## 8. Additional Demonstration Plots



Figure 4. Looseness of the objective obtained by naive gradient descent $(\gamma=1 / M)$, projected gradient descent $(\gamma=1 /(2 M)$ ) and proximal gradient descent $(\gamma=1 / M)$. Optimization starts with $\boldsymbol{m}=0$ and $C=\rho I$ where $\rho$ is a scaling factor.


Figure 5. Looseness of the objective obtained by naive gradient descent $(\gamma=1 / M)$, projected gradient descent $(\gamma=1 /(2 M)$ ) and proximal gradient descent $(\gamma=1 / M)$. Optimization starts with $\boldsymbol{m}=0$ and $C=\rho I$ where $\rho$ is a scaling factor.


Figure 6. Looseness of the objective obtained by naive gradient descent $(\gamma=1 / M)$, projected gradient descent $(\gamma=1 /(2 M)$ ) and proximal gradient descent $(\gamma=1 / M)$. Optimization starts with $\boldsymbol{m}=0$ and $C=\rho I$ where $\rho$ is a scaling factor.


Figure 7. Looseness of the objective obtained by naive gradient descent $(\gamma=1 / M)$, projected gradient descent $(\gamma=1 /(2 M)$ ) and proximal gradient descent $(\gamma=1 / M)$. Optimization starts with $\boldsymbol{m}=0$ and $C=\rho I$ where $\rho$ is a scaling factor.

## 9. Proofs for Technical Lemmas

This section gives proofs for the technical lemmas used in the main result. Firstly, we show that $\langle\cdot, \cdot\rangle_{s}$ is a valid inner-product.
Lemma 2. $\langle\boldsymbol{a}, \boldsymbol{b}\rangle_{s}=\mathbb{E}_{\mathrm{u} \sim s} \boldsymbol{a}(\mathrm{u})^{\top} \boldsymbol{b}(\mathrm{u})$ is a valid innerproduct on squared-integrable $\boldsymbol{a}: \mathbb{R}^{d} \rightarrow \mathbb{R}^{k}$.

Proof. The space of square integrable functions is $\left\{\boldsymbol{a}: \mathbb{R}^{d} \rightarrow \mathbb{R}^{k} \mid \mathbb{E}_{\mathbf{u} \sim s} a_{i}(\mathbf{u})^{2} \leq \infty \forall i \in\{1, \ldots, k\}\right\}$.
Since each component $a_{i}(\boldsymbol{u})$ and $b_{i}(\boldsymbol{u})$ is square-integrable with respect to $s(\boldsymbol{u})$ we know (by Cauchy-Schwarz) that $\mathbb{E}_{\mathbf{u} \sim s} a_{i}(\mathbf{u}) b_{i}(\mathbf{u}) \leq \sqrt{\mathbb{E}_{\mathbf{u} \sim s} a_{i}(\mathbf{u})^{2}} \sqrt{\mathbb{E}_{\mathbf{u} \sim s} b_{i}(\mathbf{u})}$ is finite and real. Therefore, we have by linearity of expectation that

$$
\begin{aligned}
\sum_{i=1}^{k} \underset{\mathrm{u} \sim s}{\mathbb{E}} a_{i}(\mathrm{u}) b_{i}(\mathrm{u}) & =\underset{\mathbf{u} \sim s}{\mathbb{E}} \sum_{i=1}^{k} a_{i}(\mathbf{u}) b_{i}(\mathbf{u}) \\
& =\underset{\mathbf{u} \sim s}{\mathbb{E}} \boldsymbol{a}(\mathbf{u})^{\top} \boldsymbol{b}(\mathbf{u}) \\
& =\langle\boldsymbol{a}, \boldsymbol{b}\rangle_{s}
\end{aligned}
$$

