# High-dimensional Bayesian asymptotics and computation 

Aguemon Yves Atchadé

University of Michigan

1 Graphical models

2 Computations

3 A Gaussian graphical example

4 Conclusion

## Graphical models

- Graphs useful to represent dependencies between random variables.
- Two main types of graphical models
- Directed acylic graph (DAG); a.k.a. Bayesian networks
- Undirected graph; known as Markov networks. Main topic.


## Graphical models

- Graphs useful to represent dependencies between random variables.
- Two main types of graphical models
- Directed acylic graph (DAG); a.k.a. Bayesian networks
- Undirected graph; known as Markov networks. Main topic.
- Useful in many applications: speech recognition, biological networks modeling, protein folding problems, etc...
- Some notation: $\mathcal{M}_{p}$ space of $p \times p$ symmetric matrices. $\mathcal{M}_{p}^{+}$its cone of spd elements,

$$
\langle A, B\rangle_{F} \stackrel{\text { def }}{=} \sum_{i \leq j} A_{i j} B_{i j}, \quad A, B \in \mathcal{M}_{p} .
$$

## Graphical models

- A parametric graphical model:
- $p$ nodes. A set $\mathrm{Y} \subset \mathbb{R}$.

■ Non-zero functions $B_{0}: \mathrm{Y} \rightarrow \mathbb{R}$, and $B: \mathrm{Y} \times \mathrm{Y} \rightarrow \mathbb{R}$ symmetric.

- Then define $\left\{f_{\theta}, \theta \in \Omega\right\}$,

$$
\begin{gathered}
f_{\theta}(y)=\frac{1}{Z(\theta)} \exp \left(\sum_{j=1}^{p} \theta_{j j} B_{0}\left(y_{j}\right)+\sum_{i<j} \theta_{i j} B\left(y_{i}, y_{j}\right)\right), \\
\Omega \stackrel{\text { def }}{=}\left\{\theta \in \mathcal{M}_{p}: Z(\theta) \stackrel{\text { def }}{=} \int e^{-\langle\theta, \bar{B}(y)\rangle_{\mathrm{F}}} y<\infty\right\} .
\end{gathered}
$$

## Graphical models

- Parametric model $\left\{f_{\theta}, \theta \in \Omega\right\}$.
- The parameter $\theta \in \Omega$ modulates the interaction. Importantly, $\theta_{i j}=0$ implies conditional independence of $y_{i}, y_{j}$ given remaining variables.


## Graphical models

- Parametric model $\left\{f_{\theta}, \theta \in \Omega\right\}$.
- The parameter $\theta \in \Omega$ modulates the interaction. Importantly, $\theta_{i j}=0$ implies conditional independence of $y_{i}, y_{j}$ given remaining variables.
- It is often very appealing to assume that $\theta$ is sparse, particularly when $p$ is large.
- Goal: estimate $\theta \in \Omega$ from multiple ( $n$ ) samples from $f_{\theta_{\star}}$ arranged in a data matrix $Z \in \mathbb{R}^{n \times p}$.


## Graphical models

■ Given a prior $\Pi$ on $\Omega$. Main object of interest:

$$
\Pi(\mathrm{d} \theta \mid Z) \propto \Pi(\mathrm{d} \theta) \prod_{i=1}^{n} f_{\theta}\left(Z_{i .}\right)
$$

- Set $\Delta$ the set of graph-skeletons (symmetric $0-1$ matrices with diagonal 1). For sparse estimation, we consider priors of the form

$$
\Pi(\mathrm{d} \theta)=\sum_{\delta \in \Delta} \pi_{\delta} \Pi(\mathrm{d} \theta \mid \delta)
$$

- where $\Pi(\mathrm{d} \theta \mid \delta)$ has support $\Omega(\delta)$.


## Graphical models

- Given a prior $\Pi$ on $\Omega$. Main object of interest:

$$
\Pi(\mathrm{d} \theta \mid Z) \propto \Pi(\mathrm{d} \theta) \prod_{i=1}^{n} f_{\theta}\left(Z_{i .}\right)
$$

- Set $\Delta$ the set of graph-skeletons (symmetric $0-1$ matrices with diagonal 1). For sparse estimation, we consider priors of the form

$$
\Pi(\mathrm{d} \theta)=\sum_{\delta \in \Delta} \pi_{\delta} \Pi(\mathrm{d} \theta \mid \delta)
$$

- where $\Pi(\mathrm{d} \theta \mid \delta)$ has support $\Omega(\delta)$.
- Difficulty with $\Pi(\cdot \mid Z)$ : Either the likelihood is intractable,
- $\operatorname{Or} \Omega(\delta)$ is a complicated space and prior is intractable.


