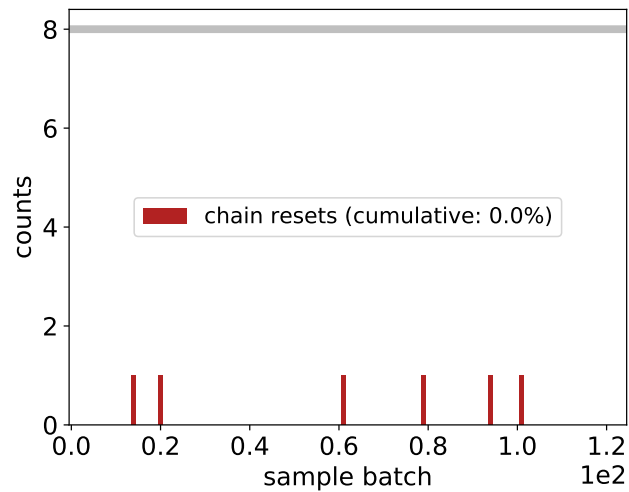

Figure 5: The report for the number of resets (re-estimation of the marginal likelihood) for stuck Markov chains, including the cumulative percentage of resets relative to the total number of samples.
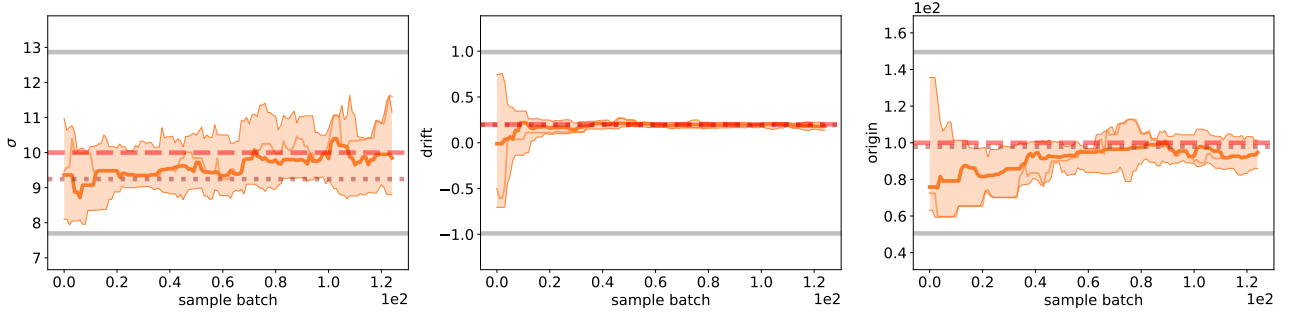
Figure 6: Markov chain parameters samples. The solid lines indicate the median and the semi-transparent spreads indicate the 5% - 95% percentiles accross multiple concurrent chains of the sampler. An auxiliary semi-transparent line indicates an example of such chain. The thick semi-transparent gray lines indicate the interval containing centererd 99% mass of the respective prior distribution. The brown dotted line indicates the estimated maximum a postriori (MAP) parameters values. The red dashed line represents the exact parameter values.
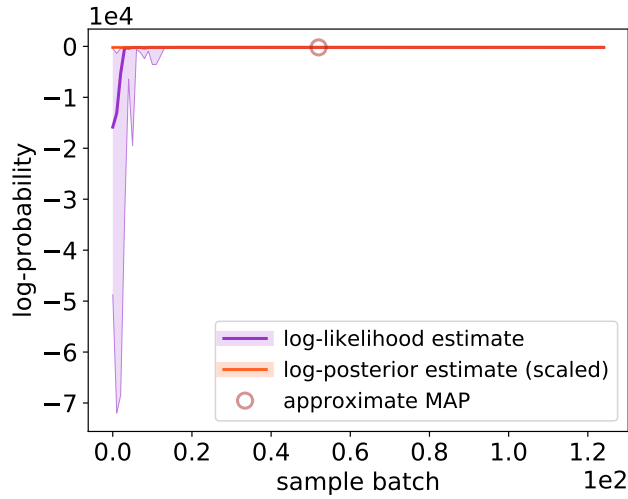

