## A Proofs

We first introduce two useful known lemmas, and prove the propositions in their order of appearance.

## A. 1 Useful lemmas

First, under Assumption 1, we note that the soft $C$-transforms are uniformly contracting on the distribution space $\mathcal{P}(\mathcal{X})$. This is clarified in the following lemma, extracted from Vialard (2019), Proposition 19. We refer the reader to the original references for proofs.

Lemma 1. Unser Assumption 1, let $\kappa=1-\exp (-L \operatorname{diam}(\mathcal{X}))$. For all $\hat{\alpha} \in \mathcal{P}(\mathcal{X})$ and $\hat{\beta} \in \mathcal{P}(\mathcal{X})$, for all $f, f^{\prime}, g, g^{\prime} \in \mathcal{C}(\mathcal{X})$,

$$
\left\|T_{\hat{\alpha}}\left(f^{\prime}\right)-T_{\hat{\alpha}}\left(f^{\prime}\right)\right\|_{\mathrm{var}} \leqslant \kappa\left\|f-f^{\prime}\right\|_{\mathrm{var}}, \quad\left\|T_{\hat{\beta}}(g)-T_{\hat{\beta}}\left(g^{\prime}\right)\right\|_{\mathrm{var}} \leqslant \kappa\left\|g-g^{\prime}\right\|_{\mathrm{var}}
$$

We will also need a uniform law of large numbers for functions. The following lemma is a consequence of Example 19.7 and Lemma 19.36 of Van der Vaart (2000), and is copied in Lemma B. 6 in Mairal (2013).
Lemma 2. Under Assumption 1, let $\left(f_{t}\right)_{t}$ be an i.i.d sequence in $\mathcal{C}(\mathcal{X})$, such that $\mathbb{E}\left[f_{0}\right]=f \in \mathcal{C}(\mathcal{X})$. Then there exists $A>0$ such that, for all $n>0$,

$$
\mathbb{E} \sup _{x \in \mathcal{X}}\left|\frac{1}{n} \sum_{i=1}^{n} f_{i}(x)-f(x)\right| \leqslant \frac{A}{\sqrt{n}}
$$

Finally, we need a result on running averages using the sequence $\left(\eta_{t}\right)_{t}$. The following result stems from a simple Abel transform of the law of large number, and is established by Mairal (2013), Lemma B.7.

Lemma 3. Let $\left(\eta_{t}\right)_{t}$ be a sequence of weights meeting Assumption 2. Let $\left(X_{t}\right)_{t}$ be an i.i.d sequence of real-valued random variables with existing first moment $\mathbb{E}\left[X_{0}\right]$. We consider the sequence $\left(\bar{X}_{t}\right)_{t}$ defined by $\bar{X}_{0} \triangleq X_{0}$ and

$$
\bar{X}_{t} \triangleq\left(1-\eta_{t}\right) \bar{X}_{t-1}+\eta_{t} X_{t} .
$$

Then $\bar{X}_{t} \rightarrow_{t \rightarrow \infty} \mathbb{E}\left[X_{0}\right]$.

## A. 2 Proof of Proposition 1

Proof. We use Theorem 1 from Diaconis and Freedman (1999). For this, we simply note that the space $\mathcal{C}(\mathcal{X}) \times \mathcal{C}(\mathcal{X})$ in which the chain $x_{t} \triangleq\left(f_{t}, g_{t}\right)_{t}$, endowed with the metric $\rho\left(\left(f_{1}, g_{1}\right),\left(f_{2}, g_{2}\right)\right)=$ $\left\|f_{1}-f_{2}\right\|_{\mathrm{var}}+\left\|g_{1}-g_{2}\right\|_{\mathrm{var}}$, is complete and separable (the countable set of polynomial functions are dense in this space, for example). We consider the operator $A_{\theta} \triangleq T_{\hat{\beta}}\left(T_{\hat{\alpha}}(\cdot)\right) . \theta \triangleq(\hat{\alpha}, \hat{\beta})$ denotes the random variable that is sampled at each iteration. We have the following recursion:

$$
x_{t+2}=A_{\theta_{t}}\left(x_{t}\right)
$$

From Lemma 1, for all $\hat{\alpha} \in \mathcal{P}(\mathcal{X}), \hat{\beta} \in \mathcal{P}(\mathcal{X})$, $A_{\theta}$ with $\theta=(\hat{\alpha}, \hat{\beta})$ is contracting, with module $\kappa_{\theta}<\kappa<1$. Therefore

$$
\int_{\theta} \kappa_{\theta} \mathrm{d} \mu(\theta)<1, \quad \int_{\theta} \log \kappa_{\theta} \mathrm{d} \mu(\theta)<0
$$

Finally, we note, for all $f \in \mathcal{C}(\mathcal{X})$

$$
\left\|T_{\beta}\left(T_{\hat{\alpha}}(f)\right)\right\|_{\infty} \leqslant\|f\|_{\infty}+2 \max _{x, y \in \mathcal{X}} C(x, y)
$$

therefore $\rho\left(A_{\theta}\left(x_{0}\right), x_{0}\right) \leqslant 2\left\|x_{0}\right\|_{\infty}+2 \max _{x, y \in \mathcal{X}} C(x, y)$ for all $\theta(\hat{\alpha}, \hat{\beta})$. The regularity condition of the theorem are therefore met. Each of the induced Markov chains $\left(f_{2 t}, g_{2 t}\right)_{t}$ and $\left(f_{2 t+1}, g_{2 t+1}\right)_{t}$ has a unique stationary distribution. These stationary distributions are the same: the stationary distribution is independent of the initialisation and both sequences differs only by their initialisation. Therefore $\left(f_{t}, g_{t}\right)_{t}$ have a unique stationary distribution $\left(F_{\infty}, G_{\infty}\right)$.

## A. 3 Proof of Proposition 2

For presentation purpose, we first show that the "slowed-down" online Sinkhorn algorithm converges in the absence of noise. We then turn to prove Proposition 2.

## A.3.1 Noise-free online Sinkhorn

Proposition 5. We suppose that $\hat{\alpha}_{t}=\alpha, \hat{\beta}_{t}=\beta$ for all $t$. Then the updates (6) yields a (deterministic) sequence $\left(f_{t}, g_{t}\right)_{t}$ such that

$$
\left\|\hat{f}_{t}-f^{\star}\right\|_{v a r}+\left\|\hat{g}_{t}-g^{\star}\right\|_{v a r} \rightarrow 0, \quad \frac{1}{2}\left\langle\alpha, f_{t}+T_{\alpha}\left(\hat{g}_{t}\right)\right\rangle+\left\langle\beta, \hat{g}_{t}+T_{\beta}\left(f_{t}\right)\right\rangle \rightarrow \mathcal{W}(\alpha, \beta)
$$

Note that, as we perform simultaneous updates, we only obtain the convergence of $f_{t} \rightarrow f^{\star}+A$, and $g_{t} \rightarrow g^{\star}$, where $f^{\star}$ and $g^{\star}$ are solutions of (1) and $A$ is a constant depending on initialization.
The "slowed-down" Sinkhorn iterations converge toward an optimal potential couple, up to a constant factor: this stems from the fact that we apply contractions in the space $\left(\mathcal{C}(\mathcal{X}),\|\cdot\|_{\text {var }}\right)$ with a contraction factor that decreases sufficiently slowly.