## Quasi-Bayesian inference

■ In large applications, it may be worth exploring less accurate but faster alternatives.

- Quasi-Bayesian inference is a framework to formulate these trade-offs.
- Think of Quasi-Bayesian inference as the Bayesian analog of M-estimation.
- General idea: instead of the model $\left\{f_{\theta}, \theta \in \Omega\right\}$, we consider a "larger pseudo-model" $\left\{\check{f}_{\theta}, \theta \in \check{\Omega}\right\}$.


## Quasi-Bayesian inference

- Pseudo-model: $z \mapsto \check{f}_{\theta}(z)$ needs not be a density. Chosen for computational convenience.
- Larger pseudo-model: $\Omega \subseteq \Omega$. Very useful to build interesting priors on $\check{\Omega}(\delta)$.


## Quasi-Bayesian inference

- Pseudo-model: $z \mapsto \check{f_{\theta}}(z)$ needs not be a density. Chosen for computational convenience.
- Larger pseudo-model: $\Omega \subseteq \Omega$. Very useful to build interesting priors on $\Omega(\delta)$.
- Quasi-posterior distributions have been used extensively in the PAC-Bayesian literature (Catoni 2004).
- ABC is a form of quasi-Bayesian inference.

■ Chernozukhov-Hong (J. Econ. 2003). Also popular in Bayesian semi-parametric inference (Yang \& He (AoS 2012), Kato (AoS 2013).

## Asymptotics of quasi-posterior distributions

Consider the quasi-posterior distribution

$$
\check{\Pi}(\mathrm{d} \theta \mid Z) \propto q_{\theta}(Z) \Pi(\mathrm{d} \theta) .
$$

## Asymptotics of quasi-posterior distributions

Consider the quasi-posterior distribution

$$
\check{\Pi}(\mathrm{d} \theta \mid Z) \propto q_{\theta}(Z) \Pi(\mathrm{d} \theta) .
$$

## Theorem

$\check{\Pi}(\cdot \mid Z)$ is a solution to the problem

$$
\min _{\mu \ll \Pi}\left[-\int_{\mathbb{R}^{d}} \log q_{\theta}(Z) \mu(\mathrm{d} \theta)+K L(\mu \mid \Pi)\right],
$$

where $K L(\mu \mid \Pi) \stackrel{\text { def }}{=} \int_{\mathbb{R}^{d}} \log (\mathrm{~d} \mu / \mathrm{d} \Pi) \mathrm{d} \mu$ is the $K L$-divergence of $\Pi$ from $\mu$.

- Proof is Easy. See e.g. T. Zhang (AoS 2006).
- If $q_{\theta}$ is good enough for a frequentist M -estimation inference, it is good enough for a quasi-Bayesian inference- upto the prior.


## Example: binary graphical models

■ Binary graphical model. $\mathrm{Y}=\{0,1\} . B(x, y)=x y$. Here $\Omega=\mathcal{M}_{p}$ and
$Z(\theta)$ is typically intractable.

## Example: binary graphical models

■ Binary graphical model. $\mathrm{Y}=\{0,1\} . B(x, y)=x y$. Here $\Omega=\mathcal{M}_{p}$ and
$Z(\theta)$ is typically intractable.

- There is a very commonly used pseudo-likelihood function to circumvent the intractable normalizing constant.

$$
\begin{aligned}
q_{\theta}(Z) & =\prod_{i=1}^{n} \prod_{j=1}^{p} \frac{\exp \left(Z_{i j}\left(\theta_{j j}+\sum_{k \neq j} \theta_{j k} Z_{i k}\right)\right)}{1+\exp \left(\theta_{j j}+\sum_{k \neq j} \theta_{j k} Z_{i k}\right)}, \theta \in \mathcal{M}_{p}, \\
& =\prod_{i=1}^{n} \prod_{j=1}^{p} f_{\theta \cdot j}^{(j)}\left(Z_{i j} \mid Z_{i,-j}\right) \theta \in \mathcal{M}_{p},
\end{aligned}
$$

- Note: $f_{\theta \cdot j}^{(j)}\left(Z_{i j} \mid Z_{i,-j}\right)$ depends only on the $j$-th column of $\theta$.


