



Monotonic Alpha-divergence Minimisation for Variational Inference

Kamélia Daudel, Randal Douc, François Roueff

► To cite this version:

Kamélia Daudel, Randal Douc, François Roueff. Monotonic Alpha-divergence Minimisation for Variational Inference. Journal of Machine Learning Research, 2023, 24 (62), pp.1-76. hal-03164338v2

HAL Id: hal-03164338

<https://telecom-paris.hal.science/hal-03164338v2>

Submitted on 27 Apr 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Monotonic Alpha-divergence Minimisation for Variational Inference

Kamélia Daudel^{1,2}

KAMELIA.DAUDEL@GMAIL.COM

1: LTCI, Télécom Paris

Institut Polytechnique de Paris, France

2: Department of Statistics, University of Oxford

Oxford OX1 3TG, United Kingdom

Randal Douc

RANDAL.DOUC@TELECOM-SUDPARIS.EU

SAMOVAR, Télécom SudParis

Institut Polytechnique de Paris, France

François Roueff

FRANCOIS.ROUEFF@TELECOM-PARIS.FR

LTCI, Télécom Paris

Institut Polytechnique de Paris, France

Editor: Shakir Mohamed

Abstract

In this paper, we introduce a novel family of iterative algorithms which carry out α -divergence minimisation in a Variational Inference context. They do so by ensuring a systematic decrease at each step in the α -divergence between the variational and the posterior distributions. In its most general form, the variational distribution is a mixture model and our framework allows us to simultaneously optimise the weights and components parameters of this mixture model. Our approach permits us to build on various methods previously proposed for α -divergence minimisation such as Gradient or Power Descent schemes and we also shed a new light on an integrated Expectation Maximization algorithm. Lastly, we provide empirical evidence that our methodology yields improved results on several multimodal target distributions and on a real data example.

Keywords: Variational Inference, Kullback-Leibler, Alpha-Divergence, Mixture Models, Bayesian Inference

1. Introduction

Bayesian inference tasks often induce intractable and hard-to-compute posterior densities which need to be approximated. Among the class of approximating methods, Variational inference methods (for example Variational Bayes Jordan et al., 1999; James Beal, 2003) have attracted a lot of attention as they have empirically been shown to be widely applicable to many high-dimensional machine-learning problems (Hoffman et al., 2013; Ranganath et al., 2014; Kingma and Welling, 2014).

These optimisation-based methods introduce a simpler variational family \mathcal{Q} and find the best approximation to the unknown posterior density among this family in terms of a

certain measure of dissimilarity, the most common choice of measure of dissimilarity being the exclusive Kullback-Leibler divergence (Blei et al., 2017; Zhang et al., 2019).

However, the exclusive Kullback-Leibler divergence is known to have some drawbacks. Indeed, its *zero-forcing* behavior is responsible for returning variational approximations with light tails that severely underestimate the posterior covariance and that are unable to capture multimodality within the posterior density (Minka, 2005; Jerfel et al., 2021; Prangle, 2021). This is especially inconvenient when the variational family \mathcal{Q} does not exactly match the posterior density (Yao et al., 2018; Campbell and Li, 2019) and even more so if one wants to appeal to Importance Sampling methods to approximate integrals of interest in a Bayesian Inference setting (Jerfel et al., 2021; Prangle, 2021).

To avoid this hurdle, advances in Variational Inference turned to more general families of divergences such as the α -divergence (Zhu and Rohwer, 1995b,a) and Rényi’s α -divergence (Rényi, 1961; van Erven and Harremoës, 2014). These families of divergences both recover the exclusive Kullback-Leibler when $\alpha \rightarrow 1$ and thanks to the hyperparameter α , they provide a more flexible framework that can bypass the difficulties associated to the exclusive Kullback-Leibler divergence by choosing $\alpha < 1$ (Minka, 2005). For this reason, they have notably been used in Minka (2004); Hernandez-Lobato et al. (2016); Li and Turner (2016); Dieng et al. (2017); Wang et al. (2018); Daudel et al. (2021); Daudel and Douc (2021).

In the spirit of Variational Inference methods based on the α -divergence, we propose in this paper to build a framework for α -divergence minimisation in a Variational Inference context. The particularity of our work is that our algorithms will ensure a monotonic decrease at each step in the α -divergence between the variational and the posterior distributions. In addition, our work will apply to variational families as wide as the class of mixture models. The paper is then organised as follows:

- In Section 2, we introduce some notation and state the optimisation problem we aim at solving in terms of the posterior density, the variational density $q \in \mathcal{Q}$ and the α -divergence.
- In Section 3, we consider the typical Variational Inference case where q belongs to a parametric family. In this particular case, we state in Theorem 1 conditions which ensure a systematic decrease in the α -divergence at each step for all $\alpha \in [0, 1)$. We then show in Corollary 1 that these conditions are satisfied for a well-chosen iterative scheme and we call the resulting approach the *maximisation approach*. This approach is particularly convenient, a fact that we illustrate over Example 1 and Lemma 1. Furthermore, we derive in Corollary 2 additional iterative schemes satisfying the conditions of Theorem 1 under the name *gradient-based approach*, which we then use to underline the links between our approach and Gradient Descent schemes for α -divergence and Rényi’s α -divergence minimisation.
- In Section 4, we extend the results from Section 3 to the more general case of mixture models. We derive in Theorems 2 and 3 conditions to simultaneously optimise the weights and the components parameters of a given mixture model, all the while maintaining the systematic decrease in the α -divergence initially enjoyed in Section 3. These conditions are then met in Corollary 4 and 5, which respectively generalise the maximisation approach of Corollary 1 and the gradient-based approach of Corollary 2 to the case of mixture models.
- Section 5 is devoted to related work. We explain in more detail the links between our framework and existing Variational Inference methods for α -divergence minimisation. We

also connect our approach to the Power Descent algorithm from Daudel et al. (2021) and provide in Proposition 2 additional monotonicity results which go beyond the case $\alpha \in [0, 1]$. In addition, we obtain that an integrated Expectation Maximization algorithm introduced in Cappé et al. (2008) can be recovered as a special case of our framework.

- In Section 6, we state generic results that solve our maximisation and gradient-based approaches when the variational family is based on the exponential family. More specifically, we introduce in Section 6.1 additional notation for the exponential family and we recall several of its useful properties. Section 6.2 and Section 6.3 then focus on the maximisation approach and gradient-based approach respectively and notably provide the theoretical justification behind some special cases mentioned earlier in the paper, such as Gaussian densities in Example 1. Lastly, Section 6.4 extends once again the maximisation approach in order to incorporate variational families that were not previously included in our framework, yet draw on the exponential family (such as multivariate Student’s t densities in Example 5).

- In Section 7, we focus on the important case of Gaussian Mixture Models (GMMs) and we discuss practical implementations of our algorithms for GMMs optimisation.

- Finally, we show in Section 8 that our enhanced framework has empirical benefits that permit us to improve on the existing algorithms presented in Section 5 when we consider a variety of multimodal targets as well as a real data example.

2. Notation and Optimisation Problem

Let (Y, \mathcal{Y}, ν) be a measured space, where ν is a σ -finite measure on (Y, \mathcal{Y}) . Assume that we have access to some observed variables \mathcal{D} generated from a probabilistic model $p(\cdot|y)$ parameterised by a hidden random variable $y \in Y$ that is drawn from a certain prior p_0 . The posterior density of the latent variable y given the data \mathcal{D} is then given by

$$p(y|\mathcal{D}) = \frac{p(y, \mathcal{D})}{p(\mathcal{D})} = \frac{p_0(y)p(\mathcal{D}|y)}{p(\mathcal{D})},$$

where $p(\mathcal{D}) = \int_Y p_0(y)p(\mathcal{D}|y)\nu(dy)$ is called the *marginal likelihood* or *model evidence*. For many interesting choices of models, sampling directly from the posterior density is impossible and the marginal likelihood is also unknown or too costly to compute.

To address this difficulty, Variational Inference methods approximate the posterior density by a simpler probability density belonging to a given variational family \mathcal{Q} (from which it is easy to sample from). They select the best probability density in \mathcal{Q} by solving an optimisation problem that involves a certain measure of dissimilarity, which is chosen to be the α -divergence in this paper due to its advantages compared to the more traditional exclusive Kullback-Leibler divergence (see for example Minka, 2005).

More precisely, let us denote by \mathbb{P} the probability measure on (Y, \mathcal{Y}) with corresponding density $p(\cdot|\mathcal{D})$ with respect to ν . As for the variational family, let us denote by \mathbb{Q} the probability measure on (Y, \mathcal{Y}) with associated density $q \in \mathcal{Q}$ with respect to ν . We let f_α be the convex function on $(0, +\infty)$ defined by $f_0(u) = -\log(u)$, $f_1(u) = u \log u$ and $f_\alpha(u) = \frac{1}{\alpha(\alpha-1)} [u^\alpha - 1]$ for all $\alpha \in \mathbb{R} \setminus \{0, 1\}$. Then, assuming that both \mathbb{Q} and \mathbb{P} are

absolutely continuous w.r.t. ν , the α -divergence between \mathbb{Q} and \mathbb{P} (extended by continuity to the cases $\alpha = 0$ and $\alpha = 1$ as for example done in Cichocki and Amari, 2010) is given by

$$D_\alpha(\mathbb{Q}||\mathbb{P}) = \int_Y f_\alpha \left(\frac{q(y)}{p(y|\mathcal{D})} \right) p(y|\mathcal{D}) \nu(dy) , \quad (1)$$

and the Variational Inference optimisation problem we aim at solving is $\inf_{q \in \mathcal{Q}} D_\alpha(\mathbb{Q}||\mathbb{P})$. Notably, it can easily be proven that this optimisation problem is equivalent to solving

$$\inf_{q \in \mathcal{Q}} \Psi_\alpha(q; p) \quad \text{with} \quad p(y) = p(y, \mathcal{D}) \quad \text{for all } y \in Y , \quad (2)$$

where, for all measurable non-negative function p on (Y, \mathcal{Y}) and for all $q \in \mathcal{Q}$, we have set

$$\Psi_\alpha(q; p) = \int_Y f_\alpha \left(\frac{q(y)}{p(y)} \right) p(y) \nu(dy) .$$

As the unknown constant $p(\mathcal{D})$ does not appear anymore in the optimisation problem (2), this formulation is often preferred. Therefore, we consider the general optimisation problem

$$\inf_{q \in \mathcal{Q}} \Psi_\alpha(q; p) , \quad (3)$$

where p is any measurable positive function on (Y, \mathcal{Y}) . Note that we may drop the dependency on p in Ψ_α for notational ease and when no ambiguity occurs.

At this stage, we are left with the choice of the variational family \mathcal{Q} appearing in the optimisation problem (3). The natural idea in Variational Inference and the starting point of our approach is then to work within a parametric family : letting (T, \mathcal{T}) be a measurable space, $K : (\theta, A) \mapsto \int_A k(\theta, y) \nu(dy)$ be a Markov transition kernel on $T \times Y$ with kernel density k defined on $T \times Y$, we consider a parametric family of the form

$$\mathcal{Q} = \{q : y \mapsto k(\theta, y) : \theta \in T\} .$$

3. An Iterative Algorithm for Optimising $\Psi_\alpha(k(\theta, \cdot); p)$

In this section, our goal is to define iterative procedures which optimise $\Psi_\alpha(k(\theta, \cdot); p)$ with respect to θ and which are such that they ensure a *systematic decrease* in Ψ_α at each step. For this purpose, we start by introducing some mild conditions on k , p and ν that will be used throughout the paper.

(A1) The kernel density k on $T \times Y$, the function p on Y and the σ -finite measure ν on (Y, \mathcal{Y}) satisfy, for all $(\theta, y) \in T \times Y$, $k(\theta, y) > 0$, $p(y) \geq 0$ and $0 < \int_Y p(y) \nu(dy) < \infty$.

From this point onwards, and as announced in the previous section, we will drop the dependency on p in Ψ_α for notational convenience. Let us now construct a sequence $(\theta_n)_{n \geq 1}$ valued in T such that the sequence $(\Psi_\alpha(k(\theta_n, \cdot)))_{n \geq 1}$ is decreasing. The core idea of our approach will rely on the following proposition.

Proposition 1 *Assume (A1). For all $\alpha \in [0, 1)$ and all $\theta, \theta' \in \mathbb{T}$ such that $\Psi_\alpha(k(\theta', \cdot)) < \infty$, it holds that*

$$\Psi_\alpha(k(\theta, \cdot)) \leq \int_{\mathcal{Y}} \frac{k(\theta', y)^\alpha p(y)^{1-\alpha}}{\alpha - 1} \log \left(\frac{k(\theta, y)}{k(\theta', y)} \right) \nu(dy) + \Psi_\alpha(k(\theta', \cdot)) . \quad (4)$$

Moreover, if $\alpha = 0$, this inequality is an equality and, if $\alpha \in (0, 1)$, equality in (4) is equivalent to having $k(\theta, y) = k(\theta', y)$ for ν -a.e. $y \in \{p > 0\}$.

Proof We treat the two cases $\alpha = 0$ and $\alpha \in (0, 1)$ separately.

(a) Case $\alpha = 0$ with $f_0(u) = -\log(u)$ for all $u > 0$. This case is immediate since

$$\Psi_0(k(\theta, \cdot)) = - \int_{\mathcal{Y}} p(y) \log \left(\frac{k(\theta, y)}{k(\theta', y)} \right) \nu(dy) + \Psi_0(k(\theta', \cdot)) .$$

(b) Case $\alpha \in (0, 1)$ with $f_\alpha(u) = \frac{1}{\alpha(\alpha-1)} [u^\alpha - 1]$ for all $u > 0$. We have that

$$\begin{aligned} \Psi_\alpha(k(\theta, \cdot)) &= \int_{\mathcal{Y}} \frac{\left[\left(\frac{k(\theta, y)}{p(y)} \right)^\alpha - 1 \right]}{\alpha(\alpha - 1)} p(y) \nu(dy) \\ &= \int_{\mathcal{Y}} \left(\frac{k(\theta', y)}{p(y)} \right)^\alpha \frac{\left[\left(\frac{k(\theta, y)}{k(\theta', y)} \right)^\alpha - 1 \right]}{\alpha(\alpha - 1)} p(y) \nu(dy) + \Psi_\alpha(k(\theta', \cdot)) \end{aligned}$$

Furthermore, using that $\log(u^\alpha) \leq u^\alpha - 1$ for all $u > 0$ with equality only if $u = 1$, and since $\alpha \in (0, 1)$, this inequality becomes

$$\frac{u^\alpha - 1}{\alpha(\alpha - 1)} \leq \frac{\log(u^\alpha)}{\alpha(\alpha - 1)} = \frac{\log(u)}{\alpha - 1} .$$

Applying this with $u = k(\theta, y)/k(\theta', y)$ in the previous expression, we get that

$$\Psi_\alpha(k(\theta, \cdot)) \leq \int_{\mathcal{Y}} \frac{k(\theta', y)^\alpha p(y)^{1-\alpha}}{\alpha - 1} \log \left(\frac{k(\theta, y)}{k(\theta', y)} \right) \nu(dy) + \Psi_\alpha(k(\theta', \cdot))$$

which is exactly (4). Moreover, the equality holds if $k(\theta, y)/k(\theta', y) = 1$ for ν -a.e. y such that $k(\theta', y)^\alpha p(y)^{1-\alpha} > 0$, which, by (A1), only happens if $p(y) > 0$. ■

This result then allows us to deduce Theorem 1 below.

Theorem 1 *Assume (A1). Let $\alpha \in [0, 1)$ and starting from an initial $\theta_1 \in \mathbb{T}$, let $(\theta_n)_{n \geq 1}$ be defined iteratively such that for all $n \geq 1$,*

$$\int_{\mathcal{Y}} \frac{k(\theta_n, y)^\alpha p(y)^{1-\alpha}}{\alpha - 1} \log \left(\frac{k(\theta_{n+1}, y)}{k(\theta_n, y)} \right) \nu(dy) \leq 0 . \quad (5)$$

Further assume that $\Psi_\alpha(k(\theta_1, \cdot)) < \infty$. Then, $\Psi_\alpha(k(\theta_{n+1}, \cdot)) \leq \Psi_\alpha(k(\theta_n, \cdot))$ for all $n \geq 1$ with an equality occurring only if (5) is also an equality.

Proof Applying Proposition 1 iteratively for $n = 1, 2, \dots$ with $\theta = \theta_{n+1}$ and $\theta' = \theta_n$ in (4) combined with (5), we get that $\Psi_\alpha(k(\theta_n, \cdot)) < \infty$ and

$$\Psi_\alpha(k(\theta_{n+1}, \cdot)) \leq \int_Y \frac{k(\theta_n, y)^\alpha p(y)^{1-\alpha}}{\alpha - 1} \log \left(\frac{k(\theta_{n+1}, y)}{k(\theta_n, y)} \right) \nu(dy) + \Psi_\alpha(k(\theta_n, \cdot)) \leq \Psi_\alpha(k(\theta_n, \cdot)) .$$

Suppose now that $\Psi_\alpha(k(\theta_{n+1}, \cdot)) = \Psi_\alpha(k(\theta_n, \cdot))$ for some $n \geq 1$. Then we have equalities in the previous display and thus an equality in (5) and in (4) with $\theta = \theta_{n+1}$ and $\theta' = \theta_n$. ■

At this point, we seek to find iterative schemes satisfying (5). To do so, for all $\alpha \in [0, 1)$, all $n \geq 1$ and all $y \in Y$, we introduce the notation:

$$\begin{aligned} \varphi_n^{(\alpha)}(y) &= k(\theta_n, y)^\alpha p(y)^{1-\alpha} \\ \check{\varphi}_n^{(\alpha)} &= \frac{\varphi_n^{(\alpha)}}{\int \varphi_n^{(\alpha)} d\nu} . \end{aligned} \tag{6}$$

Note that under (A1), the normalisation in $\check{\varphi}_n^{(\alpha)}$ satisfies $0 < \int \varphi_n^{(\alpha)} d\nu < \infty$, where the right inequality in particular follows from Jensen's inequality applied to the strictly concave function $u \mapsto u^{1-\alpha}$. We then state our first corollary.

Corollary 1 (Maximisation approach) *Assume (A1). Let $\alpha \in [0, 1)$ and let $(b_n)_{n \geq 1}$ be a non-negative sequence. Starting from an initial $\theta_1 \in \mathbb{T}$, let $(\theta_n)_{n \geq 1}$ be defined iteratively as follows*

$$\theta_{n+1} = \operatorname{argmax}_{\theta \in \mathbb{T}} \int_Y \left[\varphi_n^{(\alpha)}(y) + b_n k(\theta_n, y) \right] \log \left(\frac{k(\theta, y)}{k(\theta_n, y)} \right) \nu(dy) , \quad n \geq 1 , \tag{7}$$

where we assume that this argmax is uniquely defined at each step. Then (5) holds at all times $n \geq 1$ and we can apply Theorem 1.

Proof Under (A1) and using that $\alpha \in [0, 1)$, we can rewrite (5) as

$$\int_Y \varphi_n^{(\alpha)}(y) \log \left(\frac{k(\theta_{n+1}, y)}{k(\theta_n, y)} \right) \nu(dy) \geq 0 .$$

Since $\int_Y k(\theta_n, y) \log \left(\frac{k(\theta_{n+1}, y)}{k(\theta_n, y)} \right) \nu(dy) = -D_1(K(\theta_n, \cdot) || K(\theta_{n+1}, \cdot)) \leq 0$, we also obtain that

$$\int_Y \varphi_n^{(\alpha)}(y) \log \left(\frac{k(\theta_{n+1}, y)}{k(\theta_n, y)} \right) \nu(dy) \geq \int_Y \left[\varphi_n^{(\alpha)}(y) + b_n k(\theta_n, y) \right] \log \left(\frac{k(\theta_{n+1}, y)}{k(\theta_n, y)} \right) \nu(dy) .$$

Thus, (5) holds at all times n if the following condition is satisfied:

$$\int_Y \left[\varphi_n^{(\alpha)}(y) + b_n k(\theta_n, y) \right] \log \left(\frac{k(\theta_{n+1}, y)}{k(\theta_n, y)} \right) \nu(dy) \geq 0 , \quad n \geq 1 ,$$

which is implied by the definition of θ_{n+1} in (7). ■

We now make three comments regarding Corollary 1.

- The r.h.s. of (7) can be simplified, yielding the equivalent iterative scheme

$$\theta_{n+1} = \operatorname{argmax}_{\theta \in \mathbb{T}} \int_{\mathbb{Y}} \left[\varphi_n^{(\alpha)}(y) + b_n k(\theta_n, y) \right] \log k(\theta, y) \nu(dy) , \quad n \geq 1 ,$$

provided that $\int_{\mathbb{Y}} \left[\varphi_n^{(\alpha)}(y) + b_n k(\theta_n, y) \right] |\log k(\theta_n, y)| \nu(dy) < \infty$ for all $n \geq 1$.

- While it suffices to find any θ_{n+1} so that (5) holds to obtain a systematic decrease in Ψ_α , defining θ_{n+1} as in (7) enables us to solve this argmax problem for notable choices of kernel density k . A remarkable aspect is indeed that (7) is written as a maximisation problem involving the logarithm of the ratio $k(\theta, y)/k(\theta_n, y)$ whereas Ψ_α is not directly expressed in terms of logarithm of this ratio for $\alpha \notin \{0, 1\}$. As a result, we can use the *maximisation approach* of Corollary 1 to derive simple update rules for $(\theta_n)_{n \geq 1}$. An example and a lemma are provided thereafter to illustrate this fact and more general results regarding exponential families will follow later in Section 6.
- The sequence $(b_n)_{n \geq 1}$ appearing in (7) is arbitrary and can as a result be chosen by the practitioner. It is responsible for introducing a regularisation term in the argmax problem (7) so that the larger b_n is, the closer θ_{n+1} is to θ_n .

We next provide a motivating example where the maximisation approach is applicable.

Example 1 (Maximisation approach for a Gaussian density) *We consider a d -dimensional Gaussian density, in which case $\mathbb{Y} = \mathbb{R}^d$, ν is the d -dimensional Lebesgue measure and $k(\theta, y) = \mathcal{N}(y; m, \Sigma)$, where $\theta = (m, \Sigma) \in \mathbb{T}$ denotes the mean and covariance matrix of the Gaussian density. We let the parameter set \mathbb{T} include all possible means m in \mathbb{R}^d and all possible positive-definite covariance matrices Σ . Finally, denoting by $\|\cdot\|$ the Euclidean norm, we assume that the non-negative function p defined on \mathbb{Y} satisfies*

$$0 < \int_{\mathbb{Y}} \left(1 + \|y\|^{2/(1-\alpha)} \right) p(y) dy < \infty . \quad (8)$$

Then (A1) holds. Furthermore, starting from any $\theta_1 = (m_1, \Sigma_1) \in \mathbb{T}$ so that $\Psi_\alpha(k(\theta_1, \cdot)) < \infty$ and denoting $\theta_n = (m_n, \Sigma_n)$ for all $n \geq 1$, it holds that: for all $n \geq 1$ and all $\gamma_n \in (0, 1]$, there exists $b_n \geq 0$ such that the argmax problem (7) has a unique solution defined by

$$\begin{aligned} m_{n+1} &= \gamma_n \hat{m}_n + (1 - \gamma_n) m_n , \\ \Sigma_{n+1} &= \gamma_n \hat{\Sigma}_n + (1 - \gamma_n) \Sigma_n + \gamma_n (1 - \gamma_n) (\hat{m}_n - m_n) (\hat{m}_n - m_n)^T , \end{aligned} \quad (9)$$

where, using the definition of $\check{\varphi}_n^{(\alpha)}$ in (6), we set

$$\begin{aligned} \hat{m}_n &= \int_{\mathbb{Y}} y \check{\varphi}_n^{(\alpha)}(y) \nu(dy) , \\ \hat{\Sigma}_n &= \int_{\mathbb{Y}} yy^T \check{\varphi}_n^{(\alpha)}(y) \nu(dy) - \hat{m}_n \hat{m}_n^T . \end{aligned}$$

Here, the detailed derivation of (9) is deferred to Section 6.2 as we focus on interpreting this result in light of Corollary 1 (in particular we will see in Corollary 6 that the case

$\gamma_n = 1$ corresponding to $b_n = 0$ will require an additional non-degenerate condition on p). By Corollary 1, the systematic decrease in $(\Psi_\alpha(k(\theta_n, \cdot)))_{n \geq 1}$ holds for any choice of sequence $(\gamma_n)_{n \geq 1}$ valued in $(0, 1]$ in the updates (9). In addition, γ_n permits us to build a tradeoff at time $n \geq 1$ between selecting an update close to the current parameter (that is, taking γ_n close to zero) and choosing the Gaussian density with exactly the same mean and covariance matrix as $\check{\varphi}_n^{(\alpha)}$ (that is, setting $\gamma_n = 1$). This is a key idea that is linked to the regularisation term appearing in (7) and that we will revisit several times throughout the paper.

The maximisation approach of Corollary 1 is also applicable to the commonly-used *mean-field* variational family, which approximates the true density by a density with independent components parameterised separately. The following lemma indeed states that the global argmax problem can be separated into component-wise argmax problems in that case.

Lemma 1 (Maximisation approach for the mean-field family) *A generic member of the mean-field variational family is given by $k(\theta, y) = \prod_{\ell=1}^L k^{(\ell)}(\theta^{(\ell)}, y^{(\ell)})$ with $\theta = (\theta^{(1)}, \dots, \theta^{(L)}) \in \mathbb{T}$ and $y = (y^{(1)}, \dots, y^{(L)}) \in \mathbb{Y}$. Then, starting from $\theta_1 \in \mathbb{T}$ and denoting $\theta_n = (\theta_n^{(1)}, \dots, \theta_n^{(L)})$ for all $n \geq 1$, solving (7) yields the following update formulas:*

$$\theta_{n+1}^{(\ell)} = \operatorname{argmax}_{\theta^{(\ell)}} \int_{\mathbb{Y}} \left[\varphi_n^{(\alpha)}(y) + b_n k(\theta_n, y) \right] \log \left(\frac{k^{(\ell)}(\theta^{(\ell)}, y^{(\ell)})}{k^{(\ell)}(\theta_n^{(\ell)}, y^{(\ell)})} \right) \nu(dy), \quad \ell = 1 \dots L, \quad n \geq 1.$$

The maximisation approach is not the only way to satisfy (5). Indeed, this can also be achieved by taking a *gradient-based* approach and relying on additional smoothness conditions (see Appendix A.1 for the definition of β -smoothness), as written in Corollary 2.

Corollary 2 (Gradient-based approach) *Assume (A1). Let $\mathbb{T} \subseteq \mathbb{R}^d$ be a convex set, let $\alpha \in [0, 1)$ and let $(\gamma_n)_{n \geq 1}$ be valued in $(0, 1]$. Starting from an initial $\theta_1 \in \mathbb{T}$, let $(\theta_n)_{n \geq 1}$ be defined iteratively as follows*

$$\theta_{n+1} = \theta_n - \frac{\gamma_n}{\beta_n} \nabla g_n(\theta_n), \quad n \geq 1, \quad (10)$$

where $(g_n)_{n \geq 1}$ is the sequence of functions defined by: for all $n \geq 1$ and all $\theta \in \mathbb{T}$,

$$g_n(\theta) = \int_{\mathbb{Y}} \frac{\varphi_n^{(\alpha)}(y)}{\alpha - 1} \log \left(\frac{k(\theta, y)}{k(\theta_n, y)} \right) \nu(dy) \quad (11)$$

and g_n is assumed to be β_n -smooth. Then (5) holds for all $n \geq 1$ and we can apply Theorem 1.

Proof Since $\gamma_n \in (0, 1]$ and g_n is a β_n -smooth function by assumption, we can apply Lemma 2 and we obtain that for all $n \geq 1$,

$$g_n(\theta_n) - g_n \left(\theta_n - \frac{\gamma_n}{\beta_n} \nabla g_n(\theta_n) \right) \geq \frac{\gamma_n}{2\beta_n} \|\nabla g_n(\theta_n)\|^2.$$

Thus, by definition of θ_{n+1} in (10), we have $0 = g_n(\theta_n) \geq g_n(\theta_{n+1})$, which in turn implies (5) and the proof is concluded. \blacksquare

Let us now reflect on the implications of Corollary 2. Under common differentiability assumptions, we can write: for all $n \geq 1$ and all $\theta' \in \mathsf{T}$

$$\nabla g_n(\theta') = \int_{\mathsf{Y}} \frac{\varphi_n^{(\alpha)}(y)}{\alpha - 1} \frac{\partial \log k(\theta, y)}{\partial \theta} \Big|_{(\theta, y) = (\theta', y)} \nu(dy) .$$

Then, (10) becomes

$$\theta_{n+1} = \theta_n - \frac{\gamma_n}{\beta_n} \int_{\mathsf{Y}} \frac{\varphi_n^{(\alpha)}(y)}{\alpha - 1} \frac{\partial \log k(\theta, y)}{\partial \theta} \Big|_{(\theta, y) = (\theta_n, y)} \nu(dy) , \quad n \geq 1 . \quad (12)$$

Under this form, the iterative scheme (12) bears similarities with Gradient Descent iterations for α -divergence and Rényi's α -divergence minimisation. Indeed, given a learning rate policy $(r_n)_{n \geq 1}$ and setting $p = p(\cdot, \mathscr{D})$, such Gradient Descent iterations are given respectively by

$$\begin{aligned} \theta_{n+1} &= \theta_n - r_n \int_{\mathsf{Y}} \frac{\varphi_n^{(\alpha)}(y)}{\alpha - 1} \frac{\partial \log k(\theta, y)}{\partial \theta} \Big|_{(\theta, y) = (\theta_n, y)} \nu(dy) , \quad n \geq 1 \\ \theta_{n+1} &= \theta_n - r_n \int_{\mathsf{Y}} \check{\varphi}_n^{(\alpha)}(y) \frac{\partial \log k(\theta, y)}{\partial \theta} \Big|_{(\theta, y) = (\theta_n, y)} \nu(dy) , \quad n \geq 1 \end{aligned}$$

(we refer to Appendix A.2 for details regarding how these updates are obtained). Building on this comment, let us now give an example where the conditions on $(g_n)_{n \geq 1}$ from Corollary 2 are satisfied and show how Gradient Descent steps (in that case for Rényi's α -divergence minimisation) can originate from our gradient-based approach.

Example 2 (Gradient-based approach for a Gaussian density) *We consider the case of a d -dimensional Gaussian density with $k(\theta, y) = \mathcal{N}(y; m, \Sigma)$, where this time $\theta = m$. Further assume that $\mathsf{T} \subseteq \mathbb{R}^d$ is a convex subset and that $\Sigma = \sigma^2 \mathbf{I}_d$, where $\sigma^2 > 0$ and \mathbf{I}_d denotes the d -dimensional identity matrix. Finally, denoting by $\|\cdot\|$ the Euclidean norm, we assume that the non-negative function p defined on Y satisfies*

$$0 < \int_{\mathsf{Y}} \left(1 + \|y\|^{1/(1-\alpha)}\right) p(y) dy < \infty . \quad (13)$$

Then, (A1) holds. Furthermore, denoting $\theta_n = m_n$ and setting $\beta_n = \sigma^{-2}(1 - \alpha)^{-1} \int \varphi_n^{(\alpha)} d\nu$ for all $n \geq 1$, the conditions on $(g_n)_{n \geq 1}$ from Corollary 2 are satisfied so that the mean of the Gaussian density $\mathcal{N}(y; m, \Sigma)$ can be optimised as follows: for all $n \geq 1$ and all $\gamma_n \in [0, 1]$

$$m_{n+1} = \gamma_n \hat{m}_n + (1 - \gamma_n) m_n , \quad (14)$$

where $\hat{m}_n = \int_{\mathsf{Y}} y \check{\varphi}_n^{(\alpha)}(y) \nu(dy)$ (we refer to Corollary 8 for the detailed derivation). Setting $p = p(\cdot, \mathscr{D})$, the update (14) can notably be seen as a Gradient Descent step for Rényi's α -divergence minimisation with a learning rate $r_n = \sigma^2(1 - \alpha)\gamma_n$, where $\gamma_n \in [0, 1]$.

We now make two important comments.

- Corollary 2 sheds light on the links between our approach and the more traditional Gradient Descent methodology for optimising objective functions based on the α -divergence in Variational Inference (Li and Turner, 2016; Dieng et al., 2017).

Unlike the usual Gradient Descent methodology, Corollary 2 requires a smoothness condition on g_n . Smoothness conditions are tools that are often used to obtain stronger convergence guarantees for Variational Inference algorithms (see for example Buchholz et al., 2018; Alquier and Ridgway, 2020). Yet, it can be difficult to satisfy these smoothness assumptions in practice, even if some results have been derived for specific variational families when using the exclusive Kullback-Leibler divergence as the objective function (Domke, 2020).

To the best of our knowledge, no such results have been proven for α -divergence minimisation. In our case, we obtain that a smoothness condition on g_n translates into a systematic decrease in Ψ_α . As we shall see in the forthcoming section, being able to establish monotonicity results on $(\theta_n)_{n \geq 1}$ will come in handy as we try to go beyond the framework considered in Section 3.

- To put things into perspective with the maximisation approach from earlier, observe that the updates on the means in Examples 1 and 2 coincide. The difference between the two examples is due to the fact that the former provides an update for the covariance matrix as well, *without sacrificing the monotonicity of the overall algorithm*. It is in fact hard to derive a covariance matrix update in the Gaussian case using the gradient-based approach (see Section 6.3 for details). In that sense, the maximisation approach provides an interesting alternative that can bypass this difficult smoothness assumption on g_n . We will further delve into this aspect as we reach Section 6.

So far, the log function appearing in Theorem 1 considerably eased the derivation of iterative update formulas for well-chosen kernel densities k that can exploit this log structure (Examples 1-2 and Lemma 1). Yet, these choices of kernel densities can be too restrictive to fully capture complex and multimodal posterior densities. Since working with a larger variational family might lead to more accurate approximations of the posterior density, an idea is then to investigate whether the iterative update formulas from Section 3 can be generalised to the mixture model variational family (for example whether we can extend the updates for a Gaussian kernel from Example 1 to Gaussian Mixture Models).

As we shall see in the next section, further theoretical developments will be required in order to derive valid iterative schemes that optimise both the mixture weights and the mixture components parameters of a given mixture model.

4. Extension to Mixture Models

Let us first formally define the class of mixture models we are going to be working with. Given $J \in \mathbb{N}^*$, we introduce the simplex of \mathbb{R}^J :

$$\mathcal{S}_J = \left\{ \boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_J) \in \mathbb{R}^J : \forall j \in \{1, \dots, J\}, \lambda_j \geq 0 \text{ and } \sum_{j=1}^J \lambda_j = 1 \right\},$$

we define $\mathcal{S}_J^+ = \{\boldsymbol{\lambda} \in \mathcal{S}_J : \forall j \in \{1, \dots, J\}, \lambda_j > 0\}$ and we denote $\Theta = (\theta_1, \dots, \theta_J) \in \mathbb{T}^J$. We consider the mixture model variational family given by

$$\mathcal{Q} = \left\{ q : y \mapsto \mu_{\boldsymbol{\lambda}, \Theta} k(y) := \sum_{j=1}^J \lambda_j k(\theta_j, y) : \boldsymbol{\lambda} \in \mathcal{S}_J, \Theta \in \mathbb{T}^J \right\} \quad (15)$$

that is, we are interested in solving the optimisation problem

$$\inf_{\boldsymbol{\lambda} \in \mathcal{S}_J, \Theta \in \mathbb{T}^J} \Psi_\alpha(\mu_{\boldsymbol{\lambda}, \Theta} k),$$

with $J > 1$. Let us next denote $\boldsymbol{\lambda}_n = (\lambda_{j,n})_{1 \leq j \leq J}$ and $\Theta_n = (\theta_{j,n})_{1 \leq j \leq J}$ for all $n \geq 1$. For convenience, we also introduce the shorthand notation $\mu_n k = \mu_{\boldsymbol{\lambda}_n, \Theta_n} k$ and

$$\begin{aligned} \varphi_{j,n}^{(\alpha)}(y) &= k(\theta_{j,n}, y) \left(\frac{\mu_n k(y)}{p(y)} \right)^{\alpha-1} \\ \check{\varphi}_{j,n}^{(\alpha)} &= \frac{\varphi_{j,n}^{(\alpha)}}{\int \varphi_{j,n}^{(\alpha)} d\nu} \end{aligned} \quad (16)$$

for all $\alpha \in [0, 1)$, all $j = 1 \dots J$, all $n \geq 1$ and all $y \in \mathcal{Y}$.

Our goal in this section is to derive iterative schemes for the mixture weights and the mixture components parameters $(\boldsymbol{\lambda}_n, \Theta_n)_{n \geq 1}$ ensuring that the sequence $(\Psi_\alpha(\mu_n k))_{n \geq 1}$ is decreasing. As Theorem 1 holds for any choice of parametric family, a first idea is to apply Theorem 1 to the variational family (15), which gives Corollary 3 below.

Corollary 3 *Assume (A1). Let $J \in \mathbb{N}^*$, let $\alpha \in [0, 1)$ and starting from an initial parameter set $(\boldsymbol{\lambda}_1, \Theta_1) \in \mathcal{S}_J^+ \times \mathbb{T}^J$, let $(\boldsymbol{\lambda}_n, \Theta_n)_{n \geq 1}$ be defined iteratively such that for all $n \geq 1$,*

$$\int_{\mathcal{Y}} \frac{(\mu_n k(y))^\alpha p(y)^{1-\alpha}}{\alpha - 1} \log \left(\frac{\mu_{n+1} k(y)}{\mu_n k(y)} \right) \nu(dy) \leq 0. \quad (17)$$

Further assume that $\Psi_\alpha(\mu_1 k) < \infty$. Then, $\Psi_\alpha(\mu_{n+1} k) \leq \Psi_\alpha(\mu_n k)$ for all $n \geq 1$.

Observe that we are in a less favourable situation in Corollary 3 with $J > 1$ compared to the cases we previously studied in Section 3. Indeed, we now have a ratio of sums inside the log function in (17), meaning that the approach from Theorem 1 to derive simple iterative schemes does not directly transfer to the variational family (15) for choices of kernel densities k identified in Examples 1-2 and Lemma 1. However, by carefully exploiting the condition (17), we are able to overcome this difficulty in our second main theorem below.

