# ADAPTIVE ESTIMATION OF STATIONARY GAUSSIAN FIELDS 

By Nicolas Verzelen ${ }^{1}$<br>INRA and SUPAGRO


#### Abstract

We study the nonparametric covariance estimation of a stationary Gaussian field $X$ observed on a regular lattice. In the time series setting, some procedures like AIC are proved to achieve optimal model selection among autoregressive models. However, there exists no such equivalent results of adaptivity in a spatial setting. By considering collections of Gaussian Markov random fields (GMRF) as approximation sets for the distribution of $X$, we introduce a novel model selection procedure for spatial fields. For all neighborhoods $m$ in a given collection $\mathcal{M}$, this procedure first amounts to computing a covariance estimator of $X$ within the GMRFs of neighborhood $m$. Then it selects a neighborhood $\widehat{m}$ by applying a penalization strategy. The so-defined method satisfies a nonasymptotic oracle-type inequality. If $X$ is a GMRF, the procedure is also minimax adaptive to the sparsity of its neighborhood. More generally, the procedure is adaptive to the rate of approximation of the true distribution by GMRFs with growing neighborhoods.


1. Introduction. In this paper, we study the estimation of the distribution of a stationary Gaussian field $X=\left(X_{[i, j]}\right)_{(i, j) \in \Lambda}$ indexed by the nodes of a square lattice $\Lambda$ of size $p \times p$. This problem is often encountered in spatial statistics or in image analysis.

Various estimation methods have been proposed to handle this question. Most of them fall into two categories. On the one hand, one may consider direct covariance estimation. A traditional approach amounts to computing an empirical variogram and then fitting a suitable parametric variogram model such as the exponential or Matérn model (Cressie [10], Chapter 2). Some procedures also apply to nonregular lattices. However, a bad choice of the variogram model may lead to poor results. The issue of variogram model selection has not been completely solved yet although some procedures based on cross-validation have been proposed. See [10], Section 2.6.4, for a discussion. Most of the nonparametric (Hall, Fisher and Hoffmann [19]) and semiparametric (Im, Stein and Zhu [21]) methods are based on the spectral representation of the field. To our knowledge, these procedures have not yet been shown to achieve adaptiveness; that is, their rate of convergence does not adapt to the complexity of the correlation functions.

[^0]An alternative approach to the problem amounts to considering the conditional distribution at one node given the remaining nodes. This point of view is closely connected to the notion of Gaussian Markov random field (GMRF). Let $\mathcal{G}$ be a graph whose vertex set is $\Lambda$. The field $X$ is GMRF with respect to $\mathcal{G}$ if it satisfies the following property: for any node $(i, j) \in \Lambda$, conditionally to the set of variables $X_{[k, l]}$ such that $(k, l)$ is a neighbor of $(i, j)$ in $\mathcal{G}, X_{[i, j]}$ is independent from all the remaining variables. GMRFs are also sometimes called Gaussian graphical models. A huge literature develops around this subject since Gaussian graphical models are promising tools to analyze complex high-dimensional systems involved, for instance, in postgenomic data. In other applications, GMRFs are relevant because they allow one to perform a Markov chain Monte Carlo run quickly using Markov properties (e.g., [31]). See Lauritzen [24] or Edwards [14] for introductions to Gaussian graphical models and Markov properties. In the sequel, we assume that the node $(0,0)$ belongs to $\Lambda$. Since we assume that the field $X$ is stationary, defining a graph $\mathcal{G}$ is equivalent to defining the neighborhood $m$ of the node $(0,0)$. Indeed, the neighborhood of any node $(i, j) \in \Lambda$ is the transposition of $m$ by $(i, j)$. In the sequel, we call $m$ the neighborhood of a GMRF. If the neighborhood is empty, then the Markov property states that the components of $X$ are all independent. Alternatively, any zero-mean Gaussian stationary field is a GMRF with respect to the complete neighborhood [i.e., containing all the nodes except $(0,0)$ ].

Numerous papers have been devoted to parametric estimation for stationary GMRFs with a known neighborhood. The authors have derived their asymptotic properties of such estimators (see $[3,5,16]$ ). If the field $X$ is assumed to be a GMRF with respect to a known neighborhood, in each of these works, the issue of neighborhood selection has been less studied. Besag and Kooperberg [4], Rue and Tjelmeland [31], Song, Fuentes and Ghosh [33] and Cressie and Verzelen [11] have tackled the problem of approximating the distribution of a Gaussian field by a GMRF, but this requires the knowledge of the true distribution. Guyon and Yao have stated in [18] necessary conditions and sufficient conditions for a model selection procedure to choose asymptotically the true neighborhood of a GMRF with probability one.

In this paper, we study a nonparametric estimation procedure based on neighborhood selection. In short, we select a suitable neighborhood and estimate the distribution of $X$ in the space of stationary GMRFs with respect to this neighborhood. The objective is not to estimate the "true" neighborhood. We rather want to select a neighborhood that allows to estimate well the distribution of $X$ (i.e., to minimize a risk). In fact, we do not even assume that the true correlation of $X$ corresponds to a GMRF. This estimation procedure is relevant for two main reasons:

- To our knowledge, it is the first nonparametric estimator in a spatial setting which achieves adaptive rates of convergence.
- In most of the statistical applications where GMRFs are involved, the neighborhood is a priori unknown. Our procedure allows one to select a "good" neighborhood.

Our problem on a two-dimensional field has a natural one-dimensional counterpart in time series analysis. It is indeed known that an auto-regressive process (AR) of order $p$ is also a GMRF with $2 p$ nearest neighbors and reciprocally (see [17], Section 1.3). In this one-dimensional setting, our issue reformulates as follows: how can we select the order of an AR to estimate well the distribution of a time series? It is known that order selection by minimization of criteria like AICC, AIC or FPE satisfy asymptotically oracle inequalities (Shibata [32] and Hurvich and Tsai [20]). We refer to Brockwell and Davis [9] and McQuarrie and Tsai [26] for detailed discussions. However, one cannot readily extend these results to a spatial setting because of computational and theoretical difficulties.

In the rest of this introduction, we further describe the framework and we summarize the main results of the paper.
1.1. Conditional regression. Let us now make precise the notation and present the ideas underlying our approach. In the sequel, $\Lambda$ stands for the toroidal lattice of size $p \times p$. We consider the random field $X=\left(X_{[i, j]}\right)_{1 \leq i, j \leq p}$ indexed by the nodes of $\Lambda$. Additionally, $X^{v}$ refers to the vectorialized version of $X$ with the convention $X_{[i, j]}=X_{[(i-1) \times p+j]}^{v}$ for any $1 \leq i, j \leq p$. Using this new notation amounts to "forgetting" the spatial structure of $X$ and allows one to get into a more classical statistical framework. For the sake of simplicity, the components of $X$ are defined modulo $p$ in the remainder of the paper.

Throughout this paper, we assume the field $X$ is centered. In practice, the statistician has to first subtract some parametric form of the mean value. Hence, the vector $X^{v}$ follows a zero-mean Gaussian distribution $\mathcal{N}(0, \Sigma)$ where the $p^{2} \times p^{2}$ matrix $\Sigma$ is nonsingular but unknown. Also, we suppose that the field $X$ is stationary on the torus $\Lambda$. More precisely, for any $r>0$, any $(i, j) \in\{1, \ldots, p\}^{2}$ and any $\left(k_{1}, l_{1}\right), \ldots,\left(k_{r}, l_{r}\right) \in\{1, \ldots, p\}^{2 r}$, it holds that

$$
\left(X_{\left[k_{1}, l_{1}\right]}, \ldots, X_{\left[k_{r}, l_{r}\right]}\right) \sim\left(X_{\left[k_{1}+i, l_{1}+j\right]}, \ldots, X_{\left[k_{r}+i, l_{r}+j\right]}\right)
$$

We observe $n \geq 1$ i.i.d. replications of the vector $X^{v}$. In the sequel, $\mathbf{X}^{\mathbf{v}}$ denotes the $p^{2} \times n$ matrix of the $n$ observations of $X^{v}$. For any $1 \leq i \leq n$, the $p \times p$ matrix $\mathbf{X}_{i}$ stands for the $i$ th observation of the field $X$. All these notations are recalled in Table 1. In practice, the number of observations $n$ often equals one. Our goal is to estimate the matrix $\Sigma$.

We sometimes assume that the field $X$ is isotropic. Let $G$ be the group of vector isometries of the unit square. For any node $(i, j) \in \Lambda$ and any isometry $g \in G$, $g .(i, j)$ stands for the image of $(i, j)$ in $\Lambda$ under the action of $g$. We say that $X$ is isotropic on $\Lambda$ if for any $r>0, g \in G$, and $\left(k_{1}, l_{1}\right), \ldots,\left(k_{r}, l_{r}\right) \in\{1, \ldots, p\}^{2 r}$,

$$
\left(X_{\left[k_{1}, l_{1}\right]}, \ldots, X_{\left[k_{r}, l_{r}\right]}\right) \sim\left(X_{\left[g .\left(k_{1}, l_{1}\right)\right]}, \ldots, X_{\left[g .\left(k_{r}, l_{r}\right)\right]}\right)
$$

As mentioned earlier, we aim at estimating the distribution of the field $X$ through a conditional distribution approach. By standard Gaussian derivations (see, for instance, [24], Appendix C), there exists a unique $p \times p$ matrix $\theta$ such that $\theta_{[0,0]}=0$ and

$$
\begin{equation*}
X_{[0,0]}=\sum_{(i, j) \in \Lambda \backslash\{(0,0)\}} \theta_{[i, j]} X_{[i, j]}+\varepsilon_{[0,0]}, \tag{1}
\end{equation*}
$$

where the random variable $\varepsilon_{[0,0]}$ follows a zero-mean normal distribution and is independent from the covariates $\left(X_{[i, j]}\right)_{(i, j) \in \Lambda \backslash\{(0,0)\}}$. Equation (1) describes the conditional distribution of $X_{[0,0]}$ given the remaining variables. Since the field $X$ is stationary, the matrix $\theta$ also satisfies $\theta_{[i, j]}=\theta_{[-i,-j]}$ for any $(i, j) \in \Lambda$. Let us note $\sigma^{2}$, the conditional variance of $X_{[0,0]}$, and $I_{p^{2}}$, the identity matrix of size $p^{2}$. The matrix $\theta$ is closely related to the covariance matrix $\Sigma$ of $X^{v}$ through the following property:

$$
\begin{equation*}
\Sigma=\sigma^{2}\left(I_{p^{2}}-C(\theta)\right)^{-1} \tag{2}
\end{equation*}
$$

where the $p^{2} \times p^{2}$ matrix $C(\theta)$ is defined as $C(\theta)_{\left[i_{1}(p-1)+j_{1}, i_{2}(p-1)+j_{2}\right]}:=$ $\theta_{\left[i_{2}-i_{1}, j_{2}-j_{1}\right]}$ for any $1 \leq i_{1}, i_{2}, j_{1}, j_{2} \leq p$. The matrix $\left(I_{p^{2}}-C(\theta)\right)$ is called the partial correlation matrix of the field $X$. The so-defined matrix $C(\theta)$ is symmetric block circulant with $p \times p$ blocks as stated below. We refer to [29], Section 2.6 or the book of Gray [15] for definitions and main properties on circulant and block circulant matrices.

Lemma 1.1. Let $\theta$ be a square matrix of size $p$ such that

$$
\begin{equation*}
\text { for any } 1 \leq i, j \leq p \quad \theta_{[i, j]}=\theta_{[-i,-j]} ; \tag{3}
\end{equation*}
$$

then the matrix $C(\theta)$ is symmetric block circulant with $p \times p$ blocks. Conversely, if $B$ is a $p^{2} \times p^{2}$ symmetric block circulant matrix with $p \times p$ blocks, then there exists a square matrix $\theta$ of size $p$ satisfying (3) and such that $B=C(\theta)$.

A proof is given in the technical Appendix [36]. In conclusion, estimating the matrix $\Sigma / \sigma^{2}$ amounts to estimating the matrix $C(\theta)$ which is also equivalent to estimating the $p \times p$ matrix $\theta$. This is why we shall focus on the estimation of the matrix $\theta$.

Let us make precise the set of possible values for $\theta$. In the sequel, $\Theta$ denote the vector space of the $p \times p$ matrices that satisfy $\theta_{[0,0]}=0$ and $\theta_{[i, j]}=\theta_{[-i,-j]}$ for any $(i, j) \in \Lambda$. A matrix $\theta \in \Theta$ corresponds to the distribution of a stationary Gaussian field if and only if the $p^{2} \times p^{2}$ matrix $\left(I_{p^{2}}-C(\theta)\right)$ is positive definite. This is why we define the convex subset $\Theta^{+}$of $\Theta$ by

$$
\begin{equation*}
\Theta^{+}:=\left\{\theta \in \Theta \text { s.t. }\left(I_{p^{2}}-C(\theta)\right) \text { is positive definite }\right\} . \tag{4}
\end{equation*}
$$

The set of covariance matrices of stationary Gaussian fields on $\Lambda$ with unit conditional variance is therefore in one to one correspondence with the set $\Theta^{+}$. Let us define the corresponding set $\Theta^{\text {iso }}$ and $\Theta^{+, \text {iso }}$ for isotropic Gaussian fields:

$$
\begin{align*}
\Theta^{\text {iso }} & :=\left\{\theta \in \Theta, \theta_{[i, j]}=\theta_{[g .(i, j)]}, \forall(i, j) \in \Lambda, \forall g \in G\right\} \quad \text { and } \\
\Theta^{+, \text {iso }} & :=\Theta^{+} \cap \Theta^{\text {iso }} . \tag{5}
\end{align*}
$$

1.2. Model selection. We have the issue of covariance estimation as an estimation problem for conditional regressions (1). However, the set $\Theta^{+}$of admissible parameters for the estimation is huge. The dimension of $\Theta$ is indeed of the same order as $p^{2}$ whereas we only observe $p^{2}$ nonindependent data if $n$ equals one. In order to avoid the curse of dimensionality, it is natural to assume that the target $\theta$ is approximately sparse.

It is indeed likely that the coefficients $\theta_{[i, j]}$ are close to zero for the nodes $(i, j)$ which are far from the origin $(0,0)$. By (1), this means that $X_{[0,0]}$ is well predicted by the covariates $X_{[i, j]}$ whose corresponding nodes $(i, j)$ are close to the origin. In other terms, the true covariance is presumably well approximated by a GMRF with a reasonable neighborhood. The main difficulty is that we do not know a priori what "reasonable" means. We want to adapt to the sparsity of the matrix $\theta$.

In the sequel, $m$ refers to a subset of $\Lambda \backslash\{0,0\}$. We call it a model. By (1), the property, " $X$ is a GMRF with respect to the neighborhood $m$," is equivalent to, "the support of $\theta$ is included in $m$." We are given a nested collection $\mathcal{M}$ of models. For any of these models $m \in \mathcal{M}$, we compute $\widehat{\theta}_{m, \rho_{1}}$, the conditional least squares estimator (CLS), of $\theta$ for the model $m$ by maximizing the pseudolikelihood over a subset of matrices $\theta$ whose support is included in $m$. These estimators, as well as their dependency on the quantity $\rho_{1}$, are defined in Section 2.

The model $m$ that minimizes the risk of $\widehat{\theta}_{m, \rho_{1}}$ over the collection $\mathcal{M}$ is called an oracle and is noted $m^{*}$. In practice, this model is unknown and we have to estimate it. The art of model selection is to pick a model $m \in \mathcal{M}$ that is large enough to enable a good approximation of $\theta$ but is small enough so that the variance of $\widehat{\theta}_{m, \rho_{1}}$ is small. Let us reformulate the approach in terms of GMRFs: given a collection $\mathcal{M}$ of neighborhoods, we compute an estimator of $\theta$ in the set of GMRFs with neighborhood $m$ for any $m \in \mathcal{M}$. Our purpose is to select a suitable neighborhood $\widehat{m}$ so that the estimator $\widehat{\theta}_{\widehat{m}}$ has a risk as small as possible.

A classical method to estimate a good model $\widehat{m}$ is achieved through penalization with respect to the size of the models. In the following expression, $\gamma_{n, p}(\cdot)$ stands for the CLS empirical contrast that we shall define in Section 2. We select a model $\widehat{m}$ by minimizing the criterion,

$$
\begin{equation*}
\widehat{m}=\underset{m \in \mathcal{M}}{\arg \min }\left[\gamma_{n, p}\left(\widehat{\theta}_{m, \rho_{1}}\right)+\operatorname{pen}(m)\right], \tag{6}
\end{equation*}
$$

where pen $(\cdot)$ denotes a positive function defined on $\mathcal{M}$. In this paper, we prove that under a suitable choice of the penalty function pen $(\cdot)$, the risk of the estimator $\widehat{\theta}_{\widehat{m}}$ is as small as possible.
1.3. Risk bounds and adaptation. We shall assess our procedure using two different loss functions. First, we introduce the loss function $l(\cdot, \cdot)$ that measures how well we estimate the conditional distribution (1) of the field. For any $\theta_{1}, \theta_{2} \in \Theta$, the distance $l\left(\theta_{1}, \theta_{2}\right)$ is defined by

$$
\begin{equation*}
l\left(\theta_{1}, \theta_{2}\right):=\frac{1}{p^{2}} \operatorname{tr}\left[\left(C\left(\theta_{1}\right)-C\left(\theta_{2}\right)\right) \Sigma\left(C\left(\theta_{1}\right)-C\left(\theta_{2}\right)\right)\right] \tag{7}
\end{equation*}
$$

Let us reformulate $l\left(\theta_{1}, \theta_{2}\right)$ in terms of conditional expectation,

$$
l\left(\theta_{1}, \theta_{2}\right)=\mathbb{E}_{\theta}\left\{\left[\mathbb{E}_{\theta_{1}}\left(X_{[0,0]} \mid X_{\Lambda \backslash\{0,0\}}\right)-\mathbb{E}_{\theta_{2}}\left(X_{[0,0]} \mid X_{\Lambda \backslash\{0,0\}}\right)\right]^{2}\right\}
$$

where $\mathbb{E}_{\theta}(\cdot)$ stands for the expectation with respect to the distribution of $X^{v}$, $\mathcal{N}\left(0, \sigma^{2}\left(I_{p^{2}}-C(\theta)\right)^{-1}\right)$. Hence, $l(\widehat{\theta}, \theta)$ corresponds the mean squared prediction loss which is often used in the random design regression framework, in time series analysis [20] or in spatial statistics [33]. Moreover, the loss function $l(\widehat{\theta}, \theta)$ is also connected to the notion of kriging error. The kriging predictor (Stein [34]) of $X_{[0,0]}$ is defined as the best linear combination of the covariates $\left(X_{[k, l]}\right)_{(k, l) \in \Lambda \backslash\{(0,0\}}$ for predicting the value $X_{[0,0]}$. By (1), this predictor is exactly $\sum_{(k, l) \in \Lambda \backslash\{(0,0\}} \theta_{[k, l]} X_{[k, l]}$, and the mean squared prediction error is $\sigma^{2}$. If we do not know $\theta$ but we are given an estimator $\widehat{\theta}$, then the corresponding kriging predictor $\sum_{(k, l) \in \Lambda \backslash\{(0,0\}} \widehat{\theta}_{[k, l]} X_{[k, l]}$ has a mean squared prediction error equal to $\sigma^{2}+l(\widehat{\theta}, \theta)$. Kriging is a key concept in spatial statistics, and it is therefore interesting to consider a loss function that measures the kriging performances when one estimates $\theta$.

We shall also assess our results using the Frobenius distance noted $\|\cdot\|_{F}$ and defined by $\|A\|_{F}^{2}:=\sum_{1 \leq i, j \leq p} A_{[i, j]}^{2}$. Observe that the Frobenius distance $\| \theta_{1}-$ $\theta_{2} \|_{F}^{2}$ also equals the Frobenius distance between the partial correlation matrices $\left(I_{p^{2}}-C\left(\theta_{1}\right)\right)$ and $\left(I_{p^{2}}-C\left(\theta_{2}\right)\right)$ (up to a factor $p^{2}$ )

$$
\begin{equation*}
\left\|\theta_{1}-\theta_{2}\right\|_{F}^{2}=\frac{1}{p^{2}}\left\|\left(I_{p^{2}}-C\left(\theta_{1}\right)\right)-\left(I_{p^{2}}-C\left(\theta_{2}\right)\right)\right\|_{F}^{2} \tag{8}
\end{equation*}
$$

Our aim is then to define a suitable penalty function pen( $(\cdot)$ in (6) so that the estimator $\widehat{\theta}_{\widehat{m}, \rho_{1}}$ performs almost as well as the oracle estimator $\widehat{\theta}_{m^{*}, \rho_{1}}$. For any model $m \in \mathcal{M}$, we define $\theta_{m, \rho_{1}}$ as the matrix which minimizes the loss $l\left(\theta^{\prime}, \theta\right)$ over the sets of matrices $\theta^{\prime}$ corresponding to model $m$. The $\operatorname{loss} l\left(\theta_{m, \rho_{1}}, \theta\right)$ is called the bias. Our main result is stated in Section 3. We provide a condition on the penalty function pen $(\cdot)$, so that the selected estimator satisfies a risk bound of the form

$$
\begin{equation*}
\mathbb{E}_{\theta}\left[l\left(\widehat{\theta}_{\widehat{m}, \rho_{1}}, \theta\right)\right] \leq L \inf _{m \in \mathcal{M}}\left[l\left(\theta_{m, \rho_{1}}, \theta\right)+\varphi_{\max }(\Sigma) \frac{\operatorname{Card}(m)}{n p^{2}}\right] \tag{9}
\end{equation*}
$$

where $\varphi_{\max }(\Sigma)$ is the largest eigenvalue of $\Sigma$, and $\operatorname{Card}(\cdot)$ stands for the cardinality. Contrary to most results in a spatial setting, this upper bound on the risk is nonasymptotic and holds in a general setting. The term $\varphi_{\max }(\Sigma) \operatorname{Card}(m) /\left(n p^{2}\right)$
grows linearly with the size of $m$ and goes to 0 with $n$ and $p$. In Section 4 , we prove that the variance term of a model $m$ is of the same order as $\varphi_{\max }(\Sigma) \operatorname{Card}(m) /\left(n p^{2}\right)$. Hence, the bound (9) tells us that the risk of $\widehat{\theta}_{\widehat{m}, \rho_{1}}$ is smaller than a quantity which is the same order as the risk $\mathbb{E}_{\theta}\left[l\left(\widehat{\theta}_{m^{*}, \rho_{1}}, \theta\right)\right]$ of the oracle $m^{*}$. We say that the selected estimator achieves an oracle-type inequality.