## Example: binary graphical models

- Then very easy to set up prior of $\mathcal{M}_{p}(\delta)$.
- However, dimension of $\mathcal{M}_{p}$ grows fast. Larger than $10^{5}$, for $p \approx 500$.


## Example: binary graphical models

- Then very easy to set up prior of $\mathcal{M}_{p}(\delta)$.
- However, dimension of $\mathcal{M}_{p}$ grows fast. Larger than $10^{5}$, for $p \approx 500$.
- We can further simplify the problem by enlarging the parameter space from $\mathcal{M}_{p}$ to $\mathbb{R}^{p \times p}$ :

$$
\begin{aligned}
q_{\theta}(Z) & =\prod_{i=1}^{n} \prod_{j=1}^{p} f_{\theta \cdot j}^{(j)}\left(Z_{i j} \mid Z_{i,-j}\right) \theta \in \mathbb{R}^{p \times p} \\
& =\prod_{j=1}^{p}\left(\prod_{i=1}^{n} f_{\theta \cdot j}^{(j)}\left(Z_{i j} \mid Z_{i,-j}\right)\right), \quad \theta \in \mathbb{R}^{p \times p} .
\end{aligned}
$$

- In that case $q_{\theta}(Z)$ factorizes along the columns of $\theta$.


## Example: binary graphical models

- Take $p$ independent sparsity inducing priors on $\mathbb{R}^{p}$, and we get a posterior on $\mathbb{R}^{p \times p}$ :

$$
\check{\Pi}(\mathrm{d} \theta \mid Z)=\prod_{j=1}^{p} \check{\Pi}_{j}(\mathrm{~d} \theta \cdot \cdot j \mid Z),
$$

where

$$
\check{\Pi}_{j}(\mathrm{~d} u \mid Z)=\prod_{i=1}^{n} f_{\theta \cdot j}^{(j)}\left(Z_{i j} \mid Z_{i,-j}\right) \sum_{\delta \in \Delta_{p}} \pi_{\delta} \Pi(\mathrm{d} \theta \mid \delta) .
$$

- We can sample from the distribution $\check{\Pi}_{j}(\mathrm{~d} \theta \mid Z)$ in parallel. Potentially huge computing gain.


## Example: binary graphical models

- Very popular method for fitting large graphical models in frequentist inference.
- Initially introduced by Meinhausen \& Buhlmann (AoS 2006), for Gaussian graphical models.
- See also Ravikumar et al. (AoS 2010) for binary graphical models. Sun \& Zhang (JMLR, 2013) for a scaled-Lasso version.
- Very efficient (divide and conquer). We can fit $p=1000$ nodes in few minutes on large clusters.


## Example: binary graphical models

- Very popular method for fitting large graphical models in frequentist inference.
- Initially introduced by Meinhausen \& Buhlmann (AoS 2006), for Gaussian graphical models.
- See also Ravikumar et al. (AoS 2010) for binary graphical models. Sun \& Zhang (JMLR, 2013) for a scaled-Lasso version.
- Very efficient (divide and conquer). We can fit $p=1000$ nodes in few minutes on large clusters.
- Loss of symmetry.
- Should we worry about all the simplification involved?


## Example: binary graphical models

- Assume we build the prior $\Pi \mathbb{R}^{p}$ as follows.

$$
\begin{equation*}
\Pi(\mathrm{d} \theta)=\sum_{\delta \in \Delta_{p}} \pi_{\delta} \Pi(\mathrm{d} \theta \mid \delta) . \tag{1}
\end{equation*}
$$

- $\pi_{\delta}=\prod_{j=1}^{p} q^{\delta_{j}}(1-q)^{1-\delta_{j}}, q=p^{-u}, u>1$.

$$
\begin{gather*}
\theta_{j} \left\lvert\, \delta \sim\left\{\begin{array}{cc}
\operatorname{Dirac}(0) & \text { if } \delta_{j}=0 \\
\operatorname{Laplace}(\rho) & \text { if } \delta_{j}=1
\end{array},\right.\right.  \tag{2}\\
\rho=24 \sqrt{n \log (p)} .
\end{gather*}
$$

- See Castillo et al. (AoS 2015).


