A Generalized Maximum Pseudo-likelihood Estimator for Markov Random Fields with Noise, I: General Theory

By

David J. Barsky and Alberto Gandolfi

Technical Report No. 373
October 1992
(revised April 1993)

Department of Statistics
University of California
Berkeley, California 94720

A GENERALIZED MAXIMUM PSEUDO-LIKELIHOOD ESTIMATOR FOR MARKOV RANDOM FIELDS WITH NOISE, I: GENERAL THEORY

BY DAVID J. BARSKY AND ALBERTO GANDOLFI

University of California, Davis and University of California, Berkeley

Abstract

In this paper we present an asymptotic estimator, obtained by observing a noisy image, for the parameters of both a stationary Markov random field and an independent Bernoulli noise.

We first estimate the parameter of the noise, by solving a polynomial equation of moderate degree (about 6-7 in the one-dimensional Ising model, and about 10-15 in the two-dimensional Ising model, for instance), and then apply the maximum pseudo-likelihood method after removing the noise. Our method requires no extra simulation, and is likely to be applicable to any Markov random field, in any dimension.

Here, we present the general theory and some examples in one dimension; more interesting examples in two dimensions will be discussed at length in a companion paper.

1. Introduction. In recent applications, images, as well as other processes, have been modeled by a Markov random field, i.e., as a Gibbs state with a finite range interaction, sometimes degraded by noise (see Comets and Gidas (1992), and references contained therein).

We are interested here in statistical inference, that is the estimation of parameters, for stationary Markov random fields. This section contains a brief discussion of the models together with some previous results, and with our findings; a more formal presentation is found in the rest of the paper.

We start with a single infinite black and white *image*, which is a specification of $+1$ (*black*) or -1(*white*) at each vertex (*pixel*) of an infinite lattice; the lattice we consider is $\mathbb{Z}^d$, and typically $d = 2$. The statistical properties of the image are described by a stationary Markov random field (with stationary interaction), which depends on some parameters $\theta_0 = (\theta_0(1), \ldots, \theta_0(s))$. A *noisy image* is obtained by independently flipping the sign (i.e., the color) of the image at each pixel with probability ϵ_0. The problem here is to estimate θ_0 and ϵ_0 by observing part of the noisy image, typically a finite rectangular array, with the sole a-priori knowledge that θ_0 belongs to some subset $\Theta \subseteq \mathbb{R}^s$ and that ϵ_0 is small — typically $\epsilon_0 < \frac{1}{2}$. (We also assume sufficient information to determine s exactly; for questions related to the estimation of s in the case with no noise see Gi and Seymour (1991) and Denny and Wright (1991).) More precisely, estimators $\hat{\theta}^{(\Lambda)}$ and $\hat{\epsilon}^{(\Lambda)}$, $\Lambda \subset \mathbb{Z}^d$, are functions of the noisy image in Λ, such that if $\{\Lambda_n\}_{n\in\mathbb{N}}$ is an increasing sequence of arrays whose union equals $\mathbb{Z}^d$, then $\hat{\theta}^{(\Lambda_n)} \to \theta_0$ and $\hat{\epsilon}^{(\Lambda_n)} \to \epsilon_0$ with probability one (with respect to the joint distribution, P_{ϵ_0,θ_0}, of the Markov random field and the noise).

Various estimators have been proposed, both for specific Markov random fields and for more general models. Two of these are the Maximum Likelihood estimator (Dempster et al. (1977), Geman and McClure (1985), Younes (1989)) and the Maximum Pseudo-Likelihood estimator Chalmond (1987), Younes (1991)) both of which are based on the EM algorithm; unfortunately, these estimators are obtained by iterative methods, requiring the simulation of a Markov random field at each iteration and resulting in a complex process. Other estimators are obtained by the methods of moments (Geman and McClure (1985) and Frigessi and Piccioni (1990)); these methods do not require any extra simulation, and are based on estimating various moments of the joint distribution P_{ϵ_0,θ_0} from the noisy image; some combinations of these moments turn out to be functions of θ_0 independent of ϵ_0, so they can be used to estimate θ_0 provided they are invertible; unfortunately, inverse functions cannot be easily produced even for the two-dimensional Ising model with zero external field, for which Onsager's exact solution of the model is available (one such inverse function was remarkably obtained in Frigessi and Piccioni (1990)), and seem out of reach, if they exist at all, for all Markov random fields with no exact solution (i.e. most of them, see Baxter (1982)).

Our paper presents a new estimator for ϵ_0, whose computation requires no extra simulations, no iterations, and which is (in principle) as easy and accurate as the solution of a moderate degree polynomial equation; an estimator of θ_0 is then obtained by a method analogous to the maximum pseudo-likelihood for the Markov random field alone (which is a very effective method, see Besag (1977)). The advantages of our method lie in potentially very simple estimations; it is possible that it may also lend itself to new proofs of the identifiability of parameters. This method, however, is not without its own difficulties. The problem is now reduced to (i) producing the above mentioned polynomial equation (whose form depends on the structure of the Markov random field, and on Θ) from the noisy image, and (ii) determining a-priori which one of the roots of this equation is an estimator for ϵ_0. In this paper, we describe how to produce suitable polynomial equations for any Markov random field (sections 2 and 3). Determination of the correct root, however, is more difficult: we have made some progress in the general case, but have had enough ideas, patience or computer power to complete this programme only in some limited cases (described in section 4 and in the companion paper, Barsky and Gandolfi (1993)). We now briefly outline the main ideas in the paper; a rigorous treatment starts afresh in section 2.