In Section 4 , we bound the asymptotic expectations $\mathbb{E}\left[l\left(\widehat{\theta}_{m, \rho_{1}}, \theta\right)\right]$ and connect them to the variance terms in bound (9). As a consequence, we prove that under mild assumptions on the target $\theta$, the upper bound (9) is optimal from the asymptotic point of view (up to a multiplicative numerical constant). We discuss the assumptions in Section 5. In Section 6, we compute nonasymptotic minimax lower bounds with respect to the loss functions $l(\cdot, \cdot)$ and $\|\cdot\|_{F}^{2}$. We then derive that under mild assumptions, our estimator $\widehat{\theta}_{\widehat{m}, \rho_{1}}$ is minimax adaptive to the sparsity of $\theta$ and minimax adaptive to the decay of $\theta$.

To our knowledge, these are the first oracle-type inequalities in a spatial setting. The computation of the minimax rates of convergence is also new. Moreover, most of our results are nonasymptotic. Although we have considered a square on the two-dimensional lattice, our method straightforwardly extends to any $d$-dimensional toroidal rectangle with $d \geq 1$. In the one-dimensional setting, we retrieve a oracle-type inequality that is close to the work of Shibata [32]. Yet, he has stated an asymptotic oracle inequality for the estimation of autoregressive processes. In contrast, our result applies on a torus and is only optimal up to constants but it is nonasympotic, and, most of all, it applies for higher-dimensional lattices. In Section 7, we further discuss the advantages and the weak points of our method. Moreover, we mention the extensions and the simulations made in a subsequent paper [37]. All the proofs are postponed to Section 8 and to the Appendix [36].
1.4. Some notation. Throughout this paper, $L, L_{1}, L_{2}, \ldots$ denote constants that may vary from line to line. The notation $L(\cdot)$ specifies the dependency on some quantities. For any matrix $A, \varphi_{\max }(A)$ and $\varphi_{\min }(A)$, respectively, refer the largest eigenvalue and the smallest eigenvalues of $A$. We recall that $\|A\|_{F}$ is the Frobenius norm of $A$. For any matrix $\theta$ of size $p,\|\theta\|_{1}$ stands for the sum of of the absolute values of the components of $\theta$; we call it its $l_{1}$ norm. In the sequel, $0_{p}$ is the square matrix of size $p$ whose indices are 0 . Given $\rho>0$, the ball $\mathcal{B}_{1}\left(0_{p} ; \rho\right)$ is defined as the set of square matrices of size $p$ whose $l_{1}$ norm is smaller than $\rho$. Finally, Table 1 gathers the notation involving $X$.

Table 1
Notations for the random field and the data

| $X$ | Matrix of size $p \times p$ | Random field |
| :--- | :---: | :---: |
| $X^{v}$ | Vector of length $p^{2}$ | Vectorialized version of $X$ |
| $\mathbf{X}^{\mathbf{v}}$ | Matrix of size $p^{2} \times n$ | Observations of $X^{v}$ |
| $\mathbf{X}_{i}$ | Matrix of size $p \times p$ | $i$ th observation of the field $X$ |

2. Model selection procedure. In this section, we formally define our model selection procedure.
2.1. Collection of models. For any node $(i, j)$ belonging to the lattice $\Lambda$, let us define the toroidal norm by

$$
|(i, j)|_{t}^{2}:=[i \wedge(p-i)]^{2}+[j \wedge(p-j)]^{2}
$$

We aim at selecting a "good" neighborhood for the GMRF. Since $X$ corresponds to some "spatial" process, it is natural to assume that nodes that are close to $(0,0)$ are more likely to be significant. This is why we restrict ourselves in the sequel to the collection $\mathcal{M}_{1}$ of neighborhoods.

Definition 2.1. A subset $m \subset \Lambda \backslash\{(0,0)\}$ belongs to $\mathcal{M}_{1}$ if there exists a number $r_{m}>1$ such that

$$
\begin{equation*}
m=\left\{(i, j) \in \Lambda \backslash\{(0,0)\} \text { s.t. }|(i, j)|_{t} \leq r_{m}\right\} . \tag{10}
\end{equation*}
$$

The collection $\mathcal{M}_{1}$ is totally ordered with respect to the inclusion and we therefore order our models $m_{0} \subset m_{1} \subset \cdots \subset m_{i} \cdots$. For instance, $m_{0}$ corresponds to the empty neighborhood whereas $m_{1}$ stands for the neighborhood of size 4. See Figure 1 for other examples.

For any model $m \in \mathcal{M}_{1}$, we define the vector space $\Theta_{m}$ as the subset of the elements of $\Theta$ whose support is included in $m$. We recall that $\Theta$ is defined in Section 1.1. Similarly $\Theta_{m}^{\text {iso }}$ is the subset of $\Theta^{\text {iso }}$ whose support is included in $m$. The dimensions of $\Theta_{m}$ and $\Theta_{m}^{\text {iso }}$ are, respectively, noted $d_{m}$ and $d_{m}^{\text {iso }}$. Since we


Fig. 1. Examples of models. The four gray nodes refer to $m_{1}$. The model $m_{2}$ also contains the nodes with a cross whereas $m_{3}$ contains all the nodes except $(0,0)$.
aim at estimating the positive matrix $\left(I_{p^{2}}-C(\theta)\right)$, we shall consider the convex subsets of $\Theta_{m}^{+}$and $\Theta_{m}^{+, \text {iso }}$ that correspond to nonnegative precision matrices,

$$
\begin{equation*}
\Theta_{m}^{+}:=\Theta_{m} \cap \Theta^{+} \quad \text { and } \quad \Theta_{m}^{+, \text {iso }}:=\Theta_{m}^{\text {iso }} \cap \Theta^{+, \text {iso }} \tag{11}
\end{equation*}
$$

For instance, the set $\Theta_{m_{1}}^{+}$is in one-to-one correspondence with the sets of GMRFs whose neighborhood is made of the four nearest neighbors. Similarly, $\Theta_{m_{1}}^{+}$is in one-to-one correspondence with the GMRFs with eight nearest neighbors. In our estimation procedure, we shall restrict ourselves to precision matrices whose largest eigenvalue is upper bounded by a constant. This is why we define the subsets $\Theta_{m_{2}, \rho_{1}}^{+}$and $\Theta_{m, \rho_{1}}^{+, \text {iso }}$ for any $\rho_{1} \geq 2$ :

$$
\begin{align*}
& \Theta_{m, \rho_{1}}^{+}:=\left\{\theta \in \Theta_{m}^{+}, \varphi_{\max }\left(I_{p^{2}}-C(\theta)\right)<\rho_{1}\right\},  \tag{12}\\
& \Theta_{m, \rho_{1}}^{+, \text {iso }}:=\left\{\theta \in \Theta_{m}^{+ \text {iso }}, \varphi_{\max }\left(I_{p^{2}}-C(\theta)\right)<\rho_{1}\right\} . \tag{13}
\end{align*}
$$

Finally, we need a generating family of the spaces $\Theta_{m}$ and $\Theta_{m}^{\text {iso }}$. For any node $(i, j) \in \Lambda \backslash\{(0,0)\}$, let us define the $p \times p$ matrix $\Psi_{i, j}$ as

$$
\Psi_{i, j[k, l]}:= \begin{cases}1, & \text { if }(k, l)=(i, j) \text { or }(k, l)=-(i, j),  \tag{14}\\ 0, & \text { otherwise }\end{cases}
$$

Hence, $\Theta_{m}$ is generated by the matrices $\Psi_{i, j}$ for which $(i, j)$ belongs to $m$. Similarly, for any $(i, j) \in \Lambda \backslash\{(0,0)\}$, let us define the matrix $\Psi_{i, j}^{\text {iso }}$ by

$$
\Psi_{i, j[k, l]}^{\text {iso }}:= \begin{cases}1, & \text { if } \exists g \in G,(k, l)=g .(i, j)  \tag{15}\\ 0, & \text { otherwise } .\end{cases}
$$

2.2. Estimation by conditional least squares ( $C L S$ ). Let us turn to the conditional least squares estimator. For any $\theta^{\prime} \in \Theta^{+}$, the criterion $\gamma_{n, p}\left(\theta^{\prime}\right)$ is defined by

$$
\begin{align*}
\gamma_{n, p}\left(\theta^{\prime}\right):=\frac{1}{n p^{2}} \sum_{i=1}^{n} \sum_{1 \leq j_{1}, j_{2} \leq p}( & \mathbf{X}_{i\left[j_{1}, j_{2}\right]}  \tag{16}\\
& \left.-\sum_{\left(l_{1}, l_{2}\right) \in \Lambda \backslash\{(0,0)\}} \theta_{\left[l_{1}, l_{2}\right]}^{\prime} \mathbf{X}_{i\left[j_{1}+l_{1}, j_{2}+l_{2}\right]}\right)^{2} .
\end{align*}
$$

In a nutshell, $\gamma_{n, p}\left(\theta^{\prime}\right)$ is a least squares criterion that allows one to perform the simultaneous linear regression of all $\mathbf{X}_{i\left[j_{1}, j_{2}\right]}$ with respect to the covariates $\left(\mathbf{X}_{i\left[l_{1}, l_{2}\right]}\right)_{\left(l_{1}, l_{2}\right) \neq\left(j_{1}, j_{2}\right)}$. The advantage of this criterion is that it does not require the computation of a determinant of a huge matrix as for the likelihood. We shall often use an alternative expression of $\gamma_{n, p}\left(\theta^{\prime}\right)$ in terms of the factor $C\left(\theta^{\prime}\right)$ and the empirical covariance matrix $\overline{\mathbf{X}^{\mathbf{v}} \mathbf{X}^{\mathbf{v} *}}$,

$$
\begin{equation*}
\gamma_{n, p}\left(\theta^{\prime}\right)=\frac{1}{p^{2}} \operatorname{tr}\left[\left(I_{p^{2}}-C\left(\theta^{\prime}\right)\right) \overline{\mathbf{X}^{\mathbf{v}} \mathbf{X}^{\mathbf{v} *}}\left(I_{p^{2}}-C\left(\theta^{\prime}\right)\right)\right] \tag{17}
\end{equation*}
$$

One proves the equivalence between these two expressions by coming back to the definition of $C\left(\theta^{\prime}\right)$. Let $\rho_{1}>2$ be fixed. For any model $m \in \mathcal{M}$, we compute the CLS estimators $\widehat{\theta}_{m, \rho_{1}}$ and $\widehat{\theta}_{m, \rho_{1}}^{\text {iso }}$ by minimizing the criterion $\gamma_{n, p}(\cdot)$ as follows:

$$
\begin{equation*}
\widehat{\theta}_{m, \rho_{1}}:=\underset{\theta^{\prime} \in \overline{\Theta_{m, \rho_{1}}^{+}}}{\arg \min } \gamma_{n, p}\left(\theta^{\prime}\right) \quad \text { and } \quad \widehat{\theta}_{m, \rho_{1}}^{\text {iso }}:=\underset{\theta^{\prime} \in \overline{\Theta_{m, \rho_{1}}^{+, \text {iso }}}}{\arg \min } \gamma_{n, p}\left(\theta^{\prime}\right) \tag{18}
\end{equation*}
$$

where $\bar{A}$ stands for the closure of the set $A$. The existence and the uniqueness of $\widehat{\theta}_{m, \rho_{1}}$ and $\widehat{\theta}_{m, \rho_{1}}^{\text {iso }}$ are ensured by the following lemma.

LEMMA 2.2. For any $\theta \in \Theta^{+}, \gamma_{n, p}(\cdot)$ is almost surely strictly convex on $\overline{\Theta^{+}}$.
The proof is postponed to the Appendix [36]. We discuss the dependency of $\widehat{\theta}_{m, \rho_{1}}$ on the parameter $\rho_{1}$ in Section 5. For stationary Gaussian fields, minimizing the CLS criterion $\gamma_{n, p}(\cdot)$ over a set $\Theta_{m, \rho_{1}}^{+}$is equivalent to minimizing the product of the conditional likelihoods ( $X_{[i, j]} \mid X_{-\{i, j\}}$ ), called conditional pseudolikelihood (CPL),

$$
\begin{aligned}
p \mathcal{L}_{n}\left(\theta^{\prime}, \mathbf{X}^{\mathbf{v}}\right): & \left.=\prod_{\substack{1 \leq i \leq n,\left(j_{1}, j_{2}\right) \in \Lambda}} \mathcal{L}_{n, \theta^{\prime}}\left(\mathbf{X}_{i\left[j_{1}, j_{2}\right]}\right]\left(\mathbf{X}_{i}\right)_{-\left\{j_{1}, j_{2}\right\}}\right) \\
& =(\sqrt{2 \pi} \sigma)^{-n p^{2}} \exp \left(-\frac{1}{2} \frac{n p^{2} \gamma_{n, p}\left(\theta^{\prime}\right)}{\sigma^{2}}\right),
\end{aligned}
$$

where we recall that $\sigma^{2}$ refers to the conditional variance of any $X_{[i, j]}$. In fact, CLS estimators were first introduced by Besag [2] who call them pseudolikelihood estimators since they minimize the CPL.

Let us define the function $\gamma(\cdot)$ as an infinite sampled version of the CLS criterion $\gamma_{n, p}(\cdot)$,

$$
\begin{equation*}
\gamma\left(\theta^{\prime}\right):=\mathbb{E}_{\theta}\left[\gamma_{n, p}\left(\theta^{\prime}\right)\right]=\mathbb{E}_{\theta}\left[\left(X_{[0,0]}-\sum_{(i, j) \neq(0,0)} \theta_{[i, j]}^{\prime} X_{[i, j]}\right)^{2}\right] \tag{19}
\end{equation*}
$$

for any $\theta^{\prime}, \theta \in \Theta^{+}$. The function $\gamma\left(\theta^{\prime}\right)$ measures the prediction error of $X_{[0,0]}$ if one uses $\sum_{(i, j) \neq(0,0)} \theta_{[i, j]}^{\prime} X_{[i, j]}$ as a predictor. Moreover, it is a special case of the CMLS criterion introduced by Cressie and Verzelen (see [11], (10)) to approximate a Gaussian field by a GMRF. Hence, one may interpret the CLS criterion as a finite sampled version of their approximation method. Observe that the function $\gamma(\cdot)$ is minimized over $\Theta^{+}$at the point $\theta$ and that $\gamma(\theta)=\operatorname{Var}_{\theta}\left(X_{[0,0]} \mid X_{-\{0,0\}}\right)=\sigma^{2}$. Moreover, the difference $\gamma\left(\theta^{\prime}\right)-\gamma(\theta)$ equals the loss $l\left(\theta^{\prime}, \theta\right)$ defined by (7).

For any model $m \in \mathcal{M}$, we introduce the projections $\theta_{m, \rho_{1}}$ and $\theta_{m, \rho_{1}}^{\text {iso }}$ as the best approximation of $\theta$ in $\overline{\Theta_{m, \rho_{1}}^{+}}$and $\overline{\Theta_{m, \rho_{1}}^{+ \text {,iso }} \text { : }}$

$$
\begin{equation*}
\theta_{m, \rho_{1}}:=\underset{\theta^{\prime} \in \overline{\Theta_{m, \rho_{1}}^{+}}}{\arg \min } l\left(\theta^{\prime}, \theta\right) \quad \text { and } \quad \theta_{m, \rho_{1}}^{\text {iso }}:=\underset{\theta^{\prime} \in \overline{\Theta_{m, \rho_{1}}^{+, \text {iso }}}}{\arg \min } l\left(\theta^{\prime}, \theta\right) \tag{20}
\end{equation*}
$$

Since $\gamma(\cdot)$ is strictly convex on $\Theta^{+}$, the matrices $\theta_{m, \rho_{1}}$ and $\theta_{m, \rho_{1}}^{\text {iso }}$ are uniquely defined. By its definition (7), one may interpret $l(\cdot, \cdot)$ as an inner product on the space $\Theta$; therefore, the orthogonal projection of $\theta$ onto the convex closed set $\overline{\Theta_{m, \rho_{1}}^{+}}$ (resp., $\overline{\Theta_{m, \rho_{1}}^{+, \text {iso }}}$ ) with respect to $l(\cdot, \cdot)$ is $\theta_{m, \rho_{1}}$ (resp., $\theta_{m, \rho_{1}}^{\text {iso }}$ ). It then follows from a property of orthogonal projections that the loss of $\widehat{\theta}_{m, \rho_{1}}$ is upper bounded by

$$
\begin{equation*}
l\left(\widehat{\theta}_{m, \rho_{1}}, \theta\right) \leq l\left(\theta_{m, \rho_{1}}, \theta\right)+l\left(\widehat{\theta}_{m, \rho_{1}}, \theta_{m, \rho_{1}}\right) \tag{21}
\end{equation*}
$$

The first term $l\left(\theta_{m, \rho_{1}}, \theta\right)$ accounts for the bias whereas the second term $l\left(\widehat{\theta}_{m, \rho_{1}}\right.$, $\left.\theta_{m, \rho_{1}}\right)$ is a variance term. Observe that $\theta \in \Theta_{m}^{+}$does not necessarily imply that the $\operatorname{bias} l\left(\theta_{m, \rho_{1}}, \theta\right)$ is null because in general $\overline{\Theta_{m}^{+}} \neq \overline{\Theta_{m, \rho_{1}}^{+}}$. This will be the case only if $\theta$ satisfies the following hypothesis:

$$
\begin{equation*}
\left(\mathbb{H}_{1}\right): \quad \varphi_{\max }\left(I_{p^{2}}-C(\theta)\right)<\rho_{1} \tag{22}
\end{equation*}
$$

Assumption $\left(\mathbb{H}_{1}\right)$ is necessary to ensure the existence of a model $m \in \mathcal{M}$ such that the bias is zero (i.e., $\theta_{m, \rho_{1}}=\theta$ ). By identity (2), one observes that $\left(\mathbb{H}_{1}\right)$ is equivalent to a lower bound on the smallest eigenvalue of $\Sigma$, i.e., $\varphi_{\min }(\Sigma) \leq$ $\sigma^{2} / \rho_{1}$. We further discuss $\left(\mathbb{H}_{1}\right)$ in Section 5.

For the sake of completeness, we recall the penalization criterion introduced in (6). Given a subcollection of models $\mathcal{M} \subset \mathcal{M}_{1}$ and a positive function pen: $\mathcal{M} \rightarrow \mathbb{R}^{+}$that we call a penalty, we select a model as follows:

$$
\widehat{m}:=\underset{m \in \mathcal{M}}{\arg \min }\left[\gamma_{n, p}\left(\widehat{\theta}_{m, \rho_{1}}\right)\right]+\operatorname{pen}(m)
$$

and

$$
\widehat{m}^{\text {iso }}:=\underset{m \in \mathcal{M}}{\arg \min }\left[\gamma_{n, p}\left(\widehat{\theta}_{m, \rho_{1}}^{\mathrm{iso}}\right)\right]+\operatorname{pen}(m) .
$$

Observe that $\widehat{m}$ and $\widehat{m}^{\text {iso }}$ depend on $\rho_{1}$. For the sake clarity, we do not emphasize this dependency in the notation. In the sequel, we write $\widetilde{\theta}_{\rho_{1}}$ and $\widetilde{\theta}_{\rho_{1}}^{\text {iso }}$ for $\widehat{\theta}_{\widehat{m}, \rho_{1}}$ and $\widehat{\theta}_{\widehat{m}^{\text {iso }}}^{\text {iso }} \rho_{1}$.
3. Main result. We now provide a nonasymptotic upper bound for the risk of the estimators $\widetilde{\theta}_{\rho_{1}}$ and $\widetilde{\theta}_{\rho_{1}}^{\text {iso }}$. Let us recall that $\Sigma$ stands for the covariance matrix of $X^{v}$.

THEOREM 3.1. Let $K$ be a positive number larger than a universal constant $K_{0}$ and let $\mathcal{M}$ be a subcollection of $\mathcal{M}_{1}$. If for every model $m \in \mathcal{M}$,

$$
\begin{equation*}
\operatorname{pen}(m) \geq K \rho_{1}^{2} \varphi_{\max }(\Sigma) \frac{d_{m}}{n p^{2}} \tag{23}
\end{equation*}
$$

then for any $\theta \in \Theta^{+}$, the estimator $\tilde{\theta}_{\rho_{1}}$ satisfies

$$
\begin{equation*}
\mathbb{E}_{\theta}\left[l\left(\tilde{\theta}_{\rho_{1}}, \theta\right)\right] \leq L_{1}(K) \inf _{m \in \mathcal{M}}\left[l\left(\theta_{m, \rho_{1}}, \theta\right)+\operatorname{pen}(m)\right]+L_{2}(K) \frac{\rho_{1}^{2} \varphi_{\max }(\Sigma)}{n p^{2}} \tag{24}
\end{equation*}
$$

A similar bound holds if one replaces $\widetilde{\theta}_{\rho_{1}}$ by $\widetilde{\theta}_{\rho_{1}}^{\text {iso }}, \Theta^{+}$by $\Theta^{+, \text {iso }}, \theta_{m, \rho_{1}}$ by $\theta_{m}^{\text {iso }}$, and $d_{m}$ by $d_{m}^{\text {iso }}$.