## Example: binary graphical models

H1: There exists $\theta_{\star} \in \mathcal{M}_{p}$ such that the rows of $Z$ are i.i.d. $f_{\theta_{\star}}$.

- Set

$$
s_{\star} \stackrel{\text { def }}{=} \max _{1 \leq j \leq p} \sum_{i=1}^{p} \mathbf{1}_{\left\{\left|\theta_{i j}\right|>0\right\}}
$$

the max. degree of $\theta_{\star}$.

- For $\theta \in \mathbb{R}^{p \times p}$, define the norm

$$
\|\theta\| \stackrel{\text { def }}{=} \sup _{1 \leq j \leq p}\left\|\theta_{\cdot j}\right\|_{2} .
$$

## Example: binary graphical models

## Theorem (A.A.(2015))

With prior and assumption above, and under some regularity conditions, define

$$
r_{n, d}=\frac{1}{\underline{\kappa}\left(s_{\star}\right)} \sqrt{\frac{s_{\star} \log (p)}{n}} .
$$

There exists universal constants $M>2, A_{1}>0, A_{2}>0$ such that for $p$ large enough, and

$$
\begin{gathered}
n \geq A_{1}\left(\frac{s_{\star}}{\underline{\kappa}\left(s_{\star}\right)}\right)^{2} \log (p), \\
\mathbb{E}\left[\check{\Pi}\left(\left\{\theta \in \mathbb{R}^{p \times p}:\left\|\theta-\theta_{\star}\right\|>M_{0} r_{n, d}\right\} \mid Z\right)\right] \leq \frac{2}{e^{A_{2} n}}+\frac{12}{d} .
\end{gathered}
$$

## Example: binary graphical models

- Gives some guarantee that the method is not completely silly.
- Regularity conditions: restricted smallest eigenvalues of Fisher information matrix bounded away from 0 .
- Minimax rate. Even in full likelihood inference cannot do better in term of convergence rate.
- Extension to more general class of prior is possible.
- Similar results hold for Gaussian graphical models, and more general models.

1 Graphical models

2 Computations

3 A Gaussian graphical example

4 Conclusion

## Approximate Computations

- How to sample from

$$
\check{\Pi}(\mathrm{d} \theta \mid Z)=q_{\theta}(Z) \sum_{\delta \in \Delta_{p}} \pi_{\delta} \prod_{j: \delta_{j}=1} \phi\left(\theta_{j}\right) \mu_{p, \delta}(\mathrm{~d} \theta) ?
$$

- Rather we consider:

$$
\check{\Pi}(\delta, \mathrm{d} \theta \mid Z)=\pi_{\delta} \exp \left(\log q_{\theta}(Z)+\sum_{j=1}^{p} \delta_{j} \log \phi\left(\theta_{j}\right)\right) \mu_{p, \delta}(\mathrm{~d} \theta) .
$$

## Approximate Computations

- How to sample from

$$
\check{\Pi}(\mathrm{d} \theta \mid Z)=q_{\theta}(Z) \sum_{\delta \in \Delta_{p}} \pi_{\delta} \prod_{j: \delta_{j}=1} \phi\left(\theta_{j}\right) \mu_{p, \delta}(\mathrm{~d} \theta) ?
$$

- Rather we consider:

$$
\check{\Pi}(\delta, \mathrm{d} \theta \mid Z)=\pi_{\delta} \exp \left(\log q_{\theta}(Z)+\sum_{j=1}^{p} \delta_{j} \log \phi\left(\theta_{j}\right)\right) \mu_{p, \delta}(\mathrm{~d} \theta) .
$$

- Issue: for $\delta \neq \delta^{\prime}, \check{\Pi}(\mathrm{d} \theta \mid \delta, Z)$ and $\check{\Pi}\left(\mathrm{d} \theta \mid \delta^{\prime}, Z\right)$ are singular measures.
- We want to avoid transdimensional MCMC techniques (reversible-jump style MCMC). Poor mixing.
- We propose an approximation method using the Moreau envelops.