Our approach begins with the construction of the polynomial equations. We list the probabilities of all possible specifications of colors in some fixed finite array — as given by the Markov random field alone. This listing uses a great number of parameters: θ_0 and many other probabilities whose functional dependences on θ_0 are known only from the exact solution of the model, rarely available and in any case not used in this paper. Then we describe how these probabilities are transformed under the noise, introducing the parameter ϵ_0. Next, we invert this transformation (which amounts to the inversion of a large matrix), and apply the inverse transformation (parameterized by a new variable, ϵ) to a corresponding list of probabilities of patterns of colors in the noisy image. (An analogous matrix can be found in Meloche and Ruben (1992).) If $\epsilon = \epsilon_0$, the inversion procedure described above returns us to the original list of probabilities; but for other values of ϵ the process only gives a list of functions of ϵ (and ϵ_0 and θ_0). Using the structure of the Markov random field, we can indicate some necessary conditions, in the form of polynomial relations, which must be satisfied in order for such a list of functions to be the list of probabilities for a Markov random field. (Some similar notions can be found in Newman (1987).) Each of these necessary conditions provides a polynomial equation in ϵ, and $\epsilon = \epsilon_0$ is always a root. The idea is, therefore, to estimate from the data (i.e., the noisy image) the list of probabilities already transformed by the noise, apply the inverse transformation with the parameter ϵ to this observed list of empirical probabilities, and then solve one (or more) polynomial relations to determine for which value(s) of ϵ the inverse-transformed list satisfies some of the necessary condition(s) for being the list of probabilities for a Markov random field. Having found an estimator for ϵ_0, it is easy to "remove" the noise and use a maximum pseudo-likelihood method to estimate θ_0.

It is regrettable that we do not yet have a general method to indicate which real root of these polynomial equations is an estimator for ϵ. Some of the polynomial equations might even be identically zero, for some or for all $\theta \in \Theta$. Such equations are called *null-relations*, and we shall discard them; however, the null-relations depend on the specific models, and need to be identified on a case by case basis. The non null-relations, or *effective relations*, on the other hand, will generally have other roots besides $\epsilon = \epsilon_0$, and a-priori identification of the root estimating ϵ_0 again is done case by case. Some restrictions on ϵ_0 are demanded: for example, if the Markov random field has a global spin-flip symmetry, then ϵ_0 cannot be distinguished by $1 - \epsilon_0$ (which is reflected in the polynomial equations being invariant under the exchange $\epsilon \mapsto 1 - \epsilon$). Additionally, $\epsilon_0 = \frac{1}{2}$ cannot generally be identified (which is reflected in the above-mentioned matrix being singular when $\epsilon = \frac{1}{2}$); but even if $\epsilon_0 \in [0, \frac{1}{2})$, some polynomial equations have multiple real roots.

One possible solution to the problem of multiple roots is the simultaneous use of two or more equations, looking for common roots. However, when estimated from empirical data, such a set of equations would typically not have any common roots, and, at present, we have no good estimates on how far the roots of equations from the data can stray from their theoretical values. Additionally, it seems difficult to give a set of equations whose only common root is ϵ_0 for all $\theta_0 \in \Theta$.

In a different attempt to deal with the problem of multiple roots, we explicitly study our equations for the one-dimensional Ising models (i.e., Markov chains) in section 4, and for the two-dimensional Ising model in a companion paper, Barsky and Gandolfi (1993). There we find that for several of these equations (i) $\epsilon = \epsilon_0$ is a single root, and (ii) it is the smallest real root. It might be the case for every Markov random fields that there always are equations for which (i) and (ii) both occur, so we have formulated a theorem of consistency for the estimation in the context of this case, hoping that this will be the only consistency result required by the present theory.

The reader can now either turn directly to section 4 for a treatment (which we tried to make self-contained) of simple one-dimensional models, or else first read sections 2 and 3 for the abstract theory.

2. Definitions and the main result. Let $\mathbb{Z}^d$ be the *d-dimensional integer lattice*, and let $\Lambda \subset \mathbb{Z}^d$ be any *box* of the form $\Pi_{n=1}^{d}[i_n, i_n + j_n] \cap \mathbb{Z}^d$, for some $(i_1, \ldots, i_d) \in \mathbb{Z}^d$ and $(j_1, \ldots, j_d) \in \mathbb{Z}_+^d$. In the present paper our interest is focused on the *observed images* $y_\Lambda \in \{-1, 1\}^\Lambda$ which are the final result of some stochastic process. Our set of definitions is basically the description of this process.

Depending on the context, we indicate the *configuration space* $\{-1, 1\}^{\mathbb{Z}^d}$ by X, Y or Z; also, X_S, Y_S and Z_S will all indicate $\{-1, 1\}^S$, for $S \subseteq \mathbb{Z}^d$. We suppose that we are dealing with an *original image* $x \in X$, which has been corrupted by a *noise* $z \in Z$, resulting in an *observable image* $y \in Y$ given by $y_i = x_i \cdot z_i$ for all $i \in \mathbb{Z}^d$. Elements $i \in \mathbb{Z}^d$ are called *pixels*, and for any given pixel $i \in \mathbb{Z}^d$, x_i (respectively y_i) is called the *original* (resp. *observable*) *coloring* of i. The *observed image*, y_Λ, is the restriction of y to the box Λ. Apart from the distinction between observable and observed image, in the following we will often regard configurations in $\{-1, 1\}^{S_1}$ as the restrictions of configurations in $\{-1, 1\}^{S_2}$, if $S_1 \subseteq S_2 \subseteq \mathbb{Z}^d$.

The original image, the noise, and the observable image are described by some elements of the sets $\mathcal{P}_X$, $\mathcal{P}_Z$, and $\mathcal{P}_Y$ of the probability measures on the Borel σ-algebra of X, Z, and Y, respectively.

The original image. Let $\mathcal{C}$ be a *locally finite* (i.e., $|\{C \in \mathcal{C} : i \in C\}| < \infty$ for all $i \in \mathbb{Z}^d$, where $|A|$ denotes the cardinality of A) and *translation invariant* (i.e., $\tau_i(C) \in \mathcal{C}$ if $C \in \mathcal{C}$, where τ_i indicates the translation by the vector $i \in \mathbb{Z}^d$) collection of sets $C \subset \mathbb{Z}^d$ called *cliques*. Note that the local finiteness and translation invariance of $\mathcal{C}$ together imply that each clique is finite. An *interaction* ϕ based on $\mathcal{C}$ is a translation invariant real-valued function defined on $\cup_{C \in \mathcal{C}} X_C$. Let $\mathbf{0}$ indicate the origin of $\mathbb{Z}^d$, then the *local interactions* of ϕ are the entries of the vector $\{\phi(\eta)\}_{\eta \in X_C, \mathbf{0} \in C \in \mathcal{C}}$. We use these interactions to define Markov random fields. Later on we shall see that these models can be reparametrized using fewer parameters than the total number of local interactions; it will then be advantageous to use a different, but equivalent, notation.