is finite and real for all $\boldsymbol{a}, \boldsymbol{b} \in V_{s}$. To show that $\left(V_{s},\langle\cdot, \cdot\rangle_{s}\right)$ is a valid inner-product space, it is easy to establish all the necessary properties of the inner-product, namely for all $\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c} \in V_{s}$,
$\langle\boldsymbol{a}, \boldsymbol{b}\rangle=\langle\boldsymbol{b}, \boldsymbol{a}\rangle$
$\langle\theta \boldsymbol{a}, \boldsymbol{b}\rangle=\theta\langle\boldsymbol{a}, \boldsymbol{b}\rangle$ for $\theta \in \mathbb{R}$
$\langle\boldsymbol{a}+\boldsymbol{b}, \boldsymbol{c}\rangle=\langle\boldsymbol{a}, \boldsymbol{c}\rangle+\langle\boldsymbol{b}, \boldsymbol{c}\rangle$
$\langle\boldsymbol{a}, \boldsymbol{a}\rangle \geq 0$
$\langle\boldsymbol{a}, \boldsymbol{a}\rangle=0 \Leftrightarrow \boldsymbol{a}=\mathbf{0}$. (Where $\mathbf{0}(\boldsymbol{\varepsilon})$ is a function that always returns a vector of $k$ zeros.)

Next, we give three technical Lemmas, which do most of the work of the proof.
Lemma 3. Let $\boldsymbol{a}_{i}(\boldsymbol{u})=\frac{d}{d w_{i}} \boldsymbol{t}_{\boldsymbol{w}}(\boldsymbol{u})$. This is independent of $\boldsymbol{w}$ and $\frac{d l(\boldsymbol{w})}{d w_{i}}=\left\langle\boldsymbol{a}_{i}, \nabla f \circ \boldsymbol{t}_{\boldsymbol{w}}\right\rangle_{s}$.

Proof. Now, we can write $l(\boldsymbol{w})$ as

$$
l(\boldsymbol{w})=\underset{\mathbf{z} \sim q_{\boldsymbol{w}}}{\mathbb{E}} f(\mathbf{z})=\underset{\mathrm{u} \sim s}{\mathbb{E}} f\left(\boldsymbol{t}_{\boldsymbol{w}}(\mathrm{u})\right)
$$

Since $\boldsymbol{t}_{\boldsymbol{w}}(\boldsymbol{u})=C \boldsymbol{u}+\boldsymbol{m}$ is an affine function, it's easy to see that both $\frac{d}{d C_{i j}} \boldsymbol{t}_{\boldsymbol{w}}(\boldsymbol{u})$ and $\frac{d}{d \boldsymbol{m}_{i}} \boldsymbol{t}_{\boldsymbol{w}}(\boldsymbol{u})$ are independent of $\boldsymbol{w}$. Therefore, the gradient of $l(\boldsymbol{w})$ can be written as

$$
\begin{aligned}
\nabla_{w_{i}} l(\boldsymbol{w}) & =\nabla_{w_{i}} \underset{\mathrm{u} \sim s}{\mathbb{E}} f\left(\boldsymbol{t}_{\boldsymbol{w}}(\mathrm{u})\right) \\
& =\underset{\mathrm{u} \sim s}{\mathbb{E}} \nabla_{w_{i}} \boldsymbol{t}_{\boldsymbol{w}}(\mathrm{u})^{\top} \nabla f\left(\boldsymbol{t}_{\boldsymbol{w}}(\mathrm{u})\right) \\
& =\left\langle\boldsymbol{a}_{i}, \nabla f \circ \boldsymbol{t}_{\boldsymbol{w}}\right\rangle_{s}
\end{aligned}
$$

Lemma 4. If $s$ is standardized, then the functions $\left\{\boldsymbol{a}_{i}\right\}$ are orthonormal in $\langle\cdot, \cdot\rangle_{s}$.