Figure 7: Log-likelihood and (scaled [TODO: estimate evidence and remove scaling]) log-posterior estimates for the sampled model posterior parameters. The solid lines indicate the median and the semi-transparent spreads indicate the 10% - 90% percentiles accross multiple concurrent chains of the sampler. For log-likelihood, the estimates from the rejected proposed parameters are also taken into account. The brown "o" symbol indicates the posterior estimate at the approximate MAP parameters.
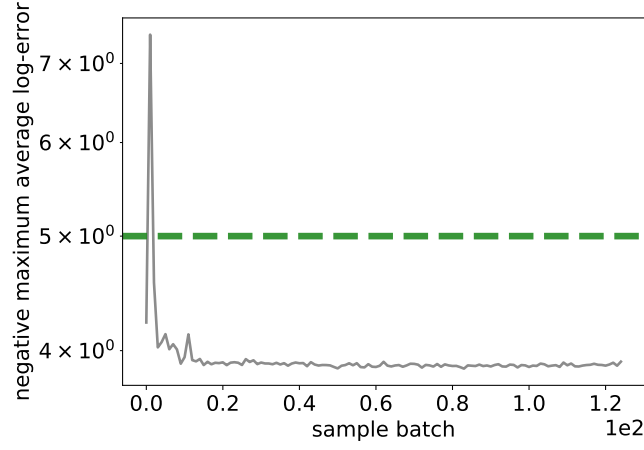
Figure 8: Maximums of the average (over dataset snapshots) marginal observational log-errors across multiple concurrent chains of the sampler. The dashed green line indicates the threshold set in the adaptive PF likelihood.
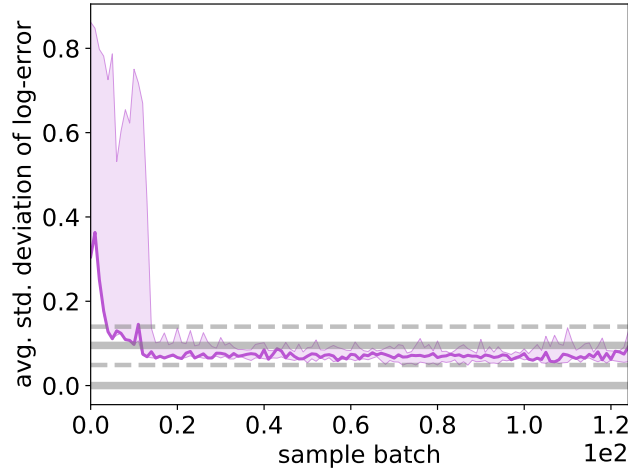