For now, let a set of cliques, $\mathcal{C}$, and an interaction, ϕ, be fixed. For each finite $S \subset \mathbb{Z}^d$, the *energy function* $U_S^\phi : X_S \to \mathbb{R}$ is defined by

$$U_S^\phi(x) = \sum_{\substack{C \in \mathcal{C} \\ C \cap S \text{ nonempty}}} \phi(x_C). \tag{2.1}$$

Note that the energy function can be thought of as a linear combination, with integer coefficients, of the local interactions.

For $\Lambda \subset \mathbb{Z}^d$ and $\bar{x} \in X$, a *finite volume Markov random field* for $\mathcal{C}$ and ϕ in Λ with boundary condition $\bar{x}$ is the probability measure

$$\mu_{\Lambda,\bar{x}}(x_\Lambda) = Z_{\Lambda,\bar{x}}^{-1} e^{-U_\Lambda^\phi(x_\Lambda \vee \bar{x})} \tag{2.2}$$

where

$$Z_{\Lambda,\bar{x}} = \sum_{x_\Lambda \in X_\Lambda} e^{-U_\Lambda^\phi(x_\Lambda \vee \bar{x})}.$$

Here $x_\Lambda \vee \bar{x} \in X$ is the configuration which agrees with x_Λ in Λ and with $\bar{x}$ in $\mathbb{Z}^d \backslash \Lambda$. A *Markov random field* for $\mathcal{C}$ and ϕ is any convex combination μ_ϕ of weak limits of $\mu_{\Lambda,\bar{x}}$ as $\Lambda \uparrow \mathbb{Z}^d$, i.e., as Λ ranges over an increasing sequence of boxes which eventually covers the whole of $\mathbb{Z}^d$ (see Ruelle (1978), chapter 1). *Phase transition* occurs if there is more than one Markov random field for the given $\mathcal{C}$ and ϕ. In this paper, we only consider *translation invariant* Markov random fields; the set $\mathcal{M}_\phi$ of all such Markov random fields for the interaction ϕ is always nonempty (Ruelle (1978), Theorem 3.7). Since our estimation scheme begins with a single infinite (noisy) image, we may assume that μ_ϕ is *ergodic* for the group of translations of $\mathbb{Z}^d$ — as any original image is a typical configuration for some ergodic component of μ_ϕ.

Markov random fields μ_ϕ satisfy the *Markov property*: for any finite $S \subset \mathbb{Z}^d$ there exists a $T \supset S$ such that $\mu_\phi(x_S|\bar{x}) = \mu_\phi(x_S|\bar{x}_{T\backslash S})$ for $x_S \in X_S$ and $\bar{x} \in X$. In particular, for $S = \{\mathbf{0}\}$ one may take $T = \cup_{C:\mathbf{0}\in C\in\mathcal{C}} C$: we denote this particular set by $\bar{N}_\mathbf{0}$, and we call it the *complete neighborhood* of the origin $\mathbf{0}$. Configurations $\xi \in X_{\bar{N}_\mathbf{0}}$ are called *complete local patterns.* For $i \in \mathbb{Z}^d$, $\bar{N}_i$ will be the translation $\tau_i(\bar{N}_\mathbf{0})$ of $\bar{N}_\mathbf{0}$ by the vector i. The *neighborhood* of i is $N_i = \bar{N}_i \backslash \{i\}$. It follows from the translation invariance of $\mathcal{C}$ that $i \in N_\mathbf{0}$ if and only if $-i \in N_\mathbf{0}$; thus $|N_\mathbf{0}|$ is even. *Local patterns* are configurations $\xi \in X_{N_\mathbf{0}}$, and each such configuration gives rise to a pair of *local characteristics*

$$\pi_\phi(x_\mathbf{0}|\xi) = \mu_\phi(x_\mathbf{0}|\xi) = Z_{\mathbf{0},\xi}^{-1} e^{-U_\mathbf{0}^\phi(x_\mathbf{0} \vee \xi)}. \tag{2.3}$$

(Our notation, here and elsewhere, in writing $U_\mathbf{0}^\phi$ (and $Z_{\mathbf{0},\xi}$) is that when S is the singleton, $\{\mathbf{0}\}$, we write $S = \mathbf{0}$ as an abbreviation.)

Note that local characteristics are functions of the $\sum_{C:\mathbf{0}\in C\in\mathcal{C}} 2^{|C|}$ local interactions. Moreover, the local characteristics are functions of ϕ independent of the specific $\mu_\phi \in \mathcal{M}_\phi$. Also, since the local characteristics are always strictly positive, the Markov property implies that $\mu_\phi(x_S) > 0$ for all $x_S \in X_S$, for every finite $S \subset \mathbb{Z}^d$.

The noise and the observable image. For some $\epsilon \in [0,1]$, the statistical properties of the configurations $z \in Z$ are described by the Bernoulli probability measure $\nu_\epsilon = \Pi_{i\in\mathbb{Z}^d} \nu_{\epsilon,i}$, defined on the Borel σ-algebra of Z, where $\nu_{\epsilon,i}(z_i = 1) = \epsilon = 1 - \nu_{\epsilon,i}(z_i = -1)$. The action of the noise is given by setting $y_i = x_i \cdot z_i$; this amounts to flipping each pixel with probability ϵ, independently of the other pixels and of x. For any given interaction ϕ and noise level ϵ, the joint probability measure $P_{\phi,\epsilon} = \mu_\phi \cdot \nu_\epsilon$, defined on the Borel σ-algebra of Y, describes the statistical properties of the observable image. Eventually, the interactions ϕ will be parameterized by a vector θ — then we will write $P_{\theta,\epsilon}$ for the joint measure.