The proof is postponed to Section 8.2. It is based on a novel concentration inequality for suprema of Gaussian chaos stated in Section 8.1. The constant $K_{0}$ is made explicit in the proof. Observe that the theorem holds for any $n$, any $p$ and that we have not performed any assumption on the target $\theta \in \Theta^{+}$(resp., $\Theta^{+ \text {,iso }}$ ). If the collection $\mathcal{M}$ does not contain the empty model, one gets the more readable upper bound,

$$
\mathbb{E}_{\theta}\left[l\left(\tilde{\theta}_{\rho_{1}}, \theta\right)\right] \leq L(K) \inf _{m \in \mathcal{M}}\left[l\left(\theta_{m, \rho_{1}}, \theta\right)+\operatorname{pen}(m)\right] .
$$

This theorem tells us that $\tilde{\theta}_{\rho_{1}}$ essentially performs as well as the best trade-off between the bias term $l\left(\theta_{m, \rho_{1}}, \theta\right)$ and $\rho_{1}^{2} \varphi_{\max }(\Sigma) \frac{d_{m}}{n p^{2}}$ that plays the role of a variance. Here are some additional comments.

REMARK 1. Consider the special case where the target $\theta$ belongs to some parametric set $\Theta_{m}^{+}$with $m \in \mathcal{M}$. Suppose that the hypothesis $\left(\mathbb{H}_{1}\right)$ defined in (22) is fulfilled. Choosing a penalty $\operatorname{pen}(m)=K \rho_{1}^{2} \varphi_{\max }(\Sigma) \frac{d_{m}}{n p^{2}}$, we get

$$
\begin{equation*}
\mathbb{E}_{\theta}\left[l\left(\tilde{\theta}_{\rho_{1}}, \theta\right)\right] \leq L(K) \rho_{1}^{2} \varphi_{\max }(\Sigma) \frac{d_{m}}{n p^{2}} \tag{25}
\end{equation*}
$$

We shall prove in Sections 4.2 and 6.1 that this rate is optimal both from an asymptotic oracle and a minimax point of view. We have mentioned in Section 2.2 that $\left(\mathbb{H}_{1}\right)$ is necessary for bound (25) to hold. If $\rho_{1}$ is chosen large enough, then assumption $\left(\mathbb{H}_{1}\right)$ is fulfilled. We do not have access to this minimal $\rho_{1}$ that ensures $\left(\mathbb{H}_{1}\right)$, since it requires the knowledge of $\theta$. Nevertheless, we argue in Section 5 that "moderate" values for $\rho_{1}$ ensure assumption $\left(\mathbb{H}_{1}\right)$ when the model $m$ is small.

REMARK 2. We have mentioned in the Introduction that our objective was to obtain oracle inequalities of the form

$$
\mathbb{E}_{\theta}\left[l\left(\widetilde{\theta}_{\rho_{1}}, \theta\right)\right] \leq L(K) \inf _{m \in \mathcal{M}} \mathbb{E}\left[l\left(\widehat{\theta}_{m, \rho_{1}}, \theta\right)\right]=L(K) \mathbb{E}\left[\left(\widehat{\theta}_{m^{*}, \rho_{1}}, \theta\right)\right]
$$

This is why we want to compare the sum $l\left(\theta_{m, \rho_{1}}, \theta\right)+\operatorname{pen}(m)$ with $\mathbb{E}\left[l\left(\widehat{\theta}_{m, \rho_{1}}, \theta\right)\right]$. First, we provide in Section 4.1 a sufficient condition so that the risk $\mathbb{E}\left[l\left(\widehat{\theta}_{m, \rho_{1}}, \theta\right)\right]$ decomposes exactly as the sum $l\left(\theta_{m, \rho_{1}}, \theta\right)+\mathbb{E}\left[l\left(\widehat{\theta}_{m, \rho_{1}}, \theta_{m, \rho_{1}}\right)\right]$. Moreover, we compute in Section 4.2 the asymptotic variance term $\mathbb{E}\left[l\left(\widehat{\theta}_{m, \rho_{1}}, \theta_{m, \rho_{1}}\right)\right]$ and compare it with the penalty term $\rho_{1}^{2} \varphi_{\max }(\Sigma) \frac{d_{m}}{n p^{2}}$. We shall then derive oracle-type inequalities and discuss the dependency of the different bounds on $\varphi_{\max }(\Sigma)$.

REMARK 3. Condition (23) gives a lower bound on the penalty function pen (•) so that the result holds. Choosing a proper penalty term according to (23) therefore requires an upper bound on the largest eigenvalue of $\Sigma$. However, such a bound is seldom known in practice. We shall mention in Section 7 a practical method to calibrate the penalty.

A bound similar to (24) holds for the Frobenius distance between the partial correlation matrices $\left(I_{p^{2}}-C(\theta)\right)$ and $\left(I_{p^{2}}-C\left(\widetilde{\theta}_{\rho_{1}}\right)\right)$.

Corollary 3.2. Assume the same as in Theorem 3.1, except that there is equality in (23). Then

$$
\begin{align*}
\mathbb{E}_{\theta}[\| C & \left.C\left(\tilde{\theta}_{\rho_{1}}\right)-C(\theta) \|_{F}^{2}\right] \\
\leq & L_{1}(K) \frac{\varphi_{\max }(\Sigma)}{\varphi_{\min }(\Sigma)} \inf _{m \in \mathcal{M}}\left[\left\|C\left(\theta_{m, \rho_{1}}\right)-C(\theta)\right\|_{F}^{2}+\frac{K \rho_{1}^{2} d_{m}}{n}\right]  \tag{26}\\
& +L_{2}(K) \frac{\varphi_{\max }(\Sigma)}{\varphi_{\min }(\Sigma)} \frac{\rho_{1}^{2}}{n}
\end{align*}
$$

A similar result holds for isotropic GMRFs.
Proof. This is a consequence of Theorem 3.1. By definition (7) of the loss function $l(\cdot, \cdot)$, the two following bounds hold:

$$
\begin{aligned}
& p^{2} l\left(\theta_{1}, \theta_{2}\right) \geq \varphi_{\min }(\Sigma)\left\|C\left(\theta_{1}\right)-C\left(\theta_{2}\right)\right\|_{F}^{2}, \\
& p^{2} l\left(\theta_{1}, \theta_{2}\right) \leq \varphi_{\max }(\Sigma)\left\|C\left(\theta_{1}\right)-C\left(\theta_{2}\right)\right\|_{F}^{2} .
\end{aligned}
$$

Gathering these bounds with (24) yields the result.
The same comments as for Theorem 3.1 hold. We may express this Corollary 3.2 in terms of the risk $\mathbb{E}\left(\left\|\tilde{\theta}_{\rho_{1}}-\theta\right\|_{F}^{2}\right)$, since $\left\|C\left(\theta_{1}\right)-C\left(\theta_{2}\right)\right\|_{F}^{2}=p^{2}\left\|\theta_{1}-\theta_{2}\right\|_{F}^{2}$ :

$$
\begin{aligned}
\mathbb{E}_{\theta}\left[\left\|\tilde{\theta}_{\rho_{1}}-\theta\right\|_{F}^{2}\right] \leq & L_{1}(K) \frac{\varphi_{\max }(\Sigma)}{\varphi_{\min }(\Sigma)} \inf _{m \in \mathcal{M}}\left[\left\|\theta_{m, \rho_{1}}-\theta\right\|_{F}^{2}+\frac{K \rho_{1}^{2} d_{m}}{n p^{2}}\right] \\
& +L_{2}(K) \frac{\varphi_{\max }(\Sigma)}{\varphi_{\min }(\Sigma)} \frac{\rho_{1}^{2}}{n p^{2}}
\end{aligned}
$$

4. Parametric risk and asymptotic oracle inequalities. In this section, we study the risk of the parametric estimators $\widehat{\theta}_{m, \rho_{1}}$ in order to assess the optimality of Theorem 3.1.
4.1. Bias-variance decomposition. The properties of the parametric estimator $\widehat{\theta}_{m, \rho_{1}}$ and of the projection $\theta_{m, \rho_{1}}$ differ slightly whether $\theta_{m, \rho_{1}}$ belongs to the open set $\Theta_{m, \rho_{1}}^{+}$or to its border. Observe that Hypothesis $\left(\mathbb{H}_{1}\right)$ defined in (22) does not
necessarily imply that projection $\theta_{m, \rho_{1}}$ belongs to $\Theta_{m}^{+}$. This is why we introduce condition $\left(\mathbb{H}_{2}\right)$ :

$$
\begin{equation*}
\theta \in \mathcal{B}_{1}\left(0_{p}, 1\right) \quad \Longleftrightarrow \quad\|\theta\|_{1}<1 \tag{27}
\end{equation*}
$$

The condition $\|\theta\|_{1}<1$ is equivalent to $\left(I_{p^{2}}-C(\theta)\right.$ ) is strictly diagonally dominant. Condition $\left(\mathbb{H}_{2}\right)$ implies that the largest eigenvalue of $\left(I_{p^{2}}-C(\theta)\right)$ is smaller than 2 and therefore that $\left(\mathbb{H}_{1}\right)$ is fulfilled since $\rho_{1}$ is supposed larger than 2 . We further discuss this assumption in Section 5.

Lemma 4.1. Let $\theta \in \Theta^{+}$such that $\left(\mathbb{H}_{2}\right)$ holds and let $m \in \mathcal{M}_{1}$. Then, the minimum of $\gamma(\cdot)$ over $\Theta_{m}$ is achieved in $\Theta_{m, 2}^{+}$. This implies that

$$
\theta_{m, \rho_{1}}=\underset{\theta^{\prime} \in \Theta_{m}}{\arg \min } \gamma\left(\theta^{\prime}\right) \quad \text { and } \quad \gamma\left(\theta_{m, \rho_{1}}\right)=\operatorname{Var}_{\theta}\left(X_{[0,0]} \mid X_{m}\right)
$$

Additionally, $\left\|\theta_{m, \rho_{1}}\right\|_{1} \leq\|\theta\|_{1}$. The same results holds for $\theta_{m, \rho_{1}}^{\text {iso }}$ if $\theta$ in $\Theta^{+, \text {iso }}$.
The proof is given in the technical Appendix [36]. The purpose of this property is threefold. First, we derive that assumption $\left(\mathbb{H}_{2}\right)$ ensures that $\theta_{m, \rho_{1}}$ belongs $\Theta_{m, \rho_{1}}^{+}$and that the smallest eigenvalue of $\left(I_{p^{2}}-C\left(\theta_{m, \rho_{1}}\right)\right)$ is larger than $1-\|\theta\|_{1}$. Second, it allows to express the projection $\theta_{m, \rho_{1}}$ in terms of conditional expectation (Corollary 4.2). Finally, we deduce a bias-variance decomposition of the estimator $\widehat{\theta}_{m, \rho_{1}}$ (Corollary 4.3). In other words, the equality holds in (21).

Corollary 4.2. Let $\theta \in \Theta^{+}$such that $\left(\mathbb{H}_{2}\right)$ holds and let $m \in \mathcal{M}_{1}$. The projection $\theta_{m, \rho_{1}}$ is uniquely defined by the equation

$$
\mathbb{E}_{\theta}\left(X_{[0,0]} \mid X_{m}\right)=\sum_{(i, j) \in m} \theta_{m, \rho_{1}[i, j]} X_{[i, j]}
$$

and $\theta_{m, \rho_{1}[i, j]}=0$ for any $(i, j) \notin m$. Similarly, if $\theta \in \Theta^{+, \text {iso }}$ satisfies $\left(\mathbb{H}_{2}\right)$, then $\theta_{m, \rho_{1}}^{\text {iso }}$ is uniquely defined by the equation

$$
\mathbb{E}_{\theta}\left(X_{[0,0]} \mid X_{m}\right)=\sum_{(i, j) \in m} \theta_{m, \rho_{1}[i, j]}^{\text {iso }} X_{[i, j]}
$$

and $\theta_{m, \rho_{1}[i, j]}^{\mathrm{iso}}=0$ for any $(i, j) \notin m$.
Consequently, $\sum_{1 \leq i, j \leq p} \theta_{m, \rho_{1}[i, j]} X_{[i, j]}$ is the best linear predictor of $X_{[0,0]}$ given the covariates $\bar{X}_{[i, j]}$ with $(i, j) \in m$. This is precisely the definition of the kriging parameters (Stein [34]). Hence, the matrix $\theta_{m, \rho_{1}}$ corresponds to the kriging parameters of $X_{[0,0]}$ with kriging neighborhood's range of $r_{m}$. The distance $r_{m}$ is introduced in Definition 2.1 and stands for the radius of $m$.

Corollary 4.3. Let $\theta \in \Theta^{+}$such that $\left(\mathbb{H}_{2}\right)$ holds and let $m \in \mathcal{M}_{1}$. The loss of $\widehat{\theta}_{m, \rho_{1}}$ decomposes as $l\left(\widehat{\theta}_{m, \rho_{1}}, \theta\right)=l\left(\theta_{m, \rho_{1}}, \theta\right)+l\left(\widehat{\theta}_{m, \rho_{1}}, \theta_{m, \rho_{1}}\right)$. If $\theta$ belongs to $\Theta_{m}^{+, \text {iso }}$ and $\left(\mathbb{H}_{2}\right)$ holds, then we also have the decomposition $l\left(\widehat{\theta}_{m, \rho_{1}}^{\text {iso }}, \theta\right)=$ $l\left(\theta_{m, \rho_{1}}^{\text {iso }}, \theta\right)+l\left(\widehat{\theta}_{m, \rho_{1}}^{\text {iso }}, \theta_{m, \rho_{1}}\right)$.

A proof is provided in the technical Appendix [36]. If $\theta$ does not satisfy assumption $\left(\mathbb{H}_{2}\right)$, then $\theta_{m, \rho_{1}}$ does not necessarily belong to $\Theta_{m, \rho_{1}}^{+}$, and there may not be such a bias variance decomposition.
4.2. Asymptotic risk. In this section, we evaluate the risk of each estimator $\widehat{\theta}_{m, \rho_{1}}$ and use it as a benchmark to assess the result of Theorem 3.1. We have mentioned in Corollary 4.3 that under $\left(\mathbb{H}_{2}\right)$ the risk $\mathbb{E}_{\theta}\left[l\left(\widehat{\theta}_{m, \rho_{1}}, \theta\right)\right]$ decomposes into the sum of the bias $l\left(\theta_{m, \rho_{1}}, \theta\right)$ and a variance term $\mathbb{E}_{\theta}\left[l\left(\widehat{\theta}_{m, \rho_{1}}, \theta_{m, \rho_{1}}\right)\right]$. If this last quantity is of the same order as the penalty pen $(m)$ introduced in (23), then Theorem 3.1 yields an oracle inequality. However, we are unable to express this variance term $\mathbb{E}_{\theta}\left[l\left(\widehat{\theta}_{m, \rho_{1}}, \theta_{m, \rho_{1}}\right)\right]$ in a simple form. This is why we restrict ourselves to study the risks when $n$ tends to infinity. Nevertheless, these results give us some hints to appreciate the strength and the weaknesses of Theorem 3.1 and the upper bound (25).

In the following proposition, we adapt a result of Guyon [17], Section 4.3.2 to obtain an asymptotic expression of the risk $\mathbb{E}_{\theta}\left[l\left(\widehat{\theta}_{m, \rho_{1}}, \theta_{m, \rho_{1}}\right)\right]$. We first need to introduce some new notation. For any model $m$ in the collection $\mathcal{M}_{1} \backslash\{\varnothing\}$, we fix a sequence $\left(i_{k}, j_{k}\right)_{k=1, \ldots, d_{m}}$ of integers such that $\left(\Psi_{i_{1}, j_{1}}, \ldots, \Psi_{i_{d_{m}}, j_{d_{m}}}\right)$ is a basis of the space $\Theta_{m}$. Then $\chi_{m[0,0]}$ stands for the random vector of size $d_{m}$ that contains the neighbors of $X_{[0,0]}$

$$
\chi_{m[0,0]}^{*}:=\left[\operatorname{tr}\left(\Psi_{i_{1}, j_{1}} X^{v}\right), \ldots, \operatorname{tr}\left(\Psi_{i_{d_{m}}, j_{d_{m}}} X^{v}\right)\right] .
$$

Additionally, for any $\theta \in \Theta^{+}$, we define the matrices $V, W$ and $I L_{m}$ as

$$
\left\{\begin{array}{l}
V:=\operatorname{cov}_{\theta}\left(\chi_{m}[0,0]\right) \\
W_{[k, l]}:=\frac{1}{p^{2}} \operatorname{tr}\left[C\left(\Psi_{i_{k}, j_{k}}\right)\left(I_{p^{2}}-C\left(\theta_{m, \rho_{1}}\right)\right)^{2}\left(I_{p^{2}}-C(\theta)\right)^{-2} C\left(\Psi_{i_{l}, j_{l}}\right)\right] \\
\quad \text { for any } k=1, \ldots, d_{m}, \\
I L_{m}:=\operatorname{Diag}\left(\left\|\Psi_{i_{k}, j_{k}}\right\|_{F}^{2}, k=1, \ldots, d_{m}\right)
\end{array}\right.
$$

where for any vector $u, \operatorname{Diag}(u)$ is the diagonal matrix whose diagonal elements are the components of $u$. We also define the corresponding quantities $\chi_{m[0,0]}^{\text {iso }}, V^{\text {iso }}$, $W^{\text {iso }}$ and $I L_{m}^{\text {iso }}$ in order to consider the isotropic estimator $\widehat{\theta}_{m, \rho_{1}}^{\text {iso }}$.

Proposition 4.4. Let $m$ be a model in $\mathcal{M}_{1} \backslash\{\varnothing\}$, and let $\theta$ be an element of $\Theta_{m}^{+}$that satisfies $\left(\mathbb{H}_{1}\right)$. Then $\widehat{\theta}_{m, \rho_{1}}$ converges to $\theta$ in probability, and

$$
\begin{equation*}
\lim _{n \rightarrow+\infty} n p^{2} \mathbb{E}_{\theta}\left[l\left(\widehat{\theta}_{m, \rho_{1}}, \theta\right)\right]=2 \sigma^{4} \operatorname{tr}\left[I L_{m} V^{-1}\right] \tag{28}
\end{equation*}
$$

Let $\theta$ in $\Theta^{+}$such that $\left(\mathbb{H}_{2}\right)$ is fulfilled. Then, $\widehat{\theta}_{m, \rho_{1}}$ converges to $\theta_{m, \rho_{1}}$ in probability and

$$
\begin{equation*}
\lim _{n \rightarrow+\infty} n p^{2} \mathbb{E}_{\theta}\left[l\left(\widehat{\theta}_{m, \rho_{1}}, \theta_{m, \rho_{1}}\right)\right]=2 \sigma^{4} \operatorname{tr}\left(W V^{-1}\right) \tag{29}
\end{equation*}
$$

Both results still hold for the estimator $\widehat{\theta}_{m, \rho_{1}}^{\text {iso }}$ if $\theta$ belongs to $\Theta^{+, \text {iso }}$ and if one replaces $V, W$, and $I L_{m}$ by $V^{\text {iso }}, W^{\text {iso }}$ and $I L_{m}^{\text {iso }}$.

In the first case, assumption $\left(\mathbb{H}_{1}\right)$ ensures that $\theta \in \Theta_{m, \rho_{1}}^{+}$whereas assumption $\left(\mathbb{H}_{2}\right)$ ensures that $\theta_{m, \rho_{1}} \in \Theta_{m, \rho_{1}}^{+}$. The proof is based on the extension of Guyon's approach in the toroidal framework.

Expressions (28) and (29) are not easily interpretable in the present form. This is why we first derive (28) when $\theta$ is zero. Observe that it is equivalent to the independence of $\left(X_{[i, j]}\right)_{(i, j) \in \Lambda}$.

EXAMPLE 4.5. Assume that $\theta$ is zero. Then for any model $m \in \mathcal{M}_{1}$, the asymptotic risks of $\widehat{\theta}_{m, \rho_{1}}$ and $\widehat{\theta}_{m, \rho_{1}}^{\text {iso }}$ satisfy

$$
\lim _{n \rightarrow+\infty} n p^{2} \mathbb{E}_{0_{p}}\left[l\left(\widehat{\theta}_{m, \rho_{1}}, 0_{p}\right)\right]=2 \sigma^{2} d_{m}
$$

and

$$
\lim _{n \rightarrow+\infty} n p^{2} \mathbb{E}_{0_{p}}\left[l\left(\widehat{\theta}_{m, \rho_{1}}^{\text {iso }}, 0_{p}\right)\right]=2 \sigma^{2} d_{m}^{\text {iso }},
$$

where we recall that $d_{m}^{\mathrm{iso}}$ is the dimension of the space $\Theta_{m}^{\mathrm{iso}}$.
Proof. Since the components of $X$ are independent, the matrix $V$ equals $\sigma^{2} I L_{m}$. We conclude by applying Proposition 4.4.

Therefore, when the variables $X_{[i, j]}$ are independent, the asymptotic risk of $\widehat{\theta}_{m, \rho_{1}}$ equals, up to a factor 2 , the variance term of the least squares estimator in the fixed design Gaussian regression framework. This quantity is of the same order as the penalty introduced in Section 3. When the matrix $\theta$ is nonzero, we can lower bound the limits (28) and (29).

