

ACCELERATING PARALLEL TEMPERING: QUANTILE TEMPERING ALGORITHM (QUANTA)

NICHOLAS G. TAWN,* AND
GARETH O. ROBERTS,* *University of Warwick*

Abstract

It is well known that traditional Markov chain Monte Carlo (MCMC) methods can fail to effectively explore the state space for multimodal problems. Parallel tempering is a well-established population approach for such target distributions involving a collection of particles indexed by temperature. However, this method can suffer dramatically from the curse of dimensionality. In this paper we introduce an improvement on parallel tempering called QuanTA. A comprehensive theoretical analysis quantifying the improved efficiency and scalability of the approach is given. Under weak regularity conditions, QuanTA gives accelerated mixing through the temperature space. Empirical evidence of the effectiveness of this new algorithm is illustrated on canonical examples.

Keywords: Parallel tempering; Metropolis-coupled MCMC; accelerated MCMC; Markov chain Monte Carlo; multimodal; population MCMC

2010 Mathematics Subject Classification: Primary 60J22
Secondary 65C05; 65C60

1. Introduction

Consider the problem of stochastic simulation from a target distribution, $\pi(x)$, on a d -dimensional state space \mathcal{X} where $\pi(\cdot)$ is known up to a scaling constant. The *gold standard* methodology for this problem uses Markov chain Monte Carlo (MCMC). However, these methods often perform poorly in the context of multimodality.

Most MCMC algorithms use localised proposal mechanisms, tuned towards local approximate optimality, e.g. [31] and [35]. Indeed, many MCMC algorithms incorporate local gradient information in the proposal mechanisms, typically attracting the chain back towards the centre of the mode. This can exacerbate the difficulties of moving between modes [17].

Popular methods used to overcome these issues include *simulated tempering* [18] and the population-based version, *parallel tempering* [8], [9]. These methods use state space augmentation to allow Markov chains to explore target distributions proportional to $\pi^\beta(x)$ for β typically in the range $(0, 1]$. For simulated tempering, this is done by introducing an auxiliary inverse temperature variable, β , and running a $(d + 1)$ -dimensional Markov chain on $\mathcal{X} \times \Delta$, where Δ consists of a discrete collection of possible inverse temperatures including 1. For the more practically applicable parallel tempering approach, a Markov chain is run on a $(|\Delta| \times d)$ -dimensional state space, $\mathcal{X}^{|\Delta|}$, where $|\Delta|$ denotes the cardinality of the set Δ .

Within this paper we will concentrate on parallel tempering as it obviates the need to approximate certain normalisation constants to work effectively. While parallel tempering has been highly successful, for example, see [5], [20], [49], etc., its efficiency declines as a

Received 9 September 2018; revision received 12 April 2019.

* Postal address: Department of Statistics, University of Warwick, Coventry, CV4 7AL, UK.

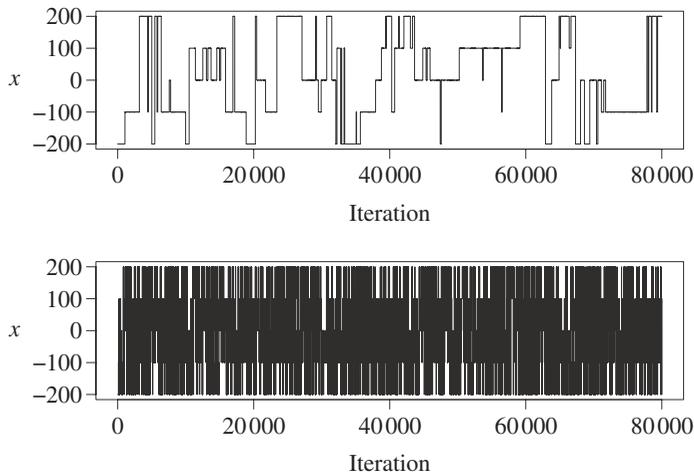


FIGURE 1: Trace plots of the target state chains for representative runs of the parallel tempering (*top*) and QuanTA schemes (*bottom*).

function of d , at least linearly and often much worse [1], [48]. This is caused by the need to set inter-inverse temperature spacings in Δ extremely small to make swaps between temperatures feasible.

In this paper we introduce and analyse the QuanTA algorithm which facilitates inter-temperature swaps by proposing moves which attempt to adjust within-mode variation appropriately for the proposed new temperature. This leads to improved temperature mixing, which in turn leads to vastly improved inter-modal mixing. Its typical improvement is demonstrated in Figure 1 with a five-mode target distribution. The construction of QuanTA resonates with the noncentering MCMC methodology described, for example, in [4], [12], [25], and [26].

Supporting theory is developed to guide setup and analyse the utility of the novel QuanTA scheme. There are two key theoretical results. The first, Theorem 1, establishes that there is an optimal temperature schedule setup for QuanTA, concluding that in general the dimensionality scaling of the distance between consecutive inverse temperature spacings should be $\mathcal{O}(d^{-1/2})$. Further to this it suggests that optimising the expected squared jumping distance between any two consecutive temperature levels induces a temperature swap move acceptance rate of 0.234, giving a useful metric for a practitioner to optimally tune QuanTA. The second key theoretical contribution, Theorem 2, of this paper shows that, under mild regularity conditions, the optimal temperature spacings of QuanTA are more ambitiously spaced than for the standard parallel tempering algorithm for cold (i.e. large) values of the inverse temperatures. The significance of this result is that QuanTA can give accelerated mixing through the cooler parts of the temperature schedule by allowing more ambitious temperature spacings.

This paper is structured into six core sections. In Sections 2 we review the parallel tempering algorithm and some of the relevant existing literature. Section 3 motivates the main idea behind the novel QuanTA scheme, which is then presented in Section 4. QuanTA utilises a population MCMC approach that requires a clustering scheme; discussion for this is found in Section 5. Section 6 contains the core theoretical contributions mentioned above. Simulation studies are detailed in Section 7 along with a discussion of the computational complexity of QuanTA.

2. The parallel tempering algorithm

There is an array of methodology available to overcome the issues of multimodality using MCMC or closely related sequential Monte Carlo (SMC) techniques, e.g. [8], [16], [18], [21], [23], [28], [37], and [46]. The majority of these methods use state space augmentation approaches. Auxiliary distributions that allow a Markov chain to explore the entirety of the state space are targeted and their mixing information is then passed on to aid inter-modal mixing in the desired target. A convenient approach for the augmentation methods is to use power-tempered target distributions i.e. the target distribution at inverse temperature level, β , for $\beta \in (0, 1]$ is defined as

$$\pi_\beta(x) \propto [\pi(x)]^\beta.$$

Such targets are the most common choice of auxiliary target when augmenting the state space for use in the popular simulated tempering (ST) and parallel tempering (PT) algorithms introduced in [18] and [8]. For each algorithm, one needs to choose a sequence of $n + 1$ ‘inverse temperatures’, $\Delta = \{\beta_0, \dots, \beta_n\}$, where $0 \leq \beta_n < \beta_{n-1} < \dots < \beta_1 < \beta_0 = 1$ with the specification that a Markov chain sampling from the target distribution $\pi_{\beta_n}(x)$ can mix well across the entire state space.

The PT algorithm runs a Markov chain on the augmented state space, $\mathcal{X}^{(n+1)}$, targeting an invariant distribution given by

$$\pi_n(x_0, x_1, \dots, x_n) \propto \pi_{\beta_0}(x_0)\pi_{\beta_1}(x_1) \cdots \pi_{\beta_n}(x_n). \tag{1}$$

From an initialisation point for the chain, the PT algorithm alternates between two types of Markovian move: *within temperature* Markov chain moves that use standard localised MCMC schemes to update each of the x_i whilst preserving marginal invariance; *temperature swap* moves that propose to swap the chain locations between a pair of adjacent temperature components. It is these swap moves that will allow mixing information from the hot, rapidly mixing temperature level to be passed to aid mixing at the cold target state.

To perform the swap move, a pair of temperatures is chosen uniformly from the set of all adjacent pairs, call this pair x_i and x_{i+1} at inverse temperatures β_{x_i} and $\beta_{x_{i+1}}$, respectively. The proposal is then

$$(x_0, \dots, x_i, x_{i+1}, \dots, x_n) \rightarrow (x_0, \dots, x_{i+1}, x_i, \dots, x_n). \tag{2}$$

To preserve detailed balance and therefore invariance to $\pi_n(\cdot)$, the swap move is accepted with probability

$$A = \min\left(1, \frac{\pi_{\beta_{x_{i+1}}}(x_i)\pi_{\beta_{x_i}}(x_{i+1})}{\pi_{\beta_{x_i}}(x_i)\pi_{\beta_{x_{i+1}}}(x_{i+1})}\right).$$

It is the combination of the suitably specified *within temperature* moves and *temperature swap* moves that ensures ergodicity of the Markov chain to the target distribution, $\pi_n(\cdot)$. Note that the *within temperature* moves certainly influence the performance of the algorithm [7]; however, the focus of the work in this article will be on designing a novel approach for the *temperature swap* move.

The novel work presented in this paper focuses on the setting where the d -dimensional state space is given by \mathbb{R}^d and the target, $\pi(\cdot)$, is the associated probability density function. Thus, herein take $\mathcal{X} = \mathbb{R}^d$, but note that natural generalisations to other state spaces and settings are possible.

3. Modal rescaling transformation

3.1. A motivating transformation move

Consider a PT algorithm that has two components x_1 and x_2 running at the neighbouring inverse temperature levels β and β' . Suppose that a temperature swap move is proposed between the two chains at the two temperature levels. Due to the dependence between the location in the state space and the temperature level, β and β' need to be close to each other to avoid the move having negligible acceptance probability. Intuitively, the problem is that the proposal from the hotter chain is likely to be an ‘unrepresentative’ location at the colder temperature, and vice versa.

So there is clearly a significant dependence between the temperature value and the location of the chain in the state space, thus explaining why temperature swap moves between arbitrarily largely spaced temperatures are generally rejected. This issue is typically exacerbated when the dimensionality grows.

Consider, for motivational purposes, a simple one-dimensional setting where the state space is given by \mathbb{R} and the target density is given by $\pi(\cdot)$. For notational convenience, letting $j = i + 1$, suppose that a temperature swap move has been proposed between adjacent levels β_i and β_j with marginal component values x_i and x_j , respectively.

Suppose that an oracle has provided a function, $t_{ij}: \mathbb{R} \rightarrow \mathbb{R}$, that is bijective, with $t_{ji}(t_{ij}(x)) = x$, and differentiable and preserves the cumulative distribution function (CDF) between the two temperature levels such that

$$F_{\beta_j}(t_{ij}(x)) = F_{\beta_i}(x), \tag{3}$$

where $F_{\beta}(\cdot)$ denotes the CDF of $\pi^{\beta}(\cdot)$.

Suppose that rather than the standard temperature swap move proposal in (2), the following is instead proposed:

$$(x_0, \dots, x_i, x_j, \dots, x_n) \rightarrow (x_0, \dots, t_{ji}(x_j), t_{ij}(x_i), \dots, x_n).$$

To preserve detailed balance, this is accepted with an acceptance ratio similar to reversible-jump MCMC [10] to account for the deterministic transformation:

$$\min\left(1, \frac{\pi^{\beta_j}(t_{ij}(x_i))\pi^{\beta_i}(t_{ji}(x_j))}{\pi^{\beta_i}(x_i)\pi^{\beta_j}(x_j)} \left| \frac{\partial t_{ij}(x_i)}{\partial x} \right| \left| \frac{\partial t_{ji}(x_j)}{\partial x} \right| \right). \tag{4}$$

By differentiating (3) with respect to x and rearranging

$$\frac{\pi_{\beta_j}(t_{ij}(x))}{\pi_{\beta_i}(x)} \left| \frac{\partial t_{ij}(x)}{\partial x} \right| = 1. \tag{5}$$

Using the result in (5), and noting the cancellation of normalising constants in (4), we can see that (4) evaluates to 1 and, hence, such a swap would always be accepted. Essentially, the transformation $t_{ij}(\cdot)$ has made the acceptance probability of a temperature swap move independent of the locations of x_i and x_j in the state space.

In practice, a CDF-preserving transformation $t_{ij}(\cdot)$ will not generally be available. Consider a simplified setting when the target is now a d -dimensional Gaussian, i.e. $\pi \sim N(\mu, \Sigma)$, and so the tempered target at inverse temperature β is given by $\pi^{\beta} \sim N(\mu, \Sigma/\beta)$. Defining a d -dimensional transformation by

$$t_{ij}(x, \mu) = \left(\frac{\beta_i}{\beta_j}\right)^{1/2} (x - \mu) + \mu, \tag{6}$$

a simple calculation shows that in this setting such a transformation, which only requires knowledge of the mode location, permits swap moves to always be accepted independently of the dimensionality and magnitude of the inverse temperature spacings.

In a broad class of applications it is not unreasonable to make a Gaussian approximation to posterior modes [36]. Indeed, this is the motivation for the similar transformation derived in [12] for use in a reversible-jump MCMC framework.

Beyond a single-dimensional setting a quantile is typically nonunique since multiple locations in the state space produce the same CDF evaluation. However, the transformation in (6), motivated via Gaussianity, is bijective and can be seen as a way of uniquely pairing two quantiles at different temperature levels allowing for ‘quantile uniqueness’ and, therefore, reversibility of the transformation. It is from this that QuanTA gets its name.