Theorem 2 *Assume (A1). Let $J \in \mathbb{N}^*$, $\alpha \in [0, 1)$ and starting from an initial parameter set $(\boldsymbol{\lambda}_1, \Theta_1) \in \mathcal{S}_J^+ \times \mathbb{T}^J$, let $(\boldsymbol{\lambda}_n, \Theta_n)_{n \geq 1}$ be defined iteratively such that for all $n \geq 1$,*

$$\int_{\mathcal{Y}} \sum_{j=1}^J \lambda_{j,n} \frac{\varphi_{j,n}^{(\alpha)}(y)}{\alpha - 1} \log \left(\frac{\lambda_{j,n+1}}{\lambda_{j,n}} \right) \nu(dy) \leq 0 \quad (18)$$

$$\int_{\mathcal{Y}} \sum_{j=1}^J \lambda_{j,n} \frac{\varphi_{j,n}^{(\alpha)}(y)}{\alpha - 1} \log \left(\frac{k(\theta_{j,n+1}, y)}{k(\theta_{j,n}, y)} \right) \nu(dy) \leq 0. \quad (19)$$

Further assume that $\Psi_\alpha(\mu_1 k) < \infty$. Then, $\Psi_\alpha(\mu_{n+1} k) \leq \Psi_\alpha(\mu_n k)$ for all $n \geq 1$.

Proof By Corollary 3, we can conclude if we show that (18) and (19) together imply (17). To show this, first observe that since the function $u \mapsto \frac{1}{\alpha-1} \log(u)$ is convex and $\alpha \in [0, 1)$, Jensen's inequality implies that: for all $y \in \mathcal{Y}$ and all $n \geq 1$,

$$\begin{aligned} \frac{1}{\alpha-1} \log \left(\frac{\mu_{n+1}k(y)}{\mu_n k(y)} \right) &= \frac{1}{\alpha-1} \log \left(\sum_{j=1}^J \frac{\lambda_{j,n} k(\theta_{j,n}, y)}{\sum_{\ell=1}^J \lambda_{\ell,n} k(\theta_{\ell,n}, y)} \frac{\lambda_{j,n+1} k(\theta_{j,n+1}, y)}{\lambda_{j,n} k(\theta_{j,n}, y)} \right) \\ &\leq \sum_{j=1}^J \frac{\lambda_{j,n} k(\theta_{j,n}, y)}{\sum_{\ell=1}^J \lambda_{\ell,n} k(\theta_{\ell,n}, y)} \frac{1}{\alpha-1} \log \left(\frac{\lambda_{j,n+1} k(\theta_{j,n+1}, y)}{\lambda_{j,n} k(\theta_{j,n}, y)} \right) \end{aligned}$$

that is:

$$\frac{1}{\alpha-1} \log \left(\frac{\mu_{n+1}k(y)}{\mu_n k(y)} \right) \leq \sum_{j=1}^J \frac{\lambda_{j,n}}{\alpha-1} \frac{k(\theta_{j,n}, y)}{\mu_n k(y)} \log \left(\frac{\lambda_{j,n+1} k(\theta_{j,n+1}, y)}{\lambda_{j,n} k(\theta_{j,n}, y)} \right).$$

Multiplying by $(\mu_n k(y))^\alpha p(y)^{1-\alpha}$ on both sides, integrating with respect to $\nu(dy)$ and using the definition of $\varphi_{j,n}^{(\alpha)}(y)$ in (16), this in turn implies that: for all $n \geq 1$,

$$\begin{aligned} \int_{\mathcal{Y}} \frac{(\mu_n k(y))^\alpha p(y)^{1-\alpha}}{\alpha-1} \log \left(\frac{\mu_{n+1}k(y)}{\mu_n k(y)} \right) \nu(dy) \\ \leq \int_{\mathcal{Y}} \sum_{j=1}^J \lambda_{j,n} \frac{\varphi_{j,n}^{(\alpha)}(y)}{\alpha-1} \log \left(\frac{\lambda_{j,n+1} k(\theta_{j,n+1}, y)}{\lambda_{j,n} k(\theta_{j,n}, y)} \right) \nu(dy). \end{aligned}$$

As a consequence, the condition

$$\int_{\mathcal{Y}} \sum_{j=1}^J \lambda_{j,n} \frac{\varphi_{j,n}^{(\alpha)}(y)}{\alpha-1} \log \left(\frac{\lambda_{j,n+1} k(\theta_{j,n+1}, y)}{\lambda_{j,n} k(\theta_{j,n}, y)} \right) \nu(dy) \leq 0 \quad (20)$$

implies (17). The condition (20) in itself is then straightforwardly implied by (18) and (19) and the proof is concluded. \blacksquare

Strikingly, (18) does not depend on Θ_{n+1} nor does (19) depend on λ_{n+1} in Theorem 2. This means that we can treat these two conditions separately and thus that the weights and components parameters of the mixture can be optimised simultaneously. This result was far from immediate by looking at the initial condition (17) and, as we shall see laterly in Section 8, will lead to reduced computational power in practice.

In addition, we have recovered in (19) the key property used in Section 3: compared to the condition (17) which involved the log of weighted sums of kernel densities, (19) considers a sum of logs of kernel densities. This suggests that we can extend the updates derived in Section 3 for well-chosen kernel densities k to the more general case of mixture models.

Observe finally that the dependency in $\lambda_{j,n+1}$ appearing in (18) is simpler than the dependency in $\theta_{j,n+1}$ appearing in (19) and that is expressed through the kernel density k . As a result, we will first study (18), in the hope of deriving iterative update formulas for the mixture weights that do not require a specific choice of kernel density k . Interestingly, while the natural idea is to perform direct optimisation of the left-hand side of (18), we will derive a more general expression for the mixture weights update, which shall induce numerical advantages later illustrated in Section 8.

4.1 Choice of $(\lambda_n)_{n \geq 1}$

In the following theorem, we identify an update formula which satisfies (18), regardless of the choice of the kernel density k .

Theorem 3 *Assume (A1). Let $\alpha \in [0, 1]$, let $(\eta_n)_{n \geq 1}$ be valued in $(0, 1]$, let $(\kappa_n)_{n \geq 1}$ be such that $(\alpha - 1)\kappa_n \geq 0$ at all times $n \geq 1$ and let $(\Theta_n)_{n \geq 1}$ be any sequence valued in \mathbb{T}^J . Starting from an initial $\lambda_1 \in \mathcal{S}_J^+$, let $(\lambda_n)_{n \geq 1}$ be defined iteratively such that for all $n \geq 1$*

$$\lambda_{j,n+1} = \frac{\lambda_{j,n} \left[\int_Y \varphi_{j,n}^{(\alpha)}(y) \nu(dy) + (\alpha - 1)\kappa_n \right]^{\eta_n}}{\sum_{\ell=1}^J \lambda_{\ell,n} \left[\int_Y \varphi_{\ell,n}^{(\alpha)}(y) \nu(dy) + (\alpha - 1)\kappa_n \right]^{\eta_n}}, \quad j = 1 \dots J. \quad (21)$$

Then (18) holds.

Proof We first check that the integrals appearing in (21) are finite:

$$\int_Y \varphi_{j,n}^{(\alpha)}(y) \nu(dy) < \infty, \quad j = 1, \dots, J, \quad n \geq 1. \quad (22)$$

Using (16), that $\lambda_{j,n} > 0$, that $\lambda_{j,n} k(\theta_{j,n}, y) \leq \mu_n k(y)$ and Jensen's inequality applied to the concave function $u \mapsto u^{1-\alpha}$, we have, for any $j = 1, \dots, J$ and $n \geq 1$,

$$\begin{aligned} \int_Y \varphi_{j,n}^{(\alpha)}(y) \nu(dy) &= \int_Y k(\theta_{j,n}, y) \left(\frac{\mu_n k(y)}{p(y)} \right)^{\alpha-1} \nu(dy) \\ &\leq \frac{1}{\lambda_{j,n}} \int_Y \mu_n k(y) \left(\frac{p(y)}{\mu_n k(y)} \right)^{1-\alpha} \nu(dy) \\ &\leq \frac{1}{\lambda_{j,n}} \left(\int_Y p(y) \nu(dy) \right)^{1-\alpha}. \end{aligned}$$

The bound (22) follows by (A1). Now, to prove (18), we treat the cases $\eta_n = 1$ and $\eta_n \in (0, 1)$ separately.

(a) Case $\eta_n = 1$. Since $(\alpha - 1)\kappa_n \geq 0$ with $\alpha \in (0, 1)$, we have that

$$\kappa_n \sum_{j=1}^J \lambda_{j,n} \log(\lambda_j / \lambda_{j,n}) \geq 0$$

where we have used that $\sum_{j=1}^J \lambda_{j,n} \log(\lambda_j / \lambda_{j,n}) \leq \sum_{j=1}^J \lambda_{j,n} (\lambda_j / \lambda_{j,n} - 1) = 0$. In other words, to obtain (18) in the particular case $\eta_n = 1$, it is enough to show

$$\int_Y \sum_{j=1}^J \lambda_{j,n} \frac{\varphi_{j,n}^{(\alpha)}(y)}{\alpha - 1} \log \left(\frac{\lambda_{j,n+1}}{\lambda_{j,n}} \right) \nu(dy) + \kappa_n \sum_{j=1}^J \lambda_{j,n} \log \left(\frac{\lambda_{j,n+1}}{\lambda_{j,n}} \right) \leq 0$$

that is, since (22) holds,

$$\sum_{j=1}^J \lambda_{j,n} \left[\int_Y \frac{\varphi_{j,n}^{(\alpha)}(y)}{\alpha - 1} \nu(dy) + \kappa_n \right] \log \left(\frac{\lambda_{j,n+1}}{\lambda_{j,n}} \right) \leq 0. \quad (23)$$

Notice then that by definition of $(\lambda_{j,n+1})_{1 \leq j \leq J}$ when $\eta_n = 1$, we can write

$$\lambda_{n+1} = \operatorname{argmin}_{\lambda \in \mathcal{S}_J^+} \sum_{j=1}^J \lambda_{j,n} \left[\int_Y \frac{\varphi_{j,n}^{(\alpha)}(y)}{\alpha - 1} \nu(dy) + \kappa_n \right] \log \left(\frac{\lambda_j}{\lambda_{j,n}} \right).$$

[Indeed, setting $\beta_j = \lambda_{j,n} \left[\int_Y \varphi_{j,n}^{(\alpha)}(y) \nu(dy) + (\alpha - 1) \kappa_n \right]$ and $\bar{\beta}_j = \beta_j / \sum_{\ell=1}^J \beta_\ell$ for all $j = 1 \dots J$, we have that $\sum_{j=1}^J \bar{\beta}_j \log \left(\bar{\beta}_j / \lambda_j \right) \geq 0$ and that this quantity is minimal when $\lambda_j = \bar{\beta}_j$ for $j = 1 \dots J$.] This implies (23) and settles the case $\eta_n = 1$.

(b) For the particular case $\eta_n \in (0, 1)$, we will use that for all $\epsilon > 0$ and all $u > 0$,

$$\log(u) = \frac{1}{\epsilon} \log(u^\epsilon) \geq \frac{1}{\epsilon} \left(1 - \frac{1}{u^\epsilon} \right).$$

Indeed, since $\int_Y \frac{\varphi_{j,n}^{(\alpha)}(y)}{\alpha - 1} \nu(dy) + \kappa_n \leq 0$ for all $j = 1 \dots J$, we can then write that for all $\epsilon > 0$,

$$\begin{aligned} \sum_{j=1}^J \lambda_{j,n} \left[\int_Y \frac{\varphi_{j,n}^{(\alpha)}(y)}{\alpha - 1} \nu(dy) + \kappa_n \right] \log \left(\frac{\lambda_{j,n+1}}{\lambda_{j,n}} \right) \\ \leq \frac{1}{\epsilon} \sum_{j=1}^J \lambda_{j,n} \left[\int_Y \frac{\varphi_{j,n}^{(\alpha)}(y)}{\alpha - 1} \nu(dy) + \kappa_n \right] \left[1 - \left(\frac{\lambda_{j,n}}{\lambda_{j,n+1}} \right)^\epsilon \right]. \end{aligned} \quad (24)$$

Now notice that by definition of $(\lambda_{j,n+1})_{1 \leq j \leq J}$ we can write

$$\lambda_{n+1} = \operatorname{argmin}_{\lambda \in \mathcal{S}_J^+} \frac{1}{\epsilon} \sum_{j=1}^J \lambda_{j,n} \left[\int_Y \frac{\varphi_{j,n}^{(\alpha)}(y)}{\alpha - 1} \nu(dy) + \kappa_n \right] \left[1 - \left(\frac{\lambda_{j,n}}{\lambda_j} \right)^\epsilon \right]$$

when ϵ satisfies $\eta_n = \frac{1}{1+\epsilon}$. [Indeed setting $\beta_j = \lambda_{j,n} \left[\int_Y \varphi_{j,n}^{(\alpha)}(y) \nu(dy) + (\alpha - 1) \kappa_n \right]^{\frac{1}{1+\epsilon}}$ and $\bar{\beta}_j = \beta_j / \sum_{\ell=1}^J \beta_\ell$ for all $j = 1 \dots J$, we have by convexity of the function $u \mapsto u^{1+\epsilon}$ that $\sum_{j=1}^J \left(\bar{\beta}_j / \lambda_j \right)^{1+\epsilon} \lambda_j \geq (\sum_{j=1}^J \bar{\beta}_j)^{1+\epsilon}$ and that this quantity is minimal when $\lambda_j = \bar{\beta}_j$ for $j = 1 \dots J$.] We then deduce that taking $\epsilon = \eta_n^{-1} - 1$ (it is always possible since $\eta_n \in (0, 1)$ by assumption) yields

$$\frac{1}{\epsilon} \sum_{j=1}^J \lambda_{j,n} \left[\int_Y \frac{\varphi_{j,n}^{(\alpha)}(y)}{\alpha - 1} \nu(dy) + \kappa_n \right] \left[1 - \left(\frac{\lambda_{j,n}}{\lambda_{j,n+1}} \right)^\epsilon \right] \leq 0$$

which in turn yields (18) [since combined with (24) it implies (23) which itself implies (18) as seen in the case $\eta_n = 1$]. This settles the case $\eta_n \in (0, 1)$. ■

Notice that as a byproduct of the proof of Theorem 3, the mixture weights update given by (21) can be rewritten under the form: for all $n \geq 1$,

$$\lambda_{n+1} = \operatorname{argmin}_{\lambda \in \mathcal{S}_J^+} h_n(\lambda)$$

where, setting $\epsilon = \eta_n^{-1} - 1$, we have defined for all $\boldsymbol{\lambda} \in \mathcal{S}_J^+$,

$$h_n(\boldsymbol{\lambda}) = \begin{cases} \sum_{j=1}^J \lambda_{j,n} \left[\int_Y \frac{\varphi_{j,n}^{(\alpha)}(y)}{\alpha-1} \nu(dy) + \kappa_n \right] \log \left(\frac{\lambda_j}{\lambda_{j,n}} \right), & \text{if } \eta_n = 1, \\ \frac{1}{\epsilon} \sum_{j=1}^J \lambda_{j,n} \left[\int_Y \frac{\varphi_{j,n}^{(\alpha)}(y)}{\alpha-1} \nu(dy) + \kappa_n \right] \left[1 - \left(\frac{\lambda_{j,n}}{\lambda_j} \right)^\epsilon \right], & \text{if } \eta_n \in (0, 1). \end{cases}$$

More specifically, $h_n(\boldsymbol{\lambda})$ acts as an upper bound of the left-hand side of (18) and we recover exactly the left-hand side of (18) in the particular case $\eta_n = 1$ and $\kappa_n = 0$.

Now that we have established the mixture weights updates (21) in Theorem 3, we are interested in deriving update formulas for the sequence $(\Theta_n)_{n \geq 1}$ satisfying (19), which we will then pair up with (21) in order to apply Theorem 2. From this point onwards, all the proofs of the coming results will be deferred to the appendices to ease the reading.

4.2 Choice of $(\Theta_n)_{n \geq 1}$

We investigate two different approaches for choosing $(\Theta_n)_{n \geq 1}$.

4.2.1 A MAXIMISATION APPROACH

As done in Corollary 1, an idea is to consider an update for $(\Theta_n)_{n \geq 1}$ of the form

$$\Theta_{n+1} = \operatorname{argmax}_{\Theta \in \mathbb{T}^J} g_n(\Theta), \quad n \geq 1,$$

where the function g_n is constructed as a lower bound on $\Theta \in \mathbb{T}^J$ of the function $\Theta \mapsto \int_Y \sum_{j=1}^J \lambda_{j,n} \varphi_{j,n}^{(\alpha)}(y) \log(k(\theta_j, y)/k(\theta_{j,n}, y)) \nu(dy)$ that satisfies $g_n(\Theta_n) = 0$. In doing so, the function $\tilde{g}_n : \Theta \mapsto g_n(\Theta)/(\alpha - 1)$ evaluated at Θ_{n+1} is an upper bound of the left-hand side of (19) and $\tilde{g}_n(\Theta_{n+1}) \leq 0$ implies (19). This leads us to Corollary 4 below.

Corollary 4 (Generalised maximisation approach) *Assume (A1). Let $\alpha \in [0, 1)$, let $(\eta_n)_{n \geq 1}$ be valued in $(0, 1]$ and let $(\kappa_n)_{n \geq 1}$ be such that $(\alpha - 1)\kappa_n \geq 0$ at all times $n \geq 1$. Furthermore, let $(b_{j,n})_{n \geq 1}$ be a non-negative sequence for all $j = 1 \dots J$. Starting from an initial parameter set $(\boldsymbol{\lambda}_1, \Theta_1) \in \mathcal{S}_J^+ \times \mathbb{T}^J$, let $(\boldsymbol{\lambda}_n, \Theta_n)_{n \geq 1}$ be defined iteratively for all $n \geq 1$ in such a way that (21) holds and*

$$\theta_{j,n+1} = \operatorname{argmax}_{\theta \in \mathbb{T}} \int_Y \left[\varphi_{j,n}^{(\alpha)}(y) + b_{j,n} k(\theta_{j,n}, y) \right] \log \left(\frac{k(\theta, y)}{k(\theta_{j,n}, y)} \right) \nu(dy), \quad j = 1 \dots J, \quad (25)$$

where we assume that this argmax is uniquely defined at each step. Then, we can apply Theorem 2.

The proof of this result is deferred to Appendix B.1. Under the assumptions of Corollary 4, Algorithm 1 leads to a systematic decrease in Ψ_α at each step. This result effectively generalises the monotonicity property from Corollary 1 to the case of mixture models and we can deduce simple iterative schemes satisfying (25) for well-chosen kernel densities k . To illustrate this, we provide below the extension of Example 1 to Gaussian Mixture Models (and Lemma 1 can be extended in the same way, see Appendix B.2).

Algorithm 1: Maximisation approach algorithm for mixture models

At iteration n ,

For all $j = 1 \dots J$, set

$$\lambda_{j,n+1} = \frac{\lambda_{j,n} \left[\int_{\mathcal{Y}} \varphi_{j,n}^{(\alpha)}(y) \nu(dy) + (\alpha - 1) \kappa_n \right]^{\eta_n}}{\sum_{\ell=1}^J \lambda_{\ell,n} \left[\int_{\mathcal{Y}} \varphi_{\ell,n}^{(\alpha)}(y) \nu(dy) + (\alpha - 1) \kappa_n \right]^{\eta_n}}$$

$$\theta_{j,n+1} = \operatorname{argmax}_{\theta \in \mathcal{T}} \int_{\mathcal{Y}} \left[\varphi_{j,n}^{(\alpha)}(y) + b_{j,n} k(\theta_{j,n}, y) \right] \log \left(\frac{k(\theta, y)}{k(\theta_{j,n}, y)} \right) \nu(dy) .$$

Example 3 (Maximisation approach for Gaussian Mixture Models) *We consider the case of d -dimensional Gaussian mixture densities, in which case $\mathcal{Y} = \mathbb{R}^d$, ν is the d -dimensional Lebesgue measure and $k(\theta_j, y) = \mathcal{N}(y; m_j, \Sigma_j)$, where $\theta_j = (m_j, \Sigma_j) \in \mathcal{T}$ denotes the mean and covariance matrix of the j -th Gaussian component density. We let the parameter set \mathcal{T} include all possible means in \mathbb{R}^d and all possible positive-definite covariance matrices. Finally, we assume that the non-negative function p defined on \mathcal{Y} satisfies (8).*

Then (A1) holds. Moreover, starting from any $(\lambda_1, \Theta_1) \in \mathcal{S}_J^+ \times \mathcal{T}^J$ so that $\Psi_\alpha(\mu_1 k) < \infty$ and denoting $\theta_{j,n} = (m_{j,n}, \Sigma_{j,n})$ for all $j = 1 \dots J$ and all $n \geq 1$, it holds that: for all $j = 1 \dots J$, all $n \geq 1$ and all $\gamma_{j,n} \in (0, 1]$, there exists $b_{j,n} \geq 0$ such that the argmax problem (25) has a unique solution defined by:

$$m_{j,n+1} = \gamma_{j,n} \hat{m}_{j,n} + (1 - \gamma_{j,n}) m_{j,n}$$

$$\Sigma_{j,n+1} = \gamma_{j,n} \hat{\Sigma}_{j,n} + (1 - \gamma_{j,n}) \Sigma_{j,n} + \gamma_{j,n} (1 - \gamma_{j,n}) (\hat{m}_{j,n} - m_{j,n}) (\hat{m}_{j,n} - m_{j,n})^T ,$$

where, using the definition of $\check{\varphi}_{j,n}^{(\alpha)}$ in (16), we set

$$\hat{m}_{j,n} = \int_{\mathcal{Y}} y \check{\varphi}_{j,n}^{(\alpha)}(y) \nu(dy) ,$$

$$\hat{\Sigma}_{j,n} = \int_{\mathcal{Y}} y y^T \check{\varphi}_{j,n}^{(\alpha)}(y) \nu(dy) - \hat{m}_{j,n} \hat{m}_{j,n}^T .$$

The detailed derivation of the mean and covariance updates above is deferred to Section 6.2. The interpretation of these updates is akin to the one already made in Example 1: $b_{j,n}$ acts as a regularisation parameter which, through $\gamma_{j,n}$, permits a tradeoff between an update close to the current parameter $\theta_{j,n}$ and choosing the Gaussian density with exactly the same mean and covariance matrix as $\check{\varphi}_{j,n}^{(\alpha)}$. As per written in Corollary 4, these updates are compatible with the mixture weights updates (21), resulting in a systematic decrease of $(\Psi_\alpha(\mu_n k))_{n \geq 1}$.

We next present another possible update formula for $(\lambda_n, \Theta_n)_{n \geq 1}$.

4.2.2 A GRADIENT-BASED APPROACH

In the spirit of Corollary 2, we now resort to Gradient Descent steps to satisfy (18).

Algorithm 2: Gradient-based approach algorithm for mixture models

At iteration n ,

For all $j = 1 \dots J$, set

$$\lambda_{j,n+1} = \frac{\lambda_{j,n} \left[\int_{\mathbf{Y}} \varphi_{j,n}^{(\alpha)}(y) \nu(dy) + (\alpha - 1) \kappa_n \right]^{\eta_n}}{\sum_{\ell=1}^J \lambda_{\ell,n} \left[\int_{\mathbf{Y}} \varphi_{\ell,n}^{(\alpha)}(y) \nu(dy) + (\alpha - 1) \kappa_n \right]^{\eta_n}}$$

$$\theta_{j,n+1} = \theta_{j,n} - \frac{\gamma_{j,n}}{\beta_{j,n}} \nabla g_{j,n}(\theta_{j,n}) .$$

Corollary 5 (Generalised gradient-based approach) *Assume (A1). Let $\mathsf{T} \subseteq \mathbb{R}^d$ be a convex set, let $\alpha \in [0, 1)$, let $(\eta_n)_{n \geq 1}$ be valued in $(0, 1]$ and let $(\kappa_n)_{n \geq 1}$ be such that $(\alpha - 1)\kappa_n \geq 0$ at all times n . Furthermore, for all $j = 1 \dots J$, let $(\gamma_{j,n})_{n \geq 1}$ be valued in $(0, 1]$. Starting from an initial parameter set $(\boldsymbol{\lambda}_1, \Theta_1) \in \mathcal{S}_J^+ \times \mathsf{T}^J$, let $(\boldsymbol{\lambda}_n, \Theta_n)_{n \geq 1}$ be defined iteratively for all $n \geq 1$ in such a way that (21) holds and*

$$\theta_{j,n+1} = \theta_{j,n} - \frac{\gamma_{j,n}}{\beta_{j,n}} \nabla g_{j,n}(\theta_{j,n}) , \quad j = 1 \dots J , \quad (26)$$

where for all $j = 1 \dots J$, $(g_{j,n})_{n \geq 1}$ is defined by: for all $n \geq 1$ and all $\theta \in \mathsf{T}$,

$$g_{j,n}(\theta) = \int_{\mathbf{Y}} \frac{\varphi_{j,n}^{(\alpha)}(y)}{\alpha - 1} \log \left(\frac{k(\theta, y)}{k(\theta_{j,n}, y)} \right) \nu(dy) .$$

and $g_{j,n}$ is assumed to be $\beta_{j,n}$ -smooth. Then, we can apply Theorem 2.

The proof of this result is deferred to Appendix B.3. Under the assumptions of Corollary 5, Algorithm 2 ensures a systematic decrease in Ψ_α at each step and Corollary 5 thus extends Corollary 2 to mixture models. Much like what we did for Corollary 2, we want to identify how our updates relate to Gradient Descent-based techniques for optimising Θ . Under common differentiability assumptions, we have: for all $n \geq 1$, all $j = 1 \dots J$ and all $\theta' \in \mathsf{T}$,

$$\nabla g_{j,n}(\theta') = \int_{\mathbf{Y}} \frac{\varphi_{j,n}^{(\alpha)}(y)}{\alpha - 1} \left. \frac{\partial \log k(\theta, y)}{\partial \theta} \right|_{(\theta, y) = (\theta', y)} \nu(dy) ,$$

so that (26) becomes

$$\theta_{j,n+1} = \theta_{j,n} - \frac{\gamma_{j,n}}{\beta_{j,n}} \int_{\mathbf{Y}} \frac{\varphi_{j,n}^{(\alpha)}(y)}{\alpha - 1} \left. \frac{\partial \log k(\theta, y)}{\partial \theta} \right|_{(\theta, y) = (\theta_{j,n}, y)} \nu(dy) , \quad j = 1 \dots J . \quad (27)$$

The link with Gradient Descent-based Variational Inference shall become apparent by (i) writing the update formulas that ensue from Gradient Descent iterations for α -divergence and Rényi's α -divergence minimisation and (ii) understanding how Example 2 generalises to Gaussian Mixture Models. Given $(\boldsymbol{\lambda}_n, \Theta_n) \in \mathcal{S}_J^+ \times \mathsf{T}^J$, an index j in $1 \dots J$ and letting

$p = p(\cdot, \mathcal{D})$, performing *one* Gradient Descent iteration w.r.t. $\theta_{j,n}$ for α -divergence and Rényi's α -divergence minimisation indeed respectively amounts to updating $\theta_{j,n+1}$ as follows

$$\begin{aligned}\theta_{j,n+1} &= \theta_{j,n} - r_{j,n} \lambda_{j,n} \int_{\mathbf{Y}} \frac{\varphi_{j,n}^{(\alpha)}(y)}{\alpha - 1} \frac{\partial \log k(\theta, y)}{\partial \theta} \Big|_{(\theta, y) = (\theta_{j,n}, y)} \nu(dy) , \\ \theta_{j,n+1} &= \theta_{j,n} - r_{j,n} \frac{\lambda_{j,n}}{\sum_{j=1}^J \lambda_{j,n} \int_{\mathbf{Y}} \varphi_{j,n}^{(\alpha)}(y) \nu(dy)} \int_{\mathbf{Y}} \frac{\varphi_{j,n}^{(\alpha)}(y)}{\alpha - 1} \frac{\partial \log k(\theta, y)}{\partial \theta} \Big|_{(\theta, y) = (\theta_{j,n}, y)} \nu(dy)\end{aligned}$$

where $r_{j,n} > 0$ is the learning rate (we refer to Appendix B.4 for details regarding these updates). There is thus a similarity between (27) and the Gradient Descent updates above. To fully comprehend the connection between our gradient-based approach and Gradient Descent steps, we present below the generalisation of Example 2 to Gaussian Mixture Models. The smoothness assumption on $g_{j,n}$ will be satisfied in that example and we will recover a Gradient Descent scheme for Rényi's α -divergence minimisation as a special case.

Example 4 (Gradient-based approach for Gaussian Mixture Models) *We consider the case of d -dimensional Gaussian mixture densities with $k(\theta_j, y) = \mathcal{N}(y; m_j, \Sigma_j)$, where this time $\theta_j = m_j$. Further assume that $\mathbf{T} \subseteq \mathbb{R}^d$ is a convex subset and that $\Sigma_j = \sigma_j^2 \mathbf{I}_d$, where $\sigma_j^2 > 0$ and \mathbf{I}_d denotes the d -dimensional identity matrix. Finally, we assume that the non-negative function p defined on \mathbf{Y} satisfies (13).*

Then, (A1) holds. Setting $\beta_{j,n} = \sigma_j^{-2} (1 - \alpha)^{-1} \int_{\mathbf{Y}} \varphi_{j,n}^{(\alpha)}(y) \nu(dy)$ and denoting $\theta_{j,n} = m_{j,n}$, the function $g_{j,n}$ is $\beta_{j,n}$ -smooth for all $j = 1 \dots J$ and all $n \geq 1$ so that the means of the Gaussian densities $(\mathcal{N}(y; m_j, \Sigma_j))_{1 \leq j \leq J}$ can be optimised as follows: for all $n \geq 1$,

$$m_{j,n+1} = \gamma_{j,n} \hat{m}_{j,n} + (1 - \gamma_{j,n}) m_{j,n} , \quad j = 1 \dots J ,$$

where $\check{\varphi}_{j,n}^{(\alpha)}$ is defined in (16), $\gamma_{j,n} \in [0, 1]$ and $\hat{m}_{j,n} = \int_{\mathbf{Y}} y \check{\varphi}_{j,n}^{(\alpha)}(y) \nu(dy)$ (we refer to Corollary 8 for details). In particular, letting $(\gamma'_{j,n})_{n \geq 1}$ be valued in $(0, 1]$ for all $j = 1 \dots J$,

$$\gamma_{j,n} = \gamma'_{j,n} \frac{\lambda_{j,n} \int_{\mathbf{Y}} \varphi_{j,n}^{(\alpha)}(y) \nu(dy)}{\int_{\mathbf{Y}} \mu_n k(y)^{\alpha} p(y)^{1-\alpha} \nu(dy)} \in (0, 1], \quad j = 1 \dots J, \quad n \geq 1,$$

(since $\int_{\mathbf{Y}} \mu_n k(y)^{\alpha} p(y)^{1-\alpha} \nu(dy) = \sum_{j=1}^J \lambda_{j,n} \int_{\mathbf{Y}} \varphi_{j,n}^{(\alpha)}(y) \nu(dy)$). The resulting iterative algorithm is given by the following update at time $n \geq 1$:

$$m_{j,n+1} = m_{j,n} + \gamma'_{j,n} \frac{\int_{\mathbf{Y}} \lambda_{j,n} \varphi_{j,n}^{(\alpha)}(y) (y - m_{j,n}) \nu(dy)}{\int_{\mathbf{Y}} \mu_n k(y)^{\alpha} p(y)^{1-\alpha} \nu(dy)} , \quad j = 1 \dots J . \quad (28)$$

Interestingly, (28) can be seen as a Gradient Descent step w.r.t. $\theta_{j,n}$ for Rényi's α -divergence minimisation with a learning rate $r_{j,n} = \sigma_j^2 (1 - \alpha) \gamma'_{j,n}$. Hence, if we were to solely rely on the Gradient Descent literature, the convergence of the iterative sequence $(\theta_{j,n})_{n \geq 1}$ defined by (28) would require the sequence $(\lambda_n)_{n \geq 1}$ to be constant. This is in contrast with Corollary 5, which allows for a simultaneous optimisation of λ and Θ according to (21) and (28).

We now add on the comments made in Section 3 for the gradient-based methodology (which we built in Corollary 2 and have since extended to mixture models in Corollary 5):

- A core insight from Corollary 5, which is exemplified in Example 4, is that under a smoothness assumption on $g_{j,n}$ our mixture weights iterative updates are compatible with gradient-based updates, themselves linked to the Gradient Descent literature. In other words, we have embedded Gradient Descent-based iterative updates, which only act on Θ , within a larger framework where simultaneous updates for the mixture components parameters and the mixture weights are well-supported theoretically.
- Putting things into perspective with the maximisation approach once again, notice that our previous conclusions from Section 3 still hold. Namely: (i) the updates on the means from Examples 3 and 4 coincide, meaning that contrary to the gradient-based approach, the maximisation approach enables covariance matrices updates on top of means updates (ii) the maximisation approach permits us to bypass the smoothness assumption made in Corollary 5. As a whole, these properties make the maximisation approach a compelling alternative to the gradient-based approach.

We presented two approaches to construct iterative schemes $(\lambda_n, \Theta_n)_{n \geq 1}$ that ensure a monotonic decrease in Ψ_α at each step and lead to simple updates formulas for Gaussian Mixture Models. We now describe how our framework can be linked to the existing literature.

5. Related Work

In this section, we detail how our work relates to and improve on previous algorithms proposed for α -Rényi's α -divergence minimisation.

5.1 Rényi Divergence Variational Inference (Li and Turner, 2016)

In Li and Turner (2016), they seek to maximise the Variational Rényi (VR) Bound via (Stochastic) Gradient Ascent. This objective function is derived from Rényi's α -divergence and is given by: for all variational density $q \in \mathcal{Q}$ and all $\alpha \in \mathbb{R} \setminus \{1\}$,

$$\mathcal{L}_\alpha(q, \mathcal{D}) := \frac{1}{1-\alpha} \log \left(\int_{\mathcal{Y}} q(y)^\alpha p(y, \mathcal{D})^{1-\alpha} \nu(dy) \right). \quad (29)$$

Contrary to us, the work from Li and Turner (2016) does not consider the case where q belongs to the mixture models variational family (15), that is $q = \mu_{\lambda, \Theta}^k$.

Yet, a parallel can be drawn in the GMM case between Li and Turner (2016) and our approach by observing that the gradient-based updates on the means in Example 4 each coincide with a Gradient Ascent step on the objective function (29) for a well-chosen learning rate (since the gradient of \mathcal{L}_α is proportional to the gradient of Rényi's α -divergence with a factor $-\alpha^{-1}$, this follows from the remarks made in Section 4.2.2 regarding Gradient Descent steps for Rényi's α -divergence minimisation in the GMM case).

Our work hence provides a theoretical framework which enables simultaneous optimisation of the mixture weights λ and of the mixture components parameters Θ . In addition, beyond the gradient-based updates, we propose the novel maximisation updates and we allow for covariance matrices optimisation. We also emphasise that our maximisation approach will apply to well-chosen kernels k beyond the Gaussian case, as we will detail in Section 6.

5.2 The Power Descent Algorithm (Daudel et al., 2021)

In order to identify the connection between our work and the Power Descent algorithm introduced in Daudel et al. (2021), let us briefly present the latter.

The Power Descent is a *gradient-based* algorithm that operates on probability measures and performs α -divergence minimisation for all $\alpha \in \mathbb{R} \setminus \{1\}$. More precisely, equipping \mathbb{T} with a σ -field \mathcal{T} and denoting by $M_1(\mathbb{T})$ the space of probability measures on $(\mathbb{T}, \mathcal{T})$, the Power Descent optimises $\Psi_\alpha(\mu k)$ with respect to $\mu \in M_1(\mathbb{T})$, where $\mu k(y) = \int_{\mathbb{T}} \mu(d\theta) k(\theta, y)$ for all $\mu \in M_1(\mathbb{T})$ and all $y \in Y$. Given an initial probability measure $\mu_1 \in M_1(\mathbb{T})$, it does so by performing several one-step transitions of the Power Descent algorithm:

$$\mu_{n+1} = \mathcal{I}_\alpha(\mu_n), \quad n \geq 1, \quad (30)$$

where, for all $\mu \in M_1(\mathbb{T})$ and all $\theta \in \mathbb{T}$,

$$b_{\mu, \alpha}(\theta) = \int_Y k(\theta, y) \frac{1}{\alpha - 1} \left[\left(\frac{\mu k(y)}{p(y)} \right)^{\alpha-1} - 1 \right] \nu(dy)$$

$$\mathcal{I}_\alpha(\mu)(d\theta) = \frac{\mu(d\theta) [(\alpha - 1)(b_{\mu, \alpha}(\theta) + \kappa) + 1]^{\frac{\eta}{1-\alpha}}}{\mu([(\alpha - 1)(b_{\mu, \alpha} + \kappa) + 1]^{\frac{\eta}{1-\alpha}})}.$$

Daudel et al. (2021) motivated the Power Descent algorithm by establishing a monotonicity result for this algorithm obtained as a particular case of (Daudel et al., 2021, Theorem 1). In their result, the monotonic decrease in Ψ_α of the scheme (30) holds for all $\mu \in M_1(\mathbb{T})$, all $\alpha \in \mathbb{R} \setminus \{1\}$, all $\eta \in (0, 1]$ and all κ such that $(\alpha - 1)\kappa \geq 0$. We provide below a more general version of their result, where the monotonic decrease in Ψ_α holds for well-chosen values of η that are larger than 1 when $\alpha < 0$ (we refer to Appendix B.5 for the proof of this result).

Proposition 2 *Assume that p and k are as in (A1). Let (α, η) belong to any of the following cases.*

- (i) $\alpha \leq -1$ and $\eta \in (0, (\alpha - 1)/\alpha]$;
- (ii) $\alpha \in (-1, 0)$ and $\eta \in (0, 1 - \alpha]$;
- (iii) $\alpha \in [0, 1)$ or $\alpha > 1$ and $\eta \in (0, 1]$.