COROLLARY 4.6. Let $m$ be a model in $\mathcal{M}_{1}$ and let $\theta \in \Theta_{m}^{+}$that satisfies $\left(\mathbb{H}_{1}\right)$. Then, the variance term is asymptotically lower bounded as follows:

$$
\begin{equation*}
\lim _{n \rightarrow+\infty} n p^{2} \mathbb{E}_{\theta}\left[l\left(\widehat{\theta}_{m, \rho_{1}}, \theta\right)\right] \geq L \sigma^{2} \varphi_{\min }\left[I_{p^{2}}-C(\theta)\right] d_{m}=L \sigma^{4} \frac{d_{m}}{\varphi_{\max }(\Sigma)} \tag{30}
\end{equation*}
$$

where $L$ is a universal constant. Let $\theta \in \Theta^{+}$that satisfies $\left(\mathbb{H}_{2}\right)$. For any model $m \in \mathcal{M}_{1}$,

$$
\begin{equation*}
\lim _{n \rightarrow+\infty} n p^{2} \mathbb{E}_{\theta}\left[l\left(\widehat{\theta}_{m, \rho_{1}}, \theta_{m, \rho_{1}}\right)\right] \geq L \sigma^{2}\left(1-\|\theta\|_{1}\right)^{3} d_{m} \tag{31}
\end{equation*}
$$

The proof is postponed to the technical Appendix [36]. Again, analogous lower bounds hold for $\widehat{\theta}_{m, \rho_{1}}^{\text {iso }}$ when $\theta$ belongs to $\Theta^{\text {iso, }+}$. This corollary states that asymptotically with respect to $n$ the variance term of $\widehat{\theta}_{m, \rho_{1}}$ is larger than the order $d_{m} /\left(n p^{2}\right)$. This expression is not really surprising since $d_{m}$ stands for the dimension of the model $m$ and $n p^{2}$ corresponds to the number of data observed. Let us define $R_{\theta, \infty}\left(\widehat{\theta}_{m, \rho_{1}}, \theta_{m, \rho_{1}}\right):=\lim _{n \rightarrow+\infty} n p^{2} \mathbb{E}_{\theta}\left[l\left(\widehat{\theta}_{m, \rho_{1}}, \theta_{m, \rho_{1}}\right)\right]$ as the asymptotic variance term for $\widehat{\theta}_{m, \rho_{1}}$ rescaled by the number $n p^{2}$ of observations.

The first part of corollary (30) states that from an asymptotic point of view the upper bound (25) is optimal. By Theorem 3.1, if we choose pen $(m)=$ $K \rho_{1}^{2} \varphi_{\max }(\Sigma) \frac{d_{m}}{n p^{2}}$, then it holds that

$$
\mathbb{E}\left[l\left(\tilde{\theta}_{\rho_{1}}, \theta\right)\right] \leq L\left(K, \rho_{1}, \varphi_{\min }\left[I_{p^{2}}-C(\theta)\right]\right) \frac{R_{\theta, \infty}\left(\widehat{\theta}_{m, \rho_{1}}, \theta\right)}{n p^{2}}
$$

for any model $m \in \mathcal{M} \backslash \varnothing$ and any $\theta \in \Theta_{m}^{+}$that satisfies $\left(\mathbb{H}_{1}\right)$. This property holds for any $n$ and any $p$. Hence, $\widetilde{\theta}_{\rho_{1}}$ performs as well as the parametric estimator $\widehat{\theta}_{m, \rho_{1}}$ if the support of $\theta$ belongs to some unknown model $m$ and if $\theta$ satisfies $\left(\mathbb{H}_{1}\right)$.

If we assume that $\|\theta\|_{1}<1$ [hypothesis $\left(\mathbb{H}_{2}\right)$ ], we are able to derive a stronger result.

Proposition 4.7. Considering ${ }_{\sim}^{K} \geq K_{0}, \rho_{1} \geq 2, \eta<1$ and a collection $\mathcal{M} \subset \mathcal{M}_{1} \backslash \varnothing$, we define the estimator $\tilde{\theta}_{\rho_{1}}$ with the penalty $\operatorname{pen}(m)=K \rho_{1}^{2} \frac{d_{m}}{n p^{2}(1-\eta)}$. Then the risk of $\widetilde{\theta}_{\rho_{1}}$ is upper bounded by

$$
\begin{equation*}
\mathbb{E}_{\theta}\left[l\left(\tilde{\theta}_{\rho_{1}}, \theta\right)\right] \leq L\left(K, \rho_{1}, \eta\right) \inf _{m \in \mathcal{M}}\left\{l\left(\theta_{m, \rho_{1}}, \theta\right)+\frac{R_{\theta, \infty}\left(\widehat{\theta}_{m, \rho_{1}}, \theta_{m, \rho_{1}}\right)}{n p^{2}}\right\} \tag{32}
\end{equation*}
$$

for any $\theta \in \Theta^{+} \cap \mathcal{B}_{1}\left(0_{p}, \eta\right)$.
Observe that this property holds for any $n$ and any $p$. If the matrix $\theta$ is strictly diagonally dominant, we therefore obtain an upper bound similar to an oracle inequality, except that the variance term $\mathbb{E}_{\theta}\left[l\left(\widehat{\theta}_{m, \rho_{1}}, \theta_{m, \rho_{1}}\right)\right]$ has been replaced by its asymptotic counterpart $R_{\theta, \infty}\left(\widehat{\theta}_{m, \rho_{1}}, \theta_{m, \rho_{1}}\right) /\left(n p^{2}\right)$. However, this inequality is not valid uniformly over any $\eta<1$ : when $\eta$ converges to one, the constant $L\left(K, \rho_{1}, \eta\right)$ tends to infinity. Indeed, if $\|\theta\|_{1}$ converges to one, the lower bound (31) on the variance term can behave like $\left(1-\|\theta\|_{1}\right)^{3} d_{m} /\left(n p^{2}\right)$ for some matrices $\theta$ whereas the penalty term $d_{m} /\left[n p^{2}\left(1-\|\theta\|_{1}\right)\right]$ tends to infinity.

In the remaining part of the section, we illustrate that the constant $L\left(K, \eta, \rho_{1}\right)$ has to go to infinity when $\eta$ goes to one. Let us consider the model $m_{1}$. It consists of GMRFs with 4-nearest neighbors.

EXAMPLE 4.8. Let $\theta$ be a nonzero element of $\Theta_{m_{1}}^{\mathrm{iso}}$; then the asymptotic risk of $\widehat{\hat{\theta}_{m_{1}, \rho_{1}}^{\text {iso }}}$ simplifies as

$$
\begin{equation*}
\lim _{n \rightarrow+\infty} n p^{2} \mathbb{E}_{\theta}\left[l\left(\widehat{\theta}_{m_{1}, \rho_{1}}^{\text {iso }}, \theta\right)\right]=2 \frac{\sigma^{4} \theta_{[1,0]}}{\operatorname{cov}\left(X_{[1,0]}, X_{[0,0]}\right)} \tag{33}
\end{equation*}
$$

If we let size $p$ of the network tend to infinity and $\theta_{[1,0]}$ go to $1 / 4$, the risk is equivalent to

$$
\lim _{p \rightarrow+\infty} \lim _{n \rightarrow+\infty} n p^{2} \mathbb{E}_{\theta}\left[l\left(\widehat{\theta}_{m_{1}, \rho_{1}}^{\text {iso }}, \theta\right)\right] \underset{\theta_{[1,0]} \rightarrow 1 / 4}{\sim} \frac{16 \sigma^{2}\left(1-4 \theta_{[1,0]}\right)}{\log (16)} .
$$

The proof is postponed to the technical Appendix [36]. If follows from the second result that the lower bound (30) is sharp since in this particular case $\varphi_{\min }\left(I_{p^{2}}-C(\theta)\right)=\sigma^{2}\left(1-4 \theta_{[1,0]}\right)$. When $\theta_{[1,0]}$ tends to $1 / 4$, then $\|\theta\|_{1}$ tends to one, and $\mathbb{E}_{\theta}\left[l\left(\widehat{\theta}_{m_{1}, \rho_{1}}^{\text {iso }}, \theta\right)\right]$ behaves like $\sigma^{2}\left(1-\|\theta\|_{1}\right) d_{m_{1}}^{\text {iso }} /\left(n p^{2}\right)$ whereas the penalty pen $\left(m_{1}\right)$ given in Theorem 3.1 has to be larger than $\sigma^{2} d_{m_{1}}^{\text {iso }} /\left[n p^{2}\left(1-\|\theta\|_{1}\right)\right]$. Hence, the variance term and the penalty pen $(\cdot)$ are not necessarily of the same order when $\|\theta\|_{1}$ tends to one. Theorem 3.1 cannot lead to an oracle inequality of the type (32) which is valid uniformly on $\eta<1$.

EXAMPLE 4.9. Let $\alpha$ be a positive number smaller than $1 / 4$. For any integer $p$ which is divisible by 4, we define the $p \times p$ matrix $\theta^{(p)}$ by

$$
\left\{\begin{array}{l}
\theta_{[p / 4, p / 4]}^{(p)}=\theta_{[-p / 4, p / 4]}^{(p)}=\theta_{[p / 4,-p / 4]}^{(p)}=\theta_{[-p / 4,-p / 4]}^{(p)}:=\alpha, \\
\theta_{[i, j]}^{(p)}:=0, \quad \text { else. }
\end{array}\right.
$$

Then the variance term is asymptotically lower bounded as follows:

$$
\lim _{p \rightarrow+\infty} \lim _{n \rightarrow+\infty} n p^{2} \mathbb{E}_{\theta^{(p)}}\left[l\left(\widehat{\theta}_{m_{1}, \rho_{1}}^{(p) \text { iso }},\left[\theta^{(p)}\right]_{m_{1}, \rho_{1}}^{\text {iso }}\right)\right] \geq \frac{L \sigma^{2}}{1-4 \alpha}
$$

The proof is postponed to the technical Appendix [36]. This variance term is of order $\sigma^{2} d_{m}^{\text {iso }} /\left[n p^{2}\left(1-\|\theta\|_{1}\right)\right]=\varphi_{\max }(\Sigma) d_{m}^{\text {iso }} /\left(n p^{2}\right)$ when $\|\theta\|_{1}$ goes to one. The penalty pen $(m)$ introduced in Proposition 4.7 is therefore a sharp upper bound of the variance terms.

On one hand, we take a penalty pen $(m)$ larger than $\sigma^{2} d_{m} /\left(n p^{2}\left(1-\|\theta\|_{1}\right)\right)$. On the other hand, the variance of $\widehat{\theta}_{m, \rho_{1}}$ is of the order $\sigma^{2}\left(1-\|\theta\|_{1}\right) d_{m} /\left(n p^{2}\right)$ in some cases. Bound (32) cannot, therefore, hold uniformly over any $\eta<1$. We think that it is intrinsic to the penalization strategy.
5. Comments on the assumptions. In this section, we discuss the dependency of the estimators $\widehat{\theta}_{m, \rho_{1}}$ on $\rho_{1}$ as well as assumptions $\left(\mathbb{H}_{1}\right)$ and $\left(\mathbb{H}_{2}\right)$.

Dependency of $\widehat{\theta}_{m, \rho_{1}}$ on $\rho_{1}$. We recall that the estimator $\widehat{\theta}_{m, \rho_{1}}$ is defined in (18) as the minimizer of the CLS empirical contrast $\gamma_{n, p}(\cdot)$ over $\Theta_{m, \rho_{1}}^{+}$. It may seem restrictive to perform the minimization over the set $\Theta_{m, \rho_{1}}^{+}$instead of $\Theta_{m}^{+}$. Nevertheless, we advocate that it is not the case, at least for small models. Let us indeed define

$$
\rho(m):=\sup _{\theta \in \Theta_{m}^{+}} \varphi_{\max }\left[I_{p^{2}}-C(\theta)\right] \quad \text { and } \quad \rho^{\text {iso }}(m):=\sup _{\theta \in \Theta_{m}^{+ \text {iso }}} \varphi_{\max }\left[I_{p^{2}}-C(\theta)\right] .
$$

TABLE 2
Approximate computation of $\rho(m)$ and $\rho^{\text {iso }}(m)$ for the four smallest models with $p=50$

| $d_{m}$ | 2 | 4 | 6 | 10 |
| :--- | :--- | :--- | :--- | :---: |
| $\rho(m)$ | 2.0 | 4.0 | 5.0 | 6.8 |
| $d_{m}^{\text {iso }}$ | 1 | 2 | 3 | 4 |
| $\rho^{\text {iso }}(m)$ | 2.0 | 4.0 | 5.0 | 6.8 |

The quantities $\rho(m)$ and $\rho^{\text {iso }}(m)$ are finite since $\Theta_{m}^{+}$is bounded. If one takes $\rho_{1}$ larger than $\rho(m)$ [resp., $\rho^{\text {iso }}(m)$ ], then the set $\Theta_{m, \rho_{1}}^{+}$(resp., $\Theta_{m, \rho_{1}}^{+ \text {,iso }}$ ) is exactly $\Theta_{m}^{+}$ (resp., $\Theta_{m}^{+, \text {iso }}$ ). We illustrate in Table 2 that $\rho(m)$ and $\rho^{\text {iso }}(m)$ are small when model $m$ is small. Consequently, choosing a moderate value for $\rho_{1}$ is not really restrictive for small models. However, when the size of model $m$ increases, the sets $\Theta_{m, \rho_{1}}^{+}$and $\Theta_{m}^{+}$become different for moderate values of $\rho_{1}$. In Section 7, we discuss the choice of $\rho_{1}$.

Assumption $\left(\mathbb{H}_{1}\right)$ defined in (22) states that the largest eigenvalue of ( $I_{p^{2}}-$ $C(\theta)$ ) is smaller than $\rho_{1}$. We have illustrated in Table 2 that if the support of $\theta$ belongs to a small model $m$, then the maximal absolute value of $\left(I_{p^{2}}-C(\theta)\right)$ is small. Hence, assumption $\left(\mathbb{H}_{1}\right)$ is ensured for "moderate" values of $\rho_{1}$ as soon as the support of $\theta$ belongs to some small model. If $\theta$ is not sparse but approximately sparse it is likely that the largest eigenvalue of $\theta$ remain moderate. In practice, we do not know in advance if a given choice of $\rho_{1}$ ensures $\left(\mathbb{H}_{1}\right)$. In Section 7, we discuss an extension of our procedure which does not require assumption $\left(\mathbb{H}_{1}\right)$.

Assumption $\left(\mathbb{H}_{2}\right)$ defined in (27) states that $\theta \in \mathcal{B}_{1}\left(0_{p}, 1\right)$ or equivalently that the matrix $\left(I_{p^{2}}-C(\theta)\right)$ is diagonally dominant. Rue and Held prove in [29], Section 2.7, that $\Theta_{m_{1}}^{+}$is included in $\mathcal{B}_{1}\left(0_{p}, 1\right)$. They also point out that a small part of $\Theta_{m_{2}}^{+}$does not belong to $\mathcal{B}_{1}\left(0_{p}, 1\right)$. In fact, assumption $\left(\mathbb{H}_{2}\right)$ becomes more and more restrictive if the support of $\theta$ becomes larger. Nevertheless, assumption $\left(\mathbb{H}_{2}\right)$ is also quite common in the literature (as, for instance, in [17]).

If one looks closely at our proofs involving assumption $\left(\mathbb{H}_{2}\right)$, one realizes that this assumption is only made to ensure the following facts:

1. The projection $\theta_{m, \rho_{1}}$ belongs to the open set $\Theta_{m, \rho_{1}}^{+}$for any model $m \in \mathcal{M}$ (Corollary 4.3).
2. The smallest eigenvalue of $\left(I_{p^{2}}-C\left(\theta_{m, \rho_{1}}\right)\right)$ is lower bounded by some positive number $\rho_{2}$, uniformly over all models $m \in \mathcal{M}$.

From empirical observations, these two last facts seem far more restrictive than $\left(\mathbb{H}_{2}\right)$. We used assumption $\left(\mathbb{H}_{2}\right)$ in the statement of our results, because we did not find any weaker but still simple condition that ensures facts 1 and 2.
6. Minimax rates. In Theorem 3.1 and Proposition 4.7 we have shown that under mild assumptions on $\theta$ the estimator $\widetilde{\theta}_{\rho_{1}}$ behaves almost as well as the best estimator among the family $\left\{\widehat{\theta}_{m, \rho_{1}}, m \in \mathcal{M}\right\}$. We now compare the risk of $\widetilde{\theta}_{\rho_{1}}$ with the risk of any other possible estimator $\widehat{\theta}$. This includes comparison with maximum likelihood methods. There is no hope to make a pointwise comparison with an arbitrary estimator. Therefore, we classically consider the maximal risk over some suitable subsets $\mathcal{T}$ of $\Theta^{+}$. The minimax risk over the set $\mathcal{T}$ is given by $\inf _{\widehat{\theta}} \sup _{\theta \in \mathcal{T}} \mathbb{E}_{\theta}[l(\widehat{\theta}, \theta)]$ where the infimum is taken over all possible estimators $\widehat{\theta}$ of $\theta$. Then the estimator $\widetilde{\theta}_{\rho_{1}}$ is said to be approximately minimax with respect to the set $\mathcal{T}$ if the ratio,

$$
\frac{\sup _{\theta \in \mathcal{T}} \mathbb{E}_{\theta}\left[l\left(\widetilde{\theta}_{\rho_{1}}, \theta\right)\right]}{\inf _{\widehat{\theta}} \sup _{\theta \in \mathcal{T}} \mathbb{E}_{\theta}[l(\widehat{\theta}, \theta)]},
$$

is smaller than a constant that does not depend on $\sigma^{2}, n$ or $p$. An estimator is said to be adaptive to a collection $\left(\mathcal{T}_{i}\right)_{i \in \mathcal{I}}$ if it is simultaneously minimax over each $\mathcal{T}_{i}$. The problem of designing adaptive estimation procedures is in general difficult. It has been extensively studied in the fixed design Gaussian regression framework. See for instance [6] for a detailed discussion. In the sequel, we adapt some of their ideas to the GMRF framework.

We prove in Section 6.1 that the estimator $\widetilde{\theta}_{\rho_{1}}$ is adaptive to the unknown sparsity of the matrix $\theta$. Moreover, it is also adaptive if we consider the Frobenius distance between partial correlation matrices. In Section 6.2, we show that $\widetilde{\theta}_{\rho_{1}}$ is also adaptive to the rates of decay of the bias.

We need to restrain ourselves to set of matrices $\theta$ such that the largest eigenvalue of the covariance matrix $\Sigma$ is uniformly bounded. This is why we define

$$
\begin{equation*}
\forall \rho_{2}>1 \quad \mathcal{U}\left(\rho_{2}\right):=\left\{\theta \in \Theta, \varphi_{\min }\left(I_{p^{2}}-C(\theta)\right) \geq \frac{1}{\rho_{2}}\right\} \tag{34}
\end{equation*}
$$

Observe that $\theta \in \mathcal{U}\left(\rho_{2}\right)$ is exactly equivalent to $\varphi_{\max }(\Sigma) \leq \sigma^{2} \rho_{2}$ since $\Sigma=$ $\sigma^{2}\left(I_{p^{2}}-C(\theta)\right)$.
6.1. Adapting to unknown sparsity. In this subsection, we prove that under mild assumptions the penalized estimator $\widetilde{\theta}_{\rho_{1}}$ is adaptive to the unknown sparsity of $\theta$. We first lower bound the minimax rate of convergence on given hypercubes.

DEFINITION 6.1. Let $m$ be a model in the collection $\mathcal{M}_{1} \backslash \varnothing$. We consider $\left(\Psi_{i_{1}, j_{1}}, \ldots, \Psi_{i_{d_{m}}}, j_{d_{m}}\right)$ a basis of the space $\Theta_{m}$ defined by (14). For any $\theta^{\prime} \in \Theta_{m}^{+}$, the hypercube $\mathcal{C}_{m}\left(\theta^{\prime}, r\right)$ is defined as

$$
\mathcal{C}_{m}\left(\theta^{\prime}, r\right):=\left\{\theta^{\prime}+\sum_{k=1}^{d_{m}} \Psi_{i_{k}, j_{k}} \phi_{k}, \phi \in\{0,1\}^{d_{m}}\right\},
$$

if the positive number $r$ is small enough so that $\mathcal{C}_{m}\left(\theta^{\prime}, r\right) \subset \Theta^{+}$. For any $\theta^{\prime} \in \Theta_{m}^{+, \text {iso }}$, we analogously define the hypercubes $\mathcal{C}_{m}^{\text {iso }}\left(\theta^{\prime}, r\right)$ using a basis $\left(\Psi_{i_{1}, j_{1}}^{\text {iso }}, \ldots, \Psi_{i_{d_{m}}, j_{d_{m}}}^{\text {iso }}\right)$.

Proposition 6.2. Let $m$ be a model in $\mathcal{M}_{1} \backslash \varnothing$ whose dimension $d_{m}$ is smaller than $p \sqrt{n}$. Then, for any estimator $\widehat{\theta}$,

$$
\begin{equation*}
\sup _{\theta \in \Theta_{m}^{+}} \mathbb{E}_{\theta}[l(\widehat{\theta}, \theta)] \geq \sup _{\theta \in \Theta_{m, 2}^{+}} \mathbb{E}_{\theta}[l(\widehat{\theta}, \theta)] \geq L \sigma^{2} \frac{d_{m}}{n p^{2}} \tag{35}
\end{equation*}
$$

Let $\theta^{\prime}$ be an element of $\Theta_{m}^{+}$that satisfies $\left(\mathbb{H}_{2}\right)$. For any estimator $\widehat{\theta}$ of $\theta$,

$$
\begin{equation*}
\sup _{\theta \in \operatorname{Co}\left[\mathcal{C}_{m}\left(\theta^{\prime},\left(1-\left\|\theta^{\prime}\right\|_{1}\right) / \sqrt{n p^{2}}\right)\right]} \mathbb{E}_{\theta}[l(\widehat{\theta}, \theta)] \geq L \sigma^{2} \varphi_{\min }^{2}\left[I_{p^{2}}-C\left(\theta^{\prime}\right)\right] \frac{d_{m}}{n p^{2}}, \tag{36}
\end{equation*}
$$

where $\operatorname{Co}\left[\mathcal{C}_{m}\left(\theta^{\prime}, r\right)\right]$ denotes the convex hull of $\mathcal{C}_{m}\left(\theta^{\prime}, r\right)$.
An analogous result holds for isotropic hypercubes. The first bound (35) means that for any estimator $\widehat{\theta}$, the supremum of the risks $\mathbb{E}_{\theta}\left[l\left(\widehat{\theta}_{m, \rho_{1}}, \theta\right)\right]$ over $\Theta_{m}^{+}$is larger than $\sigma^{2} d_{m} /\left(n p^{2}\right)$ (up to some numerical constant). This rate $\sigma^{2} d_{m} /\left(n p^{2}\right)$ is achieved by the CLS estimator by Theorem 3.1.