3.2. Transformation-aided move in a PT framework

In a multimodal setting a single Gaussian approximation to the posterior will be poor. However, it is often reasonable that the local modes may be individually approximated as Gaussian. In this paper we explore the use of the transformation in (6) applied to the local mode with the aim being to accelerate the mixing through the temperature schedule of a PT algorithm.

Now that the transformations are localised to modes we need careful specification of the transformation function. Suppose that there is a collection of K -mode points, μ_1, \dots, μ_K , and a metric, $m(x, y)$ for $x, y \in \mathbb{R}^d$, that will be used to associate locations in the state space with a mode. To this end, define the mode allocating function

$$Z(x) = \arg \min_{h \in \{1, \dots, K\}} [m(x, \mu_h)].$$

Then, with $t_{ij}(\cdot)$ from (6), define the sets

$$A_{ij} = \{x \in \mathbb{R}^d : Z(t_{ij}(x, \mu_{Z(x)})) = Z(x)\}$$

and the transformation

$$t(x, \beta_i, \beta_j) = t_{ij}(x, \mu_{Z(x)}). \quad (7)$$

Fundamentally, the set A_{ij} is the subset of the β_i -level marginal state space where the transformation in (7) is reversible, i.e. one would remain associated with the same mode point after the transformation. An illustrative example is given in Figure 2.

The aim is to use the transformation in a PT framework. So suppose that a temperature swap move proposal is made between two marginal components x_i and x_j at respective inverse temperatures β_i and β_j with $\beta_i > \beta_j$. The idea is that this swap move now utilises (7) so that the proposed move takes the form

$$(x_0, \dots, x_i, x_j, \dots, x_n) \rightarrow (x_0, \dots, t(x_j, \beta_j, \beta_i), t(x_i, \beta_i, \beta_j), \dots, x_n), \quad (8)$$

which to satisfy detailed balance is accepted with probability

$$\min \left(1, \frac{\pi(t(x_i, \beta_i, \beta_j))^{\beta_j} \pi(t(x_j, \beta_j, \beta_i))^{\beta_i}}{\pi(x_i)^{\beta_i} \pi(x_j)^{\beta_j}} \mathbf{1}_{\{x_i \in A_{ij}\}} \mathbf{1}_{\{x_j \in A_{ji}\}} \right). \quad (9)$$

Proposition 1. *Consider a Markov chain that is in stationarity with a target distribution given by (1) on a state space $\mathcal{X} = \mathbb{R}^{(n+1)d}$. Let $\mu_1, \dots, \mu_K \in \mathbb{R}^d$. If a temperature swap move of the form (8) is proposed where the transformation is given by (7) and is accepted with probability given in (9) then the chain is invariant with respect to (1).*

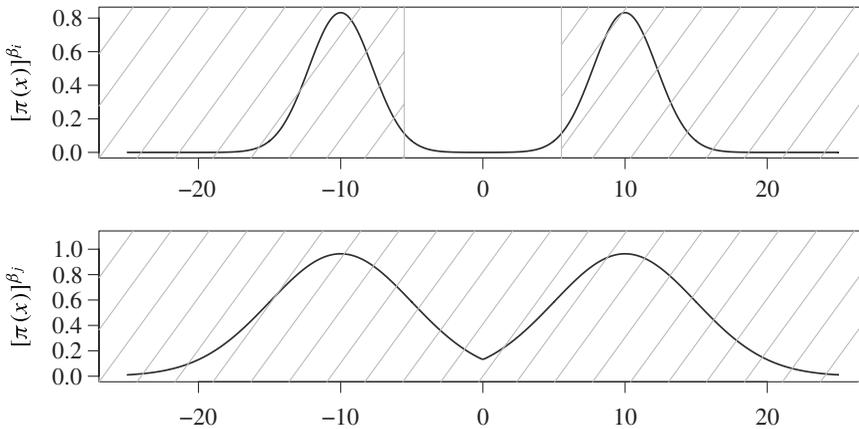


FIGURE 2: Plots of a simple bimodal Gaussian mixture distribution at two temperature levels β_i and β_j . The metric $m(\cdot, \cdot)$ that allocates locations to mode points is chosen to be the Euclidean metric on \mathbb{R} . The shaded regions represent A_{ij} and A_{ji} in the top and bottom plots, respectively.

4. Quantile tempering algorithm (QuanTA)

The aim is to utilise the transformation-aided temperature swap moves in a PT algorithm. However, to perform the transformation, we need a collection of K centring points. Typically, these are unknown *a priori* and must be estimated. There are three obvious directions to proceed in: an involved optimisation process prior to the run of the algorithm that searches for mode points which are then fixed therein; an adaptive scheme that learns the mode points during the run of the algorithm; or a population Monte Carlo approach which discovers a collection of mode points using the most recent location of the Markov chain in a way that ensures that the Markov chain remains reversible.

Fixing the centring points prior to the run of the algorithm can be highly nonrobust since it leaves no scope to adjust in situations when new regions of mass are found during the run of the Markov chain. Also, it requires careful and potentially expensive precomputation.

Using adaptive techniques that ‘learn’ the best centring points are therefore preferable. However, adaptive MCMC approaches where the proposal mechanism adapts throughout the run of the algorithm [32], [33] requires careful design to ensure ergodicity. In contrast, the population MCMC approach essentially allows the proposal mechanism to adapt at every iteration without concerns about affecting the ergodicity, albeit typically at a computational cost which can be somewhat mitigated by exploiting computer parallelism.

The algorithm presented below utilises a population MCMC approach with N ‘copies’ of the PT approach to provide a large population of points in the state space that can be used in an appropriate clustering procedure to obtain a collection of K centring points, suggestions for which are given in Section 5.

4.1. The procedure

QuanTA runs the equivalent of N parallel tempering algorithm procedures in parallel with each single procedure using the same tempering schedule. With a temperature schedule given by $\Delta = \{\beta_0, \dots, \beta_n\}$, the QuanTA approach can be seen as running a single Markov chain on the augmented state space, $(\mathbb{R}^d)^{n \times N}$. Letting $\mathbf{x} = (x_{(1,0)}, \dots, x_{(1,n)}, x_{(2,0)}, \dots, x_{(N,n)})$, the

invariant target distribution for the Markov chain induced by QuanTA is

$$\pi_Q(\mathbf{x}) \propto \prod_{i=1}^N \pi_{\beta_0}(x_{(i,0)})\pi_{\beta_1}(x_{(i,1)}) \cdots \pi_{\beta_n}(x_{(i,n)}).$$

Initialisation. To initialise the QuanTA algorithm, we need to choose initial starting values for the Markov chain components, a suitable temperature schedule (see Theorem 1 in Section 6.2 for a suggested optimality criteria for the temperature schedule), the size of N , and suitable parameters for the chosen clustering method that will be used.

Running the chain. From the start point of the chain, QuanTA alternates between two types of Markov chain moves.

1. *Within temperature* Markov chain moves that use standard localised MCMC schemes for marginal updates of each of the $x_{(i,j)}$. Essentially, this is just Metropolis-within-Gibbs MCMC and in this setting, with hugely exploitable marginal independence, this process is highly parallelisable. Denote the π_Q -invariant Markov transition kernel that performs temperature marginal updates on all components from a current point \mathbf{x} as $P_1(\mathbf{x}, d\mathbf{y})$.
2. *Temperature swap* moves that propose to swap the chain locations between a pair of adjacent temperature components. This is where QuanTA differs from the standard PT procedure and uses the new transformation-aided temperature swap move detailed in Section 3.2, in particular in (7). This follows a two-phase population-MCMC update procedure.

- *Phase 1.* Group marginal components into two collections,

$$C_1 = \{x_{(i,j)} : i = 1, \dots, \lfloor \frac{1}{2}N \rfloor \text{ and } j = 0, \dots, n\},$$

$$C_2 = \{x_{(i,j)} : i = (\lfloor \frac{1}{2}N \rfloor + 1), \dots, N \text{ and } j = 0, \dots, n\}.$$

An appropriate clustering scheme (see Section 5) is performed on C_1 , providing a set of K centres $\{c_1, \dots, c_K\}$. To enhance the effectiveness of the transformation, it is suggested that these cluster centre points are used as initialisation locations for a suitable local optimisation procedure to find K mode points $M_1 = \{\mu_1, \dots, \mu_K\}$ of $\pi(\cdot)$ (see Theorem 2 in Section 6.2).

For each $i \in \{(\lfloor N/2 \rfloor + 1), \dots, N\}$, sample $l \sim \text{Unif}\{0, 1, \dots, n - 1\}$ and select the corresponding pair of adjacent temperature marginals $(x_{(i,l)}, x_{(i,l+1)})$ for a temperature swap move proposal utilising the transformation from (7) (which is centred on the associated point from M_1). This move is accepted with probability (9).

- *Phase 2.* Repeat phase 1, but with the roles of C_1 and C_2 reversed.

Denote the π_Q -invariant Markov transition kernel that implements this temperature swap update procedure for all components using the above two-phase process by $P_2(\mathbf{x}, d\mathbf{y})$.

From the initialisation point \mathbf{x} , the Markov chain output is created by application of the kernel compilation

$$(P_2 \circ P_1^k)^T,$$

where k is the user-chosen number of within temperature Markov chain updates between each swap move proposal and T is the user-chosen number of iterations of the algorithm before stopping.

Remark 1. By Proposition 1, the transformation-aided temperature swap move established in Section 3.2 leaves the Markov chain constructed by QuanTA $\pi_Q(\cdot)$ invariant. Provided that

appropriate within-temperature MCMC moves are utilised then QuanTA establishes an irreducible and aperiodic $\pi_Q(\cdot)$ invariant Markov chain as required [42].

Remark 2. Although not made explicit in the notation above, the transformation centring points, $\{\mu_1, \dots, \mu_K\}$, are repeatedly estimated from subsets of the points in the population of the N replicate PT schemes. As such, their values are not fixed and have flexibility to change and adapt accordingly with the population. This is a major bonus of the population MCMC framework for a multimodal problem, allowing powerful and robust adaptation when new regions of mass are discovered.

5. Estimating local mode locations

The QuanTA algorithm, presented in Section 4, requires online estimation of the local mode points as centring locations for the transformation. In this section we outline a practical scheme that is used in the canonical simulation studies.

With a typically unknown number of modes and a population of chains, a principled approach would be to fit a Dirichlet process mixture model, e.g. [15] and [22]. A comprehensive Gibbs sampling approach for this can be computationally expensive, but there are alternative cheaper but approximate methods that are left for exploration in further work [29].

For the examples with well-separated modes that were studied here, it sufficed to use a cheap and fast clustering scheme [13]. To this end, a K-means approach was used [11], where although the theoretical computational complexity is NP-hard in full generality, when in a setting with well-separated clusters, the algorithm rapidly reaches a good solution as the problem is much easier.

The clustering procedure provides a collection of cluster centres that can be directly used as centring points for the transformation or as very useful initialisation points for a local optimisation method. Indeed, Theorem 2 of Section 6 shows that QuanTA can achieve accelerated mixing through the temperature levels when the centring point is chosen as the mode point, particularly at colder temperatures when the Gaussian approximation to the mode becomes increasingly accurate; see, e.g. [2] and [24].

5.1. A weighted K-means clustering

Typically, the K-means algorithm assigns all points equal leverage in determining cluster centres. A weighted K-means approach incorporates weights that can alter the leverages of points. In the tempering setting chains at the colder states, where the modes are less disperse, should have more leverage in determining the centres.

Weighted K means is an almost identical procedure to the K-means algorithm of [11] but now incorporates the weights to give points leverage. For the setting of interest, each chain location will be allocated a weight, determined by their inverse temperature value. So, for a collection of ν chain locations x_1, \dots, x_ν at inverse temperature levels, $\beta_{x_1}, \dots, \beta_{x_\nu}$ form their respective weights.

The weighted K-means algorithm gives a pragmatic way of creating a plausible particle allocation, \hat{S} , in to K partitions. Denoting a partition S by $S = \{S_1, \dots, S_K\}$, where $S_j \subset \{x_1, \dots, x_\nu\}$, then the procedure aims to establish an allocation, \hat{S} , such that

$$\hat{S} = \operatorname{argmin}_S \left\{ \sum_{i=1}^M \sum_{j=1}^{\nu} \mathbf{1}_{\{x_j \in S_i\}} \beta_{x_j} \|x_j - \mu_i\|^2 \right\}.$$

The weighted K-means algorithm begins with an initial set of K centres $\{\mu_1, \dots, \mu_K\}$. It then proceeds by alternating between two updating steps until point allocations do not change (signaling a minimum of or a pre-specified number of iterations is reached); a *point allocation step*, where each point x is assigned to the set S_j , where $j = \arg \min_j \|x - \mu_j\|^2$; a *centre point update step*, where, for the new allocation, the centring points are each updated to be the weighted mean of their respective component steps, i.e.

$$\mu_i = \frac{\sum_{j \in S_i} x_j \beta_{x_j}}{\sum_{j \in S_i} \beta_{x_j}}.$$

The weighted K-means procedure (see, e.g. [14] and [44]) can be implemented using the R package ‘FactoClass’ by Pardo and Del Campo [27], which uses a modified version of the K-means algorithm of Hartigan and Wong [11].

6. Theoretical underpinnings of QuanTA in high dimensions

In both QuanTA and the PT algorithms, the acceptance of temperature swap proposals allow the transfer of hot-state mixing information to be passed through to the cold state. The ambitiousness of the spacings between the consecutive inverse temperatures dictate the performance of the algorithm. Similarly to the problem of tuning the RWM algorithm [35], we seek the optimal balance between over- and under-ambitious proposals. This issue becomes increasingly problematic with an increase in dimensionality; hence, careful scaling of the consecutive temperature spacings is needed to prevent degeneracy of acceptance rates.