Moreover, let $\mu \in M_1(\mathbb{T})$ be such that $\Psi_\alpha(\mu k) < \infty$ and let κ be such that $(\alpha - 1)\kappa \geq 0$. Then, the two following assertions hold.

- (i) We have $\Psi_\alpha(\mathcal{I}_\alpha(\mu)k) \leq \Psi_\alpha(\mu k)$.
- (ii) We have $\Psi_\alpha(\mathcal{I}_\alpha(\mu)k) = \Psi_\alpha(\mu k)$ if and only if $\mu = \mathcal{I}_\alpha(\mu)$.

Building on the monotonicity result provided by (Daudel et al., 2021, Theorem 1) for the Power Descent algorithm - that we generalised in Proposition 2 - Daudel et al. (2021) then applied this algorithm to mixture weights optimisation by letting the initial probability measure $\mu_1 \in M_1(\mathbb{T})$ be a weighted sum of Dirac measures of the form $\mu_1 = \sum_{j=1}^J \lambda_{j,n} \delta_{\theta_j}$, with

$\Theta \in \mathbb{T}$ and $\lambda_1 \in \mathcal{S}_J$. For that choice of μ_1, μ_n in (30) can be written as $\mu_n = \sum_{j=1}^J \lambda_{j,n} \delta_{\theta_j}$ at time n and the Power Descent amounts to performing the update

$$\lambda_{j,n+1} = \frac{\lambda_{j,n} [(\alpha - 1)(b_{\mu_n, \alpha}(\theta_j) + \kappa) + 1]^{\frac{\eta}{1-\alpha}}}{\sum_{\ell=1}^J \lambda_{\ell,n} [(\alpha - 1)(b_{\mu_n, \alpha}(\theta_\ell) + \kappa) + 1]^{\frac{\eta}{1-\alpha}}}, \quad j = 1 \dots J. \quad (31)$$

Interestingly, the update (31) corresponds to the update on the mixture weights (21) we have identified in Theorem 3 for $\Theta_n = \Theta$, $\eta_n = \eta/(1 - \alpha)$ and $\kappa_n = \kappa$. Steaming from this link between our approach and Daudel et al. (2021), we now make two important comments.

- *Benefits of our approach compared to Daudel et al. (2021).* The monotonicity result from Daudel et al. (2021) generalised in Proposition 2 requires the sequence $(\Theta_n)_{n \geq 1}$ to be constant when applied to mixture weights optimisation in (31). This restricts the variational family to mixture models with fixed mixture components parameters. To remedy this problem, Daudel et al. (2021) proposed a fully-adaptive algorithm that alternates between an Exploitation step optimising the mixture weights according to (31) and an Exploration step acting on the mixture components parameters. However, they established no theoretical guarantees for their complete Exploitation-Exploration algorithm, as the choice of the Exploration step remained mostly unexplored.

A strong improvement of our approach is then that we provide theoretically-sound updates for $(\lambda_n, \Theta_n)_{n \geq 1}$, with the particularity that our mixture weights updates relate to the framework of Daudel et al. (2021). In that sense, our approach supplements the work done in Daudel et al. (2021) (albeit by an entirely different proof technique). Furthermore, we do not need to alternate between mixture weights and mixture components parameters updates in our algorithms, as both can be carried out simultaneously (as done in Algorithms 1 and 2). In practice, this will permit us to reduce the computational cost (as the samples will be shared throughout the mixture weights and the mixture components parameters approximated updates, see Section 7).

- *A gradient-based mixture weights update.* Daudel et al. (2021) establishes that the Power Descent belongs to a family of gradient-based algorithms which includes the Entropic Mirror Descent algorithm, a typical optimisation algorithm for optimisation under simplex constraints, as a special case. Viewed from this angle, the parameter η in (31) can be understood as a learning rate with $b_{\mu_n, \alpha}$ playing the role of the gradient of Ψ_α . Connecting the Power Descent to our framework thus sheds light on the gradient-based nature of the mixture weights update (21) from Theorem 3 and provides a better understanding of the role of the parameter η_n appearing in this update. This aspect will be helpful to interpret our numerical experiments in Section 8.

5.3 The M-PMC Algorithm (Cappé et al., 2008)

For any measurable positive function p on (Y, \mathcal{Y}) , the M-PMC algorithm (Cappé et al., 2008) aims at solving the optimisation problem

$$\sup_{(\lambda \in \mathcal{S}_J, \Theta \in \mathbb{T}^J)} \int_Y \log \left(\sum_{j=1}^J \lambda_j k(\theta_j, y) \right) p(y) \nu(dy),$$

or equivalently, at minimising the inclusive Kullback-Leibler divergence $D_0(\mu_{\lambda, \Theta} K || \mathbb{P})$ w.r.t. to (λ, Θ) , where $\mathbb{P}(A) = \int_A p \, d\nu / \int_Y p \, d\nu$ for all $A \in \mathcal{Y}$. This is done in (Cappé et al., 2008, Section 2) by introducing the following iterative updates: for all $j = 1 \dots J$ and all $n \geq 1$,

$$\begin{aligned} \lambda_{j,n+1} &= \int_Y \frac{\lambda_{j,n} k(\theta_{j,n}, y)}{\sum_{\ell=1}^J \lambda_{\ell,n} k(\theta_{\ell,n}, y)} \frac{p(y)}{\int_Y p(y) \nu(dy)} \nu(dy) \\ \theta_{j,n+1} &= \operatorname{argmax}_{\theta_j \in \mathcal{T}} \int_Y \frac{\lambda_{j,n} k(\theta_{j,n}, y)}{\sum_{\ell=1}^J \lambda_{\ell,n} k(\theta_{\ell,n}, y)} \log(k(\theta_j, y)) p(y) \nu(dy) . \end{aligned}$$

Cappé et al. (2008) motivated the two updates above by noticing that they can be seen as integrated versions under the target distribution of the update formulas for the Expectation-Maximisation (EM) algorithm applied to the mixture density parameter estimation problem

$$\sup_{(\lambda \in \mathcal{S}_J, \Theta \in \mathcal{T}^J)} \sum_{m=1}^M \log \left(\sum_{j=1}^J \lambda_j k(\theta_j, Y_m) \right) ,$$

meaning that these updates ensure a systematic increase in the integrated likelihood at each step. As these updates also correspond to the case $\alpha = 0$, $\eta_n = 1$, $\kappa_n = 0$, $a_{j,n} = 1$ and $b_{j,n} = 0$ in Algorithm 1, the M-PMC algorithm is in fact included in our framework and we can interpret our theoretical results in light of this algorithm.

More precisely, we have generalised an integrated EM algorithm by preserving its monotonicity property for a wide range of hyperparameters. A particularly striking fact is that the monotonicity property holds for $\alpha \in [0, 1)$, hence updates akin to an EM procedure can be derived past the traditional case of likelihood optimisation. As we shall see, the additional layers of flexibility obtained in Algorithm 1 will also have important practical consequences due to the underlying gradient-based structure behind the mixture weights (and the mixture components parameters updates in the GMMs case) we have uncovered.

Table 1 summarises the main improvements of our framework compared to the existing literature and in the coming section, we revisit the maximisation and the gradient-based approaches when the variational family is based on the exponential family.

6. Exponential Family Distributions: a Closer Look

In this section, we state generic results for the maximisation and the gradient-based approaches in the important case where the variational family is based on the exponential family. Those results will in particular enable us to show how the update formulas in Examples 3 and 4 are derived, before allowing us to incorporate novel variational families in our framework such as multivariate Student's t densities. To do so, let us first review the exponential family and some of its well-known properties.

6.1 Notation and Useful Definitions

We start with the definition of an exponential family distribution.

| | Improvements of our framework |
|--|--|
| Rényi Gradient Descent Li and Turner (2016) | Simultaneous optimisation w.r.t. $(\boldsymbol{\lambda}, \Theta)$ (prev. mixture weights $\boldsymbol{\lambda}$ optimisation not considered) For GMMs : maximisation approach encompasses Rényi Gradient Descent and provides covariance matrices updates |
| Power Descent Daudel et al. (2021) | Simultaneous optimisation w.r.t. $(\boldsymbol{\lambda}, \Theta)$ (prev. $(\Theta_n)_{n \geq 1}$ constant) |
| M-PMC algorithm Cappé et al. (2008) | Extension of an Integrated EM algorithm to: $\alpha \in [0, 1)$, $\eta_n \in (0, 1]$, $(\alpha - 1)\kappa_n \geq 0$ and $b_{j,n} \geq 0$ (prev. $\alpha = 0$, $\eta_n = 1$, $\kappa_n = 0$ and $b_{j,n} = 0$) |

Table 1: Key improvements of our framework for mixture models optimisation compared to related works (prev. stands for previously in the literature).

Definition 1 (Exponential family) *Let ν be a σ -finite measure on (Y, \mathcal{Y}) , E be a Euclidean space endowed with an inner product $\langle \cdot, \cdot \rangle_E$ and its corresponding norm $\|\cdot\|_E$, h be a non-negative function defined on Y and $S : Y \rightarrow E$ be a measurable function. In its canonical form, a member of the exponential family is defined by the parametric probability density*

$$k^{(o)}(\zeta, y) = h(y) \exp(\langle \zeta, S(y) \rangle_E - A(\zeta)) \quad , \quad y \in Y \quad , \quad \zeta \in E_0$$

where

$$E_0 := \left\{ \zeta \in E : \int_Y h(y) e^{\langle \zeta, S(y) \rangle_E} \nu(dy) < \infty \right\}$$

and where the values taken by A on the subset E_0 are obtained from the normalising constraint $\int k^{(o)}(\zeta, \cdot) d\nu = 1$. In this setting, ν is called the dominating measure and S the natural statistic. In what follows, $\text{Int}E_0$ denotes the interior of E_0 .

Recall now that S defines an exhaustive statistic for this model, that E_0 is a convex subset of E and that A is infinitely differentiable in $\text{Int}E_0$. Moreover, for all $\zeta \in \text{Int}E_0$, the expectation and covariance operator of S under the density $k^{(o)}(\zeta, \cdot)$ are the gradient ∇A and the Hessian $\nabla \nabla^T A$ of A taken at ζ , respectively. For later reference, we also recall

$$\nabla A(\zeta) = \int_Y S(y) k^{(o)}(\zeta, y) \nu(dy) \tag{32}$$

and we introduce the following assumption.

(B1) The kernel density $k^{(o)}$ defined on $E_0 \times Y$ is a member of the exponential family as in Definition 1. Moreover, we assume that:

- (a) The function h on Y is positive.
- (b) There is no affine hyperplane of E to which $S(y)$ belongs for ν -almost all $y \in Y$.

Notice that the assumption (B1)(a) can be circumvented by changing the dominating measure ν and therefore the assumptions (B1)(a) and (B1)(b) should be considered together. These

assumptions notably imply that the covariance operator of S under the density $k^{(o)}(\zeta, \cdot)$ is positive definite and thus so is $\nabla \nabla^T A(\zeta)$ for all $\zeta \in \text{Int} E_0$. In addition, one often relies on a non-canonical parameterisation (see for example Example 1), leading to the kernel density

$$k(\theta, y) = k^{(o)}(v(\theta), y) = h(y) \exp(\langle v(\theta), S(y) \rangle_E - A \circ v(\theta)) , \quad \theta \in \mathbb{T} , y \in \mathbb{Y} \quad (33)$$

where v maps the parameter space \mathbb{T} to a subset of $\text{Int} E_0$. In that case, the assumption (B1)(a) will ensure that the assumption made on k in (A1) holds.

6.2 The Maximisation Approach for the Exponential Family

Recall that the maximisation approach proposed in Corollary 1 aims at solving the optimisation problem (7) given by

$$\theta_{n+1} = \operatorname{argmax}_{\theta \in \mathbb{T}} \int_{\mathbb{Y}} \left[\varphi_n^{(\alpha)}(y) + b_n k(\theta_n, y) \right] \log \left(\frac{k(\theta, y)}{k(\theta_n, y)} \right) \nu(dy) , \quad n \geq 1 .$$

Furthermore, in the more general case of mixture models described in Corollary 4, the maximisation approach leads to the argmax problem (25). This second argmax problem is similar to (7) in the sense that it updates $\Theta_n = (\theta_{j,n})_{1 \leq j \leq J}$ by solving J component-wise argmax problems of the same form as (7) (with θ_{n+1} , $\varphi_n^{(\alpha)}$, b_n and θ_n being replaced by $\theta_{j,n+1}$, $\varphi_{j,n}^{(\alpha)}$, $b_{j,n}$ and $\theta_{j,n}$, for $j = 1 \dots J$):

$$\theta_{j,n+1} = \operatorname{argmax}_{\theta \in \mathbb{T}} \int_{\mathbb{Y}} \left[\varphi_{j,n}^{(\alpha)}(y) + b_{j,n} k(\theta_{j,n}, y) \right] \log \left(\frac{k(\theta, y)}{k(\theta_{j,n}, y)} \right) \nu(dy) , \quad j = 1 \dots J , n \geq 1 .$$

Our goal is now to solve the argmax problems above for a member of the exponential family. To do so, we first state a theorem in which the kernel k is in its canonical form, that is v is the identity mapping in (33) and $k(\theta, \cdot) = k^{(o)}(\zeta, \cdot)$. Other parameterisations will then be built using this theorem and as a result we will deduce corollaries that include non-canonical parameterisations and that are applicable to the general case of mixture models. All the proofs of the results stated in Section 6.2 are deferred to Appendix C.2.

Theorem 4 *Let $k^{(o)}$ satisfy (B1) and \mathcal{O} be an open subset of E_0 . Let $\check{\varphi}$ be a probability density function with respect to the measure ν such that*

$$\int \|S\|_E \check{\varphi} \, d\nu < \infty . \quad (34)$$

Let $\zeta_0 \in \mathcal{O}$ and $b \geq 0$. Then the argmax problem

$$\operatorname{argmax}_{\zeta \in \mathcal{O}} \int_{\mathbb{Y}} \left[\check{\varphi}(y) + b k^{(o)}(\zeta_0, y) \right] \log \left(\frac{k^{(o)}(\zeta, y)}{k^{(o)}(\zeta_0, y)} \right) \nu(dy) \quad (35)$$

admits at least a solution if and only if

$$\mathbf{s}^* := \frac{1}{1+b} \int S \check{\varphi} \, d\nu + \frac{b}{1+b} \nabla A(\zeta_0) \quad (36)$$

belongs to the image set $\nabla A(\mathcal{O})$, in which case (35) has a unique solution ζ^ defined by*

$$\nabla A(\zeta^*) = \mathbf{s}^* . \quad (37)$$

Theorem 4 shows the equivalence between the argmax problem (35) and the equation (37) under the canonical parameterisation. From there, using the expression of ∇A given in (32) and considering the non-canonical exponential family probability density (33), we can interpret the equation (37) as an equality between two means of the statistic S computed under two different distributions. Setting $\zeta^* = v(\theta^*)$ and $\zeta_0 = v(\theta_0)$, we indeed obtain

$$\int_{\mathbf{Y}} S(y) k(\theta^*, y) \nu(dy) = \int_{\mathbf{Y}} S(y) \psi(y) \nu(dy) \quad (38)$$

where ψ is the mixture

$$\psi(y) = \frac{1}{1+b} \check{\varphi}(y) + \frac{b}{1+b} k(\theta_0, y), \quad y \in \mathbf{Y}.$$

As exemplified in the following corollary, an adequate parameterisation $\zeta = v(\theta)$ may then lead to a simple solution of (38), which in turn provides a solution to the argmax problem

$$\operatorname{argmax}_{\theta \in \mathbf{T}} \int_{\mathbf{Y}} [\check{\varphi}(y) + b k(\theta_0, y)] \log \left(\frac{k(\theta, y)}{k(\theta_0, y)} \right) \nu(dy). \quad (39)$$

Corollary 6 (Gaussian density) *Let $d \geq 1$ and $\mathbf{Y} = \mathbb{R}^d$. Let $\mathcal{M}_{>0}(d)$ denote the set of symmetric positive definite $d \times d$ matrices. We consider the case of a d -dimensional Gaussian density with $k(\theta, y) = \mathcal{N}(y; m, \Sigma)$, where $\theta = (m, \Sigma) \in \mathbf{T} = \mathbb{R}^d \times \mathcal{M}_{>0}(d)$. Let $\check{\varphi}$ be a probability density function on \mathbf{Y} such that (34) holds, that is in the Gaussian case*

$$\int_{\mathbf{Y}} \|y\|^2 \check{\varphi}(y) dy < \infty.$$

Let $\theta_0 = (m_0, \Sigma_0) \in \mathbf{T}$ and $b \geq 0$. If $b = 0$, assume moreover that $\check{\varphi}$ is non-degenerate in the sense that its covariance matrix is positive definite. Then (39) has a unique solution $\theta^ = (m^*, \Sigma^*) \in \mathbf{T}$ defined by $m^* = \mathbb{E}[Y]$ and $\Sigma^* = \operatorname{Cov}(Y)$, where Y is a random variable valued in \mathbb{R}^d with density*

$$\psi(y) = \frac{1}{1+b} \check{\varphi}(y) + \frac{b}{1+b} \mathcal{N}(y; m_0, \Sigma_0), \quad y \in \mathbf{Y}.$$

It now remains to choose $\check{\varphi}$ adequately to relate (39) to (7), or rather to its generalisation (25). That way, we will in particular get that the update formulas given in Examples 1 and 3 are straightforward consequences of Corollary 6. This is the purpose of Corollary 7 below, in which we also rewrite the assumption (34) made on $\check{\varphi}$ under a more convenient form.

Corollary 7 *Let $k^{(o)}$ satisfy (B1). Further assume that the kernel k is of the form (33) with a one-to-one mapping v defined on \mathbf{T} such that its image $v(\mathbf{T})$ is an open subset of E_0 . Let $\alpha \in [0, 1)$ and let $p : \mathbf{Y} \rightarrow \mathbb{R}_+$ be such that*

$$0 < \int_{\mathbf{Y}} \left(1 + \|S(y)\|_E^{1/(1-\alpha)} \right) p(y) \nu(dy) < \infty \quad (40)$$

holds. Let $J \in \mathbb{N}$, $n \geq 1$ and $(\lambda_n, \Theta_n) \in \mathcal{S}_J^+ \times \mathbf{T}^J$. Then, for all $j = 1 \dots J$, (34) holds with $\check{\varphi} = \check{\varphi}_{j,n}^{(\alpha)}$ as defined in (16) and setting $\gamma_{j,n} = \int \varphi_{j,n}^{(\alpha)} d\nu / (\int \varphi_{j,n}^{(\alpha)} d\nu + b_{j,n})$, we have that

$$\operatorname{argmax}_{\theta \in \mathbf{T}} \int_{\mathbf{Y}} \left[\varphi_{j,n}^{(\alpha)}(y) + b_{j,n} k(\theta_{j,n}, y) \right] \log \left(\frac{k(\theta, y)}{k(\theta_{j,n}, y)} \right) \nu(dy) \quad (41)$$

admits at least a solution if and only if

$$\mathbf{s}_j^* := \gamma_{j,n} \int S \check{\varphi}_{j,n}^{(\alpha)} d\nu + (1 - \gamma_{j,n}) \nabla A \circ v(\theta_{j,n}) \quad (42)$$

belongs to the image set $\nabla A \circ v(\mathbb{T})$, in which case (41) has a unique solution θ_j^* defined by

$$\nabla A \circ v(\theta_j^*) = \mathbf{s}_j^*$$

and we can set $\theta_{j,n+1} = \theta_j^*$ for all $j = 1 \dots J$ in (25) of Corollary 4.

Since the argmax problem (41) is immune to change of variables thanks to its argmax form, the maximisation approach can be solved using the canonical parameter ζ in Theorem 4 and transported to the parameter θ via the one-to-one mapping v as written in Corollary 7.

The updates from Examples 1 and 3 then follow from Corollary 7, paired up with Corollary 6 and with the fact that (40) simplifies to (8) (since in the Gaussian case we have set $S(y) = (y, -yy^T/2)$ for all $y \in \mathbb{Y}$). More generally, we can always solve the argmax problem (41), that is the argmax problem of Corollary 4 with k belonging to the exponential family, for $b_{j,n} > 0$ large enough.¹

We now move on to the study of the gradient-based approach for the exponential family.

6.3 The Gradient-based Approach for an Exponential Family Distribution

Remember that to construct the sequence $(\theta_n)_{n \geq 1}$ according to the gradient-based approach from Corollary 2 we need to compute the gradient of the function g_n defined in (11) by

$$g_n(\theta) = \int_{\mathbb{Y}} \frac{\varphi_n^{(\alpha)}(y)}{\alpha - 1} \log \left(\frac{k(\theta, y)}{k(\theta_n, y)} \right) \nu(dy), \quad \theta \in \mathbb{T},$$

and verify that g_n satisfies a β_n -smoothness condition, which leads to updates of the form

$$\theta_{n+1} = \theta_n - \frac{\gamma_n}{\beta_n} \nabla g_n(\theta_n), \quad n \geq 1$$

with $\gamma_n \in (0, 1]$. Furthermore, the extension to mixture models in Corollary 5 involves the gradient of J functions $g_{j,n}$ that resemble g_n ($g_n, \varphi_n^{(\alpha)}$ and θ_n are replaced by $g_{j,n}, \varphi_{j,n}^{(\alpha)}$ and $\theta_{j,n}$ in the definition of g_n above) and where each function $g_{j,n}$ is assumed to be $\beta_{j,n}$ -smooth.

Let us now solve the gradient-based approach for a member of the exponential family. Like in the previous section, we start with a result handling the case where the kernel k is in its canonical form (that is, v is the identity mapping in (33) and $k(\theta, \cdot) = k^{(o)}(\zeta, \cdot)$). All the proofs of the results stated in Section 6.3 are deferred to Appendix C.3.

Proposition 3 *Let $k^{(o)}$ satisfy (B1). Let $\check{\varphi}$ be a probability density function with respect to ν satisfying (34). Let $\zeta_0 \in \text{Int} E_0$ and set*

$$g^{(o)}(\zeta) = - \int_{\mathbb{Y}} \check{\varphi}(y) \log \left(\frac{k^{(o)}(\zeta, y)}{k^{(o)}(\zeta_0, y)} \right) \nu(dy), \quad \zeta \in \text{Int} E_0. \quad (43)$$

1. Under the assumptions of Corollary 7, $v(\mathbb{T})$ is an open subset of E_0 and ∇A is a \mathcal{C}^∞ -diffeomorphism on $\text{Int} E_0$. Hence, $\nabla A \circ v(\mathbb{T})$ is an open subset of E . Since \mathbf{s}_j^* in (42) tends to $\nabla A \circ v(\theta_{j,n})$ as $b_{j,n} \rightarrow \infty$, \mathbf{s}_j^* ends up belonging to $\nabla A \circ v(\mathbb{T})$ for $b_{j,n}$ large enough. The desired result follows by using that the regularisation parameter $b_{j,n} > 0$ in Corollary 4 can be freely chosen at each step $n \geq 1$.

Then, for any convex subset $C_0 \subseteq \text{Int}E_0$, the function $g^{(o)}$ is β_0 -smooth over C_0 with

$$\nabla g^{(o)}(\zeta) = \nabla A(\zeta) - \int S \check{\varphi} \, d\nu \quad \text{and} \quad \beta_0 \geq \sup_{\zeta \in C_0} \left\| \nabla \nabla^T A(\zeta) \right\|_{\text{op}}, \quad (44)$$

where $\|\cdot\|_{\text{op}}$ denotes the operator norm.

Letting C_0 be the line segment $\{\zeta_0 - \frac{t}{\beta_0} (\nabla A(\zeta_0) - \int S \check{\varphi} \, d\nu) : t \in [0, 1]\}$ in Proposition 3, we have that C_0 stays in $\text{Int}E_0$ for β_0 large enough. Since the Hessian of A is continuous on $\text{Int}E_0$, the inequality in (44) can also be satisfied. Hence, the following gradient step

$$\zeta = \zeta_0 - \frac{\gamma}{\beta_0} \left(\nabla A(\zeta_0) - \int_Y S \check{\varphi} \, d\nu \right) \quad (45)$$

can always be performed for all $\gamma \in (0, 1]$ and we can always apply the gradient-based approach when the kernel k is in its canonical form. When considering non-canonical exponential family probability densities of the form (33) and under some reasonable assumptions on the mapping v (for instance, if v is a \mathcal{C}^2 diffeomorphism), one is left with choosing between two alternative ways of exploiting Proposition 3:

(i) Apply a gradient step on the canonical parameter ζ using the objective $g^{(o)}$ in (43) and translate it into an update of the parameter θ via a change of variable. In this case, we set $\zeta_0 = v(\theta_0)$, we apply the step (45) and we then perform the change of variable $\theta = v^{-1}(\zeta)$ to get back to a parameter in T . Overall, this leads to an update of the form

$$\theta = v^{-1} \left(v(\theta_0) - \frac{\gamma}{\beta_0} \left(\nabla A \circ v(\theta_0) - \int_Y S \check{\varphi} \, d\nu \right) \right).$$

(ii) Apply a gradient step directly on the parameter θ using the objective

$$g(\theta) = - \int \check{\varphi}(y) \log \left(\frac{k(\theta, y)}{k(\theta_0, y)} \right) dy, \quad \theta \in \mathsf{T}.$$

In that case, since $g = g^{(o)} \circ v$ where $g^{(o)}$ is again as in (43) with $\zeta_0 = v(\theta_0)$, this leads to an update of the form

$$\theta = \theta_0 - \frac{\gamma}{\beta} \nabla v(\theta_0) \cdot \left(\nabla A \circ v(\theta_0) - \int_Y S \check{\varphi} \, d\nu \right).$$

In both approaches, $\gamma \in (0, 1]$ and the smoothness indices β_0 and β have to be taken large enough in order to guarantee a decrease of the objective function. In practice, it is not clear which parameterisation leads to the simpler or more efficient algorithm. We investigate a particular setting in the following result.

Corollary 8 (Gaussian density with known covariance matrix) *Let $d \geq 1$ and $\mathsf{Y} = \mathbb{R}^d$. We consider the exponential family of a d -dimensional Gaussian density with $k(\theta, y) = \mathcal{N}(y; \theta, \Sigma)$, where $\theta \in \mathsf{T} := \mathbb{R}^d$ and Σ is a (known) covariance matrix in $\mathcal{M}_{>0}(d)$. Let $\check{\varphi}$ be a probability density function w.r.t. ν such that (34) holds, that is in the Gaussian case*

$$\int \|y\| \check{\varphi}(y) \, dy < \infty.$$

Set $v(\theta) = \Sigma^{-1}\theta$ and define the canonical kernel $k^{(o)}$ with $h(y) = (2\pi)^{-d/2} |\Sigma|^{-1/2} e^{-y^T \Sigma^{-1} y/2}$, $S(y) = y$ and $A(\zeta) = \zeta^T \Sigma \zeta / 2$. Then (33) holds and the methods (i) and (ii) lead to the the following gradient-based updates, respectively,

$$\theta = \theta_0 - \frac{\gamma}{\beta_0} \Sigma \left(\theta_0 - \int y \check{\varphi}(y) \, dy \right) \quad (46)$$

$$\theta = \theta_0 - \frac{\gamma}{\beta} \Sigma^{-1} \left(\theta_0 - \int y \check{\varphi}(y) \, dy \right) \quad (47)$$

where β_0 is the largest eigenvalue of Σ and β is the largest eigenvalue of Σ^{-1} . They correspond to the smoothness indices of $g^{(o)}$ and g over \mathbb{R}^d , respectively.

Corollary 8 illustrates how, even in a simple framework, the gradient-based updates strongly depend on the parameterisation chosen. They in fact coincide only if Σ is scalar in the above corollary, in which case Σ/β_0 and Σ^{-1}/β above are both equal to the identity matrix.

Following the reasoning of Section 6.2, we can then apply Corollary 8 with $\check{\varphi} = \check{\varphi}_n^{(\alpha)}$ (and more generally with $\check{\varphi} = \check{\varphi}_{j,n}^{(\alpha)}$). As (34) is implied by (13) (using Lemma 8 in Appendix C.2 and that here $S(y) = y$), this enables us to deduce the updates in Examples 2 and 4. Note that we did not introduce a convex subset C_0 in Corollary 8 (this is due to the fact that the smoothness index is constant over the whole parameter space for both gradient-based updates in that case), which is why C_0 does not appear in Examples 2 and 4.

Let us next put into perspective Corollary 6, in which we considered a Gaussian density with varying mean vector and covariance matrix, and Corollary 8, where the covariance matrix is assumed to be known. While Corollary 8 kept the computations straightforward and provided gradient-based updates that do not require to introduce a convex subset C_0 , this is no longer the case if we take the same exact setting as in Corollary 6. Indeed, when $\theta = (m, \Sigma) \in \mathbb{T} = \mathbb{R}^d \times \mathcal{M}_{>0}(d)$ and under the condition (34), the gradient of g is given by

$$\nabla g(\theta) = \left(\Sigma^{-1} \left(m - \int y \check{\varphi}(y) \, dy \right), \frac{1}{2} \Sigma^{-1} \left(\int y y^T \check{\varphi}(y) \, dy - \Sigma + m m^T \right) \Sigma^{-1} \right).$$

The smoothness index β is now no longer constant here, which makes it much more involved to choose a convenient convex subset C_0 and to bound the smoothness index over it.

At this stage, we can solve the maximisation and gradient-based approaches when the kernel k belongs to the exponential family for variational families as large as (finite) mixture models. The maximisation approach appears to be preferable to the gradient-based one as it does not require a smoothness assumption and does not depend on the parameterisation.

In the next section, we investigate how our maximisation approach may further generalise to include extensions of the exponential family distribution such as Student's t distributions, hence further motivating the maximisation approach over the gradient-based approach.

6.4 Extension to Linear Mixture (LM) Models

The goal of this section is to examine what becomes of the maximisation approach for a specific extension of the exponential family distribution, which we will refer to as *exponential linear mixture* (ELM) family. The ELM family will in particular encompass the Student's

t distribution with mean m , covariance matrix Σ and a degrees of freedom which can be defined as the continuous mixture

$$t(y; m, \Sigma, a) = \int_0^\infty \mathcal{N}(y; m, z^{-1}\Sigma) \check{\tau}_a(dz) ,$$

where $\check{\tau}_a$ denotes the χ^2 distribution with a degrees of freedom. As a result, building a maximisation approach within the ELM family will lead to new iterative schemes for variational families that go beyond the cases considered thus far. To this end, let us first provide a precise definition of the ELM family we want to study and introduce the main assumptions of this section. All the proofs from Section 6.4 are deferred to Appendix C.4.

6.4.1 THE (E)LM FAMILY: DEFINITION AND NOTATION

Let F be an Euclidean space endowed with the inner product $\langle \cdot, \cdot \rangle_F$ and norm $\|\cdot\|_F$. We denote by $\mathcal{L}(F, E)$ and $\mathcal{L}(F)$ the (finite-dimensional) linear spaces of linear operators from F to E and from F to itself, respectively. We now provide the definition of the (E)LM family.

Definition 2 ((E)LM family) *Let $k^{(o)} : E_0 \times \mathsf{Y} \rightarrow \mathbb{R}_+$ be a kernel density w.r.t. the dominating measure ν , with $E_0 \subseteq E$ and let $\Upsilon : \tilde{\mathsf{T}} \rightarrow F$. Set $\mathsf{Z} = \mathcal{L}(F, E)$ and denote by \mathcal{Z} its Borel σ -field. Let M be a subset of probability measures on $(\mathsf{Z}, \mathcal{Z})$ satisfying*

(i) *for all τ, τ' in M , the distributions τ and τ' are equivalent;*

(ii) *for all $\tau \in \mathsf{M}$ and τ -almost all $\ell \in \mathsf{Z}$, the image set $\ell \circ \Upsilon(\tilde{\mathsf{T}})$ is included in E_0 .*

A member of the linear mixture (LM) family with canonical kernel $k^{(o)}$, mixing class M and parameter mapping Υ , is defined for a parameter $\theta = (\vartheta, \tau) \in \mathsf{T} = \tilde{\mathsf{T}} \times \mathsf{M}$ by the density w.r.t. ν

$$k^{(1)}(\theta, y) = \int_{\mathsf{Z}} k^{(o)}(\ell \circ \Upsilon(\vartheta), y) \tau(d\ell) , \quad y \in \mathsf{Y} . \quad (48)$$

If moreover $k^{(o)}$ satisfies (B1), we say that $k^{(1)}(\theta, \cdot)$ is a member of the exponential linear mixture (ELM) family.

Here, Assumption (i) implies that we only need to check Assumption (ii) for one τ in M . As for Assumption (ii), it ensures that for all $\vartheta \in \tilde{\mathsf{T}}$ and all $\tau \in \mathsf{M}$, the function $(y, \ell) \mapsto k^{(o)}(\ell \circ \Upsilon(\vartheta), y)$ appearing in (48) is a probability density function w.r.t. $\nu \otimes \tau$, so that for all $\theta \in \mathsf{T}$, the function $y \mapsto k^{(1)}(\theta, y)$ is a probability density function w.r.t. ν .

Furthermore, the mapping Υ in Definition 2 allows us to introduce various parameterisations. It is also important to note that for $\theta = (\vartheta, \tau) \in \mathsf{T}$, the probability density function $k^{(1)}(\theta, \cdot)$ only depends on $\xi = \Upsilon(\vartheta) \in \Upsilon(\tilde{\mathsf{T}})$ and $\tau \in \mathsf{M}$. As we shall see later, Student's t distributions in particular will fit this general setting. Before that, let us state conditions leading to a systematic decrease in Ψ_α when the kernel $k^{(1)}$ is as in Definition 2 and from there, let us see how a maximisation approach for the ELM family can be derived.

6.4.2 MONOTONIC DECREASE CONDITIONS FOR THE (E)LM FAMILY

In the case of an ELM family, we are not able to directly solve the argmax problem (25) as we did for the exponential family in Section 6.2. Instead, given an (E)LM family $k^{(1)}$ as in Definition 2, we come back to a monotonic decrease condition of the form

$$\int_Y \check{\varphi}(y) \log \left(\frac{k^{(1)}(\theta, y)}{k^{(1)}(\theta_0, y)} \right) \nu(dy) \geq 0. \quad (49)$$

As detailed in the remark below, for selected choices of $\check{\varphi}$ and of (θ, θ_0) , (49) can indeed be linked to the conditions (5) and (19) appearing in Theorem 1 and Theorem 2 respectively.

Remark 1 *Let $\alpha \in [0, 1)$. Taking $\check{\varphi} = \check{\varphi}_n^{(\alpha)}$ defined by (6) and $(\theta, \theta_0) = (\theta_{n+1}, \theta_n)$, (49) becomes the condition (5) of Theorem 1. Applying (49) for $j = 1, \dots, J$ with $\check{\varphi} = \check{\varphi}_{j,n}^{(\alpha)}$ defined by (16) and $(\theta, \theta_0) = (\theta_{j,n+1}, \theta_{j,n})$ yields the condition (19) of Theorem 2.*

Let $D_{\tau', \tau} = \frac{d\tau'}{d\tau}$ denote the Radon-Nikodym derivative of τ' w.r.t. τ . We then have the following proposition (we defer the proofs of the results from Section 6.4.2 to Appendix C.4.1).

Proposition 4 (Conditions for a monotonic decrease within the (E)LM family)

Let $k^{(1)}$ be an LM family as in Definition 2. Let $\check{\varphi} : Y \rightarrow \mathbb{R}_+$ be a probability density function with respect to ν . Let $\vartheta_0, \vartheta \in \tilde{T}$ and $\tau_0, \tau \in M$ such that

$$\int_Y \left(\int_Z \varphi(\ell, y) \log(D_{\tau, \tau_0}(\ell)) \tau_0(d\ell) \right) \nu(dy) \geq 0, \quad (50)$$

$$\int_Y \left(\int_Z \varphi(\ell, y) \log \left(\frac{k^{(o)}(\ell \circ \Upsilon(\vartheta), y)}{k^{(o)}(\ell \circ \Upsilon(\vartheta_0), y)} \right) \tau_0(d\ell) \right) \nu(dy) \geq 0 \quad (51)$$

where we set

$$\varphi(\ell, y) = \check{\varphi}(y) \frac{k^{(o)}(\ell \circ \Upsilon(\vartheta_0), y)}{k^{(1)}(\theta_0, y)}. \quad (52)$$

Then (49) holds with $\theta = (\vartheta, \tau)$ and $\theta_0 = (\vartheta_0, \tau_0)$.

One possible way to obtain (50) is to observe that since the left-hand side of (50) is zero for $\tau = \tau_0$, (50) is fulfilled by setting

$$\tau = \operatorname{argmax}_{\tau' \in M} \int_Y \left(\int_Z \varphi(\ell, y) \log(D_{\tau', \tau_0}(\ell)) \tau_0(d\ell) \right) \nu(dy), \quad (53)$$

assuming that this argmax is well-defined. This will notably be done later in order to get the update formula (70) of Example 5 (see the proof of Corollary 10 in Appendix C.4).

As for the updating of ϑ , we can again adopt a maximisation approach to derive a convenient ϑ . This is the purpose of the following result, where we directly update $\xi = \Upsilon(\vartheta)$ since Υ is a known mapping in Definition 2.

Corollary 9 *Consider an LM family as in Definition 2 and $\theta_0 = (\vartheta_0, \tau_0) \in \tilde{\mathbf{T}} \times \mathbf{M}$. Set $\xi_0 = \Upsilon(\vartheta_0)$, let $\tilde{\varphi} : \mathbf{Y} \rightarrow \mathbb{R}_+$ be a probability density function with respect to ν and define $\varphi : \mathbf{Z} \times \mathbf{Y} \rightarrow \mathbb{R}_+$ by (52). Let $\vartheta \in \tilde{\mathbf{T}}$ be such that $\Upsilon(\vartheta)$ is a solution of the argmax problem*

$$\operatorname{argmax}_{\xi \in \Upsilon(\tilde{\mathbf{T}})} \int_{\mathbf{Y}} \left(\int_{\mathbf{Z}} [\varphi(\ell, y) + b_0(\ell) k^{(o)}(\ell(\xi_0), y)] \log \left(\frac{k^{(o)}(\ell(\xi), y)}{k^{(o)}(\ell(\xi_0), y)} \right) \tau_0(d\ell) \right) \nu(dy), \quad (54)$$

where b_0 is any non-negative function defined on \mathbf{Z} such that $\int b_0 d\tau_0 < \infty$. Then (51) holds.