The second lower bound (36) is of independent interest. It implies that in a small neighborhood of $\theta^{\prime}$ the risk $\mathbb{E}_{\theta}\left[l\left(\widehat{\theta}_{m, \rho_{1}}, \theta\right)\right]$ is larger than $\sigma^{2} \varphi_{\min }^{2}\left[I_{p^{2}}-\right.$ $\left.C\left(\theta^{\prime}\right)\right] d_{m} /\left(n p^{2}\right)$. This confirms the lower bound (30) of Corollary 4.6 in a nonasymptotic way. Indeed, these two expressions match up to a factor $\varphi_{\min }\left[I_{p^{2}}-\right.$ $\left.C\left(\theta^{\prime}\right)\right]$. This difference comes from the fact that the lower bound (36) holds for any estimator $\widehat{\theta}$. Bound (36) is sharp in the sense that the maximum likelihood estimator $\widehat{\theta}_{m_{1}}^{\text {iso,mle }}$ of isotropic GMRF in $m_{1}$ exhibits an asymptotic risk of order $\sigma^{2} \varphi_{\min }^{2}\left[I_{p^{2}}-C(\theta)\right] /\left(n p^{2}\right)$ for the parameter $\theta$ studied in Example 4.8. It is shown using the methodology introduced in the proof of Example 4.8. We now state that $\widetilde{\theta}_{\rho}$ is adaptive to the sparsity of $m$.

COROLLARY 6.3. Considering $K \geq K_{0}, \rho_{1} \geq 2, \rho_{2}>2$ and a collection $\mathcal{M} \subset \mathcal{M}_{1}$, we define the estimator $\widetilde{\theta}_{\rho_{1}}$ with the penalty $\operatorname{pen}(m)=K \sigma^{2} \rho_{1}^{2} \rho_{2} \frac{d_{m}}{n p^{2}}$. For any nonempty model $m$,

$$
\begin{equation*}
\sup _{\theta \in \Theta_{m, \rho_{1}}^{+} \cap \mathcal{U}\left(\rho_{2}\right)} \mathbb{E}_{\theta}\left[l\left(\tilde{\theta}_{\rho_{1}}, \theta\right)\right] \leq L\left(K, \rho_{1}, \rho_{2}\right) \inf _{\widehat{\theta}} \sup _{\theta \in \Theta_{m, \rho_{1}}^{+} \cap \mathcal{U}\left(\rho_{2}\right)} \mathbb{E}[l(\widehat{\theta}, \theta)] \tag{37}
\end{equation*}
$$

where $\mathcal{U}\left(\rho_{2}\right)$ is defined in (34).
A similar result holds for $\widetilde{\theta}_{\rho_{1}}^{\text {iso }}$ and $\Theta_{m, \rho_{1}}^{+ \text {,iso }}$. Corollary 6.3 is nonasymptotic and applies for any $n$ and any $p$. If $\theta$ belongs to some model $m$, then the optimal risk from a minimax point of view is of order $\frac{d_{m}}{n p^{2}}$. In practice, we do not know the
true model $m$. Nevertheless, the procedure simultaneously achieves the minimax rates for all supports $m$ possible. This means that $\widetilde{\theta}_{\rho_{1}}$ reaches this minimax rate $\frac{d_{m}}{n p^{2}}$ without knowing in advance the true model $m$.

The procedure is not adaptive to the smallest or the largest eigenvalue of ( $I_{p^{2}}-$ $C(\theta))$ which correspond to $\rho_{1}$ and $\rho_{2}$. Indeed, the constant $L\left(K, \rho_{1}, \rho_{2}\right)$ depends on $\rho_{1}$ and $\rho_{2}$. We are not aware of any other covariance estimation procedure which is really adaptive to the smallest or the largest eigenvalue of the matrix.

Finally, $\widetilde{\theta}_{\rho_{1}}$ exhibits the same adaptive properties with respect to the Frobenius norm.

Corollary 6.4. Under the same assumptions as Corollary 6.3,

$$
\begin{aligned}
& \sup _{\theta \in \Theta_{m, \rho_{1}}^{+} \cap \mathcal{U}\left(\rho_{2}\right)} \mathbb{E}_{\theta}\left[\left\|C\left(\widetilde{\theta}_{\rho_{1}}\right)-C(\theta)\right\|_{F}^{2}\right] \\
& \quad \leq L\left(K, \rho_{1}, \rho_{2}\right) \inf _{\hat{\theta}} \sup _{\theta \in \Theta_{m, \rho_{1}}^{+} \cap \mathcal{U}\left(\rho_{2}\right)} \mathbb{E}\left[\|C(\widehat{\theta})-C(\theta)\|_{F}^{2}\right] .
\end{aligned}
$$

Proof. As in the proof of Corollary 3.2, we observe that

$$
\left\|C\left(\theta_{1}\right)-C\left(\theta_{2}\right)\right\|_{F} \geq \frac{p^{2} \rho_{1}}{\sigma^{2}} l\left(\theta_{1}, \theta_{2}\right)
$$

if $\theta$ satisfies assumption $\left(\mathbb{H}_{1}\right)$. We conclude by applying Proposition 6.2 and Corollary 3.2.
6.2. Adapting to the decay of the bias. In this section, we prove that the estimator $\widetilde{\theta}_{\rho_{1}}$ is adaptive to a range of sets that we call pseudo-ellipsoids.

DEFINITION 6.5 (Pseudo-ellipsoids). Let $\left(a_{j}\right)_{1 \leq j \leq \operatorname{Card}\left(\mathcal{M}_{1}\right)}$ be a nonincreasing sequence of positive numbers. Then, $\theta \in \Theta^{+}$belongs to the pseudo-ellipsoid $\mathcal{E}(a)$ if and only if

$$
\begin{equation*}
\sum_{i=1}^{\operatorname{Card}\left(\mathcal{M}_{1}\right)} \frac{\operatorname{var}_{\theta}\left(X_{[0,0]} \mid X_{\mathcal{N}\left(m_{i-1}\right)}\right)-\operatorname{var}_{\theta}\left(X_{[0,0]} \mid X_{\mathcal{N}\left(m_{i}\right)}\right)}{a_{i}^{2}} \leq 1 \tag{38}
\end{equation*}
$$

Condition (38) measures how fast $\operatorname{Var}_{\theta}\left(X_{[0,0]} \mid X_{\mathcal{N}\left(m_{i}\right)}\right)$ tends to $\operatorname{Var}_{\theta}\left(X_{[0,0]} \mid\right.$ $\left.X_{\Lambda \backslash\{(0,0)\}}\right)$. Suppose that assumption $\left(\mathbb{H}_{2}\right)$ defined in (27) is fulfilled. By Corollary 4.2, $\operatorname{Var}_{\theta}\left(X_{[0,0]} \mid X_{\mathcal{N}\left(m_{i}\right)}\right)$ is the sum of $l\left(\theta_{m_{i}}, \theta\right)$ and $\sigma^{2}$, and condition (38) is equivalent to

$$
\begin{equation*}
\sum_{i=1}^{\operatorname{Card}\left(\mathcal{M}_{1}\right)} \frac{l\left(\theta_{m_{i-1}}, \theta\right)-l\left(\theta_{m_{i}}, \theta\right)}{a_{i}^{2}} \leq 1 \tag{39}
\end{equation*}
$$

Hence, the sequence $\left(a_{i}\right)$ gives some condition on the rate of decay of the bias when the dimension of the model increases. These sets $\mathcal{E}(a)$ are not true ellipsoids. Nevertheless, one may consider them as counterparts of the classical ellipsoids studied in the fixed design Gaussian regression framework (see, for instance, [25], Section 4.3).

To prove adaptivity, we shall need the equivalence between conditions (38) and (39). This equivalence holds if $\operatorname{Var}_{\theta}\left(X_{[0,0]} \mid X_{\mathcal{N}\left(m_{i}\right)}\right)$ decomposes as $l\left(\theta_{m_{i}}, \theta\right)+\sigma^{2}$ for any model $m \in \mathcal{M}_{1}$. As mentioned earlier, assumption $\left(\mathbb{H}_{2}\right)$ is sufficient (but not necessary) for this property to hold. This is why we restrict ourselves to study sets of the type $\mathcal{E}(a) \cap \mathcal{B}_{1}\left(0_{p}, 1\right)$. We shall also perform the following assumption on the ellipsoids $\mathcal{E}(a)$ :

$$
\left(\mathbb{H}_{a}\right): \quad a_{i}^{2} \leq \frac{\sigma^{2}}{d_{m_{i}}} \quad \text { for any } 1 \leq i \leq\left|\mathcal{M}_{1}\right| .
$$

It essentially means that the sequence $\left(a_{i}\right)$ converges fast enough toward 0 . For instance, all the sequences $a_{i}=\sigma\left(d_{m_{i}}\right)^{-s}$ with $s \geq 1 / 2$ satisfy $\left(\mathbb{H}_{a}\right)$.

PROPOSITION 6.6. Under assumption $\left(\mathbb{H}_{a}\right)$, the minimax rate of estimation on $\mathcal{E}(a) \cap \mathcal{B}_{1}\left(0_{p}, 1\right) \cap \mathcal{U}(2)$ is lower bounded by

$$
\begin{equation*}
\inf _{\hat{\theta}} \sup _{\theta \in \mathcal{E}(a) \cap \mathcal{B}_{1}\left(0_{p}, 1\right) \cap \mathcal{U}(2)} \mathbb{E}_{\theta}[l(\widehat{\theta}, \theta)] \geq L \sup _{1 \leq i \leq \operatorname{Card}\left(\mathcal{M}_{1}\right)}\left(a_{i}^{2} \wedge \sigma^{2} \frac{d_{m_{i}}}{n p^{2}}\right) . \tag{40}
\end{equation*}
$$

This lower bound is analogous to the minimax rate of estimation for ellipsoids in the Gaussian sequence model. Gathering Theorem 3.1 and Proposition 6.6 enables to derive adaptive properties for $\widetilde{\theta}_{\rho_{1}}$.

Proposition 6.7. Considering $K \geq K_{0}, \rho_{1} \geq 2, \rho_{2}>2$ and the collection $\mathcal{M}_{1}$, we define the estimator $\widetilde{\theta}_{\rho_{1}}$ with the penalty pen $(m)=K \sigma^{2} \rho_{1}^{2} \rho_{2} \frac{d_{m}}{n p^{2}}$. For any ellipsoid $\mathcal{E}(a)$ that satisfies $\left(\mathbb{H}_{a}\right)$ and such that $a_{1}^{2} \geq 1 /\left(n p^{2}\right)$, the estimator $\widetilde{\theta}_{\rho_{1}}$ is minimax over the set $\mathcal{E}(a) \cap \mathcal{B}_{1}\left(0_{p}, 1\right) \cap \mathcal{U}\left(\rho_{2}\right)$,

$$
\begin{align*}
& \sup _{\theta \in \mathcal{E}(a) \cap \mathcal{B}_{1}\left(0_{p}, 1\right) \cap \mathcal{U}\left(\rho_{2}\right)} \mathbb{E}_{\theta}\left[l\left(\tilde{\theta}_{\rho_{1}}, \theta\right)\right] \\
& \leq L\left(K, \rho_{1}, \rho_{2}\right) \inf _{\hat{\theta}} \sup _{\theta \in \mathcal{E}(a) \cap \mathcal{B}_{1}\left(0_{p}, 1\right) \cap \mathcal{U}\left(\rho_{2}\right)} \mathbb{E}_{\theta}[l(\widehat{\theta}, \theta)] . \tag{41}
\end{align*}
$$

Let us first illustrate this result. We have mentioned earlier that assumption $\left(\mathbb{H}_{a}\right)$ is satisfied for all sequences $a_{i}=\sigma\left(d_{m_{i}}\right)^{-s}$ with $s \geq 1 / 2$. We note that $\mathcal{E}^{\prime}(s)$ such a pseudo-ellipsoid. By Propositions 6.6 and 6.7, the minimax rate over one pseudo ellipsoid $\mathcal{E}^{\prime}(s)$ is $\sigma^{2}\left(n p^{2}\right)^{-2 s /(1+2 s)}$. The larger $s$ is, the faster the minimax rates is. The estimator $\widetilde{\theta}_{\rho_{1}}$ achieves simultaneously the rate $\sigma^{2}\left(n p^{2}\right)^{-2 s /(1+2 s)}$ for all
$s \geq 1 / 2$. Consequently, $\widetilde{\theta}_{\rho_{1}}$ is adaptive to the rate $s$ of decay of the bias: it achieves the optimal rates without knowing $s$ in advance.

Let us further comment on Proposition 6.7. By (41), the estimator $\tilde{\theta}_{\rho_{1}}$ is adaptive over $\mathcal{E}(a) \cap \mathcal{B}_{1}\left(0_{p}, 1\right) \cap \mathcal{U}\left(\rho_{2}\right)$ for all sequences $(a)$ such that $\left(\mathbb{H}_{a}\right)$ is satisfied and such that $a_{1}^{2} \geq 1 /\left(n p^{2}\right)$. Again, the result applies for any $n$ and any $p$. The condition $a_{1}^{2} \geq 1 /\left(n p^{2}\right)$ is classical. It ensures that the pseudo-ellipsoid $\mathcal{E}(a)$ is not degenerate, that is, that the minimax rates of estimation is not smaller than $\sigma^{2} /\left(n p^{2}\right)$. We explained earlier that we restrict ourselves to parameters $\theta$ in $\mathcal{B}_{1}\left(0_{p}, 1\right)$ only because this enforces the equivalence between (38) and (39). In contrast, the hypothesis $\varphi_{\max }(\Sigma) \leq \sigma^{2} \rho_{2}$ is really necessary because we fail to be adaptive to $\rho_{2}$.

COROLLARY 6.8. Under assumption $\left(\mathbb{H}_{a}\right)$, the minimax rate of estimation over $\mathcal{E}(a) \cap \mathcal{U}(2) \cap \mathcal{B}_{1}\left(0_{p}, 1\right)$ is lower bounded by

$$
\inf _{\hat{\theta}} \sup _{\theta \in \mathcal{E}(a) \cap \mathcal{B}_{1}\left(0_{p}, 1\right) \cap \mathcal{U}(2)} \mathbb{E}_{\theta}\left[\|C(\widehat{\theta})-C(\theta)\|_{F}^{2}\right] \geq L \sup _{1 \leq i \leq \operatorname{Card}\left(\mathcal{M}_{1}\right)}\left(a_{i}^{2} p^{2} \wedge \frac{d_{m_{i}}}{n}\right)
$$

Under the same assumptions as Proposition 6.7,

$$
\begin{aligned}
& \sup _{\theta \in \mathcal{E}(a) \cap \mathcal{B}_{1}\left(0_{p}, 1\right) \cap \mathcal{U}\left(\rho_{2}\right)} \mathbb{E}_{\theta}\left[\|C(\widehat{\theta})-C(\theta)\|_{F}^{2}\right] \\
& \quad \leq L\left(K, \rho_{1}, \rho_{2}\right) \inf _{\hat{\theta}} \sup _{\theta \in \mathcal{E}(a) \cap \mathcal{B}_{1}\left(0_{p}, 1\right) \cap \mathcal{U}\left(\rho_{2}\right)} \mathbb{E}_{\theta}\left[\|C(\widehat{\theta})-C(\theta)\|_{F}^{2}\right] .
\end{aligned}
$$

Proof. As in the proof of Corollary 3.2, we observe that

$$
\begin{aligned}
\left\|C\left(\theta_{1}\right)-C\left(\theta_{2}\right)\right\|_{F} & \geq p^{2}\left[\varphi_{\max }(\Sigma)\right]^{-1} l\left(\theta_{1}, \theta_{2}\right) \geq \frac{p^{2}}{\rho_{2} \sigma^{2}} l\left(\theta_{1}, \theta_{2}\right) \\
\left\|C\left(\theta_{1}\right)-C\left(\theta_{2}\right)\right\|_{F} & \leq p^{2}\left[\varphi_{\min }(\Sigma)\right]^{-1} l\left(\theta_{1}, \theta_{2}\right) \leq p^{2} \frac{\varphi_{\max }\left[I_{p^{2}}-C(\theta)\right]}{\sigma^{2}} l\left(\theta_{1}, \theta_{2}\right) \\
& \leq \frac{\rho_{2} p^{2}}{\sigma^{2}} l\left(\theta_{1}, \theta_{2}\right)
\end{aligned}
$$

if $\theta \in \mathcal{B}_{1}\left(0_{p}, 1\right) \cap \mathcal{B}_{\text {op }}\left(\rho_{2}\right)$. We conclude by applying Propositions 6.6 and 6.7.
Again, $\widetilde{\theta}_{\rho_{1}}$ satisfies the same minimax properties with respect to the Frobenius norm. All these properties easily extend to isotropic fields if one defines the corresponding sets $\mathcal{E}^{\text {iso }}(a) \cap \mathcal{B}_{1}\left(0_{p}, 1\right) \cap \mathcal{U}\left(\rho_{2}\right)$ of isotropic GMRFs.

## 7. Discussion.

7.1. Comparison with maximum likelihood estimation. Let us first compare the computational cost, the CLS estimation method and the maximum likelihood
estimator (MLE). For toroidal lattices, fast algorithms based on two-dimensional fast-Fourier transformation (see, for instance, [31]) allow to compute the MLE as fast as the CLS estimator. More details on the computation of the CLS estimators for toroidal lattices are given in [37], Section 2.3. When the lattice is not a torus, the MLE becomes intractable because it involves the optimization of a determinant of size $p^{2}$. In contrast, the CLS criterion $\gamma_{n, p}(\cdot)$ defined in (16) is a quadratic function of $\theta$. Consequently, CLS estimators are still computationally amenable. We extend our model selection to nontoroidal lattices in [37].

Let us compare the risk of CLS estimators and MLE. Given a small-dimensional model $m$, the risk of the parametric CLS estimator and the parametric MLE have been compared from an asymptotic point of view ([17], Section 4.3). It is generally accepted (see, for instance, Cressie [10], Section 7.3.1) and that parametric CLS estimators are almost as efficient as parametric MLE for the major part of the parameter spaces $\Theta_{m}^{+}$. We have nonasymptotically assessed this statement in Proposition 6.2 by minimax arguments. Nevertheless, for some parameters $\theta$ that are close to the border of $\Theta_{m}^{+}$, Kashyap and Chellappa [22] have pointed out that CLS estimators are less efficient than MLE. If we have proved nonasymptotic bounds for CLS-based model selection method, we are not aware of any such result for model selection procedures based on MLE.
7.2. Concluding remarks. We have developed a model selection procedure for choosing the neighborhood of a GMRF. In Theorem 3.1, we have proven a nonasymptotic upper bound for the risk of the estimator $\widetilde{\theta}_{\rho_{1}}$ with respect to the prediction error $l(\cdot, \cdot)$. Under assumption $\left(\mathbb{H}_{1}\right)$, this bound is shown to be optimal from an asymptotic point of view if the support of $\theta$ belongs to one of the models in the collection. If assumption $\left(\mathbb{H}_{2}\right)$ is fulfilled, we are able to obtain an oracle-type inequality for $\widetilde{\theta}_{\rho_{1}}$. Moreover, $\widetilde{\theta}_{\rho_{1}}$ is minimax adaptive to the sparsity of $\theta$ under $\left(\mathbb{H}_{1}\right)$. Finally, it simultaneously achieves the minimax rates of estimation over a large class of sets $\mathcal{E}(a)$ if $\left(\mathbb{H}_{2}\right)$ holds. Some of these properties still hold if we use the Frobenius loss function. The case of isotropic Gaussian fields is handled similarly.

However, in the oracle inequality (32) and in the minimax bounds (37) and (41), we either perform an assumption on the $l_{1}$ norm of $\theta$ or on the smallest eigenvalue of $\left(I_{p^{2}}-C(\theta)\right)$. When $\|\theta\|_{1}$ tends to one or $\varphi_{\min }\left[I_{p^{2}}-C(\theta)\right]$ tends to 0 , there is a distortion between the upper bound $\mathbb{E}_{\theta}\left[l\left(\widetilde{\theta}_{\rho_{1}}, \theta\right)\right]$ provided by Theorem 3.1 and the lower bounds given by Corollary 4.6 or Proposition 6.2. This limitation seems intrinsic to our penalization method which is linear with respect to the dimension, whereas the asymptotic variance term $\mathbb{E}_{\theta}\left[l\left(\widehat{\theta}_{m, \rho_{1}}, \theta\right)\right]$ depends in a complex way on the dimension of the model $m$ and on the target $\theta$. In our opinion, achieving adaptivity with respect to the smallest eigenvalue of $\left(I_{p^{2}}-C(\theta)\right)$ (or, equivalently, the largest value of $\Sigma$ ) would require a different penalization technique. Nevertheless, we are not aware of any procedure in a covariance estimation setting that is adaptive to the largest eigenvalues of $\Sigma$.

So far, we have provided an estimation procedure for $\left(I_{p^{2}}-C(\theta)\right)=\sigma^{2} \Sigma^{-1}$. If we aim at estimating the precision matrix $\Sigma^{-1}$, we also have to take into account the quantity $\sigma^{2}$. It is natural to estimate it by $\widetilde{\sigma}^{2}:=\gamma_{n, p^{2}}\left(\widetilde{\theta}_{\rho_{1}}\right)$ as done for instance by Guyon in [17], Section 4.3, in the parametric setting. Then, we obtain the estimate $\widetilde{\Sigma}^{-1}:=\widetilde{\sigma}^{2}\left(I_{p^{2}}-C\left(\widetilde{\theta}_{\rho_{1}}\right)\right.$. It is of interest to study the adaptive properties of this estimator with respect to loss functions such as the Frobenius or operator norm as is done in [28] in the nonstationary setting. Nevertheless, let us mention that the matrix $\widetilde{\Sigma}^{-1}$ is not necessarily invertible since the estimator $\widetilde{\theta}_{\rho_{1}}$ belongs to the closure of $\Theta^{+}$.