The work in [1] sought an optimal scaling result for temperature spacings in a PT algorithm. This section takes a similar approach to derive an equivalent result for QuanTA. It will be shown in Theorem 1 that consecutive spacings inducing swap rates of approximately 0.234 are optimal, thus giving guidance for practitioners to tune towards an optimal setup. Complementary to this, Theorem 2 justifies the use of QuanTA outside the Gaussian setting, showing that, under mild conditions, the transformation move allows for larger spacings in the temperature schedule than the PT algorithm does.

6.1. Optimal scaling of QuanTA—the setup and assumptions

As the dimensionality d of the target distribution tends to infinity, the problem of selecting temperature spacings for QuanTA is investigated. Suppose that a swap move between two consecutive temperature levels β and $\beta' = \beta + \epsilon$ for some $\epsilon > 0$ is proposed. As in [1], the measure of the efficiency of the inverse temperature spacing will be the expected squared jumping distance, $ESJD_\beta$, defined as

$$ESJD_\beta = \mathbb{E}_{\pi_n}[(\gamma - \beta)^2], \tag{10}$$

where $\gamma = \beta + \epsilon$ if the proposed swap is accepted and $\gamma = \beta$ otherwise. Note the assumption that the Markov chain has reached invariance and so the expectation is taken with respect to the invariant distribution $\pi_n(\cdot)$.

The $ESJD_\beta$ is a natural quantity to consider ([38]) since maximising this would appear to ensure that one is being sufficiently ambitious with spacings but not inducing degenerate acceptance rates. However, it is worth noting that it is only truly justified when there is an associated diffusion limit for the chain; see [34].

The aim is to establish the limiting behaviour of the $ESJD_\beta$ as $d \rightarrow \infty$ and then optimise this limiting form. To this end, for tractability, the form of the d -dimensional target is restricted

to distributions of the form

$$\pi(x) \propto f_d(x) = \prod_{i=1}^d f(x_i), \tag{11}$$

and to achieve a nondegenerate acceptance rate as $d \rightarrow \infty$ the spacings are necessarily scaled as $\mathcal{O}(d^{-1/2})$, i.e.

$$\epsilon = \frac{\ell}{d^{1/2}},$$

where ℓ a positive constant that one tunes to attain an optimal ESJD $_{\beta}$.

Furthermore, assume that the univariate marginal components $f(\cdot)$ are C^4 and have a global maxima at μ , i.e.

$$f'(\mu) = 0 \quad \text{and} \quad f(\mu) > f(z) \quad \text{for all } z \neq \mu. \tag{12}$$

The point μ is the centring point for the transformation-aided temperature swap move analysed in the forthcoming Theorem 1. Note that the result of Theorem 1 still holds even if μ is chosen as an arbitrary point in the state space. However, the global mode point μ is the canonical transformation centring point, as will be seen in the subsequent result in Theorem 2.

Furthermore, the marginal components $f(\cdot)$ are assumed to be of the form

$$f(x) = e^{-H(x)} \quad \text{for all } x \in \mathbb{R}, \tag{13}$$

where $H(x) := -\log(f(x))$ is regularly varying [3], i.e. there exists an $\alpha > 0$ such that, for $x > 0$,

$$\frac{H(tx)}{H(t)} \rightarrow x^{\alpha} \quad \text{as } |t| \rightarrow \infty. \tag{14}$$

This is a sufficient condition for Theorem 1 and ensures that the moments and integrals required for the proof are all well defined. Furthermore, assume that the fourth derivatives of $(\log f)(\cdot)$ are bounded, i.e. there exists an $M > 0$ such that

$$|(\log f)^{''''}(z)| < M \quad \text{for all } z \in \mathbb{R}. \tag{15}$$

This condition is sufficient for proving Theorem 1 but not necessary. The proof still works if the condition is weakened so that, for some $k \geq 4$, the k th derivative of the logged density is bounded.

Finally, for notational convenience, the following are defined, with the subscript β indicating that the expectation is with respect to $f^{\beta}(\cdot)$:

$$\begin{aligned} V(\beta) &= \frac{1}{\beta^2}, \\ I(\beta) &= \text{var}_{\beta}[(\log f)(x)], \\ R(\beta) &= \mathbb{E}_{\beta}[(x - \mu)^2 (\log f)''(x) - (x - \mu)(\log f)'(x)]. \end{aligned}$$

Theorem 1 below deals with a single-mode setting which is a very good approximation of a multimodal setting with homogeneous modal structure and where the mode weights remain stable at consecutive temperatures [40], [48]. Outside this setting other methods are needed to stabilise the modal weights at different temperatures [41]; then the QuanTA transformation-aided swap moves can be successfully utilised alongside such an approach.

Since we assume only one transformation-centring point, there is a simplified form of the acceptance probability that no longer requires the indicator functions. Denote the acceptance probability of the QuanTA-style swap move by $\alpha_\beta(x, y)$, so

$$\alpha_\beta(x, y) = \min\left(1, \frac{f_d^{\beta'}(g(x, \beta, \beta'))f_d^\beta(g(y, \beta', \beta))}{f_d^{\beta'}(y)f_d^\beta(x)}\right). \tag{16}$$

Then a simple calculation shows that the ESJD_β from (10) becomes

$$\text{ESJD}_\beta = \epsilon^2 \mathbb{E}_{\pi_n}[\alpha_\beta(x, y)], \tag{17}$$

which will be maximised with respect to ℓ in the limit as $d \rightarrow \infty$.

6.2. Scaling results and interpretation

Under the setting of Section 6.1 and with $\Phi(\cdot)$ denoting the CDF of a standard Gaussian, the following optimal scaling result is derived.

Theorem 1. (Optimal scaling for the QuanTA algorithm.) *Consider QuanTA targeting a distribution $\pi(\cdot)$ satisfying (11). Assume that the marginal components $f(\cdot)$ satisfy (12), are regularly varying satisfying (13) and (14), and that $\log f(\cdot)$ satisfies (15). Assuming that $\epsilon = \ell/d^{1/2}$ for some $\ell \in \mathbb{R}_+$ then in the limit as $d \rightarrow \infty$, the ESJD_β , given in (17) is maximised when ℓ is chosen to maximise*

$$2\ell^2 \Phi\left(-\frac{\ell[V(\beta)/2 - I(\beta) + R(\beta)/4\beta]^{1/2}}{\sqrt{2}}\right). \tag{18}$$

Furthermore, for the optimal ℓ , the corresponding swap move acceptance rate induced between two consecutive temperatures is given by 0.234 (to three significant figures).

Proof. The details of the proof of Theorem 1 are deferred to Appendix A. The strategy comprises three key stages: establishing a Taylor series expansion of the logged swap move acceptance ratio (i.e. the log of (16)); establishing limiting Gaussianity of this logged acceptance ratio; and, finally, achieving a tractable form of the limiting ESJD_β which is then optimised with respect to ℓ , giving rise to an associated optimal acceptance rate. \square

Remark 3. In the special case that the marginal targets are Gaussian, i.e. $f(x) = \phi(x; \mu, \sigma^2)$, then the transformation swap move should permit arbitrarily ambitious spacings. This is verified by observing that in this case

$$\left[\frac{1}{2}V(\beta) - I(\beta) + \frac{1}{4\beta}R(\beta)\right] = 0,$$

and so, with respect to ℓ , (18) becomes proportional to ℓ^2 which has no finite maximal value; thus demonstrating consistency with what is known in the Gaussian case.

Remark 4. The optimality criterion given in (18) is very similar to that derived in [1] and [34]. Indeed, both QuanTA and the PT algorithm require the same dimensionality spacing scaling, and both are optimised when a 0.234 acceptance rate is induced. However, there will be a difference in the behaviour of the optimal ℓ , which is where QuanTA can be shown to give accelerated mixing versus the PT approach; see Theorem 2 below.

Remark 5. Theorem 1 gives an explicit formula for derivation of the optimal $\hat{\ell}$ between consecutive temperatures, but this is usually intractable in a real problem. However, for a practitioner, the associated 0.234 optimal swap acceptance rate gives useful setup guidelines. In fact, the theorem suggests a strategy for optimal setup starting with a chain at the hottest level and tuning the spacing to successively colder temperature levels based on the swap acceptance rate to attain consecutive swap rates close to 0.234. Indeed, using a stochastic approximation algorithm (see [30]), Miasojedow *et al.* [19] took an adaptive MCMC approach [33] to do this for the PT algorithm, but their framework also extends naturally to QuanTA.

6.2.1. *Higher-order scalings at cold temperatures.* For any univariate Gaussian distribution at inverse temperature level β , $I(\beta) = 1/(2\beta^2)$. It was shown in [1] that the optimal choice for the scaling parameter takes the form

$$\hat{\ell} \propto I(\beta)^{-1/2} \propto \beta$$

resulting in a geometrically spaced temperature schedule.

Assuming appropriate smoothness for the marginal components $f(\cdot)$, for a sufficiently cold temperature, the local mode can be well approximated by a Gaussian [43]. So, for sufficiently cold temperatures (i.e. large β), we expect $I(\beta) \approx 1/(2\beta^2)$, thus spacings become (approximately) $\mathcal{O}(\beta)$ (note that a rigorous derivation that $I(\beta) \approx 1/(2\beta^2)$ is contained in the proof of Theorem 2). Defining the ‘order of the spacing with respect to the inverse temperature β ’ as the value of ζ such that the optimal spacing is $\mathcal{O}(\beta^\zeta)$, then the standard PT algorithm is order 1 for sufficiently cold temperatures.

In the Gaussian setting, QuanTA exhibits ‘infinitely’ high-order behaviour since there is no restriction on the size of the temperature spacings with regards the value of β . It is hoped that some of this higher-order behaviour is inherited in a more general target distribution setting when the modes can be considered to be sufficiently close to Gaussian.

It will be seen in Theorem 2 below that, under the setting of Theorem 1 but with a single additional condition, QuanTA does exhibit higher-order behaviour than the PT algorithm at cold temperatures.

With $f(\cdot)$ as in Theorem 1 (but now, without loss of generality, the mode point is at $\mu = 0$), define the normalised density $g_\beta(\cdot)$ as

$$g_\beta(y) \propto f^\beta\left(\mu + \frac{y}{\sqrt{-\beta(\log f)''(\mu)}}\right) = f^\beta\left(\frac{y}{\sqrt{-\beta(\log f)''(0)}}\right).$$

The additional assumption required to prove the higher-order behaviour of QuanTA is that there exists $\gamma > 0$ such that, as $\beta \rightarrow \infty$,

$$|\text{var}_{g_\beta}(Y^2) - 2| = \mathcal{O}\left(\frac{1}{\beta^\gamma}\right). \tag{19}$$

This assumption essentially guarantees the convergence to Gaussianity about the mode as $\beta \rightarrow \infty$. This assumption appears to be reasonable with studies of both gamma and student-t distributions, demonstrating a value of $\gamma = 1$; details can be found in [39].

Theorem 2. (Cold temperature scalings.) *For marginal targets $f(\cdot)$ satisfying the conditions of Theorem 1 and (19), for sufficiently large β ,*

$$\left[\frac{1}{2}V(\beta) - I(\beta) + \frac{1}{4\beta}R(\beta)\right] = \mathcal{O}\left(\frac{1}{\beta^k}\right),$$

where

- $k = \min\{2 + \gamma, 3\} > 2$ if f is symmetric about the mode point 0,
- $k = \min\{2 + \gamma, \frac{5}{2}\} > 2$ otherwise.

This induces an optimising value $\hat{\ell}$ such that

$$\hat{\ell} = \mathcal{O}(\beta^{k/2}),$$

showing that at the colder temperatures *QuanTA* permits higher-order behaviour than the standard *PT* scheme which has $\hat{\ell} = \mathcal{O}(\beta)$.

Proof. Since the optimal ℓ derived in Theorem 1 is given by

$$\hat{\ell} \propto \left[\frac{1}{2}V(\beta) - I(\beta) + \frac{1}{4\beta}R(\beta) \right]^{-1/2},$$

the proof of Theorem 2 follows immediately if it can be shown that

$$\left[\frac{1}{2}V(\beta) - I(\beta) + \frac{1}{4\beta}R(\beta) \right] = \mathcal{O}\left(\frac{1}{\beta^k}\right). \tag{20}$$

Indeed, two key lemmata are derived in Appendix A.2: Lemma 4 establishes that $\frac{1}{2}V(\beta) - I(\beta) = \mathcal{O}(1/\beta^k)$ and Lemma 5 establishes that $R(\beta)/4\beta = \mathcal{O}(1/\beta^k)$. Thus, the result in (20) holds and the proof is complete. \square

Remark 6. Theorem 2 studies the behaviour of the optimal spacings for large values of β , where β is larger than 1. Even though the *QuanTA* and *PT* approaches use temperatures in the range $[0, 1]$, the result still proves insightful in many situations that *QuanTA* can encounter.

This is due to the arbitrary nature of the ‘ $\beta = 1$ ’ temperature level which is a result of the tempering being relative rather than absolute for the modes. Hence, the modes of the target distribution of interest may exhibit a within mode structure that makes it appear that a mode has been cooled to an inverse temperature much larger than 1.

