SANDIA REPORT
SAND2020-3166
Printed March, 2020

Transitional Markov Chain Monte Carlo Sampler in UQTk

Mohammad Khalil, Jacky Lao, Cosmin Safta, Habib Najm

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185
Livermore, California 94550

Issued by Sandia National Laboratories, operated for the United States Department of Energy by National Technology & Engineering Solutions of Sandia, LLC.

Available to DOE and DOE contractors from

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@osti.gov
Online ordering: http://www.osti.gov/scitech

Available to the public from

U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Road
Alexandria, VA 22312

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.gov
Online order: https://classic.ntis.gov/help/order-methods

ABSTRACT

Transitional Markov Chain Monte Carlo (TMCMC) is a variant of a class of Markov Chain Monte Carlo algorithms known as tempering-based methods. In this report, the implementation of TMCMC in the Uncertainty Quantification Toolkit is investigated through the sampling of high-dimensional distributions, multi-modal distributions, and nonlinear manifolds. Furthermore, the Bayesian model evidence estimates obtained from TMCMC are tested on problems with known analytical solutions and shown to provide consistent results.

ACKNOWLEDGMENT

The authors acknowledge the helpful feedback that Xun Huan, Khachik Sargsyan, Bert Debusschere, and Tiernan Casey have provided on various aspects of UQTk.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

PREFACE

SUMMARY

This report examines several applications that employ Transitional Markov chain Monte Carlo (TMCMC) algorithm implemented in the Uncertainty Quantification Toolkit. The performance of this parallel sampling algorithm is assessed on several test problems, including high dimensional and multi-modal distributions, with a good degree of success in extracting 1st and 2nd moments. The suite of tests includes verification studies based on estimating the fractional weight of modes in multi-modal distributions as well as sampling from complex nonlinear manifolds. TMCMC's model evidence estimates are examined for linear regression test cases, exhibiting increasing accuracy with increasing number of samples. Overall, TMCMC is able to provide samples from multi-modal and high-dimensional distributions as well as consistent Bayesian model evidence estimates.

1. INTRODUCTION

Markov chain Monte Carlo (MCMC) methods are widely used for sampling from general probability density functions (PDFs) in various settings. Traditional MCMC methods are inefficient or ineffective in sampling from multi-modal or high dimensional distributions. Transitional Markov Chain Monte Carlo (TMCMC) [4] is an effective sampling method that transitions samples from the prior PDF to the posterior PDF through several intermediary distributions. In the context of Bayesian inversion, TMCMC techniques can be particularly useful when the model evaluations are expensive and there are sufficient parallel computing resources to perform multiple runs concurrently.

In this report we examine the performance of TMCMC on several canonical problems that simulate challenges related to high-dimensionality and clustering of samples along low-dimensional manifolds or regions.

This report is organized as follows. Chapter 2 provides a brief summary of TMCMC as implemented in the Uncertainty Quantification Toolkit (UQTk) [6]. Chapters 3 through 5 show the results of tests relating to sampling of (a) multi-modal posterior PDFs, (b) moment estimation for high-dimensional posterior PDFs, and (c) capturing probability mass of modes in multi-modal posterior PDFs, respectively. Chapter 6 shows results related to sampling of manifolds. Chapter 7 shows the results of the TMCMC model evidence estimator on a linear regression test case. Chapter 8 provides a summary of results and insight into the mechanism behind TMCMC and how that manifests itself for the problems investigated in this report.

2. TRANSITIONAL MARKOV CHAIN MONTE CARLO ALGORITHM

TMCMC belongs to a class of tempering sampling algorithms. The algorithm starts with samples corresponding to a "high temperature" distribution, essentially samples from a prior distribution in a Bayesian framework. Then, as "temperature" decreases the target distribution "cools down" to concentrated regions, or regions where the posterior distribution is concentrated by data in a Bayesian framework. During this process samples are drawn to high-probability region through a standard accept-reject procedure.

TMCMC combines aspects of simulated annealing optimization with MCMC, creating an algorithm that parallelizes the model evaluation. It also provides an estimate for model evidence in a Bayesian context [4]. The algorithm can be summarized as follows.

Algorithm 2.0.1 [TMCMC] Start by generating N samples from the prior distribution $p(\theta|M)$, $\theta_k^{(0)}$, for $k = 0, \ldots, N$ and set initial likelihood exponent $\beta_0 = 0$. For each iteration i of the algorithm, repeat the following steps until $\beta_i = 1$

1. Compute the likelihood of the samples, $p(D|\theta_k^{(i)})$, $k = 1, \ldots, N$
2. Calculate the plausibility weights of each sample $w(\theta_k^{(i)}) = p(D|\theta_k^{(i)})^{\beta_i - \beta_{i-1}}$, with β_i chosen such that coefficient of variation of those weights is equal to or less than a set value, CoV.
3. Resample (with replacement) N samples from the current set of samples $\theta_k^{(i)}$ with corresponding probabilities $\frac{w(\theta_k^{(i)})}{\sum_{k=1}^{N} w(\theta_k^{(i)})}$.
4. Run Markov chains with standard Metropolis updates for each sample that was chosen in the above step, with a chain length equal to the count t that the sample was selected (in resampling). The proposal PDF is a multivariate Gaussian centered at the chosen sample, with covariance matrix constructed using samples from the previous stage (see below) and γ^2 a scaling parameter.