The choice of the quantity $\rho_{1}$ is problematic. On one hand, $\rho_{1}$ should be large enough so that assumption $\left(\mathbb{H}_{1}\right)$ is fulfilled. On the other hand, a large value of $\rho_{1}$ yields worse bounds in Theorem 3.1. Moreover, the largest eigenvalue of ( $I_{p^{2}}-$ $C(\theta)$ ) is unknown in practice, which makes more difficult the choice of $\rho_{1}$. We see two possible answers to this issue:

- First, moderate values of $\rho_{1}$ are sufficient to enforce $\left(\mathbb{H}_{1}\right)$ if the target $\theta$ is sparse as illustrated in Table 2.
- Second, we believe that the bounds for the risk are pessimistic with respect to $\rho_{1}$. A future direction of research is to derive risk bounds for $\widetilde{\theta}_{\rho_{1}}$ with $\rho_{1}=+\infty$. In [37], we illustrate that such a procedure gives rather good results in practice.

In Theorem 3.1, we only provide a lower bound of the penalty so that the procedure performs well. However, this bound depends on the largest eigenvalue of $\Sigma$ which is seldom known in practice and we did not give any advice for choosing a "reasonable" constant $K$ in practice. This is why we introduce in [37] a data-driven method based on the slope heuristics of Birgé and Massart [7] for calibrating the penalty. We also provide numerical evidence of its performances on simulated data. For instance, the procedure outperforms variogram-based methods for estimating Matérn correlations.

We have mentioned in the Introduction that the toroidal assumption for the lattice is somewhat artificial in several applications. Nevertheless, we needed to neglect the edge effects in order to derive nonasymptotic properties for $\widetilde{\theta}_{\rho_{1}}$ as in Theorem 3.1. In practice, it is often more realistic to suppose that we observe a small window of a Gaussian field defined on the whole plane $\mathbb{Z}^{2}$. The previous nonasymptotic properties do not extend to this new setting. Nevertheless, Lakshman and Derin have shown in [23] that there is no phase transition within the valid parameter space for GMRFs defined on the plane $\mathbb{Z}^{2}$. In short, this implies that the distribution of a field observed in a fixed window of a GMRF does not asymptotically depend on the bound condition. Therefore, it is reasonable to think that our estimation procedure performs well if it was adapted to this new setting. In [37], we describe such an extension and we provide numerical evidence of its performances.
7.3. Possible extensions. In many statistical applications stationary Gaussian fields (or Gaussian Markov random fields) are not directly observed. For instance,

Aykroyd [1] or Dass and Nair [13] use compound Gaussian Markov random fields to account for nonstationarity and steep variations. The wavelet transform has emerged as a powerful tool in image analysis. the wavelet coefficients of an image are sometimes modeled using hidden Markov models [12, 27]. More generally, the success of the GMRF is mainly due to the use of hierarchical models involving latent GMRFs [30]. The study and the implementation of our penalization strategy for selecting the complexity of the latent Markov models is an interesting direction of research.

## 8. Proofs.

8.1. A concentration inequality. In this section, we prove a new concentration inequality for suprema of Gaussian chaos of order 2. It will be useful for proving Theorem 3.1.

Proposition 8.1. Let $F$ be a compact set of symmetric matrices of size $r$, $\left(Y^{1}, \ldots, Y^{n}\right)$ be a $n$-sample of a standard Gaussian vector of size $r$ and $Z$ be the random variable defined by

$$
Z:=\sup _{R \in F} \operatorname{tr}\left[R\left(\overline{Y Y^{*}}-I_{r}\right)\right] .
$$

Then

$$
\begin{equation*}
\mathbb{P}(Z \geq \mathbb{E}(Z)+t) \leq \exp \left[-\left(\frac{t^{2}}{L_{1} \mathbb{E}(W)} \wedge \frac{t}{L_{2} B}\right)\right] \tag{42}
\end{equation*}
$$

where the quantities $B$ and $W$ are such that

$$
\begin{aligned}
B & :=\frac{2}{n} \sup _{R \in F} \varphi_{\max }(R), \\
W & :=\frac{4}{n} \sup _{R \in F} \operatorname{tr}\left(R \overline{Y Y^{*}} R^{\prime}\right)
\end{aligned}
$$

The main argument of this proof is to transfer a deviation inequality for suprema of Rademacher chaos of order 2 to suprema of Gaussian chaos. Talagrand [35] has first given in Theorem 1.2 a concentration inequality for such suprema of Rademacher chaos. Boucheron et al. [8] have recovered the upper bound applying a new methodology based on the entropy method. We adapt their proof to consider nonnecessarily homogeneous chaos of order 2 . More details are found in the technical Appendix [36].

### 8.2. Proof of Theorem 3.1.

Proof. We only consider the case of anisotropic estimators. The proofs and lemma are analogous for isotropic estimators. We first fix a model $m \in \mathcal{M}$. By
definition, the model $\widehat{m}$ satisfies

$$
\gamma_{n, p}\left(\widetilde{\theta}_{\rho_{1}}\right)+\operatorname{pen}(\widehat{m}) \leq \gamma_{n, p}\left(\theta_{m, \rho_{1}}\right)+\operatorname{pen}(m) .
$$

For any $\theta^{\prime} \in \Theta^{+}, \bar{\gamma}_{n, p}\left(\theta^{\prime}\right)$ stands for the difference between $\gamma_{n, p}\left(\theta^{\prime}\right)$ and its expectation $\gamma\left(\theta^{\prime}\right)$. Then, the previous inequality turns into

$$
\gamma\left(\widetilde{\theta}_{\rho_{1}}\right) \leq \gamma\left(\theta_{m, \rho_{1}}\right)+\bar{\gamma}_{n, p}\left(\theta_{m, \rho_{1}}\right)-\bar{\gamma}_{n, p}\left(\widetilde{\theta}_{\rho_{1}}\right)+\operatorname{pen}(m)-\operatorname{pen}(\widehat{m}) .
$$

Subtracting the quantity $\gamma(\theta)$ to both sides of this inequality yields

$$
\begin{equation*}
l\left(\widetilde{\theta}_{\rho_{1}}, \theta\right) \leq l\left(\theta_{m, \rho_{1}}, \theta\right)+\bar{\gamma}_{n, p}\left(\theta_{m, \rho_{1}}\right)-\bar{\gamma}_{n, p}\left(\widetilde{\theta}_{\rho_{1}}\right)+\operatorname{pen}(m)-\operatorname{pen}(\widehat{m}) \tag{43}
\end{equation*}
$$

The proof is based on the control of the random variable $\bar{\gamma}_{n, p}\left(\theta_{m, \rho_{1}}\right)-\bar{\gamma}_{n, p}\left(\widetilde{\theta}_{\rho_{1}}\right)$.
Lemma 8.2. For any positive number $\alpha, \xi$, and $\delta>1$ the event $\Omega_{\xi}$ defined by

$$
\Omega_{\xi}=\left\{\begin{array}{l}
\bar{\gamma}_{n, p}\left(\theta_{m, \rho_{1}}\right)-\bar{\gamma}_{n, p}\left(\widetilde{\theta}_{\rho_{1}}\right) \\
\leq \frac{1}{\sqrt{\delta}} l\left(\widetilde{\theta}_{\rho_{1}}, \theta\right)+\frac{\sqrt{\delta}}{\sqrt{\delta}-1} l\left(\theta_{m, \rho_{1}}, \theta\right) \\
\quad+\frac{K_{0} \delta^{2} \rho_{1}^{2} \varphi_{\max }(\Sigma)}{n p^{2}}\left[(1+\alpha / 2)\left(d_{m}+d_{\widehat{m}}\right)+\frac{\xi^{2}}{\delta-1}\right]
\end{array}\right\}
$$

satisfies

$$
\begin{aligned}
\mathbb{P}\left(\Omega_{\xi}^{c}\right) \leq & \exp \left\{-L_{1} \xi\left[\frac{\alpha}{\sqrt{1+\alpha / 2}} \wedge \sqrt{n}\right]\right\} \\
& \times \sum_{m^{\prime} \in \mathcal{M}} \exp \left\{-L_{2} \sqrt{d_{m^{\prime}}}\left(\frac{\alpha}{\sqrt{1+\alpha / 2}} \wedge \frac{\alpha^{2}}{1+\alpha / 2}\right)\right\}
\end{aligned}
$$

A similar lemma holds in the isotropic case. In particular, we choose $\alpha=(K-$ $\left.K_{0}\right) / K_{0}$ and $\delta=\sqrt{(1+\alpha) /(1+\alpha / 2)}$. Lemma 8.2 implies that on the event $\Omega_{\xi}$,

$$
\begin{aligned}
\bar{\gamma}_{n, p}\left(\theta_{m, \rho_{1}}\right)-\bar{\gamma}_{n, p}\left(\tilde{\theta}_{\rho_{1}}\right) \leq & \frac{1}{\sqrt{\delta(\alpha)}} l\left(\tilde{\theta}_{\rho_{1}}, \theta\right)+\frac{\sqrt{\delta(\alpha)}}{\sqrt{\delta(\alpha)}-1} l\left(\theta_{m, \rho_{1}}, \theta\right)+\operatorname{pen}(m) \\
& +\operatorname{pen}(\widehat{m})+\frac{K_{0} \xi^{2} \delta(\alpha)^{2} \rho_{1}^{2} \varphi_{\max }(\Sigma)}{n p^{2}(\delta(\alpha)-1)}
\end{aligned}
$$

Thus, gathering this bound with inequality (43) yields

$$
\begin{aligned}
\frac{\delta(\alpha)^{1 / 2}-1}{\delta(\alpha)^{1 / 2}} l\left(\tilde{\theta}_{\rho_{1}}, \theta\right) \leq & {\left[1+\delta(\alpha)^{-1 / 2}\left(\delta(\alpha)^{1 / 2}-1\right)^{-1}\right] l\left(\theta_{m, \rho_{1}}, \theta\right)+2 \operatorname{pen}(m) } \\
& +\frac{K_{0} \xi^{2} \rho_{1}^{2} \varphi_{\max }(\Sigma) \delta(\alpha)^{2}}{n p^{2}(\delta(\alpha)-1)}
\end{aligned}
$$

with probability larger than $1-\mathbb{P}\left(\Omega_{\xi}\right)$. Integrating this inequality with respect to $\xi>0$ leads to

$$
\begin{align*}
& \frac{\delta(\alpha)^{1 / 2}-1}{\delta(\alpha)^{1 / 2}} \mathbb{E}_{\theta}\left[l\left(\widetilde{\theta}_{\rho_{1}}, \theta\right)\right] \\
& \leq {\left[1+\delta(\alpha)^{-1 / 2}\left(\delta(\alpha)^{1 / 2}-1\right)^{-1}\right] l\left(\theta_{m, \rho_{1}}, \theta\right) }  \tag{44}\\
&+2 \operatorname{pen}(m)+\frac{\delta(\alpha)^{2} L(\alpha)}{(\delta(\alpha)-1)\left[\alpha^{2} /(1+\alpha / 2) \wedge n\right]} \frac{\rho_{1}^{2} \varphi_{\max }(\Sigma)}{n p^{2}}
\end{align*}
$$

We upper bound $\left[\left(\alpha^{2} /(1+\alpha / 2)\right) \wedge n\right]^{-1}$ by $\left[\left(\alpha^{2} /(1+\alpha / 2)\right) \wedge 1\right]^{-1}$. Since $\alpha=$ $\frac{K-K_{0}}{K_{0}}$, it follows that

$$
\mathbb{E}_{\theta}\left[l\left(\tilde{\theta}_{\rho_{1}}, \theta\right)\right] \leq L_{1}(K)\left[l\left(\theta_{m, \rho_{1}}, \theta\right)+\operatorname{pen}(m)\right]+L_{2}(K) \frac{\rho_{1}^{2} \varphi_{\max }(\Sigma)}{n p^{2}}
$$

Taking the infimum over the models $m \in \mathcal{M}$ allows us to conclude.
Proof of Lemma 8.2. Throughout this proof, it is more convenient to express the quantities $\bar{\gamma}_{n, p}(\cdot)$ and $l(\cdot)$ in terms of covariance and precision matrices. Thanks to (19), we also provide a matricial expression for $\gamma(\cdot)$ :

$$
\begin{equation*}
\gamma\left(\theta^{\prime}\right)=\frac{1}{p^{2}} \operatorname{tr}\left[\left(I-C\left(\theta^{\prime}\right)\right) \Sigma\left(I-C\left(\theta^{\prime}\right)\right)\right] \tag{45}
\end{equation*}
$$

Gathering identities (45) and (17), we get

$$
\begin{aligned}
& \bar{\gamma}_{n, p}\left(\theta_{m, \rho_{1}}\right)-\bar{\gamma}_{n, p}\left(\tilde{\theta}_{\rho_{1}}\right) \\
& \quad=\frac{1}{p^{2}} \operatorname{tr}\left[\left(\left[I_{p^{2}}-C\left(\theta_{m, \rho_{1}}\right)\right]^{2}-\left[I_{p^{2}}-C\left(\tilde{\theta}_{\rho_{1}}\right)\right]^{2}\right)\left(\overline{\mathbf{X}^{\mathbf{v}} \mathbf{X}^{\mathbf{v *}}}-\Sigma\right)\right]
\end{aligned}
$$

Since the matrices $\Sigma$, $\left(I_{p^{2}}-C\left(\theta_{m, \rho_{1}}\right)\right)$ and $\left(I_{p^{2}}-C\left(\tilde{\theta}_{\rho_{1}}\right)\right)$ correspond to covariance or precision matrices of stationary fields on the two-dimensional torus, they are symmetric block circulant. By Lemma A.1, they are jointly diagonalizable in the same orthogonal basis. In the sequel, $P$ stands for an orthogonal matrix associated with this basis. Then the matrices $C\left(\theta_{m, \rho_{1}}\right), C\left(\widetilde{\theta}_{\rho_{1}}\right)$ and $\Sigma$, respectively, decompose in

$$
C\left(\theta_{m, \rho_{1}}\right)=P^{*} D\left(\theta_{m, \rho_{1}}\right) P, \quad C\left(\widetilde{\theta}_{\rho_{1}}\right)=P^{*} D\left(\widetilde{\theta}_{\rho_{1}}\right) P, \quad \Sigma=P^{*} D_{\Sigma} P
$$

where the matrices $D\left(\theta_{m, \rho_{1}}\right), D\left(\widetilde{\theta}_{\rho_{1}}\right)$ and $D_{\Sigma}$ are diagonal. Let the $p^{2} \times n$ matrix $\mathbf{Y}$ be defined by $\mathbf{Y}:=\sqrt{\Sigma^{-1}} \mathbf{X}^{\mathbf{v}}$. Clearly, the components of $\mathbf{Y}$ follow independent standard normal distributions. Gathering these new notation, we get

$$
\begin{align*}
& \bar{\gamma}_{n, p}\left(\theta_{m, \rho_{1}}\right)-\bar{\gamma}_{n, p}\left(\widetilde{\theta}_{\rho_{1}}\right)  \tag{46}\\
& \quad=\frac{1}{p^{2}} \operatorname{tr}\left[\left(\left[I_{p^{2}}-D\left(\theta_{m, \rho_{1}}\right)\right]^{2}-\left[I_{p^{2}}-D\left(\widetilde{\theta}_{\rho_{1}}\right)\right]^{2}\right) D_{\Sigma}\left(\overline{\mathbf{Y}} \mathbf{Y}^{*}-I_{p^{2}}\right)\right]
\end{align*}
$$

Except $\overline{\mathbf{Y} \mathbf{Y}^{*}}$ all the matrices in this last expression are diagonal and we may therefore commute them in the trace.

Let $\langle\cdot, \cdot\rangle_{\mathcal{H}}$ and $\langle\cdot, \cdot\rangle_{\mathcal{H}^{\prime}}$ be two inner products in the space of square matrices of size $p^{2}$, respectively, defined by

$$
\langle A, B\rangle_{\mathcal{H}}:=\frac{\operatorname{tr}\left(A^{*} \Sigma B\right)}{p^{2}} \quad \text { and } \quad\langle A, B\rangle_{\mathcal{H}^{\prime}}:=\frac{\operatorname{tr}\left(A^{*} D_{\Sigma} B\right)}{p^{2}}
$$

This first inner product is related to the loss function $l(\cdot, \cdot)$ through the identity

$$
l\left(\theta^{\prime}, \theta\right)=\left\|C\left(\theta^{\prime}\right)-C(\theta)\right\|_{\mathcal{H}}^{2} .
$$

Moreover, these two inner products clearly satisfy $\left\|C\left(\theta^{\prime}\right)\right\|_{\mathcal{H}}=\left\|D\left(\theta^{\prime}\right)\right\|_{\mathcal{H}^{\prime}}$ for any $\theta^{\prime} \in \Theta^{+}$. Gathering these new notation, we may upper bound (46) by

$$
\left.\begin{array}{l}
\bar{\gamma}_{n, p}\left(\theta_{m, \rho_{1}}\right)-\bar{\gamma}_{n, p}\left(\tilde{\theta}_{\rho_{1}}\right) \\
\quad \leq\left\|\left[I_{p^{2}}-D\left(\theta_{m, \rho_{1}}\right)\right]^{2}-\left[I_{p^{2}}-D\left(\widetilde{\theta}_{\rho_{1}}\right)\right]^{2}\right\|_{\mathcal{H}^{\prime}} \\
\times \sup _{\theta_{1} \in \Theta_{m}, \theta_{2} \in \Theta_{\widehat{m}},} \quad\left\langle\left[I_{p^{2}}-D\left(\theta_{1}\right)\right]^{2}\right.  \tag{47}\\
\quad\left\|\left[I_{p^{2}}-D\left(\theta_{1}\right)\right]^{2}-\left[I_{p^{2}}-D\left(\theta_{2}\right)\right]^{2}\right\|_{\mathcal{H}^{\prime}} \leq 1
\end{array}\right] \begin{aligned}
& \\
& \\
& \\
& \left.\quad-\left[I_{p^{2}}-D\left(\theta_{2}\right)\right]^{2},\left[\overline{\mathbf{Y} \mathbf{Y}^{*}}-I_{p^{2}}\right]\right\rangle_{\mathcal{H}^{\prime}} .
\end{aligned}
$$

The first term in this product is easily bounded as these matrices are diagonal:

$$
\begin{aligned}
\|\left[I_{p^{2}}\right. & \left.-D\left(\theta_{m, \rho_{1}}\right)\right]^{2}-\left[I_{p^{2}}-D\left(\widetilde{\theta}_{\rho_{1}}\right)\right]^{2} \|_{\mathcal{H}^{\prime}} \\
& =\operatorname{tr}\left[\left(\left[I_{p^{2}}-D\left(\theta_{m, \rho_{1}}\right)\right]^{2}-\left[I_{p^{2}}-D\left(\widetilde{\theta}_{\rho_{1}}\right)\right]^{2}\right)^{2} \frac{D_{\Sigma}}{p^{2}}\right]^{1 / 2} \\
& =\operatorname{tr}\left[\left[D\left(\theta_{m, \rho_{1}}\right)-D\left(\widetilde{\theta}_{\rho_{1}}\right)\right]^{2} \frac{D_{\Sigma}}{p^{2}}\left[2 I_{p^{2}}-D\left(\theta_{m, \rho_{1}}\right)-D\left(\widetilde{\theta}_{\rho_{1}}\right)\right]^{2}\right]^{1 / 2} \\
& \leq \varphi_{\max }\left[2 I_{p^{2}}-D\left(\theta_{m, \rho_{1}}\right)-D\left(\widetilde{\theta}_{\rho_{1}}\right)\right]\left\|D\left(\theta_{m, \rho_{1}}\right)-D\left(\widetilde{\theta}_{\rho_{1}}\right)\right\|_{\mathcal{H}^{\prime}}
\end{aligned}
$$

Since $\theta_{m, \rho_{1}}$ and $\tilde{\theta}_{\rho_{1}}$, respectively, belong to $\Theta_{m, \rho_{1}}^{+}$and $\Theta_{\widehat{m}, \rho_{1}}^{+}$, the largest eigenvalues of the matrices $I_{p^{2}}-C\left(\theta_{m, \rho_{1}}\right)$ and $I_{p^{2}}-C\left(\widetilde{\theta}_{\rho_{1}}\right)$ are smaller than $\rho_{1}$. Hence, we get

$$
\begin{aligned}
\varphi_{\max } & {\left[2 I_{p^{2}}-D\left(\theta_{m, \rho_{1}}\right)-D\left(\widetilde{\theta}_{\rho_{1}}\right)\right] } \\
& =\varphi_{\max }\left[I_{p^{2}}-C\left(\theta_{m, \rho_{1}}\right)\right]+\varphi_{\max }\left[I_{p^{2}}-C\left(\widetilde{\theta}_{\rho_{1}}\right)\right] \leq 2 \rho_{1} .
\end{aligned}
$$

Let us turn to the second term in (47). First, we embed the set of matrices over which the supremum is taken in a ball of a vector space. For any model $m^{\prime} \in \mathcal{M}$, let $U_{m^{\prime}}$ be the space generated by the matrices $D\left(\theta^{\prime}\right)^{2}$ and $D\left(\theta^{\prime}\right)$ for $\theta^{\prime} \in \Theta_{m^{\prime}}$. In the sequel, we note $d_{m^{\prime 2}}$ the dimension of $U_{m^{\prime}}$. The space $U_{m, m^{\prime}}$ is defined as the sum of $U_{m}$ and $U_{m^{\prime}}$ whereas $d_{m^{2}, m^{\prime 2}}$ stands for its dimension. Finally, we note
$\mathcal{B}_{m^{2}, m^{\prime 2}}^{\mathcal{H}^{\prime}}$ the unit ball of $U_{m, m^{\prime}}$ with respect to the inner product $\langle\mid\rangle_{\mathcal{H}^{\prime}}$. Gathering these notation, we get

$$
\sup _{\substack{R=\left[I-D\left(\theta_{1}\right)\right]^{2}-\left[I_{p^{2}}-D\left(\theta_{2}\right)\right]^{2}, \theta_{1} \in \Theta_{m}, \theta_{2} \in \Theta_{\widehat{m}} \text { and }\|R\|_{\mathcal{H}^{\prime}} \leq 1}}\left\langle R, \overline{\mathbf{Y} \mathbf{Y}^{*}}-I_{p^{2}}\right\rangle_{\mathcal{H}^{\prime}} \leq \sup _{R \in \mathcal{B}_{\mathcal{B}^{2}, \hat{m}^{\prime}}^{\mathcal{H}^{\prime}}} \frac{1}{p^{2}} \operatorname{tr}\left[R D_{\Sigma}\left(\overline{\mathbf{Y Y}^{*}}-I_{p^{2}}\right)\right] .
$$

Applying the classical inequality $a b \leq \delta a^{2}+\delta^{-1} b^{2} / 4$ and gathering inequalities (47) and (48) yields

$$
\begin{align*}
& \bar{\gamma}_{n, p}\left(\theta_{m, \rho_{1}}\right)-\bar{\gamma}_{n, p}\left(\widetilde{\theta}_{\rho_{1}}\right) \leq \delta^{-1}\left\|C\left(\theta_{m, \rho_{1}}\right)-C\left(\widetilde{\theta}_{\rho_{1}}\right)\right\|_{\mathcal{H}}^{2}  \tag{49}\\
&+\rho_{1}^{2} \delta \sup _{\substack{\mathcal{B}_{m^{\prime}, \tilde{m}^{2}}}} \frac{1}{p^{2}} \operatorname{tr}^{2}\left[R D _ { \Sigma } \left(\overline{\mathbf{Y}} \mathbf{Y}^{*}\right.\right. \\
&\left.\left.\mathcal{H}_{p^{2}}\right)\right] .
\end{align*}
$$

For any model $m^{\prime} \in \mathcal{M}$, we define the random variable $Z_{m^{\prime}}$ as

$$
Z_{m^{\prime}}:=\sup _{R \in \mathcal{B}_{m^{2}, m^{\prime 2}}^{\mathcal{H}^{\prime}}} \frac{1}{p^{2}} \operatorname{tr}\left[R D_{\Sigma}\left(\overline{\mathbf{Y} \mathbf{Y}^{*}}-I_{p^{2}}\right)\right]
$$

The variables $Z_{m^{\prime}}$ turn out to be suprema of Gaussian chaos of order 2. In order to bound $Z_{\widehat{m}}$, we simultaneously control the deviations of $Z_{m^{\prime}}$ for any model $m^{\prime} \in \mathcal{M}$ thanks to the following lemma.

