# Non-reversibly updating a uniform $[0,1]$ value for accept / reject decisions

Radford M. Neal, University of Toronto, Vector Institute Affiliate

## Summary

Metropolis accept/reject decisions can be done by comparing a probability ratio with a uniform $[0,1]$ value. This value can be updated non-reversibly, as part of the Markov chain state, rather than being sampled independently each iteration. For Langevin updates with persistent momentum, non-reversible updates cause rejections to be clustered together, with long stretches of acceptances, which improves sampling by reducing random walk behaviour. A small benefit for simple random-walk Metropolis updates in high dimensions is also seen.

## Metropolis and Persistent Langevin Methods

Simple random-walk Metropolis:

1) Propose $x^* \sim N(x, \sigma^2 I)$.
2) Accept $x' = x^*$ as next state if $u < \frac{\pi(x^*)}{\pi(x)}$, with $u \sim U(0,1)$; otherwise let $x' = x$.

Langevin with persistent momentum (Horowitz):

1) Set $p' = \alpha p + \sqrt{1-\alpha^2}\, n$, where $n \sim N(0,I)$, and $x' = x$.
2) Find $(x^*, p^*)$ from $(x', p')$ with one "leapfrog" step (as in HMC), with stepsize $\eta$.
3) Accept $(x'', p'') = (x^*, -p^*)$ if $u < \frac{\pi(x^*, -p^*)}{\pi(x', p')}$ with $u \sim U(0,1)$; otherwise $(x'', p'') = (x', p')$.
4) Let $p''' = -p''$ and $x''' = x''$.

If step (3) accepts, the negations of $p$ in steps (3) and (4) cancel, but rejections reverse the motion.

## Non-reversibly updating $u$

We can make $u \sim U(0,1)$ be part of the Markov chain state, independent of $x$. Or we can represent $u$ as $u = |v|$ with $v \sim U(-1,+1)$, independent of $x$, or as $u = s/\pi(x)$ with $s|x \sim U(0, \pi(x))$.

We can update $v$ non-reversibly as follows:

$$\begin{aligned} & v \;\leftarrow\; v + \delta + \text{noise} \\ & \text{while } v > +1 \text{ do } \; v \;\leftarrow\; v - 2 \\ & \text{while } v < -1 \text{ do } \; v \;\leftarrow\; v + 2 \end{aligned}$$

For any $\delta$, and any noise distribution, this leaves the $U(-1,+1)$ distribution for $v$ invariant. We can alternate updates of $v$ as above with Metropolis (or Langevin) updates for $x$ (or $x$ and $p$). When a proposal is accepted, we need to keep $s$ constant, so the reverse update accepts, which requires updating $v$ to $v\,\pi(x)/\pi(x^*)$ (or to $v\,\pi(x,p)/\pi(x^*, -p^*)$).

## Results: Metropolis updates, high-D Gaussians

40-dimensional Gaussian, identity covariance, r-w Metropolis with stepsize ($\sigma$) scaled down by $40^{1/2}$. Autocorrelation times for one coordinate and for log probability of state, for groups of 40 iterations.

Top: Standard Metropolis.
Bottom: Non-reversible with $\delta = 0.2$, no noise.