$$\Sigma = \gamma^2 \sum_{k=1}^{N} w(\theta_k^{(i)}) \left(\theta_k^{(i)} - \bar{\theta}^{(i)}\right) \left(\theta_k^{(i)} - \bar{\theta}^{(i)}\right)^T$$

$$\bar{\theta}^{(i)} = \sum_{j=1}^{N} \theta_j^{(i)} \frac{w(\theta_j^{(i)})}{\sum_{l=1}^{N} w(\theta_l^{(i)})}$$

Markov chain Monte Carlo algorithms require tuning of algorithmic knobs for optimal performance. Generally, the optimal parameter values will change depending on the state of the samples. Since TMCMC transitions through intermediate PDFs, optimal parameter values will change from one iteration to the next. Adaptive tuning is a necessity in improving efficiency and reducing total number of likelihood evaluations (or forward model simulations).

The TMCMC algorithm can be adjusted through several control parameters. The β parameter is the temperature parameter that generates the intermediate PDFs to transition through, and takes values $0 = \beta_0 < ... < \beta_m = 1$, with m being the number of stages or iterations required in TMCMC. The CoV is the ratio of the sample weight standard deviation and the sample weight mean. It's useful as a relative measure of dispersion, and in the TMCMC context, relates to whether the samples have conformed to the intermediate PDF. A lower threshold results in more intermediate PDFs and a smoother transition to posterior, but also more computational expense. A higher threshold creates a sharper transition to posterior, at the risk of a small effective sample size and limited ability to search the space. The threshold value CoV is typically held constant, and Beck and Zuev [19] show that the optimal cooling rate from prior to posterior distributions is $CoV = 1.0$ under certain assumptions.

Another control parameter in TMCMC is a scaling factor that adjusts the covariance matrix of the MVN proposal PDF. In the original paper, a parameter $c_m^2 = 0.04$ is proposed as a reasonable constant value. However, this is likely to be suboptimal depending on the problem involved, particularly in highly correlated parameters. The general motivation for adapting the control parameter is to produce larger proposals and encourage exploration of the space when acceptance rate R is high, and produce smaller proposals with a higher acceptance when the rate is low. Minson *et al.* [12] suggest a rescaling based on the observed acceptance rate, in the form

$$c_m = a + bR \tag{2.0.1}$$

where $a = \frac{1}{9}, b = \frac{8}{9}$ are selected constants. This is an improvement, but is also relatively limiting in that the control parameter is limited in the range $[a, 1]$. In [3], Catanach uses control theory to show that the control parameter is best adapted using

$$\ln c_{m+1} = \ln c_m + G(R - R^*) \tag{2.0.2}$$

$$c_{m+1} = c_m exp\left(G(R - R^*)\right) \tag{2.0.3}$$

.

The current implementation of TMCMC uses this adaptive form, with an optimal value of $R^* = 0.234$, based on Roberts and Rosenthal [14].

3. SAMPLING OF MULTI-MODAL POSTERIORS

Classical single-chain MCMC methods are challenged by multi-modal probability distributions, as Markov chains will rarely propose **and** accept jumps between modes. This challenge is further exacerbated when the dimensionality of the parameter space is large and typically high-probability regions become more concentrated. All these make the selection of a good proposal distribution difficult. Various algorithms have been developed to deal with these situations, including mode-jumping MCMC [17], nested sampling [8], and geometric MCMC methods [9].

TMCMC's ability to sample from multi-modal distributions will be examined using canonical test functions for optimization, including the Rosenbrock, Griewank, Himmelblau, and Schwefel functions [16]. Those functions will serve as the negative log-likelihood functions. All tests employ uniform prior PDFs. We will consider the following functions

$$f(\mathbf{x}) = (1-x_1)^2 + 100(x_2 - x_1^2)^2 \qquad \text{(Rosenbrock)}$$

$$f(\mathbf{x}) = \sum_{i=1}^{d} \frac{x_i^2}{4000} - \prod_{i=1}^{d} \cos\left(\frac{x_i}{\sqrt{i}}\right) + 1 \qquad \text{(Griewank)}$$

$$f(\mathbf{x}) = 418.9829d - \sum_{i=1}^{d} x_i \sin\left(\sqrt{|x_i|}\right) \qquad \text{(Schwefel)}$$

$$f(\mathbf{x}) = (x^2 + y - 11)^2 + (x + y^2 - 7)^2 \qquad \text{(Himmelblau)}$$

The Rosenbrock function, is a common optimization test function with a global minimum at $(1,1)$ in 2 dimensions. The Griewank function exhibits several regularly spaced local minima and can be extended to d dimensions. The global minimum is at $(0,\ldots,0)$. The Schwefel function exhibits multiple local minima, with a global minimum is at (420.9687, ..., 420.9687). The Himmelblau function has four local minima, at (3.0, 2.0), (-2.805118, 3.131312), (-3.779310, -3.283186), and (3.584428, -1.848126). The motivation for using these examples to is illustrate the ability of TMCMC to easily sample multi-modal posterior PDFs. Each region centered on local or global minima for these functions translates into a region of high-probability when used as a negative log-likelihood.