Figure 9: Average (over dataset snapshots) standard deviations for the estimated marginal observational log-error using the PF. The semi-transparent spread indicates the range (minimum and maximumum) accross multiple concurrent chains of the sampler and the solid line indicates the value of the chain with the largest estimated log-likelihood. The solid thick gray line above the same line for a zero value reference indicates the specified accuracy and the dashed thick gray lines indicate the specified margins - all specified within the adaptive PF likelihood.
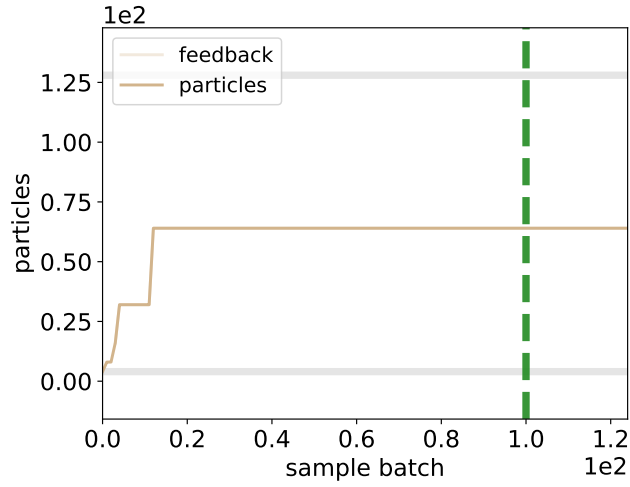
Figure 10: The adaptavity of the number of particles in the PF likelihood. The brighter line indicates the feedback (recommendation) of the adaptation algorithm, and the darker line indicates the actual number of used particles. The semi-transparent thick gray lines indicate the limits for minimum and the maximum number of allowed particles in the PF likelihood.
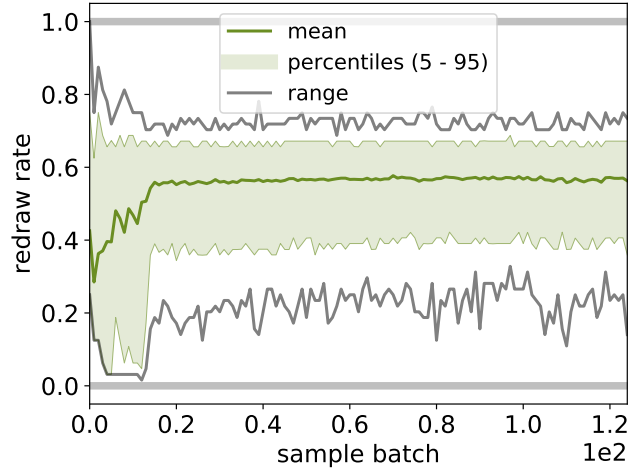

Figure 11: Particle redraw rates (the fraction of surviving particles) in the PF likelihood estimator. The solid line indicates the mean, the semi-transparent spreads indicate the 5% - 95% percentiles, and the dotted lines indicate the range (minimum and maximum) accross multiple concurrent chains of the sampler.
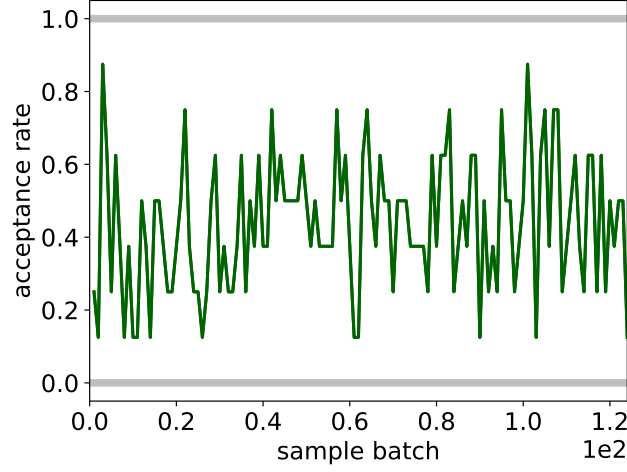
Figure 12: Acceptance rate (accross multiple concurrent chains of the sampler) for the proposed parameters samples.
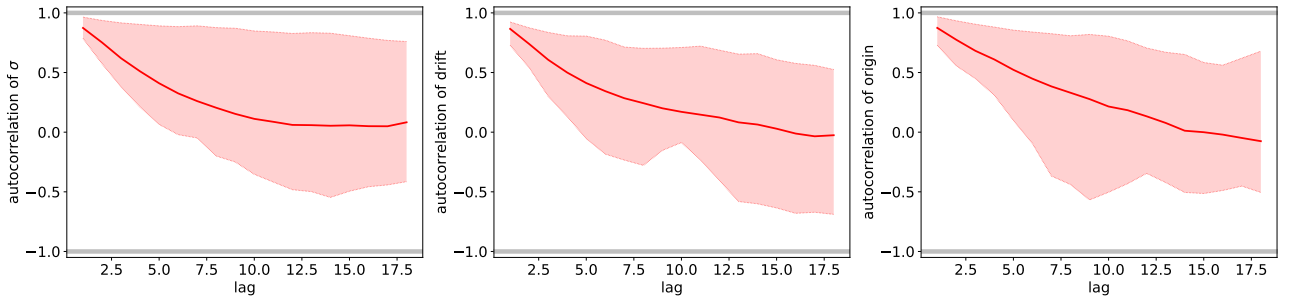