## Approximate Computations

■ Suppose $h: \mathbb{R}^{p} \rightarrow(-\infty,+\infty]$ is convex (possibly not smooth).
■ For $\gamma>0$, the Moreau-Yosida approximation of $h$ is:

$$
h_{\gamma}(\theta)=\min _{u \in \mathbb{R}^{p}}\left[h(u)+\frac{1}{2 \gamma}\|u-\theta\|^{2}\right] .
$$

- $h_{\gamma}$ is convex, class $\mathcal{C}^{1}$ with Lip. gradient, and $h_{\gamma} \uparrow h$ pointwise as $\gamma \rightarrow 0$.
- Well-studied approximation method.
- Leads to the proximal algorithm.


## Approximate Computations

- In many cases, $h_{\gamma}$ cannot be computed/evaluated.
- If $h=f+g$, and $f$ is smooth, one can use the forward-backward approximation

$$
\tilde{h}_{\gamma}(x)=\min _{u \in \mathbb{R}^{d}}\left[f(x)+\langle\nabla f(x), u-x\rangle+g(u)+\frac{1}{2 \gamma}\|u-x\|^{2}\right] .
$$

- $\tilde{h}_{\gamma} \leq h_{\gamma} \leq h$, and has similar properties as $h_{\gamma}$.
- $h_{\gamma}$ is easy to compute when $g$ is simple enough.

■ Explored by (Pereyra (Stat. Comp. (2015), Schrek et al. (2014)) as proposal mechanism in MCMC.


Figure: Figure showing the function $h(x)=-a x+\log \left(1+e^{a x}\right)+b|x|$ for $a=0.8, b=0.5$ (blue/solid line), and the approximations $h_{\gamma}$ and $\tilde{h}_{\gamma}$ $\left(h_{\gamma} \leq \tilde{h}_{\gamma}\right)$, for $\gamma \in\{5,1,0.1\}$.

## Approximate Computations

- For $\gamma>0$, the Moreau-Yosida approximation of $h$ is:

$$
h_{\gamma}(\theta)=\min _{u \in \mathbb{R}^{p}}\left[h(u)+\frac{1}{2 \gamma}\|u-\theta\|^{2}\right]
$$

- Notice that even if $\operatorname{dom}(h) \neq \mathbb{R}^{p}, h_{\gamma}$ is still finite everywhere.
- Hence if $h(x)=-\log \pi(x)$ for some log-concave density $\pi$

$$
\pi_{\gamma}(x)=\frac{1}{Z_{\gamma}} e^{-h_{\gamma}(x)}, x \in \mathbb{R}^{p},
$$

is an approximation of $\pi$ (assume $Z_{\gamma}<\infty$ ), and $\pi_{\gamma} \ll \operatorname{Leb}_{\mathbb{R}^{d}}$.

- We show that $\pi_{\gamma}$ converges weakly to $\pi$ as $\gamma \rightarrow 0$.


## Approximate Computations

- Back to $\check{\Pi}(\cdot \mid Z)$.

$$
\begin{aligned}
\check{\Pi}(\delta, \mathrm{d} \theta \mid Z) & \propto \\
& \pi_{\delta} \exp \left(\log q_{\theta}(Z)+\sum_{j=1}^{p} \delta_{j} \log \phi\left(\theta_{j}\right)\right) \mu_{p, \delta}(\mathrm{~d} u), \\
& \propto \pi_{\delta} \exp [\underbrace{\log q_{\theta}(Z)+\sum_{j=1}^{d} \delta_{j} \log \phi\left(\theta_{j}\right)-\iota_{\Theta_{\delta}}(\theta)}_{-h(\theta \mid \delta)}] \mu_{p, \delta}(\mathrm{~d} u) .
\end{aligned}
$$

- Leads to

$$
\check{\Pi}_{\gamma}(\delta, \mathrm{d} \theta) \propto \pi_{\delta}(2 \pi \gamma)^{\|\delta\|_{1} / 2} e^{-h_{\gamma}(\theta \mid \delta)} \mathrm{d} \theta,
$$

where $h_{\gamma}(\cdot \mid \delta)$ is the forward-backward approx. of $h$.