As an example, take a $\text{Gamma}(\alpha, \gamma)$ distribution, where, for some $\alpha > 0$, $\gamma > 0$, and $x > 0$,

$$f(x) \propto x^{\alpha-1} e^{-\gamma x}.$$

Then, with slight abuse of notation, $f^{\beta^*} \sim \text{Gamma}(\alpha\beta^* - \beta^* + 1, \gamma\beta^*)$. Suppose that the target distribution of interest was given by $\pi \sim f^{\beta^*}$ and one raises π to a power $\beta < 1$. Then $\pi^\beta \sim \text{Gamma}(\alpha\beta\beta^* - \beta\beta^* + 1, \gamma\beta\beta^*)$. Thus, if β^* is large and $\beta < 1$ is such that $\beta\beta^*$ remains large, then one is in the setting of Theorem 2.

Furthermore, ongoing work is hoping to develop novel methodology for sampling from multimodal distributions that utilises temperature levels with $\beta > 1$ where necessarily $\beta \rightarrow \infty$ as $d \rightarrow \infty$. The results established here are directly applicable to that work.

Remark 7. The result in Theorem 2 does not imply that *QuanTA* is not useful outside the Gaussian or super cold settings. The *QuanTA* approach will be practically useful in settings where the mode can be well approximated by a Gaussian and thus the transformation in the transformation-aided swap move approximately preserves the mode’s CDF. What Theorem 2 does show is that, for a large class of distributions that exhibit appropriate smoothness, *QuanTA* is sensible, and is arguably the canonical approach to take at the colder temperature levels, since it enables accelerated mixing speed through the temperature schedule.

7. Examples of implementation

In this section we give illustrative examples for the canonical setting of a Gaussian mixture to illustrate the potential gains of QuanTA over the standard PT approach.

The QuanTA transformation move does not solve all the issues inherent in the PT framework. This will be highlighted with the final example in this section. In fact, Woodard *et al.* [47], [48] showed that, for most ‘interesting’ examples, the mixing speed decays exponentially slowly with dimension. Prototype approaches to navigating this problem can be found in [39] and [41]. Also, note that the examples all focus on the setting when the target distribution is a Gaussian mixture, which is often a good approximation to a range of multimodal distributions. Outside the setting of Gaussian modes the transformation-aided swap move will not be as effective since the within-mode CDF preservation transformation will only be an approximation to preservation of the CDF.

In each of the examples given, both the new QuanTA and standard (PT) parallel schemes will be run for comparison of performance. In all examples the following applies.

1. Both the new QuanTA and PT versions were run 10 times to ensure replicability.
2. Both the PT and QuanTA algorithms were run so that 20 000 swap moves would be attempted. For QuanTA, this would be 20 000 swaps for each of the N individual parallel tempering schemes in parallel of which there were $N = 100$ in this example. Also, all schemes had the same within to swap move ratio (3 : 1).
3. Both versions use the same set of (geometrically generated) temperature spacings, chosen to be overly ambitious for the PT setup but demonstrably under ambitious for the new QuanTA scheme.
4. Also presented is the optimal temperature schedule for the PT setup generated under the optimal acceptance rate of 0.234 for the PT algorithm suggested in [1]. This demonstrates the extra complexity needed to produce a functioning algorithm for the PT approach.
5. For all runs, the within temperature level proposals were made with Gaussian RWM moves tuned to an optimal 0.234 acceptance rate (see [35]).

7.1. One-dimensional example

The target distribution given by

$$\pi(x) \propto \sum_{k=1}^5 w_k \phi(x; \mu_k, \sigma^2), \quad (21)$$

where $\phi(\cdot; \mu, \sigma^2)$ is the density function of a univariate Gaussian with mean μ and variance σ^2 . In this example, $\sigma = 0.01$, the mode centres are given by $(\mu_1, \mu_2, \mu_3, \mu_4, \mu_5) = (-200, -100, 0, 100, 200)$ and all modes are equally weighted with $w_1 = w_2 = \dots = w_5$.

The temperature schedule for this example is given by a geometric schedule with an ambitious 0.0002 common ratio for the spacings. Only three levels are used and so the temperature schedule is given by $\Delta = \{1, 0.0002, 0.0002^2\}$; see Figure 3.

In all runs all the chains were started from a start location of -200 . In Figure 1, from the introductory section, we presented two representative trace plots of the target state chain for a run of the PT algorithm and a single scheme from QuanTA. There is a clear improvement in the inter-modal mixing for the QuanTA.

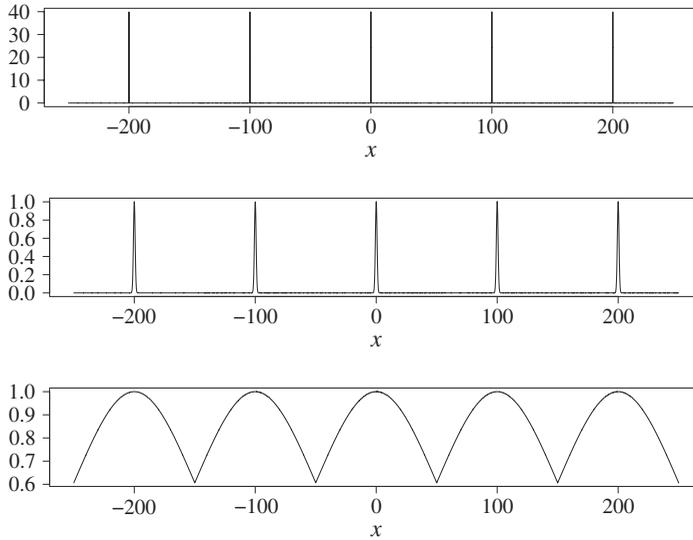


FIGURE 3: The (non-normalised) tempered target distributions for (21) for inverse temperatures $\Delta = \{1, 0.0002, 0.0002^2\}$, respectively.

TABLE 1: Comparison of the acceptance rates of swap moves for the PT algorithm and QuanTA targeting the one-dimensional distribution given in (21) and setup with the ambitious inverse temperature schedule given by $\Delta = \{1, 0.0002, 0.0002^2\}$.

	Swap location	
	1	2
PT	0.06	0.07
QuanTA	0.99	0.99

In Table 1 we give the associated acceptance rates. Clearly, the rate of transfer of mixing information from the hot states to the cold state is significantly higher for QuanTA.

In Figure 4 we compare the running modal weight approximation for the mode centred on 200 when using the standard PT and QuanTA schemes. This used the cold state chains from 10 individual runs of the PT algorithm and 10 single schemes selected randomly from 10 separate runs of the QuanTA algorithm.

Denoting the estimator of the k th mode’s weight by \hat{w}_k and the respective cold state chain’s i th value as X_i ,

$$\hat{w}_k = \frac{1}{N - B + 1} \sum_{i=B}^N \mathbf{1}_{\{c_k < X_i \leq C_k\}},$$

where c_k and C_k are the chosen upper and lower boundary points for allocation to the k th mode, and B is the length of the burn-in removed.

Figure 4 reveals that the QuanTA approach has a vastly improved rate of convergence; with the PT runs still exhibiting bias from the chain initialisation locations.

An interesting comparison between the approaches is to observe how many extra temperature levels would be required to make the PT scheme work optimally (i.e. with consecutive

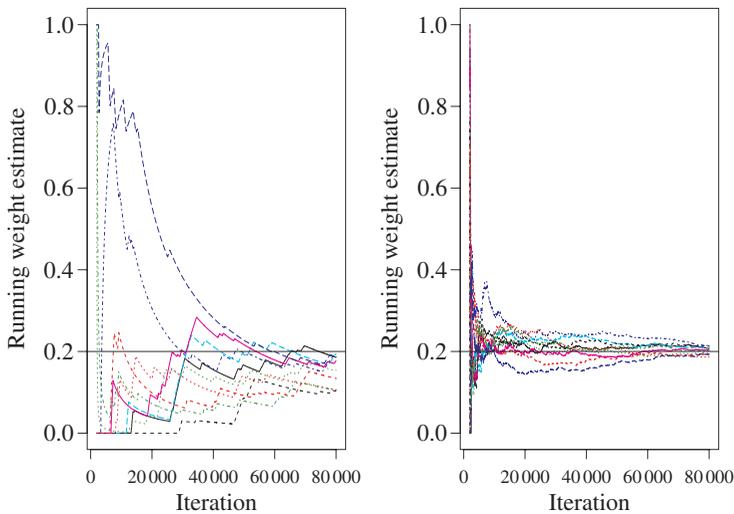


FIGURE 4: For the target given in (21), the running weight approximations for the mode centred on 200 with target weight $w_5 = 0.2$ for 10 separate runs of the PT and QuanTA schemes. *Left*: the PT runs showing slow and variable estimates for w_5 . *Right*: the new QuanTA scheme showing fast, unbiased convergence to the true value for w_5 .

0.234 swap acceptance rates). This gives a clearer idea of the reduction in the number of intermediate levels that can be achieved using the QuanTA.

With the same hottest state level of $\beta = 0.0002^2$, a geometrical inverse temperature schedule was tuned to give a swap rate of approximately 0.234 between consecutive levels for the PT algorithm in this example. In fact, a 0.04 geometric ratio suggested optimality for the PT scheme. Hence, to reach the stated hottest level, seven temperatures are needed, as opposed to the three that were evidently unambitious for QuanTA.

7.2. Twenty-dimensional example

The target distribution is a twenty-dimensional tri-modal Gaussian:

$$\pi(x) \propto \sum_{k=1}^3 w_k \left[\prod_{j=1}^{20} \phi(x_j; \mu_k, \sigma^2) \right]. \quad (22)$$

In this example, $\sigma = 0.01$, the marginal mode centres are given by $(\mu_1, \mu_2, \mu_3) = (-20, 0, 20)$, and all modes are equally weighted with $w_1 = w_2 = w_3$. The temperature schedule for this example is derived from a geometric schedule with an ambitious 0.002 common ratio for the spacings. Only four levels are used and so the temperature schedule is given by $\{1, 0.002, 0.002^2, 0.002^3\}$.

In all runs all the chains were started from a start location of $(-20, \dots, -20)$. In Figure 5 we show two representative trace plots of the target state chain for a run of the PT algorithm and QuanTA. There is a clear improvement in the inter-modal mixing for the new QuanTA scheme. There is a stark contrast between the two algorithmic performances. The run using the standard PT scheme entirely fails to improve the mixing of the cold chain. In contrast, the QuanTA scheme establishes a chain that is very effective at escaping the initialising mode and then mixes rapidly throughout the state space between the three modes.

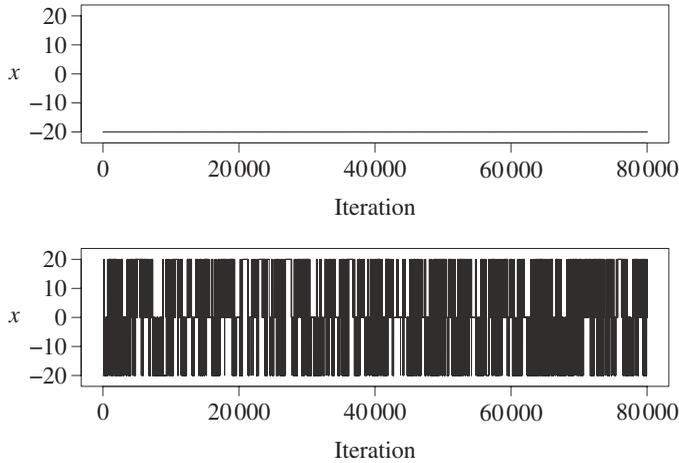


FIGURE 5: Trace plots of the first component of the twenty dimensional cold state chains for representative runs of the PT (top) and new QuanTA (bottom) schemes. Note the fast inter-modal mixing of the new QuanTA scheme, allowing rapid exploration of the target distribution. In contrast the PT scheme never escapes the initialising mode.

TABLE 2: Comparison of the acceptance rates of swap moves for the PT algorithm and QuanTA targeting the twenty-dimensional distribution given in (22) and setup with the ambitious inverse temperature schedule given by $\{1, 0.002, 0.002^2, 0.002^3\}$.

	Swap location		
	1	2	3
PT	0	0	0
QuanTA	0.99	0.99	0.99

The consecutive swap acceptance rates between the four levels are given in Table 2. Clearly, there is no transfer of mixing information from the hot states to the cold state for the PT algorithm, but there is excellent mixing in the QuanTA scheme.

The temperature schedule choice that induces a 0.234 swap acceptance rate between consecutive temperature levels for this example using the PT algorithm indicates a geometric schedule with a 0.58 common ratio. This is in stark contrast to the 0.002 ratio that is evidently underambitious for QuanTA. Indeed, to reach the allocated hot state of $\beta = 0.002^3$ then the PT algorithm would need 36 temperature levels in contrast to the 4 that sufficed for QuanTA.

7.3. Five-dimensional noncanonical example

Leaving the canonical symmetric mode setting, the following example has a five-dimensional Gaussian mixture target with even weight to the modes but with different covariance scaling within each mode. The target distribution is given by

$$\pi(x) \propto \sum_{k=1}^3 w_k \left[\prod_{j=1}^5 \phi(x_j; \mu_k, \sigma_k^2) \right]. \tag{23}$$

In this example, $(\sigma_1, \sigma_2, \sigma_3) = (0.02, 0.01, 0.015)$, the marginal mode centres are given by $(\mu_1, \mu_2, \mu_3) = (-20, 0, 20)$, and all modes are equally weighted with $w_1 = w_2 = w_3$.