Proof. We write $\left(f_{t}, g_{t}\right)_{t}$ the sequence of iterates. Given a pair of optimal potentials $\left(f^{\star}, g^{\star}\right)$, we write $u_{t} \triangleq f_{t}-f^{\star}, v_{t} \triangleq g_{t}-g^{\star}, u_{t}^{T} \triangleq T_{\alpha}\left(f_{t}\right)-g^{\star}$ and $v_{t}^{T} \triangleq T_{\alpha}\left(g_{t}\right)-f^{\star}$. For all $t>0$, we observe that

$$
\begin{aligned}
\max u_{t+1} & =-\log \min \exp \left(-u_{t+1}\right) \\
& =-\log \left(\min \left(\left(1-\eta_{t}\right) \exp \left(-u_{t}\right)+\eta_{t} \exp \left(-v_{t}^{T}\right)\right)\right) \\
& \leqslant-\log \left(\left(1-\eta_{t}\right) \min \exp \left(-u_{t}\right)+\eta_{t} \min \exp \left(-v_{t}^{T}\right)\right) \\
& \leqslant-\left(1-\eta_{t}\right) \log \min \exp \left(-u_{t}\right)-\eta_{t} \log \min \exp \left(-v_{t}^{T}\right) \\
& =\left(1-\eta_{t}\right) \max u_{t}+\eta_{t} \max v_{t}^{T}
\end{aligned}
$$

where we have used the algorithm recursion on the second line, $\min f+g \geqslant \min f+\min g$ on the third line and Jensen inequality on the fourth line. Similarly

$$
\min u_{t+1} \geqslant\left(1-\eta_{t}\right) \min u_{t}+\eta_{t} \min v_{t}^{T}
$$

and mirror inequalities hold for $v_{t}$. Summing the four inequalities, we obtain

$$
\begin{aligned}
e_{t+1} & \triangleq\left\|u_{t+1}\right\|_{\mathrm{var}}+\left\|v_{t+1}\right\|_{\mathrm{var}} \\
& =\max u_{t+1}-\min u_{t+1}+\max v_{t+1}-\min v_{t+1} \\
& \leqslant\left(1-\eta_{t}\right)\left(\left\|u_{t}\right\|_{\mathrm{var}}+\left\|v_{t}\right\|_{\mathrm{var}}\right)+\eta_{t}\left(\left\|u_{t}^{T}\right\|_{\mathrm{var}}+\left\|v_{t}^{T}\right\|_{\mathrm{var}}\right) \\
& \leqslant\left(1-\eta_{t}\right)\left(\left\|u_{t}\right\|_{\mathrm{var}}+\left\|v_{t}\right\|_{\mathrm{var}}\right)+\eta_{t} \kappa\left(\left\|u_{t}\right\|_{\mathrm{var}}+\left\|v_{t}\right\|_{\mathrm{var}}\right),
\end{aligned}
$$

where we use the contractivity of the soft- $C$-transform, that guarantees that there exists $\kappa<1$ such that $\left\|v_{t}^{T}\right\|_{\text {var }} \leqslant \kappa\left\|v_{t}\right\|_{\text {var }}$ and $\left\|u_{t}^{T}\right\|_{\text {var }} \leqslant \kappa\left\|u_{t}\right\|_{\text {var }}$ (Peyré and Cuturi, 2019).
Unrolling the recursion above, we obtain

$$
\log e_{t}=\sum_{s=1}^{t} \log \left(1-\eta_{t}(1-\kappa)\right)+\log \left(e_{0}\right) \rightarrow-\infty
$$

provided that $\sum \eta_{t}=\infty$. The proposition follows.
Proof of Proposition 2. For discrete realizations $\hat{\alpha}$ and $\hat{\beta}$, we define the perturbation terms

$$
\varepsilon_{\hat{\beta}}(\cdot) \triangleq f^{\star}-T_{\hat{\beta}}\left(g^{\star}\right), \quad \iota_{\hat{\alpha}}(\cdot) \triangleq g^{\star}-T_{\hat{\alpha}}\left(f^{\star}\right)
$$

so that the updates can be rewritten as

$$
\begin{aligned}
& \exp \left(-f_{t+1}+f^{\star}\right)=\left(1-\eta_{t}\right) \exp \left(-f_{t}+f^{\star}\right)+\eta_{t} \exp \left(-T_{\hat{\beta}_{t}}\left(g_{t}\right)+T_{\hat{\beta}_{t}}\left(g^{\star}\right)+\varepsilon_{\hat{\beta}_{t}}\right) \\
& \exp \left(-g_{t+1}+g^{\star}\right)=\left(1-\eta_{t}\right) \exp \left(-g_{t}+g^{\star}\right)+\eta_{t} \exp \left(-T_{\hat{\alpha}_{t}}\left(f_{t}\right)+T_{\hat{\alpha}_{t}}\left(f^{\star}\right)+\iota_{\hat{\alpha}_{t}}\right) .
\end{aligned}
$$

We denote $u_{t} \triangleq-f_{t}+f^{\star}, v_{t} \triangleq-g_{t}+g^{\star}, u_{t}^{T} \triangleq T_{\hat{\beta}_{t}}\left(f_{t}\right)-T_{\hat{\beta}_{t}}\left(f^{\star}\right), v_{t}^{T} \triangleq T_{\hat{\beta}_{t}}\left(g_{t}\right)-T_{\hat{\beta}_{t}}\left(g^{\star}\right)$. Reusing the same derivations as in the proof of Proposition 5, we obtain

$$
\begin{aligned}
\left\|u_{t+1}\right\|_{\mathrm{var}} \leqslant & \left(1-\eta_{t}\right)\left\|u_{t}\right\|_{\mathrm{var}} \\
& +\eta_{t} \log \left(\max _{x, y \in \mathcal{X}} \exp \left(\varepsilon_{\hat{\beta}_{t}}(x)-\varepsilon_{\hat{\beta}_{t}}(y)\right) \exp \left(v_{t}^{T}(x)-v_{t}^{T}(y)\right)\right) \\
\leqslant & \left(1-\eta_{t}\right)\left\|u_{t}\right\|_{\mathrm{var}}+\eta_{t}\left\|v_{t}^{T}\right\|_{\mathrm{var}}+\eta_{t}\left\|\varepsilon_{\hat{\beta}_{t}}\right\|_{\mathrm{var}}
\end{aligned}
$$