## Approximate Computations

$$
\check{\Pi}_{\gamma}(\delta, \mathrm{d} \theta) \propto \pi_{\delta}(2 \pi \gamma)^{\frac{\|\delta\|_{1}}{2}} e^{-h_{\gamma}(\theta \mid \delta)} \mathrm{d} \theta
$$

## Approximate Computations

$$
\check{\Pi}_{\gamma}(\delta, \mathrm{d} \theta) \propto \pi_{\delta}(2 \pi \gamma)^{\frac{\|\delta\|_{1}}{2}} e^{-h_{\gamma}(\theta \mid \delta)} \mathrm{d} \theta
$$

- Assume: $-\log q_{\theta}(Z)$ is convex, has L-Lip. gradient, and

$$
-\log q_{\theta}(Z) \geq \frac{1}{2 L}\left\|\nabla \log q_{\theta}(Z)\right\|^{2}
$$

- Assume: $-\log \phi$ is convex.


## Theorem

Take $\gamma=\gamma_{0} / L, \gamma_{0} \in(0,1 / 4]$. Then $\Pi_{\gamma}$ is a well-defined p.m. on $\Delta_{p} \times \mathbb{R}^{p}$, and there exists a finite constant (in $p$ ) $C$ such that

$$
\beta\left(\check{\Pi}_{\gamma}, \check{\Pi}\right) \leq \sqrt{\gamma_{0}}+C \gamma_{0} p,
$$

where $\beta(\cdot, \cdot)$ is the $\beta$-metric between p.m. (metricizes weak convergence).

## Approximate Computations

- In theory, we get better bound by taking for e.g.

$$
\gamma=\frac{\gamma_{0}}{L p} .
$$

■ However as $\check{\Pi}_{\gamma}$ gets very close to $\check{\Pi}$, sampling from $\check{\Pi}_{\gamma}$ becomes hard.

- The theorem above is a worst case analysis. What is the behavior for typical data realizations?

Relative Error


Structure Recovery


Figure: Sparse Bayesian linear regression example. $p=500, n=200$.

## Approximate Computations

$$
\check{\Pi}_{\gamma}(\delta, \mathrm{d} \theta) \propto \pi_{\delta}(2 \pi \gamma)^{\frac{\|\delta\|_{1}}{2}} e^{-h_{\gamma}(\theta \mid \delta)} \mathrm{d} \theta
$$

- Linear regression: $-\log q_{\theta}(Z)=\|Z-X \theta\|^{2} / 2 \sigma^{2}$.
- Assume $Z \sim \mathbf{N}\left(X \theta_{\star}, \sigma^{2} I_{n}\right)$.
- Assume: the sparse prior assumption in Theorem 1.


## Theorem

Take $\gamma=\gamma_{0} / L, \gamma_{0} \in(0,1 / 4]$. There exists a finite constant (in $p$ ) $C$ such that

$$
\mathbb{E}\left[\beta\left(\check{\Pi}_{\gamma}, \check{\Pi}\right)\right] \leq \sqrt{\gamma_{0}}+C\left(1+\gamma_{0} \log (p)\right) .
$$

## Approximate Computations

$$
\check{\Pi}_{\gamma}(\delta, \mathrm{d} \theta) \propto \pi_{\delta}(2 \pi \gamma)^{\frac{\|\delta\|_{1}}{2}} e^{-h_{\gamma}(\theta \mid \delta)} \mathrm{d} \theta
$$

■ We can sample from $\Pi$ Ǐ using "standard" MCMC methods.

- Key advantage: given $\theta$, the comp. of $\delta$ are conditionally indep. Bernoulli.
- Given $\delta$, do a Metropolis-Langevin approach that takes adv. of the smoothness of $h_{\gamma}$.
- The gradient of $\theta \mapsto h_{\gamma}(\theta \mid \delta)$ is related to the proximal map of $h$.

1 Graphical models

2 Computations

3 A Gaussian graphical example

4 Conclusion

## A Gaussian graphical example

- Example: sparse estimation of large Gaussian graphical models.
- We compare the quasi-posterior mean and g-lasso estimator

$$
\hat{\vartheta}_{\text {glasso }}=\operatorname{Argmin}_{\theta \in \mathcal{M}_{p}^{+}}\left[-\log \operatorname{det} \theta+\operatorname{Tr}(\theta S)+\lambda \sum_{i, j}\left(\alpha\left|\theta_{i j}\right|+\frac{(1-\alpha)}{2} \theta_{i j}^{2}\right)\right],
$$

where $S=(1 / n) Z^{\prime} Z$.

