Bayesian Inference for Optimal Transport with Stochastic Cost
Supplementary

1. Hamiltonian Monte Carlo Diagnostics

We present some diagnostics for the HMC sampling carried out on a toy example. In the experiment, we place the two measures μ and ν, both with n atoms, on the circle. See Fig. 1. Then, we compute $C(x, y) = \|x-y\|+\sigma^2$, where $\sigma \sim \mathcal{N}(0, 1)$ is a noise variable for the cost. Varying the amount of cost matrices sampled, we then run the HMC sampling with 1000 warm-up iterations and sample 5000 elements from the posterior. We then compute the $\hat{r}$ (Fig. 2) and N_{eff} (Fig. 3) statistics, which are standard diagnostics of the chain (Gelman et al., 2013).

The $\hat{r}$ value describes the convergence of the chain. By definition, the value is higher that 1. A standard practice is to accept the convergence of the chain, if $\hat{r} \leq 1.05$, which we still observe when we have $n = 10$ atoms per each measure in Fig. 2. However, larger problems would require increasing the length of the chain to yield an acceptable $\hat{r}$ value.

The N_{eff} describes the effective number of simulation draws, i.e., approximately the number of independent simulation draws. As can be seen in Fig. 3, this value decreases rapidly as the problem size is increased, reflecting on the difficulty of MCMC sampling in high dimensions.

References

Andrew Gelman, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin. *Bayesian Data Analysis*. CRC press, 2013.

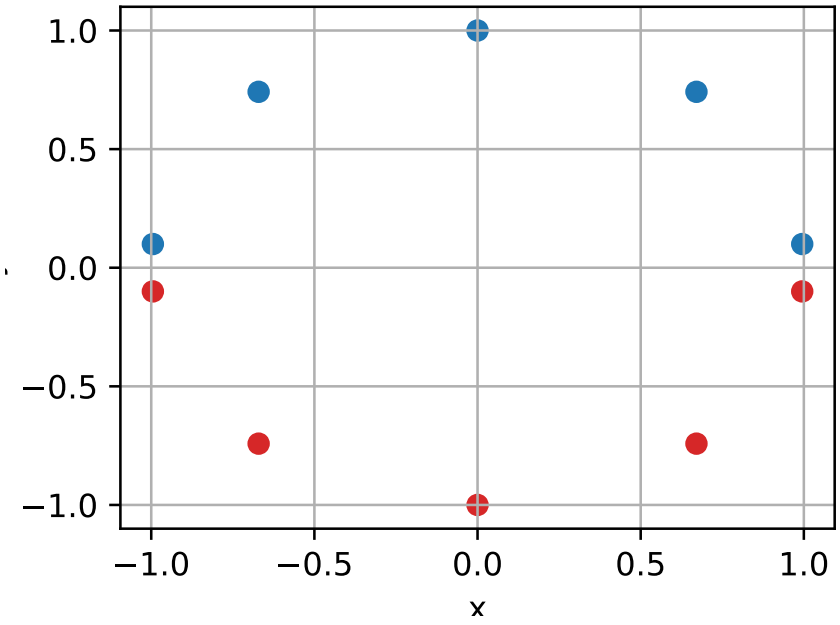

Figure 1: Simulation study setting, where μ (blue) and ν (red) with uniform distributions are placed on the sphere.

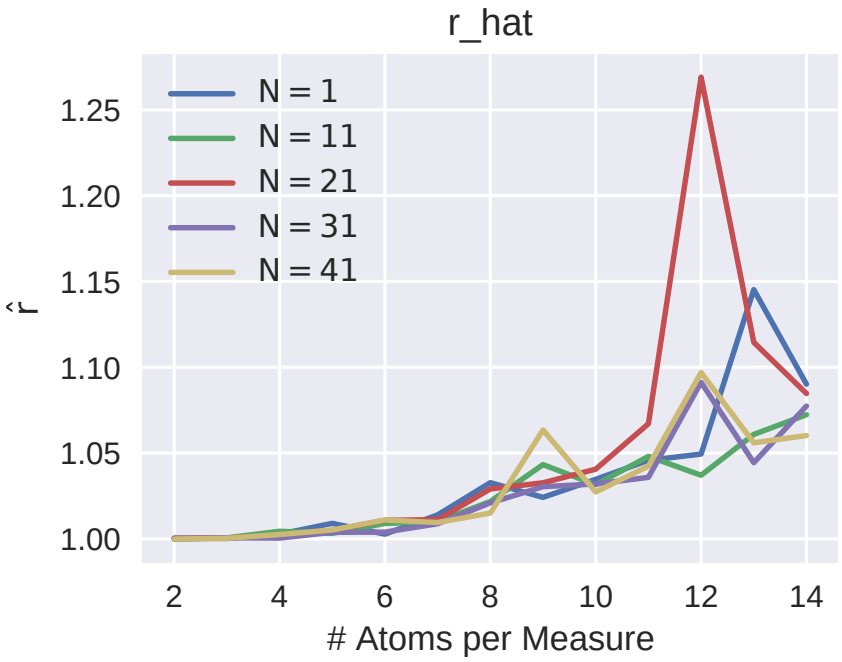

Figure 2: $\hat{r}$ as a function of the size of the problem. For example, if both measures have n atoms, then HMC is carried out in a $(n-1)^2$ dimensional space. N describes the amount of cost matrices sampled.

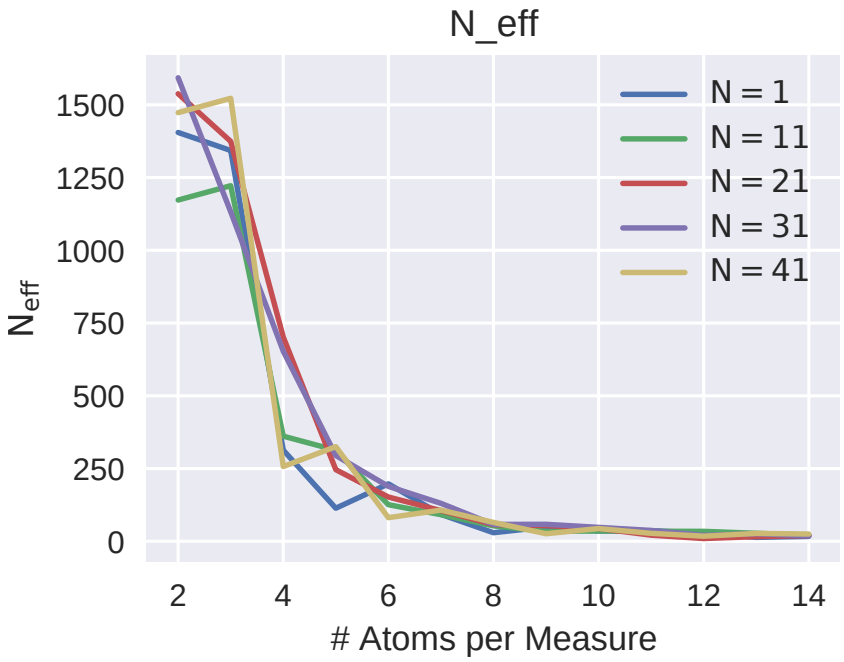

Figure 3: N_{eff} as a function of the size of the problem. For example, if both measures have n atoms, then HMC is carried out in a $(n-1)^2$ dimensional space. N describes the amount of cost matrices sampled.