Estimation of parameters. Suppose now that the single infinite black and white observable image $y \in Y$ is fixed, and we observe y_Λ as $\Lambda \uparrow \mathbb{Z}^d$. The statistical properties of y are described by P_{ϕ_0,ϵ_0} for a known $\mathcal{C}$, but with both ϵ_0 and ϕ_0 unknown. We want to estimate ϵ_0 and ϕ_0, but this is only feasible if both are identifiable. The restriction of ϵ_0 to $[0, \frac{1}{2})$ is sufficient for identification of ϵ_0, but ϕ_0 can only be identified modulo the following equivalence relation (see also Gidas (1987), Appendix, for a related discussion).

Two interactions ϕ_1 and ϕ_2 are *equivalent*, $\phi_1 \approx \phi_2$, if $\mathcal{M}_{\phi_1} = \mathcal{M}_{\phi_2}$ (or, equivalently, if $\mathcal{M}_{\phi_1}$ and $\mathcal{M}_{\phi_2}$ have nonempty intersection; see Gidas (1991)). In our setting, the equivalence relation is better described by the following lemma. Note that, by translation invariance, an interaction ϕ is identified by the

$$r = \sum_{C:\mathbf{0}\in C\in\mathcal{C}} 2^{|C|}$$

local interactions, so it can be treated as a vector in $\mathbb{R}^r$. The energies corresponding to the choice of any s interactions $\bar{\phi}_1, \ldots, \bar{\phi}_s \in \mathbb{R}^r$ can be written as a function $U_\mathbf{0} = (U_\mathbf{0}^{\bar{\phi}_1}, \ldots, U_\mathbf{0}^{\bar{\phi}_s})$ defined on $X_{\bar{N}_\mathbf{0}}$ and taking values in $\mathbb{R}^s$.

Lemma 1. a. *Two interactions ϕ_1 and ϕ_2 are equivalent iff*

$$U_\mathbf{0}^{\phi_1}(x_\mathbf{0} \vee \xi) - U_\mathbf{0}^{\phi_1}(-x_\mathbf{0} \vee \xi) = U_\mathbf{0}^{\phi_2}(x_\mathbf{0} \vee \xi) - U_\mathbf{0}^{\phi_2}(-x_\mathbf{0} \vee \xi) \tag{2.4}$$

for all $\xi \in X_{N_\mathbf{0}}$. The two interactions are also equivalent iff

$$U_\mathbf{0}^{\phi}(x_\mathbf{0} \vee \xi) - U_\mathbf{0}^{\phi}(-x_\mathbf{0} \vee \xi) = 0 \tag{2.5}$$

for all $\xi \in X_{N_\mathbf{0}}$, with $\phi = \phi_1 - \phi_2$. In both (2.4) and (2.5), $x_\mathbf{0}$ can be taken to be either $+1$ or -1.
b. *Define $\mathcal{N} = \{\phi \in \mathbb{R}^r : \phi$ satisfies (2.5) for all $\xi \in X_{N_\mathbf{0}}\}$, and let $\bar{\phi}_1, \ldots, \bar{\phi}_s \in \mathbb{R}^r$ be a basis of some linear space $\mathcal{S}$ which is linearly independent of $\mathcal{N}$. If $\theta = (\theta_1, \ldots, \theta_s)$, and $\theta \cdot U_\mathbf{0}$ is the standard inner product of θ and $U_\mathbf{0}$ in $\mathbb{R}^s$, then*

$$\theta \cdot U_\mathbf{0} = U_\mathbf{0}^{(\sum_{i=1}^s \theta_i \bar{\phi}_i)}, \tag{2.6}$$

$$\theta \cdot [U_\mathbf{0}(x_\mathbf{0} \vee \xi) - U_\mathbf{0}(-x_\mathbf{0} \vee \xi)] = 0 \text{ for all } \xi \in X_{N_\mathbf{0}} \text{ iff } \theta = 0, \tag{2.7}$$

and

$$s \leq \min\{2^{|N_\mathbf{0}|}, r\}. \tag{2.8}$$

Proof: Two interactions ϕ_1 and ϕ_2 are equivalent iff all of the finite volume Markov random fields for ϕ_1 and ϕ_2 coincide, and this holds iff all ratios

$$\frac{e^{-U_\Lambda^\phi(x_\Lambda \vee \bar{x})}}{e^{-U_\Lambda^\phi(y_\Lambda \vee \bar{x})}} \tag{2.9}$$

are the same for ϕ_1 and ϕ_2, whenever $\Lambda \subset \mathbb{Z}^d$, $\bar{x} \in X_{\mathbb{Z}^d \setminus \Lambda}$, and $x_\Lambda, y_\Lambda \in X_\Lambda$. It may further be assumed in (2.9) that x_Λ and y_Λ differ in exactly one pixel. Now all ratios (2.9) are of the form

$$e^{-U_\mathbf{0}^\phi(x_\mathbf{0} \vee \xi) + U_\mathbf{0}^\phi(-x_\mathbf{0} \vee \xi)}$$

for some $\xi \in X_{N_0}$ and $x_0 = \pm 1$, so $\phi_1 \approx \phi_2$ iff (2.4) holds for all $\xi \in X_{N_0}$. Note that U_0^{ϕ} is linear in $\phi \in \mathbb{R}^r$, so that (2.5) holds (with $\phi = \phi_1 - \phi_2$) whenever (2.4) is satisfied. This linearity also yields (2.6), as $\theta \cdot U_0 = \sum_{i=1}^{s} \theta_i U_0^{\bar{\phi}_i} = U_0^{(\sum_{i=1}^{s} \theta_i \bar{\phi}_i)}$. Moreover, $\theta \cdot [U_0(x_0 \vee \xi) - U_0(-x_0 \vee \xi)] = 0$ for all $\xi \in X_{N_0}$ iff $U_0^{(\sum \theta_i \bar{\phi}_i)}(x_0 \vee \xi) - U_0^{(\sum \theta_i \bar{\phi}_i)}(-x_0 \vee \xi) = 0$ for all $\xi \in X_{N_0}$ iff $\sum_{i=1}^{s} \theta_i \bar{\phi}_i \approx 0$ (the interaction which is identically zero) iff $\theta = 0$ (the zero vector in $\mathbb{R}^s$), which proves (2.7). Finally, the nontrivial part of (2.8) (that $s \leq 2^{|N_0|}$) follows from the facts that the dimension of $\mathcal{S}$ cannot exceed the number of linearly independent equations of type (2.5), and that there are only $2^{|N_0|}$ linear relations of this type — before checking for linear independence. ∎