Proof. It is easy to calculate that

$$
\begin{aligned}
\frac{d}{d m_{i}} \boldsymbol{t}_{\boldsymbol{w}}(\boldsymbol{u}) & =\boldsymbol{e}_{i} \\
\frac{d}{d C_{i j}} \boldsymbol{t}_{\boldsymbol{w}}(\boldsymbol{u}) & =\boldsymbol{e}_{i} u_{j}
\end{aligned}
$$

where $\boldsymbol{e}_{i}$ is the indicator vector in the $i$-th component. Therefore, we have that

$$
\begin{aligned}
& \underset{\mathrm{u} \sim s}{\mathbb{E}}\left(\frac{d}{d m_{i}} \boldsymbol{t}_{\boldsymbol{w}}(\mathrm{u})\right)^{\top}\left(\frac{d}{d m_{j}} \boldsymbol{t}_{\boldsymbol{w}}(\mathrm{u})\right) \\
& =\underset{\mathrm{u} \sim s}{\mathbb{E}} \boldsymbol{e}_{i}^{\top} \boldsymbol{e}_{j} \\
& =I[i=j] \\
& \underset{\mathbf{u} \sim s}{\mathbb{E}}\left(\frac{d}{d C_{i j}} \boldsymbol{t}_{\boldsymbol{w}}(\mathbf{u})\right)^{\top}\left(\frac{d}{d m_{k}} \boldsymbol{t}_{\boldsymbol{w}}(\mathbf{u})\right) \\
& =\underset{\mathrm{u} \sim s}{\mathbb{E}} \mathbf{u}_{j} \boldsymbol{e}_{i}^{\top} \boldsymbol{e}_{k} \\
& =I[i=k] \underset{\mathrm{u} \sim s}{\mathbb{E}} \mathrm{u}_{j} \\
& =0
\end{aligned}
$$

(since zero mean)

$$
\begin{aligned}
& \underset{\mathbf{u} \sim s}{\mathbb{E}}\left(\frac{d}{d C_{i j}} \boldsymbol{t}_{\boldsymbol{w}}(\mathbf{u})\right)^{\top}\left(\frac{d}{d C_{k l}} \boldsymbol{t}_{\boldsymbol{w}}(\mathbf{u})\right) \\
& \quad=\mathbb{E} \mathbf{u}_{j} \mathbf{u}_{l} \boldsymbol{e}_{i}^{\top} \boldsymbol{e}_{k} \\
& \quad=I[i=k] \underset{\substack{ \\
\mathbb{U} \sim s}}{\mathbb{E}} \mathbf{u}_{j} \mathbf{u}_{l} \\
& \quad=I[i=k] I[j=l]
\end{aligned}
$$

(since unit variance and zero mean)

These three identities are equivalent to stating that $\left\{\boldsymbol{a}_{i}\right\}$ are orthonormal in $\langle\cdot, \cdot\rangle_{s}$.

Lemma 5. If $s$ is standardized, then $\mathbb{E}_{\mathrm{u} \sim s}\left\|\boldsymbol{t}_{\boldsymbol{w}}(\mathrm{u})-\boldsymbol{t}_{\boldsymbol{v}}(\mathrm{u})\right\|_{2}^{2}=\|\boldsymbol{w}-\boldsymbol{v}\|_{2}^{2}$.

Proof. Let $\Delta \boldsymbol{m}$ and $\Delta S$ denote the difference of the $\boldsymbol{m}$ and $S$ parts of $\boldsymbol{w}$, respectively. We want to calculate

$$
\begin{aligned}
& \underset{\mathrm{u} \sim s}{\mathbb{E}}\left\|\boldsymbol{t}_{\boldsymbol{w}}(\mathrm{u})-\boldsymbol{t}_{\boldsymbol{v}}(\mathrm{u})\right\|_{2}^{2} \\
& =\underset{\mathrm{u} \sim s}{\mathbb{E}}\|\Delta C \boldsymbol{\varepsilon}+\Delta \boldsymbol{m}\|_{2}^{2} \\
& =\underset{\mathrm{u} \sim s}{\mathbb{E}}\left(\|(\Delta C) \mathrm{u}\|_{2}^{2}+2 \Delta \boldsymbol{m}^{\top} \Delta C \mathrm{u}+\|\Delta \boldsymbol{m}\|_{2}^{2}\right) .
\end{aligned}
$$