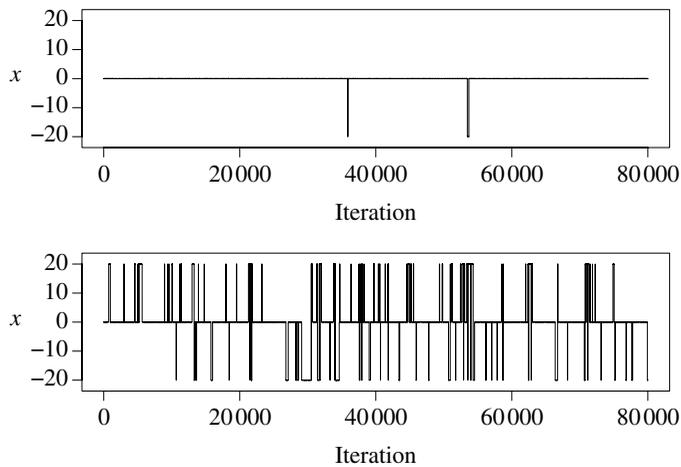


FIGURE 6: Trace plots of the first component of the five dimensional cold state chains for representative runs of the PT and QuanTA schemes respectively. Note the difference in inter-modal mixing between the QuanTA scheme and the PT scheme which struggles to escape the initialisation mode.

Although at first glimpse this does not sound like a significantly harder problem, or even far from the canonical setting, the differing modal scalings make this a much more complex example. This is due to the lack of preservation of modal weight through power-based tempering [48]. Indeed, Tawn and Roberts [40] and Tawn *et al.* [41] presented a novel modification to the auxiliary tempered targets that remain compatible with the QuanTA scheme.

The temperature schedule for this example cannot be a simple geometric schedule as in the previous example due to the scaling indifference between the modes. By using an ambitious geometric schedule, the clustering was very unstable early on and this often led to an inability to establish mode centres for the run. Instead, a mixture of geometric schedules was used with an ambitious spacing for the coldest levels and then a less ambitious spacing for the hotter levels. For the four coldest states, an ambitious geometric schedule with 0.08 common ratio was used. A further eight hotter levels were added using a conservative geometric schedule with ratio 0.4. Hence, the schedule was given by

$$\Delta = \{1, 0.08, 0.08^2, 0.08^3, 0.4^9, 0.4^{10}, \dots, 0.4^{15}, 0.4^{16}\}. \quad (24)$$

For the QuanTA scheme, the transformation moves were used for swap moves between the coldest seven levels and standard swap moves were used otherwise.

In Figure 6 we show two representative trace plots of the target state chain for a run of the PT and QuanTA algorithms. There is a clear improvement in the inter-modal mixing for the QuanTA scheme; albeit far less stark than that in the canonical one-dimensional and twenty-dimensional examples already shown. The run using the standard PT scheme fails to explore the state space. The QuanTA scheme establishes a chain that is able to explore the state space but does appear to have a bit of trouble during burn-in; mixing is good therein.

The consecutive swap acceptance rates between the 12 levels are given in Table 3. Clearly, there is very poor mixing through the four coldest states for the PT algorithm. In contrast, the QuanTA scheme has solid swap acceptance rates through the coldest levels but, unlike the previous examples, they are not all close to 1.

TABLE 3: Comparison of the acceptance rates of swap moves for the PT and new QuanTA algorithms targeting the five-dimensional distribution given in (23) and setup with the ambitious inverse temperature schedule given in (24). Note that, for QuanTA, the reparametrised swap move was only used for swaps in the coldest seven levels.

	Swap location					
	1	2	3	4	5	6
PT	0.001	0.0161	0.0138	0.469	0.317	0.348
QuanTA	0.446	0.970	0.997	0.999	0.999	0.999
	7	8	9	10	11	–
PT	0.328	0.334	0.359	0.324	0.327	–
QuanTA	0.285	0.285	0.285	0.285	0.302	–

This example is both positive (showing the improved mixing using the QuanTA scheme on a hard example), but also serves as a warning for the degeneracy of both the PT and new QuanTA schemes when using power-based tempering on a target outside of the canonical symmetric mode setting.

7.4. The computational cost of QuanTA

It is important to analyse the computational cost of QuanTA. To be an effective algorithm, the inferential gains of QuanTA per iteration should not be outweighed by the increase in runtime.

The analysis uses the runs of the one-dimensional and twenty-dimensional examples given above, using both the QuanTA and PT approaches. The algorithms were setup the same as in the ambitious versions of the spacing schedules in each case.

The key idea is to first establish the total runtime, denoted by R , in each case. Typically, one looks to compare the time-standardised effective sample size. In this case it is natural to take the acceptance rate as a direct proxy for the effective sample size. This is due to the fact that the target distributions have symmetric modes with equal weights. Hence, the acceptance rate between consecutive temperature levels dictates the performance of the algorithm, in particular the quality of inter-modal mixing.

To this end, taking the first level temperature swap acceptance rate, denoted by A , the runs are compared using runtime standardised acceptance rates, i.e. A/R .

Note that in both dimensional cases, the output from QuanTA is 100 times larger due to the use of 100 schemes running in parallel. Hence, for a standardised comparison, the time was divided by 100. Therefore, in what follows in this section, when the runtime R of the QuanTA approach is referred to, this means the full runtime divided by 100. The fairness of this is discussed below.

As can be seen in Tables 4 and 5, the QuanTA approach has a longer runtime to generate the same amount of output, as would be expected due to the added cost of clustering. Indeed, it takes approximately 1.5 times longer to generate the ‘same amount of output’. However, the temperature swap move acceptance rates are 16.5 and ∞ times better, respectively, when using the QuanTA approach. Using the acceptance rate as a proxy for the effective sample size, the quantity A/R is the fundamental value for comparison. In both cases the QuanTA approach shows a significant improvement over the PT approach.

TABLE 4: Complexity comparisons between QuanTA and PT for the one-dimensional example.

Algorithm	PT	QuanTA
Runtime (seconds)	5.60	8.01
Swap rate	0.06	0.99
A/R	0.01	0.12

TABLE 5: Complexity comparisons between QuanTA and PT for the twenty-dimensional example.

Algorithm	PT	QuanTA
Runtime (seconds)	8.00	12.79
Swap rate	0.00	0.99
A/R	0.00	0.08

There are the following issues with the fairness of this comparison.

- Standardising the runtime of QuanTA by the number of parallel schemes is not fully fair since it is sharing out the clustering expense between schemes.
- The spacings are too ambitious for the PT approach, meaning that the acceptance rates are very low. For a complete analysis, one should run the PT algorithm on its optimal temperature schedule and then use the time-standardised effective sample size from each of the optimised algorithms.

The empirical computational studies are favourable to the QuanTA approach. This is for a couple of examples that are canonical for QuanTA. Outside of this canonical setting the improvements from running QuanTA will be less obvious.

8. Conclusion and further work

The prototype QuanTA approach utilises a noncentred transformation approach to accelerate the transfer of mixing information from the rapidly mixing ‘hot’ state to aid the inter-modal mixing in the target ‘cold’ state. Examples show that this novel algorithm has the potential to dramatically improve the inferential gains, particularly in settings where the modes are similar to a Gaussian in structure.

The accompanying theoretical results that are given in Section 6 show that in a generic non-Gaussian setting the QuanTA approach can still exhibit accelerated mixing through the temperature schedule. Although the inverse temperature spacings are generally still $\mathcal{O}(d^{-1/2})$ there is a higher-order behaviour exhibited in the mixing for large (i.e. cold) values of the inverse temperature β . This suggests that the QuanTA approach will be powerful for accelerating the mixing through the colder levels of the temperature schedule for a typical smooth target.

It is clear that there are interesting questions to be addressed and further work needed before QuanTA can be considered practical in a real-data problem. In terms of optimising the computational expense, it has been shown that parallelisation of the PT algorithm can give significant practical gains [45]; by design, QuanTA is also highly parallelisable. The procedure described in Section 4.1 has huge scope to be parallelised. This requires care since

communication costs and synchronisation between parallel collections of the PT schemes could negate the effectiveness of the parallelisation.

A criticism of the current clustering method used is that it requires prior specification of the number of modes K which is likely to be unknown and would need online estimation as part of the clustering process.

An interesting question is whether using extra cold levels (with $\beta > 1$) along with the weighted clustering would help to aid the stability of the clustering procedure once invariance is reached for the population. Indeed, the mixing at these auxiliary extra cold levels should be very fast due to the QuanTA exhibiting higher-order behaviour in these modes. The other interesting question is regarding the robustness of the method outside of Gaussian modes, e.g. in heavier tailed modes, when the Gaussian approximation to the mode can be poor. Consider the setting of a univariate Laplace distribution and observe that the QuanTA-style transformation never agrees with the true CDF preserving transformation. It would be interesting to search, both empirically and theoretically, for settings where QuanTA fails to outperform the standard PT swap move. Some initial ideas and details of this further work can be found in [39].

Appendix A. Proofs

In this section we give the proof details of the results in Section 6.2. Firstly, some key notation is introduced that will be useful throughout this section.

Definition 1. Let

- $B = \log \left(\frac{f_d^{\beta'}(t(x, \beta, \beta')) f_d^\beta(t(y, \beta', \beta))}{f_d^{\beta'}(y) f_d^\beta(x)} \right)$;
- $h(x) := \log(f(x))$;
- $k(x) := (x - \mu)h'(x)$; and
- $r(x) := (x - \mu)^2 h''(x)$.

Then define

$$\begin{aligned}
 M(\beta) &= \mathbb{E}_\beta(h(z)), \\
 S(\beta) &= \mathbb{E}_\beta(k(z)), \\
 R(\beta) &= \mathbb{E}_\beta(r(z) - k(z)),
 \end{aligned}
 \tag{25}$$

where all expectations are with respect to the distribution $f^\beta(x)/Z_\beta$ with $Z_\beta = \int f^\beta(z) dz$.

Proposition 2. Under the notation and assumptions of Theorem 1 and Definition 1, it can be shown that

$$I(\beta) := M'(\beta) = \text{var}_\beta(h(x)) \tag{26}$$

and

$$S(\beta) = -\frac{1}{\beta}, \tag{27}$$

which trivially gives

$$V(\beta) := S'(\beta) = \frac{1}{\beta^2}.$$

Proof. The proof of (26) is routine and can be found in [1]. The derivation of (27) is less obvious using integration by parts:

$$\begin{aligned}
 S(\beta) &= \int (x - \mu)(\log f)'(x) \frac{f^\beta(x)}{Z_\beta} dx \\
 &= \int (x - \mu) f'(x) \frac{f^{\beta-1}(x)}{Z_\beta} dx \\
 &= \left[\frac{(x - \mu) f^\beta(x)}{\beta Z_\beta} \right]_{-\infty}^{\infty} - \int \frac{1}{\beta} \frac{f^\beta(x)}{Z_\beta} dx \\
 &= - \int \frac{1}{\beta} \frac{f^\beta(x)}{Z_\beta} dx. \quad \square
 \end{aligned}$$

A.1. Proof of Theorem 1

In this section we derive three key results that are specific to deriving the result in Theorem 1. Lemma 1 will establish a Taylor-expanded form of the log acceptance ratio of a temperature swap move that will prove to be asymptotically useful. Lemma 2 will then establish the limiting Gaussianity of this logged acceptance ratio and, finally, Lemma 3 completes the proof of Theorem 1 by establishing the optimal spacings and associated optimal acceptance rates required.

Lemma 1. (QuanTA log-acceptance ratio.) *Under the notation and assumptions of Theorem 1 and Definition 1,*

$$\begin{aligned}
 B &= \epsilon \left[\sum_{i=1}^d h(x_i) - h(y_i) + \frac{1}{2}(k(y_i) - k(x_i)) \right] \\
 &\quad + \frac{\epsilon^2}{8\beta} \left[\sum_{i=1}^d r(x_i) - k(x_i) + r(y_i) - k(y_i) \right] + (T_x + T_y),
 \end{aligned}$$

where both $T_x \rightarrow 0$ and $T_y \rightarrow 0$ in probability as $d \rightarrow \infty$.

Proof. By taking logarithms, it is immediate that

$$\begin{aligned}
 B &= \sum_{i=1}^d [\beta' h(t(x_i, \beta, \beta')) - \beta h(x_i)] + \sum_{i=1}^d [\beta h(t(y_i, \beta', \beta)) - \beta' h(y_i)] \\
 &=: H_{\beta'}^{\beta'}(\mathbf{x}) + H_{\beta'}^{\beta}(\mathbf{y})
 \end{aligned}$$

With the aim being to derive the asymptotic behaviour of the log-acceptance ratio, the next step is to use Taylor expansions (in ϵ) to appropriate order so that the asymptotic behaviour of B can be understood.