Later on we show that actually a strict inequality holds in (2.8). It is now convenient to fix a basis $\bar{\phi}_1, \ldots, \bar{\phi}_s$ of a maximal linear space, $\mathcal{S} \subset \mathbb{R}^r$, which is independent of $\mathcal{N}$, and to replace the energy in (2.1) - (2.3) by $\theta \cdot U_\Lambda(x_\Lambda \vee \bar{x})$, with $\theta \in \mathbb{R}^s$ and $U_\Lambda = (U_\Lambda^{\bar{\phi}_1}, \ldots, U_\Lambda^{\bar{\phi}_s})$. The parameters $\theta \in \mathbb{R}^s$ are now identifiable, and they will replace ϕ in our various notations: μ_ϕ and $P_{\phi,\epsilon}$ become μ_θ and $P_{\theta,\epsilon}$, respectively. We may also assume, for simplicity, that each $\bar{\phi}_i$ has integer entries.

Fix $\Theta \subseteq \mathbb{R}^s$, $\theta_0 \in \Theta$ and $\epsilon_0 \in [0, \frac{1}{2})$. We want to define functions $\hat{\epsilon}^{(\Lambda)}(y)$ and $\hat{\theta}^{(\Lambda)}(y)$ such that

$$\hat{\epsilon}^{(\Lambda)}(y) \to \epsilon_0 \text{ as } \Lambda \uparrow \mathbb{Z}^d \tag{2.10}$$

and

$$\hat{\theta}^{(\Lambda)}(y) \to \theta_0 \text{ as } \Lambda \uparrow \mathbb{Z}^d \tag{2.11}$$

for P_{θ_0,ϵ_0}-almost all $y \in Y$.

The reason for not necessarily taking $\Theta = \mathbb{R}^s$ is that the extra information provided by the knowledge of Θ can make the estimation of ϵ_0 easier, as will be seen in the Consistency Theorem.

The main quantity which we will be estimating from the data is most conveniently introduced as a function of the empirical process generated by an original (resp., observable) image, or more generally, as a function of probability measures in $\mathcal{P}_X$ (resp., $\mathcal{P}_Y$). For each box $\Lambda \subset \mathbb{Z}^d$ and image $x \in X$ (resp., $y \in Y$), define $x^{(\Lambda)} \in X$ (resp., $y^{(\Lambda)} \in Y$) to be the periodic extension of the restricted image x_Λ (resp., y_Λ). The *empirical process* $R_{\Lambda,x} \in \mathcal{P}_X$ is defined by

$$R_{\Lambda,x}(f) = \frac{1}{|\Lambda|} \sum_{i \in \Lambda} f(\tau_{-i} x^{(\Lambda)})$$

for all continuous functions $f : X \to \mathbb{R}$; the empirical process $R_{\Lambda,y}$ is similarly defined. For $P \in \mathcal{P}_X$ (resp., $P \in \mathcal{P}_Y$), let M_P be the vector whose $2^{|\bar{N}_0|}$ entries, indexed by the complete local patterns $\bar{\xi} \in X_{\bar{N}_0}$ (resp., $\bar{\xi} \in Y_{\bar{N}_0}$), are given by

$$M_P(\bar{\xi}) = E^P(\mathbf{1}[x_{\bar{N}_0} = \bar{\xi}]),$$

where E^P indicates the expectation with respect to P, and $\mathbf{1}$ is the indicator function of the event in the brackets. The components of M are thus the probabilities of the various local patterns; in the case of the empirical processes, the components are just the relative frequencies in some portion of the image. For simplicity we use the notation

$$M_{\Lambda,x} = M_{R_{\Lambda,x}},$$

$$M_{\Lambda,y} = M_{R_{\Lambda,y}}, \tag{2.12}$$

and

$$M_\theta = M_{\mu_\theta}, \tag{2.13}$$

where in (2.13), μ_θ is some Markov random field for the interaction θ — we suppress the dependence of M_θ on μ_θ as we shall eventually work (see Lemma 3 below) with properties of the vector which are independent of the particular choice of the measure $\mu_\theta \in \mathcal{M}_\theta$. Additionally,

$$M_{\theta,\epsilon} = M_{P_{\theta,\epsilon}} = M_\theta A_\epsilon,$$

where the second equality makes reference to the $2^{|\bar{N}_0|} \times 2^{|\bar{N}_0|}$ matrix A_ϵ defined in (2.14) below.

We comment here on our vector notation. We generally do not distinguish between row vectors and column vectors, and use whichever notation seems to be most natural for the purpose at hand, as it will usually be clear from the context which is meant. For example, in writing $M_\theta A_\epsilon$ above, we are using the usual probabilistic notation and regarding M_θ as a row vector. Later, in section 4, part (I), it will be equally evident that the vectors ϕ and $\mathbf{U}(\phi)$ are column vectors and that we are using the usual linear-algebraic notation in writing $\mathbf{U}(\phi) = \mathbf{U}\phi$.