where we have used $\max _{x} f(x) g(x) \leqslant \max _{x} f(x) \max _{x} f(x)$ on the second line. Therefore, using the contractivity of the soft $C$-transform,

$$
\begin{equation*}
e_{t+1} \leqslant\left(1-\tilde{\eta}_{t}\right) e_{t}+\frac{\tilde{\eta}_{t}}{1-\kappa}\left(\left\|\varepsilon_{\hat{\beta}_{t}}\right\|_{\mathrm{var}}+\left\|\iota_{\hat{\alpha}_{t}}\right\|_{\mathrm{var}}\right) \tag{8}
\end{equation*}
$$

where we set $e_{t} \triangleq\left\|u_{t}\right\|_{\mathrm{var}}+\left\|v_{t}\right\|_{\mathrm{var}}, \tilde{\eta}_{t}=\eta_{t}(1-\kappa)$ and $\kappa$ is set to be the biggest contraction factor over all empirical realizations $\hat{\alpha}_{t}, \hat{\beta}_{t}$ of the distributions $\alpha$ and $\beta$. It is upper bounded by $1-e^{-L \operatorname{diam}(\mathcal{X})}$, thanks to Assumption 1 and Lemma 1.
The realizations $\hat{\beta}_{t}$ and $\hat{\alpha}_{t}$ are sampled according to the same distribution $\hat{\alpha}$ and $\hat{\beta}$. We define the sequence $r_{t}$ to be the running average of the variational norm of the (functional) error term:

$$
r_{t+1} \triangleq\left(1-\tilde{\eta}_{t}\right) r_{t}+\frac{\tilde{\eta}_{t}}{1-\kappa}\left(\left\|\varepsilon_{\hat{\beta}_{t}}\right\|_{\mathrm{var}}+\left\|\iota_{\hat{\alpha}_{t}}\right\|_{\mathrm{var}}\right)
$$

We thus have, for all $t>0, e_{t} \leqslant r_{t}$. Using Lemma 3, the sequence $\left(r_{t}\right)_{t}$ converges towards the scalar expected value

$$
\begin{equation*}
r_{\infty} \triangleq \frac{1}{1-\kappa} \mathbb{E}_{\hat{\alpha}, \hat{\beta}}\left[\left\|\varepsilon_{\hat{\beta}}\right\|_{\mathrm{var}}+\left\|\iota_{\hat{\alpha}}\right\|_{\mathrm{var}}\right]>0 \tag{9}
\end{equation*}
$$

We now relate $r_{\infty}$ to the number of samples $n$ using a uniform law of large number result on parametric functions. We write $\hat{\beta}=\hat{\beta}_{n}$ to make explicit the dependency of the quantities on the batch size $n$.
Using Lemma 2, we bound the quantity

$$
\begin{aligned}
& E_{n} \triangleq \mathbb{E}_{\hat{\beta}_{n}}\left\|\varepsilon_{\hat{\beta}_{n}}\right\|_{\text {var }}=\mathbb{E}_{\hat{\beta}_{n}}\left\|\exp \left(-T_{\beta}\left(g_{0}^{\star}\right)\right)-\exp \left(-T_{\hat{\beta}_{n}}\left(g_{0}^{\star}\right)\right)\right\|_{\infty} \\
&\left.=\mathbb{E}_{Y_{1}, \ldots Y_{n} \sim \beta} \sup _{x \in \mathcal{X}} \left\lvert\, \frac{1}{n} \sum_{i=1}^{n} \exp \left(g^{\star}\left(Y_{i}\right)\right)-C\left(x, Y_{i}\right)\right.\right) \\
& \quad-\mathbb{E}_{Y \sim \beta}\left[\exp \left(g_{0}^{\star}(Y)\right)-C(x, Y)\right] \mid \\
&=\mathbb{E} \sup _{x \in \mathcal{X}}\left|\frac{1}{n} \sum_{i=1}^{n} \varphi_{i}(x)-\varphi(x)\right|
\end{aligned}
$$

where we have defined $\varphi_{i}: x \rightarrow \exp \left(g^{\star}\left(Y_{i}\right)-C\left(x, Y_{i}\right)\right)$ and set $\varphi$ to be the expected value of each $\varphi_{i}$. The compactness of $\mathcal{X}$ ensures that the functions are square integrable and uniformly bounded. Lemma 2 ensures that there exists $S\left(g^{\star}\right)$ such that

$$
E_{n} \leqslant \frac{S\left(g^{\star}\right)}{\sqrt{n}}
$$

We now bound $\mathbb{E}_{\hat{\beta}_{n}}\left\|\varepsilon_{\hat{\beta}_{n}}\right\|_{\text {var }}$ using the quantity $E_{n}$. First, we observe that $\|_{\text {var }}=g_{\min }^{\star}<g^{\star}<0$, and there exists $C_{\max }>0$ such that $0 \leqslant C(x, y) \leqslant C_{\max }$ for all $x, y \in \mathcal{X}$, thanks to the Assumption 1.

$$
\begin{aligned}
& \delta \triangleq \exp \left(-\left\|g^{\star}\right\|_{\text {var }}-C_{\max }\right) \leqslant \exp \left(-T_{\beta}\left(g^{\star}\right)\right) \leqslant 1 \\
& \quad \quad \exp \left(-\left\|g^{\star}\right\|_{\text {var }}-C_{\max }\right) \leqslant \exp \left(-T_{\hat{\beta}_{n}}\left(g^{\star}\right)\right) \leqslant 1
\end{aligned}
$$

where we have used $g^{\star}=\left\|g^{\star}\right\|_{\text {var }}$. For all $x \in \mathcal{X}$,

$$
\begin{equation*}
\left|\varepsilon_{\hat{\beta}_{n}}\right|=\left|\log \frac{\exp \left(-T_{\hat{\beta}_{n}}\left(g^{\star}\right)\right)}{\exp \left(-T_{\beta}\left(g^{\star}\right)\right)}\right|=\left|\log \left(1+\frac{\exp \left(-T_{\hat{\beta}_{n}}\left(g^{\star}\right)\right)-\exp \left(-T_{\beta}\left(g^{\star}\right)\right)}{\exp \left(-T_{\beta}\left(g^{\star}\right)\right)}\right)\right| . \tag{10}
\end{equation*}
$$

We first obtain an upper-bound independent of $n$ with the first equality in (10):

$$
\begin{equation*}
\left\|\varepsilon_{\hat{\beta}_{n}}\right\|_{\mathrm{var}} \leqslant\left\|\varepsilon_{\hat{\beta}_{n}}\right\|_{\infty} \leqslant\left\|g^{\star}\right\|_{\mathrm{var}}+C_{\max } \tag{11}
\end{equation*}
$$