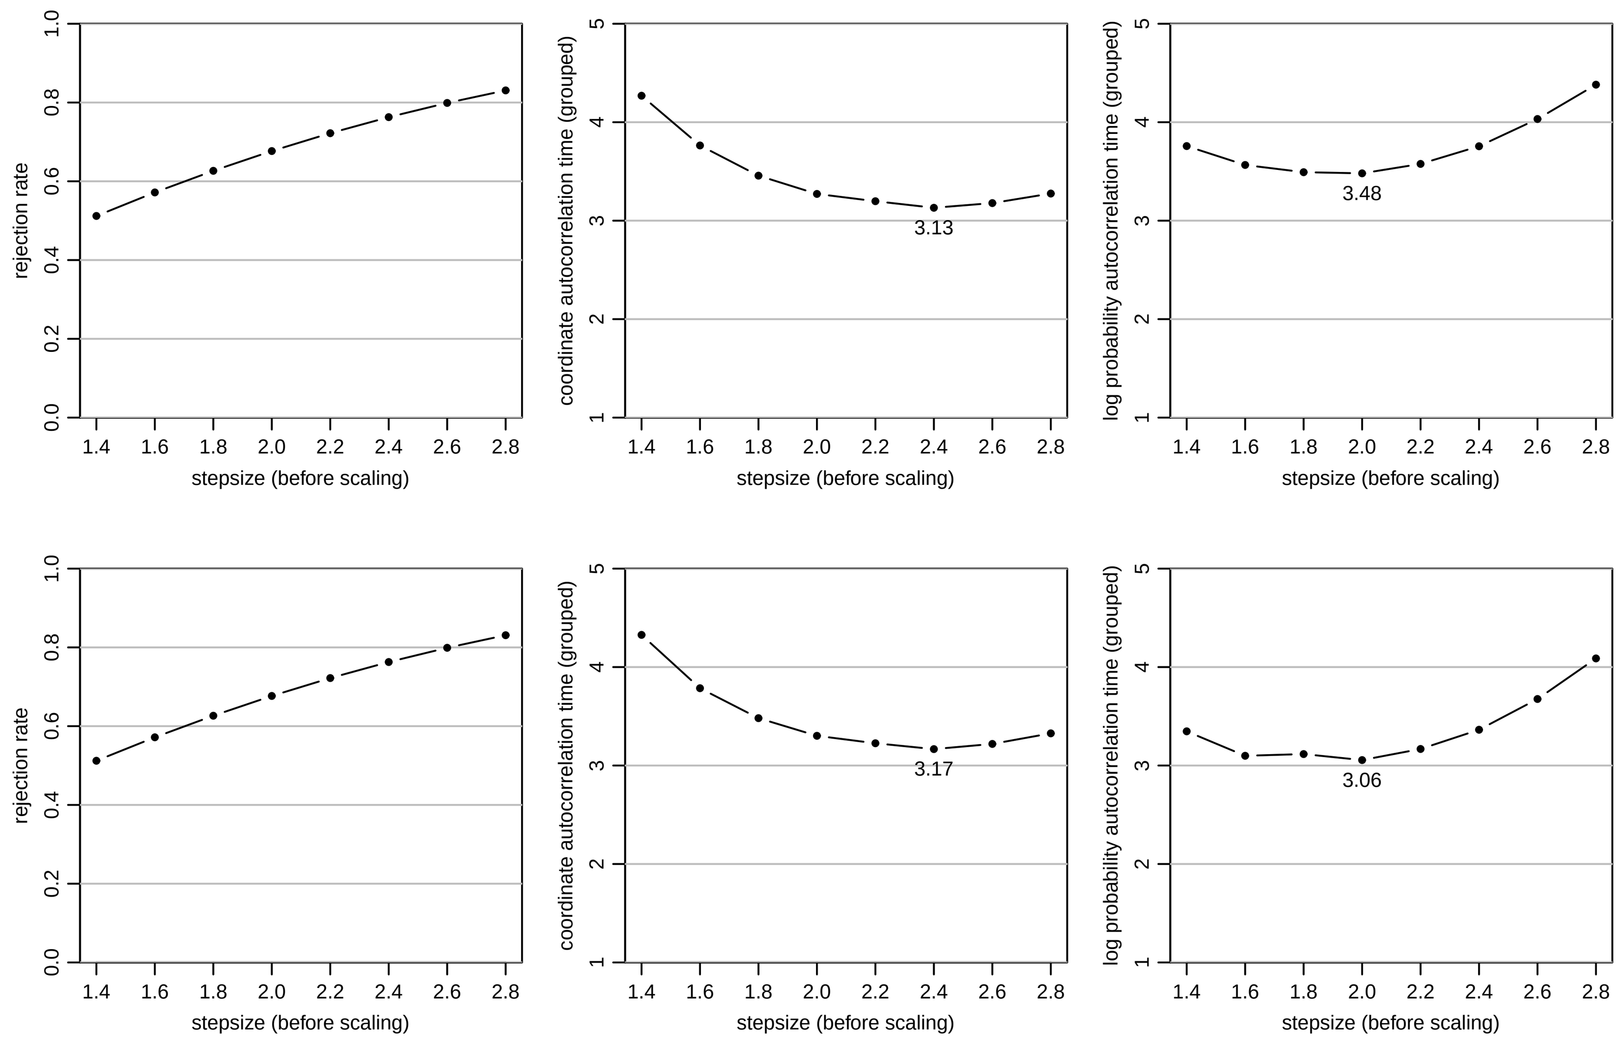


Non-reversible is 1.14 times better for log prob.

## Results: Persistent Langevin, high correlation

16 independent pairs of bivariate Gaussian variables, with variances of 1 and correlations of 0.99. Persistent Langevin, stepsize ($\eta$) scaled down by $32^{1/6}$, $\alpha \in \{0.1^\eta, \ldots, 0.8^\eta\}$. 31 iterations in group.

Top: Standard Persistent Langevin.
Bottom: Non-reversible with $\delta = 0.03$, no noise.

