

Supplementary Material

Confidence ellipsoid

Lemma 3. Let $\mathbf{V}_t = \mathbf{X}_{t-1}\mathbf{X}_{t-1}^\top + \Lambda$ and $\|\alpha^*\|_\Lambda \leq C$. For any \mathbf{x} and $t \geq 1$, with probability at least $1 - \delta$:

$$|\mathbf{x}^\top(\hat{\alpha}_t - \alpha^*)| \leq \|\mathbf{x}\|_{\mathbf{V}_t^{-1}} \left(R \sqrt{2 \log \left(\frac{|\mathbf{V}_t|^{1/2}}{\delta |\Lambda|^{1/2}} \right)} + C \right)$$

Proof of Lemma 3. We have:

$$\begin{aligned} |\mathbf{x}^\top(\hat{\alpha}_t - \alpha^*)| &= |\mathbf{x}^\top(-\mathbf{V}_t^{-1}\Lambda\alpha^* + \mathbf{V}_t^{-1}\xi_t)| \\ &\leq |\mathbf{x}^\top\mathbf{V}_t^{-1}\Lambda\alpha^*| + |\mathbf{x}^\top\mathbf{V}_t^{-1}\xi_t| \\ &\leq \langle \mathbf{x}^\top, \Lambda\alpha^* \rangle_{\mathbf{V}_t^{-1}} + \langle \mathbf{x}, \xi_t \rangle_{\mathbf{V}_t^{-1}} \\ &\leq \|\mathbf{x}\|_{\mathbf{V}_t^{-1}} \left(\|\xi_t\|_{\mathbf{V}_t^{-1}} + \|\Lambda\alpha^*\|_{\mathbf{V}_t^{-1}} \right), \end{aligned}$$

where we used Cauchy-Schwarz inequality in the last step. Now we bound $\|\xi_t\|_{\mathbf{V}_t^{-1}}$ by Lemma 1 and notice that:

$$\begin{aligned} \|\Lambda\alpha^*\|_{\mathbf{V}_t^{-1}} &= \sqrt{(\alpha^*)^\top \Lambda \mathbf{V}_t^{-1} \Lambda \alpha^*} \\ &\leq \sqrt{(\alpha^*)^\top \Lambda \alpha^*} = \|\alpha^*\|_\Lambda \leq C \end{aligned}$$

□

Effective dimension

Lemma 11. For any real positive-definite matrix A with only simple eigenvalue multiplicities and any vector \mathbf{x} such that $\|\mathbf{x}\|_2 \leq 1$ we have that the determinant $|\mathbf{A} + \mathbf{x}\mathbf{x}^\top|$ is maximized by a vector \mathbf{x} which is aligned with an eigenvector of \mathbf{A} .

Proof of Lemma 11. Using Sylvester's determinant theorem, we have:

$$|\mathbf{A} + \mathbf{x}\mathbf{x}^\top| = |\mathbf{A}| |\mathbf{I} + \mathbf{A}^{-1}\mathbf{x}\mathbf{x}^\top| = |\mathbf{A}| (1 + \mathbf{x}^\top \mathbf{A}^{-1} \mathbf{x})$$

From the spectral theorem, there exists an orthonormal matrix \mathbf{U} , the columns of which are the eigenvectors of \mathbf{A} ; such that $\mathbf{A} = \mathbf{U}\mathbf{D}\mathbf{U}^\top$ with \mathbf{D} being a diagonal matrix with the positive eigenvalues of \mathbf{A} on the diagonal. Thus:

$$\begin{aligned} \max_{\|\mathbf{x}\|_2 \leq 1} \mathbf{x}^\top \mathbf{A}^{-1} \mathbf{x} &= \max_{\|\mathbf{x}\|_2 \leq 1} \mathbf{x}^\top \mathbf{U} \mathbf{D}^{-1} \mathbf{U}^\top \mathbf{x} \\ &= \max_{\|\mathbf{y}\|_2 \leq 1} \mathbf{y}^\top \mathbf{D}^{-1} \mathbf{y}, \end{aligned}$$

since \mathbf{U} is a bijection from $\{\mathbf{x}, \|\mathbf{x}\|_2 \leq 1\}$ to itself.