We now use the second expression in (10): for $n$ large enough, $E_{n}<\delta$

$$
\begin{equation*}
\left\|\varepsilon_{\hat{\beta}_{n}}\right\|_{\mathrm{var}} \leqslant \max \left(\log \left(1+\frac{E_{n}}{\delta}\right),-\log \left(1-\frac{E_{n}}{\delta}\right)\right)=-\log \left(1-\tilde{E}_{n}\right) \tag{12}
\end{equation*}
$$

where we have set $\tilde{E}_{n} \triangleq \frac{E_{n}}{\delta}$. On the event $\Omega_{n}=\left\{\tilde{E}_{n} \leqslant \frac{1}{2}\right\}$, a simple calculation gives $-\log (1-$ $\left.\tilde{E}_{n}\right) \leqslant(2 \log 2) \tilde{E}_{n} \leqslant 2 \tilde{E}_{n}$. Thanks to Markov inequality, $\mathbb{P}\left[\tilde{E}_{n}>\frac{1}{2}\right] \leqslant 2 \mathbb{E}\left[\tilde{E}_{n}\right]$. We then split the expectation over the event $\Omega_{n}$, and use inequalities (12) and (11) on each conditional expectation:

$$
\begin{align*}
\mathbb{E}\left\|\varepsilon_{\hat{\beta}_{n}}\right\|_{\mathrm{var}}= & \mathbb{P}\left[\tilde{E}_{n} \leqslant \frac{1}{2}\right] \mathbb{E}\left[\left\|\varepsilon_{\hat{\beta}_{n}}\right\|_{\mathrm{var}} \left\lvert\, \tilde{E}_{n} \leqslant \frac{1}{2}\right.\right]  \tag{13}\\
& +\mathbb{P}\left[\tilde{E}_{n}>\frac{1}{2}\right] \mathbb{E}\left[\left\|\varepsilon_{\hat{\beta}_{n}}\right\|_{\mathrm{var}} \left\lvert\, \tilde{E}_{n}>\frac{1}{2}\right.\right] \\
\leqslant & \frac{2 \varphi\left(\left\|g^{\star}\right\|_{\mathrm{var}}+C_{\max }\right) S\left(g^{\star}\right)}{\sqrt{n}} \\
\leqslant & \frac{4 \exp \left(\left\|g^{\star}\right\|_{\mathrm{var}}+C_{\max }\right) S\left(g^{\star}\right)}{\sqrt{n}} \triangleq \frac{A\left(g^{\star}\right)}{\sqrt{n}}
\end{align*}
$$

The constants $S$ depends on the complexity of estimating the functional $x \rightarrow \int_{y} \exp \left(g^{\star}(y)-\right.$ $C(x, y)) \mathrm{d} \beta(y)$ with samples from $\beta$. A parallel result holds for $\mathbb{E}_{\hat{\alpha}_{n}}\left\|\iota_{\hat{\alpha}_{n}}\right\|_{\text {var }}$. Therefore, there exists $A\left(f^{\star}\right), A\left(g^{\star}\right)>0$ such that $r_{\infty} \leqslant \frac{A\left(f^{\star}\right)+A\left(g^{\star}\right)}{\sqrt{n}}$. As for all $t>0, e_{t} \leqslant r_{t} \rightarrow_{t \rightarrow \infty} r_{\infty}$, the proposition follows, writing $A=A\left(f^{\star}\right)+A\left(g^{\star}\right)$.
The constant $A$ is larger than $\exp \left(C_{\max }\right)$ when $C_{\max } \rightarrow \infty$; Hence it behaves at least like $\exp \left(\frac{1}{\varepsilon}\right)$ when $\varepsilon \rightarrow 0$.

Note that we have used twice a corollary of the law of large numbers: once when averaging over $t$ with $t \rightarrow \infty$ (Eq. (9)), and once when averaging over $n$ with $n$ finite (Eq. (13)).

## A. 4 Proof of Proposition 3

In the proof of Proposition 2 and in particular Eq. (8), the term that prevents the convergence of $e_{t}$ is

$$
\eta_{t}\left(\left\|\varepsilon_{\hat{\beta}_{t}}\right\|_{\mathrm{var}}+\left\|\iota_{\hat{\alpha}_{t}}\right\|_{\mathrm{var}}\right)
$$

which is not summable in general. We can control this term by increasing the size of $\hat{\alpha}_{t}$ and $\hat{\beta}_{t}$ with time, at a sufficient rate: this is what Assumption 3 ensures.

Proof. From Eq. (8), for all $t>0$, we have

$$
\begin{equation*}
0 \leqslant e_{t+1} \leqslant\left(1-\tilde{\eta}_{t}\right) e_{t}+\eta_{t}\left(\left\|\varepsilon_{\hat{\beta}_{t}}\right\|_{\mathrm{var}}+\left\|\iota_{\hat{\alpha}_{t}}\right\|_{\mathrm{var}}\right) \tag{14}
\end{equation*}
$$

Taking the expectation and using the uniform law of large number (13),

$$
\begin{align*}
\mathbb{E} e_{t+1} & \leqslant\left(1-(1-\kappa) \eta_{t}\right) \mathbb{E} e_{t}+\eta_{t} \frac{A}{\sqrt{n(t)}}  \tag{15}\\
& =\left(1-(1-\kappa) \eta_{t}\right) \mathbb{E} e_{t}+A \eta_{t} w_{t}
\end{align*}
$$

where we have used the definition of $n(t)$ from Assumption 3 in the last line.
The proof follows from a simple asymptotic analysis of the sequence $\left(\mathbb{E} e_{t}\right)_{t}$, following recursion (15). For all $t>0$,

$$
\begin{equation*}
\mathbb{E} e_{t+1}-\mathbb{E} e_{t}=-(1-\kappa) \eta_{t} \mathbb{E} e_{t}+A \eta_{t} w_{t} \leqslant A \eta_{t} w_{t} \tag{16}
\end{equation*}
$$

Therefore, from Assumption 3, $\left(\mathbb{E} e_{t+1}-\mathbb{E} e_{t}\right)_{t}$ is summable and $\mathbb{E} e_{t} \rightarrow_{t \rightarrow \infty} \ell \geqslant 0$. Let's assume $\ell>0$. Summing (16) over $t$, we obtain

$$
\mathbb{E} e_{t} \leqslant \mathbb{E} e_{1}-(1-\kappa) \sum_{s=1}^{t-1} \eta_{s} \mathbb{E}_{s}+A \sum_{s=1}^{t-1} \eta_{s} w_{s} \rightarrow_{t \rightarrow \infty}-\infty
$$

which leads to a contradiction. Therefore $\mathbb{E} e_{t} \rightarrow_{t \rightarrow \infty} 0$. As $e_{t} \geqslant 0$ for all $t>0$, this implies that $e_{t} \rightarrow_{t \rightarrow \infty} 0$ almost surely.