Figure 13: Autocorrelations of Markov chain parameters samples. The solid lines indicate the mean and the semi-transparent spreads indicate the range (minimum and maximum) accross multiple concurrent chains of the sampler.
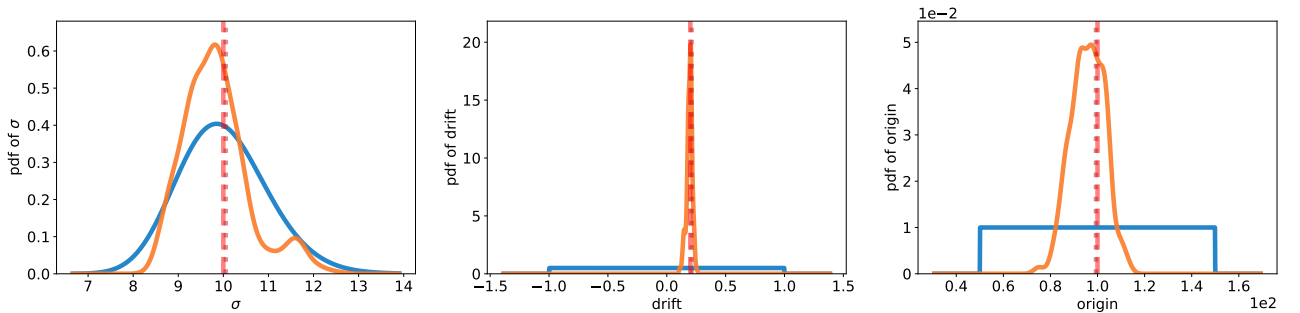