We will qualitatively examine the progression of the algorithm as it transitions through the intermediate PDFs from prior to posterior, as illustrated in Figure 3. The Rosenbrock function test case is shown in 3-1(a), proceeding from the uniform distribution (prior PDF) to a thin, curved, posterior. Note the intermediate distribution is a diffused version of the target posterior. Eventually, the resulting samples are found to be distributed according to the target posterior PDF,

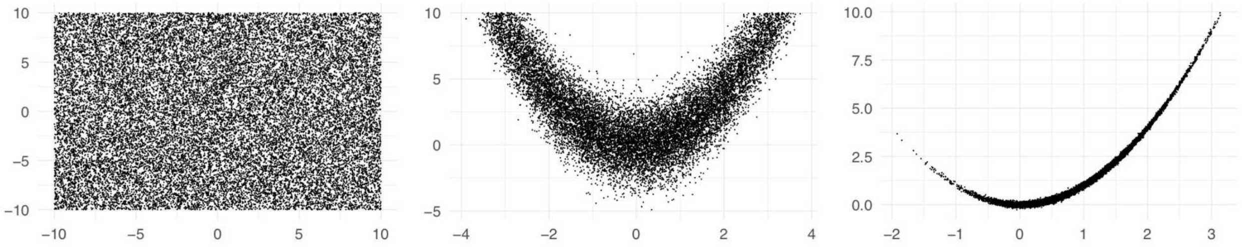

(a) Rosenbrock

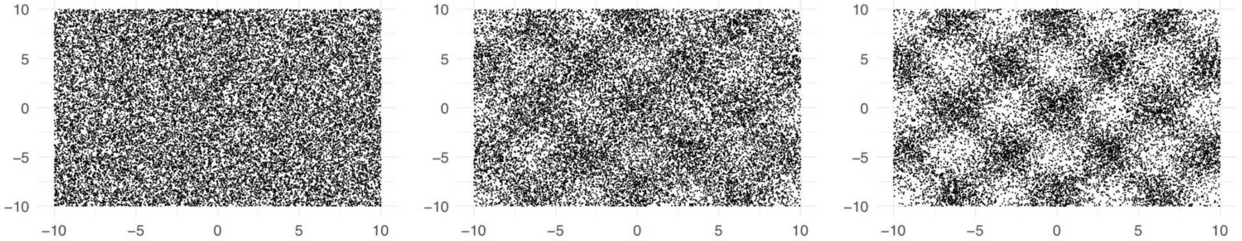

(b) Griewank

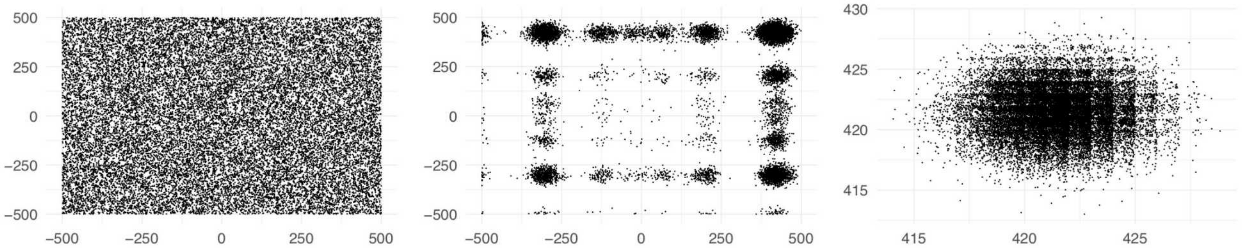

(c) Schwefel

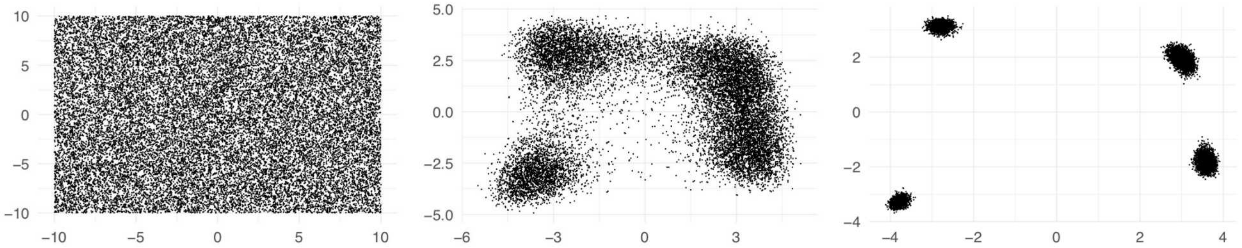

(d) Himmelblau

Figure 3-1. TMCMC Performance on Optimization Functions

as expected. The Griewank function in Fig. 3-1(b) displays multiple local minima and the regular spacing between the modes. Even though there is no unique global maximum for this posterior PDF, TMCMC is able to capture all modes starting with a flat (uniform) prior. The Schwefel function has irregularly spaced local minima, and those are clearly visible in this situation, compared to the Griewank function. Despite the relatively diffuse prior PDF on $[-500,500]^2$, TMCMC can identify the mode at a corner of this 2-dimensional setting. The Himmelblau function shown in Fig. 3-1(d) has four modes, all of which are clearly distinguished. In the next section we will present a quantitative analysis of the TMCMC algorithms.