As we will see, (53) and (54) can be achieved for continuous mixtures such as the one used to define Student's t distributions but they require a thorough analysis. To this end, we next formulate and solve a maximisation approach for a kernel $k^{(1)}$ belonging to the ELM family.

6.4.3 THE MAXIMISATION APPROACH FOR THE ELM FAMILY

Let us show that we can solve an argmax of the form (54) for an adequate choice of b_0 by fully exploiting the fact that $k^{(o)}(\zeta, y)$ is a canonical kernel of the exponential family distribution which satisfies (B1) and by relying on additional assumptions to be introduced alongside some helpful notation. Our first assumption is the following.

(C1) The image set $\Upsilon(\tilde{\mathbf{T}})$ is an open subset of F and for all $\tau \in \mathbf{M}$, τ -a.e. ℓ is full rank.

(C1) guarantees that, for all $\vartheta \in \tilde{\mathbf{T}}$, $\tau \in \mathbf{M}$ and τ -almost all ℓ in \mathbf{Z} , $\ell \circ \Upsilon(\vartheta)$ belongs to $\operatorname{Int} E_0$, the interior set of E_0 (see Lemma 9 in Appendix C.4 for details). This will come in handy in the upcoming derivations and we now introduce some helpful notation, before presenting our next two assumptions. We let $\|\cdot\|_{\text{op}}$ denote operator norms, for instance, for any $\ell \in \mathbf{Z}$,

$$\|\ell\|_{\text{op}} = \sup \{ \|\ell(x)\|_F : x \in E, \|x\|_E \leq 1 \}.$$

We further denote the adjoint of the linear operator ℓ by ℓ^T , for instance, in the case $\ell \in \mathbf{Z}$, $\ell^T \in \mathcal{L}(E, F)$ is defined by

$$\langle \ell(x), y \rangle_E = \langle x, \ell^T(y) \rangle_F \quad \text{for all } (x, y) \in F \times E \quad (55)$$

and, for convenience, we also denote

$$\ell^\nabla := \ell^T \circ \nabla A \circ \ell. \quad (56)$$

Here, A is the function appearing in the definition of an exponential family (Definition 1). Recall that A is infinitely differentiable on $\operatorname{Int} E_0$ so that by (C1), ℓ^∇ is well-defined on $\Upsilon(\tilde{\mathbf{T}})$ for τ -a.e. ℓ with $\tau \in \mathbf{M}$. Furthermore, we use $\mathbb{P}_{\vartheta, \tau}$ and $\mathbb{E}_{\vartheta, \tau}$ to denote the probability and its corresponding expectation under which the pair (Y, L) has density $(y, \ell) \mapsto k^{(o)}(\ell \circ \Upsilon(\vartheta), y)$ with respect to $\nu \otimes \tau$. Namely, for any non-negative measurable g on $\mathbf{Y} \times \mathbf{Z}$,

$$\mathbb{E}_{\vartheta, \tau} [g(Y, L)] = \int g(y, \ell) k^{(o)}(\ell \circ \Upsilon(\vartheta), y) \nu(dy) \tau(d\ell). \quad (57)$$

We now introduce, for any $(\vartheta_0, \tau_0) \in \tilde{\mathbf{T}} \times \mathbf{M}$, two auxiliary functions $\tilde{m}_{\vartheta_0, \tau_0}$ and $\tilde{s}_{\vartheta_0, \tau_0}$, both defined from \mathbf{Y} to \mathbb{R}_+ , such that the two following conditions hold.

(C2) For all $\vartheta_0 \in \tilde{\mathbf{T}}$ and $\tau_0 \in \mathbf{M}$, we have

$$\mathbb{E}_{\vartheta_0, \tau_0} \left[\|L\|_{\text{op}} \mid Y \right] \leq \tilde{m}_{\vartheta_0, \tau_0}(Y) \quad \mathbb{P}_{\vartheta_0, \tau_0} - \text{a.s.}$$

(C3) For all $\vartheta_0 \in \tilde{\mathbf{T}}$ and $\tau_0 \in \mathbf{M}$, we have

$$\forall \xi \in \Upsilon(\tilde{\mathbf{T}}), \exists \epsilon, C > 0, \mathbb{E}_{\vartheta_0, \tau_0} \left[e^{\epsilon \|L^\nabla(\xi)\|_F} \mid Y \right] \leq C e^{\tilde{s}_{\vartheta_0, \tau_0}(Y)} \quad \mathbb{P}_{\vartheta_0, \tau_0} - \text{a.s.}$$

(C2) and (C3) are made to ensure that the integrals used to solve the argmax problem from Corollary 9 will be well-defined. We finally present our last assumption, which corresponds to some sort of identifiability assumption (see Lemma 10 in Appendix C.4 for details). It is used to obtain the uniqueness for the argmax problem we will solve.

(C4) For all $\xi \neq \xi' \in \Upsilon(\tilde{\mathbf{T}})$ and $\tau_0 \in \mathbf{M}$, we have $\tau_0(\{\ell \in \mathbf{Z} : \ell(\xi) \neq \ell(\xi')\}) > 0$.

We can now state the following result, whose proof can be found in Appendix C.4.2.

Theorem 5 *Consider an ELM family as in Definition 2 and let $\theta_0 = (\vartheta_0, \tau_0) \in \tilde{\mathbf{T}} \times \mathbf{M}$. Assume (C1)–(C4) hold. Let $\tilde{\varphi}$ be a probability density function w.r.t. ν such that*

$$\int_Y \tilde{\varphi}(y) \left(\|S(y)\|_E \tilde{m}_{\vartheta_0, \tau_0}(y) + e^{\tilde{s}_{\vartheta_0, \tau_0}(y)} \right) \nu(dy) < \infty, \quad (58)$$

define $\varphi : \mathbf{Z} \times \mathbf{Y} \rightarrow \mathbb{R}_+$ by (52) and set for all $(\ell, y, \xi) \in \mathbf{Z} \times \mathbf{Y} \times \Upsilon(\tilde{\mathbf{T}})$ such that $\ell \circ \Upsilon(\vartheta_0) \in \text{Int} E_0$

$$\tilde{\varphi}(\ell) = \int_Y \varphi(\ell, y) \nu(dy), \quad (59)$$

$$\mathbf{w}_{\tau_0}^\nabla(\xi) = \int_{\mathbf{Z}} \tilde{\varphi}(\ell) \ell^\nabla(\xi) \tau_0(d\ell). \quad (60)$$

Then $\tilde{\varphi} < \infty$ τ_0 -a.s., and $\mathbf{w}_{\tau_0}^\nabla$ is well-defined from $\Upsilon(\tilde{\mathbf{T}})$ to F and one-to-one on $\Upsilon(\tilde{\mathbf{T}})$, hence bijective from $\Upsilon(\tilde{\mathbf{T}})$ to its image $\mathbf{w}_{\tau_0}^\nabla \circ \Upsilon(\tilde{\mathbf{T}})$, and satisfies, for all $\vartheta \in \tilde{\mathbf{T}}$,

$$\mathbf{w}_{\tau_0}^\nabla \circ \Upsilon(\vartheta) = \mathbb{E}_{\vartheta, \tau_0} \left[L^T \circ S(Y) \tilde{\varphi}(L) \right]. \quad (61)$$

Moreover, setting $\xi_0 = \Upsilon(\vartheta_0)$, for any $b \in \mathbb{R}_+$, the argmax problem

$$\operatorname{argmax}_{\xi \in \Upsilon(\tilde{\mathbf{T}})} \int_Y \left(\int_{\mathbf{Z}} \left[\varphi(\ell, y) + b \tilde{\varphi}(\ell) k^{(o)}(\ell(\xi_0), y) \right] \log \left(\frac{k^{(o)}(\ell(\xi), y)}{k^{(o)}(\ell(\xi_0), y)} \right) \tau_0(d\ell) \right) \nu(dy), \quad (62)$$

has at least a solution if and only if

$$\mathbf{w}^* = \frac{1}{1+b} \int \varphi(\ell, y) \ell^T \circ S(y) \nu(dy) \tau_0(d\ell) + \frac{b}{1+b} \mathbf{w}_{\tau_0}^\nabla(\xi_0) \quad (63)$$

belongs to $\mathbf{w}_{\tau_0}^\nabla \circ \Upsilon(\tilde{\mathbf{T}})$, in which case (62) has a unique solution ξ^* defined by

$$\mathbf{w}_{\tau_0}^\nabla(\xi^*) = \mathbf{w}^*. \quad (64)$$

By setting $\xi^* = \Upsilon(\vartheta^*)$ and $\xi_0 = \Upsilon(\vartheta_0)$ as well as using (61), (64) can be interpreted as

$$\mathbb{E}_{\vartheta^*, \tau_0} [L^T \circ S(Y) \tilde{\varphi}(L)] = \mathbb{E}_{\vartheta_0, \tilde{\tau}_0} [L^T \circ S(\tilde{Y})] , \quad (65)$$

where $\tilde{\tau}_0$ denotes the probability having density $\tilde{\varphi}$ with respect to τ_0 and \tilde{Y} denotes a random variable valued in Y such that, under $\mathbb{P}_{\vartheta_0, \tilde{\tau}_0}$, the pair (\tilde{Y}, L) is distributed according to the mixture density

$$(y, \ell) \mapsto \psi(y, \ell) = \frac{1}{1+b} \varphi(\ell, y) + \frac{b}{1+b} k^{(o)}(\ell \circ \Upsilon(\vartheta_0), y) \tilde{\varphi}(\ell)$$

with respect to $\nu \otimes \tau_0$. Interestingly, the second marginals of these two mixands are $\tilde{\tau}_0$. The distribution of (\tilde{Y}, L) under $\mathbb{P}_{\vartheta_0, \tilde{\tau}_0}$ can thus equivalently be defined by saying that $L \sim \tilde{\tau}_0$ and the conditional distribution of \tilde{Y} given $L = \ell$ is the mixture with density

$$y \mapsto \psi[y | \ell] = \frac{1}{1+b} \frac{\varphi(\ell, y)}{\tilde{\varphi}(\ell)} + \frac{b}{1+b} k^{(o)}(\ell \circ \Upsilon(\vartheta_0), y) .$$

Using the definition of $\tilde{\tau}_0$ on the left-hand side of (65) as well, the latter reads

$$\mathbb{E}_{\vartheta^*, \tilde{\tau}_0} [L^T \circ S(Y)] = \mathbb{E}_{\vartheta_0, \tilde{\tau}_0} [L^T \circ S(\tilde{Y})] . \quad (66)$$

Hence, Theorem 5 along with (66) can be used to solve the argmax problem (54) in the case where (B1) and (C1)–(C4) hold and for $b_0(\ell) = b \tilde{\varphi}(\ell)$ with b being a non-negative constant. From there, it only remains to find ways to check that the condition (58) on $\check{\varphi}$ holds. This can notably be done when $\check{\varphi}$ takes one of the two forms listed in Remark 1 (see Lemma 8 in Appendix C.2 for details) and we now provide an example of particular interest.

Example 5 (Finite mixture of Student's t distributions) *Let ν be the Lebesgue measure on $\mathsf{Y} = \mathbb{R}^d$, with $d \geq 1$. Let $\alpha \in [0, 1)$ and $p : \mathbb{R}^d \rightarrow \mathbb{R}_+$ such that*

$$0 < \int_{\mathbb{R}^d} (1 + \|y\|^2)^{1/(1-\alpha)} p(y) \, dy < \infty . \quad (67)$$

Let $J \geq 1$ and consider a J -mixture family \mathcal{Q} as in (15), with $k(\theta, \cdot)$ being the density of the Student's t distribution with parameter $\theta = (m, \Sigma, a) \in \mathsf{T} = \mathbb{R}^d \times \mathcal{M}_{>0}(d) \times (0, \infty)$, given by

$$k(\theta, y) = \frac{\Gamma((a+d)/2)}{\Gamma(a/2) (a\pi)^{d/2} |\Sigma|^{1/2}} \left(1 + \frac{1}{a} (y-m)^T \Sigma^{-1} (y-m) \right)^{-(a+d)/2} .$$

Define a sequence $(\mu_n k)_{n \geq 1}$ valued in \mathcal{Q} where $\mu_n = \mu_{\lambda_n, \Theta_n}$ with $\lambda_n \in \mathcal{S}_J^+$ and $\Theta_n = (m_{j,n}, \Sigma_{j,n}, a_{j,n})_{j=1 \dots J} \in \mathsf{T}^J$ such that, for all $n \geq 1$,

(i) given $(\eta_n)_{n \geq 1}$ valued in $(0, 1]$, $(\kappa_n)_{n \geq 1}$ valued in $(-\infty, 0]$ and $\lambda_1 \in \mathcal{S}_J^+$, $(\lambda_n)_{n \geq 2}$ is defined by the iterative formula (21);

(ii) given $(\gamma_{j,n})_{j=1,\dots,J,n \geq 1}$ valued in $(0, 1]$ and $\Theta_1 \in \mathbb{T}^J$, $(\Theta_n)_{n \geq 2}$ is defined by

$$m_{j,n+1} = \left(\int_{\mathbb{R}^d \times \mathbb{R}_+} z \psi_{j,n}(z, y) \, dy \, \check{\tau}_{a_{j,n}}(dz) \right)^{-1} \int_{\mathbb{R}^d \times \mathbb{R}_+} y z \psi_{j,n}(z, y) \, dy \, \check{\tau}_{a_{j,n}}(dz), \quad (68)$$

$$\Sigma_{j,n+1} = \int_{\mathbb{R}^d \times \mathbb{R}_+} (y - m_{j,n+1})(y - m_{j,n+1})^T z \psi_{j,n}(z, y) \, dy \, \check{\tau}_{a_{j,n}}(dz), \quad (69)$$

$$a_{j,n+1} = 2 \kappa^{-1} \left(\int_{\mathbb{R}^d \times \mathbb{R}_+} (z - \ln(z)) \varphi_{j,n}(z, y) \, dy \, \check{\tau}_{a_{j,n}}(dz) \right), \quad (70)$$

where, defining $\check{\varphi}_{j,n}^{(\alpha)}$ as in (16) and denoting $\kappa^{-1} : \mathbb{R} \rightarrow (0, \infty)$ the inverse mapping of $\kappa(x) = \log(x) + \Gamma'(x)/\Gamma(x)$, we set

$$\varphi_{j,n}(z, y) = \check{\varphi}_{j,n}^{(\alpha)}(y) \frac{e^{\frac{-z}{2}(y-m_{j,n})^T \Sigma_{j,n}^{-1}(y-m_{j,n})}}{(2\pi/z)^{d/2} |\Sigma_{j,n}|^{1/2} k(\theta_{j,n}, y)}, \quad (71)$$

$$\psi_{j,n}(z, y) = \gamma_{j,n} \varphi_{j,n}(z, y) + (1 - \gamma_{j,n}) \frac{e^{\frac{-z}{2}(y-m_{j,n})^T \Sigma_{j,n}^{-1}(y-m_{j,n})}}{(2\pi/z)^{d/2} |\Sigma_{j,n}|^{1/2}} \left(\int \varphi_{j,n}(z, y) \, dy \right), \quad (72)$$

$$\check{\tau}_a(dz) = \mathbb{1}_{\mathbb{R}_+}(z) \frac{(a/2)^{a/2}}{\Gamma(a/2)} z^{a/2-1} e^{-za} \, dz \quad \text{for any } a > 0, \quad (73)$$

Here, $\check{\tau}_a$ in (73) is the χ^2 distribution with a degrees of freedom and κ^{-1} in (70) is well-defined (this follows from Lemma 13 in Appendix C.4.3).

We then have the following result, whose proof is postponed to Appendix C.4.3.

Corollary 10 *In the setting of Example 5, if $\Psi_\alpha(\mu_1 k) < \infty$, then, at any time $n \geq 1$, we have $\Psi_\alpha(\mu_{n+1} k) \leq \Psi_\alpha(\mu_n k)$.*

We have stated results that solve our maximisation and gradient-based approaches when the variational family is based on the exponential family. However, all the update formulas we have obtained so far involve intractable integrals. In the next section, we focus on the case of Gaussian Mixture Models optimisation seen in Examples 3 and 4 and we investigate how our framework can be used to derive tractable algorithms in that case.

7. Gaussian Mixture Models (GMMs) Optimisation

We consider d -dimensional Gaussian mixture densities with $k(\theta_j, y) = \mathcal{N}(y; m_j, \Sigma_j)$, where $\theta_j = (m_j, \Sigma_j) \in \mathbb{T}$ denotes the mean and covariance matrix of the j -th Gaussian component density. By Theorem 2, a valid update formula at time $n \geq 1$ for $(\lambda_n)_{n \geq 1}$ is

$$\lambda_{j,n+1} = \frac{\lambda_{j,n} \left[\int_{\mathbb{Y}} \varphi_{j,n}^{(\alpha)}(y) \nu(dy) + (\alpha - 1) \kappa_n \right]^{\eta_n}}{\sum_{\ell=1}^J \lambda_{\ell,n} \left[\int_{\mathbb{Y}} \varphi_{\ell,n}^{(\alpha)}(y) \nu(dy) + (\alpha - 1) \kappa_n \right]^{\eta_n}}, \quad j = 1 \dots J. \quad (74)$$

Furthermore, as underlined in Example 3, a valid update at time $n \geq 1$ for the means $(m_{j,n+1})_{1 \leq j \leq J}$ and the covariances matrices $(\Sigma_{j,n+1})_{1 \leq j \leq J}$ is given by

$$\begin{aligned} \forall j = 1 \dots J, \quad m_{j,n+1} &= (1 - \gamma_{j,n}) m_{j,n} + \gamma_{j,n} \hat{m}_{j,n} \\ \Sigma_{j,n+1} &= (1 - \gamma_{j,n}) \Sigma_{j,n} + \gamma_{j,n} \hat{\Sigma}_{j,n} + \gamma_{j,n} (1 - \gamma_{j,n}) (\hat{m}_{j,n} - m_{j,n}) (\hat{m}_{j,n} - m_{j,n})^T \end{aligned} \quad (75)$$

where $\hat{m}_{j,n} = \int_Y y \check{\varphi}_{j,n}^{(\alpha)}(y) \nu(dy)$, $\hat{\Sigma}_{j,n} = \int_Y yy^T \check{\varphi}_{j,n}^{(\alpha)}(y) \nu(dy) - \hat{m}_{j,n} \hat{m}_{j,n}^T$ with $\check{\varphi}_{j,n}^{(\alpha)} = \varphi_{j,n}^{(\alpha)} / \int \varphi_{j,n}^{(\alpha)} \nu(dy)$ and $\gamma_{j,n} \in (0, 1]$. In addition, another possible update for the means is

$$m_{j,n+1} = m_{j,n} + \gamma_{j,n} \frac{\int_Y \lambda_{j,n} \varphi_{j,n}^{(\alpha)}(y)(y - m_{j,n}) \nu(dy)}{\int_Y \mu_n k(y)^\alpha p(y)^{1-\alpha} \nu(dy)}, \quad j = 1 \dots J. \quad (76)$$

By virtue of Example 4, this update can be used in lieu of the means update written in (75) with $\gamma_{j,n}$ still in $(0, 1]$. Since the update (76) is linked to Gradient Descent steps for Rényi's α -divergence minimisation (see Example 4), we will refer to this update as the *Rényi Gradient Descent* (RGD) update in the following. On the other hand, the update on the means written in (75) will be referred to as the *Maximisation Gradient* (MG) update.

Notice that the updates (74), (75) and (76) all involve intractable integrals. However, we can resort to approximate update rules and take advantage of the fact that the integrals appearing in these update formulas are of the form

$$\int_Y \varphi_{j,n}^{(\alpha)}(y) g(y) \nu(dy), \quad (77)$$

where g is a well-chosen function defined on Y and where $\varphi_{j,n}^{(\alpha)}$ is defined in (16) by $\varphi_{j,n}^{(\alpha)}(y) = k(\theta_{j,n}, y) (\mu_n k(y)/p(y))^{\alpha-1}$ for all $y \in Y$. Many choices are indeed possible to approximate (77) and for simplicity we restrict ourselves to typical Importance Sampling estimation. Looking at the definition of $\varphi_{j,n}^{(\alpha)}$, a first idea would be to sample M i.i.d. random variables $(Y_{m,n}^{(j)})_{1 \leq m \leq M}$ according to $k(\theta_{j,n}, \cdot)$ for $j = 1 \dots J$ and to use the unbiased estimator of (77)

$$\frac{1}{M} \sum_{m=1}^M \left(\frac{\mu_n k(Y_{m,n}^{(j)})}{p(Y_{m,n}^{(j)})} \right)^{\alpha-1} g(Y_{m,n}^{(j)}).$$

As this sampler depends on j , this option becomes very computationally heavy as J and M increase due to the *sampling cost* ($J \times M$) and to the *evaluation cost* ($\mu_n k$ and p need to be evaluated $M \times J$ times each, where $\mu_n k$ is a sum over J). To reduce this computational bottleneck, it is then preferable to consider sequences of samplers $(q_n)_{n \geq n}$ that do not depend on j at time n . More specifically, we discuss two possibilities.

(i) **Best sampler at time n (IS-n).** Approximated update rules can be obtained by sampling M i.i.d. random variables $(Y_{m,n})_{1 \leq m \leq M}$ according to the best approximation of p we have at time n , that is $q_n = \mu_n k$. Then, the samples are shared throughout the J mixture weights and mixture components parameters updates. This eases the computational burden, with both a sampling and an evaluation costs of M , as an unbiased estimator of (77) is

$$\frac{1}{M} \sum_{m=1}^M \hat{\varphi}_{j,n}^{(\alpha)}(Y_{m,n}) g(Y_{m,n}) \quad \text{where} \quad \hat{\varphi}_{j,n}^{(\alpha)}(y) = \frac{k(\theta_{j,n}, y)}{q_n(y)} \left(\frac{\mu_n k(y)}{p(y)} \right)^{\alpha-1}, \quad y \in Y. \quad (78)$$

(ii) **Uniform sampler (IS-unif).** Approximated update rules can be obtained by sampling according to the uniform sampler $q_n = J^{-1} \sum_{j=1}^J k(\theta_{j,n}, \cdot)$ and using the unbiased estimator (78). That way, the sampling and evaluation costs are also M each. Contrary to IS-n, this sampler ensures a fair sampling among all components, as it entails sampling an index vector $[i_1, i_2, \dots, i_M]$ uniformly from $\{1, \dots, J\}$ whereas IS-n samples an index vector $[i_1, i_2, \dots, i_M]$ according to the mixture weights λ_n .

Algorithm 3: GMMs optimisation with the IS-n/IS-unif sampler

At iteration n ,

1. Draw independently M samples $(Y_{m,n})_{1 \leq m \leq M}$ from the proposal q_n .

2. For all $j = 1 \dots J$,

(a) Set $\hat{m}_{j,n} = \frac{\sum_{m=1}^M \hat{\varphi}_{j,n}^{(\alpha)}(Y_{m,n}) \cdot Y_{m,n}}{\sum_{m=1}^M \hat{\varphi}_{j,n}^{(\alpha)}(Y_{m,n})}$ and $\hat{\Sigma}_{j,n} = \frac{\sum_{m=1}^M \hat{\varphi}_{j,n}^{(\alpha)}(Y_{m,n}) \cdot Y_{m,n} Y_{m,n}^T}{\sum_{m=1}^M \hat{\varphi}_{j,n}^{(\alpha)}(Y_{m,n})} - \hat{m}_{j,n} \hat{m}_{j,n}^T$.

(b) Choose one between the (MG) and (RGD) approaches

$$(MG) \quad m_{j,n+1} = (1 - \gamma_n) m_{j,n} + \gamma_n \hat{m}_{j,n}$$

$$(RGD) \quad m_{j,n+1} = m_{j,n} + \gamma_n \frac{\lambda_{j,n} \sum_{m=1}^M \hat{\varphi}_{j,n}^{(\alpha)}(Y_{m,n}) \cdot (Y_{m,n} - m_{j,n})}{\sum_{j=1}^J \sum_{m=1}^M \lambda_{j,n} \hat{\varphi}_{j,n}^{(\alpha)}(Y_{m,n})}$$

and set

$$\lambda_{j,n+1} = \frac{\lambda_{j,n} \left[\sum_{m=1}^M \hat{\varphi}_{j,n}^{(\alpha)}(Y_{m,n}) + (\alpha - 1) \kappa_n \right]^{\eta_n}}{\sum_{\ell=1}^J \lambda_{\ell,n} \left[\sum_{m=1}^M \hat{\varphi}_{\ell,n}^{(\alpha)}(Y_{m,n}) + (\alpha - 1) \kappa_n \right]^{\eta_n}}$$

$$\Sigma_{j,n+1} = (1 - \gamma_n) \Sigma_{j,n} + \gamma_n \hat{\Sigma}_{j,n} + \gamma_n (1 - \gamma_n) (\hat{m}_{j,n} - m_{j,n}) (\hat{m}_{j,n} - m_{j,n})^T.$$

Remark 2 *Approximated update rules for GMMs can be obtained by sampling according to $q_n = \mathcal{N}(\cdot; \mathbf{0}_d, \mathbf{I}_d)$, where $\mathbf{0}_d$ is the null vector of dimension d . These can be deduced by applying the reparameterisation trick: for all $m = 1 \dots M$, $Y_{m,n}^{(j)} \sim k(\theta_{j,n}, \cdot)$ if $Y_{m,n}^{(j)} = m_{j,n} + \Sigma_{j,n}^{-1/2} \varepsilon_{m,n}$ with $\varepsilon_{m,n} \sim q_n$ and by observing that an unbiased estimator of (77) is $M^{-1} \sum_{m=1}^M \left(\mu_n k(Y_{m,n}^{(j)}) / p(Y_{m,n}^{(j)}) \right)^{\alpha-1} g(Y_{m,n}^{(j)})$. This choice of q_n relies on the existence of a reparameterisation and while it incurs a gain in terms of sampling cost as we only need M samples, the evaluation cost remains $M \times J$.*

We thus have access to tractable algorithms that iteratively update both the weights and components parameters of a GMM by optimising the α -divergence between the mixture distribution and the targeted distribution. Our framework also allows the use of learning rates $\gamma_{j,n}$ that are dependent on j . This means that the learning rates can differ for each component parameters update and it paves the way for adaptive learning rates per components. This aspect is left for future work as we will focus on the case $\gamma_{j,n} := \gamma_n$ according to Algorithm 3 and move on to presenting our numerical experiments.

Remark 3 *Computing the estimator (78) can be expensive for two reasons: (i) $\mu_n k$ involves a sum over J so that evaluating this function requires heavier computations as the number of components J grows and (ii) bayesian applications with $p = p(\cdot, \mathcal{D})$ often consider large-scale data sets \mathcal{D} . To alleviate this computational burden, one can (i) approximate the summation*

appearing in $\mu_n k$ using subsampling and (ii) use a mini-batch approach to approximate $p = p(\cdot, \mathcal{D})$, as detailed in (Li and Turner, 2016, Section 4.3).

8. Numerical Experiments

In our numerical experiments, we are interested in understanding the behaviour of Algorithm 3 in practice. We first investigate the challenging case of multimodal targets, as our framework is designed to handle multimodality thanks to the hyperparameter $\alpha \in [0, 1)$ and to the use of mixture models as a variational family.

8.1 Toy multimodal targets

Let \mathbf{u}_d be the d -dimensional vector whose coordinates are all equal to 1, \mathbf{I}_d be the identity matrix and c be a positive constant (we set $c = 2$). We consider three multimodal examples:

(i) *Equally-weighted Gaussian Mixture Model*. The target p is a mixture density of two equally-weighted d -dimensional Gaussian distributions multiplied by c such that

$$p(y) = c \times [0.5\mathcal{N}(y; -2\mathbf{u}_d, \mathbf{I}_d) + 0.5\mathcal{N}(y; 2\mathbf{u}_d, \mathbf{I}_d)] .$$

(ii) *Imbalanced Gaussian Mixture Model*. The target p is a mixture density of three d -dimensional Gaussian distributions with unequal weights and multiplied by c such that

$$p(y) = c \times [0.35\mathcal{N}(y; -2\mathbf{u}_d, \mathbf{I}_d) + 0.25\mathcal{N}(y; 2\mathbf{u}_d, \mathbf{I}_d) + 0.4\mathcal{N}(y; \mathbf{u}_d, \mathbf{I}_d)]$$

(iii) *Equally-weighted Student's t Mixture Model*. The target p is a mixture density of two equally-weighted d -dimensional Student's t distributions with two degree of freedom (i.e. $a = 2$) multiplied by c such that

$$p(y) = c \times [0.5 \, \mathfrak{t}(y; -2\mathbf{u}_d, \mathbf{I}_d, a) + 0.5 \, \mathfrak{t}(y; 2\mathbf{u}_d, \mathbf{I}_d, a)] .$$

8.1.1 COMPARING THE RGD AND THE MG APPROACHES WITH $\eta_n = 0$

Setting $\eta_n = 0$ at all times $n \geq 1$ keeps the mixture weights fixed in Algorithm 3. In that case, our RGD approach relates to Gradient Ascent steps on the VR Bound (Li and Turner, 2016) w.r.t. the means of our Gaussian components (see Section 5.1). We then want to compare the performances of the RGD approach (that can be derived from Li and Turner, 2016) to the novel MG approach we have introduced in our work.

Implementation details. The covariance matrices of the mixture components are fixed and equal to $\sigma^2 \mathbf{I}_d$ with $\sigma^2 = 1$, $\alpha = 0.2$, $J \in \{10, 50\}$, $d = 16$, $M = 200$, the total number of iterations N is equal to 100, Θ_1 is generated by sampling from a centered normal distribution with covariance matrix $10\mathbf{I}_d$, $\boldsymbol{\lambda}_1 = [1/J, \dots, 1/J]$ and for all time $n = 1 \dots N$, $\kappa_n = 0$, $\eta_n = 0$. and $\gamma_n = \gamma$ with $\gamma \in \{0.1, 0.5, 1.\}$. The experiments are replicated independently 30 times and the convergence of the RGD and of the MG approaches is monitored for the three multimodal examples (i), (ii) and (iii) by computing a Monte Carlo estimator of the VR Bound (since we have sampled M samples from q_n at time n , these samples can readily be used to obtain an estimate the VR Bound with no additional computational cost).

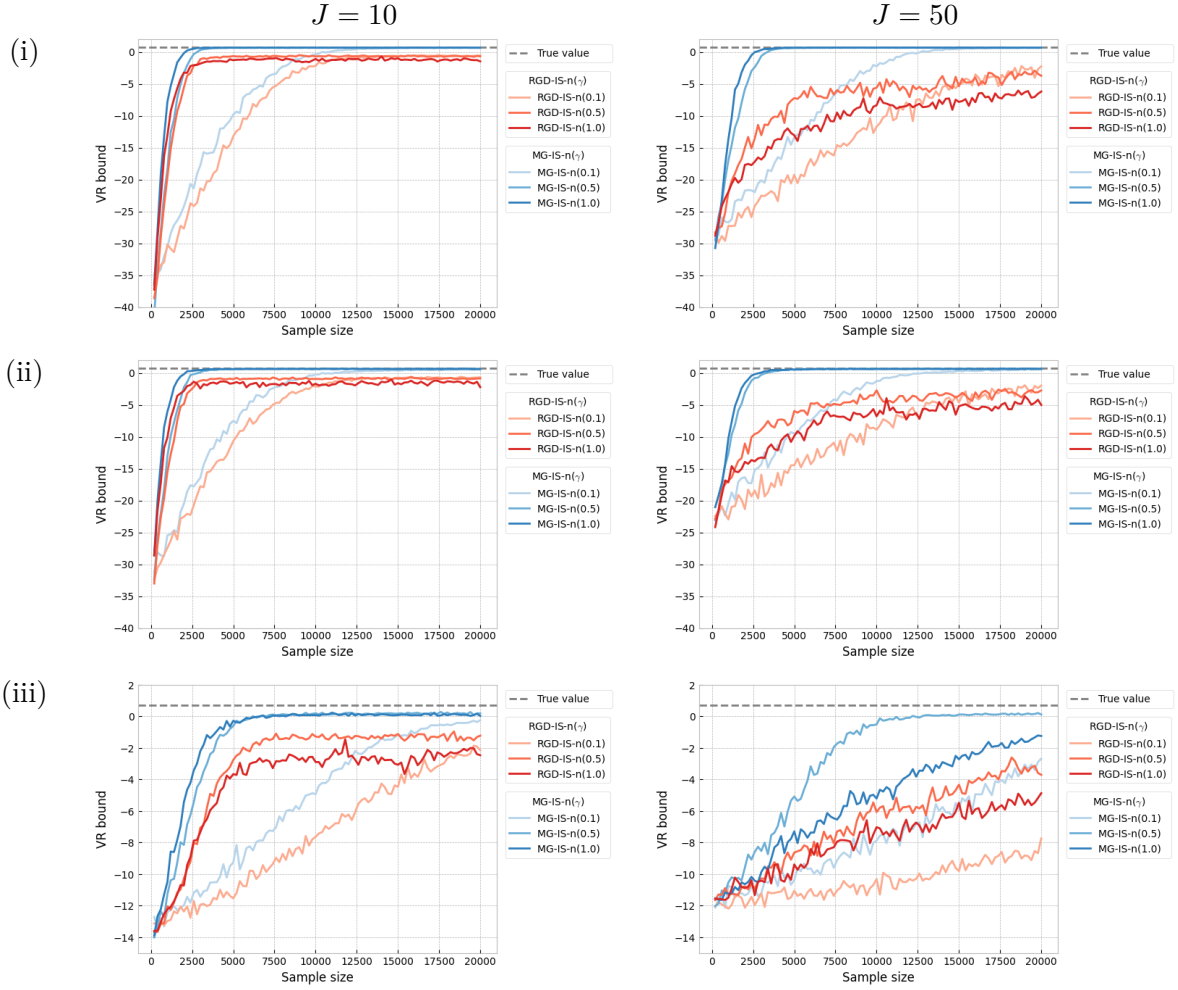


Figure 1: Monte Carlo estimate of the VR Bound for the RGD and the MG approaches (fixed mixture weights) when considering each of the target distributions (i), (ii) and (iii).

Our results are plotted on Figure 1, where $\text{RGD-IS-n}(\gamma)$ and $\text{MG-IS-n}(\gamma)$ denote the algorithms originating from the RGD and MG approaches (the IS-n and IS-unif samplers are equivalent here), and error bounds plots can be found in Figures 5 and 6 of Appendix D. As minimising Ψ_α is equivalent to maximising the VR Bound when $\alpha \in (0, 1)$, our theoretical results predict a systematic increase in the VR Bound for the algorithms considered. This is what we observe in Figure 1 and we also see that the choice of γ plays the role of a learning rate our algorithms: while increasing γ mostly improves the speed of convergence, it may deteriorate the performances if chosen too large, such as when $J = 50$ with the target (iii).

Furthermore, $\text{MG-IS-n}(\gamma)$ leads to better performances in all but one case compared to $\text{RDG-IS-n}(\gamma)$. This renders the $\text{MG-IS-n}(\gamma)$ algorithm more suitable overall for capturing the multimodality of the various multimodal targets p we considered, a fact that is further

| | | $J = 10$ | | | $J = 50$ | | |
|-------|----------------------|----------------|----------------|----------------|----------------|----------------|----------------|
| | | $\gamma = 0.1$ | $\gamma = 0.5$ | $\gamma = 1.0$ | $\gamma = 0.1$ | $\gamma = 0.5$ | $\gamma = 1.0$ |
| (i) | RGD-IS-n(γ) | -0.081 | -0.076 | -0.218 | -1.640 | -1.673 | -1.560 |
| | MG-IS-n(γ) | -3.702 | -1.875 | -2.711 | -2.760 | -2.771 | -2.788 |
| (ii) | RGD-IS-n(γ) | -0.211 | -0.072 | -0.015 | -1.401 | -1.437 | -1.515 |
| | MG-IS-n(γ) | -2.581 | -2.101 | -1.742 | -2.611 | -2.328 | -1.933 |
| (iii) | RGD-IS-n(γ) | -0.108 | -0.008 | -0.111 | -1.652 | -1.654 | -1.634 |
| | MG-IS-n(γ) | -0.913 | -1.489 | -1.846 | -2.036 | -2.530 | -0.717 |

Table 2: LogMSE averaged for the RGD and the MG approaches (fixed mixture weights) when considering each of the target distributions (i), (ii) and (iii).

supported in Table 2, in which we compare the log of the Mean-Squared Error (logMSE) returned by each algorithm. (The MSE is computed as the average of $\|m_{\text{approx}} - m_{\text{true}}\|^2$ over 30 independent runs of each algorithm, where $\|\cdot\|$ stands for the Euclidean norm, $m_{\text{true}} = \mathbb{E}_{\mathbb{P}}[Y]$ with $\mathbb{P}(A) = \int_A p(y)\nu(dy) / \int_{\mathcal{Y}} p(y)\nu(dy)$ for all $A \in \mathcal{Y}$ and where m_{approx} is the mean of the optimised variational distribution, that is in our setting $m_{\text{approx}} = \sum_{j=1}^J \lambda_{j,N} m_{j,N}$.)

Let us next pair up the means optimisation with mixture weights optimisation and investigate the impact of the sequence $(\eta_n)_{n \geq 1}$ on our algorithms.

8.1.2 COMPARING THE RGD AND THE MG APPROACHES WITH $\eta_n > 0$.

So far, we have demonstrated the viability of our novel MG approach compared to the more usual Rényi’s α -divergence-based RGD approach. Yet, we have not taken advantage of the fact that we can perform mixture weights optimisation on top of means optimisation, which is another novelty of our framework compared to traditional Variational Inference methods.

Implementation details. We work with the same implementation details as those described in Section 8.1.1, except for the fact that we will now let $\eta_n = \eta$ at time n and we will vary the value of η . In addition to the RGD-IS-n(γ) and the MG-IS-n(γ) algorithms, we also include the RGD-IS-unif(γ) and the MG-IS-unif(γ) algorithms in our results (that is, the algorithms resulting from pairing up the RGD and MG approaches with the IS-unif sampler).