Figure 14: Marginal posterior (orange) and prior (blue) distributions of model parameters. The brown dotted line indicates the estimated maximum a postriori (MAP) parameters values. The red dashed line represents the exact parameter values.
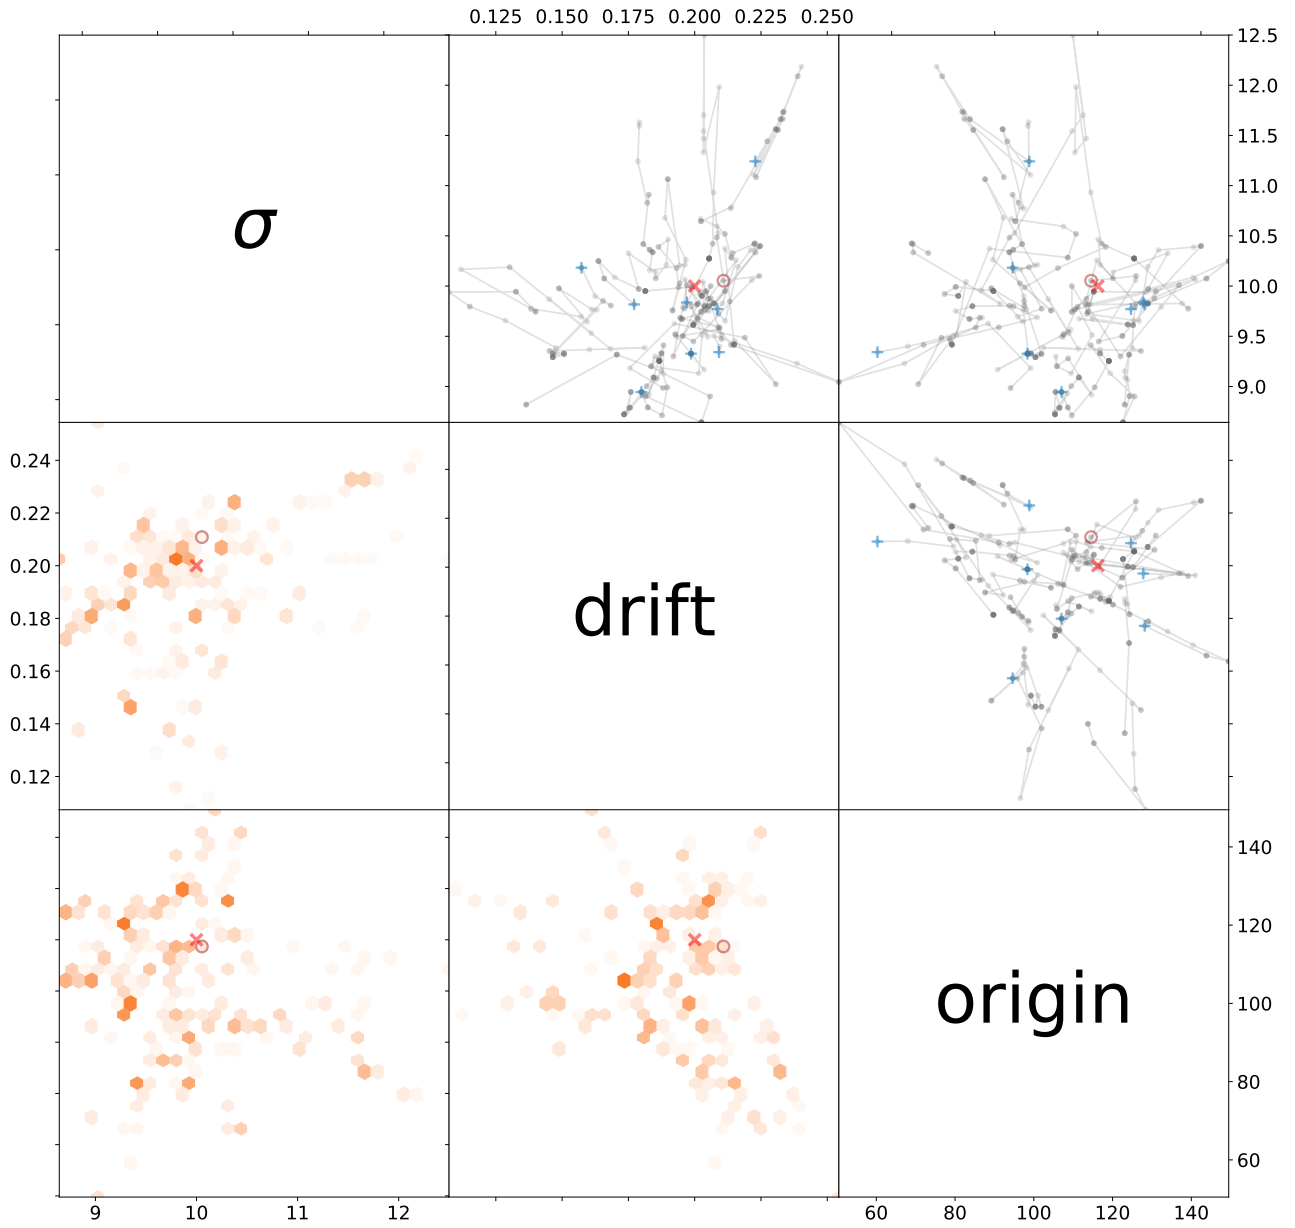
Figure 15: Joint pairwise marginal posterior distribution of all model parameters, including the corresponding Markov chains from the sampler. Legend: thick semi-transparent gray lines - intervals containing centererd 99% mass of the respective prior distribution, blue "+" - initial parameters, brown "o" - approximate MAP parameters, red "x" - the exact parameters, thin semi-transparent gray lines and dots - concurrent chains, orange hexagons - histogram of the joint pairwise marginal posterior parameters samples.
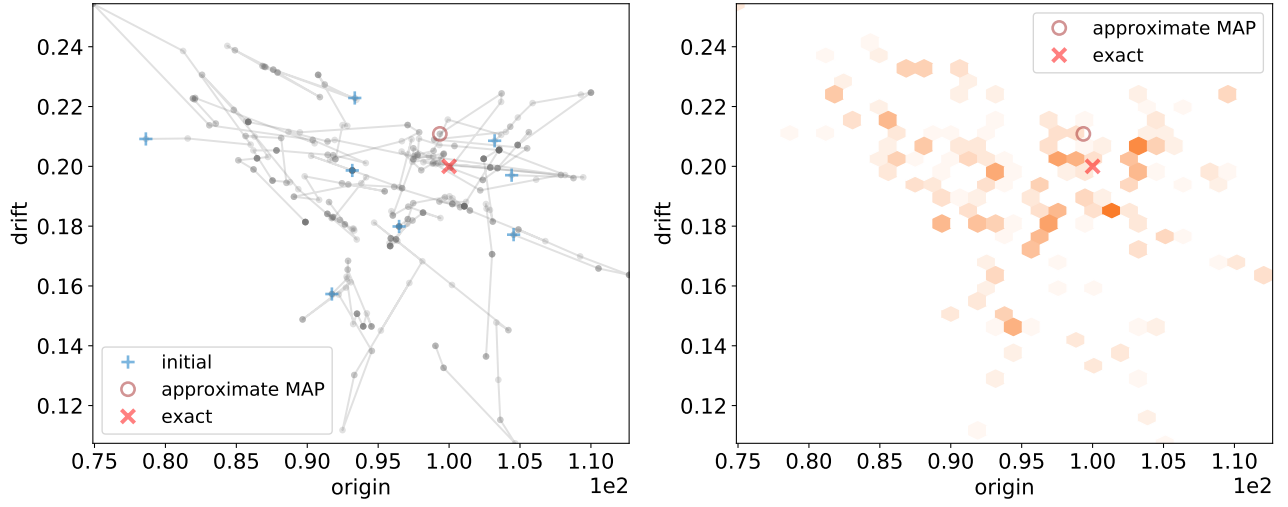
Figure 16: Joint pairwise marginal posterior distribution of origin and drift, including the corresponding Markov chains from the sampler. Legend: thin semi-transparent gray lines and dots - concurrent chains, orange hexagons - histogram of the joint pairwise marginal posterior parameters samples, blue "+" - initial parameters, brown "o" - approximate MAP parameters, red "x" - the exact parameters.
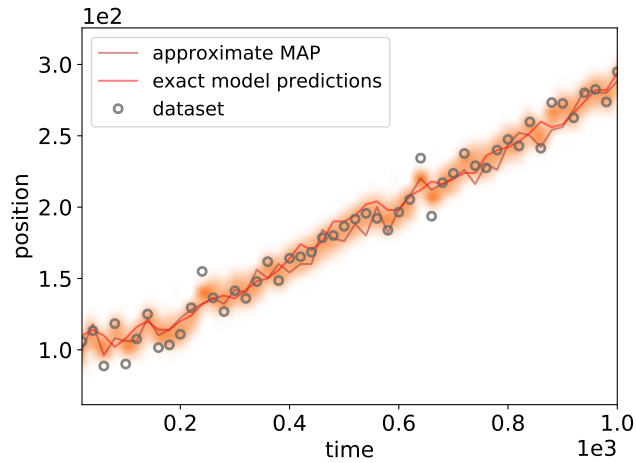