4. SAMPLING OF HIGH-DIMENSIONAL POSTERIORS

Next, we consider the efficacy at which TMCMC samples high-dimensional multivariate normal (MVN) PDFs towards accurate estimation of parameter mean (1st order moment) and variance (2nd order central moment). MVN parameter likelihood PDFs in 2, 5, 10, 20, 50, and 100 dimensions are sampled by TMCMC (20 repeated experiments for each dimensionality) and the Normalized Mean Absolute Error (NMAE) and Normalized Root Mean Squared Error (NRMSE) are obtained for the first and second moments, as in

$$\text{NMAE} \triangleq \frac{\frac{\sum_{i=1}^{d} |\text{Estimated Moment}(x_i) - \text{Target}|}{d}}{\text{Target}} \tag{4.0.1}$$

$$\text{NRMSE} \triangleq \frac{\sqrt{\frac{\sum_{i=1}^{d} (\text{Estimated Moment}(x_i) - \text{Target})^2}{d}}}{\text{Target}} \tag{4.0.2}$$

The MVN likelihood parameter PDFs all have means at $(5,\ldots,5)$, and standard deviation of $(2,\ldots,2)$. The MVN priors have have means at $(0,\ldots,0)$, and standard deviation of $(5,\ldots,5)$. The resulting posterior also follows an MVN density with mean and variances available analytically (due to conjugacy). The experiments are repeated with three algorithmic choices: (a) 10,000 samples and CoV of 0.1, (b) 25,000 samples and CoV of 0.1, and (a) 10,000 samples and CoV of 0.05, with results presented in Figs. 4-1-4-3. For each pair of CoV and sample number, both errors tend to increase with increasing dimensionality, regardless of the number of samples or CoV used. Furthermore, the error in the 2nd moment is generally greater than that in the 1st moment, as expected. We also observe that a decrease in allowable CoV for the intermediate TMCMC sample weights from 0.1 to 0.05 (Fig. 4-1 to Fig. 4-3) results in a reduced error in both mean and variance estimates, having a similar effect as increasing the number of samples.

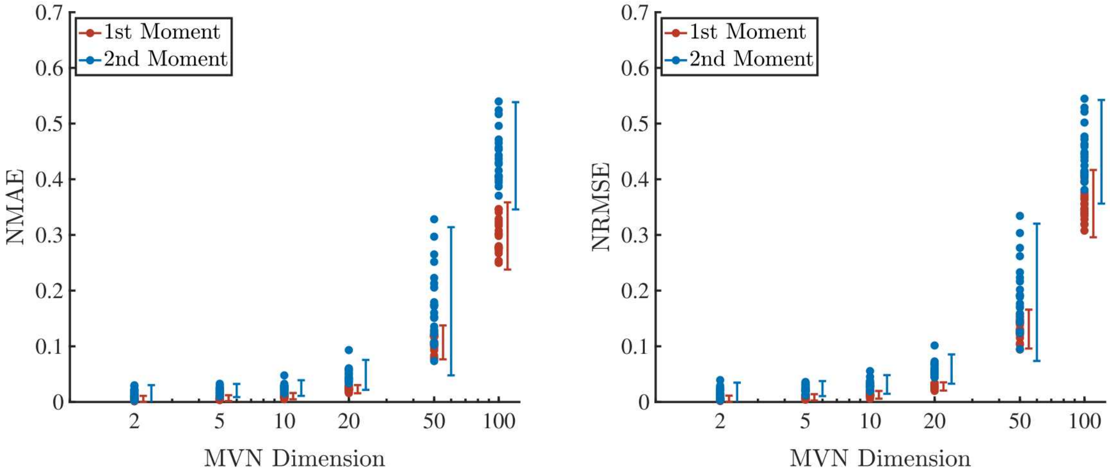

(a) Normalized Mean Absolute Error (b) Normalized Root Mean Squared Error

Figure 4-1. TMCMC for moment estimation of high-dimensional Gaussian posteriors using 10,000 samples and CoV of 0.1

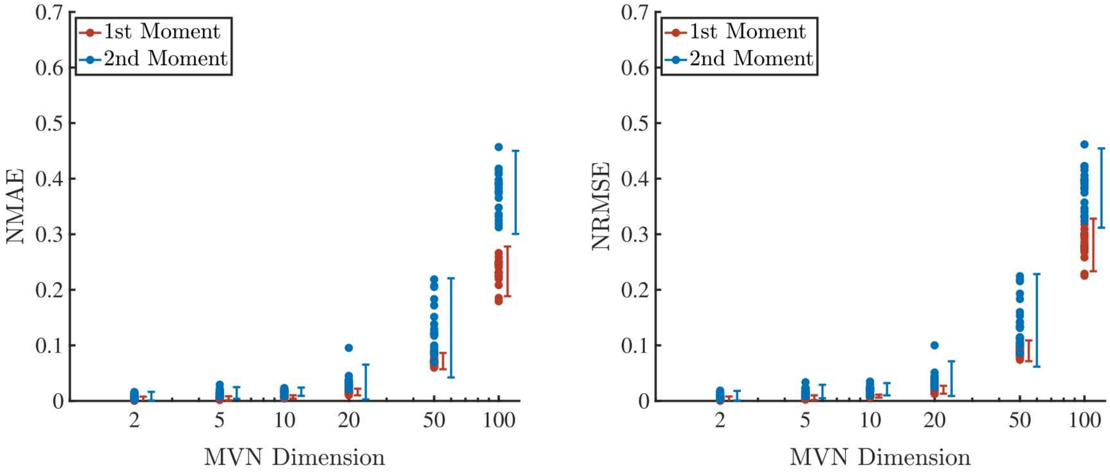

(a) Normalized Mean Absolute Error (b) Normalized Root Mean Squared Error

Figure 4-2. TMCMC for moment estimation of high-dimensional Gaussian posteriors using 25,000 samples and CoV of 0.1