Since there are no multiplicities, it is easy to see that the quadratic mapping $\mathbf{y} \mapsto \mathbf{y}^\top \mathbf{D}^{-1} \mathbf{y}$ is maximized (under the constraint $\|\mathbf{y}\|_2 \leq 1$) by a canonical vector \mathbf{e}_I corresponding to the lowest diagonal entry I of \mathbf{D} . Thus the maximum of $\mathbf{x} \mapsto \mathbf{x}^\top \mathbf{A}^{-1} \mathbf{x}$ is reached for $\mathbf{U}\mathbf{e}_I$, which is the eigenvector of \mathbf{A} corresponding to its lowest eigenvalue. □

Lemma 4. Let $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_N)$ be any diagonal matrix with strictly positive entries. Then for any vectors $(\mathbf{x}_t)_{1 \leq t \leq T}$, such that $\|\mathbf{x}_t\|_2 \leq 1$ for all $1 \leq t \leq T$, we have that the determinant $|\mathbf{V}_T|$ of $\mathbf{V}_T = \Lambda + \sum_{t=1}^T \mathbf{x}_t \mathbf{x}_t^\top$ is maximized when all \mathbf{x}_t are aligned with the axes.

Proof of Lemma 4. Let us write $d(\mathbf{x}_1, \dots, \mathbf{x}_T) = |\mathbf{V}_T|$ the determinant of \mathbf{V}_T . We want to characterize:

$$\max_{\mathbf{x}_1, \dots, \mathbf{x}_T: \|\mathbf{x}_t\|_2 \leq 1, \forall 1 \leq t \leq T} d(\mathbf{x}_1, \dots, \mathbf{x}_T)$$

For any $1 \leq t \leq T$, let us define:

$$\mathbf{V}_{-t} = \Lambda + \sum_{\substack{s=1 \\ s \neq t}}^T \mathbf{x}_s \mathbf{x}_s^\top$$

We have that $\mathbf{V}_T = \mathbf{V}_{-t} + \mathbf{x}_t \mathbf{x}_t^\top$. Consider the case with only simple eigenvalue multiplicities. In this case, Lemma 11 implies that $\mathbf{x}_t \mapsto d(\mathbf{x}_1, \dots, \mathbf{x}_t, \dots, \mathbf{x}_T)$ is maximized when \mathbf{x}_t is aligned with an eigenvector of \mathbf{V}_{-t} . Thus all \mathbf{x}_t , for $1 \leq t \leq T$, are aligned with an eigenvector of \mathbf{V}_{-t} and therefore also with an eigenvector of \mathbf{V}_T . Consequently, the eigenvectors of $\sum_{t=1}^T \mathbf{x}_t \mathbf{x}_t^\top$ are also aligned with \mathbf{V}_T . Since $\Lambda = \mathbf{V}_T - \sum_{t=1}^T \mathbf{x}_t \mathbf{x}_t^\top$ and Λ is diagonal, we conclude that \mathbf{V}_T is diagonal and all \mathbf{x}_t are aligned with the canonical axes.

Now in the case of eigenvalue multiplicities, the maximum of $|\mathbf{V}_T|$ may be reached by several sets of vectors $\{(\mathbf{x}_t^m)_{1 \leq t \leq T}\}_m$ but for some m^* , the set $(\mathbf{x}_t^{m^*})_{1 \leq t \leq T}$ will be aligned with the axes. In order to see that, consider a perturbed matrix $\mathbf{V}_{-t}^\varepsilon$ by a random perturbation of amplitude at most ε , i.e. such that $\mathbf{V}_{-t}^\varepsilon \rightarrow \mathbf{V}_{-t}$ when $\varepsilon \rightarrow 0$. Since the perturbation is random, then the probability that Λ^ε , as well as all other $\mathbf{V}_{-t}^\varepsilon$ possess an eigenvalue of multiplicity bigger than 1 is zero. Since the mapping $\varepsilon \mapsto \mathbf{V}_{-t}^\varepsilon$ is continuous, we deduce that any adherent point $\bar{\mathbf{x}}_t$ of the sequence $(\mathbf{x}_t^\varepsilon)_\varepsilon$ (there exists at least one since the sequence is bounded in ℓ_2 -norm) is aligned with the limit \mathbf{V}_{-t} and we can apply the previous reasoning. □

Lemma 5. For any T , let $\mathbf{V}_T = \sum_{t=1}^T \mathbf{x}_t \mathbf{x}_t^\top + \Lambda$. Then:

$$\log \frac{|\mathbf{V}_T|}{|\Lambda|} \leq \max \sum_{i=1}^N \log \left(1 + \frac{t_i}{\lambda_i} \right),$$

where the maximum is taken over all possible positive real numbers $\{t_1, \dots, t_N\}$, such that $\sum_{i=1}^N t_i = T$.