The plots for the averaged Monte Carlo estimate of the VR Bound with $\eta = 0.1$ and $\gamma = 0.5$ are then available in Figure 2 (and similar plots for $\eta \in \{0.05, 0.5\}$ are provided in Figures 7 and 8 of Appendix D respectively). According to Figure 2, MG-IS-unif(0.5) outperforms most of the time the other methods (this trend for MG-IS-unif(γ) is further confirmed by looking at the logMSE in Table 3 of Appendix D, which includes the cases $\gamma \in \{0.1, 1.\}$ for completeness).

We also plot in Figure 3 the final mixture weights λ_N obtained after one typical run of MG-IS-unif(0.5) for $\eta \in \{0.05, 0.1, 0.5\}$ and $J \in \{10, 50\}$ (and additional LogMSE results can be found in Table 4 of Appendix D). This is done in order to delve into how η impacts the variational density returned at the end of Algorithm 3 (MG-IS-unif(0.5) was chosen among all four options since it enjoys good empirical performances according to Figure 2).

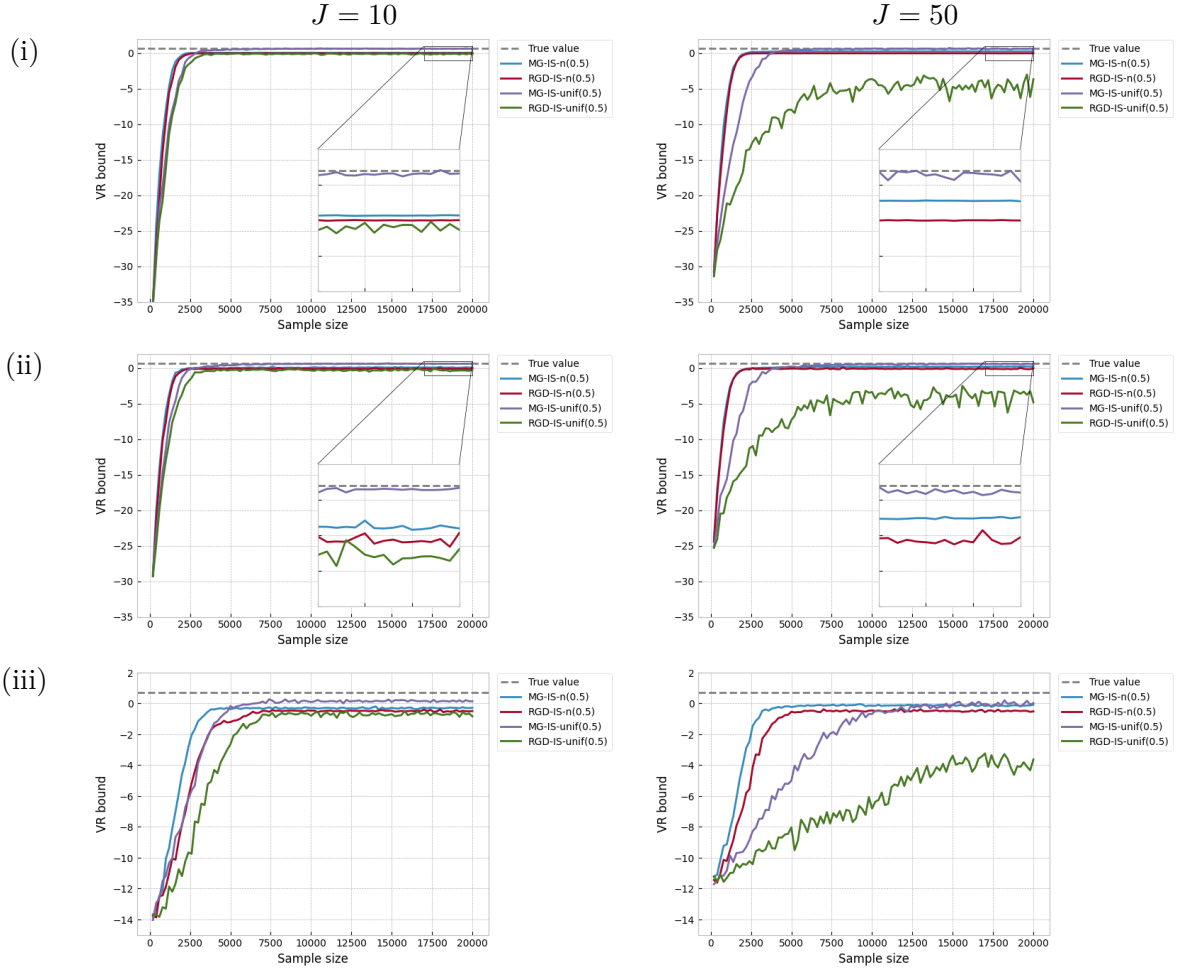


Figure 2: Monte Carlo estimate of the VR Bound for the RGD and the MG approaches ($\eta = 0.1$) when considering each of the target distributions (i), (ii) and (iii).

A key insight from Figure 3 is that optimising the mixture weights permits us to select the components appearing in our mixture model according to their overall contribution in constructing a good approximation of the targeted density.

On the one hand, we enable more flexibility in our variational approximation as we optimise over a set of component parameters Θ with $J > 1$ instead of just one component parameter θ ($J = 1$). On the other hand, we avoid complexifying unnecessarily the variational distribution returned at the end of the optimisation procedure since we perform mixture weights optimisation alongside components parameters optimisation. That way, we can bypass the limitation of the case $\eta = 0$, which may use more components than needed when approximating the targeted density (this is the case here since the targeted densities considered have three modes at best while $J = \{10, 50\}$).

However, there is a tradeoff to find between the simplicity of the variational density that is returned and its accuracy at describing the targeted density, which is expressed via

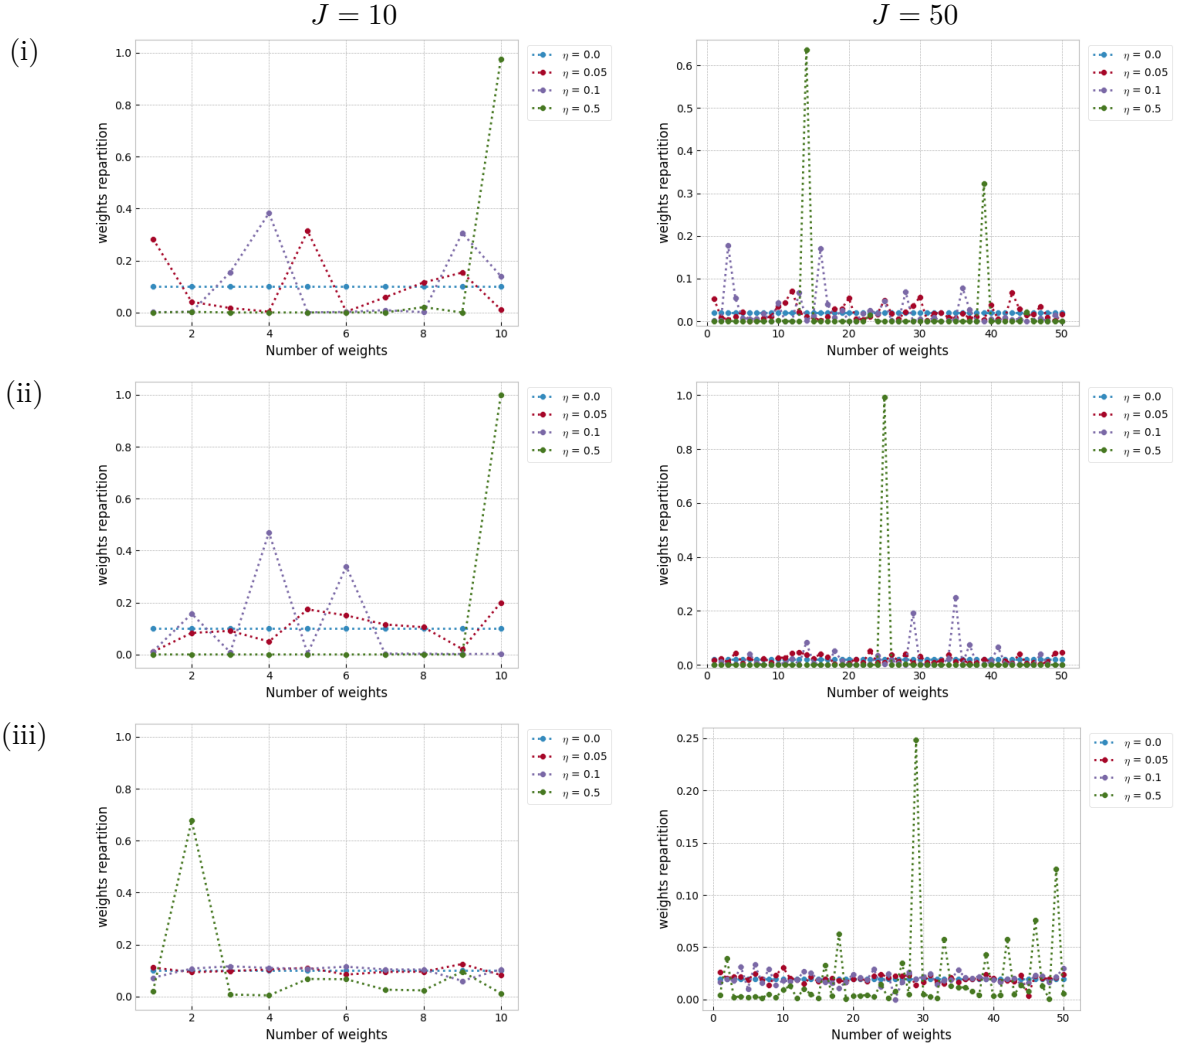


Figure 3: Final mixture weights λ_N for one run of the MG-IS-unif(0.5) algorithm (varying mixture weights) when considering each of the target distributions (i), (ii) and (iii).

the choice of the learning rate η . Indeed, choosing η too large might lead to missing some of the modes while having η too small may not discriminate quickly enough between the components. Observe in particular that the number of components J too plays a role in how fast the components selection process occurs, as for the same value of η , this selection process is slower as J increases.

8.2 Bayesian Logistic Regression

We consider the Bayesian Logistic Regression setting from Gershman et al. (2012). Namely, the data $\mathcal{D} = \{\mathbf{c}, \mathbf{x}\}$ is made of I binary class labels, $c_i \in \{-1, 1\}$, and of L covariates for each datapoint, $\mathbf{x}_i \in \mathbb{R}^L$. The latent variable $\mathbf{y} = \{\mathbf{w}, \beta\}$ consists of L regression coefficients

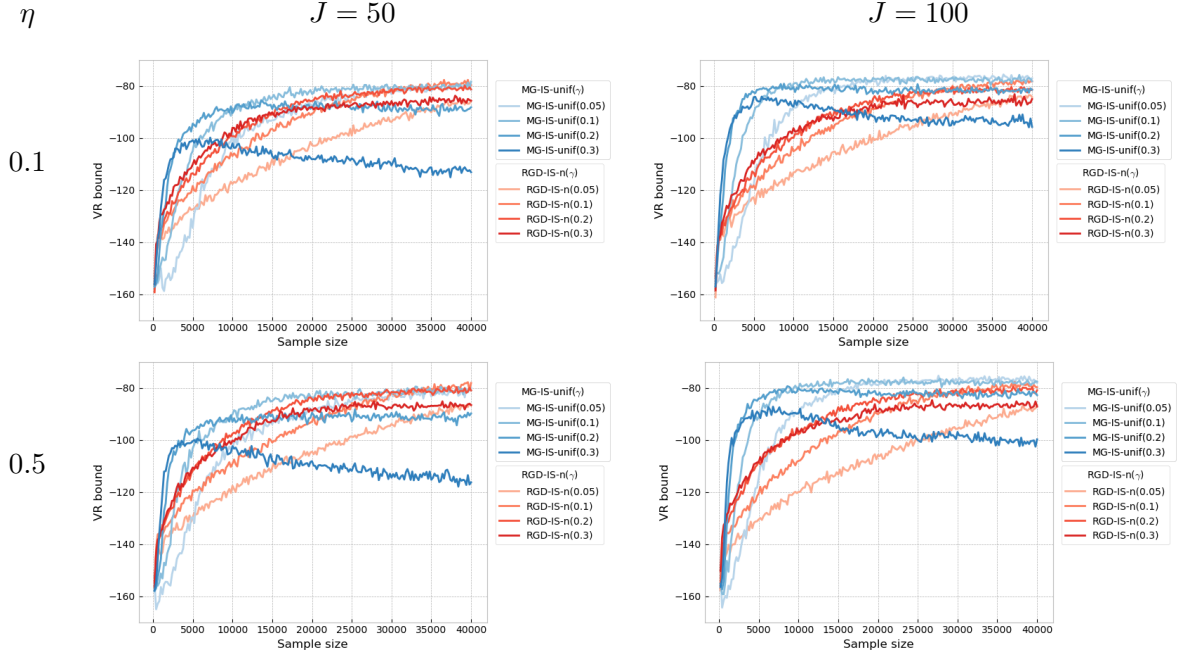


Figure 4: Monte Carlo estimate of the VR Bound for the RGD and the MG approaches when considering the Bayesian Logistic Regression on the Covertypes data set.

$w_l \in \mathbb{R}$, and a precision parameter $\beta \in \mathbb{R}^+$. Furthermore, the following model is assumed

$$p_0(\beta) = \text{Gamma}(\beta; a, b), \quad p_0(\mathbf{w}|\beta) = \mathcal{N}(\mathbf{w}; 0, \beta^{-1} \mathbf{I}_L), \quad p(c_i = 1|\mathbf{x}_i, \mathbf{w}) = \frac{1}{1 + e^{-\mathbf{w}^T \mathbf{x}_i}},$$

where $a = 1$ and $b = 0.01$, so that $p(y, \mathcal{D}) = p_0(\beta)p_0(\mathbf{w}|\beta) \prod_{i=1}^I p(c_i|\mathbf{x}_i, y)$. We now want to demonstrate the practicability of our framework in a real data scenario. To this end, we select the *Covertypes* data set (581,012 data points and 54 features, available at <https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/binary.html>) and we compare the RGD and the MG approaches for this choice of target density.

Implementation details. The covariance matrices of the mixture components are fixed and equal to $\sigma^2 \mathbf{I}_d$ with $\sigma^2 = 1$, $\alpha = 0.2$, $J \in \{50, 100\}$, $d = 56$, $M = 200$, the total number of iterations N is equal to 200, Θ_1 is generated by sampling from a centered normal distribution with covariance matrix $5\mathbf{I}_d$, $\lambda_1 = [1/J, \dots, 1/J]$ and for all time $n = 1 \dots N$, $\kappa_n = 0$, $\eta_n = \eta$, $\gamma_n = \gamma$ where $\eta \in \{0.1, 0.5\}$ and $\gamma \in \{0.05, 0.1, 0.2, 0.3\}$. Computing $p(y, \mathcal{D})$ constitutes the major computation bottleneck here since $I = 581,012$. To address this problem, $p(y, \mathcal{D})$ is approximated according to Remark 3 with subsampled mini-batches of size 100. The experiments are replicated independently 30 times and the convergence of the RGD and of the MG approaches is monitored by computing a Monte Carlo estimator of the VR Bound.

Our results are plotted in Figure 4, in which we focus on RGD-IS-n(γ) and MG-IS-unif(γ) (those two versions of Algorithm 3 are the most interesting to compare in this particular setting since MG-IS-n(γ) enjoys similar performances to RGD-IS-n(γ) and RGD-IS-unif(γ)).

underperforms, see Figure 9 of Appendix D for details). We observe that both algorithms are able to learn in a real data scenario for a proper tuning of (η, γ) and that selecting either η or γ too small/large can deteriorate the performance (as per the learning rate behaviour associated to those parameters). Furthermore, MG-IS-unif(γ) can be tuned to outperform RGD-IS-n(γ), which confirms our previous empirical findings regarding the relevance of the novel MG approach compared to the more traditional RGD one.

9. Conclusion

We introduced a novel methodology to build algorithms ensuring a monotonic decrease in the α -divergence at each step. Our methodology enabled simultaneous updates for both the weights and components parameters of a given mixture model, making it suitable for capturing complex multimodal target densities. Our work also connected and improved on different approaches: Gradient Descent for α -divergence minimisation, Power Descent and an Integrated EM algorithm. By investigating variational families based on the exponential family, we applied our framework to important classes of models such as Gaussian Mixture Models and Student’s t Mixture Models. Finally, we provided empirical evidence that our methodology can be used to enhance the aforementioned existing algorithms.

To conclude, we state several directions to extend our work. Now that we have established a systematic decrease for our iterative schemes, the next step could be to derive convergence results and to compare them with those obtained using typical Gradient Descent schemes. Based on Proposition 2, another direction is to generalise the monotonicity property from Theorem 3 beyond the case $\alpha \in [0, 1)$. Lastly, much like it is already the case in traditional Black-Box Variational Inference, we expect that fine-tuning our hyperparameters and resorting to more advanced Monte Carlo methods in the estimation of the intractable integrals appearing in our ideal algorithms will lead to further improved numerical results.

Acknowledgments and Disclosure of Funding

We would like to thank the action editor and the reviewers for helpful comments and suggestions on the paper. Kamélia Daudel acknowledges support of the UK Defence Science and Technology Laboratory (Dstl) and Engineering and Physical Research Council (EPSRC) under grant EP/R013616/1. This is part of the collaboration between US DOD, UK MOD and UK EPSRC under the Multidisciplinary University Research Initiative.

A. Deferred Proofs and Results for Section 3

A.1 Quantifying the Improvement in one Step of Gradient Descent

Let $\mathbb{T} \subseteq \mathbb{R}^d$ be a convex set. Here, $\langle \cdot, \cdot \rangle$ is the standard inner product on \mathbb{R}^d and $\|\cdot\|$ is the Euclidean norm.

Definition 3 *A continuously differentiable function g defined on \mathbb{T} is said to be β -smooth if for all $\theta, \theta' \in \mathbb{T}$,*

$$\|\nabla g(\theta) - \nabla g(\theta')\| \leq \beta \|\theta - \theta'\| .$$

Lemma 2 *Let $\gamma \in (0, 1]$, let g be a β -smooth function defined on \mathbb{T} . Then, for all $\theta \in \mathbb{T}$ it holds that*

$$g(\theta) - g\left(\theta - \frac{\gamma}{\beta} \nabla g(\theta)\right) \geq \frac{\gamma}{2\beta} \|\nabla g(\theta)\|^2.$$

Proof By assumption on g , we have that for all $\theta, \theta' \in \mathbb{T}$

$$g(\theta') - g(\theta) - \langle \nabla g(\theta), \theta' - \theta \rangle \leq \frac{\beta}{2} \|\theta - \theta'\|^2.$$

In particular, setting $\theta' = \theta - \frac{\gamma}{\beta} \nabla g(\theta)$ yields

$$\begin{aligned} g(\theta) - g\left(\theta - \frac{\gamma}{\beta} \nabla g(\theta)\right) &\geq \frac{\gamma}{\beta} \|\nabla g(\theta)\|^2 - \frac{\gamma^2}{2\beta} \|\nabla g(\theta)\|^2 \\ &\geq \frac{\gamma}{\beta} \left(1 - \frac{\gamma}{2}\right) \|\nabla g(\theta)\|^2. \end{aligned}$$

Since $\gamma \in (0, 1]$, we can deduce the desired result. ■

A.2 Gradient Descent for α - / Rényi's α -divergence Minimisation in Section 3

Let us first write the definition of the α -divergence (resp. of Rényi's α -divergence) between the two absolutely continuous probability measures $K(\theta, \cdot)$ and \mathbb{P}

$$\begin{aligned} D_\alpha(K(\theta, \cdot) || \mathbb{P}) &= \int_{\mathcal{Y}} \frac{1}{\alpha(\alpha-1)} \left[\left(\frac{k(\theta, y)}{p(y|\mathcal{D})} \right)^\alpha - 1 \right] p(y|\mathcal{D}) \nu(dy) \\ D_\alpha^{(\text{AR})}(K(\theta, \cdot) || \mathbb{P}) &= \frac{1}{\alpha(\alpha-1)} \log \left(\int_{\mathcal{Y}} k(\theta, y)^\alpha p(y|\mathcal{D})^{1-\alpha} \nu(dy) \right). \end{aligned}$$

Here, the Rényi divergence is defined following the convention from Cichocki and Amari (2010), alternative definitions may use a different scaling factor.

A.2.1 GRADIENT DESCENT FOR α -DIVERGENCE MINIMISATION

As underlined in the introduction, minimising the α -divergence $D_\alpha(\mathbb{Q} || \mathbb{P})$ w.r.t q is equivalent to minimising $\Psi_\alpha(q; p)$ with $p = p(\cdot, \mathcal{D})$ w.r.t q , where we have gotten rid of $p(\mathcal{D})$ in the optimisation problem as written in (2). Letting \mathcal{Q} be of the form $\mathcal{Q} = \{q : y \mapsto k(\theta, y) : \theta \in \mathbb{T}\}$, the traditional Variational Inference way to optimise $\Psi_\alpha(k(\theta, \cdot); p)$ with $p = p(\cdot, \mathcal{D})$ w.r.t θ corresponds to performing Gradient Descent steps on θ to construct a sequence $(\theta_n)_{n \geq 1}$ that converges towards a local optimum of the function $\theta \mapsto \Psi_\alpha(k(\theta, \cdot); p)$. This procedure involves a well-chosen learning rate policy $(r_n)_{n \geq 1}$ and sets:

$$\theta_{n+1} = \theta_n - r_n \nabla \Psi_\alpha(k(\theta_n, \cdot); p), \quad n \geq 1. \quad (79)$$

Under common differentiability assumptions,

$$\begin{aligned}
 \nabla \Psi_\alpha(k(\theta_n, \cdot); p) &= \nabla \left(\int_Y f_\alpha \left(\frac{k(\theta_n, y)}{p(y)} \right) p(y) \nu(dy) \right) \\
 &= \int_Y \frac{\partial}{\partial \theta} \left(f_\alpha \left(\frac{k(\theta, y)}{p(y)} \right) \right) \Big|_{(\theta, y) = (\theta_n, y)} p(y) \nu(dy) \\
 &= \int_Y f'_\alpha \left(\frac{k(\theta_n, y)}{p(y)} \right) \frac{\partial k(\theta, y)}{\partial \theta} \Big|_{(\theta, y) = (\theta_n, y)} \nu(dy) \\
 &= \int_Y \frac{k(\theta_n, y)^\alpha p(y)^{1-\alpha}}{\alpha - 1} \frac{\partial \log k(\theta, y)}{\partial \theta} \Big|_{(\theta, y) = (\theta_n, y)} \nu(dy)
 \end{aligned}$$

so that (79) becomes

$$\theta_{n+1} = \theta_n - r_n \int_Y \frac{\varphi_n^{(\alpha)}(y)}{\alpha - 1} \frac{\partial \log k(\theta, y)}{\partial \theta} \Big|_{(\theta, y) = (\theta_n, y)} \nu(dy), \quad n \geq 1.$$

A.2.2 GRADIENT DESCENT FOR RÉNYI'S α -DIVERGENCE MINIMISATION

Considering yet again $\mathcal{Q} = \{q : y \mapsto k(\theta, y) : \theta \in \mathbb{T}\}$, minimising Rényi's α -divergence

$$D_\alpha^{(\text{AR})}(K(\theta, \cdot) || \mathbb{P}) = \frac{1}{\alpha(\alpha - 1)} \log \left(\int_Y k(\theta, y)^\alpha p(y|\mathcal{D})^{1-\alpha} \nu(dy) \right)$$

w.r.t θ can be done by performing Gradient Descent steps on θ to construct a sequence $(\theta_n)_{n \geq 1}$ that converges towards a local optimum of the function $\theta \mapsto D_\alpha^{(\text{AR})}(K(\theta, \cdot) || \mathbb{P})$. This procedure involves a well-chosen learning rate policy $(r_n)_{n \geq 1}$ and sets:

$$\theta_{n+1} = \theta_n - r_n \nabla D_\alpha^{(\text{AR})}(K(\theta_n, \cdot) || \mathbb{P}), \quad n \geq 1. \tag{80}$$

Under common differentiability assumptions and setting $p = p(\cdot, \mathcal{D})$,

$$\begin{aligned}
 \nabla D_\alpha^{(\text{AR})}(K(\theta_n, \cdot) || \mathbb{P}) &= \nabla \left(\frac{1}{\alpha(\alpha - 1)} \log \left(\int_Y k(\theta_n, y)^\alpha p(y|\mathcal{D})^{1-\alpha} \nu(dy) \right) \right) \\
 &= \frac{1}{\alpha(\alpha - 1)} \frac{\nabla \left(\int_Y k(\theta_n, y)^\alpha p(y, \mathcal{D})^{1-\alpha} \nu(dy) \right)}{\int_Y k(\theta_n, y)^\alpha p(y, \mathcal{D})^{1-\alpha} \nu(dy)} \\
 &= \int_Y \frac{\check{\varphi}_n^{(\alpha)}(y)}{\alpha - 1} \frac{\partial \log k(\theta, y)}{\partial \theta} \Big|_{(\theta, y) = (\theta_n, y)} \nu(dy)
 \end{aligned}$$

so that (80) becomes

$$\theta_{n+1} = \theta_n - r_n \int_Y \frac{\check{\varphi}_n^{(\alpha)}(y)}{\alpha - 1} \frac{\partial \log k(\theta, y)}{\partial \theta} \Big|_{(\theta, y) = (\theta_n, y)} \nu(dy), \quad n \geq 1.$$

B. Deferred Proofs and Results for Section 4

B.1 Proof of Corollary 4

Proof of Corollary 4 Following the proof of Corollary 1, we obtain that (19) holds by using the definition of θ_{n+1} combined with the fact that $\alpha \in [0, 1)$, $D_1(K(\theta_{j,n}, \cdot) || K(\theta_{j,n+1}, \cdot)) \geq 0$

and $\lambda_{j,n} > 0$ for all $j = 1 \dots J$. (18) holds by Theorem 3 and we can thus apply Theorem 2. ■

B.2 Extension of Lemma 1 to Mixture Models

Lemma 3 (Generalised maximisation approach for the mean-field family) *Let each component of the mixture be a member of the same mean-field variational family so that $k(\theta_j, y) = \prod_{\ell=1}^L k^{(\ell)}(\theta_j^{(\ell)}, y^{(\ell)})$ with $\theta_j = (\theta_j^{(1)}, \dots, \theta_j^{(L)}) \in \mathbb{T}$. Then, starting from $\Theta_1 \in \mathbb{T}$ and denoting $\Theta_n = (\theta_{1,n}, \dots, \theta_{J,n})$ with $\theta_{j,n} = (\theta_{j,n}^{(1)}, \dots, \theta_{j,n}^{(L)})$ for all $n \geq 1$ and all $j = 1 \dots J$, solving (25) yields the following update formulas: for all $n \geq 1$ and all $j = 1 \dots J$,*

$$\theta_{j,n+1}^{(\ell)} = \operatorname{argmax}_{\theta^{(\ell)}} \int_{\mathcal{Y}} \left[\varphi_{j,n}^{(\alpha)}(y) + b_{j,n} k(\theta_n, y) \right] \log \left(\frac{k^{(\ell)}(\theta^{(\ell)}, y^{(\ell)})}{k^{(\ell)}(\theta_{j,n}^{(\ell)}, y^{(\ell)})} \right) \nu(dy), \quad \ell = 1 \dots L.$$

B.3 Proof of Corollary 5

Proof of Corollary 5 Following the proof of Corollary 2, we will use that $\gamma_{j,n} \in (0, 1]$ and that $g_{j,n}$ is a $\beta_{j,n}$ -smooth function. Indeed, we can thus apply Lemma 2 and we obtain by definition of $\theta_{j,n+1}$ in (26) that $0 = g_{j,n}(\theta_{j,n}) \geq g_{j,n}(\theta_{j,n+1})$ for all $n \geq 1$ and all $j = 1 \dots J$, which in turn implies (19). In addition, (18) holds by Theorem 3, hence we can apply Theorem 2. ■

B.4 Gradient Descent for α - / Rényi's α -divergence Minimisation in Section 4

We obtain the desired results by adapting the reasoning from Appendix A.2 to the case

$$\mathcal{Q} = \left\{ q : y \mapsto \mu_{\lambda, \Theta} k(y) = \sum_{j=1}^J \lambda_j k(\theta_j, y) : \Theta \in \mathbb{T}^J \right\},$$

with $\lambda \in \mathcal{S}_J$, meaning that we consider the α -divergence (resp. Rényi's α -divergence) between the two absolutely continuous probability measures $\mu_{\lambda, \Theta} K$ and \mathbb{P} :

$$D_{\alpha}(\mu_{\lambda, \Theta} K \| \mathbb{P}) = \int_{\mathcal{Y}} \frac{1}{\alpha(\alpha-1)} \left[\left(\frac{\mu_{\lambda, \Theta} k(y)}{p(y|\mathcal{D})} \right)^{\alpha} - 1 \right] p(y|\mathcal{D}) \nu(dy)$$

$$D_{\alpha}^{(\text{AR})}(\mu_{\lambda, \Theta} K \| \mathbb{P}) = \frac{1}{\alpha(\alpha-1)} \log \left(\int_{\mathcal{Y}} (\mu_{\lambda, \Theta} k(y))^{\alpha} p(y|\mathcal{D})^{1-\alpha} \nu(dy) \right),$$

(we use the convention from Cichocki and Amari (2010) for Rényi's α -divergence).

B.5 Monotonicity Property for the Power Descent

B.5.1 PRELIMINARY REMARKS

For convenience, we redefine in this section the function $b_{\mu, \alpha}$ for all $\mu \in \mathcal{M}_1(\mathbb{T})$ by

$$b_{\mu, \alpha}(\theta) = \int_{\mathcal{Y}} k(\theta, y) \frac{1}{\alpha-1} \left(\frac{\mu k(y)}{p(y)} \right)^{\alpha-1} \nu(dy), \quad \theta \in \mathbb{T}.$$

Then, for all $\eta > 0$, the iteration $\mu \mapsto \mathcal{I}_\alpha(\mu)$ is well-defined if we have

$$0 < \mu(|b_{\mu,\alpha} + \kappa|^{\frac{\eta}{1-\alpha}}) < \infty. \quad (81)$$

Furthermore, Daudel et al. (2021) already established that one transition of the Power Descent algorithm ensures a monotonic decrease in the α -divergence at each step for all $\eta \in (0, 1]$ and all κ such that $(\alpha - 1)\kappa \geq 0$ under the assumption of Proposition 2, which settles the case (iii). Finally, while we establish our results for (i) and (ii) in the general case where $\mu \in \mathcal{M}_1(\mathcal{T})$, the particular case of mixture models follows immediately by choosing μ as a weighted sum of dirac measures.

B.5.2 EXTENDING THE MONOTONICITY

Let (ζ, μ) be a couple of probability measures where ζ is dominated by μ , which we denote by $\zeta \preceq \mu$. A first lower-bound for the difference $\Psi_\alpha(\mu k) - \Psi_\alpha(\zeta k)$ was derived in Daudel et al. (2021) and was used to establish that the Power Descent algorithm diminishes Ψ_α for all $\eta \in (0, 1]$. We now prove a novel lower-bound for the difference $\Psi_\alpha(\mu k) - \Psi_\alpha(\zeta k)$ which will allow us to extend the monotonicity results from Daudel et al. (2021) beyond the case $\eta \in (0, 1]$ when $\alpha < 0$. This result relies on the existence of an exponent ϱ satisfying condition (A2) below, which will later on be used to specify a range of values for η ensuring that Ψ_α is decreasing after having applied one transition $\mu \mapsto \mathcal{I}_\alpha(\mu)$

(A2) We have $\varrho \in \mathbb{R} \setminus [0, 1]$ and the function $f_{\alpha,\varrho} : u \mapsto f_\alpha(u^{1/\varrho})$ is non-decreasing and concave on $\mathbb{R}_{>0}$.

Proposition 5 Assume (A1). Let $\alpha \in \mathbb{R} \setminus \{1\}$, assume that ϱ satisfies (A2) and let κ be such that $(\alpha - 1)\kappa \geq 0$. Then, for all $\mu, \zeta \in \mathcal{M}_1(\mathcal{T})$ such that $\mu(|b_{\mu,\alpha}|) < \infty$ and $\zeta \preceq \mu$,

$$|\varrho|^{-1} \{ \mu(|b_{\mu,\alpha} + \kappa|) - \mu(|b_{\mu,\alpha} + \kappa|g^\varrho) \} \leq \Psi_\alpha(\mu k) - \Psi_\alpha(\zeta k), \quad (82)$$

where g is the density of ζ wrt μ , i.e. $\zeta(d\theta) = \mu(d\theta)g(\theta)$. Moreover, equality holds in (82) if and only if $\zeta = \mu$.

Proof First note that for all $\alpha \in \mathbb{R} \setminus \{1\}$, we have by (A2) that $f'_{\alpha,\varrho}(u) \geq 0$ for all $u > 0$, and thus that $sg(\varrho) = sg(\alpha - 1)$ where $sg(v) = 1$ if $v \geq 0$ and -1 otherwise. Since $sg(f'_\alpha(u)) = sg(\alpha - 1) = sg(\kappa)$ for all $u > 0$, this implies that $\varrho^{-1}f'_\alpha(u) = |\varrho|^{-1}|f'_\alpha(u)|$, $\varrho^{-1}\kappa = |\varrho|^{-1}\kappa$ and finally that $\varrho^{-1}(b_{\mu,\alpha}(\theta) + \kappa) = |\varrho|^{-1}|b_{\mu,\alpha}(\theta) + \kappa|$ for all $\theta \in \mathcal{T}$, which will be used later in the proof.

Write by definition of $f_{\alpha,\varrho}$ in (A2) and ζ ,

$$\begin{aligned} \Psi_\alpha(\zeta k) &= \int_{\mathcal{Y}} f_\alpha \left(\frac{\zeta k(y)}{p(y)} \right) p(y) \nu(dy) \\ &= \int_{\mathcal{Y}} f_{\alpha,\varrho} \left(\left[\frac{\zeta k(y)}{p(y)} \right]^\varrho \right) p(y) \nu(dy) \\ &= \int_{\mathcal{Y}} f_{\alpha,\varrho} \left(\left[\int_{\mathcal{T}} \mu(d\theta) \frac{k(\theta, y)}{\mu k(y)} \left(\frac{g(\theta) \mu k(y)}{p(y)} \right) \right]^\varrho \right) p(y) \nu(dy) \\ &\leq \int_{\mathcal{Y}} f_{\alpha,\varrho} \left(\int_{\mathcal{T}} \mu(d\theta) \frac{k(\theta, y)}{\mu k(y)} \left(\frac{g(\theta) \mu k(y)}{p(y)} \right)^\varrho \right) p(y) \nu(dy) \end{aligned} \quad (83)$$

where the last inequality follows from Jensen's inequality applied to the convex function $u \mapsto u^\varrho$ (since $\varrho \in \mathbb{R} \setminus [0, 1]$) and the fact that $f_{\alpha, \varrho}$ is non-decreasing. Now set

$$\begin{aligned} u_y &= \int_{\mathbb{T}} \mu(d\theta) \frac{k(\theta, y)}{\mu k(y)} \left(\frac{g(\theta) \mu k(y)}{p(y)} \right)^\varrho \\ v_y &= \left(\frac{\mu k(y)}{p(y)} \right)^\varrho \end{aligned}$$

and note that

$$u_y - v_y = \left(\frac{\mu k(y)}{p(y)} \right)^\varrho \left(\int_{\mathbb{T}} \mu(d\theta) \frac{k(\theta, y)}{\mu k(y)} g^\varrho(\theta) - 1 \right). \quad (84)$$

Since $f_{\alpha, \varrho}$ is concave, $f_{\alpha, \varrho}(u_y) \leq f_{\alpha, \varrho}(v_y) + f'_{\alpha, \varrho}(v_y)(u_y - v_y)$. Combining with (83), we get

$$\begin{aligned} \Psi_\alpha(\zeta k) &\leq \int_Y f_{\alpha, \varrho}(u_y) p(y) \nu(dy) \\ &\leq \int_Y f_{\alpha, \varrho}(v_y) p(y) \nu(dy) + \int_Y f'_{\alpha, \varrho}(v_y)(u_y - v_y) p(y) \nu(dy) \end{aligned} \quad (85)$$

Note that the first term of the rhs can be written as

$$\int_Y f_{\alpha, \varrho}(v_y) p(y) \nu(dy) = \int_Y f_\alpha \left(\frac{\mu k(y)}{p(y)} \right) p(y) \nu(dy) = \Psi_\alpha(\mu k). \quad (86)$$

Using now $f'_{\alpha, \varrho}(v_y) = \varrho^{-1} v_y^{1/\varrho - 1} f'_\alpha(v_y^{1/\varrho})$ and (84), the second term of the rhs of (85) may be expressed as

$$\begin{aligned} &\int_Y f'_{\alpha, \varrho}(v_y)(u_y - v_y) p(y) \nu(dy) \\ &= \varrho^{-1} \int_Y \left(\frac{\mu k(y)}{p(y)} \right)^{1-\varrho} f'_\alpha \left(\frac{\mu k(y)}{p(y)} \right) \\ &\quad \left(\frac{\mu k(y)}{p(y)} \right)^\varrho \left(\int_{\mathbb{T}} \mu(d\theta) \frac{k(\theta, y)}{\mu k(y)} g^\varrho(\theta) - 1 \right) p(y) \nu(dy) \\ &= \varrho^{-1} \int_{\mathbb{T}} \mu(d\theta) \left(\int_Y k(\theta, y) f'_\alpha \left(\frac{\mu k(y)}{p(y)} \right) \nu(dy) \right) g^\varrho(\theta) \\ &\quad - \varrho^{-1} \int_Y \mu k(y) f'_\alpha \left(\frac{\mu k(y)}{p(y)} \right) \nu(dy) \\ &= \varrho^{-1} \{ \mu(b_{\mu, \alpha} \cdot g^\varrho) - \mu(b_{\mu, \alpha}) \} \\ &= |\varrho|^{-1} \{ \mu(|b_{\mu, \alpha} + \kappa| g^\varrho) - \mu(|b_{\mu, \alpha} + \kappa|) \} + |\varrho|^{-1} \kappa |1 - \mu(g^\varrho)|, \end{aligned}$$

where we have used that $\varrho^{-1}(b_{\mu, \alpha}(\theta) + \kappa) = |\varrho|^{-1} |b_{\mu, \alpha}(\theta) + \kappa|$ for all $\theta \in \mathbb{T}$ and that $\varrho^{-1} \kappa = |\varrho|^{-1} \kappa$. In addition, Jensen's inequality applied to the convex function $u \mapsto u^\varrho$ implies that $\mu(g^\varrho) \geq 1$ and thus

$$\int_Y f'_{\alpha, \varrho}(v_y)(u_y - v_y) p(y) \nu(dy) \leq |\varrho|^{-1} \{ \mu(|b_{\mu, \alpha} + \kappa| g^\varrho) - \mu(|b_{\mu, \alpha} + \kappa|) \}. \quad (87)$$

Combining this inequality with (85) and (86) finishes the proof of the inequality. Furthermore, if the equality holds in (82), then the equality in Jensen's inequality (87) shows that g is constant μ -a.e. so that $\zeta = \mu$, and the proof is completed. \blacksquare

Remark 4 *The proof of Proposition 5 relies on f'_α being of constant sign. Notice however that the definition of the α -divergence in (1) is invariant with respect to the transformation $\tilde{f}_\alpha(u) = f_\alpha(u) + \kappa(u - 1)$ for any arbitrary constant κ , that is f_α can be equivalently replaced by \tilde{f}_α in (1). This aspect is in fact expressed through the constant κ appearing in the update formula, that we however need to assume to satisfy $(\alpha - 1)\kappa \geq 0$ in our proofs.*

We now plan on setting $\zeta = \mathcal{I}_\alpha(\mu)$ in Proposition 5 and obtain that one iteration of the Power Descent yields $\Psi_\alpha(\mathcal{I}_\alpha(\mu)k) \leq \Psi_\alpha(\mu k)$. For this purpose and based on the upper bound obtained in Proposition 5, we strengthen the condition (81) as follows to take into account the exponent ϱ

$$0 < \mu(|b_{\mu,\alpha} + \kappa|^{\frac{\eta}{1-\alpha}}) < \infty \text{ and } \mu(|b_{\mu,\alpha} + \kappa|g^\varrho) \leq \mu(|b_{\mu,\alpha} + \kappa|) \quad \text{with } g = \frac{|b_{\mu,\alpha} + \kappa|^{\frac{\eta}{1-\alpha}}}{\mu(|b_{\mu,\alpha} + \kappa|^{\frac{\eta}{1-\alpha}})} . \quad (88)$$

This leads to the following result.