It is easy to see that the expectation of the middle term is zero, and the last is a constant. The expectation of the first
term is

$$
\begin{aligned}
\underset{\mathrm{u} \sim s}{\mathbb{E}}\|(\Delta C) \mathrm{u}\|_{2}^{2}= & \underset{\mathrm{u} \sim s}{\mathbb{E}} \mathrm{u}^{\top}(\Delta C)^{\top}(\Delta C) \mathrm{u} \\
= & \underset{\mathrm{u} \sim s}{\mathbb{E}} \operatorname{tr}\left(\mathrm{u}^{\top}(\Delta C)^{\top}(\Delta C) \mathrm{u}\right) \\
= & \underset{\mathrm{u} \sim s}{\mathbb{E}} \operatorname{tr}\left((\Delta C)^{\top}(\Delta C) \mathrm{uu}^{\top}\right) \\
= & \operatorname{tr}\left((\Delta C)^{\top}(\Delta C)\right)=\|\nabla C\|_{F}^{2} . \\
& \text { (since zero mean and unit variance) }
\end{aligned}
$$

Putting this together gives that

$$
\begin{aligned}
\underset{\mathrm{u} \sim s}{\mathbb{E}}\left\|\boldsymbol{t}_{\boldsymbol{w}}(\mathrm{u})-\boldsymbol{t}_{\boldsymbol{v}}(\mathrm{u})\right\|_{2}^{2} & =\|\Delta C\|_{F}^{2}+\|\Delta \boldsymbol{m}\|_{2}^{2} \\
& =\|\boldsymbol{w}-\boldsymbol{v}\|_{2}^{2}
\end{aligned}
$$

## 10. Proof for Example Function

Theorem 6. Let $q_{\boldsymbol{w}}=\operatorname{LocScale}(\boldsymbol{m}, C, s)$ with parameters $\boldsymbol{w}=(\boldsymbol{m}, C)$ and a standardized base distribution $s$ and let $f(\boldsymbol{z})=\frac{a}{2}\left\|\boldsymbol{z}-\boldsymbol{z}^{*}\right\|_{2}^{2}$. Then $l(\boldsymbol{w})=\mathbb{E}_{\mathbf{z} \sim q_{\boldsymbol{w}}} f(\mathbf{z})=$ $\frac{a}{2}\left(\left\|\boldsymbol{m}-\boldsymbol{z}^{*}\right\|_{2}^{2}+\|C\|_{F}^{2}\right)$.

Proof. For a general distribution, we have that

$$
\begin{aligned}
\mathbb{E} f(\mathbf{z})= & \frac{a}{2} \mathbb{E}\left\|\mathbf{z}-\mathbb{E}[\mathrm{z}]+\mathbb{E}[\mathbf{z}]-\boldsymbol{z}^{*}\right\|_{2}^{2} \\
= & \frac{a}{2} \mathbb{E}\left(\|\mathbf{z}-\mathbb{E}[\mathrm{z}]\|_{2}^{2}\right. \\
& \left.+2(\mathbf{z}-\mathbb{E}[\mathbf{z}])^{\top}\left(\mathbb{E}[\mathbf{z}]-\boldsymbol{z}^{*}\right)+\left\|\mathbb{E}[\mathbf{z}]-\boldsymbol{z}^{*}\right\|_{2}^{2}\right) \\
= & \frac{a}{2}\left(\operatorname{tr} \mathbb{V}[\mathbf{z}]+\left\|\mathbb{E}[\mathbf{z}]-\mathbf{z}^{*}\right\|_{2}^{2}\right)
\end{aligned}
$$

Now, if $q_{\boldsymbol{w}}$ is a location-scale family, we have that $\mathrm{z}=$ $C \mathrm{u}+\boldsymbol{m}$. Thus,

$$
\begin{aligned}
\operatorname{tr} \mathbb{V}[\mathbf{z}] & =\operatorname{tr} \mathbb{V}[C \mathbf{u}+\boldsymbol{m}] \\
& =\operatorname{tr} \mathbb{V}[C \mathbf{u}] \\
& =\operatorname{tr} C \mathbb{V}[\mathbf{u}] C^{\top} \\
& =\operatorname{tr} C C^{\top} \mathbb{V}[\mathbf{u}] .
\end{aligned}
$$