For notational convenience, we

- make $h(t(x, \beta, \beta'))$ explicitly dependent on ϵ :

$$\alpha_x(\epsilon) := h(t(x, \beta, \beta')) = \log \left[f \left(\left(\frac{\beta}{\beta + \epsilon} \right)^{1/2} (x - \mu) + \mu \right) \right];$$

- define

$$d_x(\epsilon) := \left(\frac{\beta}{\beta + \epsilon} \right)^{1/2} (x - \mu) + \mu.$$

By a Taylor series expansion in ϵ , for fixed x , with Taylor remainder correction term denoted by ξ_x such that $0 < \xi_x < \epsilon$,

$$h(g(x, \beta, \beta')) = \alpha_x(\epsilon) = \alpha_x(0) + \epsilon \alpha'_x(0) + \frac{\epsilon^2}{2} \alpha''_x(0) + \frac{\epsilon^3}{6} \alpha'''_x(\xi_x), \tag{28}$$

where

$$\alpha'_x(\epsilon) = -\frac{(x - \mu)}{2} \frac{\beta^{1/2}}{(\beta + \epsilon)^{3/2}} (\log f)'(d_x(\epsilon)), \tag{29}$$

$$\begin{aligned} \alpha''_x(\epsilon) = & \frac{(x - \mu)^2}{4} \frac{\beta}{(\beta + \epsilon)^3} (\log f)''(d_x(\epsilon)) \\ & + \frac{3(x - \mu)}{4} \frac{\beta^{1/2}}{(\beta + \epsilon)^{5/2}} (\log f)'(d_x(\epsilon)), \end{aligned} \tag{30}$$

$$\begin{aligned} \alpha'''_x(\epsilon) = & -\frac{(x - \mu)^3}{8} \frac{\beta^{3/2}}{(\beta + \epsilon)^{9/2}} (\log f)'''(d_x(\epsilon)) \\ & - \frac{9(x - \mu)^2}{8} \frac{\beta}{(\beta + \epsilon)^4} (\log f)''(d_x(\epsilon)) \\ & - \frac{15(x - \mu)}{8} \frac{\beta^{1/2}}{(\beta + \epsilon)^{7/2}} (\log f)'(d_x(\epsilon)). \end{aligned} \tag{31}$$

As a preview to the later stages of this proof, the terms up to second order in ϵ dictate the asymptotic distribution of B . However, to show that the higher-order terms ‘disappear’ in the limit as $\epsilon \rightarrow 0$, a careful analysis is required. Thus, the next step is to establish that, under the assumptions made above, the higher-order terms converge to 0 in probability.

To this end, a careful analysis of $\alpha'''_x(\cdot)$ is undertaken. Firstly, it will be shown that $|\mathbb{E}_\beta[\alpha'''_x(\xi_x)]|$ is bounded; then application of Markov’s inequality will establish that the higher-order terms converge to 0 in probability as $d \rightarrow \infty$. Define

$$\eta_\epsilon := \left[\left(\frac{\beta}{\beta + \epsilon} \right)^{1/2} - 1 \right],$$

so that

$$d_x(\epsilon) - x = \left[\left(\frac{\beta}{\beta + \epsilon} \right)^{1/2} - 1 \right] (x - \mu) := \eta_\epsilon (x - \mu),$$

which has the property that $\eta_\epsilon \rightarrow 0$ as $d \rightarrow \infty$ and $|\eta_\epsilon| \leq 1$.

Then, with Taylor remainder correction terms denoted by ξ_1^ϵ , ξ_2^ϵ , and ξ_3^ϵ such that $0 < |\xi_k^\epsilon - x| < |d_x(\epsilon) - x|$,

$$\begin{aligned} (\log f)'(d_x(\epsilon)) = & (\log f)'(x) + \eta_\epsilon (x - \mu) (\log f)''(x) + \frac{\eta_\epsilon^2 (x - \mu)^2}{2} (\log f)'''(x) \\ & + \frac{\eta_\epsilon^3 (x - \mu)^3}{6} (\log f)''''(\xi_1^\epsilon), \end{aligned} \tag{32}$$

$$(\log f)''(d_x(\epsilon)) = (\log f)''(x) + \eta_\epsilon(x - \mu)(\log f)'''(x) + \frac{\eta_\epsilon^2(x - \mu)^2}{2}(\log f)''''(\xi_2^\epsilon), \tag{33}$$

$$(\log f)'''(d_x(\epsilon)) = (\log f)'''(x) + \eta_\epsilon(x - \mu)(\log f)''''(\xi_3^\epsilon). \tag{34}$$

Recall the assumptions (14) and (15). Substituting (32), (33), and (34) into (31), evaluating the expectation with respect to $X \sim f^\beta$, and, for convenience, denoting $|x - \mu|$ by S , then there exists $C \in \mathbb{R}_+$ such that

$$\begin{aligned} |\mathbb{E}_\beta[\alpha_x'''(\xi_x)]| &\leq \mathbb{E}_\beta[|\alpha_x'''(\xi_x)|] \\ &\leq \mathbb{E}_\beta \left[\frac{S^3}{8} \beta^{-3} |(\log f)''''(d(\xi_x))| + \frac{9S^2}{8} \beta^{-3} |(\log f)''(d(\xi_x))| \right. \\ &\quad \left. + \frac{15S}{8} \beta^{-3} |(\log f)'(d(\xi_x))| \right] \\ &\leq \mathbb{E}_\beta \left[\frac{S^3}{8} \beta^{-3} (|(\log f)''''(x)| + S |(\log f)''''(\xi_3^{\xi_x})|) \right. \\ &\quad \left. + \frac{9S^2}{8} \beta^{-3} \left(|(\log f)''(x)| + S |(\log f)''(x)| + \frac{|x|^2}{2} |(\log f)''''(\xi_2^{\xi_x})| \right) \right. \\ &\quad \left. + \frac{15S}{8} \beta^{-3} \left(|(\log f)'(x)| + S |(\log f)''(x)| + \frac{S^2}{2} |(\log f)''''(x)| \right. \right. \\ &\quad \left. \left. + \frac{S^3}{6} |(\log f)''''(\xi_1^{\xi_x})| \right) \right] \\ &\leq C, \end{aligned} \tag{35}$$

where the first three inequalities are from the direct application of the triangle inequality (with the second also using the boundedness of η_ϵ); whereas the final inequality arises from both the finiteness of expectations of the terms involving derivatives of order three or below (this is due to the regularly varying tails of $\log(f(\cdot))$) and the assumption that $|(\log f)''''(\cdot)| < M$.

Using (28), with substitution of terms from (29), (30), and (31), $H_\beta^{\beta'}(x)$ can be expressed as

$$\begin{aligned} H_\beta^{\beta'}(x) &= \sum_{i=1}^d (\beta + \epsilon) [\alpha_{x_i}(\epsilon) - \beta \alpha_{x_i}(0)] \\ &= \epsilon \sum_{i=1}^d [\alpha_{x_i}(0) + \beta \alpha'_{x_i}(0)] + \epsilon^2 \sum_{i=1}^d \left[\frac{\beta}{2} \alpha''_{x_i}(0) + \alpha'_{x_i}(0) \right] \\ &\quad + \epsilon^3 \sum_{i=1}^d \left[\frac{1}{2} \alpha''_{x_i}(0) + \frac{\beta}{6} \alpha'''_{x_i}(\xi_{x_i}) \right] + \epsilon^4 \sum_{i=1}^d \frac{1}{6} \alpha'''_{x_i}(\xi_{x_i}). \end{aligned}$$

By (35) and using the i.i.d. nature of the x_i and Markov's inequality, for all $\delta > 0$,

$$\begin{aligned} \delta \mathbb{P} \left(\left| \epsilon^3 \sum_{i=1}^d \left[\frac{1}{2} \alpha''_{x_i}(0) + \frac{\beta}{6} \alpha'''_{x_i}(\xi_{x_i}) \right] \right| > \delta \right) &< \mathbb{E} \left(\left| \frac{\ell^3}{d^{3/2}} \sum_{i=1}^d \left[\frac{1}{2} \alpha''_{x_i}(0) + \frac{\beta}{6} \alpha'''_{x_i}(\xi_{x_i}) \right] \right| \right) \\ &\leq \frac{\ell^3}{d^{1/2}} \left[\frac{1}{2} \mathbb{E}(|\alpha''_{x_i}(0)|) + \frac{\beta}{6} C \right] \\ &\rightarrow 0 \quad \text{as } d \rightarrow \infty. \end{aligned}$$

Thus,

$$\epsilon^3 \sum_{i=1}^d \left[\frac{1}{2} \alpha''_{x_i}(0) + \frac{\beta}{6} \alpha'''_{x_i}(\xi_{x_i}) \right] \rightarrow 0 \quad \text{in probability as } d \rightarrow \infty.$$

By identical methodology, as $d \rightarrow \infty$,

$$\epsilon^4 \sum_{i=1}^d \frac{1}{6} \alpha'''_{x_i}(\xi_{x_i}) \rightarrow 0 \quad \text{in probability.}$$

Consequently,

$$\begin{aligned} H_{\beta}^{\beta'}(\mathbf{x}) &= \epsilon \left[\sum_{i=1}^d h(x_i) - \frac{1}{2}(x_i - \mu)h'(x_i) \right] \\ &+ \frac{\epsilon^2}{8\beta} \left[\sum_{i=1}^d (x_i - \mu)^2 h''(x_i) - (x_i - \mu)h'(x_i) \right] + T_x, \end{aligned}$$

where

$$T_x = \epsilon^3 \sum_{i=1}^d \left[\frac{1}{2} \alpha''_{x_i}(0) + \frac{\beta}{6} \alpha'''_{x_i}(\xi_{x_i}) \right] + \epsilon^4 \sum_{i=1}^d \frac{1}{6} \alpha'''_{x_i}(\xi_{x_i}),$$

with $T_x \rightarrow 0$ in probability as $d \rightarrow \infty$.

Now denoting $h(t(y, \beta', \beta))$ as

$$\alpha_y(\epsilon) := h(t(y, \beta', \beta)) = \log \left[f \left(\left(\frac{\beta + \epsilon}{\beta} \right)^{1/2} (y - \mu) + \mu \right) \right],$$

the Taylor series expansion in ϵ , for a fixed y , with Taylor truncation term denoted by ξ_y such that $0 < \xi_y < \epsilon$ is given by

$$h(t(y, \beta', \beta)) = \alpha_y(\epsilon) = \alpha_y(0) + \epsilon \alpha'_y(0) + \frac{\epsilon^2}{2} \alpha''_y(0) + \frac{\epsilon^3}{6} \alpha_y(\xi_y).$$

By identical methodology to the above calculation in (35) for $\alpha_x(\cdot)$, it can be shown that there exists $C_y \in \mathbb{R}_+$ such that

$$|\mathbb{E}_{\beta}[\alpha'''_y(\xi_y)]| \leq C_y.$$

Hence, using exactly the same methodology as for the x_i above,

$$\begin{aligned} H_{\beta'}^{\beta}(y) &= -\epsilon \left[\sum_{i=1}^d h(y_i) - \frac{1}{2}(y_i - \mu)h'(y_i) \right] \\ &+ \frac{\epsilon^2}{8\beta} \left[\sum_{i=1}^d (y_i - \mu)^2 h''(y_i) - (y_i - \mu)h'(y_i) \right] + T_y, \end{aligned}$$

where $T_y \rightarrow 0$ in probability as $d \rightarrow \infty$.

Using the notation from Definition 1, the desired form of B in Lemma 1 is reached. □

Lemma 2. (Asymptotic Gaussianity of the log-acceptance ratio for QuanTA.) *Under the notation and assumptions of Theorem 1 and Definition 1, B is asymptotically Gaussian of the form $B \rightsquigarrow N(-\sigma^2, 2, \sigma^2)$, where*

$$\sigma^2 = 2\ell^2 \left[\frac{1}{2}V(\beta) - I(\beta) + \frac{1}{4\beta}R(\beta) \right].$$

Proof. Recall the form of B from Lemma 1. Then, making the dimensionality dependence explicit, write $B = W(d) + (T_x + T_y)$, where

$$\begin{aligned} W(d) := & \epsilon \left[\sum_{i=1}^d h(x_i) - h(y_i) + \frac{1}{2}(k(y_j) - k(x_j)) \right] \\ & + \frac{\epsilon^2}{8\beta} \left[\sum_{i=1}^d r(x_i) - k(x_i) + r(y_i) - k(y_i) \right] \end{aligned}$$

and $(T_x + T_y) \rightarrow 0$ in probability as $d \rightarrow \infty$. Hence, if it can be shown that $W(d)$ converges in distribution to a Gaussian of the form $N(-c, 2c)$ then by Slutsky’s theorem we can conclude that B converges in distribution to the same Gaussian as the W .