## A. 5 Proof of Proposition 4

Proof. The proof of Proposition 3 allows us to derive non-asymptotic rates for potential estimations using the online Sinkhorn algorithm. Let us set $\eta_{t}=\frac{\lambda}{t^{a}}, n(t)=\left\lceil B t^{2 b}\right\rceil$ in (14), so that Assumption 3 is met. $\lceil\cdot\rceil$ denotes the ceiling function. We are left to study the recursion (15):

$$
\delta_{t+1} \triangleq \mathbb{E} e_{t+1} \leqslant\left(1-\frac{\lambda(1-\kappa)}{t^{a}}\right) \delta_{t}+\frac{A \lambda}{\sqrt{B} t^{a+b}}
$$

Following the derivations of Moulines and Bach (2011, Theorem 2), we have the following biasvariance decomposed upper-bound, provided that $0 \leqslant a<1$ and $a+b>1$. For all $t>0$,

$$
\begin{equation*}
\delta_{t} \leqslant\left(\delta_{0}+\frac{A S}{(a+b-1) \sqrt{B}}\right) \exp \left(-\frac{S(1-\kappa)}{2} t^{1-a}\right)+\frac{2 A S}{\sqrt{B}(1-\kappa) t^{a}} \tag{17}
\end{equation*}
$$

Let us now relate the iteration number $t$ to the number of seen sample $N$. By definition

$$
n_{t}=\sum_{s=1}^{t} n(s) \leqslant B \sum_{s=1}^{t} s^{2 b}+t \leqslant t+\frac{(t+1)^{2 b+1}-1}{2 b+1} \leqslant(2 t)^{2 b+1}
$$

Therefore, when we have seen $N$ samples, the iteration number is superior to $t(N)$, and the expected error $\delta_{N}$ is of the order of $\delta_{t(N)}$, with

$$
\begin{equation*}
t(N)=(N / 2)^{\frac{1}{2 b+1}} \tag{18}
\end{equation*}
$$

We write $\delta_{N}=\delta_{t(N)}$. Replacing (18) in (17) yields

$$
\begin{equation*}
\delta_{n} \leqslant\left(\delta_{0}+\frac{A \lambda}{(a+b-1) \sqrt{B}}\right) \exp \left(-\frac{\lambda(1-\kappa)}{2}(n / 2)^{\frac{1-a}{2 b+1}}\right)+\frac{2 A \lambda}{\sqrt{B}(1-\kappa)(n / 2)^{\frac{a}{2 b+1}}} . \tag{19}
\end{equation*}
$$

We note that $b$ and $a$ should be as close to 0 as possible to reduce the bias term, while $a$ should be as close to 1 and $b$ as close to 0 as possible to reduce the variance term. Of course, $b$ should remain larger than $1-a$ to ensure convergence.
To obtain the best asymptotical rates (the error is always dominated by the variance term), we set $a=1-\iota, b=2 \iota$, with $\iota \succsim 0$. This yields

$$
\begin{aligned}
\delta_{n} & \leqslant\left(\delta_{0}+\frac{A \lambda}{\iota \sqrt{B}}\right) \exp \left(-\frac{\lambda(1-\kappa)}{2}(n / 2)^{\frac{\iota}{1+4 \iota}}\right)+\frac{2 A \lambda}{\sqrt{B}(1-\kappa)(n / 2)^{\frac{1-\iota}{1+4 \iota}}} \\
& =\mathcal{O}\left(n^{-\frac{1-\iota}{1+4 \iota}}\right) .
\end{aligned}
$$

This rate is as close to the rate $\mathcal{O}\left(\frac{1}{n}\right)$ as desired. We may then perform a last soft $C$-transform (using the $n_{t}$ seen samples) over the estimated $f_{t(n)}, g_{t(n)}$ to obtain a estimated solution of the dual optimisation problem (2). The Sinkhorn potentials can therefore be estimated with fast rates. Note that the upper bound explodes when $\varepsilon \rightarrow 0$, as $C_{\max } \rightarrow \infty$, hence $A \rightarrow \infty$, and $(1-\kappa) \rightarrow 0$.

Estimating the Sinkhorn distance. The Sinkhorn distance requires to estimate the integral

$$
\mathcal{W}(\alpha, \beta)=\int_{x} f^{\star}(x) \mathrm{d} \alpha(x)+\int_{y} g^{\star}(y) \mathrm{d} \beta(y)
$$

At iteration $t(n)$, with empirical realization $\bar{\alpha}_{t}$ and $\bar{\beta}_{t}$, containing $n$ samples, we use the estimator

$$
\hat{\mathcal{W}}(\alpha, \beta)=\frac{1}{n} \sum_{i=1}^{n} f_{t(n)}\left(x_{i}\right)+\frac{1}{n} \sum_{i=1}^{n} g_{t(n)}\left(y_{i}\right)
$$

We can bound the estimation error $|\hat{\mathcal{W}}(\alpha, \beta)-\mathcal{W}(\alpha, \beta)|=\mathcal{O}\left(\frac{1}{\sqrt{n}}\right)$, dominated by the integral evaluation noise. We thus recover a new estimator of the Sinkhorn distance with the same sample complexity as the batch Sinkhorn estimator (Genevay, Chizat, et al., 2019). Our estimator enjoys an original rate for estimating the potentials in $\|\cdot\|_{\text {var }}$.

```
Algorithm 2 Fully-corrective online Sinkhorn
    Input: Distribution \(\alpha\) and \(\beta\), learning weights \(\left(\eta_{t}\right)_{t}\) and batch-sizes \((n(t))_{t}\). Set \(p_{i, 1}=q_{i, 1}=0\)
    for \(i \in\left(0, n_{1}\right]\)
    for \(t=0, \ldots, T-1\) do
        Sample \(\left(x_{i}\right)_{\left(n_{t}, n_{t+1}\right]} \sim \alpha,\left(y_{j}\right)_{\left(n_{t}, n_{t+1}\right]} \sim \beta\).
        Evaluate \(\left(\hat{f}_{t}\left(x_{i}\right)\right)_{i=\left(0, n_{t+1}\right]},\left(\hat{g}_{t}\left(y_{i}\right)\right)_{i=\left(0, n_{t+1}\right]}\) using \(\left(q_{i, t}, p_{i, t}, x_{i}, y_{i}\right)_{i=\left(0, n_{t}\right]}\) in (7).
        \(q_{\left(0, n_{t+1}\right], t+1} \leftarrow \log \frac{1}{n}+\left(\hat{g}_{t}\left(y_{i}\right)\right)_{\left(0, n_{t+1}\right]}, \quad p_{\left(n_{t}, n_{t+1}\right], t+1} \leftarrow \log \frac{1}{n}+\left(\hat{f}_{t}\left(x_{i}\right)\right)_{\left(n_{t}, n_{t+1}\right]}\).
    Returns: \(\hat{f}_{T}:\left(q_{i, T}, y_{i}\right)_{\left(0, n_{T}\right]}\) and \(\hat{g}_{T}:\left(p_{i, T}, x_{i}\right)_{\left(0, n_{T}\right]}\)
```

```
Algorithm 3 Online Sinkhorn potentials in the discrete setting
    Input: Distribution \(\alpha \in \triangle^{N}\) and \(\beta \in \triangle^{N}, x \in \mathbb{R}^{n \times d}, y \in \mathbb{R}^{n \times d}\), learning weights \(\left(\eta_{t}\right)_{t}\)
    Set \(p=q=-\infty \in \mathbb{R}^{n}\).
    for \(t=1, \ldots, T\) do
        \(q \leftarrow q+\log \left(1-\eta_{t}\right), p \leftarrow p+\log \left(1-\eta_{t}\right)\).
        Sample \(J_{t} \subset[1, N], I_{t} \subset[1, N]\) of size \(n(t)\).
        for \(i \in J_{t}\) do
            \(q_{i} \leftarrow \log \left(\exp \left(q_{i}\right)+\exp \left(\log \left(\eta_{t}\right)-\log \frac{1}{N} \sum_{j=1}^{N} \exp \left(p_{j}-C\left(x_{j}, y_{i}\right)\right)\right)\right.\).
        for \(i \in I_{t}\) do
            \(p_{i} \leftarrow \log \left(\exp \left(q_{i}\right)+\exp \left(\log \left(\eta_{t}\right)-\log \frac{1}{M} \sum_{j=1}^{M} \exp \left(q_{j}-C\left(x_{i}, y_{j}\right)\right)\right)\right.\).
    Returns \(f_{T}:(q, y)\) and \(g_{T}:(p, x)\)
```