Meanwhile, we have that

$$
\begin{aligned}
\left\|\mathbb{E}[\mathrm{z}]-\boldsymbol{z}^{*}\right\|_{2}^{2} & =\left\|\mathbb{E}[C \mathbf{u}+\boldsymbol{m}]-\boldsymbol{z}^{*}\right\|_{2}^{2} \\
& =\left\|C \mathbb{E}[\mathbf{u}]+\boldsymbol{m}-\boldsymbol{z}^{*}\right\|_{2}^{2}
\end{aligned}
$$

Thus,

$$
\mathbb{E} f(\mathbf{z})=\frac{a}{2}\left(\operatorname{tr} C \mathbb{V}[\mathbf{u}] C^{\top}+\left\|C \mathbb{E}[\mathbf{u}]+\boldsymbol{m}-\boldsymbol{z}^{*}\right\|_{2}^{2}\right)
$$

The case where $s$ is standardized follows from substituting $\mathbb{E}[\mathbf{u}]=0$ and $\mathbb{V}[\mathbf{u}]=I$ and applying the fact that $\operatorname{tr} C C^{\top}=$ $\|C\|_{F}^{2}$.

## 11. Proofs for Solution Guarantees

Lemma 8. Let $q_{\boldsymbol{w}}=\operatorname{LocScale}(\boldsymbol{m}, C, s)$ with parameters $\boldsymbol{w}=(\boldsymbol{m}, C)$ and a standardized and spherically symmetric base distribution s. Let $l(\boldsymbol{w})=\mathbb{E}_{\mathbf{z} \sim q_{w}} f(\mathbf{z})$. Suppose $C$ is diagonal and $f$ is $M$-smooth. Then, $\left|\frac{d l(\boldsymbol{w})}{d C_{i i}}\right| \leq M\left|C_{i i}\right|$.

Proof. Define $\boldsymbol{w}^{\prime}$ to be $\boldsymbol{w}$ but with $C_{i i}$ set to zero. We will first show that $\frac{d l\left(\boldsymbol{w}^{\prime}\right)}{d C_{i i}}=0$. Using the definition of $\boldsymbol{t}_{w}$ and the fact that $\frac{d}{d C_{i j}} \boldsymbol{t}_{w}(u)=\boldsymbol{e}_{i} u_{j}$ gives that

$$
\begin{align*}
\frac{d}{d C_{i i}} l\left(\boldsymbol{w}^{\prime}\right) & =\underset{\mathrm{u} \sim s}{\mathbb{E}} \frac{d}{d C_{i i}} f\left(\boldsymbol{t}_{\boldsymbol{w}^{\prime}}(\mathbf{u})\right)  \tag{11}\\
& =\underset{\mathrm{u} \sim s}{\mathbb{E}} \mathbf{u}_{i} \boldsymbol{e}_{i}^{\top} \nabla f\left(\boldsymbol{t}_{\boldsymbol{w}^{\prime}}(\mathbf{u})\right)  \tag{12}\\
& =0 . \tag{13}
\end{align*}
$$

The final equality above follows from the facts that $\mathbb{E} \mathbf{u}_{i}=0$ and $\mathbf{u}_{i} \perp \boldsymbol{e}_{i}^{\top} \nabla f\left(\boldsymbol{t}_{\boldsymbol{w}^{\prime}}(\mathrm{u})\right)\left(\right.$ Since $\boldsymbol{t}_{\boldsymbol{w}^{\prime}}(\boldsymbol{u})$ ignores $\left.\boldsymbol{u}_{i}\right)$ so the expectation in Eq. (11) is over two independent random variables, one with mean zero. Now, by Thm. $1, l$ is also $M$-smooth, thus

$$
\begin{aligned}
\left|\frac{d l(\boldsymbol{w})}{d C_{i i}}\right| & =\left|\frac{d l\left(\boldsymbol{w}^{\prime}\right)}{d C_{i i}}-\frac{d l(\boldsymbol{w})}{d C_{i i}}\right| \\
& \leq\left\|\nabla l\left(\boldsymbol{w}^{\prime}\right)-\nabla l(\boldsymbol{w})\right\|_{2} \\
& \leq M\left\|\boldsymbol{w}^{\prime}-\boldsymbol{w}\right\|_{2} \\
& =M\left|C_{i i}\right|
\end{aligned}
$$