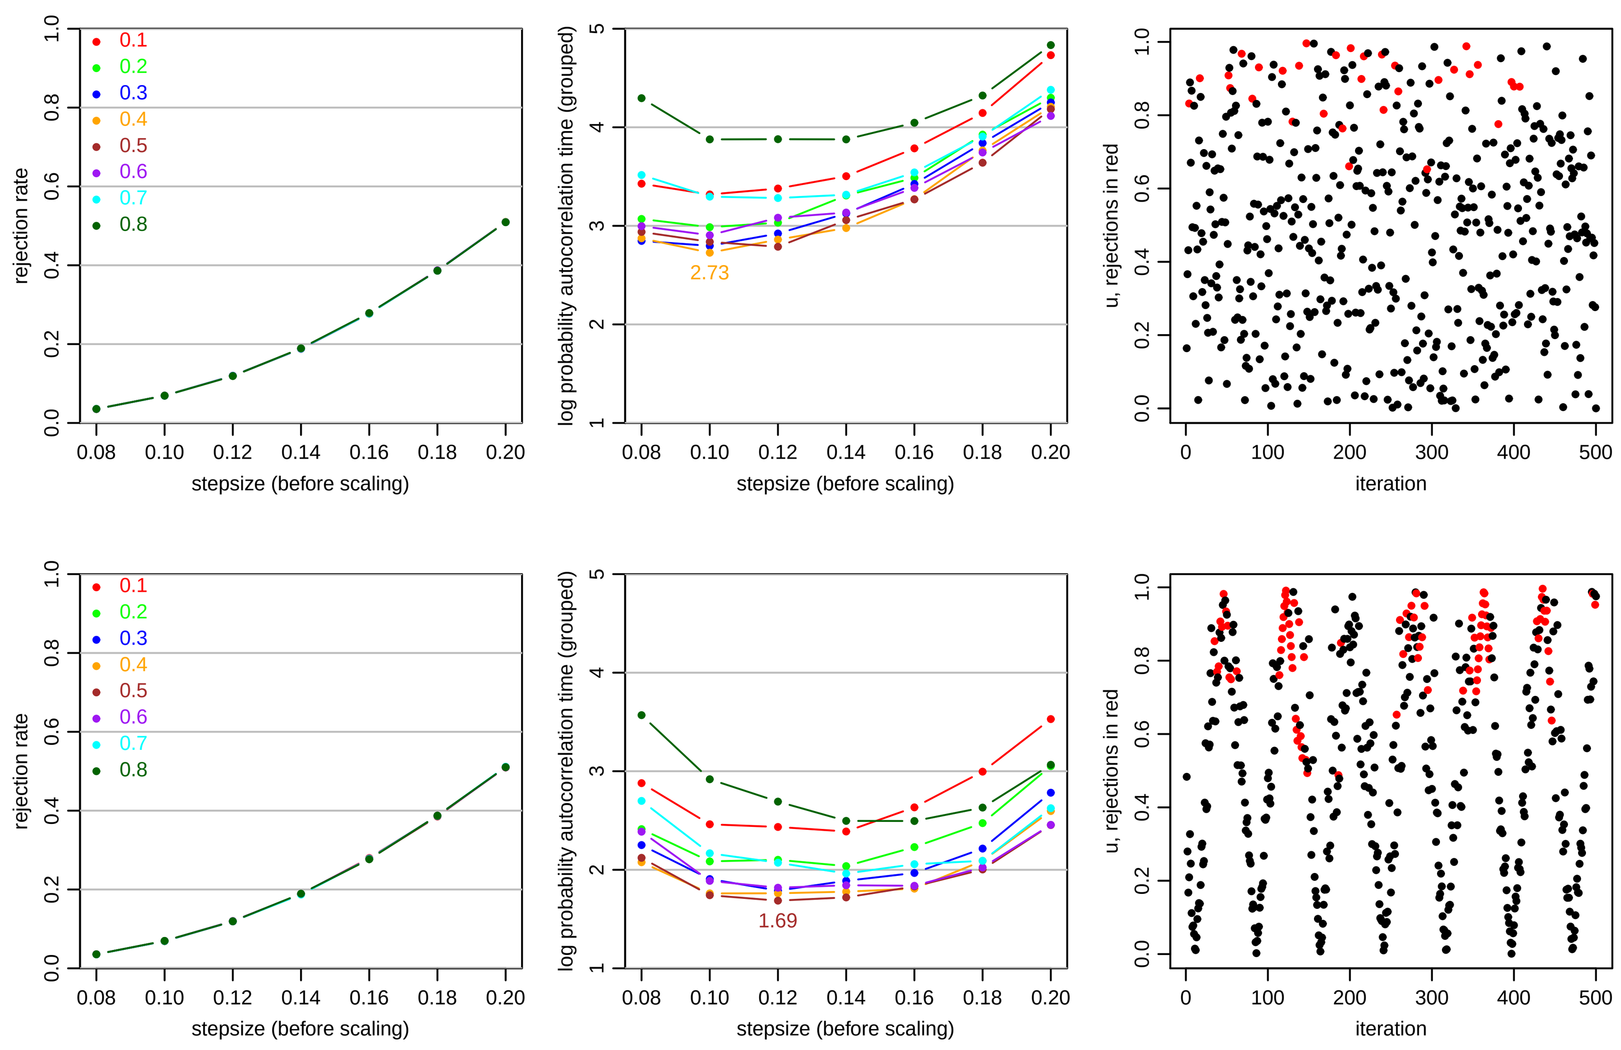


Non-reversible method is 1.62 times better for log prob. Right plots show clustering of rejections with non-reversible updates, reducing random walks.

## Bayesian neural networks

Compared to HMC, sampling a NN posterior with persistent Langevin allows more frequent hyperparameter updates with Gibbs sampling. Non-reversible updates of $v$ make this work better.

Advances in Approximate Bayesian Inference, Vancouver, 2019-12-08