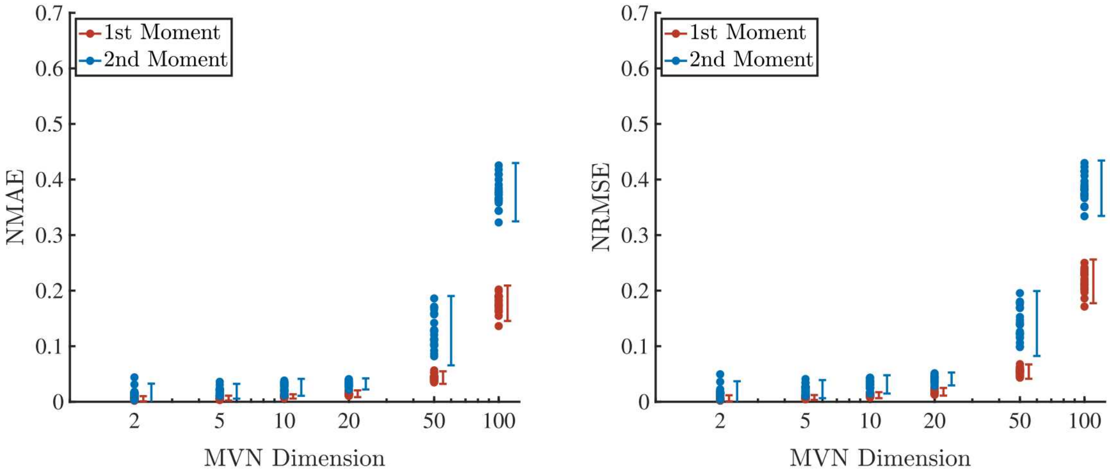

(a) Normalized Mean Absolute Error (b) Normalized Root Mean Squared Error

Figure 4-3. TMCMC for moment estimation of high-dimensional Gaussian posteriors using 10,000 samples and CoV of 0.05

5. CAPTURING PROBABILITY MASS OF MODES IN MULTI-MODAL POSTERIORS

Next we test the sampling process, for a multi-modal posterior case, by accounting for the proportion of samples that correspond to each mode. We compare the results against the true value given by the probability mass associated with each mode. This tests the ability of TMCMC to sample multi-modal densities beyond the simpler challenge of identifying all relevant modes. Here, we consider a Gaussian mixture model with two modes in 2, 5, and 10 dimensions, and we use 1000, 5000, and 25000 samples. The prior distribution is a uniform prior with support on $[0,1]^d$. The modes means are located at $(0.25,\ldots,0.25)$ and $(0.75,\ldots,0.75)$, with a diagonal covariance and standard deviation of 0.05. This results in the same probability mass for both modes. Each test was run 50 times with different random seeds. Modal probability mass for the first mode was estimated by counting the number of samples with all coordinates less than 0.5. Fig. 5-1 and Table 5 summarize this set of results. With increasing sample size, the error (standard deviation) in the estimated mass (target value of 0.5) decreases, as expected. This is also seen with decreasing dimensionality.

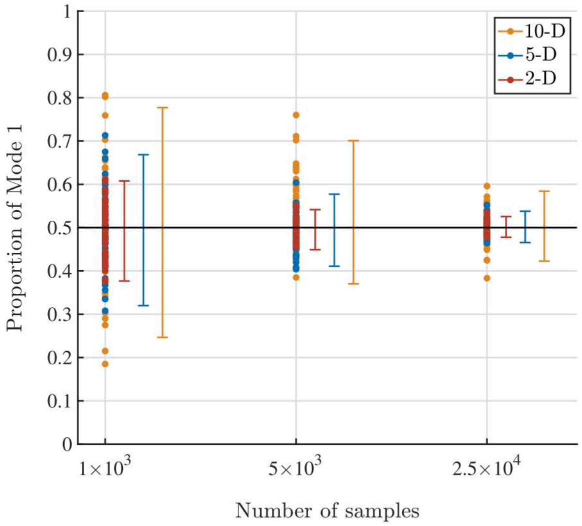

Figure 5-1. Proportions of a Mode in a Bimodal Distribution

Number of Samples	**Dimension**	**# Over 50%**	**# Under 50%**	**Mean**	**SD**
1000	2	28	22	0.5119	0.1326
1000	5	22	28	0.4942	0.0871
1000	10	25	25	0.4923	0.0578
5000	2	33	17	0.5355	0.0827
5000	5	22	28	0.4941	0.0416
5000	10	21	29	0.4953	0.0231
25000	2	30	20	0.5035	0.0404
25000	5	30	20	0.5017	0.0181
25000	10	24	26	0.5017	0.0120