Figure 17: Posterior distribution of model predictions for the observational dataset. The shaded orange regions indicate the log-density of the posterior model predictions distribution at the respective time points, the brown line indicates the approximate MAP model prediction., the red line represents the exact model prediction values.
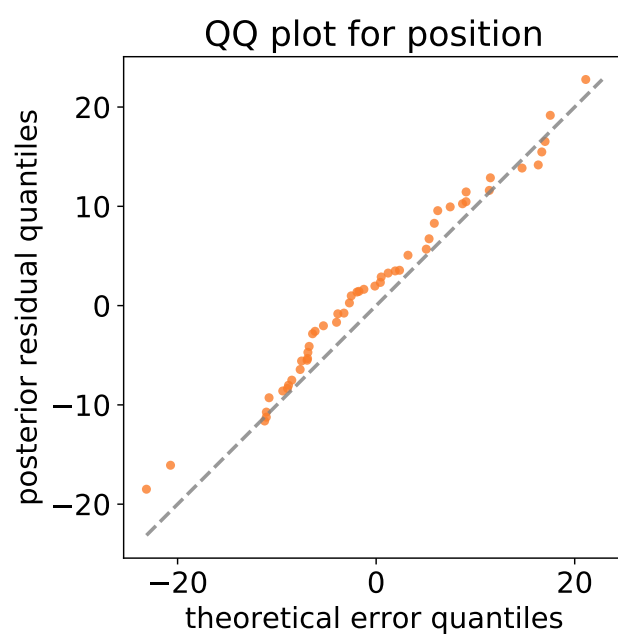
10