## B Online Sinkhorn variants

## B. 1 Fully-corrective scheme

We report the fully-corrective online Sinkhorn algorithm in Algorithm 2. This algorithm also enjoys almost sure convergence, provided that the following assumption is met.
Assumption 4. For all $t>0$, the total batch-size $n_{t}=\frac{B}{w_{t}^{2}}$ is an integer. The step-size $\eta_{t}$ and the batch-size $n_{t}$ grows so that $\sum w_{t} \eta_{t}<\infty$ and $\sum \eta_{t}=\infty$.

With full correction, the total number of observed samples $n_{t}$ needs to grow at the same rate as the single-iteration batch-size $n(t)$ in Assumption 3. For $\eta_{t}=\frac{1}{t^{a}}, a \in(1 / 2,1]$, it is sufficient to use a constant batch-size $n(t)=B$ to meet Assumption 4. We then have the following property
Proposition 6. Under Assumption 1 and 4, the fully-corrective online Sinkhorn algorithm converges almost surely:

$$
\left\|\hat{f_{t}}-f^{\star}\right\|_{\mathrm{var}}+\left\|\hat{g}_{t}-g^{\star}\right\|_{\mathrm{var}} \rightarrow 0
$$

Proof. Using the fully-corrective scheme allows to replace $n(t)$ by $n_{t}=\sum_{s=0}^{t} n(s)$ in (15). The proposition is then obtained in the same way as Proposition 4.

## B. 2 Online Sinkhorn for discrete distributions

The online Sinkhorn algorithm takes a simpler form with discrete distributions. We derive it in Algorithm 3. We set $\alpha$ and $\beta$ to have size $N$ and $M$, respectively. We evaluate the potentials as

$$
\begin{aligned}
& g_{t}(y)=-\log \sum_{j=1}^{N} \exp \left(p_{j}-C\left(x_{j}, y\right)\right) \\
& f_{t}(x)=-\log \sum_{j=1}^{M} \exp \left(q_{j}-C\left(x, y_{j}\right)\right)
\end{aligned}
$$

Table 2: Schedules of batch-sizes and learning rates that ensures online Sinkhorn convergence.

| Param. schedule | Online Sinkhorn | Fully-corrective online Sinkhorn |
| :---: | :---: | :---: |
| Batch size $n(t)=B t^{b}$ | $0<b$ | $0 \leqslant b$ |
| Step size $\eta_{t}=\frac{1}{t^{a}}$ | $a \geqslant 1-\frac{b}{2}$ | $\left\{\begin{array}{lll}a>\frac{1}{2}-\frac{b}{2} & \text { and } \quad b<1 \\ a \geqslant 0 & \text { and } & b \geqslant 1\end{array}\right.$ |

where $\left(p_{j}\right)_{J \in[1, N]}$ and $\left(q_{j}\right)_{J \in[1, M]}$ are fixed-size vectors. Note that the computations written in Algorithm 3 are in log-space, as they should be implemented to prevent numerical overflows. The sets $|I|$ and $|J|$ can have varying sizes along the algorithm, which allows for example to speed-up the initial Sinkhorn iteration (§5.2). In this case, the cost matrix $\left.\hat{C}=C\left(x_{i}, y_{j}\right)\right)_{i, j}$ should be progressively recorded along the algorithm iterations.

## B. 3 Recapitulation on batch-sizes and learning rates

To provide practical guidance on choosing rates in batch-sizes $n(t)$ and step-sizes $\eta_{t}$, we can parametrize $\eta_{t}=\frac{1}{t^{a}}$ and $n(t)=B t^{b}$ and study what is implied by Assumption 3 and Assumption 4. We summarize the schedules for which convergence is guarantees in Table 2. Note that in practice, it is useful to replace $t$ by $(1+r t)$ in these schedules. We set $r=0.1$ in all experiments.

$$
\varepsilon=0.1
$$



$$
\varepsilon=0.0001
$$



Figure 4: Performance of online Sinkhorn for various $\varepsilon$.

## C Extra numerical experiments

We display and describe the supplementary figures mentionned in the main text, as well as experimental details useful for reproduction.

## C. 1 Online Sinkhorn and variants

Grids and details for §5.1. We set $\left(\eta_{t}, n(t)\right)=\left(\frac{1}{(1+0.1 t)^{a}}, 100(1+0.1 t)^{b}\right)$, with $(a, b)=(0,2)$, $(a, b)=\left(\frac{1}{2}, 1\right)$ and $(a, b=1,0)$ (constant batch-sizes). Batch Sinkhorn algorithms uses

$$
\varepsilon=0.1
$$




$\varepsilon=0.01$




$$
\varepsilon=0.001
$$






$$
\varepsilon=0.0001
$$


O-S $n(t)=100$
O-S $n(t) \propto t^{0.5}$
$=$ O-S $n(t) \propto t^{1.0}$
Sinkhorn $n=10^{4}$

- Sinkhorn $n=100$
- Sinkhorn $n=1000$

Figure 5: Performance of fully-corrective online Sinkhorn (O-S) for various $\varepsilon$.
$N=100,1000,10000$. We train Sinkhorn on $t=5000$ iterations, and train online Sinkhorn long enough to match the number of computations of the large Sinkhorn reference.