- We do comparison along:

$$
\begin{align*}
& \mathcal{E}=\frac{\|\hat{\vartheta}-\vartheta\|_{F}}{\|\vartheta\|_{F}}, \quad \text { SEN }=\frac{\sum_{i<j} \mathbf{1}_{\left\{\left|\vartheta_{i j}\right|>0\right\}} \mathbf{1}_{\left\{\operatorname{sign}\left(\hat{\vartheta}_{i j}\right)=\operatorname{sign}\left(\vartheta_{i j}\right)\right\}}}{\sum_{i<j}} ; \\
& \text { and } \quad \text { PREC }=\frac{\sum_{i<j} \mathbf{1}_{\left\{\left|\hat{\vartheta}_{i j}\right|>0\right\}} \mathbf{1}_{\left\{\operatorname{sign}\left(\hat{\vartheta}_{i j}\right)=\operatorname{sign}\left(\vartheta_{i j}\right)\right\}}}{\sum_{i<j} \mathbf{1}_{\left\{\left|\hat{v}_{i j}\right|>0\right\}}} . \tag{3}
\end{align*}
$$

## A Gaussian graphical example

|  | $\vartheta_{j j}^{2}$ known | Empirical Bayes | Glasso |
| :--- | :---: | :---: | :---: |
| Relative Error (E in \%) | 19.2 | 21.6 | 63.1 |
| Sensitivity (SEN in \%) | 68.4 | 69.0 | 40.5 |
| Precision (PREC in \%) | 100.0 | 100.0 | 74.9 |

Table: Table showing the relative error, sensitivity and precision (as defined in (3)) for Setting (a), with $p=100$ nodes. Based on 20 simulation replications.

Each MCMC run is $5 \times 10^{4}$ iterations.

## A Gaussian graphical example

|  | $\vartheta_{j j}^{2}$ known | Empirical Bayes | Glasso |
| :--- | :---: | :---: | :---: |
| Relative Error ( $\mathcal{E}$ in \%) | 23.1 | 26.2 | 45.2 |
| Sensitivity (SEN in \%) | 44.6 | 45.4 | 87.9 |
| Precision (PREC in \%) | 100 | 99.9 | 56.1 |

Table: Table showing the relative error, sensitivity and precision (as defined in (3)) for Setting (b), with $p=500$ nodes. Based on 20 simulation replications. Each MCMC run is $5 \times 10^{4}$ iterations.

## A Gaussian graphical example

|  | $\vartheta_{j j}^{2}$ known | Empirical Bayes | Glasso |
| :--- | :---: | :---: | :---: |
| Relative Error (E in \%) | 30.8 | 35.2 | 66.9 |
| Sensitivity (SEN in \%) | 16.3 | 16.4 | 6.6 |
| Precision (PREC in \%) | 99.9 | 99.8 | 94.7 |

Table: Table showing the relative error, sensitivity and precision (as defined in (3)) for Setting (c), with $p=1,000$ nodes. Based on 20 simulation replications. Each MCMC run is $5 \times 10^{4}$ iterations.

## A Gaussian graphical example



Figure: Figure showing the confidence interval bars (obtained from one MCMC run), for the non-diagonal entries of $\vartheta$ in Setting (a). The dots represent the true values.

1 Graphical models

2 Computations

3 A Gaussian graphical example

## 4 Conclusion

## Conclusion

- Quasi-posterior inference is consistent in high-dimensional setting.
- On the approx. computation, how to formalize the trade-off between good approx. and fast MCMC computation.
- Joint statistical and computational asymptotics.
- Matlab code available from website.


## Conclusion

- Quasi-posterior inference is consistent in high-dimensional setting.
- On the approx. computation, how to formalize the trade-off between good approx. and fast MCMC computation.
- Joint statistical and computational asymptotics.
- Matlab code available from website.
- Postdoc opening available at the University of Michigan: www.stat.lsa.umich.edu/~yvesa


## Conclusion

- Quasi-posterior inference is consistent in high-dimensional setting.
- On the approx. computation, how to formalize the trade-off between good approx. and fast MCMC computation.
- Joint statistical and computational asymptotics.
- Matlab code available from website.
- Postdoc opening available at the University of Michigan: www.stat.lsa.umich.edu/~yvesa
- Thanks for your attention... and patience !