To this end, the asymptotic Gaussianity of $W(d)$ is established. First note that, due to the i.i.d. nature of the x_i and y_i respectively, by the standard central limit theorem (see, e.g. [6]) for a sum of i.i.d. variables, asymptotic Gaussianity is immediate, where

$$W(d) \Rightarrow N(\mu_W, \sigma_W^2) \quad \text{as } d \rightarrow \infty$$

and

$$\mu_W = \lim_{d \rightarrow \infty} \mathbb{E}[W(d)], \quad \sigma_W^2 = \lim_{d \rightarrow \infty} \text{var}[W(d)].$$

To this end, the terms $\mathbb{E}[W(d)]$ and $\text{var}[W(d)]$ are computed. We obtain

$$\begin{aligned} \mathbb{E}[W(d)] := & \epsilon \left[\sum_{i=1}^d M(\beta) - M(\beta + \epsilon) - \frac{1}{2}(S(\beta) - S(\beta + \epsilon)) \right] \\ & + \frac{\epsilon^2}{8\beta} \left[\sum_{i=1}^d R(\beta) + R(\beta + \epsilon) \right] \\ = & \epsilon \left[\sum_{i=1}^d -\epsilon M'(\beta) + \frac{\epsilon}{2} S'(\beta) \right] + \frac{\epsilon^2}{8\beta} \left[\sum_{i=1}^d 2R(\beta) \right] + \mathcal{O}(d^{-1/2}) \\ \rightarrow & \ell^2 \left[\frac{1}{2}V(\beta) - I(\beta) + \frac{1}{4\beta}R(\beta) \right] \quad \text{as } d \rightarrow \infty, \end{aligned}$$

where the final line utilises the results in Proposition 2. Similarly,

$$\text{var}(W(d)) \rightarrow 2\ell^2 \text{var}_\beta \left(h(x) - \frac{1}{2}k(x) \right) \quad \text{as } d \rightarrow \infty.$$

Hence, by Slutsky’s theorem, B is asymptotically Gaussian such that

$$B \rightsquigarrow N \left(\ell^2 \left[\frac{1}{2}V(\beta) - I(\beta) + \frac{1}{4\beta}R(\beta) \right], 2\ell^2 \text{var}_\beta \left(h(x) - \frac{1}{2}k(x) \right) \right). \tag{36}$$

However, this does not obviously have the form required with $B \sim N(-\sigma^2/2, \sigma^2)$ for some σ^2 . This form is verified with the following proposition, which completes the proof of Lemma 2. □

Proposition 3. *Under the notation and assumptions of Theorem 1 and Definition 1,*

$$\ell^2 \left[\frac{1}{2}V(\beta) - I(\beta) + \frac{1}{4\beta}R(\beta) \right] = -\ell^2 \text{var}_\beta \left(h(x) - \frac{1}{2}k(x) \right). \tag{37}$$

Proof. From (36), let

$$\mu = \ell^2 \left[\frac{1}{2}V(\beta) - I(\beta) + \frac{1}{4\beta}R(\beta) \right] \tag{38}$$

and

$$\sigma^2 = 2\ell^2 \text{var}_\beta \left(h(x) - \frac{1}{2}k(x) \right).$$

Then, by using the standard properties of variance, it is routine to show that

$$-\frac{\sigma^2}{2} = \ell^2 \left[-I(\beta) - \frac{1}{4} \text{var}_\beta(k(x)) + V(\beta) \right]. \tag{39}$$

Consequently, equating the terms on the right-hand side of (38) and (39) shows that if the following can be shown to hold then the required identity in (37) is validated:

$$\frac{1}{4\beta}R(\beta) = -\frac{1}{4} \text{var}_\beta(k(x)) + \frac{1}{2}V(\beta). \tag{40}$$

The left-hand side (LHS) and right-hand side (RHS) of (40) will be considered separately. The following integration by parts is well defined due to the assumption that $-\log(f(\cdot))$ has regularly varying tails. Starting with the RHS and recalling that, from (27), $\mathbb{E}_\beta(k(x)) = -1/\beta$,

$$\begin{aligned} -\frac{1}{4} \text{var}_\beta(k(x)) + \frac{1}{2}V(\beta) &= -\frac{1}{4} [\mathbb{E}_\beta(k(x)^2) - \mathbb{E}_\beta(k(x))^2] + \frac{1}{2\beta^2} \\ &= -\frac{1}{4} \mathbb{E}_\beta(k(x)^2) + \frac{3}{4\beta^2}. \end{aligned}$$

Then, noting that $(\log f)'(x) f^\beta(x) = f'(x) f^{\beta-1}(x)$, and using integration by parts (by first integrating $f'(x) f^{\beta-1}(x)$),

$$\begin{aligned} \mathbb{E}_\beta(k(x)^2) &= \int (x - \mu)^2 [(\log f)'(x)]^2 \frac{f^\beta(x)}{Z_\beta} dx \\ &= \left[\frac{(x - \mu)^2}{\beta} (\log f)'(x) \frac{f^\beta(x)}{Z_\beta} \right]_{-\infty}^{\infty} \\ &\quad - \frac{1}{\beta} \int [(x - \mu)^2 (\log f)''(x) + 2(x - \mu)(\log f)'(x)] \frac{f^\beta(x)}{Z_\beta} dx \\ &= 0 - \frac{1}{\beta} \mathbb{E}_\beta(r(x)) - \frac{2}{\beta} \mathbb{E}_\beta(k(x)) \\ &= -\frac{1}{\beta} \mathbb{E}_\beta(r(x)) + \frac{2}{\beta^2}. \end{aligned} \tag{41}$$

Collating the above in (40) and (41), we obtain

$$-\frac{1}{4}\text{var}_\beta(k(x)) + \frac{1}{2}V(\beta) = \frac{1}{4\beta}\mathbb{E}_\beta(r(x)) + \frac{1}{4\beta^2} = \frac{1}{4\beta}R(\beta),$$

where the final equality simply comes from the definition of $R(\beta)$ from (25). □

Lemma 3. (Optimisation of the ESJD $_{\beta,\cdot}$) *Under the notation and assumptions of Theorem 1 and Definition 1, the ESJD $_{\beta}$ is maximised when ℓ is chosen to maximise*

$$2\ell^2\Phi\left(-\frac{\ell[V(\beta)/2 - I(\beta) + R(\beta)/4\beta]^{1/2}}{\sqrt{2}}\right).$$

Furthermore, for the optimal ℓ , the corresponding swap move acceptance rate induced between two consecutive temperatures is given by 0.234 (to three significant figures).

Proof. Let $\phi_{(m,\sigma^2)}$ denote the density function of a Gaussian with mean m and variance σ^2 , and suppose that $G \sim N(-\sigma^2/2, \sigma^2)$. Then a routine calculation (which can be found in, e.g. [35]) shows that

$$\mathbb{E}(1 \wedge e^G) = 2\Phi\left(-\frac{\sigma}{2}\right). \tag{42}$$

Using the result in (42) and Lemma 2, in the limit as $d \rightarrow \infty$,

$$\lim_{d \rightarrow \infty} (d\text{ESJD}_\beta) = 2\ell^2\Phi\left(-\frac{\ell[V(\beta)/2 - I(\beta) + R(\beta)/4\beta]^{1/2}}{\sqrt{2}}\right).$$

Substituting

$$u = \ell \left[\frac{1}{2}V(\beta) - I(\beta) + \frac{1}{4\beta}R(\beta) \right]^{1/2},$$

and then maximising with respect to u attains an optimising value u^* that does not depend on

$$\left[\frac{1}{2}V(\beta) - I(\beta) + \frac{1}{4\beta}R(\beta) \right].$$

Recalling the form of the ESJD $_{\beta}$ from (17), it is clear that the associated acceptance rate, denoted (ACC $_{\beta}$), induced by choosing any value of u is

$$\text{ACC}_\beta = \mathbb{E}_{\pi_n}[\alpha_\beta(x, y)],$$

which, as established above, in the limit as $d \rightarrow \infty$ is asymptotically given by

$$\text{ACC}_\beta = 2\Phi\left(-\frac{u}{\sqrt{2}}\right).$$

Now it can be shown numerically that the optimising value u^* induces

$$\text{ACC}_\beta = 0.234 \quad (\text{three significant figures}). \tag{□}$$

A.2. Proof of Theorem 2

Note that in Theorem 2 the conditions on $f(\cdot)$ are inherited from the conditions on $f(\cdot)$ from Theorem 1. This includes the bounded fourth derivatives of $\log(f)$ and the existence of eighth moments, i.e. $\mathbb{E}_\beta[X^8]$, which is due to the assumption of regularly varying tails. These will be assumed for the following lemmata.

Lemma 4. *Under the notation and assumptions of Theorems 1 and 2 and Definition 1,*

$$\frac{1}{2}V(\beta) - I(\beta) = \mathcal{O}(\beta^{-k}),$$

where, in general, $k = \min\{2 + \gamma, 5/2\}$, but if $h'''(0) = 0$ then $k = \min\{2 + \gamma, 3\}$.

Proof. It has already been established that $V(\beta) = 1/\beta^2$ for all distributions. Also, for a Gaussian density $f(\cdot)$, $I(\beta) = 1/(2\beta^2)$. Since $g_\beta(\cdot)$ approaches the density of a standard Gaussian, $\phi(\cdot)$, as $\beta \rightarrow \infty$, then we expect that $I(\beta)$ would approach $1/(2\beta^2)$ too. Hence, a rigorous analysis of this convergence needs to be established. Note that

$$\begin{aligned} I(\beta) &= \text{var}_\beta[h(X)] \\ &= \int (h(x) - \mathbb{E}_\beta[h(X)])^2 \frac{f^\beta(x)}{Z(\beta)} dx \\ &= \int \left(h\left(\frac{y}{\sqrt{\beta(-h''(0))}}\right) - \mathbb{E}_{g_\beta}\left[h\left(\frac{y}{\sqrt{\beta(-h''(0))}}\right)\right] \right)^2 g_\beta(y) dy, \end{aligned} \tag{43}$$

using the change of variable, $X = Y/\sqrt{\beta(-h''(0))}$. By Taylor expansion of h about the mode point 0 up to fourth order,

$$h\left(\frac{y}{\sqrt{\beta(-h''(0))}}\right) = h(0) - \frac{y^2}{2\beta} + \frac{y^3 h'''(0)}{6(\beta(-h''(0)))^{3/2}} + \frac{y^4 h''''(\xi_1(y))}{24(\beta(-h''(0)))^2}, \tag{44}$$

where $\xi_1(\cdot)$ is the truncation term for the Taylor expansion such that $0 < |\xi_1(y)| < |y/\sqrt{\beta(-h''(0))}|$ for all y . Using the Taylor expansion form of h and the assumption of bounded fourth derivatives,

$$\begin{aligned} &\left| \mathbb{E}_{g_\beta} \left[h\left(\frac{Y}{\sqrt{\beta(-h''(0))}}\right) - h(0) + \frac{Y^2}{2\beta} - \frac{Y^3 h'''(0)}{6(\beta(-h''(0)))^{3/2}} \right] \right| \\ &\leq \mathbb{E}_{g_\beta} \left[\left| \frac{Y^4 h''''(\xi_1(Y))}{24(\beta(-h''(0)))^2} \right| \right] \\ &\leq \frac{M}{24(\beta(-h''(0)))^2} \mathbb{E}_{g_\beta}[Y^4] \\ &= \mathcal{O}\left(\frac{1}{\beta^2}\right), \end{aligned}$$

where $\mathbb{E}_{g_\beta}[Y^4] < \infty$ due to the assumption on the existence of moments up to the eighth moment. Thus,

$$\begin{aligned} \mathbb{E}_{g_\beta} \left[h\left(\frac{Y}{\sqrt{\beta(-h''(0))}}\right) \right] &= h(0) - \frac{\mathbb{E}_{g_\beta}[Y^2]}{2\beta} + \frac{\mathbb{E}_{g_\beta}[Y^3]h'''(0)}{6(\beta(-h''(0)))^{3/2}} \\ &\quad + \frac{\mathbb{E}_{g_\beta}[Y^4 h''''(\xi_1(Y))]}{24(\beta(-h''(0)))^2}, \end{aligned}$$

and substituting this into (43), along with the Taylor expansion of h to the fourth order given in (44), gives

$$\begin{aligned}
 I(\beta) &= \int \left(h(0) - \frac{y^2}{2\beta} + \frac{y^3 h'''(0)}{6(\beta(-h''(0)))^{3/2}} + \frac{y^4 h''''(\xi_1(y))}{24(\beta(-h''(0)))^2} \right. \\
 &\quad \left. - \left[h(0) + \frac{\mathbb{E}_{g_\beta}[Y^2] h''(0)}{2\beta(-h''(0))} + \frac{\mathbb{E}_{g_\beta}[Y^3] h'''(0)}{6(\beta(-h''(0)))^{3/2}} \right. \right. \\
 &\quad \left. \left. + \frac{\mathbb{E}_{g_\beta}[Y^4 h''''(\xi_1(Y))]}{24(\beta(-h''(0)))^2} \right] \right)^2 g_\beta(y) dy \\
 &= \frac{1}{4\beta^2} \int (y^2 - \mathbb{E}_{g_\beta}[Y^2])^2 g_\beta(y) dy \\
 &\quad + \frac{2h'''(0)}{24\beta^{5/2}(-h''(0))^{3/2}} \int (y^2 - \mathbb{E}_{g_\beta}[Y^2])(y^3 - \mathbb{E}_{g_\beta}[Y^3]) g_\beta(y) dy + \mathcal{O}\left(\frac{1}{\beta^3}\right),
 \end{aligned}$$

which is finite and well defined due to (14) and (15). Consequently, in general,

$$I(\beta) = \frac{1}{4\beta^2} \text{var}_{g_\beta}(Y^2) + \mathcal{O}\left(\frac{1}{\beta^{5/2}}\right),$$

but in the case that $h'''(0) = 0$, which indeed holds in the case that f is symmetric about the mode point, then

$$I(\beta) = \frac{1}{2\beta^2} \text{var}_{g_\beta}(Y^2) + \mathcal{O}\left(\frac{1}{\beta^3}\right),$$

and so, under the key assumption given in (19),

$$I(\beta) = \frac{1}{2\beta^2} + \mathcal{O}\left(\frac{1}{\beta^k}\right),$$

where, in general, $k = \min\{2 + \gamma, 5/2\}$, but if $h'''(0) = 0$ then $k = \min\{2 + \gamma, 3\}$, and so $\frac{1}{2}V(\beta) - I(\beta) = \mathcal{O}(1/\beta^k)$. □

Lemma 5. *Under the notation and assumptions of Theorems 1 and 2 and Definition 1,*

$$\frac{1}{4\beta} R(\beta) = \mathcal{O}(\beta^{-k}),$$

where, in general, $k = 5/2$, but if $h'''(0) = 0$ then $k = 3$.