Figure 18: Quantile-quantile distribution comparison between the prediction residuals and the specified error model.

3 Performance

The performance metrics of the SPUX inference are available in this section.

Timer	Value
Total runtime (excl. checkpointer)	2 hours 43 minutes 2 seconds
Average runtime per sample	0 hours 0 minutes 10 seconds
Equivalent serial runtime	2 hours 43 minutes 2 seconds

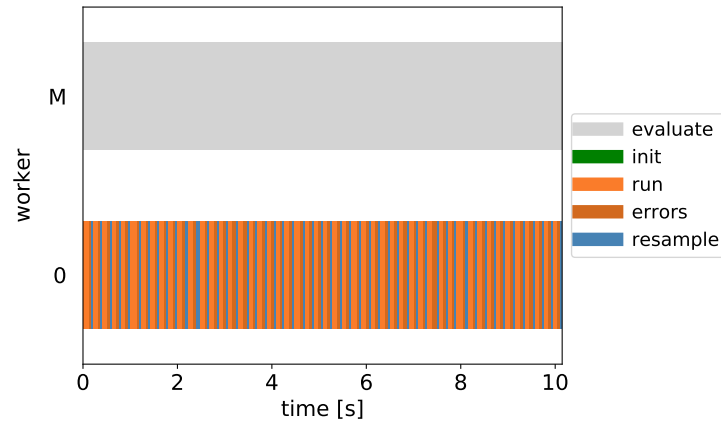Table 8: Runtimes (excl. checkpointer).

Figure 19: Timestamps of key methods within a single estimation of the PF likelihood across all parallel workers.
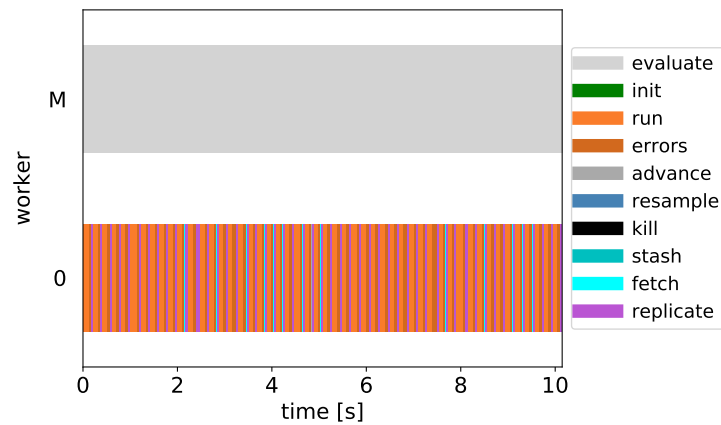

Figure 20: Timestamps of key methods within a single estimation of the PF likelihood across all parallel workers.

4 Miscellaneous

In this section we include all tables and figures found in the respective directories but not yet assigned to any of the sections above.

Component	Class	Fields	Iterators for infos
Sampler	EMCEE	proposes, likelihoods, resets, timing, index, priors, infos, feedback, feedbacks, timings, accepts, parameters, posteriors	0 - 8
Likelihood	PF	variances, avg_deviation, particles, MAP, timing, variance, errors_prior, estimates, weights, successful, traffic, sources, redraw, timings, predictions, predictions_prior	-

Table 9: SPUX infos structure.

No figures found.