Proposition 6 *Assume (A1). Let $\alpha \in \mathbb{R} \setminus \{1\}$, assume that ϱ satisfies (A2) and let κ be such that $(\alpha - 1)\kappa \geq 0$. Let $\mu \in \mathcal{M}_1(\mathcal{T})$ be such that $\mu(|b_{\mu,\alpha}|) < \infty$ and assume that η satisfies (88). Then, the two following assertions hold.*

- (i) *We have $\Psi_\alpha(\mathcal{I}_\alpha(\mu)k) \leq \Psi_\alpha(\mu k)$.*
- (ii) *We have $\Psi_\alpha(\mathcal{I}_\alpha(\mu)k) = \Psi_\alpha(\mu k)$ if and only if $\mu = \mathcal{I}_\alpha(\mu)$.*

Proof We treat the case $\kappa = 0$ in the proof below (the case $\kappa \neq 0$ unfolds similarly). We apply Proposition 5 with $\zeta = \mathcal{I}_\alpha(\mu)$ so that $\zeta(d\theta) = \mu(d\theta)g(\theta)$ with $g = |b_{\mu,\alpha}|^{\eta/(1-\alpha)} / \mu(|b_{\mu,\alpha}|^{\eta/(1-\alpha)})$. Then,

$$\Psi_\alpha(\mathcal{I}_\alpha(\mu)k) \leq \Psi_\alpha(\mu k) + |\varrho|^{-1} \{ \mu(|b_{\mu,\alpha}|g^\varrho) - \mu(|b_{\mu,\alpha}|) \} \leq \Psi_\alpha(\mu k) , \quad (89)$$

where the last inequality follows from condition (88). Let us now show (ii). The *if* part is obvious. As for the *only if* part, $\Psi_\alpha(\mathcal{I}_\alpha(\mu)k) = \Psi_\alpha(\mu k)$ combined with (89) yields

$$\Psi_\alpha(\mathcal{I}_\alpha(\mu)k) = \Psi_\alpha(\mu k) + |\varrho|^{-1} \{ \mu(|b_{\mu,\alpha}|g^\varrho) - \mu(|b_{\mu,\alpha}|) \} ,$$

which is the case of equality in Proposition 5. Therefore, $\mathcal{I}_\alpha(\mu) = \mu$. \blacksquare

While Proposition 6 resembles (Daudel et al., 2021, Theorem 1) in its formulation and in the properties on the iteration $\mu \mapsto \mathcal{I}_\alpha(\mu)$ it establishes, it is important to note that the proof techniques used, and thus the conditions on η obtained, are different. This brings us to the proof of Proposition 2. The proof of this theorem requires intermediate results, which are derived in Appendix B.5.3 alongside the proof of Proposition 2.

B.5.3 PROOF OF PROPOSITION 2

For the sake of readability, we only treat the case $\kappa = 0$ in the proofs below (the case $\kappa \neq 0$ unfolds similarly). In Proposition 5, the difference $\Psi_\alpha(\zeta k) - \Psi_\alpha(\mu k)$ is split into two terms

$$\Psi_\alpha(\zeta k) - \Psi_\alpha(\mu k) = A(\mu, \zeta) + |\varrho|^{-1} \{ \mu(|b_{\mu,\alpha}|g^\varrho) - \mu(|b_{\mu,\alpha}|) \} ,$$

where $g = d\zeta/d\mu$. Moreover, Proposition 5 states that $A(\mu, \zeta)$ is always non-positive. It turns out that the second term is minimal over all positive probability densities g when it is proportional to $|b_{\mu,\alpha}|^{1/(1-\varrho)}$, as we show in Lemma 4 below.

Lemma 4 *Let $\varrho \in \mathbb{R} \setminus [0, 1]$. Then, for any positive probability density g w.r.t μ , we have*

$$\mu(|b_{\mu,\alpha}|g^\varrho) \geq \left[\mu(|b_{\mu,\alpha}|^{1/(1-\varrho)}) \right]^{1-\varrho} ,$$

with equality if and only if $g \propto |b_{\mu,\alpha}|^{1/(1-\varrho)}$.

Proof The function $x \mapsto x^{1-\varrho}$ is strictly convex for $\varrho \in \mathbb{R} \setminus [0, 1]$. Thus Jensen's inequality yields, for any positive probability density g w.r.t. μ ,

$$\mu(|b_{\mu,\alpha}|g^\varrho) = \int_{\mathbb{T}} \mu(d\theta) \left(\frac{|b_{\mu,\alpha}(\theta)|^{1/(1-\varrho)}}{g(\theta)} \right)^{1-\varrho} g(\theta) \geq \left[\mu(|b_{\mu,\alpha}|^{1/(1-\varrho)}) \right]^{1-\varrho} \quad (90)$$

which finishes the proof of the inequality. The next statement follows from the case of equality in Jensen's inequality: g must be proportional to $|b_{\mu,\alpha}|^{1/(1-\varrho)}$. \blacksquare

The next lemma shows that this choice leads to a non-positive second term, thus implying that $\Psi_\alpha(\zeta k) \leq \Psi_\alpha(\mu k)$.

Lemma 5 *Assume (A1). Let $\alpha \in \mathbb{R} \setminus \{1\}$ and assume that ϱ satisfies (A2). Then $\eta = (1 - \alpha)/(1 - \varrho)$ satisfies (88) for any $\mu \in M_1(\mathbb{T})$ such that $\mu(|b_{\mu,\alpha}|) < \infty$.*

Proof We apply (90) with $g = 1$ and get that

$$\left[\mu(|b_{\mu,\alpha}|^{1/(1-\varrho)}) \right]^{1-\varrho} \leq \mu(|b_{\mu,\alpha}|) < \infty . \quad (91)$$

Then (88) can be readily checked with $\eta = (1 - \alpha)/(1 - \varrho)$. Set $\phi = \eta/(1 - \alpha)$. Using that $\mu(|b_{\mu,\alpha}|) < \infty$ when $\phi < 0$ and (A1) for $\phi > 0$, we obtain $\mu(|b_{\mu,\alpha}|^\phi) > 0$, which concludes the proof. \blacksquare

While Lemma 5 seems to advocate for $g = d\zeta/d\mu$ to be proportional to $|b_{\mu,\alpha}|^{1/(1-\varrho)}$, notice that this choice of g might not be optimal to minimize $\Psi_\alpha(\zeta k) - \Psi_\alpha(\mu k)$, as $A(\mu, \zeta)$ also depends on g through ζ . In the next lemma, we thus propose another choice of the tuning parameter η , which also satisfies (88) for any $\mu \in M_1(\mathbb{T})$ such that $\mu(|b_{\mu,\alpha}|) < \infty$.

Lemma 6 *Assume (A1). Let $\alpha \in \mathbb{R} \setminus \{1\}$ and assume that ϱ satisfies (A2). Let $\mu \in M_1(\mathbb{T})$ be such that $\mu(|b_{\mu,\alpha}|) < \infty$. Further assume that $|\varrho| \geq 1$, then $\eta = (\alpha - 1)/\varrho$ satisfies (88).*

Proof Setting $g \propto |b_{\mu,\alpha}|^{-1/\varrho}$, we get

$$\mu(|b_{\mu,\alpha}|g^\varrho) = \mu(|b_{\mu,\alpha}|^{1-\varrho/\varrho})[\mu(|b_{\mu,\alpha}|^{-1/\varrho})]^{-\varrho} = [\mu(|b_{\mu,\alpha}|^{-1/\varrho})]^{-\varrho} \leq \mu(|b_{\mu,\alpha}|)$$

where the last inequality follows from Jensen's inequality applied to the convex function $u \mapsto u^{-\varrho}$ (since $|\varrho| \geq 1$). Since $\mu(|b_{\mu,\alpha}|) < \infty$, the parameter $\eta = (\alpha - 1)/\varrho$ satisfies (88). Set $\phi = \eta/(1 - \alpha)$. Using that $\mu(|b_{\mu,\alpha}|) < \infty$ when $\phi < 0$ and (A1) for $\phi > 0$, we obtain $\mu(|b_{\mu,\alpha}|^\phi) > 0$, which concludes the proof. \blacksquare

Lemma 5 and Lemma 6 allow us to define a range of values for η that decreases Ψ_α after one transition of the Power Descent, under the assumption that ϱ satisfies (A2). Now, in order to prove Proposition 2 and given $\alpha \in \mathbb{R} \setminus \{1\}$, we need to check which values of ϱ satisfy the conditions expressed in (A2).

Proof of Proposition 2 The proof consists in verifying that we can apply Proposition 6, that is, given $\alpha \in \mathbb{R} \setminus \{1\}$, we must find a range of constants ϱ which satisfy (A2). We then use Lemma 5 or Lemma 6 to deduce that, for the provided constants η , (88) holds.

(i) Assumption (A2) holds for all $\varrho < 0$, with $f_{\alpha,\varrho}(u) = -\log(u)/\varrho$. Moreover, by definition of $b_{\mu,\alpha}$, we get for all $n \geq 1$,

$$\mu(|b_{\mu,\alpha}|) = \int_Y \mu k(y) \frac{p(y)}{\mu k(y)} \nu(dy) = \int_Y p(y) \nu(dy) < \infty.$$

Combining with Lemma 5 and Lemma 6, (88) holds for all $\mu \in M_1(T)$ and for any $\eta \in (0, 1]$.

(ii) Observing that for $\alpha \notin \{0, 1\}$,

$$f_{\alpha,\varrho}(u) = \frac{1}{\alpha(\alpha - 1)} \left(u^{\alpha/\varrho} - 1 \right),$$

we get that (A2) holds for $\varrho \leq \alpha$ if $\alpha < 0$. Lemmas 5 and 6 provide the corresponding ranges for η in Cases (i) and (ii). To finish the proof, we now show that for all $\mu \in M_1(T)$, $\mu(|b_{\mu,\alpha}|)$ is finite, so that Lemmas 5 and 6 can indeed be applied.

Since $uf'_\alpha(u) = \alpha f_\alpha(u) + 1/(\alpha - 1)$, we have, for all $n \geq 1$,

$$\begin{aligned} \mu(|b_{\mu,\alpha}|) &= \int_Y \left| \left(\frac{\mu k(y)}{p(y)} \right) f'_\alpha \left(\frac{\mu k(y)}{p(y)} \right) \right| p(y) \nu(dy) \\ &\leq |\alpha| \int_Y \left| f_\alpha \left(\frac{\mu k(y)}{p(y)} \right) \right| p(y) \nu(dy) + \frac{1}{|\alpha - 1|} \end{aligned} \quad (92)$$

Using that $\Psi_\alpha(\mu k) > -\infty$ (which is a consequence of (A1) and of Jensen's inequality applied to the convex function $u \mapsto uf_\alpha(1/u)$), the r.h.s is finite if and only if $\Psi_\alpha(\mu k)$ is finite, which is what we have assumed and thus the proof is finished. \blacksquare

C. Deferred Proofs and Results of Section 6

We start with a useful lemma.

C.1 A Useful Lemma

The following lemma will be used for the proofs of Theorems 4 and 5 and Proposition 3.

Lemma 7 *Let $k^{(o)}$ satisfy (B1). We have, for all $\zeta, \zeta' \in \text{Int}E_0$,*

$$\log \left(\frac{k^{(o)}(\zeta', y)}{k^{(o)}(\zeta, y)} \right) \geq \langle S(y) - \nabla A(\zeta'), \zeta' - \zeta \rangle_E, \quad \text{with equality if and only if } \zeta' = \zeta, \quad (93)$$

$$0 \geq \int_Y \log \left(\frac{k^{(o)}(\zeta', y)}{k^{(o)}(\zeta, y)} \right) k^{(o)}(\zeta, y) \nu(dy) \geq \langle \nabla A(\zeta) - \nabla A(\zeta'), \zeta' - \zeta \rangle_E. \quad (94)$$

Proof Let $\zeta, \zeta' \in \text{Int}E_0$. We have, since $\text{Int}E_0$ is convx,

$$\begin{aligned} \log \left(\frac{k^{(o)}(\zeta', y)}{k^{(o)}(\zeta, y)} \right) &= \langle S(y), \zeta' - \zeta \rangle_E + A(\zeta) - A(\zeta') \\ &= \langle S(y), \zeta' - \zeta \rangle_E + \langle \nabla A(\zeta'), \zeta - \zeta' \rangle_E \\ &\quad + (\zeta - \zeta')^T \left(\int_0^1 (1-t) \left[\nabla \nabla^T A(t\zeta + (1-t)\zeta') \right] dt \right) (\zeta - \zeta'). \end{aligned}$$

The inequality and the equality case of (93) then follow from the fact that the hessian $\nabla \nabla^T A$ is positive definite in $\text{Int}E_0$. Using (93) and (32), we further obtain the lower bound displayed in (94) on the negated exclusive Kullback-Leibler divergence between the two (absolutely continuous w.r.t. ν) probability distributions with probability density functions $k^{(o)}(\zeta, \cdot)$ and $k^{(o)}(\zeta', \cdot)$ respectively. \blacksquare

C.2 Proofs of Section 6.2

The following lemma will be useful.

Lemma 8 *Let p and g be functions measurable from (Y, \mathcal{Y}) to the Borel sets of \mathbb{R}_+ such that*

$$0 < \int_Y p(y) (1 + g(y))^{1/(1-\alpha)} \nu(dy) < \infty. \quad (95)$$

Let $\alpha \in [0, 1)$ and $k : \mathbb{T} \times Y \rightarrow \mathbb{R}_+$ be a positive kernel density with respect to ν . Let $\theta_0 \in \mathbb{T}$ and μ be a probability on \mathbb{T} such that

$$\inf_{y \in Y} \frac{k(\theta_0, y)}{\mu k(y)} > 0. \quad (96)$$

Then, setting

$$\varphi(y) = k(\theta_0, y)^\alpha \left(\frac{p(y)}{\mu k(y)} \right)^{1-\alpha},$$

we have $\int \varphi \, d\nu > 0$ and the probability density function $\check{\varphi} = \varphi / (\int \varphi \, d\nu)$ satisfies

$$\int \check{\varphi}(y) g(y) \nu(dy) < \infty . \quad (97)$$

Proof We have $\mu k > 0$ by assumption on k and μ . Thus φ is well defined and $\int \varphi \, d\nu > 0$ as a consequence of $p \geq 0$ with the left-hand side of (95). We now check (97) in which $\check{\varphi}$ can equivalently be replaced by φ . We have

$$\begin{aligned} \int \varphi(y) g(y) \nu(dy) &= \int k(\theta_0, y) \left(\frac{p(y)}{\mu k(y)} \right)^{1-\alpha} g(y) \nu(dy) \\ &\leq \left(\int k(\theta_0, y) \frac{p(y)}{\mu k(y)} g(y)^{1/(1-\alpha)} \nu(dy) \right)^\alpha , \end{aligned}$$

by Jensen's inequality and concavity of $x \mapsto x^{1-\alpha}$ for $\alpha \in [0, 1)$. Hence we obtain (97) as a consequence of (95) and (96). \blacksquare

Proof of Theorem 4 Let $\zeta_0 \in \mathcal{O}$ and $b \geq 0$. In the following, we denote by $\mathcal{C}_0(\zeta)$ the value of the integral in the argmax in (35). Using that $\check{\varphi}$ is a probability density function paired up with (32) and (34), we have

$$\begin{aligned} \int_Y [\check{\varphi}(y) + b k^{(o)}(\zeta_0, y)] \nu(dy) &= 1 + b > 0 , \\ \int_Y [\check{\varphi}(y) + b k^{(o)}(\zeta_0, y)] S(y) \nu(dy) &= \int_Y \check{\varphi} S \, d\nu + b \nabla A(\zeta_0) \in E . \end{aligned}$$

By (93) in Lemma 7, for all $\zeta' \neq \zeta$ in \mathcal{O} and all $y \in Y$, we have

$$\log \left(\frac{k^{(o)}(\zeta', y)}{k^{(o)}(\zeta, y)} \right) > \langle S(y) - \nabla A(\zeta'), \zeta' - \zeta \rangle_E .$$

Hence, combining with the three previous identities, we get that, for all $\zeta' \neq \zeta$ in \mathcal{O} ,

$$\begin{aligned} \mathcal{C}_0(\zeta') - \mathcal{C}_0(\zeta) &= \int_Y [\check{\varphi}(y) + b k^{(o)}(\zeta_0, y)] \log \left(\frac{k^{(o)}(\zeta', y)}{k^{(o)}(\zeta, y)} \right) \nu(dy) \\ &> \left\langle \int_Y \check{\varphi} S \, d\nu + b \nabla A(\zeta_0) - (1 + b) \nabla A(\zeta'), \zeta' - \zeta \right\rangle_E \\ &= (1 + b) \langle \mathbf{s}^* - \nabla A(\zeta'), \zeta' - \zeta \rangle_E , \end{aligned} \quad (98)$$

where we used the definition of \mathbf{s}^* in (36). Now, since ∇A is a \mathcal{C}^∞ -diffeomorphism on $\text{Int} E_0$, if \mathbf{s}^* belongs to $\nabla A(\mathcal{O})$ then there exists a unique $\zeta^* \in \mathcal{O}$ such that $\nabla A(\zeta^*) = \mathbf{s}^*$. As a result, plugging this into (98) and setting $\zeta' = \zeta^*$, we have that $\mathcal{C}_0(\zeta^*) > \mathcal{C}_0(\zeta)$ for all $\zeta \neq \zeta^*$. This shows that ζ^* is the (unique) solution to the argmax problem (35).

We conclude this proof with the proof of the reciprocal implication. Suppose that the argmax problem (35) has a solution $\zeta \in \mathcal{O}$ so that, for all $\zeta' \in \mathcal{O}$, $\mathcal{C}_0(\zeta') - \mathcal{C}_0(\zeta) \leq 0$. Using (98), this would imply $\langle \mathbf{s}^* - \nabla A(\zeta'), \zeta' - \zeta \rangle < 0$ for all $\zeta' \in \mathcal{O} \setminus \{\zeta\}$. Since \mathcal{O} is an open set, we can take $\zeta' = \zeta + \epsilon(\mathbf{s}^* - \nabla A(\zeta))$ with $\epsilon > 0$ small enough, which gives

$$\langle \mathbf{s}^* - \nabla A(\zeta'), \mathbf{s}^* - \nabla A(\zeta) \rangle < 0 .$$

Now letting $\epsilon \downarrow 0$ and since ∇A is continuous, we get that $\mathbf{s}^* - \nabla A(\zeta) = 0$ and thus that \mathbf{s}^* belongs to $\nabla A(\mathcal{O})$. \blacksquare

Proof of Corollary 6 First note that letting ν be the d -dimensional Lebesgue measure and setting $S(y) = (y, -yy^T/2)$, the condition (34) written in Theorem 4 takes the form

$$\int_Y \|y\|^2 \check{\varphi}(y) \, dy < \infty. \quad (99)$$

Furthermore, setting $v(\theta) = (\Sigma^{-1}m, \Sigma^{-1})$ with $\theta = (m, \Sigma)$, we obtain that $v(\mathbb{T}) = \mathbb{R}^d \times \mathcal{M}_{>0}(d)$, which is the whole interior set $\text{Int}E_0$, seen as an open set of $E = \mathbb{R}^d \times \mathbb{R}^{d \times d}$. Hence Theorem 4 applies and we can use the interpretation (38) of the equation (37) with $k(\theta^*, y) = \mathcal{N}(y; m^*, \Sigma^*)$ and $k(\theta_0, y) = \mathcal{N}(y; m_0, \Sigma_0)$, that is

$$\int_Y S(y) \mathcal{N}(y; m^*, \Sigma^*) \nu(dy) = \int_Y S(y) \psi(y) \nu(dy), \quad (100)$$

where ψ is the mixture

$$\psi(y) = \frac{1}{1+b} \check{\varphi}(y) + \frac{b}{1+b} \mathcal{N}(y; m_0, \Sigma_0), \quad y \in Y.$$

Using now that $S(y) = (y, -yy^T/2)$, (100) is equivalent to saying that the distributions on both sides of this equation have the same mean and covariance matrix. Since the left-hand side distribution is a Gaussian distribution, we directly obtain that the solution of (100) is given by $\theta^* = (m^*, \Sigma^*)$ with m^* and Σ^* being the mean and covariance matrix of the random variable Y with density ψ .

Observe that these mean and covariance matrix are well-defined due to the condition (99). This means that we only have to check that the covariance matrix is positive definite in order to have a solution in $\mathbb{T} = \mathbb{R}^d \times \mathcal{M}_{>0}(d)$. If $b > 0$, then $1/(1+b) < 1$ and the fact that $\theta \in \mathbb{T}$ guaranties that the covariance matrix of Y is positive definite. If $b = 0$, then Y has density $\check{\varphi}$ and the non-degenerate condition on $\check{\varphi}$ guaranties that the covariance is also positive definite in this case. This concludes the proof. \blacksquare

Proof of Corollary 7 Recall from (16) that $\varphi_{j,n}^{(\alpha)}(y) = k(\theta_{j,n}, y)(p(y)/\mu_n k(y))^{1-\alpha}$ with $k(\theta_{j,n}, y) > 0$. Then Lemma 8 gives that

$$0 < \int_Y (1 + \|S(y)\|_E) \varphi_{j,n}^{(\alpha)}(y) \nu(dy) < \infty. \quad (101)$$

We can thus apply Theorem 4 with $\zeta_0 = v(\theta_{j,n})$, $\mathcal{O} = v(\mathbb{T})$, $b = b_{j,n}/\int_Y \varphi_{j,n}^{(\alpha)} \, d\nu$ and $\check{\varphi} = \varphi_{j,n}^{(\alpha)}/\int_Y \varphi_{j,n}^{(\alpha)} \, d\nu$, which satisfies (34). \blacksquare

C.3 Proofs of Section 6.3

Proof of Proposition 3 Following the proof of Theorem 4 but with $b = 0$, we have using (98) that: for all $\zeta', \zeta \in \text{Int}E_0$,

$$\left\langle \nabla A(\zeta) - \int_Y S \check{\varphi} \, d\nu, \zeta' - \zeta \right\rangle_E \leq g^{(o)}(\zeta') - g^{(o)}(\zeta) \leq \left\langle \nabla A(\zeta') - \int_Y S \check{\varphi} \, d\nu, \zeta' - \zeta \right\rangle_E.$$

Since ∇A is continuous in $\text{Int}E_0$, we get that $g^{(o)}$ is \mathcal{C}^1 over $\text{Int}E_0$ and, for all $\zeta \in \text{Int}E_0$,

$$\nabla g^{(o)}(\zeta) = \nabla A(\zeta) - \int_{\mathbf{Y}} S \check{\varphi} \, \mathrm{d}\nu .$$

The β_0 -smoothness condition follows from the mean value theorem applied to $\nabla g^{(o)}$ whose Jacobian is $\nabla \nabla^T A$. \blacksquare

Proof of Corollary 8 The context is similar to the proof of Corollary 6 except that the covariance matrix is known and thus the sufficient statistic reduces to $S(y) = y$ and the canonical parameter to $v(\theta) = \Sigma^{-1}\theta$. The given expressions of h and A follow. Then we have, for all $\zeta \in \mathbb{R}^d$, $\nabla A(\zeta) = \Sigma\zeta$. Applying Proposition 3, we get that

$$\nabla g^{(o)}(\zeta) = \Sigma\zeta - \int_{\mathbf{Y}} y \check{\varphi}(y) \, \mathrm{d}y .$$

This gives that $g^{(o)}$ is β_0 -smooth over \mathbb{R}^d with β_0 equal to the maximal eigenvalue of Σ and leads to the gradient step

$$\zeta = \zeta_0 - \frac{\gamma}{\beta_0} \left(\Sigma\zeta_0 - \int_{\mathbf{Y}} y \check{\varphi}(y) \, \mathrm{d}y \right) .$$

Multiplying on both sides by Σ leads to the gradient-based update (46).

We now derive (47). Using that $g = g^{(o)} \circ v$, we get that $\nabla g(\theta) = \Sigma^{-1} \nabla g^{(o)}(\Sigma^{-1}\theta)$ and thus

$$\nabla g(\theta) = \Sigma^{-1} \left(\theta - \int_{\mathbf{Y}} y \check{\varphi}(y) \, \mathrm{d}y \right) .$$

We obtain that g is β -smooth over \mathbb{R}^d with β equal to the maximal eigenvalue of Σ^{-1} , which leads to the gradient update (47). \blacksquare

C.4 Proofs of Section 6.4

C.4.1 PROOFS OF PRELIMINARY RESULTS

It will be convenient to use the shorthand notation: for all $(\ell, y, \xi) \in \mathbf{Z} \times \mathbf{Y} \times F$,

$$\begin{aligned} k_\xi(\ell, y) &:= k^{(o)}(\ell(\xi), y) , \\ \check{k}_{\xi, \tau}(\ell, y) &:= \frac{k_\xi(\ell, y)}{\tau k_\xi(y)} , \end{aligned}$$

where τk_ξ denotes the mapping $y \mapsto \int k_\xi(\ell, y) \, \tau(\mathrm{d}\ell)$.

Proof of Proposition 4 We have, for all $\tau_0, \tau \in \mathcal{M}$, $\xi_0, \xi \in \mathcal{F}$ and $y \in \mathcal{Y}$, using the definitions above and that of D_{τ, τ_0} ,

$$\begin{aligned} \log \frac{\tau k_\xi(y)}{\tau_0 k_{\xi_0}(y)} &= \log \int_{\mathcal{Z}} \left(\frac{k_\xi(\ell, y)}{k_{\xi_0}(\ell, y)} \right) \check{k}_{\xi_0, \tau_0}(\ell, y) \tau(d\ell) \\ &= \log \int_{\mathcal{Z}} \left(\frac{k_\xi(\ell, y)}{k_{\xi_0}(\ell, y)} D_{\tau, \tau_0}(\ell) \right) \check{k}_{\xi_0, \tau_0}(\ell, y) \tau_0(d\ell) \\ &\geq \int_{\mathcal{Z}} \check{k}_{\xi_0, \tau_0}(\ell, y) \log \left(\frac{k_\xi(\ell, y)}{k_{\xi_0}(\ell, y)} D_{\tau, \tau_0}(\ell) \right) \tau_0(d\ell), \end{aligned}$$

where we used the concavity of the log function and the fact that $\check{k}_{\xi_0, \tau_0}(\cdot, y)$ is a probability density function with respect to τ_0 . Let $\xi = \Upsilon(\vartheta)$ and $\xi_0 = \Upsilon(\vartheta_0)$ and set $\theta = (\vartheta, \tau)$ and $\theta_0 = (\vartheta_0, \tau_0)$, so that, by Definition 2, we have $\tau k_\xi(y) = k^{(1)}(\theta, y)$ and $\tau_0 k_{\xi_0}(y) = k^{(1)}(\theta_0, y)$. Applying the previous bound and integrating w.r.t. $\check{\varphi}(y)\nu(dy)$ thus yields

$$\begin{aligned} \int_{\mathcal{Y}} \check{\varphi}(y) \log \left(\frac{k^{(1)}(\theta, y)}{k^{(1)}(\theta_0, y)} \right) \nu(dy) \\ \geq \int_{\mathcal{Y}} \left(\int_{\mathcal{Z}} \varphi(\ell, y) \log \left(\frac{k_\xi(\ell, y)}{k_{\xi_0}(\ell, y)} D_{\tau, \tau_0}(\ell) \right) \tau_0(d\ell) \right) \nu(dy), \end{aligned} \quad (102)$$

where we used that $\check{\varphi}(y)\check{k}_{\xi_0, \tau_0}(\ell, y) = \varphi(\ell, y)$ as a consequence of (52). We conclude by noting that (50), (51) and (102) imply (49). \blacksquare

Proof of Corollary 9 Suppose that $\Upsilon(\vartheta)$ belongs to the argmax (54), and let us show that (51) holds. Note that the objective function in the argmax is zero for $\xi = \xi_0$ and thus must be non-negative for $\xi = \Upsilon(\vartheta)$ in the argmax. Note also that the left-hand side of (51) is the sum of this objective function at $\xi = \Upsilon(\vartheta)$ with the term

$$- \int_{\mathcal{Y}} \left(\int_{\mathcal{Z}} b_0(\ell) k_{\xi_0}(\ell, y) \log \left(\frac{k_{\Upsilon(\vartheta)}(\ell, y)}{k_{\xi_0}(\ell, y)} \right) \tau_0(d\ell) \right) \nu(dy), \quad (103)$$

and it only remains to show that this term is also non-negative. To conclude, we observe that, for τ_0 -a.e. $\ell \in \mathcal{Z}$, Condition (ii) of Definition 1 implies that $k_{\xi_0}(\ell, \cdot)$ and $k_{\Upsilon(\vartheta)}(\ell, \cdot)$ are probability densities, so that

$$- \int_{\mathcal{Y}} k_{\xi_0}(\ell, y) \log \left(\frac{k_{\Upsilon(\vartheta)}(\ell, y)}{k_{\xi_0}(\ell, y)} \right) \nu(dy) \geq 0,$$

and thus, b_0 being non-negative,

$$- \int_{\mathcal{Y}} \left(\int_{\mathcal{Z}} b_0(\ell) k_{\xi_0}(\ell, y) \log \left(\frac{k_{\Upsilon(\vartheta)}(\ell, y)}{k_{\xi_0}(\ell, y)} \right) \nu(dy) \right) \tau_0(d\ell) \geq 0.$$

Using that $\log_+(x) \leq x$ for all $x \geq 0$, the positive part inside this double integral is at most $b_0(\ell) k_{\Upsilon(\vartheta)}(\ell, y)$, whose integral w.r.t. $\nu(dy)\tau_0(d\ell)$ is $\int b_0 d\tau_0 < \infty$ by assumption. Therefore,

the same inequality holds when reversing the order of integration in the previous display. We get that the expression in (103) is non-negative, which concludes the proof. \blacksquare

The following lemma is used to illustrate (C1).

Lemma 9 *Consider an ELM family satisfying (C1). Then, for all $\vartheta \in \tilde{T}$, $\tau \in M$ and τ -a.e. ℓ , $\ell \circ \Upsilon(\tilde{T})$ is an open subset included in $\text{Int}E_0$.*

Proof In (C1), the linear operator $\ell \in Z$ being full rank implies that the image of an open set is an open subset of E . Indeed, take $F' \subset F$ a linear supplementary space of the kernel of ℓ . Then $\ell|_{F'}$ is a bijective linear mapping from F' to E and thus, for any open set O included in F , $\ell(O) = \ell(O \cap F') = \ell|_{F'}(O \cap F')$ is an open subset of E since $O \cap F'$ is an open subset of F' and $\ell|_{F'}$ is bicontinuous. Hence (C1) implies that $\ell \circ \Upsilon(\tilde{T})$ is an open subset of E for all $\tau \in M$ and τ -almost all $\ell \in Z$, which, by Assumption (ii) from Definition 2, must be included in $\text{Int}E_0$. \blacksquare

We conclude with the following lemma, which relates Assumption (C4) to a kind of identifiability property.

Lemma 10 *Consider an ELM family and let $\mathbb{P}_{\vartheta, \tau}$ denote the probability of (Y, L) defined on $Y \times Z$ as above. Then (C4) is equivalent to have that for all $(\vartheta, \tau), (\vartheta', \tau') \in \tilde{T} \times M$, $\mathbb{P}_{\vartheta, \tau} = \mathbb{P}_{\vartheta', \tau'}$ if and only if $\Upsilon(\vartheta) = \Upsilon(\vartheta')$ and $\tau = \tau'$.*

Proof Let $(\vartheta, \tau), (\vartheta', \tau') \in \tilde{T} \times M$. Note that $\Upsilon(\vartheta) = \Upsilon(\vartheta')$ and $\tau = \tau'$ is always sufficient to have $\mathbb{P}_{\vartheta, \tau} = \mathbb{P}_{\vartheta', \tau'}$. Also, since τ is the marginal over Z of $\mathbb{P}_{\vartheta, \tau}$ and $k^{(o)}(\ell \circ \Upsilon(\vartheta), \cdot)$ is the conditional density of Y given $L = \ell$ with respect to ν , it is clear that $\mathbb{P}_{\vartheta, \tau} = \mathbb{P}_{\vartheta', \tau'}$ is equivalent to have $\tau = \tau'$ and

$$k^{(o)}(\ell \circ \Upsilon(\vartheta), y) = k^{(o)}(\ell \circ \Upsilon(\vartheta'), y) \quad \text{for } \nu\text{-a.e. } y \text{ and } \tau\text{-a.e. } \ell. \quad (104)$$

By (B1), which is assumed for an ELM family in Definition 2, for all $\zeta, \zeta' \in E_0$, we have $k^{(o)}(\zeta, y) = k^{(o)}(\zeta', y)$ for ν -a.e. y if and only if $\zeta = \zeta'$. By (ii) in Definition 2, we have $\ell \circ \Upsilon(\vartheta) \in E_0$ for all $\vartheta \in \tilde{T}$, $\tau \in M$ and τ -a.e. ℓ . Thus, for all $(\vartheta, \vartheta', \tau) \in \tilde{T}^2 \times M$, Eq. (104) is equivalent to have

$$\ell \circ \Upsilon(\vartheta) = \ell \circ \Upsilon(\vartheta') \quad \text{for } \tau\text{-a.e. } \ell. \quad (105)$$

Hence, to prove the result, we only need to show that (C4) is equivalent to have that, for all $(\vartheta, \vartheta', \tau) \in \tilde{T}^2 \times M$, Eq. (105) implies $\Upsilon(\vartheta) = \Upsilon(\vartheta')$. In fact (C4) exactly says that for all $\vartheta, \vartheta' \in \tilde{T}$, $\Upsilon(\vartheta) \neq \Upsilon(\vartheta')$ implies the opposite of (105). \blacksquare

C.4.2 PROOF OF THEOREM 5

The proof of Theorem 5 relies on a general proposition (Proposition 7) and two technical lemmas (Lemmas 11 and 12). The following definition will be useful to state the proposition.

Definition 4 (Directional continuity) *Let g be defined on an open subset \mathcal{O} of the Euclidean space F and valued in a topological space. Let $\xi \in \mathcal{O}$ and $x \in F$ such that $\|x\|_F = 1$. We say that g is continuous at ξ in direction x if $\lim_{t \downarrow 0} g(\xi + tx) = g(\xi)$.*

Proposition 7 Consider an ELM family as in Definition 2 and Υ satisfying (C1). Let $\check{\varphi}$ be a probability density function w.r.t. ν . Let $\tau_0 \in \mathcal{M}$ and $\varphi : \mathcal{Z} \times \mathcal{Y} \rightarrow \mathbb{R}_+$ be measurable function such that

$$\int_{\mathcal{Y} \times \mathcal{Z}} \varphi(\ell, y) \|\ell\|_{\text{op}} \|S(y)\|_E \nu(dy) \tau_0(d\ell) < \infty \quad (106)$$

$$\int_{\mathcal{Z}} \tilde{\varphi}(\ell) \|\ell^\nabla(\xi)\|_F \tau_0(d\ell) < \infty \quad \text{for all } \xi \in \Upsilon(\tilde{\mathcal{T}}), \quad (107)$$

where we set $\varphi(\ell, y)$ and $\tilde{\varphi}(\ell)$ as in (52) and (59). Define $w_{\tau_0}^\nabla : \Upsilon(\tilde{\mathcal{T}}) \rightarrow F$ as in (60). Let $b \in \mathbb{R}_+$ and ξ_0 in $\Upsilon(\tilde{\mathcal{T}})$ and set \mathbf{w}^* as in (63). Then the reciprocal set

$$\Xi^* = \left\{ \xi \in \Upsilon(\tilde{\mathcal{T}}) : w_{\tau_0}^\nabla(\xi) = \mathbf{w}^* \right\}$$

is included in the set of solutions of the argmax problem

$$\operatorname{argmax}_{\xi \in \Upsilon(\tilde{\mathcal{T}})} \int_{\mathcal{Y}} \left(\int_{\mathcal{Z}} [\varphi(\ell, y) + b \tilde{\varphi}(\ell) k^{(o)}(\ell(\xi_0), y)] \log \left(\frac{k^{(o)}(\ell(\xi), y)}{k^{(o)}(\ell(\xi_0), y)} \right) \tau_0(d\ell) \right) \nu(dy). \quad (108)$$

Moreover the two following assertions hold.