All OS convergence curves. To complete Fig. 1, Fig. 4 report the performance of online Sinkhorn for $\left.\varepsilon \in\left\{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}\right]\right\}$. The comparison of performance remains similar to the one produced in the main text.

Fully-corrective online Sinkhorn. Fig. 5 reports the performance of fully-corrected online Sinkhorn (FCOS). We observe that the fully-corrective scheme is less noisy than the non-corrected one. It is less efficient than OS on low-dimensional problems, but faster on the 10 dimensional



$$
\begin{aligned}
& \text { R-S } n=100 \\
& \text { R-S } n=1000 \\
& \text { Sinkhorn } n=10^{4} \\
& \text { Sinkhorn } n=100 \\
& \text { Sinkhorn } n=1000
\end{aligned}
$$

$$
\varepsilon=0.01
$$




R-S $n=100$
R-S $n=1000$

- Sinkhorn $n=10^{4}$
Sinkhorn $n=100$
- Sinkhorn $n=1000$
$\varepsilon=0.001$

R-S $n=100$
$\mathrm{R}-\mathrm{S} n=1000$
Sinkhorn $n=10^{4}$
Sinkhorn $n=100$
Sinkhorn $n=1000$
$\varepsilon=0.0001$


Figure 6: Performance of randomized Sinkhorn (R-S) for various $\varepsilon$.
problem. For GMM-10D, it outperforms the batch Sinkhorn algorithm with $N=100,1000$. Note that we interrupt FCOS for $n_{t}>20,000$, as our implementation of the $C$-transform has a quadratic memory cost in $n_{t}$-this cost can be reduced to a linear cost with more careful implementation ${ }^{1}$.

Randomized Sinkhorn. Fig. 6 reports the performance of randomized Sinkhorn. In low dimension, randomized Sinkhorn is a reasonable alternative to batch Sinkhorn, as it often outperforms it on average, for the same memory complexity (compare purple to orange curve for instance). In high dimension, batch Sinkhorn tend to perform slightly better.

[^0]
## C. 2 OT between Gaussians

We measure the performance of online Sinkhorn to transport one Gaussian distribution $\alpha$ to another $\beta$. The potentials $f^{\star}, g^{\star}$ are known exactly for this problem, which allows to have a strong golden standard. More precisely, adapting the formulae from Janati et al., 2020, assuming $\alpha \sim \mathcal{N}(\mu, A)$ and $\beta \sim \mathcal{N}(\nu, \beta)$ and writing $I$ the identity matrix in $\mathbb{R}^{d}$, we have

$$
\begin{aligned}
& C \triangleq\left(A B+\frac{\varepsilon^{2}}{4} I\right)^{1 / 2}, \quad U \triangleq B\left(C+\frac{\varepsilon}{2} I\right)^{-1}-I, \quad V \triangleq A\left(C+\frac{\varepsilon}{2} I\right)^{-1}-I \\
& f^{\star}: x \rightarrow-\frac{1}{2}(x-\mu)^{\top} U(x-\mu)+x^{\top}(\mu-\nu) \\
& g^{\star}: y \rightarrow-\frac{1}{2}(y-\nu)^{\top} V(y-\nu)+y^{\top}(\nu-\mu)
\end{aligned}
$$

We compare batch Sinkhorn $(N=100,1000,10000)$ to (non fully-corrected) online Sinkhorn, with $n(t)=B$, and $n(t)=B(1+0.1 t)^{1 / 2}, B=100$, and $\varepsilon \in\left\{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}\right\}$.

Results. As displayed in Fig. 7, online Sinkhorn outperforms batch Sinkhorn for all tested batch sizes and all $\varepsilon$. It is faster and does not converge towards biased potentials. This suggests that the performance of online Sinkhorn may be underestimated in the previous analyses due to poor potential reference.

## C. 3 Illustration of online Sinkhorn potentials on a 2D GMM

The estimate $\hat{f}_{t}$ is useful to compute the gradient of the Sinkhorn distance $\mathcal{W}(\alpha, \beta)$ with respect to the distribution $\alpha$. This is useful when $\alpha$ is a parametric distribution $\alpha_{\theta}$, as it allows to compute the gradient of the Sinkhorn distance with respect to $\theta$ using backpropagation. For simplicity, let us assume that $\alpha=\frac{1}{n} \sum_{i=1}^{n} \delta_{x_{i}}$. Then, for all $i \in[1, n]$,

$$
\frac{\partial \mathcal{W}(\alpha, \beta)}{\partial x_{i}}=\nabla_{x}\left(x \rightarrow f^{\star}(\alpha, \beta)\right)\left(x_{i}\right)
$$

so that $\nabla_{x} f^{\star}(\alpha, \beta)$ provides a displacement field that can be descended to minimize $\alpha \rightarrow \mathcal{W}(\alpha, \beta)$. Such point of view can be extended to general distributions using the mean-field point of view, see e.g. Chizat, 2019; Santambrogio, 2015. Estimating $\nabla_{x} f^{\star}(\alpha, \beta)$ is therefore crucial to train e.g. generator networks. Both the online Sinkhorn and the batch Sinkhorn algorithm allow to estimate this vector field, through the plug-in estimator $x \rightarrow \nabla_{x} \hat{f}_{t}$, easily computed using the form (7) of $\hat{f}_{t}$.

Experiment. With 2D GMMs, we estimate a reference vector field $\nabla f_{0}^{\star}$ using Sinkhorn on $N=$ 10,000 samples and qualitatively compare the estimations provided by online Sinkhorn and batch Sinkhorn $(N=1,000)$, for the same number of computations.

Results. We represent the estimations $\nabla_{x} \hat{f}_{t}$ in Fig. 8, for $10^{8}$ computations. We compare them to a reference displacement field, estimated wityh $10^{1} 0$ computations. We observe that online Sinkhorn estimates a smoother displacement field than batch Sinkhorn for the same computational budget, that is closer to the reference displacement field. In particular, it is less noisy in low-mass areas. This suggest that online Sinkhorn would be a interesting replacement for batch Sinkhorn in training generative architectures (used by e.g. Genevay, Peyré, et al. (2018)). $\alpha_{\theta}$ is then defined as the push-forward of some simple measure with a neural network $g_{\theta}$. We leave this direction for future work.

## C. 4 Online Sinkhorn as a warmup process

Grids and details for §5.2. We set $\left(\eta_{t}, n(t)\right)=\left(\frac{1}{(1+0.1 t)^{a}}, 100(1+0.1 t)^{b}\right)$, with $(a, b)=(0,2)$, $(a, b)=\left(\frac{1}{2}, 1\right)$ and $(a, b=1,0)$ (constant batch-sizes). The batch Sinkhorn algorithm that is used for reference and after warmup uses $N=10000$. In the reference algorithm, we precompute the distance matrix to save computation. In the warmup algorithm, this distance matrix is filled progressively and then kept in memory to perform $C$-transforms.

$$
\varepsilon=0.1
$$



Figure 7: Performance of online-Sinkhorn to estimate OT between two Gaussians. Online Sinkhorn systematically outperforms batch Sinkhorn, but in term of speed and correction.