Proof of Lemma 5. We want to bound the determinant $|\mathbf{V}_T|$ under the coordinate constraints $\|\mathbf{x}_t\|_2 \leq 1$. Let:

$$M(\mathbf{x}_1, \dots, \mathbf{x}_T) = \left| \mathbf{\Lambda} + \sum_{t=1}^T \mathbf{x}_t \mathbf{x}_t^\top \right|$$

From Lemma 4 we deduce that the maximum of M is reached when all \mathbf{x}_t are aligned with the axes:

$$\begin{aligned} M &= \max_{\mathbf{x}_1, \dots, \mathbf{x}_T; \mathbf{x}_t \in \{\mathbf{e}_1, \dots, \mathbf{e}_N\}} \left| \mathbf{\Lambda} + \sum_{t=1}^T \mathbf{x}_t \mathbf{x}_t^\top \right| \\ &= \max_{t_1, \dots, t_N \text{ positive integers}, \sum_{i=1}^N t_i = T} \left| \text{diag}(\lambda_i + t_i) \right| \\ &\leq \max_{t_1, \dots, t_N \text{ positive reals}, \sum_{i=1}^N t_i = T} \prod_{i=1}^N (\lambda_i + t_i), \end{aligned}$$

from which we obtain the result. \square

Lemma 6. *Let d be the effective dimension. Then:*

$$\log \frac{|\mathbf{V}_T|}{|\mathbf{\Lambda}|} \leq 2d \log \left(1 + \frac{T}{\lambda} \right)$$

Proof of Lemma 6. Using Lemma 5 and Definition 1:

$$\begin{aligned} \log \frac{|\mathbf{V}_T|}{|\mathbf{\Lambda}|} &\leq \sum_{i=1}^d \log \left(1 + \frac{T}{\lambda} \right) + \sum_{i=d+1}^N \log \left(1 + \frac{t_i}{\lambda_d} \right) \\ &\leq d \log \left(1 + \frac{T}{\lambda} \right) + \sum_{i=1}^N \frac{t_i}{\lambda_{d+1}} \\ &\leq d \log \left(1 + \frac{T}{\lambda} \right) + \frac{T}{\lambda_{d+1}} \\ &\leq 2d \log \left(1 + \frac{T}{\lambda} \right) \end{aligned}$$

SPECTRALELIMINATOR \square

Lemma 7. *For any fixed $\mathbf{x} \in \mathbb{R}^N$ and any $\delta > 0$, we have that if $\beta(\boldsymbol{\alpha}^*, \delta) = 2R\sqrt{14 \log(2/\delta)} + \|\boldsymbol{\alpha}^*\|_{\mathbf{\Lambda}}$, then at time t_j (beginning of phase j):*

$$\mathbb{P} \left(|\mathbf{x}^\top (\hat{\boldsymbol{\alpha}}_j - \boldsymbol{\alpha}^*)| \leq \|\mathbf{x}\|_{\mathbf{V}_j^{-1}} \beta(\boldsymbol{\alpha}^*, \delta) \right) \geq 1 - \delta$$

Proof of Lemma 7. Defining $\boldsymbol{\xi}_j = \sum_{s=t_{j-1}+1}^{t_j} \mathbf{x}_s \varepsilon_s$, we have:

$$\begin{aligned} |\mathbf{x}^\top (\hat{\boldsymbol{\alpha}}_j - \boldsymbol{\alpha}^*)| &= |\mathbf{x}^\top (-\mathbf{V}_j^{-1} \mathbf{\Lambda} \boldsymbol{\alpha}^* + \mathbf{V}_j^{-1} \boldsymbol{\xi}_j)| \\ &\leq |\mathbf{x}^\top \mathbf{V}_j^{-1} \mathbf{\Lambda} \boldsymbol{\alpha}^*| + |\mathbf{x}^\top \mathbf{V}_j^{-1} \boldsymbol{\xi}_j| \end{aligned} \quad (4)$$

The first term in the right hand side of (4) is bounded as:

$$\begin{aligned} |\mathbf{x}^\top \mathbf{V}_j^{-1} \mathbf{\Lambda} \boldsymbol{\alpha}^*| &\leq \|\mathbf{x}^T \mathbf{V}_j^{-1} \mathbf{\Lambda}^{1/2}\| \|\mathbf{\Lambda}^{1/2} \boldsymbol{\alpha}^*\| \\ &= \|\boldsymbol{\alpha}^*\|_{\mathbf{\Lambda}} \sqrt{\mathbf{x}^\top \mathbf{V}_j^{-1} \mathbf{\Lambda} \mathbf{V}_j^{-1} \mathbf{x}} \\ &\leq \|\boldsymbol{\alpha}^*\|_{\mathbf{\Lambda}} \sqrt{\mathbf{x}^\top \mathbf{V}_j^{-1} \mathbf{x}} = \|\boldsymbol{\alpha}^*\|_{\mathbf{\Lambda}} \|\mathbf{x}\|_{\mathbf{V}_j^{-1}} \end{aligned}$$