The entries of A_ϵ are indexed by $\bar{\xi}^{(1)} \in X_{\bar{N}_0}$, $\bar{\xi}^{(2)} \in Y_{\bar{N}_0}$ and given by

$$A_\epsilon(\bar{\xi}^{(1)}, \bar{\xi}^{(2)}) = \epsilon^D (1-\epsilon)^{(|\bar{N}_0|-D)}, \tag{2.14}$$

where $D = D(\bar{\xi}^{(1)}, \bar{\xi}^{(2)}) = \sum_{i\in\bar{N}_0} \frac{1}{2}|\bar{\xi}_i^{(1)} - \bar{\xi}_i^{(2)}|$ is the Hamming distance between $\bar{\xi}^{(1)}$ and $\bar{\xi}^{(2)}$. As shown below in Lemma 5, A_ϵ is invertible for $\epsilon \neq \frac{1}{2}$, and its inverse A_ϵ^{-1} has components

$$A_\epsilon^{-1}(\bar{\xi}^{(1)}, \bar{\xi}^{(2)}) = (2\epsilon-1)^{-|\bar{N}_0|} \epsilon^D (\epsilon-1)^{(|\bar{N}_0|-D)}.$$

(Properties of related matrices appear in Barsky (1993) and Meloche and Ruben (1992).)

The estimator $\hat{\epsilon}^{(\Lambda)}$ in (2.10) is one of the roots of a polynomial equation in ϵ constructed by relating the entries of $M_{\Lambda,y}A_\epsilon^{-1}$ to the probabilities of complete local patterns in μ_{θ_0}. After this, $\hat{\epsilon}^{(\Lambda)}$ is used to remove the noise from the data so that an estimator $\hat{\theta}^{(\Lambda)}$ satisfying (2.11) can be determined.

Before we can give the exact form of the polynomial equations, it is necessary to study Markov random fields in greater detail.

Structure of Markov random fields. We shall eventually see (in Lemma 3 below), that it is possible to produce $2^{|N_0|} - s$ polynomial equations (although several of these may be null). To be certain that we have any equations at all we must first show that the upper bound for s given in (2.8) can be improved.

Lemma 2. *Let $\mathcal{C}$ be a locally finite, translation invariant set of cliques, not all of size one. Then if $\mathcal{S}$ is as in Lemma 1,*

$$s = \dim(\mathcal{S}) \leq \sum_{C:0\in C\in\mathcal{C},|C|\geq 2} \frac{1}{|C|}(2^{|C|}-1) < 2^{|N_0|}. \tag{2.15}$$

Proof: From (2.8) we have that $s \leq r = \sum_{C:0\in C\in\mathcal{C}} 2^{|C|}$, but s is in fact smaller for three reasons. We present these reasons as linear conditions which can (actually, the first condition must) be satisfied by the vectors of $\mathcal{S}$, thereby giving successive upper bounds to s.

In the first place, interactions are translation invariant. If τ_i stands both for the translation by the vector $\in \mathbb{Z}^d$ and for the map induced by this translation on the configurations of X, then interactions satisfy $\phi(\tau_i\eta_C) = \phi(\eta_C)$, for all $\eta_C \in X_C$ and $C \in \mathcal{C}$. Roughly speaking, the translation invariance implies that a clique C and all of its translates can contribute at most $2^{|C|}$ parameters to the sum which is the dimension of $\mathcal{S}$. More formally,

$$s \leq \sum_{C:0\in C\in\mathcal{C}} \frac{1}{|C|} 2^{|C|}.$$

In the second place, for each $C \in \mathcal{C}$ and for any interaction ϕ based on C, we may assume that $\phi(\bar{\eta}) = 0$ for at least one $\bar{\eta} \in X_C$. In fact, for any fixed $\bar{\eta} \in X_C$ (with $C \in \mathcal{C}$), we can define

$$\bar{\phi}(\eta) = \begin{cases} \phi(\eta) - \phi(\bar{\eta}) & \text{for } \eta \in X_C \\ \phi(\eta) & \text{for } \eta \in X_{C'}, C' \neq C. \end{cases}$$

Then $\bar{\phi} \approx \phi$ by (2.4), which shows that

$$s \leq \sum_{C:0\in C\in\mathcal{C}} \frac{1}{|C|}(2^{|C|}-1).$$

Finally, we can also assume that $\phi(\eta) = 0$ for all $\eta \in X_C$ whenever $|C| = 1$. Indeed, we have just shown above that if $|C| = 1$, then it may be assumed that (acting on the color at that single pixel) $\phi(-1) = 0$. Assuming that $\phi(+1) \neq 0$, and that there is a clique $\tilde{C} \in \mathcal{C}$ with $|\tilde{C}| \geq 2$, define

$$\tilde{\phi}(\eta) = \begin{cases} \phi(\eta) + \phi(+1)|\{i \in \tilde{C} : \eta_i = +1\}|/|\tilde{C}| & \text{for } \eta \in X_{\tilde{C}} \\ 0 & \text{if } \eta \in X_C, |C| = 1 \\ \phi(\eta) & \text{otherwise.} \end{cases}$$

By (2.4), $\tilde{\phi} \approx \phi$, which concludes the proof of the first inequality in (2.15).

We next prove the second inequality in (2.15) for $d = 1$; afterwards we will show how to reduce higher dimensional cases to this setting. Given $\mathcal{C}$ and its associated complete neighborhood of $\mathbf{0}$, $\bar{N}_{\mathbf{0}} = \bar{N}_{\mathbf{0}}(\mathcal{C}) = \{i_{\frac{-|N_{\mathbf{0}}|}{2}}, i_{\frac{-|N_{\mathbf{0}}|}{2}+1}, \ldots, i_{\frac{|N_{\mathbf{0}}|}{2}}\}$, define $\psi : \bar{N}_{\mathbf{0}} \to \bar{N} = \{\frac{-|N_{\mathbf{0}}|}{2}, \frac{-|N_{\mathbf{0}}|}{2} + 1, \ldots, \frac{|N_{\mathbf{0}}|}{2}\}$ by $\psi(i_n) = n$. In particular, note that $\psi(i_0) = \mathbf{0}$ since $i_0 = \mathbf{0}$. Now define $\mathcal{C}' = \{k + \psi(C) : k \in \mathbb{Z}, \mathbf{0} \in C \in \mathcal{C}\}$. It is readily seen that $\mathcal{C}'$ is a locally finite, translation invariant set of cliques and that ψ induces a cardinality-preserving bijection between $\{C \in \mathcal{C} : \mathbf{0} \in C\}$ and $\{C' \in \mathcal{C}' : \mathbf{0} \in C'\}$. Therefore,

$$\sum_{C:\mathbf{0}\in C\in\mathcal{C},|C|\geq 2} \frac{1}{|C|}(2^{|C|} - 1) = \sum_{C':\mathbf{0}\in C'\in\mathcal{C}',|C'|\geq 2} \frac{1}{|C'|}(2^{|C'|} - 1).$$