We evaluated OS and fully-corrective OS, and found that fully-corrective was less efficient (due to its higher cost in the early iterations). We evaluated sampling with and without replacement in the warmup phase, and found sampling without replacement to be more efficient.

All warmup convergence curves. To complete Fig. 3, we report convergence curves for different $\varepsilon$ in Fig. 9. We find that speed-up increased with $\varepsilon$ and both the 2D and 3D problems, but remains limited for the 10D problem.


Sinkhorn $N=1000$, $10^{8}$ computations



Reg. OT displacement field on empirical samples
$10^{11}$ computations


Reg. OT displacement field on a regular grid

Figure 8: Displacement field as defined by the potentials estimated by online-Sinkhorn and Sinkhorn on a 2D GMM. With the same computational budget, online Sinkhorn finds smoother displacement fields than Sinkhorn. Those are closer to the true reference displacement field (we use Sinkhorn on $N=10000$ to estimate this reference). $\alpha$ and $\beta$ log-likelihood level-lines are displayed in red and blue, while the arrows are proportional to $\nabla_{x} \hat{f}_{t}(x) \mathrm{d} \alpha(x)$.


Figure 9: Performance of online-Sinkhorn as warmup for various $\varepsilon$.

## D Stochastic mirror descent interpretation

The online Sinkhorn can be understood as a stochastic mirror descent algorithm for a non-convex problem. This equivalence is obtained by applying a change of variable in (1), defining

$$
\begin{equation*}
\mu \triangleq \alpha \exp (f) \quad \text { and } \quad \nu \triangleq \beta \exp (g) \tag{20}
\end{equation*}
$$

The dual problem (2) rewrites as a minimisation problem over positive measures on $\mathcal{X}$ and $\mathcal{Y}$ :

$$
\begin{equation*}
-\min _{(\mu, \nu) \in \mathcal{M}^{+}(\mathcal{X})^{2}} \operatorname{KL}(\alpha \mid \mu)+\mathrm{KL}(\beta \mid \nu)+\left\langle\mu \otimes \nu, e^{-C}\right\rangle-1 \tag{21}
\end{equation*}
$$

where the function $\mathrm{KL}: \mathcal{P}(\mathcal{X}) \times \mathcal{M}^{+}(\mathcal{X}) \triangleq\left\langle\alpha, \log \frac{\mathrm{d} \alpha}{\mathrm{d} \mu}\right\rangle$ is the Kullback-Leibler divergence between $\alpha$ and $\mu$. This objective is block convex in $\mu, \nu$, but not jointly convex. As we now detail, this problem can be solved using a stochastic mirror descent (Beck and Teboulle, 2003), applied here over the Banach space of Radon measures on $\mathcal{X}$, equipped with the total variation norm.

Mirror maps and gradient. For this, we define the (convex) distance generating function $\mathcal{M}^{+}(\mathcal{X})^{2} \rightarrow \mathbb{R}:$

$$
\omega(\mu, \nu) \triangleq \mathrm{KL}(\alpha \mid \mu)+\operatorname{KL}(\beta \mid \nu)
$$

The gradient of this function and of its Fenchel conjugate $\omega^{\star}: \mathcal{C}(\mathcal{X})^{2} \rightarrow \mathbb{R}$ yields two mirror maps. For all $(\mu, \nu) \in \mathcal{M}^{+}(\mathcal{X})^{2},(\varrho, \varphi) \in \mathcal{C}(\mathcal{X})^{2}, \varrho<0, \varphi<0$,

$$
\nabla \omega(\mu, \nu)=\left(-\frac{\mathrm{d} \alpha}{\mathrm{~d} \mu},-\frac{\mathrm{d} \beta}{\mathrm{~d} \nu}\right) \quad \nabla \omega^{\star}(\varrho, \varphi)=\left(-\frac{\alpha}{\varrho},-\frac{\beta}{\varphi}\right) .
$$

The gradient $\nabla F(\mu, \nu)$ of the objective $F$ appearing in (21) is a continuous function

$$
\nabla_{\mu} F(\mu, \nu)=-\frac{1}{\frac{\mathrm{~d} \mu}{\mathrm{~d} \alpha}}+\int_{y \in \mathcal{X}} \frac{\mathrm{~d} \nu}{\mathrm{~d} \beta}(y) \exp (-C(\cdot, y)) \mathrm{d} \beta(y)
$$

and similarly for $\nabla_{\nu} F$.
Stochastic mirror descent. To define stochastic mirror descent iterations, we may replace integration over $\beta$ is by an integration over a sampled measure $\hat{\beta}$. This in turn defines an unbiased gradient estimate $\tilde{\nabla} F$ of $\nabla F$, which has bounded second order moments. This absence of bias is crucial to prove convergence of SMD with high probability. Using the mirror maps and the stochastic estimation of the gradient, one has the following equivalence result, whose proofs stems from direct computations.
Proposition 7. The stochastic mirror descent iterations

$$
\left(\mu_{t}, \nu_{t}\right)=\nabla \omega^{\star}\left(\nabla \omega\left(\mu_{t}, \nu_{t}\right)-\eta_{t} \tilde{\nabla} F\left(\mu_{t}, \nu_{t}\right)\right)
$$

are equal to the updates (6) under the change of variable (20).
Interpretation. It is important to realize that $\mu_{t}$ and $\nu_{t}$ do not need to be stored in memory. Instead, their associated potentials $f_{t}$ and $g_{t}$ are parametrized as (7). In particular, $\mu_{t}$ and $\nu_{t}$ remain absolutely continuous with respect to $\alpha$ and $\beta$ respectively, so that the Kullbach-Leibler divergence terms are always finite. Note that the mirror descent we consider operates in an infinite-dimensional space, as in Hsieh et al. (2018).

Finally, we mention that when computing exact gradients (in the absence of noise) and when using constant step-size of $\eta_{t}=1$, the algorithm matches exactly Sinkhorn iterations with simultaneous updates of the dual variables. This provides a novel interpretation on the Sinkhorn algorithm, that differs from the usual Bregman projection (Benamou et al., 2015), and the related understanding of Sinkhorn as a constant step-size mirror descent on the primal objective (Mishchenko, 2019) and on a semi-dual formulation (Léger, 2019).

Note that one can not directly apply the proofs of convergence of mirror descent to our problem, as the lack of convexity of problem (21) prevents their use.


[^0]:    ${ }^{1}$ Using e.g. https://www.kernel-operations.io/keops/index.html