Table 5-1. Statistics on Proportions of samples in a Bimodal Distribution

6. MANIFOLD SAMPLING

Sampling from manifolds is a difficult problem that occurs in a variety of situations, such as in sampling constrained parameter spaces (see Diaconis et al [7] for an overview). In the case of data-free inference [13, 11], samples are drawn from data space, generating datasets that satisfy constraints on the resulting posterior parameter PDFs. From these datasets, we can fit the data to multiple models, and run inferences, in the absence of the original data. The sampling of these datasets is an ill-posed problem since the dimensionality of the data space is normally much greater than the number of available constraints. Furthermore, these constraints compose a manifold in this space, with asymptotically thin regions of probability mass. Traditional MCMC methods have a negligible probability of proposing jumps along these manifolds and, even when they do, the exploration of the manifold is difficult since it is associated with an invariant density.

TMCMC is tested for its capability for sampling manifolds for several canonical test cases. Due to the nature of the algorithm, samples explore the data space freely at high temperatures (low β values), and slowly aggregate on the manifold as the temperature decreases (β approaching 1). However, manifolds are asymptotically thin, requiring a further adjustment. We use a loose constraint construction as in Berry et al. [2], where the manifold is smoothed out using Gaussian kernels, and successive iterations using decreasing δ values to enforce the constraints more strictly.

$$\pi(\mathbf{x}) = \exp\left(\frac{-\|f(\mathbf{x}) - f^*(\mathbf{x})\|^2}{\delta^2}\right) \tag{6.0.1}$$

For each value of δ, TMCMC is run on the previous iteration's data, and utilizes the previous posterior as the current iteration's prior distribution. In effect, δ is a meta-temperature such that the posterior distribution converges to the manifold as δ approaches 0. The motivation for this construction was to reduce time spent in excessive model evaluations for samples far from the manifold. A single run of TMCMC is insufficient as, depending on the prior, the majority of samples will not propose jumps into the manifold at small δ values. By using this construction, samples are guaranteed to be sufficiently close to the manifold, such that Markov chains spawned from these samples have a reasonable acceptance ratio. This construction is also flexible enough to expand to multiple simultaneous constraints as well as high dimensionality.

Two test cases are proposed to illustrate this nested sampling algorithm. The first is a circle with center at $(5,10)$ and radius of 10. The second is a hyperbola with the constraint $x_1x_2 = 4$. In the outer loop, the meta-temperature δ decreases according to the power law schedule, with a rate $r < 1$ and t representing the number of times TMCMC has been run. The rate for the two examples is $r = 0.9$. Fig. 6-1(a) is shown at $t = 5,25,49$. Fig. 6-1(b) is shown at $t = 0,2,8$.

$$\delta = \delta_0 \times r^t \tag{6.0.2}$$

$$(x_1 - 5)^2 + (x_2 - 10)^2 - 100 = \delta \tag{6.0.3}$$

$$x_1 x_2 - 4 = \delta \tag{6.0.4}$$

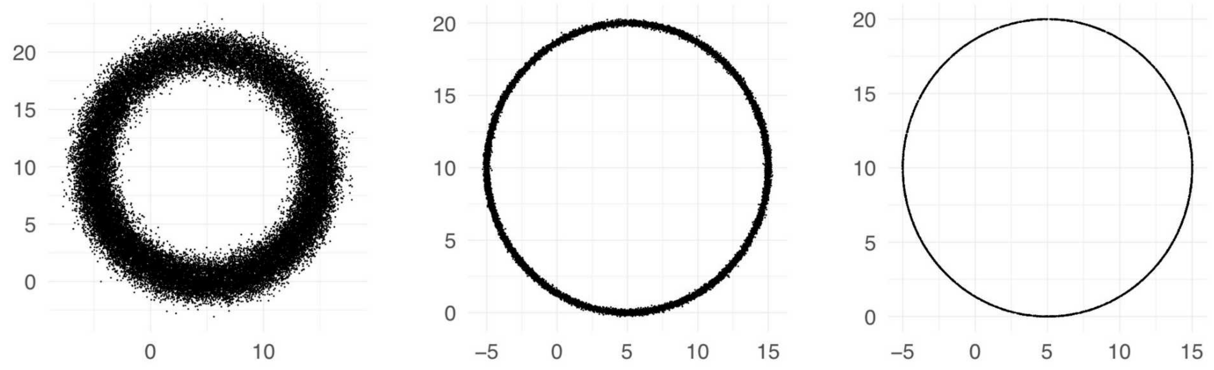

(a) Circle

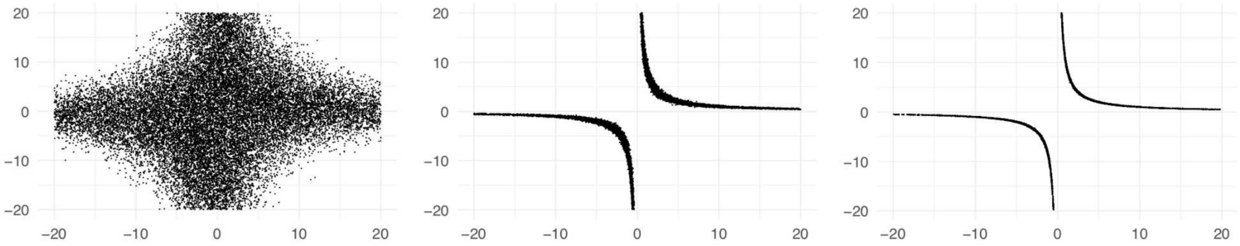

(b) Nonlinear Function: $x_1 x_2 = 4$

Figure 6-1. TMCMC Performance on manifolds in two dimensions.