Now consider the second term in the r.h.s. of (4). We have:

$$|\mathbf{x}^\top \mathbf{V}_j^{-1} \boldsymbol{\xi}_j| = \left| \sum_{s=t_{j-1}+1}^{t_j} (\mathbf{x}^\top \mathbf{V}_j^{-1} \mathbf{x}_s) \varepsilon_s \right|$$

Let us notice that the points (\mathbf{x}_s) selected by the algorithm during phase $j-1$ only depend on their width $\|\mathbf{x}\|_{\mathbf{V}_s^{-1}}$ which does not depend on the rewards received during the phase $j-1$. Thus, given \mathcal{F}_{j-2} , the sequence $(\mathbf{x}_s)_{t_{j-1}+1 \leq s < t_j}$ is deterministic. Consequently, one may use a variant of Azuma's inequality (Shamir, 2011):

$$\begin{aligned} \mathbb{P} \left(|\mathbf{x}^\top \mathbf{V}_j^{-1} \boldsymbol{\xi}_j|^2 \leq 28R^2 2 \log(2/\delta) \times \right. \\ \left. \times \mathbf{x}^\top \mathbf{V}_j^{-1} \left(\sum_{s=t_{j-1}+1}^{t_j} \mathbf{x}_s \mathbf{x}_s^\top \right) \mathbf{V}_j^{-1} \mathbf{x} \middle| \mathcal{F}_{j-2} \right) \geq 1 - \delta, \end{aligned}$$

from which we deduce:

$$\mathbb{P} \left(|\mathbf{x}^\top \mathbf{V}_j^{-1} \boldsymbol{\xi}_j|^2 \leq 56R^2 \mathbf{x}^\top \mathbf{V}_j^{-1} \mathbf{x} \log(2/\delta) \middle| \mathcal{F}_{j-2} \right) \geq 1 - \delta,$$

since $\sum_{s=t_{j-1}+1}^{t_j} \mathbf{x}_s \mathbf{x}_s^\top \prec \mathbf{V}_j$. Thus:

$$\mathbb{P} \left(|\mathbf{x}^\top \mathbf{V}_j^{-1} \boldsymbol{\xi}_j| \leq 2R \|\mathbf{x}\|_{\mathbf{V}_j^{-1}} \sqrt{14 \log(2/\delta)} \right) \geq 1 - \delta$$

Lemma 8. *For all $\mathbf{x} \in A_j$, we have:* \square

$$\|\mathbf{x}\|_{\mathbf{V}_j^{-1}}^2 \leq \frac{1}{t_j - t_{j-1}} \sum_{s=t_{j-1}+1}^{t_j} \|\mathbf{x}_s\|_{\mathbf{V}_{s-1}^{-1}}^2$$

Proof of Lemma 8. We have:

$$\begin{aligned} (t_j - t_{j-1}) \|\mathbf{x}\|_{\mathbf{V}_j^{-1}}^2 &\leq \max_{\mathbf{x} \in A_j} \sum_{s=t_{j-1}+1}^{t_j} \|\mathbf{x}\|_{\mathbf{V}_{s-1}^{-1}}^2 \\ &\leq \sum_{s=t_{j-1}+1}^{t_j} \max_{\mathbf{x} \in A_j} \|\mathbf{x}\|_{\mathbf{V}_{s-1}^{-1}}^2 \\ &\leq \sum_{s=t_{j-1}+1}^{t_j} \max_{\mathbf{x} \in A_{j-1}} \|\mathbf{x}\|_{\mathbf{V}_{s-1}^{-1}}^2 \\ &= \sum_{s=t_{j-1}+1}^{t_j} \|\mathbf{x}_s\|_{\mathbf{V}_{s-1}^{-1}}^2, \end{aligned}$$

since the algorithm selects (during phase $j-1$) the arms with largest width. \square

Lemma 9. *For each phase j , we have:*

$$\sum_{s=t_{j-1}+1}^{t_j} \min \left(1, \|\mathbf{x}_s\|_{\mathbf{V}_{s-1}^{-1}}^2 \right) \leq \log \frac{|\mathbf{V}_j|}{|\mathbf{\Lambda}|}.$$

Proof of Lemma 9. This lemma is proved by instantiating Lemma 2 for each phase. \square