Lemma 8.3. For any positive numbers $\alpha$ and $\xi$ and any model $m^{\prime} \in \mathcal{M}$,

$$
\begin{aligned}
& \mathbb{P}\left(Z_{m^{\prime}} \geq \sqrt{\frac{2 \varphi_{\max }(\Sigma)}{n}}\left\{\sqrt{1+\alpha / 2} \sqrt{d_{m^{2}, m^{\prime 2}}}+\xi\right\}\right) \\
& \quad \leq \exp \left\{-L_{2} \sqrt{d_{m^{\prime}}}\left(\frac{\alpha}{\sqrt{1+\alpha / 2}} \wedge \frac{\alpha^{2}}{1+\alpha / 2}\right)-L_{1} \xi\left[\frac{\alpha}{\sqrt{1+\alpha / 2}} \wedge \sqrt{n}\right]\right\}
\end{aligned}
$$

This result is a consequence from a general concentration inequality for suprema Gaussian chaos of order 2 stated in Proposition 8.1. Its proof is postponed to the technical Appendix [36]. Let us fix the positive numbers $\alpha$ and $\xi$. Applying Lemma 8.3 to any model $m^{\prime} \in \mathcal{M}$, the event $\Omega_{\xi}^{\prime}$ defined by

$$
\Omega_{\xi}^{\prime}=\left\{Z_{\widehat{m}} \leq \sqrt{\frac{2 \varphi_{\max }(\Sigma)}{n}}\left[\sqrt{1+\alpha / 2} \sqrt{d_{m^{2}, \widehat{m}^{2}}}+\xi\right]\right\}
$$

satisfies

$$
\begin{aligned}
\mathbb{P}\left(\Omega_{\xi}^{\prime c}\right) \leq & \exp \left\{-L_{1} \xi\left[\frac{\alpha}{\sqrt{1+\alpha / 2}} \wedge \sqrt{n}\right]\right\} \\
& \times \sum_{m^{\prime} \in \mathcal{M}} \exp \left\{-L_{2} \sqrt{d_{m^{\prime}}}\left(\frac{\alpha}{\sqrt{1+\alpha / 2}} \wedge \frac{\alpha^{2}}{1+\alpha / 2}\right)\right\}
\end{aligned}
$$

From inequality (49), it follows that

$$
\begin{aligned}
& \bar{\gamma}_{n, p}\left(\theta_{m, \rho_{1}}\right)-\bar{\gamma}_{n, p}\left(\tilde{\theta}_{\rho_{1}}\right) \\
& \quad \leq \delta^{-1}\left\|C\left(\theta_{m, \rho_{1}}\right)-C\left(\tilde{\theta}_{\rho_{1}}\right)\right\|_{\mathcal{H}}^{2} \\
& \quad+\frac{2 \delta \rho_{1}^{2} \varphi_{\max }(\Sigma)}{n p^{2}}\left\{\sqrt{1+\alpha / 2} \sqrt{d_{m^{2}, \widehat{m}^{2}}}+\xi\right\}^{2}
\end{aligned}
$$

conditionally to $\Omega_{\xi}^{\prime}$. By the triangle inequality,

$$
\left\|C\left(\theta_{m, \rho_{1}}\right)-C\left(\widetilde{\theta}_{\rho_{1}}\right)\right\|_{\mathcal{H}} \leq\left\|C\left(\theta_{m, \rho_{1}}\right)-C(\theta)\right\|_{\mathcal{H}}+\left\|C\left(\widetilde{\theta}_{\rho_{1}}\right)-C(\theta)\right\|_{\mathcal{H}} .
$$

We recall that the loss function $l\left(\theta^{\prime}, \theta\right)$ equals $\left\|C\left(\theta^{\prime}\right)-C(\theta)\right\|_{\mathcal{H}}^{2}$. We apply twice the inequality $(a+b)^{2} \leq(1+\beta) a^{2}+\left(1+\beta^{-1}\right) b^{2}$. Setting the first $\beta$ to $\sqrt{\delta}-1$, it follows that

$$
\begin{aligned}
& \bar{\gamma}_{n, p}\left(\theta_{m, \rho_{1}}\right)-\bar{\gamma}_{n, p}\left(\tilde{\theta}_{\rho_{1}}\right) \\
& \quad \leq \frac{1}{\sqrt{\delta}} l\left(\widetilde{\theta}_{\rho_{1}}, \theta\right)+\frac{\sqrt{\delta}}{\sqrt{\delta}-1} l\left(\theta_{m, \rho_{1}}, \theta\right) \\
& \quad+\frac{2 \delta \rho_{1}^{2} \varphi_{\max }(\Sigma)}{n p^{2}}\left[d_{m^{2}, \widehat{m}^{2}}(1+\beta)(1+\alpha / 2)+\xi^{2}\left(1+\beta^{-1}\right)\right] .
\end{aligned}
$$

By the definition of $U_{m, \widehat{m}}$, its dimension $d_{m^{2}, \widehat{m}^{2}}$ is bounded by $d_{m^{2}}+d_{\widehat{m}^{2}}$. Choosing $\beta=\delta-1$ yields

$$
\begin{align*}
& \bar{\gamma}_{n, p}\left(\theta_{m, \rho_{1}}\right)-\bar{\gamma}_{n, p}\left(\tilde{\theta}_{\rho_{1}}\right) \\
& \leq \frac{1}{\sqrt{\delta}} l\left(\tilde{\theta}_{\rho_{1}}, \theta\right)+\frac{\sqrt{\delta}}{\sqrt{\delta}-1} l\left(\theta_{m, \rho_{1}}, \theta\right) \\
&+\frac{2 \delta^{2} \rho_{1}^{2} \varphi_{\max }(\Sigma)}{n p^{2}}\left[d_{m^{2}}(1+\alpha / 2)+d_{\widehat{m}^{2}}(1+\alpha / 2)\right]  \tag{50}\\
&+\frac{8 \xi^{2} \varphi_{\max }(\Sigma) \delta^{2}}{n p^{2}(\delta-1)}
\end{align*}
$$

To conclude, we need to compare the dimension $d_{m^{\prime 2}}$ of the space $U_{m^{\prime}}$ with $d_{m^{\prime}}$.

LEMMA 8.4. For any model $m \in \mathcal{M}$, it holds that

$$
d_{m^{2}} \leq L d_{m},
$$

where $L$ is a numerical constant between 4 and 5.48.

The proof is postponed to the technical Appendix [36]. Defining the universal constant $K_{0}:=2 L$, we derive from (50) that

$$
\begin{aligned}
& \bar{\gamma}_{n, p}\left(\theta_{m, \rho_{1}}\right)-\bar{\gamma}_{n, p}\left(\tilde{\theta}_{\rho_{1}}\right) \\
& \quad \leq \frac{1}{\sqrt{\delta}} l\left(\widetilde{\theta}_{\rho_{1}}, \theta\right)+\frac{\sqrt{\delta}}{\sqrt{\delta}-1} l\left(\theta_{m, \rho_{1}}, \theta\right) \\
& \quad+\frac{K_{0} \delta^{2} \rho_{1}^{2} \varphi_{\max }(\Sigma)}{n p^{2}}\left[d_{m}(1+\alpha / 2)+d_{\widehat{m}}(1+\alpha / 2)+\frac{\xi^{2}}{\delta-1}\right]
\end{aligned}
$$

with probability larger than $\mathbb{P}\left(\Omega_{\xi}^{\prime}\right)$. The isotropic case is analogous if we replace $d_{m}$ by $d_{m}^{\text {iso }}$.
8.3. Proofs of the minimax results. Let us first prove a minimax lower bound on hypercubes $\mathcal{C}_{m}\left(\theta^{\prime}, r\right)$. We recall that these hypercubes are introduced in Definition 6.1.

LEMMA 8.5. Let $m$ be a model in $\mathcal{M}_{1}$ that satisfies $d_{m} \leq \sqrt{n} p$, and let $\theta^{\prime}$ be a matrix in $\Theta_{m} \cap \mathcal{B}_{1}\left(0_{p}, 1\right)$. Then, for any positive number $r$ such that $\left(1-\left\|\theta^{\prime}\right\|_{1}-\right.$ $2 r d_{m}$ ) is positive,

$$
\inf _{\widehat{\theta}} \sup _{\theta \in \operatorname{Co}\left[\mathcal{C}_{m}\left(\theta^{\prime}, r\right)\right]} \mathbb{E}_{\theta}[l(\widehat{\theta}, \theta)] \geq L \sigma^{2}\left(r \wedge \frac{1-\left\|\theta^{\prime}\right\|_{1}}{\sqrt{n p^{2}}}\right)^{2} d_{m}
$$

where $\operatorname{Co}\left[\mathcal{C}_{m}\left(\theta^{\prime}, r\right)\right]$ denotes the convex hull of $C_{m}\left(\theta^{\prime}, r\right)$. Similarly, let $m$ be a model in $\mathcal{M}_{1}$ such $d_{m}^{\text {iso }} \leq \sqrt{n} p$, and let $\theta^{\prime}$ be a matrix in $\Theta_{m}^{\text {iso }} \cap \mathcal{B}_{1}\left(0_{p}, 1\right)$. Then, for any positive number $r$ such that $\left(1-\left\|\theta^{\prime}\right\|_{1}-8 r d_{m}^{\text {iso }}\right)$ is positive,

$$
\inf _{\widehat{\theta}} \sup _{\theta \in \operatorname{Co}\left[\mathcal{C}_{m}^{\text {iso }}\left(\theta^{\prime}, r\right)\right]} \mathbb{E}_{\theta}[l(\widehat{\theta}, \theta)] \geq L \sigma^{2}\left(r \wedge \frac{1-\left\|\theta^{\prime}\right\|_{1}}{\sqrt{n p^{2}}}\right)^{2} d_{m}^{\text {iso }}
$$

Proof of Proposition 6.2. The first result derives from Lemma 8.5 applied to the hypercube $\mathcal{C}_{m}\left(0_{p},\left(n p^{2}\right)^{-1 / 2}\right)$. We prove the second result using the same lemma with $\mathcal{C}_{m}\left[\theta^{\prime},\left(1-\|\theta\|_{1}\right) /(\sqrt{n} p)\right]$.

Proof of Lemma 8.5. This lower bound is based on an application of Fano's approach. See [38] for a review of this method and comparisons with Le Cam's and Assouad's lemma. The proof follows three main steps: first, we upper bound the Kullback-Leibler entropy between distributions corresponding to $\theta_{1}$ and $\theta_{2}$ in the hypercube. Second, we find a set of points in the hypercube well separated with respect to the Hamming distance. Finally, we conclude by applying Birgé's version of Fano's lemma. More details can be found in the technical Appendix [36].

Proof of Proposition 6.6. First, observe that the set $\mathcal{E}(a) \cap \mathcal{B}_{1}\left(0_{p}, 1 / 2\right)$ is included in $\mathcal{E}(a) \cap \mathcal{B}_{1}\left(0_{p}, 1\right) \cap \mathcal{U}(2)$. We then derive minimax lower bounds on $\mathcal{E}(a) \cap \mathcal{B}_{1}\left(0_{p}, 1 / 2\right)$ from the lower bounds on hypercubes.

Let $m_{i}$ be a model in $\mathcal{M}_{1}$ such that $d_{m}$ is smaller than $\sqrt{n} p$. Let us look for positive numbers $r$ such that the hypercube $\left[\mathcal{C}_{m_{i}}\left(0_{p}, r\right)\right]$ is included in the set $\mathcal{E}(a) \cap \mathcal{B}_{1}\left(0_{p}, 1 / 2\right)$.

LEMMA 8.6. Let $m$ be a model in $\mathcal{M}_{1}$ and $r$ be a positive number smaller than $1 /\left(4 d_{m}\right)$. For any $\theta \in \operatorname{Co}\left[\mathcal{C}_{m}\left(0_{p}, r\right)\right]$,

$$
\operatorname{var}_{\theta}\left(X_{[0,0]}\right) \leq \sigma^{2}\left(1+16 d_{m} r^{2}\right)
$$

The proof is postponed to the technical Appendix [36]. If we choose

$$
r \leq \frac{a_{i}}{16 \sigma \sqrt{d_{m_{i}}}}
$$

then $2 r d_{m_{i}}$ is smaller than $1 / 8$ by assumption $\left(\mathbb{H}_{a}\right)$. Applying Lemma 8.6, we then derive that $\operatorname{Var}_{\theta}\left(X_{[0,0]}\right) \leq \sigma^{2}+a_{i}^{2}$. Hence, we get the upper bound $\sum_{j=1}^{i}\left[\operatorname{Var}\left(X_{[0,0]} \mid X_{m_{j-1}}\right)-\operatorname{Var}\left(X_{[0,0]} \mid X_{m_{j}}\right)\right] \leq a_{i}^{2}$ and it follows that

$$
\sum_{j=1}^{\operatorname{Card}\left(\mathcal{M}_{1}\right)} \frac{\operatorname{Var}\left(X_{[0,0]} \mid X_{m_{k-1}}\right)-\operatorname{Var}\left(X_{[0,0]} \mid X_{m_{j}}\right)}{a_{j}^{2}} \leq 1
$$

since the sequence $\left(a_{j}\right)_{1 \leq j \leq \operatorname{Card}\left(\mathcal{M}_{1}\right)}$ is nonincreasing. Consequently, $\operatorname{Co}\left[\mathcal{C}_{m}\left(0_{p}\right.\right.$, $r)]$ is a subset of $\mathcal{E}(a) \cap \mathcal{B}_{1}\left(0_{p}, 1 / 2\right)$. By Lemma 8.5 , we get

$$
\begin{align*}
\inf _{\hat{\theta}} \sup _{\theta \in \mathcal{E}(a) \cap \mathcal{B}_{1}\left(0_{p}, 1 / 2\right)} \mathbb{E}_{\theta}[l(\widehat{\theta}, \theta)] & \geq L \sigma^{2}\left(\frac{a_{i}^{2}}{16 \sigma^{2}} \wedge \frac{d_{m_{i}}}{n p^{2}}\right) \\
& \geq L\left(a_{i}^{2} \wedge \frac{\sigma^{2} d_{m_{i}}}{n p^{2}}\right) . \tag{51}
\end{align*}
$$

Considering all models $m \in \mathcal{M}_{1}$ such that $d_{m} \leq \sqrt{n} p$ yields

$$
\begin{equation*}
\inf _{\widehat{\theta}} \sup _{\theta \in \mathcal{E}(a) \cap \mathcal{B}_{1}\left(0_{p}, 1 / 2\right)} \mathbb{E}_{\theta}[l(\widehat{\theta}, \theta)] \geq L \sup _{i \leq \operatorname{Card}\left(\mathcal{M}_{1}\right), d_{m_{i}} \leq \sqrt{n} p}\left(a_{i}^{2} \wedge \frac{\sigma^{2} d_{m_{i}}}{n p^{2}}\right) \tag{52}
\end{equation*}
$$

If the maximal dimension $d_{m_{\operatorname{Card}\left(\mathcal{M}_{1}\right)}}$ is smaller than $\sqrt{n} p$, the proof is complete. In the opposite case, we need to show that the supremum (40) over all models $m \in \mathcal{M}_{1}$ is achieved at some model $m$ of dimension less than $\sqrt{n} p$.

Lemma 8.7. For any integer $1 \leq i \leq \operatorname{Card}\left(\mathcal{M}_{1}\right)-1$, the ratio $d_{m_{i+1}} / d_{m_{i}}$ is less than 2.

The proof of Lemma 8.7 is postponed to the technical Appendix [36]. Let $i^{\prime}$ be the largest integer such that $d_{m_{i^{\prime}}} \leq \sqrt{n} p$. Since $i^{\prime}$ is smaller than $\operatorname{Card}\left(\mathcal{M}_{1}\right)$, we know from Lemma 8.7 that $\sqrt{n} p / 2 \leq d_{m_{i^{\prime}}} \leq \sqrt{n} p$. By assumption $\left(\mathbb{H}_{a}\right), a_{i^{\prime}}^{2}$ is smaller than $\sigma^{2} / d_{m^{\prime}}$. Gathering these bounds yields

$$
a_{i^{\prime}}^{2} \leq \frac{\sigma^{2}}{d_{m_{i^{\prime}}}} \leq \frac{4 d_{m_{i^{\prime}}} \sigma^{2}}{n p^{2}}
$$

Since the sequence $\left(a_{i}\right)_{1 \leq i \leq \operatorname{Card}\left(\mathcal{M}_{1}\right)}$ is nonincreasing, the supremum (40) over all models in $\mathcal{M}_{1}$ is either achieved for some $i \leq i^{\prime}$ or is smaller than $4\left(a_{i^{\prime}}^{2} \wedge\right.$ $\left.\sigma^{2} d_{m_{i^{\prime}}}\left(n p^{2}\right)\right)$.

Proof of Corollary 6.3. Observe that $\operatorname{Co}\left[\mathcal{C}_{m}\left(0_{p}, 1 /\left(4 d_{m}\right)\right]\right.$ is included in $\Theta_{m} \cap \mathcal{B}_{1}\left(0_{p}, 1 / 2\right)$. This last set is, itself, included in $\Theta_{m, \rho_{1}}^{+} \cap \mathcal{U}\left(\rho_{2}\right)$. Applying Lemma 8.5, we get the following minimax lower bound:

$$
\inf _{\hat{\theta}} \sup _{\theta \in \Theta_{m, \rho_{1}}^{+} \cap \mathcal{U}\left(\rho_{2}\right)} \mathbb{E}[l(\widehat{\theta}, \theta)] \geq L \sigma^{2} \frac{d_{m}}{n p^{2}}
$$

since the dimension $d_{m}$ is smaller than $n p^{2}$. Applying Theorem 3.1, we derive that

$$
\begin{aligned}
\sup _{\theta \in \Theta_{m, \rho_{1}}^{+} \cap \mathcal{U}\left(\rho_{2}\right)} \mathbb{E}\left[l\left(\tilde{\theta}_{\rho_{1}}, \theta\right)\right] \leq & L(K) \sigma^{2} \rho_{1}^{2} \rho_{2} \frac{d_{m}}{n p^{2}} \\
& +L_{2}(K) \frac{\rho_{1}^{2}}{n p^{2}} \sup _{\theta \in \Theta_{m, \rho_{1}}^{+} \mathfrak{\mathcal { U } ( \rho _ { 2 } )}} \varphi_{\max }(\Sigma) \\
\leq & L\left(K, \rho_{1}, \rho_{2}\right) \sigma^{2} \frac{d_{m}}{n p^{2}}
\end{aligned}
$$

We conclude by combining the two different bounds.
Proof of Proposition 6.7. This result derives from the upper bound of the risk of $\widetilde{\theta}_{\rho_{1}}$ stated in Theorem 3.1 and the minimax lower bound stated in Proposition 6.6. For details, we refer to the technical Appendix [36].

### 8.4. Proofs of the asymptotic risk bounds.

Proof of Proposition 4.4. This result is closely related to Proposition 4.11 in [17]. In fact, we extend his proof to stationary fields on a torus. In the sequel, we shall only consider nonisotropic GMRFs, the isotropic case being similar. Let us fix a model $m$ in the collection $\mathcal{M}_{1}$ and let us assume $\left(\mathbb{H}_{1}\right)$.