As a proof of concept, Fig. 6 shows that TMCMC can sample from "narrow" distributions along low-dimensional embeddings. The circular manifold would enforce a circular proposal, and traditional MCMC would only be able to explore a small percentage of the circumference. The nonlinear manifold would be intractable due to the bimodality and the strong nonlinear correlations. TMCMC succeeds in exploring these proof-of-concept models. Further explorations are needed to gauge the applicability of this algorithm to higher dimensional problems.

7. MODEL EVIDENCE

Model evidence is the likelihood that a set of data was generated by a specific model and is given by the expected likelihood, weighted by the prior PDF. This is an important component in Bayesian inference, as model evidence is used to compare different models. In general, the calculation of the evidence value is not available analytically and its numerical evaluation relies on a large number of model evaluations. Further notes on model selection and model updating can be found in [5, 15].

TMCMC, in addition to being a robust sampler, provides a theoretically asymptotically unbiased estimator for model evidence, being the the product of the mean of the plausibility weights across all intermediate distributions. We gauge the convergence properties of the TMCMC-based evidence on for a polynomial model. For this test, 10 data points were generated from a cubic polynomial (in Eq. (7.0.1)) with x values uniformly spaced in $(-1.2, 1.2)$. The y-values were then contaminated by additive Gaussian noise $\varepsilon \sim N(0, 0.2)$.

$$y = x^3 + x^2 - 6 + \varepsilon \tag{7.0.1}$$

TMCMC was then used to sample the posterior distributions for the coefficients of polynomial models of increasing order, from 1 to 7. For these tests, the polynomial coefficients employ independent Gaussian priors with mean 0 and standard deviation 5. Similarly, the likelihood is presumed to be a product of independent Gaussians, centered at each data point and a standard deviation of 0.2 - the same as the value used to corrupt the original model evaluations. Given the settings for prior and likelihood, the model evidence can be estimated analytically. Several TMCMC cases were run, with different CoV values and number of samples per stage. For each test case, a number of 10 replicas were run to gauge the spread in results due to the choice in the random number seed.

Error bars in Fig 7-1 show 2 standard deviations above and below the mean of the 10 TMCMC runs, and are offset for visibility. The results in this figure show increased agreement between TMCMC results and the analytical values with smaller CoV values and with increasing number of samples per stage. These obervations are consistent with trends displayed earlier in this report about the effects of CoV and number of samples in the accuracy of the posterior distribution in several canonical test cases.

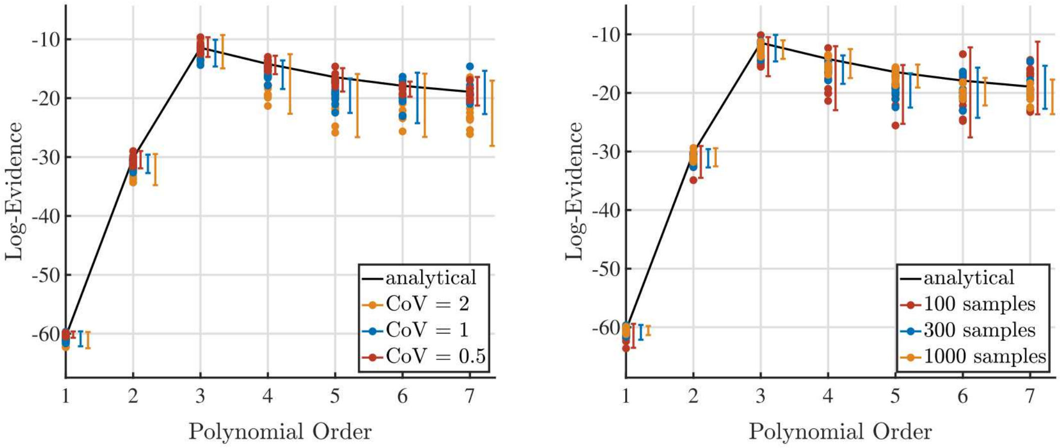

(a) Varying CoV (b) Varying number of samples

Figure 7-1. TMCMC Model Evidence Estimates: 10 repeated TMCMC runs per experiment with varying random number generator seed

8. CONCLUSION

The Transitional Markov Chain Monte Carlo (TMCMC) algorithm has been demonstrated to be an effective MCMC algorithm for sampling complex distributions. Due to the temperature adaptation and the resulting intermediate densities constructed by the sampling algorithm, multiple modes in a posterior PDF can be robustly identified. Manifold sampling with TMCMC was exhibited with low-dimsnional yet complex (strongly nonlinear) manifolds. Tuning parameters, specifically the threshold CoV and the resulting likelihood exponent β, have strong effects on the performance of TMCMC. By design, β is adaptive, based on the uniformity of samples for a particular intermediate PDF. This ensures that samples are transitioning smoothly through the intermediate PDFs, which results in desirable properties such as the capability to sample from multi-modal distributions. The covariance scaling parameter is also adaptive, based on the acceptance ratio. This assists in maintaining a suitably optimal acceptance ratio and allowing efficient search of the parameter space.