(i) If $w_{\tau_0}^\nabla$ is continuous on $\Upsilon(\tilde{\mathcal{T}})$ in every direction, the opposite inclusion is true, namely, the solution set of the argmax problem (108) is included in Ξ^* .

(ii) If we have, for all $\xi \neq \xi' \in \Upsilon(\tilde{\mathcal{T}})$,

$$\int_{\{\ell \in \mathcal{Z} : \ell(\xi - \xi') \neq 0\}} \tilde{\varphi}(\ell) \tau_0(d\ell) > 0, \quad (109)$$

then $\xi \mapsto w_{\tau_0}^\nabla(\xi)$ is one-to-one on $\Upsilon(\tilde{\mathcal{T}})$ and thus Ξ^* is either a singleton or empty.

Proof In this proof, we take $b \geq 0$ and $\xi_0 \in \Upsilon(\tilde{\mathcal{T}})$, and we denote the integral defining the objective function in the argmax problem (108) by $\mathcal{C}(\xi)$.

A first remark is that, by (106), $\int \varphi(\ell, y) \ell^T \circ S(y) \nu(dy) \tau_0(d\ell)$ is well-defined in F and $w_{\tau_0}^\nabla(\xi)$ in (60) is well-defined in F for all $\xi \in \Upsilon(\tilde{\mathcal{T}})$ as a consequence of (107). In particular we get that \mathbf{w}^* in (63) is well-defined in F (but does not necessarily belong to $w_{\tau_0}^\nabla \circ \Upsilon(\tilde{\mathcal{T}})$).

Next we show that the objective function $\mathcal{C}(\xi)$ is well defined and finite for all $\xi \in \Upsilon(\tilde{\mathcal{T}})$ and that Fubini's theorem applies (so that the double integral in (108) can be computed in the opposite order or as a single integral over $\mathcal{Y} \times \mathcal{Z}$ w.r.t. the product measure $\nu \otimes \tau_0$). To this end, we show successively that

$$\int_{\mathcal{Y} \times \mathcal{Z}} \varphi(\ell, y) \left| \log \left(\frac{k^{(o)}(\ell(\xi), y)}{k^{(o)}(\ell(\xi_0), y)} \right) \right| \nu(dy) \tau_0(d\ell) < \infty, \quad (110)$$

$$\int_{\mathcal{Y} \times \mathcal{Z}} \tilde{\varphi}(\ell) k^{(o)}(\ell(\xi_0), y) \log_+ \left(\frac{k^{(o)}(\ell(\xi), y)}{k^{(o)}(\ell(\xi_0), y)} \right) \nu(dy) \tau_0(d\ell) < \infty, \quad (111)$$

$$-\infty < \int_{\mathcal{Z}} \tilde{\varphi}(\ell) \left(\int_{\mathcal{Y}} k^{(o)}(\ell(\xi_0), y) \log \left(\frac{k^{(o)}(\ell(\xi), y)}{k^{(o)}(\ell(\xi_0), y)} \right) \nu(dy) \right) \tau_0(d\ell) \leq 0, \quad (112)$$

$$\int_{\mathcal{Y} \times \mathcal{Z}} \tilde{\varphi}(\ell) k^{(o)}(\ell(\xi_0), y) \left| \log \left(\frac{k^{(o)}(\ell(\xi), y)}{k^{(o)}(\ell(\xi_0), y)} \right) \right| \nu(dy) \tau_0(d\ell) < \infty. \quad (113)$$

Let us now show that (110) holds. Let $\ell \in \mathbf{Z}$ such that $\ell \circ \Upsilon(\tilde{\mathbf{T}}) \subseteq \text{Int}E_0$, which by Condition (ii) and (C1), is true for τ_0 -a.e. $\ell \in \mathbf{Z}$. Applying (93) in Lemma 7 with $(\zeta, \zeta') = (\ell(\xi_0), \ell(\xi))$ as well as $(\zeta, \zeta') = (\ell(\xi), \ell(\xi_0))$ and using the definition of ℓ^T in (55) and that of ℓ^∇ in (56), we get that

$$\left\langle \ell^T \circ S(y) - \ell^\nabla(\xi), \xi - \xi_0 \right\rangle_F \leq \log \left(\frac{k^{(o)}(\ell(\xi), y)}{k^{(o)}(\ell(\xi_0), y)} \right) \leq \left\langle \ell^T \circ S(y) - \ell^\nabla(\xi_0), \xi - \xi_0 \right\rangle_F .$$

It follows that, for all $\xi_0, \xi \in \Upsilon(\tilde{\mathbf{T}})$,

$$\left| \log \left(\frac{k^{(o)}(\ell(\xi), y)}{k^{(o)}(\ell(\xi_0), y)} \right) \right| \leq \left(\left\| \ell^T \circ S(y) \right\|_F + \max(\left\| \ell^\nabla(\xi_0) \right\|_F, \left\| \ell^\nabla(\xi) \right\|_F) \right) \left\| \xi - \xi_0 \right\|_F . \quad (114)$$

Using $\left\| \ell^T \circ S(y) \right\|_F \leq \left\| \ell^T \right\|_{\text{op}} \|S(y)\|_E = \|\ell\|_{\text{op}} \|S(y)\|_E$, (106) and (107), we obtain (110).

To show (111), we use the bound $\log_+(x) \leq x$ for all $x \geq 0$ so that the left-hand side of (111) is upper-bounded by $\int_{\mathbf{Y} \times \mathbf{Z}} \tilde{\varphi}(\ell) k^{(o)}(\ell(\xi), y) \nu(dy) \tau_0(d\ell) = \int_{\mathbf{Z}} \tilde{\varphi}(\ell) \tau_0(d\ell) = 1$.

Next, to show (112), we observe that, as a consequence of (94), (55) and (56), we have

$$\begin{aligned} 0 &\geq \int_{\mathbf{Y}} k^{(o)}(\ell(\xi), y) \log \left(\frac{k^{(o)}(\ell(\xi), y)}{k^{(o)}(\ell(\xi_0), y)} \right) \nu(dy) \geq \left\langle \ell^\nabla(\xi_0) - \ell^\nabla(\xi), \xi - \xi_0 \right\rangle_F \\ &\geq - \left\| \ell^\nabla(\xi_0) - \ell^\nabla(\xi) \right\|_F \left\| \xi - \xi_0 \right\|_F . \end{aligned}$$

Hence, using (107), we get (112).

Finally, Tonelli's theorem gives us that the same positive part as in (111) has a finite integral when integrating in the same order as in (112), so that (111) and (112) imply (113).

We now prove the inclusion of Ξ^* in the set of argmax solutions. Since Fubini's theorem applies in the definition of the objective function, we can write, for all $\xi, \xi' \in \Upsilon(\tilde{\mathbf{T}})$,

$$\mathcal{C}(\xi') - \mathcal{C}(\xi) = \int_{\mathbf{Y} \times \mathbf{Z}} \left[\varphi(\ell, y) + b \tilde{\varphi}(\ell) k^{(o)}(\ell(\xi_0), y) \right] \log \left(\frac{k^{(o)}(\ell(\xi'), y)}{k^{(o)}(\ell(\xi), y)} \right) \nu(dy) \tau_0(d\ell) .$$

Using (93) in Lemma 7 this time with $\zeta = \ell(\xi)$ and $\zeta' = \ell(\xi')$, we obtain that

$$\mathcal{C}(\xi') - \mathcal{C}(\xi) \geq \int_{\mathbf{Y} \times \mathbf{Z}} \left[\varphi(\ell, y) + b \tilde{\varphi}(\ell) k^{(o)}(\ell(\xi_0), y) \right] \left\langle \ell^T \circ S(y) - \ell^\nabla(\xi'), \xi' - \xi \right\rangle_E \nu(dy) \tau_0(d\ell) .$$

Using (106), (107) to integrate in any convenient order, we have

$$\begin{aligned} \int_{\mathbf{Y} \times \mathbf{Z}} \varphi(\ell, y) \left\langle \ell^T \circ S(y) - \ell^\nabla(\xi'), \xi' - \xi \right\rangle_E \nu(dy) \tau_0(d\ell) \\ = \left\langle \int_{\mathbf{Y} \times \mathbf{Z}} \varphi(\ell, y) \ell^T \circ S(y) \nu(dy) \tau_0(d\ell) - \mathbf{w}_{\tau_0}^\nabla(\xi'), \xi' - \xi \right\rangle_E , \end{aligned}$$

where we have used the definition of $\mathbf{w}_{\tau_0}^\nabla$ in (60), and

$$\begin{aligned} \int_{\mathbf{Y} \times \mathbf{Z}} \tilde{\varphi}(\ell) k^{(o)}(\ell(\xi_0), y) \left\langle \ell^T \circ S(y) - \ell^\nabla(\xi'), \xi' - \xi \right\rangle_E \nu(dy) \tau_0(d\ell) \\ = \left\langle \mathbf{w}_{\tau_0}^\nabla(\xi_0) - \mathbf{w}_{\tau_0}^\nabla(\xi'), \xi' - \xi \right\rangle_E . \end{aligned}$$

where we have used once again the definition of $w_{\tau_0}^\nabla$ in (60) as well as the expression of ∇A given in (32). Plugging these two equalities in the previous inequality, we finally get that, for all $\xi, \xi' \in \Upsilon(\tilde{T})$,

$$\begin{aligned} \mathcal{C}(\xi') - \mathcal{C}(\xi) &\geq \left\langle \int_{Y \times Z} \varphi(\ell, y) \ell^T \circ S(y) \nu(dy) \tau_0(d\ell) - (1+b) w_{\tau_0}^\nabla(\xi') + b w_{\tau_0}^\nabla(\xi_0), \xi' - \xi \right\rangle_E \\ &= (1+b) \left\langle \mathbf{w}^* - w_{\tau_0}^\nabla(\xi'), \xi' - \xi \right\rangle_E, \end{aligned} \quad (115)$$

Hence any $\xi' \in \Xi^*$ must be a solution of the argmax (108).

Let us now prove Assertion (i). We use again (115) this time with ξ in the set of the argmax solutions and ξ' in the neighborhood of ξ (remember that $\Upsilon(\tilde{T})$ is assumed to be open) such that $\xi' - \xi = \epsilon \left(\mathbf{w}^* - w_{\tau_0}^\nabla(\xi) \right)$ with $\epsilon > 0$ small enough. We obtain that

$$\left\langle \mathbf{w}^* - w_{\tau_0}^\nabla(\xi'), \mathbf{w}^* - w_{\tau_0}^\nabla(\xi) \right\rangle_E \leq 0.$$

The assumption that $w_{\tau_0}^\nabla$ is continuous at ξ in direction $\left(\mathbf{w}^* - w_{\tau_0}^\nabla(\xi) \right)$ gives us that $\mathbf{w}^* = w_{\tau_0}^\nabla(\xi)$, that is, $\xi \in \Xi^*$ and the opposite inclusion is shown.

We conclude with the proof of Assertion (ii). By (60), for all $\xi, \xi' \in \Upsilon(\tilde{T})$, we can write

$$w_{\tau_0}^\nabla(\xi) - w_{\tau_0}^\nabla(\xi') = \int_Z \tilde{\varphi}(\ell) \left(\ell^\nabla(\xi) - \ell^\nabla(\xi') \right) \tau_0(d\ell).$$

Since ∇A has a positive definite Jacobian on the convex set $\text{Int}E_0$, we have, for all $\zeta \neq \zeta' \in \text{Int}E_0$, $\langle \nabla A(\zeta) - \nabla A(\zeta'), \zeta - \zeta' \rangle_E > 0$.

Letting $\xi \neq \xi' \in \Upsilon(\tilde{T})$ and $\ell \in Z$ be such that $\ell(\xi), \ell(\xi') \in \text{Int}E_0$ and $\ell(\xi - \xi') \neq 0$, we can apply the previous inequality with $\zeta = \ell(\xi)$ and $\zeta' = \ell(\xi')$ leading to

$$\left\langle \ell^\nabla(\xi) - \ell^\nabla(\xi'), \xi - \xi' \right\rangle_F > 0.$$

Combining this result with the condition (109) gives us that for all $\xi \neq \xi' \in \Upsilon(\tilde{T})$

$$\left\langle w_{\tau_0}^\nabla(\xi) - w_{\tau_0}^\nabla(\xi'), \xi - \xi' \right\rangle_F > 0.$$

In particular, $w_{\tau_0}^\nabla$ is one-to-one and the proof is concluded. ■

Let d_F denote the dimension of F and let us now introduce and prove two key lemmas.

Lemma 11 *Let $\xi \in \Upsilon(\tilde{T})$ and let $x \in F$ be such that $\|x\|_F = 1$. Let us denote, for any $\epsilon > 0$,*

$$\begin{aligned} B_F(\xi, \epsilon) &= \{ \xi' \in F : \|\xi - \xi'\|_F < \epsilon \}, \\ D_x(\xi, \epsilon) &= \{ \xi + t x : 0 \leq t < 1 \}. \end{aligned}$$

Let $\ell \in Z$ be such that $\ell \circ \Upsilon(\tilde{T}) \subset \text{Int}E_0$. Then the two following assertions hold.

(i) *For any $0 < \epsilon \leq \epsilon' < \epsilon''$ such that $D_x(\xi, \epsilon'') \subset \Upsilon(\tilde{T})$, we have*

$$\sup_{\xi' \in D_x(\xi, \epsilon)} |A \circ \ell(\xi') - A \circ \ell(\xi)| \leq \epsilon \max \left(\left\| \ell^\nabla(\xi) \right\|_F, \left\| \ell^\nabla(\xi + \epsilon' x) \right\|_F \right). \quad (116)$$

(ii) For any $\epsilon > 0$ such that $B_F(\xi, 2(1 + \sqrt{d_F})\epsilon) \subset \Upsilon(\tilde{T})$, we have, for all $\epsilon_0 \in (0, \epsilon]$,

$$\int_Y e^{\epsilon_0 \|\ell^T \circ S(y)\|_F} \left(\sup_{\xi' \in D_x(\xi, \epsilon_0)} k^{(o)}(\ell(\xi'), y) \right) \nu(dy) \leq 2 d_F \max_{\xi' \in \Delta_{x, \epsilon}(\xi)} \left(e^{\epsilon'_0 \|\ell^\nabla(\xi')\|_F} \right), \quad (117)$$

where $\epsilon'_0 = (1 + 2\sqrt{d_F})\epsilon_0$ and $\Delta_{x, \epsilon}(\xi)$ is a set of $2(d_F + 1)$ points in $B_x(\xi, 2(1 + \sqrt{d_F})\epsilon)$ only depending on ξ , x and ϵ .

Proof We prove Assertions (i) and (ii) successively.

Proof of Assertion (i) Let $0 < \epsilon \leq \epsilon' < \epsilon''$ be such that $D_x(\xi, \epsilon'') \subset \Upsilon(\tilde{T})$. First note that $t \mapsto A \circ \ell(\xi + tx)$ has its first and second derivative on $t \in [0, \epsilon']$ given respectively by

$$\begin{aligned} h_1 : t &\mapsto \langle \nabla A \circ \ell(\xi + tx), \ell(x) \rangle_E = \langle \ell^\nabla(\xi + tx), x \rangle_F \\ h_2 : t &\mapsto \left\langle \left[\nabla \nabla^T A \circ \ell(\xi + tx) \right] \ell(x), \ell(x) \right\rangle_E. \end{aligned}$$

Since $\nabla \nabla^T A$ is positive definite on $\text{Int} E_0$, h_2 is non-negative, thus h_1 is monotonous on $[0, \epsilon']$ and its maximal absolute value is obtained at $t = 0$ or $t = \epsilon'$, meaning that $|h_1(t)| \leq \max(|h_1(0)|, |h_1(\epsilon')|)$ for all $t \in [0, \epsilon']$. Therefore, the bound in (116) follows from the mean value theorem combined with the fact that: for all $t \in [0, \epsilon']$ and $\xi' \in D_x(\xi, \epsilon)$,

$$\begin{aligned} |h_1(t)| &\leq \|\ell^\nabla(\xi + tx)\|_F \|x\|_F = \|\ell^\nabla(\xi + tx)\|_F \\ \|\xi' - \xi\|_F &\leq \epsilon \|x\|_F \leq \epsilon \end{aligned}$$

since $\|x\|_F = 1$ by assumption.

Proof of Assertion (ii) Take $0 < \epsilon_0 \leq \epsilon$ such that $B_x(\xi, 2(1 + \sqrt{d_F})\epsilon) \subset \Upsilon(\tilde{T})$, let ϵ'_0 be as stated and set $\epsilon' = (1 + 2\sqrt{d_F})\epsilon$. Let x_1, \dots, x_{d_F} be an orthonormal basis of F so that, for any $w \in E$, we have

$$\begin{aligned} \|\ell^T(w)\|_F &= \left(\sum_{k=1}^{d_F} \left| \langle \ell^T(w), x_k \rangle_F \right|^2 \right)^{1/2} \leq \sqrt{d_F} \max_{k,s} \left(s \langle \ell^T(w), x_k \rangle_F \right) \\ &= \max_{k,s} \langle w, \ell(s \sqrt{d_F} x_k) \rangle_E, \end{aligned} \quad (118)$$

where, here and in the following, $\max_{k,s}$ is the maximum over $s \in \{\pm 1\}$ and $k = 1 \dots d_F$. By definition of $k^{(o)}$ in Definition 1, we also have: for all $\xi' \in D_x(\xi, \epsilon_0)$,

$$\begin{aligned} e^{\epsilon_0 \|\ell^T \circ S(y)\|_F} k^{(o)}(\ell(\xi'), y) &= h(y) e^{\epsilon_0 \|\ell^T \circ S(y)\|_F + \langle S(y), \ell(\xi') \rangle_E - A \circ \ell(\xi')} \\ &= h(y) e^{\epsilon_0 \|\ell^T \circ S(y)\|_F + \langle S(y), \ell(\xi' - \xi) \rangle_E + \langle S(y), \ell(\xi) \rangle_E - A \circ \ell(\xi')} \end{aligned}$$

Now using that

$$|\langle S(y), \ell(\xi' - \xi) \rangle_E| = \left| \langle \ell^T \circ S(y), \xi' - \xi \rangle_F \right| \leq \epsilon_0 \|\ell^T \circ S(y)\|_F$$

(since $\|x\|_F = 1$ and $\xi' \in D_x(\xi, \epsilon_0)$) paired up with (118), we can deduce that

$$\begin{aligned} e^{\epsilon_0 \|\ell^T \circ S(y)\|_F} k^{(o)}(\ell(\xi'), y) &\leq h(y) e^{2\epsilon_0 \|\ell^T \circ S(y)\|_F + \langle S(y), \ell(\xi) \rangle_E - A \circ \ell(\xi')} \\ &\leq \max_{k,s} h(y) \exp \left(\left\langle S(y), \ell(\xi + 2s \sqrt{d_F} \epsilon_0 x_k) \right\rangle_E - A \circ \ell(\xi') \right) \\ &\leq \max_{k,s} k^{(o)}(\ell(\xi''_{s,k}), y) e^{|A \circ \ell(\xi''_{s,k}) - A \circ \ell(\xi')|}, \end{aligned} \quad (119)$$

where $\xi''_{s,k} = \xi + 2s \sqrt{d_F} \epsilon_0 x_k$. Note that we took ϵ small enough so that for all s, k ,

$$D_x(\xi, \epsilon' + \epsilon) \cup D_{s x_k}(\xi, \epsilon' + \epsilon) \subset B_F(\xi, 2(1 + \sqrt{d_F})\epsilon) \subset \Upsilon(\tilde{\Gamma}).$$

ξ' and $\xi''_{s,k}$ are in $D_x(\xi, \epsilon_0)$ and $D_{s x_k}(\xi, 2\sqrt{d_F}\epsilon_0)$, respectively. Using (116) of Assertion (i) twice with $\epsilon \leq \epsilon' < \epsilon''$ successively replaced by $2\sqrt{d_F}\epsilon_0 \leq \epsilon' \leq 2(1 + \sqrt{d_F})\epsilon$ and $\epsilon_0 \leq \epsilon' \leq 2(1 + \sqrt{d_F})\epsilon$, we get that

$$\begin{aligned} |A \circ \ell(\xi''_{s,k}) - A \circ \ell(\xi')| &\leq |A \circ \ell(\xi''_{s,k}) - A \circ \ell(\xi)| + |A \circ \ell(\xi) - A \circ \ell(\xi')| \\ &\leq 2\sqrt{d_F}\epsilon_0 \max \left(\|\ell^\nabla(\xi)\|_F, \|\ell^\nabla(\xi + \epsilon' s x_k)\|_F \right) \\ &\quad + \epsilon_0 \max \left(\|\ell^\nabla(\xi)\|_F, \|\ell^\nabla(\xi + \epsilon' x)\|_F \right) \\ &\leq \epsilon'_0 C_{\xi,x}(\ell), \end{aligned}$$

where the last line follows by setting

$$C_{\xi,x}(\ell) = \max \left\{ \|\ell^\nabla(\xi + \epsilon' x')\|_F : x' \in \{0, x, s x_k\}, \text{ with } s \in \{\pm 1\}, k = 1 \dots d_F \right\}.$$

Combining this inequality with (119) then yields : for all $\xi' \in D_x(\xi, \epsilon_0)$,

$$\begin{aligned} e^{\epsilon_0 \|\ell^T \circ S(y)\|_F} k^{(o)}(\ell(\xi'), y) &\leq e^{\epsilon'_0 C_{\xi,x}(\ell)} \max_{k,s} k^{(o)}(\ell(\xi''_{s,k}), y) \\ &\leq e^{\epsilon'_0 C_{\xi,x}(\ell)} \sum_{k,s} k^{(o)}(\ell(\xi''_{s,k}), y). \end{aligned}$$

Observe that the right-hand side of the previous display does not depend on ξ' and that the summands in the sum (which has $2d_F$ terms) all have integrals equal w.r.t. ν equal to 1. The bound (117) follows and Assertion (ii) is proved. \blacksquare

Lemma 12 *Consider an ELM family as in Definition 2. Assume (C1) and (C3). Let $\vartheta_0 \in \tilde{\Gamma}$, $\tau_0 \in \mathcal{M}$ and $\check{\varphi} : \mathcal{Y} \rightarrow \mathbb{R}_+$ be a density function such that*

$$\int_{\mathcal{Y}} \check{\varphi}(y) e^{\tilde{s}\vartheta_0, \tau_0(y)} \nu(dy) < \infty. \quad (120)$$

Define $\varphi(\ell, y)$, $\tilde{\varphi}(\ell)$ and $w_{\tau_0}^\nabla$ as in (52), (59) and (60). Then for all $\xi \in \Upsilon(\tilde{\Gamma})$ and $x \in F$ such that $\|x\|_F = 1$, there exists $\epsilon_0 > 0$ such that

$$\int_{\mathcal{Y} \times \mathcal{Z}} \|\ell^T \circ S(y)\|_F \sup_{\xi' \in D_x(\xi, \epsilon_0)} \left(k^{(o)}(\ell(\xi'), y) \right) \tilde{\varphi}(\ell) \nu(dy) \tau_0(d\ell) < \infty. \quad (121)$$

Moreover, $w_{\tau_0}^\nabla$ is continuous on $\Upsilon(\tilde{T})$ in every direction, and, for all $\vartheta \in \tilde{T}$, it holds that

$$w_{\tau_0}^\nabla \circ \Upsilon(\vartheta) = \mathbb{E}_{\vartheta, \tau_0} \left[L^T \circ S(Y) \tilde{\varphi}(L) \right]. \quad (122)$$

Proof Let $\ell \in \mathbf{Z}$ is such that $\ell \circ \Upsilon(\tilde{T}) \subset \text{Int}E_0$, $\xi \in \Upsilon(\tilde{T})$ and $x \in F$ such that $\|x\|_F = 1$. Then for $\epsilon > 0$ small enough, we have $B_x(\xi, 2(1 + \sqrt{d_F})\epsilon) \subset \Upsilon(\tilde{T})$, and for all $\epsilon_0 \in (0, \epsilon]$,

$$\int_Y e^{\epsilon_0 \|\ell^T \circ S(y)\|_F} \sup_{\xi' \in D_x(\xi, \epsilon_0)} \left(k^{(o)}(\ell(\xi'), y) \right) \nu(dy) \leq 2 d_F \sum_{\xi' \in \Delta} e^{\epsilon'_0 \|\ell^\nabla(\xi')\|_F},$$

where we used Assertion (ii) of Lemma 11 with $\epsilon' = (1 + 2\sqrt{d_F})\epsilon$, $\epsilon'_0 = (1 + 2\sqrt{d_F})\epsilon_0$ and Δ is a set of $2(d_F + 1)$ points in $B_x(\xi, 2(1 + \sqrt{d_F})\epsilon)$ only depending on ξ , x and ϵ . Since $\|\ell^T \circ S(y)\|_F \leq \frac{1}{\epsilon_0} e^{\epsilon_0 \|\ell^T \circ S(y)\|_F}$, we get that:

$$\begin{aligned} G_{\xi, \epsilon_0}(\ell) &:= \int_Y \|\ell^T \circ S(y)\|_F \sup_{\xi' \in D_x(\xi, \epsilon_0)} \left(k^{(o)}(\ell(\xi'), y) \right) \nu(dy) \\ &\leq \frac{2 d_F}{\epsilon_0} \sum_{\xi' \in \Delta} e^{\epsilon'_0 \|\ell^\nabla(\xi')\|_F}. \end{aligned}$$

In particular by (C1), with Lemma 9, this bounds hold for τ_0 -almost all $\ell \in \mathbf{Z}$. By (C3), since Δ is a finite set only depending on ξ , x and ϵ and satisfying $\Delta \subset B_x(\xi, 2(1 + \sqrt{d_F})\epsilon) \subset \Upsilon(\tilde{T})$, there exist $\epsilon_\Delta, C_\Delta > 0$ only depending on ξ , x and ϵ such that, for all $\xi' \in \Delta$,

$$\mathbb{E}_{\vartheta_0, \tau_0} \left[e^{\epsilon_\Delta \|\ell^\nabla(\xi')\|_F} \mid Y \right] \leq C_\Delta e^{\tilde{s}_{\vartheta_0, \tau_0}(Y)} \quad \mathbb{P}_{\vartheta_0, \tau_0} - \text{a.s.}$$

Since Δ has cardinal $2(d_F + 1)$, taking $\epsilon_0 \in (0, \epsilon]$ small enough to have $\epsilon'_0 \leq \epsilon_\Delta$, this gives us that

$$g_{\xi, \epsilon_0}(Y) := \mathbb{E}_{\vartheta_0, \tau_0} [G_{\xi, \epsilon_0}(L) \mid Y] \leq \frac{4 C_\Delta d_F (d_F + 1)}{\epsilon_0} e^{\tilde{s}_{\vartheta_0, \tau_0}(Y)} \quad \mathbb{P}_{\vartheta_0, \tau_0} - \text{a.s.}$$

Now observing that the ratio in (52) is the conditional density of L applied to ℓ given $Y = y$ under $\mathbb{P}_{\vartheta_0, \tau_0}$ and that the marginal distribution of $\mathbb{P}_{\vartheta_0, \tau_0}$ on \mathbf{Y} is equivalent to ν , we obtain that

$$\begin{aligned} \int_{\mathbf{Z}} G_{\xi, \epsilon_0}(\ell) \tilde{\varphi}(\ell) \tau_0(d\ell) &= \int_{\mathbf{Z} \times \mathbf{Y}} G_{\xi, \epsilon_0}(\ell) \varphi(\ell, y) \tau_0(d\ell) \nu(dy) \\ &= \int_Y g_{\xi, \epsilon_0}(y) \check{\varphi}(y) \nu(dy) \\ &\leq \frac{4 C_\Delta d_F (d_F + 1)}{\epsilon_0} \int_Y e^{\tilde{s}_{\vartheta_0, \tau_0}(y)} \check{\varphi}(y) \nu(dy), \end{aligned}$$

Hence, using condition (120), we deduce that

$$\int_{\mathbf{Z}} G_{\xi, \epsilon_0}(\ell) \tilde{\varphi}(\ell) \tau_0(d\ell) < \infty,$$

which is exactly the claimed bound (121). This bound gives in particular that

$$\xi' \mapsto \int_{Y \times Z} \ell^T \circ S(y) k^{(\circ)}(\ell(\xi'), y) \tilde{\varphi}(\ell) \nu(dy) \tau_0(d\ell)$$

is well-defined in F for $\xi' \in D_x(\xi, \epsilon_0)$, and continuous at ξ in direction x . Moreover, the Fubini theorem applies so that (i) integrating first with respect to y and (ii) using the expression of ∇A given in (32), we obtain that this mapping is in fact $\xi' \mapsto w_{\tau_0}^\nabla(\xi')$. This being true for all $\xi \in \Upsilon(\tilde{T})$ and $x \in F$ such that $\|x\|_F = 1$, we get the last claim by observing that, by (57), the expectation in (122) is another way to express the same integral in the case where $\xi = \Upsilon(\vartheta)$. ■

We can now prove Theorem 5.

Proof of Theorem 5 We use Proposition 7 with $\varphi(\ell, y)$ given by (52). We first prove that Conditions (106) and (107) hold in this setting, under the assumptions of Theorem 5.

We use again that the ratio in (52) is the conditional density of L applied to ℓ given $Y = y$ under $\mathbb{P}_{\vartheta_0, \tau_0}$, and that the marginal $\mathbb{P}_{\vartheta_0, \tau_0}$ on Y is equivalent to the dominating measure ν . Hence we can write the integral in (106) as

$$\int_Y \check{\varphi}(y) \|S(y)\|_E g(y) \nu(dy) \quad \text{with} \quad g(Y) = \mathbb{E}_{\vartheta_0, \tau_0} [\|L\|_{\text{op}} \mid Y] .$$

By (C2) and (58), we obtain Condition (106). Similarly the integral in (107) reads, for all $\xi \in \Upsilon(\tilde{T})$,

$$\int_Y \check{\varphi}(y) g(y) \nu(dy) \quad \text{with} \quad g(Y) = \mathbb{E}_{\vartheta_0, \tau_0} [\|L^\nabla(\xi)\|_F \mid Y] .$$

Clearly, for this g , (C3) gives that, for some $\epsilon, C > 0$ $g(y) \leq C\epsilon^{-1} e^{\tilde{s}_{\vartheta_0, \tau_0}(y)}$ for ν -a.e. y , and (58) implies (107).

We can thus apply Proposition 7 and to conclude, we only need to check that the conditions of Assertions (i) and (ii) of Proposition 7 also applies, that is, we need to show that $w_{\tau_0}^\nabla$ is continuous on $\Upsilon(\tilde{T})$ in every direction and that (109) holds for all $\xi \neq \xi' \in \Upsilon(\tilde{T})$. The continuity of $w_{\tau_0}^\nabla$ follows from Lemma 12 since (120) holds by (58) (and we obtain in passing (61), which is exactly (122)). Finally, Condition (109) follows from (C4), since in the setting of Theorem 5, $\tilde{\varphi}(\ell) > 0$ for all $\ell \in Z$. ■

C.4.3 PROOF OF COROLLARY 10

The proof of Corollary 10 relies on some preliminary results. We first show in Proposition 8 below that Theorem 5 applies to a general mixture of Gaussian distributions (which notably includes the Student's t distribution family considered in Example 5).

Proposition 8 *Let \check{M} be a class of probability distributions on $(0, \infty)$ and suppose that for all $\check{\tau}_0 \in \check{M}$, there exists $\epsilon_0 > 0$ such that*

$$\int_0^\infty e^{\epsilon_0 z} \check{\tau}_0(dz) < \infty . \tag{123}$$

Let $d \geq 1$ and define for any $\vartheta = (m, \Sigma) \in \tilde{\mathbf{T}} = \mathbb{R}^d \times \mathcal{M}_{>0}(d)$

$$\check{k}(\vartheta, y) = \int_0^\infty \mathcal{N}(y; m, z^{-1}\Sigma) \check{\tau}_0(dz), \quad (124)$$

seen as a kernel density with respect to ν being the Lebesgue measure on \mathbb{R}^d . Then the following assertions hold.

(i) Let $k^{(1)}$ be ELM family given by Definition 2 with $F = E = \mathbb{R}^d \times \mathbb{R}^{d \times d}$, $E_0 = \tilde{\mathbf{T}} = \mathbb{R}^d \times \mathcal{M}_{>0}(d)$, $\Upsilon(m, \Sigma) = (\Sigma^{-1}m, \Sigma^{-1})$,

$$k^{(o)}(\xi, y) = \mathcal{N}(y; m, \Sigma) \quad \text{for all } \xi = \Upsilon(\vartheta) \text{ with } \vartheta = (m, \Sigma), \quad (125)$$

and \mathbf{M} the class of all push-forward probabilities τ_0 of $\check{\tau}_0 \in \check{\mathbf{M}}$ through the mapping $z \mapsto z\text{Id}$ defined from $(0, \infty)$ to $\mathcal{L}(E)$, where Id is the identity operator on E . Then, \mathbf{M} is a class of distributions on positive scalar operators on E and, for all $\vartheta = (m, \Sigma) \in \tilde{\mathbf{T}}$, we have $\check{k}(\vartheta, y) = k^{(1)}((\vartheta, \tau_0), y)$.

(ii) Setting, for all $y \in \mathbb{R}^d$ and $\vartheta_0 = (m_0, \Sigma_0) \in \tilde{\mathbf{T}}$,

$$q_{\vartheta_0}(y) = \frac{1}{2}(y - m_0)^T \Sigma_0^{-1}(y - m_0), \quad (126)$$

$$g_{\check{\tau}_0}(u, v) = \int_0^\infty z^u e^{-vz} \check{\tau}_0(dz) \quad \text{for all } u > -1 \text{ and } v > -\epsilon_0, \quad (127)$$

this ELM family satisfies (C1)–(C4) with

$$\tilde{m}_{\vartheta_0, \tau_0}(y) = \frac{g_{\check{\tau}_0}(d/2 + 1, q_{\vartheta_0}(y))}{g_{\check{\tau}_0}(d/2, q_{\vartheta_0}(y))}, \quad (128)$$

$$\tilde{s}_{\vartheta_0, \tau_0}(y) = \log \left(\frac{g_{\check{\tau}_0}(d/2, q_{\vartheta_0}(y) - \epsilon_0)}{g_{\check{\tau}_0}(d/2, q_{\vartheta_0}(y))} \right). \quad (129)$$

(iii) Theorem 5 applies for any $\theta_0 = (\vartheta_0, \tau_0) \in \tilde{\mathbf{T}} \times \mathbf{M}$ and any probability density function $\check{\varphi} : \mathbb{R}^d \rightarrow \mathbb{R}_+$ satisfying

$$\int_Y \check{\varphi}(y) \left((\|y\| + \|y\|^2) \tilde{m}_{\vartheta_0, \tau_0}(y) + e^{\tilde{s}_{\vartheta_0, \tau_0}(y)} \right) \nu(dy) < \infty, \quad (130)$$

in which case the argmax problem (62) has a unique solution $\xi^* = \Upsilon(\vartheta^*)$ with $\vartheta^* = (m^*, \Sigma^*)$ given by

$$m^* = \frac{\mathbb{E}_{\vartheta_0, \check{\tau}_0} [Z \tilde{Y}]}{\mathbb{E}_{\vartheta_0, \check{\tau}_0} [Z]}, \quad (131)$$

$$\Sigma^* = \mathbb{E}_{\vartheta_0, \check{\tau}_0} [Z \tilde{Y} \tilde{Y}^T] - \frac{\mathbb{E}_{\vartheta_0, \check{\tau}_0} [Z \tilde{Y}] \mathbb{E}_{\vartheta_0, \check{\tau}_0} [Z \tilde{Y}]^T}{\mathbb{E}_{\vartheta_0, \check{\tau}_0} [Z]}, \quad (132)$$

where, under $\mathbb{P}_{\vartheta_0, \check{\tau}_0}$, Z has density with respect to $\check{\tau}_0$ given by

$$\tilde{\varphi}(z) = \int_Y \check{\varphi}(y) \frac{z^{d/2} e^{-z q_{\vartheta_0}(y)}}{g_{\check{\tau}_0}(d/2, q_{\vartheta_0}(y))} dy \quad (133)$$

and the conditional distribution of \tilde{Y} given $Z = z$ has density

$$\psi[y|z] = \frac{1}{1+b} \frac{z^{d/2} e^{-z q_{\vartheta_0}(y)} \check{\varphi}(y)}{g_{\check{\tau}_0}(d/2, q_{\vartheta_0}(y)) \check{\varphi}(z)} + \frac{b}{1+b} \mathcal{N}(y; m_0, z^{-1} \Sigma_0) . \quad (134)$$

Proof We prove the claimed assertions successively.

Proof of Assertion (i) The spaces $F = E = \mathbb{R}^d \times \mathbb{R}^{d \times d}$ are endowed with the usual inner product

$$\langle (x, M), (x', M') \rangle_E = x^T x' + \text{Tr}(M^T M') .$$

We have $\text{Int} E_0 = E_0 = \tilde{\mathbf{T}} = \mathbb{R}^d \times \mathcal{M}_{>0}(d)$, and set, for all $y \in \mathbb{R}^d$, $S(y) = (y, -yy^T/2)$ and, for all $(x, M) \in \text{Int} E_0$, $A(x, M) = \frac{1}{2} (x^T M^{-1} x - \log |M|)$. In this setting, $\ell^T = \ell$, $\ell = \|\ell\|_{\text{op}} \text{Id}$ for τ_0 -almost all $\ell \in \mathbf{Z}$, and $\|L\|_{\text{op}}$ has distribution $\check{\tau}_0$ for $L \sim \tau_0$. Then, for all $\vartheta = (m, \Sigma) \in \tilde{\mathbf{T}}$ and τ_0 -almost all $\ell \in \mathbf{Z}$, we have $\ell \circ \Upsilon(m, \Sigma) = \Upsilon(m, \Sigma / \|\ell\|_{\text{op}})$. Consequently, for all $\vartheta \in \tilde{\mathbf{T}}$, $y \in \mathbb{R}^d$ and τ_0 -almost all $\ell \in \mathbf{Z}$, we have

$$\begin{aligned} k^{(o)}(\ell \circ \Upsilon(\vartheta), y) &= \left(\frac{1}{2\pi} \left| \|\ell\|_{\text{op}}^{-1} \Sigma \right| \right)^{1/2} e^{-(y-m)^T (\|\ell\|_{\text{op}}^{-1} \Sigma)^{-1} (y-m)/2} \\ &= \mathcal{N}(y; m, \|\ell\|_{\text{op}}^{-1} \Sigma) . \end{aligned} \quad (135)$$

Thus the obtained ELM family $k^{(1)}$ of Definition 2 satisfies $\check{k}(\vartheta, y) = k^{(1)}((\vartheta, \tau_0), y)$ for all $\vartheta = (m, \Sigma) \in \tilde{\mathbf{T}}$.