Theorem 7. Let $q_{\boldsymbol{w}}=\operatorname{LocScale}(\boldsymbol{m}, C, s)$ with parameters $\boldsymbol{w}=(\boldsymbol{m}, C)$ and a standardized and spherically symmetric base distribution s. Suppose $\boldsymbol{w}$ minimizes $l(\boldsymbol{w})+h(\boldsymbol{w})$ from Eq. (1) and $\log p(\boldsymbol{z}, \boldsymbol{x})$ is $M$-smooth over $\boldsymbol{z}$. Then, $\boldsymbol{w} \in \mathcal{W}_{M}$.

Proof. First, suppose that $C$ is diagonal. Since $\boldsymbol{w}$ minimizes $l+h, \nabla l(\boldsymbol{w})=-\nabla h(\boldsymbol{w})$. The gradient of $h$ with respect to $C$ is $-C^{-\top}$. Thus, $\left|\frac{d l(\boldsymbol{w})}{d C_{i i}}\right|=\left|\frac{d h(\boldsymbol{w})}{d C_{i i}}\right|=\frac{1}{\mid C_{i i}}$. But by Lem. 8, $\left|\frac{d l(\boldsymbol{w})}{d C_{i i}}\right| \leq M\left|C_{i i}\right|$. This establishes the claim for diagonal $C$.

Now, consider some non-diagonal $C$. Let the singular value decomposition be $C=U S V^{\top}$. Define $f_{U}(\boldsymbol{z})=f(U \boldsymbol{z})$ and define $l_{U}$ with respect to $f_{U}$. Let $\boldsymbol{w}^{\prime}=\left(S, U^{\top} \boldsymbol{m}\right)$. Then, the following statements are equivalent to $\boldsymbol{w} \in$ $\operatorname{argmin}_{\boldsymbol{w}} l(\boldsymbol{w})+h(\boldsymbol{w}):$

$$
\begin{aligned}
& (C, \boldsymbol{m}) \in \underset{(C, \boldsymbol{m})}{\operatorname{argmin}} \underset{\mathrm{u} \sim s}{\mathbb{E}} f(C \mathbf{u}+\boldsymbol{m})-\log |C| \\
& \Leftrightarrow(S, \boldsymbol{m}) \in \underset{(S, \boldsymbol{m})}{\operatorname{argmin}} \underset{\mathrm{u} \sim s}{\mathbb{E}} f\left(U S V^{\top} \mathbf{u}+\boldsymbol{m}\right)-\log \left|U S V^{\top}\right| \\
& \Leftrightarrow(S, \boldsymbol{m}) \in \underset{(S, \boldsymbol{m})}{\operatorname{argmin}} \underset{\mathbf{u} \sim s}{\mathbb{E}} f(U S \mathbf{u}+\boldsymbol{m})-\log |S| \\
& \Leftrightarrow(S, \boldsymbol{m}) \in \underset{(S, \boldsymbol{m})}{\operatorname{argmin}} \underset{\mathbf{u} \sim s}{\mathbb{E}} f_{U}\left(S \mathbf{u}+U^{\top} \boldsymbol{m}\right)-\log |S| \\
& \Leftrightarrow \boldsymbol{w}^{\prime} \in \underset{\boldsymbol{w}}{\operatorname{argmin}} l_{U}(\boldsymbol{w})+h(\boldsymbol{w}) .
\end{aligned}
$$

Thus, $\boldsymbol{w}$ minimizing $l+h$ is equivalent to $\boldsymbol{w}^{\prime}$ minimizing $l_{U}+h$. Since $f_{U}$ is $M$-smooth and $S$ is diagonal, we know that $S_{i i} \geq \frac{1}{\sqrt{M}}$ for all.