Model evidence is a central quantity in Bayesian model selection and model averaging [10]. Standard methods for calculating model evidence involve numerical quadrature and specialized methods. TMCMC produces an estimator for model evidence as a part of the algorithm. The model evidence estimates appear to be unbiased for a polynomial fitting problem having analytical results, while the precision increases with increasing number of samples or decreasing CoV, as expected.

High dimensionality poses a challenge for all MCMC algorithms. Since TMCMC gradually adapts from prior to posterior, it can be an effective algorithm for searching low-dimensional embeddings in high dimensional spaces. Certain geometric MCMC methods, such as Metropolis-Adjusted Langevin algorithm (MALA), and Riemannian Manifold Hamiltonian Monte Carlo (RMHMC) use the geometry of the posterior distribution to inform the proposals, which improve exploration the probability distribution [18, 1, 9]. In the context of TMCMC, these methods can replace the MVN proposal to improve efficiency. This will be explored in subsequent versions of UQTk.

REFERENCES

[1] Yves F. Atchadé. An adaptive version for the metropolis adjusted langevin algorithm with a truncated drift. *Methodology and Computing in Applied Probability*, 8(2):235–254, 2006.

[2] R.D. Berry, H.N. Najm, B.J. Debusschere, H. Adalsteinsson, and Y.M. Marzouk. Data-free inference of the joint distribution of uncertain model parameters. *Journal of Computational Physics*, 231:2180–2198, 2012.

[3] Thomas Anthony Catanach. *Computational Methods for Bayesian Inference in Complex Systems*. PhD thesis, California Institute of Technology, 2017.

[4] J. Ching and Y.-C. Chen. Transitional Markov Chain Monte Carlo Method for Bayesian Model Updating, Model Class Selection, and Model Averaging. *Journal of Engineering Mechanics*, 133(7):816–832, 2007.

[5] Jianye Ching, Matthew Muto, and James L. Beck. Structural Model Updating and Health Monitoring with Incomplete Modal Data Using Gibbs Sampler. *Computer-Aided Civil and Infrastructure Engineering*, 21(4):242–257, 2006.

[6] B. Debusschere, C. Safta, K. Sargsyan, K. Johnston, P. Rai, M. Khalil, T. Casey, and X. Zeng. UQ Toolkit. http://www.sandia.gov/UQToolkit, 2020.

[7] Persi Diaconis, Susan Holmes, and Mehrdad Shahshahani. *Sampling from a Manifold*, volume Volume 10 of *Collections*, pages 102–125. Institute of Mathematical Statistics, Beachwood, Ohio, USA, 2013.

[8] F. Feroz and M. P. Hobson. Multimodal nested sampling: an efficient and robust alternative to Markov Chain Monte Carlo methods for astronomical data analyses. *Monthly Notices of the Royal Astronomical Society*, 384(2):449–463, 01 2008.

[9] Mark Girolami and Ben Calderhead. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2):123–214, 2011.

[10] Jennifer A. Hoeting, David Madigan, Adrian E. Raftery, and Chris T. Volinsky. Bayesian model averaging: a tutorial. *Statist. Sci.*, 14(4):382–417, 11 1999.

[11] M. Khalil and H.N. Najm. Probabilistic Inference of Reaction Rate Parameters from Summary Statistics. *Combustion Theory and Modeling*, 22(4):635–665, 2018.

[12] S. E. Minson, M. Simons, and J. L. Beck. Bayesian inversion for finite fault earthquake source models (I) - theory and algorithm. *Geophysical Journal International*, 194(3):1701–1726, 2013.

[13] Habib N. Najm, Robert D. Berry, Cosmin Safta, Khachik Sargsyan, and Bert J. Debusschere. Data-free inference of uncertain parameters in chemical models. *International Journal for Uncertainty Quantification*, 4(2):111–132, 2014.

[14] Gareth O. Roberts and Jeffrey S. Rosenthal. Optimal scaling for various Metropolis-Hastings algorithms. *Statist. Sci.*, 16(4):351–367, 11 2001.

[15] Rimple Sandhu, Chris Pettit, Mohammad Khalil, Dominique Poirel, and Abhijit Sarkar. Bayesian model selection using automatic relevance determination for nonlinear dynamical systems. *Computer Methods in Applied Mechanics and Engineering*, 320(C):237–260, 2017.

[16] Sonja Surjanovic and Derek Bingham. Virtual Library of Simulation Experiments: Test Functions and Datasets. http://www.sfu.ca/˜ssurjano, 2013.

[17] Hakon Tjelmeland and Bjorn Kare Hegstad. Mode Jumping Proposals in MCMC. *Scandinavian Journal of Statistics*, 28(1):205–223, 2001.

[18] T. Xifara, C. Sherlock, S. Livingstone, S. Byrne, and M. Girolami. Langevin diffusions and the Metropolis-adjusted Langevin algorithm. *Statistics & Probability Letters*, 91:14 – 19, 2014.

[19] Konstantin M. Zuev and James L. Beck. *Asymptotically Independent Markov Sampling: A New MCMC Scheme for Bayesian Inference*, pages 2022–2031.

Sandia
National
Laboratories