Proof. Recall that

$$\begin{aligned}
 \frac{1}{4\beta} R(\beta) &= \frac{1}{4\beta} \mathbb{E}_\beta[X^2 h''(X) - X h'(X)] \\
 &= \frac{1}{4\beta} \mathbb{E}_{g_\beta} \left[\left(\frac{Y}{\sqrt{\beta(-h''(0))}} \right)^2 h''\left(\frac{Y}{\sqrt{\beta(-h''(0))}} \right) \right. \\
 &\quad \left. - \frac{Y}{\sqrt{\beta(-h''(0))}} h'\left(\frac{Y}{\sqrt{\beta(-h''(0))}} \right) \right]. \tag{45}
 \end{aligned}$$

Using Taylor’s expansion about the mode at 0,

$$h' \left(\frac{y}{\sqrt{\beta(-h''(0))}} \right) = h'(0) + \frac{y}{\sqrt{\beta(-h''(0))}} h''(0) + \frac{y^2}{2\beta(-h''(0))} h'''(0) + \frac{y^3}{6\beta^{3/2}(-h''(0))^{3/2}} h''''(\xi_2(y)),$$

where $\xi_2(\cdot)$ is the truncation term for the Taylor expansion such that $0 < |\xi_2(y)| < |y/\sqrt{\beta(-h''(0))}|$ for all y . Also,

$$h'' \left(\frac{y}{\sqrt{\beta(-h''(0))}} \right) = h''(0) + \frac{y}{\sqrt{\beta(-h''(0))}} h'''(0) + \frac{y^2}{2\beta^{3/2}(-h''(0))^{3/2}} h''''(\xi_3(y)),$$

where $\xi_3(\cdot)$ is the truncation term for the Taylor expansion such that $0 < |\xi_3(y)| < |y/\sqrt{\beta(-h''(0))}|$ for all y . Hence,

$$\begin{aligned} & \frac{y^2}{2\beta(-h''(0))} h'' \left(\frac{y}{\sqrt{\beta(-h''(0))}} \right) - \frac{y}{\sqrt{\beta(-h''(0))}} h' \left(\frac{y}{\sqrt{\beta(-h''(0))}} \right) \\ &= \frac{y^3}{2(\beta(-h''(0)))^{3/2}} h'''(0) + \frac{y^4}{(\beta(-h''(0)))^2} \left[\frac{1}{2} h''''(\xi_3(y)) - \frac{1}{6} h''''(\xi_2(y)) \right]. \end{aligned}$$

Substituting this into the $R(\beta)/4\beta$ term in (45),

$$\begin{aligned} \frac{1}{4\beta} R(\beta) &= \frac{1}{4\beta} \mathbb{E}_{g\beta} \left[\frac{Y^3}{2(\beta(-h''(0)))^{3/2}} h'''(0) \right. \\ &\quad \left. + \frac{Y^4}{(\beta(-h''(0)))^2} \left[\frac{1}{2} h''''(\xi_3(Y)) - \frac{1}{6} h''''(\xi_2(Y)) \right] \right] \\ &= \frac{h'''(0)}{8\beta^{5/2}(-h''(0))^{3/2}} \mathbb{E}_{g\beta} [Y^3] \\ &\quad + \frac{1}{4\beta^3(-h''(0))^2} \mathbb{E}_{g\beta} \left[Y^4 \left[\frac{1}{2} h''''(\xi_3(Y)) - \frac{1}{6} h''''(\xi_2(Y)) \right] \right], \end{aligned}$$

where

$$\mathbb{E}_{g\beta} \left[Y^4 \left[\frac{1}{2} h''''(\xi_3(Y)) - \frac{1}{6} h''''(\xi_2(Y)) \right] \right] < \infty$$

due to the assumptions of boundedness of the fourth derivatives of $\log f(X)$ and the existence of moments. Hence, in general,

$$\frac{1}{4\beta} R(\beta) = \mathcal{O} \left(\frac{1}{\beta^{5/2}} \right),$$

but in the case that $h'''(\cdot) = 0$, which is the case when $f(\cdot)$ is symmetric about the mode point 0,

$$\frac{1}{4\beta} R(\beta) = \mathcal{O} \left(\frac{1}{\beta^3} \right).$$

Consequently,

$$\frac{1}{4\beta} R(\beta) = \mathcal{O} \left(\frac{1}{\beta^k} \right)$$

where, in general, $k = 5/2$, but in the case that $h'''(0) = 0$ then $k = 3$. □

Acknowledgements

Both authors would like to acknowledge the funding from the EPSRC grants EP/L505110/1 and EP/KO14463/1 that made this work possible. We would also like to thank both Professors Jonathan Tawn (Lancaster) and Jeff Rosenthal (Toronto) for the helpful suggestions and support during the work.

References

- [1] ATCHADÉ, Y. F., ROBERTS, G. O. AND ROSENTHAL, J. S. (2011). Towards optimal scaling of Metropolis-coupled Markov chain Monte Carlo. *Statist. Comput.* **21**, 555–568.
- [2] BARNDORFF-NIELSEN, O. E. AND COX, D. R. (1989). *Asymptotic Techniques for use in Statistics*. Chapman & Hall, London.
- [3] BINGHAM, N. H., GOLDIE, C. M. AND TEUGELS, J. L. (1989). *Regular Variation* (Encyclopedia Math. **27**). Cambridge University Press.
- [4] BROOKS, S. P., GIUDICI, P. AND ROBERTS, G. O. (2003). Efficient construction of reversible jump Markov chain Monte Carlo proposal distributions. *J. R. Statist. Soc. B (Statist. Methodology)* **65**, 3–55.
- [5] CARTER, J. N. AND WHITE, D. A. (2013). History matching on the imperial college fault model using parallel tempering. *Comput. Geosci.* **17**, 43–65.
- [6] DURRETT, R. (2010). *Probability: Theory and Examples*. Cambridge University Press.
- [7] GE, R., LEE, H. AND RISTESKI, A. (2018). Beyond Log-concavity: provable guarantees for sampling multimodal distributions using simulated tempering Langevin Monte Carlo. In *32nd Internat. Conf. on Neural Information Processing Systems*, pp. 7858–7867.
- [8] GEYER, C. J. (1991). Markov chain Monte Carlo maximum likelihood. In *Computing Science and Statistics, Proceedings of the 23rd Symposium on the Interface*, pp. 156–163.
- [9] GEYER, C. J. AND THOMPSON, E. A. (1995). Annealing Markov chain Monte Carlo with applications to ancestral inference. *J. Amer. Statist. Assoc.* **90**, 909–920.
- [10] GREEN, P. J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* **82**, 711–732.
- [11] HARTIGAN, J. A. AND WONG, M. A. (1979). Algorithm AS 136: A K-means clustering algorithm. *J. R. Statist. Soc. C (Appl. Statist.)* **28**, 100–108.
- [12] HASTIE, D. (2005). Towards automatic reversible jump Markov chain Monte Carlo. Doctoral Thesis, University of Bristol.
- [13] HASTIE, T., TIBSHIRANI, R. AND FRIEDMAN, J. (2001). *The Elements of Statistical Learning*. Springer, New York.
- [14] KERDPRASOP, K., KERDPRASOP, N. AND SATTAYATHAM, P. (2005). Weighted K-means for density-biased clustering. In *Data Warehousing and Knowledge Discovery*, Springer, Berlin, pp. 488–497.
- [15] KIM, S., TADESSE, M. G. AND VANNUCCI, M. (2006). Variable selection in clustering via Dirichlet process mixture models. *Biometrika* **93**, 877–893.
- [16] KOU, S. C., ZHOU, Q. AND WONG, W. H. (2006). Equi-energy sampler with applications in statistical inference and statistical mechanics. *Ann. Statist.* **34**, 1581–1652.
- [17] MANGOUBI, O., PILLAI, N. S. AND SMITH, A. (2018). Does Hamiltonian Monte Carlo mix faster than a random walk on multimodal densities? Preprint. Available at [arXiv:1808.03230v2](https://arxiv.org/abs/1808.03230v2).
- [18] MARINARI, E. AND PARISI, G. (1992). Simulated tempering: a new Monte Carlo scheme. *Europhys. Lett.* **19**, 451–458.
- [19] MIASOJEDOW, B., MOULINES, E. AND VIHOLA, M. (2013). An adaptive parallel tempering algorithm. *J. Comput. Graph. Statist.* **22**, 649–664.
- [20] MOHAMED, L. *et al.* (2012). Population MCMC methods for history matching and uncertainty quantification. *Comput. Geosci.* **16**, 423–436.
- [21] NEAL, R. M. (1996). Sampling from multimodal distributions using tempered transitions. *Statist. Comput.* **6**, 353–366.
- [22] NEAL, R. M. (2000). Markov chain sampling methods for Dirichlet process mixture models. *J. Comput. Graph. Statist.* **9**, 249–265.
- [23] NEMETH, C., LINDSTEN, F., FILIPPONE, M. AND HENSMAN, J. (2019). Pseudo-extended Markov chain Monte Carlo. Preprint. Available at [arXiv:1708.05239v2](https://arxiv.org/abs/1708.05239v2).
- [24] OLVER, F. W. J. (1968). Error bounds for the Laplace approximation for definite integrals. *J. Approximation Theory* **1**, 293–313.
- [25] PAPASPILIOPOULOS, O. AND ROBERTS, G. O. (2003). Non-centered parameterisations for hierarchical models and data augmentation. In *Bayesian Statistics 7*, Oxford University Press, New York, pp. 307–326.

- [26] PAPANASTASIIOPOULOS, O., ROBERTS, G. O. AND SKÖLD, M. (2007). A general framework for the parametrization of hierarchical models. *Statist. Sci.* **22**, 59–73.
- [27] PARDO, C. E. AND DEL CAMPO, P. C. (2007). Combinación de métodos factoriales y de análisis de conglomerados en R: el paquete FactoClass. *Revista Colombiana de Estadística* **30**, 231–245.
- [28] PAULIN, D., JASRA, A., THIERY, A. (2019). Error bounds for sequential Monte Carlo samplers for multimodal distributions. *Bernoulli* **25**, 310–340.
- [29] RAYKOV, Y. P., BOUKOUVALAS, A., LITTLE, M. A. (2016). Simple approximate MAP inference for Dirichlet processes mixtures. *Electron. J. Statist.* **10**, 3548–3578.
- [30] ROBBINS, H. AND MONRO, S. (1951). A stochastic approximation method. *Ann. Math. Statist.* **22**, 400–407.
- [31] ROBERTS, G. O. AND ROSENTHAL, J. S. (2001). Optimal scaling for various Metropolis-Hastings algorithms. *Statist. Sci.* **16**, 351–367.
- [32] ROBERTS, G. O. AND ROSENTHAL, J. S. (2007). Coupling and ergodicity of adaptive Markov chain Monte Carlo algorithms. *J. Appl. Prob.* **44**, 458–475.
- [33] ROBERTS, G. O. AND ROSENTHAL, J. S. (2009). Examples of adaptive MCMC. *J. Comput. Graph. Statist.* **18**, 349–367.
- [34] ROBERTS, G. O. AND ROSENTHAL, J. S. (2014). Minimising MCMC variance via diffusion limits, with an application to simulated tempering. *Ann. Appl. Prob.* **24**, 131–149.
- [35] ROBERTS, G. O., GELMAN, A. AND GILKS, W. R. (1997). Weak convergence and optimal scaling of random walk Metropolis algorithms. *Ann. Appl. Prob.* **7**, 110–120.
- [36] RUE, H., MARTINO, S. AND CHOPIN, N. (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *J. R. Statist. Soc. B (Statist. Methodology)* **71**, 319–392.
- [37] SCHWEIZER, N. (2012). Non-asymptotic error bounds for sequential MCMC methods. Doctoral Thesis, Universitäts- und Landesbibliothek Bonn.
- [38] SHERLOCK, C. (2006). Methodology for inference on the Markov modulated Poisson process and theory for optimal scaling of the random walk metropolis. Doctoral Thesis, Lancaster University.
- [39] TAWN, N. (2017). Towards optimality of the parallel tempering algorithm. Doctoral Thesis, University of Warwick.
- [40] TAWN, N. G. AND ROBERTS, G. O. (2018). Optimal temperature spacing for regionally weight-preserving tempering. Preprint. Available at [arXiv:1810.05845v1](https://arxiv.org/abs/1810.05845v1).
- [41] TAWN, N. G., ROBERTS, G. O. AND ROSENTHAL, J. S. (2019). Weight preserving simulated tempering. *Statist. Comput.*, 15pp.
- [42] TIERNEY, L. (1994). Markov chains for exploring posterior distributions. *Ann. Statist.* **22**, 1701–1762.
- [43] TIERNEY, L. AND KADANE, J. B. (1986). Accurate approximations for posterior moments and marginal densities. *J. Amer. Statist. Assoc.* **81**, 82–86.
- [44] TSENG, G. C. (2007). Penalized and weighted K-means for clustering with scattered objects and prior information in high-throughput biological data. *Bioinformatics* **23**, 2247–2255.
- [45] VANDERWERKEN, D. N. AND SCHMIDLER, S. C. (2013). Parallel Markov chain Monte Carlo. Preprint. Available at [arXiv:1312.7479v1](https://arxiv.org/abs/1312.7479v1).
- [46] WANG, J.-S. AND SWENDSEN, R. H. (1990). Cluster Monte Carlo algorithms. *Phys. A* **167**, 565–579.
- [47] WOODARD, D. B., SCHMIDLER, S. C. AND HUBER, M. (2009). Conditions for rapid mixing of parallel and simulated tempering on multimodal distributions. *Ann. Appl. Prob.* **19**, 617–640.
- [48] WOODARD, D. B., SCHMIDLER, S. C. AND HUBER, M. (2009). Sufficient conditions for torpid mixing of parallel and simulated tempering. *Electron. J. Prob.* **14**, 780–804.
- [49] XIE, Y., ZHOU, J. AND JIANG, S. (2010). Parallel tempering Monte Carlo simulations of lysozyme orientation on charged surfaces. *J. Chem. Phys.* **132**, 065101.