Proof of Assertion (ii) Assumption (C1) is obvious in the above setting. We then check (C4), (C2) and (C3), successively. For all $\xi \neq \xi' \in \Upsilon(\tilde{\mathbf{T}})$ and $\tau_0 \in \mathbf{M}$, we easily see that

$$\tau_0(\{\ell \in \mathbf{Z} : \ell(\xi - \xi') = 0\}) = \check{\tau}_0(\{0\}) = 0 ,$$

and (C4) follows. To check (C2), we compute, for all $\vartheta_0 = (m_0, \Sigma_0)$ and $\tau_0 \in \mathbf{M}$,

$$\begin{aligned} \mathbb{E}_{\vartheta_0, \tau_0} [\|L\|_{\text{op}} | Y] &= \int_0^\infty z \frac{\mathcal{N}(Y; m_0, z^{-1} \Sigma_0)}{\check{k}(\vartheta_0, Y)} \check{\tau}_0(dz) \\ &= \frac{\int z^{d/2+1} e^{-z q_{\vartheta_0}(Y)} \check{\tau}_0(dz)}{\int z^{d/2} e^{-z q_{\vartheta_0}(Y)} \check{\tau}_0(dz)} \end{aligned}$$

where we successively used (57) with (135) and (126). Hence (C2) follows from (127) and (128). Let us now check (C3). Straightforward computations yield, for all $\xi = (x, M) \in \Upsilon(\tilde{\mathbf{T}})$ and τ_0 -almost all $\ell \in \mathbf{Z}$,

$$\begin{aligned} \nabla A \circ \ell(\xi) &= (M^{-1} x, -(M^{-1} x x^T M^{-1} + \|\ell\|_{\text{op}}^{-1} M^{-1})/2) \\ \ell^\nabla(\xi) &= (\|\ell\|_{\text{op}} M^{-1} x, -(\|\ell\|_{\text{op}} M^{-1} x x^T M^{-1} + M^{-1})/2) \\ \|\ell^\nabla(\xi)\|_F &\leq (1 + \|\ell\|_{\text{op}}) \left\| (M^{-1} x, M^{-1} x x^T M^{-1} + M^{-1}) \right\|_F . \end{aligned}$$

It follows, that, for all $\vartheta_0 = (m_0, \Sigma_0) \in \tilde{\mathcal{T}}$, $\tau_0 \in \mathcal{M}$, $\xi = (x, M) \in \Upsilon(\tilde{\mathcal{T}})$ and $\epsilon > 0$,

$$\begin{aligned} \mathbb{E}_{\vartheta_0, \tau_0} \left[e^{\epsilon \|L^\nabla(\xi)\|_F} \mid Y \right] &\leq C_1 \mathbb{E}_{\vartheta_0, \tau_0} \left[e^{C_2 \epsilon \|L\|_{\text{op}}} \mid Y \right] \quad \mathbb{P}_{\vartheta_0, \tau_0} - \text{a.s.} \\ &= C_1 \log \left(\frac{g_{\tilde{\tau}_0}(d/2, q_{\vartheta_0}(y) - C_2 \epsilon)}{g_{\tilde{\tau}_0}(d/2, q_{\vartheta_0}(y))} \right) \quad \mathbb{P}_{\vartheta_0, \tau_0} - \text{a.s.}, \end{aligned}$$

where $C_1, C_2 > 0$ only depend on ξ . Hence, since we assumed (123), the definition of $g_{\tilde{\tau}_0}$ in (127) holds for all $v > -\epsilon_0$, and (C3) holds with $\tilde{s}_{\vartheta_0, \tau_0}$ as in (129) by taking $\epsilon < \epsilon_0/C_2$.

Proof of Assertion (iii) From the previous assertion, we can apply Theorem 5 and use (66) to solve the argmax problem (62) for a density $\check{\varphi}$ with respect to the Lebesgue measure on \mathbb{R}^d satisfying (58), which, since $S(y) = (y, -yy^T/2)$ and

$$\|S(y)\|_E = \left(\|y\|^2 + \frac{1}{4} \text{Tr}(yy^T yy^T) \right)^{1/2} = O(\|y\| + \|y\|^2),$$

is equivalent to (130). Let us compute both sides of the equation (66) which characterizes the solution $\xi^* = \Upsilon(\vartheta^*)$. The left-hand side of (66) reads

$$\mathbb{E}_{\vartheta^*, \tilde{\tau}_0} [L^T \circ S(Y)] = \mathbb{E}_{\vartheta^*, \tilde{\tau}_0} [Z S(Y)],$$

where $Z = \|L\|_{\text{op}}$ has density with respect to $\tilde{\tau}_0$ given by (133). The right-hand side of (66) reads

$$\mathbb{E}_{\vartheta_0, \tilde{\tau}_0} [L^T \circ S(\tilde{Y})] = \mathbb{E}_{\vartheta_0, \tilde{\tau}_0} [Z S(\tilde{Y})],$$

where Z is the same as above and the conditional distribution of \tilde{Y} given $Z = z$ has density given by (134). Finally, the solution $\xi^* = \Upsilon(\vartheta^*)$ is given by the equation

$$\mathbb{E}_{\vartheta^*, \tilde{\tau}_0} [Z S(Y)] = \mathbb{E}_{\vartheta_0, \tilde{\tau}_0} [Z S(\tilde{Y})] \quad (136)$$

and this equation leads to the solutions (131) and (132). Indeed, note that

$$\mathbb{E}_{\vartheta^*, \tilde{\tau}_0} [Y|Z] = m^* \quad \text{and} \quad \mathbb{E}_{\vartheta^*, \tilde{\tau}_0} [YY^T|Z] - m^*(m^*)^T = \Sigma^*/Z \quad (137)$$

Applying the first identity in (137), then the tower property and the identity (136) applied to the first component of $S(Y) = (Y, -YY^T/2)$,

$$m^* = \frac{\mathbb{E}_{\vartheta^*, \tilde{\tau}_0} [Z m^*]}{\mathbb{E}_{\vartheta^*, \tilde{\tau}_0} [Z]} = \frac{\mathbb{E}_{\vartheta^*, \tilde{\tau}_0} [ZY]}{\mathbb{E}_{\vartheta^*, \tilde{\tau}_0} [Z]} = \frac{\mathbb{E}_{\vartheta_0, \tilde{\tau}_0} [Z \tilde{Y}]}{\mathbb{E}_{\vartheta^*, \tilde{\tau}_0} [Z]} = \frac{\mathbb{E}_{\vartheta_0, \tilde{\tau}_0} [Z \tilde{Y}]}{\mathbb{E}_{\vartheta_0, \tilde{\tau}_0} [Z]}$$

where the last equality follows from the fact that under $\mathbb{P}_{\vartheta^*, \tilde{\tau}_0}$ or $\mathbb{P}_{\vartheta_0, \tilde{\tau}_0}$, Z has the density $\tilde{\varphi}$ with respect to $\tilde{\tau}_0$. Similarly, using the second equality in (137), then the tower property and the identity (136) applied to the second component of $S(Y) = (Y, -YY^T/2)$

$$\begin{aligned} \Sigma^* &= \mathbb{E}_{\vartheta^*, \tilde{\tau}_0} [Z(\Sigma^*/Z)] = \mathbb{E}_{\vartheta^*, \tilde{\tau}_0} [ZYY^T] - \mathbb{E}_{\vartheta^*, \tilde{\tau}_0} [Z] m^*(m^*)^T \\ &= \mathbb{E}_{\vartheta_0, \tilde{\tau}_0} [Z \tilde{Y} \tilde{Y}^T] - \mathbb{E}_{\vartheta^*, \tilde{\tau}_0} [Z] m^*(m^*)^T \end{aligned}$$

This proves (131) and (132).

■

Remark 5 *The probability density function (124) is called an elliptical distribution. In particular if $\check{\tau}_0$ is the χ^2 distribution with a_0 degrees of freedom, then it is the density of the Student's t distribution with parameter (m, Σ, a_0) . In this case we have*

$$g_{\check{\tau}_0}(u, v) = \frac{(a_0/2)^{a_0/2} \Gamma(a_0/2 + u)}{(a_0/2 + v)^{a_0/2+u} \Gamma(a_0/2)},$$

and $\tilde{m}_{\vartheta_0, \tau_0}$ and $\tilde{s}_{\vartheta_0, \tau_0}$ in (128) and (129) read, for any $\epsilon_0 \in (0, a_0/2]$,

$$\begin{aligned} \tilde{m}_{\vartheta_0, \tau_0}(y) &= \frac{a_0/2 + d/2}{a_0/2 + q_{\vartheta_0}(y)} = O(1), \\ \tilde{s}_{\vartheta_0, \tau_0}(y) &= (a_0/2 + d/2) \log \left(\frac{a_0/2 + q_{\vartheta_0}(y)}{a_0/2 - \epsilon_0 + q_{\vartheta_0}(y)} \right) = O(1), \end{aligned}$$

so that Condition (130) boils down to $\int_Y \check{\varphi}(y) \|y\|^2 \nu(dy) < \infty$.

The following lemma is used in Example 5.

Lemma 13 *Define, for all $x \in (0, \infty)$,*

$$\kappa(x) = \log(x) + \Gamma'(x)/\Gamma(x).$$

Then κ is increasing and bijective from $(0, \infty)$ to \mathbb{R} .

Proof This result follows from the fact that the digamma function $x \mapsto \Gamma'(x)/\Gamma(x)$ is increasing from $(0, \infty)$ to \mathbb{R} . ■

Proof of Corollary 10 We apply Theorem 2 in the setting of Example 5. Condition (18) holds by Theorem 3 and the update (i) in Example 5.

The remainder of the proof consists in proving Condition (19) for given $n \geq 1$ and $j = 1, \dots, J$. We apply Proposition 8, whose Assertion (i) with Remark 5 provides an ELM family $k^{(1)}$ such that, for any $\vartheta = (m, \Sigma) \in \mathbb{R}^d \times \mathcal{M}_{>0}(d)$ and $a > 0$, $\check{k}(\vartheta, y)$ in (124) with $\check{\tau}_0 = \check{\tau}_a$ given by (73) is $k^{(1)}((\vartheta, \tau_0), y)$ but also $k(\theta, y)$ in Example 5 for $\theta = (m, \Sigma, a)$. We will use Proposition 4 to this ELM family with $\check{\varphi} = \check{\varphi}_{j,n}^{(\alpha)}$, $\theta = (\vartheta, \tau)$ defined with $\vartheta = (m_{j,n+1}, \Sigma_{j,n+1})$ and τ the measure on obtained by pushing forward $\check{\tau}_{a_{j,n+1}}$ through the mapping $z \mapsto z\text{Id}$ and $\theta_0 = (\vartheta_0, \tau_0)$ defined similarly but using the parameters $(m_{j,n}, \Sigma_{j,n}, a_{j,n})$. Then conclusion (49) of Proposition 4 is exactly Condition (19) with k as in Example 5. Hence, to conclude the proof we only need to show that Proposition 4 applies in this setting. Namely, we need to prove (50) and (51), which respectively read

$$\int_Y \left(\int_0^\infty \varphi_{j,n}(z, y) \log \left(\frac{d\check{\tau}_{a_{j,n+1}}(z)}{d\check{\tau}_{a_{j,n}}(z)} \right) \check{\tau}_{a_{j,n}}(dz) \right) dy \geq 0, \quad (138)$$

$$\int_Y \left(\int_Z \varphi_{j,n}(\|\ell\|_{\text{op}}, y) \log \left(\frac{k^{(\alpha)}(\ell \circ \Upsilon(\vartheta), y)}{k^{(\alpha)}(\ell \circ \Upsilon(\vartheta_0), y)} \right) \tau_{a_{j,n}}(d\ell) \right) dy \geq 0, \quad (139)$$

where $\varphi_{j,n}$ is defined by (71) and $k^{(o)}$ is the Gaussian canonical kernel defined by (125). We thus conclude with the proof of these two conditions in the opposite order.

Proof of (139) Assertion (ii) of Proposition 8 with Remark 5 gives that the ELM family $k^{(1)}$ satisfies (C1)–(C4) with $\tilde{m}_{\vartheta_0, \tau_0}$ and $\tilde{s}_{\vartheta_0, \tau_0}$ such that Condition (130) boils down to $\int \check{\varphi}(y) \|y\|^2 d\nu < \infty$. This latter condition is satisfied with $\check{\varphi} = \check{\varphi}_{j,n}^{(\alpha)}$ as a consequence of (67) by using Lemma 8. Therefore Condition (130) holds and Assertion (iii) of Proposition 8 provides the solution of the argmax problem (62). A careful comparison of (131) and (132) with (68) and (69) tells us that $\vartheta = (m_{j,n+1}, \Sigma_{j,n+1})$ is the unique argmax of (62). Applying Corollary 9 in this setting finally gives us (139).

Proof of (138) Since the left-hand side of (138) is zero for $a = a_{j,n}$, it suffices to show that $a_{j,n+1}$ in (70) satisfies

$$a_{j,n+1} = \operatorname{argmax}_{a \in (0, \infty)} \int_Y \left(\int_0^\infty \varphi_{j,n}(z, y) \log \left(\frac{d\check{\tau}_a}{d\check{\tau}_{a_{j,n}}}(z) \right) \check{\tau}_{a_{j,n}}(dz) \right) dy. \quad (140)$$

By (73), we have, for all $z > 0$,

$$\log \left(\frac{d\check{\tau}_a}{d\check{\tau}_{a_{j,n}}}(z) \right) = \frac{a}{2} \log \left(\frac{a}{2} \right) + \log \Gamma \left(\frac{a}{2} \right) + \left(\frac{a}{2} - \frac{a_{j,n}}{2} \right) (\log z - z) + C_{j,n},$$

where $C_{j,n}$ is a positive constant not depending on a . The derivative of $x \mapsto x \log(x) + \log \Gamma(x)$ is κ as defined in Step (ii) of Example 5 and in Lemma 13. It is easy to check that $\int \varphi_{j,n}(z, y) \check{\tau}_{a_{j,n}}(dz) dy = \int \check{\varphi}(y) dy = 1$. Thus (140) follows from (70) with Lemma 13 if we can show that

$$\int_Y \left(\int_0^\infty \varphi_{j,n}(z, y) (|\log z| + z) \check{\tau}_{a_{j,n}}(dz) \right) dy < \infty.$$

In fact, we will show that there exists $\epsilon > 0$ such that for all $t \in (-\epsilon, 1]$,

$$\int_Y \left(\int_0^\infty \varphi_{j,n}(z, y) z^t \check{\tau}_{a_{j,n}}(dz) \right) dy < \infty, \quad (141)$$

which indeed implies the previous display. By definition of $\varphi_{j,n}$ in (71), we find that (141) is implied by

$$I(t) := \int_Y \left(\int_0^\infty \check{\varphi}_{j,n}^{(\alpha)}(y) z^{t+d/2} e^{-zq_{j,n}(y)} \left(1 + \frac{2}{a_{j,n}} q_{j,n}(y) \right)^{(a_{j,n}+d)/2} \check{\tau}_{a_{j,n}}(dz) \right) dy < \infty,$$

where we set $q_{j,n}(y) = \frac{1}{2}(y - m_{j,n})^T \Sigma_{j,n}(y - m_{j,n})$. Since, as in Remark 5, for all $t > -d/2$,

$$\int_0^\infty z^{t+d/2} e^{-zq_{j,n}(y)} \check{\tau}_{a_{j,n}}(dz) = \frac{(a_{j,n}/2)^{a_{j,n}/2} \Gamma(a_{j,n}/2 + t + d/2)}{(a_{j,n}/2 + q_{j,n}(y))^{a_{j,n}/2 + t + d/2} \Gamma(a_{j,n}/2)},$$

we get that for all $t > -d/2$, there exists $C'_{j,n}(t) > 0$ such that

$$I(t) = C'_{j,n}(t) \int_Y \check{\varphi}_{j,n}^{(\alpha)}(y) \left(1 + \frac{2}{a_{j,n}} q_{j,n}(y) \right)^{-t} dy.$$

Since $q_{j,n}(y)$ is a quadratic form and we already checked that $\int \check{\varphi}_{j,n}^{(\alpha)}(y) (1 + \|y\|^2) d\nu < \infty$ in the proof of (139), we finally find that $I(t) < \infty$ if $t > -d/2$ and $t \geq -2$, which concludes the proof of (141). \blacksquare

D. Additional Numerical Experiments

In this section we provide further plots based on the numerical experiments from Section 8.




Figure 5: Error bounds for the Monte Carlo estimate of the VR Bound in the MG-IS-n approach (fixed mixture weights) when considering each of the target distributions (i), (ii) and (iii).

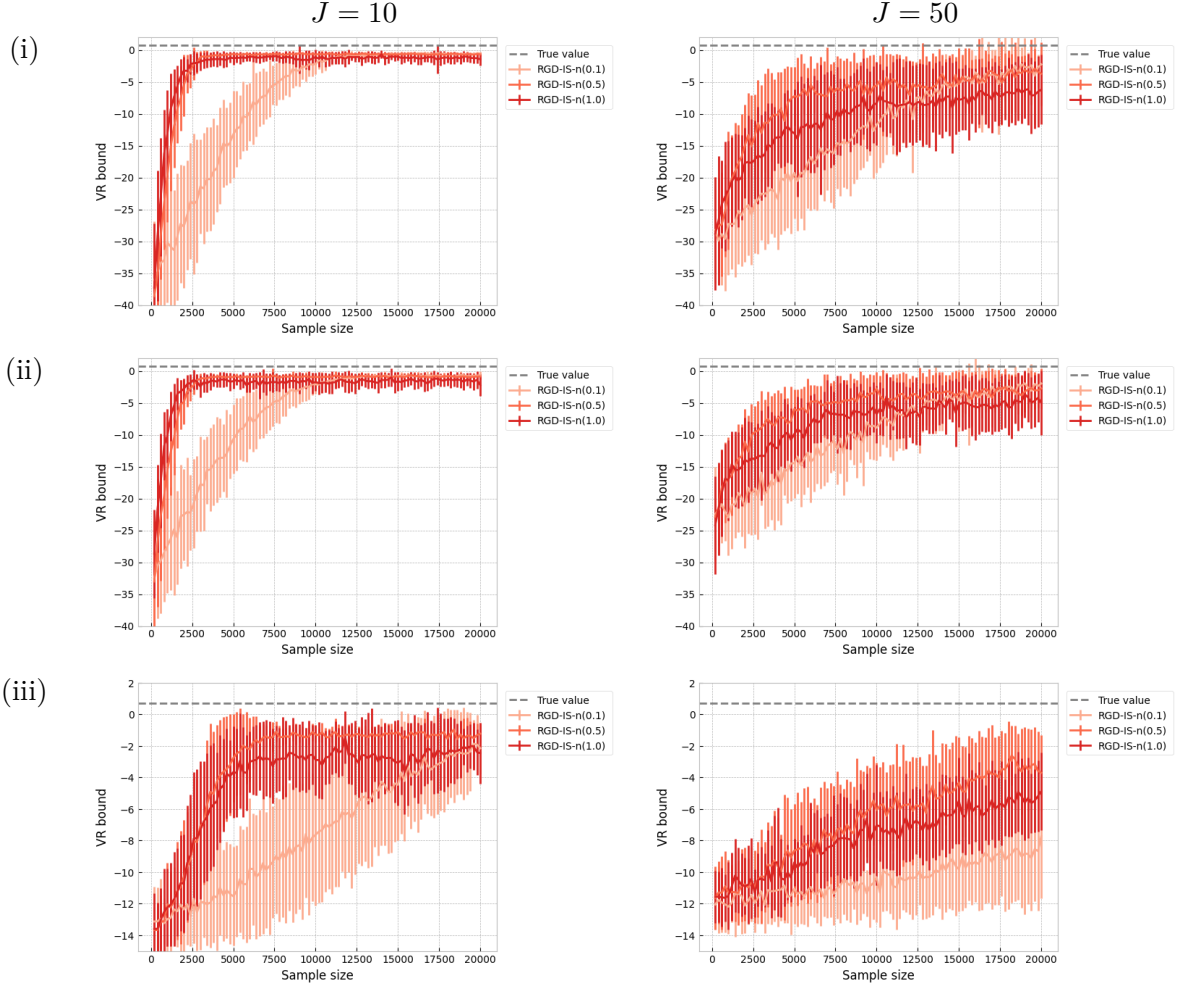


Figure 6: Error bounds for the Monte Carlo estimate of the VR Bound in the RGD-IS-n approach (fixed mixture weights) when considering each of the target distributions (i), (ii) and (iii).

References

Pierre Alquier and James Ridgway. Concentration of tempered posteriors and of their variational approximations. *The Annals of Statistics*, 48(3):1475 – 1497, 2020. doi: 10.1214/19-AOS1855. URL <https://doi.org/10.1214/19-AOS1855>.

David M. Blei, Alp Kucukelbir, and Jon D. McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.

Alexander Buchholz, Florian Wenzel, and Stephan Mandt. Quasi-Monte Carlo variational inference. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*,

| | | $J = 10$ | | | $J = 50$ | | |
|-------|-------------------------|----------------|----------------|----------------|----------------|----------------|----------------|
| | | $\gamma = 0.1$ | $\gamma = 0.5$ | $\gamma = 1.0$ | $\gamma = 0.1$ | $\gamma = 0.5$ | $\gamma = 1.0$ |
| (i) | RGD-IS-n(γ) | 0.372 | 0.510 | 0.384 | -0.616 | -0.713 | -0.778 |
| | MG-IS-n(γ) | 1.104 | 1.074 | 0.387 | 1.135 | -0.077 | -0.060 |
| | RGD-IS-unif(γ) | 0.359 | 0.469 | 0.458 | -0.688 | -0.670 | -0.583 |
| | MG-IS-unif(γ) | -0.200 | -0.229 | -0.515 | -1.500 | -1.462 | -1.246 |
| (ii) | RGD-IS-n(γ) | -0.025 | -0.056 | -0.087 | -1.027 | -0.997 | -0.969 |
| | MG-IS-n(γ) | -0.270 | -0.126 | 0.235 | -0.269 | -0.417 | -0.487 |
| | RGD-IS-unif(γ) | -0.121 | -0.111 | 0.052 | -1.097 | -0.966 | -0.883 |
| | MG-IS-unif(γ) | -1.120 | -0.938 | -0.957 | -1.764 | -1.889 | -1.192 |
| (iii) | RGD-IS-n(γ) | -0.329 | -0.197 | -0.238 | -1.691 | -1.612 | -1.637 |
| | MG-IS-n(γ) | 1.101 | 0.758 | 0.524 | 0.181 | -0.181 | 0.893 |
| | RGD-IS-unif(γ) | -0.370 | -0.224 | -0.212 | -1.708 | -1.627 | -1.649 |
| | MG-IS-unif(γ) | -1.211 | -1.313 | -1.083 | -2.013 | -1.882 | -0.491 |

Table 3: LogMSE for the RGD and the MG approaches ($\eta = 0.1$) when considering each of the target distributions (i), (ii) and (iii).

| | | $J = 10$ | | | $J = 50$ | | |
|-------|-------------------------|---------------|---------------|---------------|---------------|---------------|---------------|
| | | $\eta = 0.05$ | $\eta = 0.1$ | $\eta = 0.5$ | $\eta = 0.05$ | $\eta = 0.1$ | $\eta = 0.5$ |
| (i) | RGD-IS-n(γ) | 0.045 | 0.510 | 1.299 | -1.355 | -0.713 | 0.924 |
| | MG-IS-n(γ) | 0.087 | 1.074 | 1.343 | -1.205 | -0.077 | 1.329 |
| | RGD-IS-unif(γ) | -0.018 | 0.469 | 1.328 | -1.385 | -0.670 | 0.928 |
| | MG-IS-unif(γ) | -1.244 | -0.229 | 1.100 | -2.524 | -1.462 | 0.309 |
| (ii) | RGD-IS-n(γ) | -0.096 | -0.056 | 0.522 | -1.509 | -0.997 | 0.542 |
| | MG-IS-n(γ) | -0.629 | -0.126 | 0.100 | -1.430 | -0.417 | 0.348 |
| | RGD-IS-unif(γ) | -0.195 | -0.111 | 0.489 | -1.542 | -0.966 | 0.529 |
| | MG-IS-unif(γ) | -1.814 | -0.938 | -0.149 | -1.711 | -1.889 | -0.282 |
| (iii) | RGD-IS-n(γ) | -0.091 | -0.197 | 0.339 | -1.596 | -1.612 | -0.282 |
| | MG-IS-n(γ) | -0.772 | 0.758 | 1.358 | -0.878 | -0.181 | 0.927 |
| | RGD-IS-unif(γ) | -0.113 | -0.224 | 0.322 | -1.611 | -1.627 | -0.300 |
| | MG-IS-unif(γ) | -1.608 | -1.313 | -0.253 | -1.879 | -1.882 | -0.716 |

Table 4: LogMSE for the RGD and the MG approaches ($\gamma = 0.5$) when considering each of the target distributions (i), (ii) and (iii).

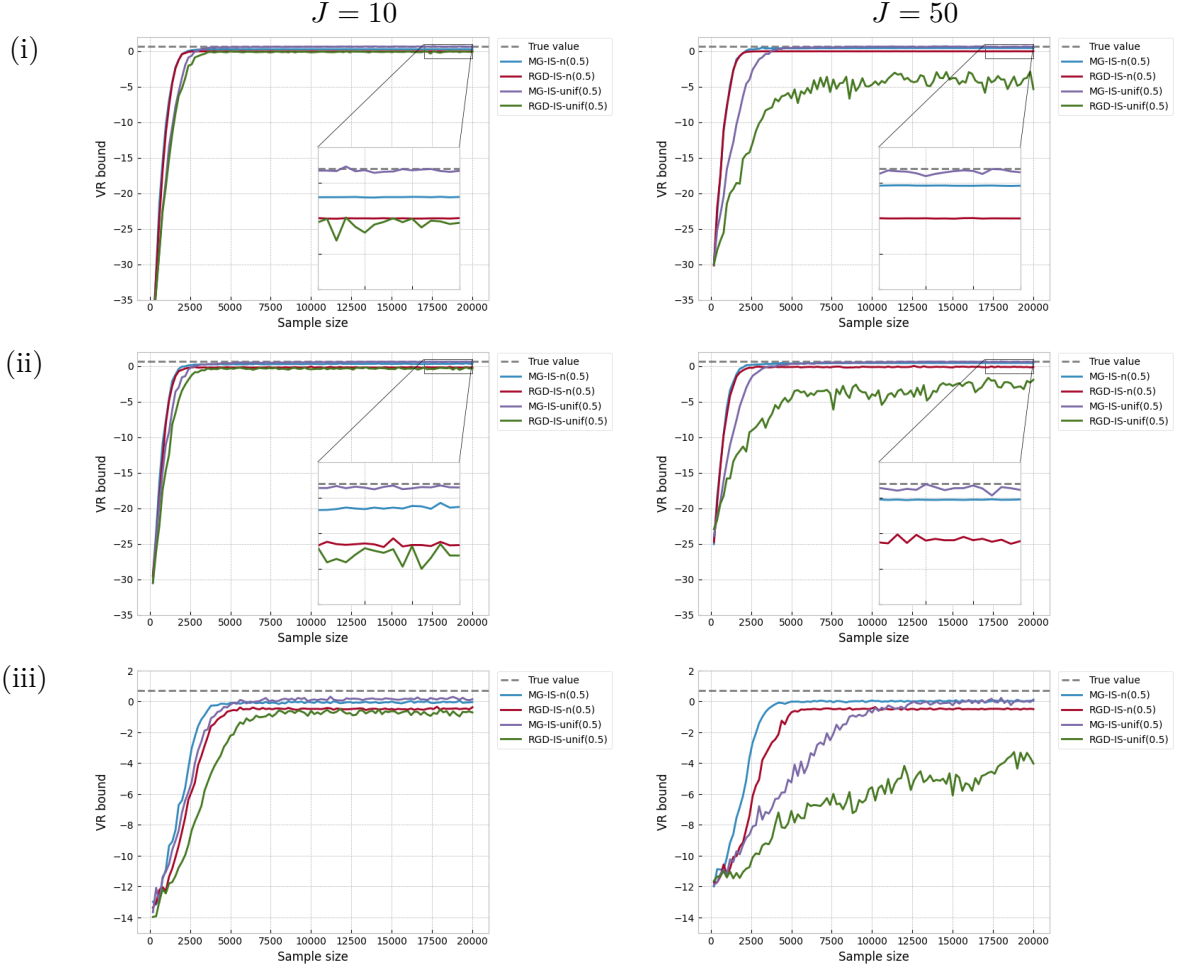


Figure 7: Monte Carlo estimate of the VR Bound for the RGD and the MG approaches ($\eta = 0.05$) when considering each of the target distributions (i), (ii) and (iii).

pages 668–677. PMLR, 10–15 Jul 2018. URL <https://proceedings.mlr.press/v80/buchholz18a.html>.

Trevor Campbell and Xinglong Li. Universal boosting variational inference. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

Olivier Cappé, Randal Douc, Arnaud Guillin, Jean-Michel Marin, and Christian P Robert. Adaptive importance sampling in general mixture classes. *Statistics and Computing*, 18(4):447–459, 2008.

Andrzej Cichocki and Shun-ichi Amari. Families of alpha- beta- and gamma- divergences: Flexible and robust measures of similarities. *Entropy*, 12(6):1532–1568, 2010.

Kamélia Daudel and Randal Douc. Mixture weights optimisation for alpha-divergence variational inference. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and

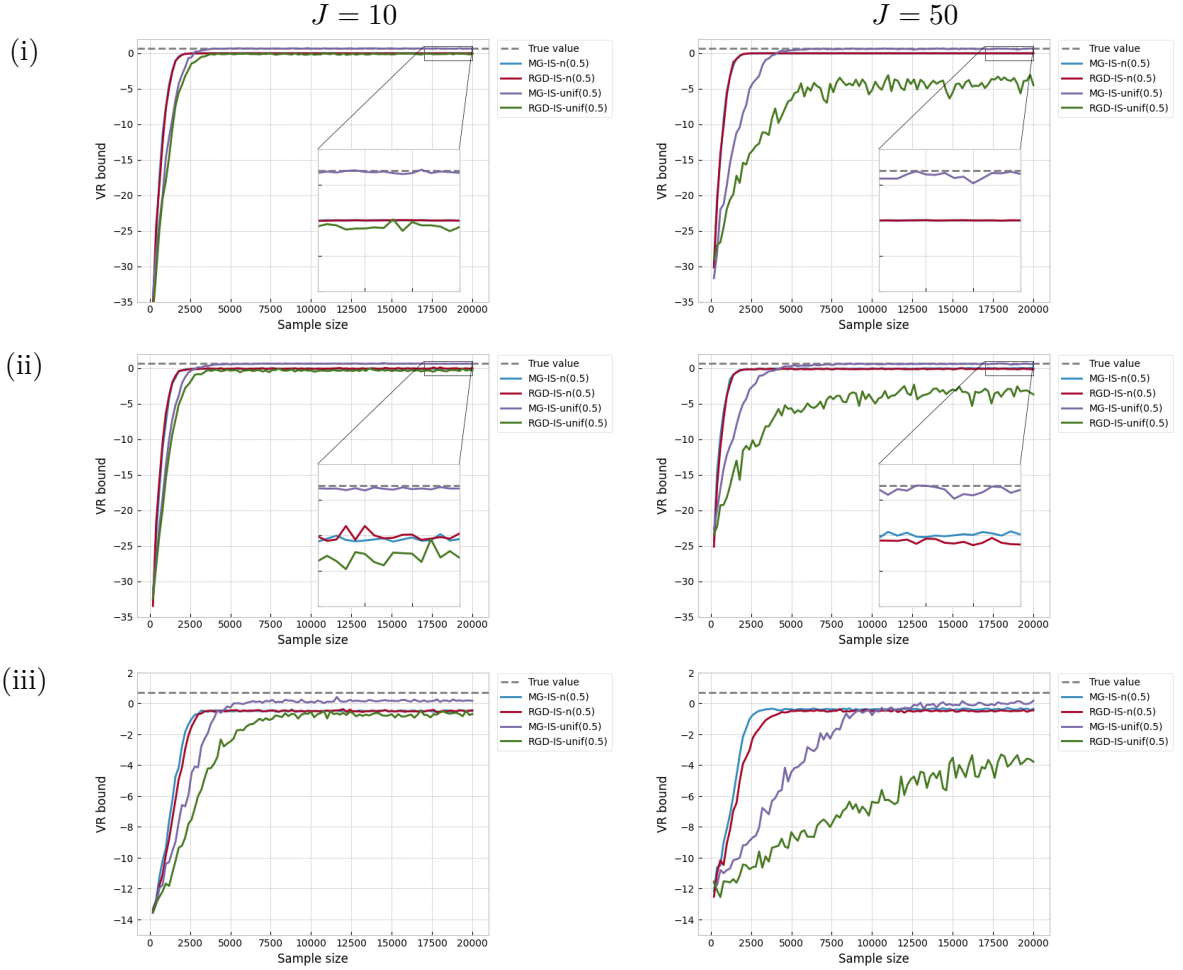


Figure 8: Monte Carlo estimate of the VR Bound for the RGD and the MG approaches ($\eta = 0.5$) when considering each of the target distributions (i), (ii) and (iii).

J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 4397–4408. Curran Associates, Inc., 2021.

Kamélia Daudel, Randal Douc, and François Portier. Infinite-dimensional gradient-based descent for alpha-divergence minimisation. *The Annals of Statistics*, 49(4):2250 – 2270, 2021.

Adj Bousso Dieng, Dustin Tran, Rajesh Ranganath, John Paisley, and David Blei. Variational inference via χ upper bound minimization. In *Advances in Neural Information Processing Systems 30*, pages 2732–2741, 2017.

Justin Domke. Provable smoothness guarantees for black-box variational inference. In *International Conference on Machine Learning*, pages 2587–2596. PMLR, 2020.




Figure 9: Monte Carlo estimate of the VR Bound for the RGD and the MG approaches when considering the Bayesian Logistic Regression on the Covertypes data set.

Samuel Gershman, Matt Hoffman, and David Blei. Nonparametric variational inference. In *Proceedings of the 29th International Conference on Machine Learning*, Edinburgh, Scotland, UK, 2012.

Jose Hernandez-Lobato, Yingzhen Li, Mark Rowland, Thang Bui, Daniel Hernandez-Lobato, and Richard Turner. Black-box alpha divergence minimization. In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1511–1520, New York, USA, 2016.

Matthew D. Hoffman, David M. Blei, Chong Wang, and John Paisley. Stochastic variational inference. *Journal of Machine Learning Research*, 14(4):1303–1347, 2013.

Matthew James Beal. *Variational algorithms for approximate Bayesian inference*, PhD thesis, The Gatsby Computational Neuroscience Unit, University of London, UK, 2003.

Ghassen Jerfel, Serena Wang, Clara Fannjiang, Katherine A. Heller, Yian Ma, and Michael I. Jordan. Variational refinement for importance sampling using the forward kullback-leibler divergence. In *37th Conference on Uncertainty in Artificial Intelligence (UAI 2021)*, forthcoming, 2021.

Michael I. Jordan, Zoubin Ghahramani, Tommi S. Jaakkola, and Lawrence K. Saul. An introduction to variational methods for graphical models. *Machine Learning*, 37(2): 183–233, 1999.

- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *Proceedings of the 2nd International Conference on Learning Representations*, 2014.
- Yingzhen Li and Richard E Turner. Rényi divergence variational inference. In *Advances in Neural Information Processing Systems 29*, pages 1073–1081, 2016.
- Tom Minka. Power ep. Technical Report MSR-TR-2004-149, Microsoft Research Ltd, Cambridge, UK, 2004.
- Tom Minka. Divergence measures and message passing. Technical Report MSR-TR-2005-173, Microsoft Research Ltd., Cambridge, UK, 2005.
- Dennis Prangle. Distilling importance sampling. *Available on arXiv:1910.03632v3*, last accessed on 10 Sep 2021, 2021.
- Rajesh Ranganath, Sean Gerrish, and David Blei. Black Box Variational Inference. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*, volume 33 of *Proceedings of Machine Learning Research*, pages 814–822, Reykjavik, Iceland, 2014. PMLR.
- Alfréd Rényi. On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, pages 547–561, Berkeley, Calif., 1961. University of California Press.
- Tim van Erven and Peter Harremoës. Rényi divergence and kullback-leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, 2014.
- Dilin Wang, Hao Liu, and Qiang Liu. Variational inference with tail-adaptive f-divergence. In *Advances in Neural Information Processing Systems 31*, pages 5737–5747, 2018.
- Yuling Yao, Aki Vehtari, Daniel Simpson, and Andrew Gelman. Yes, but did it work?: Evaluating variational inference. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 5581–5590, Stockholmsmässan, Stockholm Sweden, 2018. PMLR.
- Cheng Zhang, Judith Butepage, Hedvig Kjellstrom, and Stephan Mandt. Advances in variational inference. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(8):2008–2026, 2019.
- Huaiyu Zhu and Richard Rohwer. Information geometric measurements of generalisation. Technical Report NCRG/4350, Neural Computing Research Group Dept. of Computer Science and Applied Mathematics Aston University, Aston Triangle, UK, 1995a.
- Huaiyu Zhu and Richard Rohwer. Bayesian invariant measurements of generalization. *Neural Processing Letters*, 2:28–31, 1995b.