## 12. Proofs with Convexity

Theorem 10. Let $q_{\boldsymbol{w}}=\operatorname{LocScale}(\boldsymbol{m}, C, s)$ with parameters $\boldsymbol{w}=(\boldsymbol{m}, C)$ and a standardized and spherically symmetric base distribution s. Suppose $\boldsymbol{w}$ minimizes $l(\boldsymbol{w})+h(\boldsymbol{w})$ from Eq. (1) and $-\log p(\boldsymbol{z}, \boldsymbol{x})$ is $c$-strongly convex over $\boldsymbol{z}$. Then, $\|C\|_{F}^{2}+\left\|\boldsymbol{m}-\boldsymbol{z}^{*}\right\|_{2}^{2} \leq \frac{d}{c}$, where $\boldsymbol{z}^{*}=\operatorname{argmax}_{\boldsymbol{z}} \log (\boldsymbol{z}, \boldsymbol{x})$.

It's easy to see that $l$ is minimized by $\overline{\boldsymbol{w}}=\left(\boldsymbol{z}^{*}, \mathbf{0}_{d \times d}\right)$. By Thm. $9, l(\boldsymbol{w})$ is $c$-strongly convex. Thus applying a standard inner-product result on strong convexity (Nesterov, 2014, Thm. 2.1.9),
$c\|\boldsymbol{w}-\overline{\boldsymbol{w}}\|_{2}^{2} \leq\langle\nabla l(\boldsymbol{w})-\nabla l(\overline{\boldsymbol{w}}), \boldsymbol{w}-\overline{\boldsymbol{w}}\rangle$
(since $l$ is strongly convex)

$$
\begin{aligned}
= & \langle\nabla l(\boldsymbol{w}), \boldsymbol{w}-\overline{\boldsymbol{w}}\rangle \\
& (\text { since } \nabla l(\overline{\boldsymbol{w}})=0) \\
= & -\langle\nabla h(\boldsymbol{w}), \boldsymbol{w}-\overline{\boldsymbol{w}}\rangle \\
& (\text { since } \nabla l(\boldsymbol{w})+\nabla h(\boldsymbol{w})=0) \\
= & \operatorname{tr}\left(C^{-\top} C\right) \\
& \left(\text { since } \nabla_{C} h(\boldsymbol{w})=-C^{-\top}, \nabla_{\boldsymbol{m}} h(\boldsymbol{w})=0\right) \\
= & \operatorname{tr} I=d
\end{aligned}
$$

The result follows from observing that $\|\boldsymbol{w}-\overline{\boldsymbol{w}}\|_{2}^{2}=$ $\|C\|_{F}^{2}+\left\|\boldsymbol{m}-\boldsymbol{z}^{*}\right\|_{2}^{2}$.

## 13. Convergence Considerations

Lemma 12. Let $q_{\boldsymbol{w}}=\operatorname{LocScale}(\boldsymbol{m}, C, s)$ with parameters $\boldsymbol{w}=(\boldsymbol{m}, C)$. Then, $h(\boldsymbol{w})=\mathbb{E}_{\mathbf{z} \sim q_{\boldsymbol{w}}}\left[\log q_{\boldsymbol{w}}(\mathrm{z})\right]$ is $M-$ smooth over $\mathcal{W}_{M}$.

Proof. Take $\boldsymbol{w}=(C, \boldsymbol{m}) \in \mathcal{W}_{M}$ and $\boldsymbol{v}=(B, \boldsymbol{n}) \in \mathcal{W}_{M}$. We write $h(C)$ since $h(\boldsymbol{w})$ is independent of $\boldsymbol{m}$. The gradient is $\nabla h(C)=C^{-T}$. Now, use that $\|A X\|_{F} \leq$ $\|A\|_{2}\|X\|_{F}$ to get that

$$
\begin{aligned}
\|\nabla h(B)-\nabla h(C)\|_{F} & =\left\|B^{-1}-C^{-1}\right\|_{F} \\
& =\left\|B^{-1}(B-C) C^{-1}\right\|_{F} \\
& \leq\left\|B^{-1}\right\|_{2}\left\|C^{-1}\right\|_{2}\|B-C\|_{F}
\end{aligned}
$$

But, since $\boldsymbol{w} \in \mathcal{W}_{M},\left\|C^{-1}\right\|_{2}=\frac{1}{\sigma_{\min }(C)} \leq \sqrt{M}$ and similarly for $C$. This establishes that $\|\nabla h(B)-\nabla h(C)\|_{F} \leq$ $M\|B-C\|_{F}$, equivalent to the result.