Since $\bar{N}_{\mathbf{0}}(\mathcal{C}') = \bar{N}$, it now suffices to prove that the second inequality in (2.15) holds for all $\mathcal{C}$ with $\bar{N}_{\mathbf{0}}(\mathcal{C}) = \bar{N}$. For such a $\mathcal{C}$, observe that the subset of those cliques containing $\mathbf{0}$ can be partitioned into equivalence classes by declaring a pair of cliques to be equivalent if they are translates of one another. Each equivalence class has a "least" representative: a clique for which the origin is the maximal pixel. Let $\mathcal{D}$ denote the set of least representatives for the equivalence classes of those cliques C having $\mathbf{0} \in C$ and $|C| \geq 2$. Then

$$\sum_{C:\mathbf{0}\in C\in\mathcal{C},|C|\geq 2} \frac{1}{|C|}(2^{|C|} - 1) = \sum_{C\in\mathcal{D}} (2^{|C|} - 1).$$

The number of least representatives in $\mathcal{D}$ having cardinality n cannot exceed $\binom{|N_{\mathbf{0}}|/2}{n-1}$ since any such representative must contain the origin and exactly $n - 1$ pixels in $\{\frac{-|N_{\mathbf{0}}|}{2}, \frac{-|N_{\mathbf{0}}|}{2} + 1, \ldots, -1\}$. Thus

$$\begin{aligned}\sum_{C\in\mathcal{D}} (2^{|C|} - 1) &= \sum_{n=2}^{N_{\mathbf{0}}/2+1} \binom{|N_{\mathbf{0}}|/2}{n-1} (2^n - 1) \\ &= 2 \cdot 3^{|N_{\mathbf{0}}|/2} - 2^{|N_{\mathbf{0}}|/2} - 1 \\ &< 2^{|N_{\mathbf{0}}|},\end{aligned}$$

for $|N_{\mathbf{0}}| = 2, 4, \ldots$.

For the higher dimensional cases, we construct a map $\psi : \mathbb{Z}^d \to \mathbb{Z}$ which induces a cardinality-preserving bijection between $\{C \in \mathcal{C} : \mathbf{0} \in C\}$ and $\{C' \in \mathcal{C}' : \mathbf{0} \in C'\}$, where $\mathcal{C}'$ is some locally finite, translation invariant collection of cliques in $\mathbb{Z}$. Given $\mathcal{C}$, we note that it is always possible to choose integers $n_1, \ldots, n_d$ so that $\psi(i_1, \ldots, i_d) = \sum_{r=1}^{d} n_r i_r$ is injective on $\bar{N}_{\mathbf{0}}(\mathcal{C})$. The collection $\mathcal{C}' = \{k + \psi(C) : k \in \mathbb{Z}, \mathbf{0} \in C \in \mathcal{C}\}$ is a translation invariant, locally finite set of cliques in $\mathbb{Z}$ with $N_{\mathbf{0}}(\mathcal{C}') = \psi(N_{\mathbf{0}}(\mathcal{C}))$, and hence $|N_{\mathbf{0}}(\mathcal{C}')| = |N_{\mathbf{0}}(\mathcal{C})|$. Therefore,

$$\begin{aligned}\sum_{C:\mathbf{0}\in C\in\mathcal{C},|C|\geq 2} \frac{1}{|C|}(2^{|C|} - 1) &\leq \sum_{C':\mathbf{0}\in C'\in\mathcal{C}',|C'|\geq 2} \frac{1}{|C'|}(2^{|C'|} - 1) \\ &< 2^{|N_{\mathbf{0}}(\mathcal{C}')|} \\ &= 2^{|N_{\mathbf{0}}(\mathcal{C})|}\end{aligned}$$

as desired. ■

It is clear that $\dim(\mathcal{S})$ might even be smaller than proven in Lemma 2. For instance, the first inequality in (2.15) yields $s \leq 3$ for the one-dimensional nearest-neighbor Ising model; but it is easy to verify that actually $s = 2$, with the two parameters corresponding to the two parameters of the 2×2 stochastic matrix defining the one-dimensional two state Markov chain. Also, the choice of the parameters made in this proof might not be the most convenient for a physical description of the model; in the Ising model, for example, one often uses the external field, $h = \frac{1}{2}(\phi(+1) + \phi(-1))$, as a parameter.

We now fix $\mathcal{C}$ and Θ, and turn to the issue of showing that the probabilities of the local patterns in a Markov random field satisfy certain polynomial equations. A *polynomial relation* for the vector $M = (M(\bar{\xi}))_{\bar{\xi} \in X_{\bar{N}_0}}$ is any homogeneous polynomial $Q = Q(M)$ of the $2^{\bar{N}_0}$ entries of M. Some particularly relevant polynomial relations are of the form

$$Q_\alpha(M) = \prod_{\xi \in X_{N_0}} [M(+1 \vee \xi)]^{\alpha_+(\xi)} [M(-1 \vee \xi)]^{-\alpha_-(\xi)} - \prod_{\xi \in X_{N_0}} [M(-1 \vee \xi)]^{\alpha_+(\xi)} [M(+1 \vee \xi)]^{-\alpha_-(\xi)}, \tag{2.16}$$

where $\alpha = (\alpha(\xi))_{\xi \in X_{N_0}}$ is a vector with integer entries, $\alpha_+(\xi) = \max\{\alpha(\xi), 0\}$, and $\alpha_-(\xi) = \min\{\alpha(\xi), 0\}$. We are especially interested in polynomial relations for the vector M_θ defined in (2.13); some of these are described in the next lemma.