We define the $d_{m} \times p^{2}$ matrix $\chi_{m}^{v}$ as

$$
\left(\chi_{m}^{v}\right)^{*}:=\left(\left[C\left(\Psi_{i_{k}, j_{k}}\right) X^{v}\right], k=1, \ldots, d_{m}\right)
$$

For any $(i, j) \in\{1, \ldots, p\}^{2}$, the $[(i-1) p+j]$ th row of $\chi_{m}^{v}$ corresponds to the list of covariates used when performing the regression of $X_{[i, j]}$ with respect to its neighbors in the model $m$. Contrary to the previous proofs, we need to express the $n \times p^{2}$ matrix $\mathbf{X}^{\mathbf{v}}$ in terms of a vector. This is why we define the vector $\mathbf{X}^{\mathbf{V}}$ of size $n p^{2}$ as

$$
\mathbf{X}_{\left[p^{2}(j-1)+p\left(i_{1}-1\right)+i_{2}\right]}^{\mathbf{V}}:=\mathbf{X}_{\left[i_{1}, i_{2}\right]}^{j}
$$

for any $\left(i_{1}, i_{2}\right) \in\{1, \ldots, p\}^{2}$ and any $j \leq n$. Similarly, let $\chi_{m}^{V}$ be the $d_{m} \times n p^{2}$ matrix defined as

$$
\chi_{m\left[k, p^{2}(j-1)+p\left(i_{1}-1\right)+i_{2}\right]}^{\mathbf{V}}:=\chi_{m\left[p\left(i_{1}-1\right)+i_{2}\right]}^{j}
$$

for any $\left(i_{1}, i_{2}\right) \in\{1, \ldots, p\}^{2}$ and any $j \leq n$.
We are not able to work out directly the asymptotic risk of $\widehat{\theta}_{m, \rho_{1}}$. This is why we introduce a new estimator $\check{\theta}_{m}$ whose asymptotic distribution is easier to derive. Afterward, we shall prove that $\breve{\theta}_{m}$ and $\widehat{\theta}_{m, \rho_{1}}$ have the same asymptotic distribution. Let us, respectively, define the estimators $\check{a}_{m}$ in $\mathbb{R}^{d_{m}}$ and $\check{\theta}_{m}$ as

$$
\begin{align*}
& \check{a}_{m}:=\left(\left(\boldsymbol{\chi}_{m}^{\mathbf{V}}\right)^{*} \chi_{m}^{\mathbf{V}}\right)^{-1} \boldsymbol{\chi}_{m}^{\mathbf{V}} \mathbf{X}^{\mathbf{V}},  \tag{53}\\
& \check{\theta}_{m}:=\sum_{k=1}^{d_{m}} \check{a}_{m[k]} \Psi_{i_{k}, j_{k}},
\end{align*}
$$

where we recall that $\left(\Psi_{i_{1}, j_{1}}, \ldots, \Psi_{i_{d_{m}}, j_{d_{m}}}\right.$ ) is a basis of $\Theta_{m}$. Obviously, $\check{\theta}_{m}$ is a conditional least squares estimator since it minimizes the expression (16) of $\gamma_{n, p}(\cdot)$ over the whole space $\Theta_{m}$. Consequently, $\check{\theta}_{m}$ coincides with $\widehat{\theta}_{m, \rho_{1}}$ if $\check{\theta}_{m}$ belongs to $\Theta_{m, \rho_{1}}^{+}$.

For the second result, we assume that assumption $\left(\mathbb{H}_{2}\right)$ holds. Applying Corollary 4.2, we know that for any $(k, l) \in \Lambda, X_{[k, l]}$ decomposes as

$$
\begin{equation*}
X_{[k, l]}=\sum_{(i, j) \in m} \theta_{m, \rho_{1}[i, j]} X_{[k+i, l+j]}+\varepsilon_{m[k, l]}, \tag{54}
\end{equation*}
$$

where $\varepsilon_{m[k, l]}$ is independent from $\left\{X_{[k+i, l+j]},(i, j) \in m\right\}$. For the first result, the same decomposition holds since $\theta$ is assumed to belong to $\Theta_{m, \rho_{1}}^{+}$, and $\theta_{m, \rho_{1}}$, therefore, equals $\theta$.

Let $a_{m} \in \mathbb{R}^{d_{m}}$ be the unique vector such that $\theta_{m, \rho_{1}}=\sum_{k=1}^{d_{m}} a_{m[k]} \Psi_{i_{k}, j_{k}}$. Then, the previous decomposition becomes

$$
X^{v}=a_{m}^{*} \chi_{m}^{v}+\varepsilon_{m}^{v}
$$

Gathering this last identity with (53) yields

$$
\check{a}_{m}-a_{m}=\left(\frac{1}{n p^{2}}\left(\chi_{m}^{\mathbf{V}}\right)^{*} \chi_{m}^{\mathbf{V}}\right)^{-1}\left(\frac{1}{n p^{2}} \chi_{m}^{\mathbf{V}} \boldsymbol{\varepsilon}_{m}^{\mathbf{V}}\right)
$$

where the vector $\boldsymbol{\varepsilon}_{m}^{\mathbf{V}}$ of size $n p^{2}$ corresponds to the $n$ observations of the vector $\varepsilon_{m}^{v}$. When $n$ goes to the infinity, $1 /\left(n p^{2}\right)\left(\chi_{m}^{\mathbf{V}}\right)^{*} \chi_{m}^{\mathbf{V}}$ converges almost surely to the covariance matrix $V$ by the law of large numbers. By definition, the variable $\varepsilon_{m[i, j]}$ is independent from the $[(i-1) p+j]$ th row of $\chi_{m[i, j]}^{v}$. It follows that $\mathbb{E}_{\theta}\left(\chi_{m}^{\mathbf{V}} \boldsymbol{\varepsilon}^{\mathbf{V}}\right)=0$. Applying again the law of large numbers we conclude that $\check{a}_{m}$ converges almost surely toward $a_{m}$ and that $\check{\theta}_{m}$ converges almost surely toward $\theta_{m, \rho_{1}}$. Additionally, the central limit theorem states that the random vector $1 /(\sqrt{n} p) \chi_{m}^{\mathbf{V}} \varepsilon^{\mathbf{V}}$ converges in distribution toward a zero mean Gaussian vector whose covariance matrix equals $1 / p^{2} \operatorname{Var}_{\theta}\left(\chi_{m}^{v} \varepsilon_{m}^{v}\right)$. By decomposition (54), $\varepsilon_{m}^{v}=\left(I-C\left(\theta_{m, \rho_{1}}\right)\right) X^{v}$ while the $k$ th row of $\chi_{m}^{v}$ equals $\left[C\left(\Psi_{i_{k}, j_{k}}\right) X^{v}\right]^{*}$. Thus for any $1 \leq k, l \leq d_{m}$,

$$
\begin{array}{r}
\frac{1}{p^{2}} \operatorname{Var}_{\theta}\left(\chi_{m}^{v} \varepsilon_{m}^{v}\right)_{[k, l]}=\frac{1}{p^{2}} \operatorname{cov}_{\theta}\left[\left(X^{v}\right)^{*} C\left(\Psi_{i_{k}, j_{k}}\right)\left[I-C\left(\theta_{m, \rho_{1}}\right)\right] X^{v}\right. \\
\left.\left(X^{v}\right)^{*} C\left(\Psi_{i_{l}, j_{l}}\right)\left[I-C\left(\theta_{m, \rho_{1}}\right)\right] X^{v}\right] .
\end{array}
$$

As the covariance matrix of $X^{v}$ is $\sigma^{2}(I-C(\theta))^{-1}$, we obtain, by standard Gaussian properties,

$$
\begin{aligned}
\frac{1}{p^{2}} & \operatorname{Var}_{\theta}\left(\chi_{m}^{v} \varepsilon_{m}^{v}\right)_{[k, l]} \\
=\frac{2 \sigma^{4}}{p^{2}} \operatorname{cov}_{\theta} & {\left[[I-C(\theta)]^{-1} C\left(\Psi_{i_{k}, j_{k}}\right)\left[I-C\left(\theta_{m, \rho_{1}}\right)\right]\right.} \\
& \left.\times[I-C(\theta)]^{-1} C\left(\Psi_{i_{l}, j_{l}}\right)\left[I-C\left(\theta_{m, \rho_{1}}\right)\right]\right]
\end{aligned}
$$

By Lemma A.1, all these matrices are diagonalizable in the same basis and, therefore, commute with each other. We conclude that $\frac{1}{p^{2}} \operatorname{Var}_{\theta}\left(\chi_{m}^{v} \varepsilon_{m}^{v}\right)=2 \sigma^{4} W$, and

$$
\sqrt{n} p\left(\check{a}_{m}-a_{m}\right) \rightarrow \mathcal{N}\left(0, V^{-1} W V^{-1}\right)
$$

As $\widehat{\theta}_{m, \rho_{1}}$ belongs to $\Theta_{m, \rho_{1}}^{+}$, there exists a unique vector $\widehat{a}_{m} \in \mathbb{R}^{d_{m}}$ such that $\widehat{\theta}_{m, \rho_{1}}=\sum_{k=1}^{d_{m}} \widehat{a}_{m[k]} \Psi_{i_{k}, j_{k}}$. The matrix $\theta_{m, \rho_{1}}$ belongs to the open set $\Theta_{m, \rho_{1}}^{+}$for the two cases of the propositions. Indeed, $\theta_{m, \rho_{1}}$ equals $\theta$ in the first situation. In the second situation, this is due to the fact that $\theta$ satisfies $\left(\mathbb{H}_{2}\right)$ and to Lemma 4.1.

Since $\check{\theta}_{m}$ converges almost surely to $\theta_{m, \rho_{1}}$, the matrix $\check{\theta}_{m}$ belongs to $m$ with probability going to one when $n$ goes to infinity. If follows that the estimators $\check{a}_{m}$ and $\widehat{a}_{m}$ coincide with probability going to one. By Slutsky's lemma, we obtain that

$$
\sqrt{n} p\left(\widehat{a}_{m}-a_{m}\right) \rightarrow \mathcal{N}\left(0, V^{-1} W V^{-1}\right)
$$

Let us express the risk of $\widehat{\theta}_{m, \rho_{1}}$ with respect to the distribution of $\widehat{a}_{m}$ :

$$
\begin{aligned}
l\left(\widehat{\theta}_{m, \rho_{1}}, \theta_{m, \rho_{1}}\right) & =\mathbb{E}_{\theta}\left[\sum_{k=1}^{d_{m}}\left(\widehat{a}_{m[k]}-a_{m[k]}\right) \operatorname{tr}\left(\Psi_{i_{k}, j_{k}} X\right)\right]^{2} \\
& =\operatorname{tr}\left[V\left(\widehat{a}_{m}-a_{m}\right)^{*}\left(\widehat{a}_{m}-a_{m}\right)\right]
\end{aligned}
$$

By Portmanteau's lemma, $n p^{2} l\left(\widehat{\theta}_{m, \rho_{1}}, \theta_{m, \rho_{1}}\right)$ converges in distribution toward a random variable whose expectation is $\operatorname{tr}\left(W V^{-1}\right)$. In order to conclude, it remains to prove that the sequence $\left[n p^{2} l\left(\widehat{\theta}_{m, \rho_{1}}, \theta\right)\right]_{n \geq 1}$ is asymptotically uniformly integrable.

Let us consider a model selection procedure with the collection $\mathcal{M}=\{m\}$ and a penalty term satisfying the assumptions of Theorem 3.1. Arguing as in the proof of this theorem, we derive from identity (44) the following property. For any $\xi>0$, with probability larger than $1-L_{1} \exp \left[-L_{2} \xi\right]$,

$$
n p^{2} l\left(\widehat{\theta}_{m, \rho_{1}}, \theta_{m, \rho_{1}}\right) \leq L_{3} d_{m} \varphi_{\max }(\Sigma)+L_{4} \xi^{2} \varphi_{\max }(\Sigma)
$$

This clearly implies that the sequence $\left[n p^{2} l\left(\widehat{\theta}_{m, \rho_{1}}, \theta_{m, \rho_{1}}\right)\right]_{n \geq 1}$ is asymptotically uniformly integrable and the first part of the result follows.

For the first result of the proposition, we have stated that $\theta$ equals $\Theta_{m}$. As a consequence,

$$
\lim _{n \rightarrow+\infty} \mathbb{E}_{\theta}\left[l\left(\widehat{\theta}_{m, \rho_{1}}, \theta\right)\right]=2 \sigma^{4} \operatorname{tr}\left[W V^{-1}\right]
$$

Also, the term $W_{[k, l]}$ here equals $\operatorname{tr}\left[C\left(\Psi_{i_{k}, j_{k}}\right) C\left(\Psi_{i_{l}, j_{l}}\right)\right]$. This last quantity is zero if $k \neq l$ and equals $\left\|C\left(\Psi_{i_{k}, j_{k}}\right)\right\|_{F}^{2}$ if $k=l$.

Proof of Proposition 4.7. As $\theta$ belongs to $\Theta^{+} \cap \mathcal{B}_{1}\left(0_{p}, \eta\right)$, the largest eigenvalue of $\Sigma$ is smaller than $\sigma^{2} /(1-\eta)$. Applying Theorem 3.1, we get

$$
\begin{aligned}
\mathbb{E}_{\theta}\left[l\left(\tilde{\theta}_{\rho_{1}}, \theta\right)\right] & \leq L(K) \inf _{m \in \mathcal{M}}\left[l\left(\theta_{m, \rho_{1}}, \theta\right)+K \frac{\sigma^{2}}{n p^{2}(1-\eta)}\right] \\
& \leq L(K, \eta) \inf _{m \in \mathcal{M}}\left[l\left(\theta_{m, \rho_{1}}, \theta\right)+K \frac{\sigma^{2}}{n p^{2}}(1-\eta)^{3}\right]
\end{aligned}
$$

Gathering this bound with the result of Corollary 4.6 enables us to conclude.

## APPENDIX

Lemma A.1. There exists an orthogonal matrix $P$ which simultaneously diagonalizes every $p^{2} \times p^{2}$ symmetric block circulant matrices with $p \times p$ blocks. Conversely, if $\theta$ is a square matrix of size $p$ which satisfies (3), then the matrix $D(\theta)=P C(\theta) P^{*}$ is diagonal and satisfies

$$
\begin{equation*}
D(\theta)_{[(i-1) p+j,(i-1) p+j]}=\sum_{k=1}^{p} \sum_{l=1}^{p} \theta_{[k, l]} \cos (2 \pi(k i / p+l j / p)) \tag{55}
\end{equation*}
$$

for any $1 \leq i, j \leq p$.
This lemma is proved in [29], Section 2.6 .2 when is $P$ a unitary matrix. A slight modification of their proof allows to show that $P$ is orthogonal in our case. The difference comes from the fact that contrary to Rue and Held we also assume that $C(\theta)$ is symmetric.

This lemma states that all symmetric block circulant matrices are simultaneously diagonalizable. Moreover, expression (55) explicitly provides the eigenvalues of the $C(\theta)$ as the two-dimensional discrete Fourier transform of the $p \times p$ matrix $\theta$.

Acknowledgments. I am grateful to Pascal Massart for many fruitful discussions. I also thank the referees and the associate editor for their suggestions that led to an improvement of the manuscript.

## REFERENCES

[1] Aykroyd, R. (1998). Bayesian estimation for homogeneous and inhomogeneous Gaussian random fields. IEEE Trans. Pattern Anal. Machine Intell. 20 533-539.
[2] Besag, J. E. (1975). Statistical analysis of non-lattice data. Statistica 24 179-195.
[3] Besag, J. E. (1977). Efficiency of pseudolikelihood estimation for simple Gaussian fields. Biometrika 64 616-618. MR0494640
[4] Besag, J. E. and Kooperberg, C. (1995). On conditional and intrinsic autoregressions. Biometrika 82 733-746. MR1380811
[5] Besag, J. E. and Moran, P. A. P. (1975). On the estimation and testing of spatial interaction in Gaussian lattice processes. Biometrika 62 555-562. MR0391451
[6] Birgé, L. and Massart, P. (2001). Gaussian model selection. J. Eur. Math. Soc. (JEMS) 3 203-268. MR1848946
[7] Birgé, L. and Massart, P. (2007). Minimal penalties for Gaussian model selection. Probab. Theory Related Fields 138 33-73. MR2288064
[8] Boucheron, S., Bousquet, O., Lugosi, G. and Massart, P. (2005). Moment inequalities for functions of independent random variables. Ann. Probab. 33 514-560. MR2123200
[9] Brockwell, P. J. and Davis, R. A. (1991). Time Series: Theory and Methods, 2nd ed. Springer, New York. MR1093459
[10] Cressie, N. A. C. (1993). Statistics for Spatial Data. Wiley, New York. MR1239641
[11] Cressie, N. A. C. and Verzelen, N. (2008). Conditional-mean least-squares of Gaussian Markov random fields to Gaussian fields. Comput. Statist. Data Anal. 52 2794-2807. MR2419542
[12] Crouse, M., Nowak, R. and Baraniuk, R. (1998). Wavelet-based statistical signal processing using hidden Markov models. IEEE Trans. Signal Process. 46 886-902. MR1665651
[13] Dass, S. C. and NAIR, V. N. (2003). Edge detection, spatial smoothing, and image reconstruction with partially observed multivariate data. J. Amer. Statist. Assoc. 98 77-89. MR1977200
[14] Edwards, D. (2000). Introduction to Graphical Modelling, 2nd ed. Springer, New York. MR1880319
[15] Gray, R. (2006). Toeplitz and Circulant Matrices: A Review, rev. ed. Now Publishers, Norwell, MA.
[16] Guyon, X. (1987). Estimation d'un champ par pseudo-vraisemblance conditionnelle: Étude asymptotique et application au cas Markovien. In Spatial processes and spatial time series analysis (Brussels, 1985). Travaux Rech. 11 15-62. Publ. Fac. Univ. Saint-Louis, Brussels. MR0947996
[17] Guyon, X. (1995). Random Fields on a Network. Springer, New York. MR1344683
[18] GUYON, X. and YaO, J. (1999). On the underfitting and overfitting sets of models chosen by order selection criteria. J. Multivariate Anal. 70 221-249. MR1711522
[19] Hall, P., Fisher, N. and Hoffmann, B. (1994). On the nonparametric estimation of covariance functions. Ann. Statist. 22 2115-2134. MR1329185
[20] Hurvich, C. and Tsai, C.-L. (1989). Regression and time series model selection in small samples. Biometrika 76 297-307. MR1016020
[21] Im, H., Stein, M. and Zhu, Z. (2007). Semiparametric estimation of spectral density with irregular observations. J. Amer. Statist. Assoc. 102 726-735. MR2381049
[22] Kashyap, R. and Chellapa, R. (1984). Estimation and choice of neighbors in spatialinteraction models of images. IEEE Trans. Inform. Theory 29 60-72. MR0781270
[23] Lakshmanan, S. and Derin, H. (1993). Valid parameter space for 2-D Gaussian Markov random fields. IEEE Trans. Inform. Theory 39 703-709. MR1224361
[24] Lauritzen, S. L. (1996). Graphical Models. Oxford Statistical Science Series 17. Oxford Univ. Press, New York. MR1419991
[25] Massart, P. (2007). Concentration Inequalities and Model Selection. Lecture Notes in Math. 1896. Springer, Berlin. MR2319879
[26] McQuarrie, A. D. R. and Tsai, C.-L. (1998). Regression and Time Series Model Selection. World Scientific, River Edge, NJ. MR1641582
[27] Portilla, J., Strela, V., Wainwright, M. J. and Simoncelli, E. P. (2003). Image denoising using scale mixtures of Gaussians in the wavelet domain. IEEE Trans. Image Process. 12 1338-1351. MR2026777
[28] Rothman, A. J., Bickel, P. J., Levina, E. and Zhu, J. (2008). Sparse permutation invariant covariance estimation. Electron. J. Stat. 2 494-515. MR2417391
[29] Rue, H. and Held, L. (2005). Gaussian Markov Random Fields: Theory and Applications. Monographs on Statistics and Applied Probability 104. Chapman \& Hall/CRC, London. MR2130347
[30] Rue, H., Martino, S. and Chopin, N. (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. J. R. Stat. Soc. Ser. B Stat. Methodol. 71 319-392.
[31] Rue, H. and Tjelmeland, H. (2002). Fitting Gaussian Markov random fields to Gaussian fields. Scand. J. Statist. 29 31-49. MR1894379
[32] Shibata, R. (1980). Asymptotically efficient selection of the order of the model for estimating parameters of a linear process. Ann. Statist. 8 147-164. MR0557560
[33] Song, H.-R., Fuentes, M. and Ghosh, S. (2008). A comparative study of Gaussian geostatistical models and Gaussian Markov random field models. J. Multivariate Anal. 99 1681-1697.
[34] Stein, M. L. (1999). Interpolation of Spatial Data: Some Theory for Kriging. Springer, New York. MR1697409
[35] Talagrand, M. (1996). New concentration inequalities in product spaces. Invent. Math. 126 505-563. MR1419006
[36] Verzelen, N. (2009). Technical Appendix to "Adaptive estimation of stationary Gaussian fields." Available at arXiv:0908.4586.
[37] Verzelen, N. (2010). Data-driven neighborhood selection of a Gaussian field. Comput. Statist. Data Anal. To appear.
[38] Yu, B. (1997). Assouad, Fano and Le Cam. In Festschrift for Lucien Le Cam 423-435. Springer, New York. MR1462963
INRA AND SUPAGRO
UMR 729 MISTEA
BÂTIMENT 29
2, PLACE PIERRE VIALA
F-34060 MONTPELLIER
FRANCE
E-MAIL: nicolas.verzelen@supagro.inra.fr


[^0]:    Received January 2009; revised September 2009.
    ${ }^{1}$ Research mostly carried out at Univ. Paris-Sud (Laboratoire de Matématiques, CNRS-UMR 8628).

    AMS 2000 subject classifications. Primary 62H11; secondary 62M40.
    Key words and phrases. Gaussian field, Gaussian Markov random field, model selection, pseudolikelihood, oracle inequalities, minimax rate of estimation.