Theorem 13. Suppose $h(\boldsymbol{w})$ corresponds to a locationscale family with a standardized s, and $\boldsymbol{w}=(\boldsymbol{m}, C)$.

- If $C$ has singular value decomposition $C=U S V^{\top}$, then $\operatorname{proj}_{\mathcal{W}_{M}}(\boldsymbol{w})=\left(\boldsymbol{m}, U T V^{\top}\right)$, where $T$ is a diagonal matrix with $T_{i i}=\max \left(S_{i i}, \frac{1}{\sqrt{M}}\right)$.
- If $C$ is triangular with a positive diagonal, then $\operatorname{prox}_{\gamma}(\boldsymbol{w})=(\boldsymbol{m}, C+\Delta C)$, where $\Delta C$ is a diagonal matrix with $\Delta C_{i i}=\frac{1}{2}\left(\sqrt{C_{i i}^{2}+4 \gamma}-C_{i i}\right)$.

Proof. (Proximal Operator) We know that $h(\boldsymbol{w})=$ Const. $-\log |C|$. Write $\boldsymbol{w}=(\boldsymbol{m}, C)$ and $\boldsymbol{v}=(\boldsymbol{n}, B)$. Then, we can write the proximal operator as

$$
\underset{\lambda}{\operatorname{prox}}(\boldsymbol{w})=\underset{\boldsymbol{v}}{\operatorname{argmin}}-\log |B|+\frac{1}{2 \lambda}\|\boldsymbol{v}-\boldsymbol{w}\|_{2}^{2}
$$

Now, assuming that $C$ is triangular, the solution will leave all entries of $\boldsymbol{w}$ other than the diagonal entries of $C$ unchanged. Then, we will have that $\log |B|=\sum_{i=1}^{d} \log B_{i i}$. Since

$$
\underset{x>0}{\operatorname{argmin}}-\log x+\frac{1}{2 \lambda}(x-y)^{2}=\frac{y+\sqrt{y^{2}+4 \lambda}}{2}
$$

The solution is to set

$$
\begin{aligned}
B_{i i} & =\frac{1}{2}\left(C_{i i}+\sqrt{C_{i i}^{2}+4 \lambda}\right) \\
& =C_{i i}+\frac{1}{2}\left(\sqrt{C_{i i}^{2}+4 \lambda}-C_{i i}\right) .
\end{aligned}
$$

(Projection Operator) Von-Neumann's trace inequality states that $\left|\operatorname{tr} A^{\top} B\right| \leq \sum_{i} \sigma_{i}(A) \sigma_{i}(B)$. Consider any candidate solution $B$ with SVD $Q T P^{\top}$. Then, we can write that

$$
\begin{aligned}
\|B-C\|_{F}^{2} & =\operatorname{tr}(B-C)^{\top}(B-C) \\
& =\|B\|_{F}^{2}-2 \operatorname{tr}\left(B^{\top} C\right)+\|C\|_{F}^{2} \\
& \geq\|T\|_{F}^{2}-2 \sum_{i} T_{i i} S_{i i}+\sum_{i} S_{i i}^{2} \\
& =\sum_{i}\left(T_{i i}-S_{i i}\right)^{2} .
\end{aligned}
$$

We can minimize this lower bound by choosing $T_{i i}=\max \left(1 / \sqrt{M}, S_{i i}\right)$, with a corresponding value of $\sum_{i} \max \left(0,1 / \sqrt{M}-S_{i i}\right)^{2}$. Thus any valid solution will have $\|B-C\|_{F}^{2}$ at least this large.
However, suppose we choose $B=U T_{i i} V^{\top}$ with $T_{i i}$ as above. Then,

$$
\|B-C\|_{F}^{2}=\left\|U T V^{\top}-U S V^{\top}\right\|_{F}^{2}=\sum_{i}\left(T_{i i}-S_{i i}\right)^{2}
$$

so this value $B$ is optimal.