Lemma 3. *Given $\mathcal{C}$ and Θ, there exist linearly independent integer vectors $\alpha^{(n)} \in \mathbb{Z}^{2^{|N_0|}}$, $n = 1, \ldots, 2^{|N_0|} - s$, each with entries indexed by $\xi \in X_{N_0}$, such that*

$$Q_{\alpha^{(n)}}(M_\theta) = 0 \tag{2.17}$$

for all $\theta \in \Theta$, and for all Markov random fields $\mu_\theta \in \mathcal{M}_\theta$.

Proof: A simple computation using (2.16), (2.13) and (2.3) shows that for any $\theta \in \mathbb{R}^s$ (and any $\mu_\theta \in \mathcal{M}_\theta$), $Q_\alpha(M_\theta) = 0$ iff

$$\sum_{\xi \in X_{N_0}} \alpha(\xi) \log \frac{M_\theta(+1 \vee \xi)}{M_\theta(-1 \vee \xi)} = 0$$

iff

$$\sum_{\xi \in X_{N_0}} \alpha(\xi)[U_0^\theta(+1 \vee \xi) - U_0^\theta(-1 \vee \xi)] = 0. \tag{2.18}$$

Asking that (2.18) be satisfied for *all* $\theta \in \mathbb{R}^s$ is the same as seeking solutions to

$$\alpha \mathbf{U} = 0$$

where $\mathbf{U}$ is the $2^{|\mathcal{N}_0|} \times s$ matrix having entries

$$\mathbf{U}(\xi, i) = U_{\mathbf{0}}^{\bar{\phi}_i}(+1 \vee \xi) - U_{\mathbf{0}}^{\bar{\phi}_i}(-1 \vee \xi) \tag{2.19}$$

for $\xi \in X_{\mathcal{N}_0}$ and $i \in \{1, \ldots, s\}$. By the assumption that $\bar{\phi}_1, \ldots, \bar{\phi}_s$ is a basis for a subspace of $\mathbb{R}^r$ which is linearly independent of $\mathcal{N}$, we know that the columns of $\mathbf{U}$ must be linearly independent. Since rank$(\mathbf{U}) = s$, we have that the nullity of the mapping $\alpha \mapsto \alpha\mathbf{U}$ is $2^{|\mathcal{N}_0|} - s \geq 1$, where the inequality follows from Lemma 2. Finally, as the entries of $\mathbf{U}$ are all integers, it is possible to find a basis $\{\alpha^{(n)}\}_{n=1}^{2^{|\mathcal{N}_0|}-s}$ for the nullspace of this mapping with each $\alpha^{(n)}$ having only integer entries. ■

Definition of an estimator. In the estimation of ϵ_0, we will use some polynomial equations of the form $Q(M) = 0$ — provided they satisfy some suitable conditions. To better understand these conditions, let us consider the application of the relations found in Lemma 3 to the vector

$$M_{\theta,\epsilon',\epsilon} = M_{\theta,\epsilon'} A_\epsilon^{-1} = M_\theta A_{\epsilon'} A_\epsilon^{-1}. \tag{2.20}$$

The case when $\epsilon' = \epsilon_0$ is of special interest; also setting $\epsilon = \epsilon_0$ shows that $M_{\theta,\epsilon_0,\epsilon_0}$ is just a scalar multiple of M_θ, and hence $\epsilon = \epsilon_0$ is a solution of

$$Q(M_{\theta,\epsilon',\epsilon}) = 0 \tag{2.21}$$

when $\epsilon' = \epsilon_0$.

As already mentioned in the introduction, for any $\epsilon' = \epsilon_0 \in [0, \frac{1}{2})$, there will typically be several values of ϵ satisfying (2.21). The examples presented in section 4 suggest that perhaps it is always possible to find equations (2.21) for which ϵ_0 is the smallest real root in ϵ when $\epsilon' = \epsilon_0$. Motivated by these examples, we will take our "suitable conditions" on the relations (2.17) to be some variant of requiring that ϵ_0 is the smallest real root in ϵ of (2.21) when $\epsilon' = \epsilon_0$. Actually, it will turn out to be sufficient to require that 0 is a single root in ϵ, and the only root in $(-\infty, 0]$, of (2.21) when $\epsilon' = 0$. We mention here, primarily to establish terminology which will be used in the treatment of the examples (see section 4), that one way for a relation Q to fail to satisfy this condition is for $Q(M_{\theta,\epsilon_0,\epsilon})$ to be identically zero in ϵ for some $\theta \in \Theta$ and $\epsilon_0 \in [0, \frac{1}{2})$; whenever this occurs, we say that Q is a *null-relation.*

If we already knew the interaction θ_0, then we could further specialize (2.21) to the polynomial equation $Q(M_{\theta_0,\epsilon',\epsilon}) = 0$ for $\epsilon' = \epsilon_0$ (or 0). However, it is because we suppose ourselves to be initially ignorant of the choice of $\theta_0 \in \Theta$ that we use the relations of Lemma 3 which are valid for *all* $\theta \in \Theta$. Actually, an examination of the proof of Lemma 3 shows that the relations (2.17) are satisfied for all $\theta \in \mathbb{R}^s$ — the reason for keeping track of Θ is that it is simpler to verify the above-mentioned suitable conditions if we are allowed to use the a-priori knowledge that the interaction belongs to the region Θ in the interaction-space $\mathbb{R}^s$. In fact, when we subsequently turn to the estimation of θ_0, we will often make that estimation in the context of the larger space, $\mathbb{R}^s$.

Consistency Theorem. *Let C and Θ be given, let $\epsilon_0 \in [0, \frac{1}{2})$ and $\theta_0 \in \Theta$, and let Q be a polynomial relation. Suppose that $Q(M_\theta) = 0$ for all $\theta \in \Theta$, and define*

$$M_{\theta,\epsilon',\epsilon} = M_\theta A_{\epsilon'} A_\epsilon^{-1}.$$

Further suppose that, for all $\theta \in \Theta$, the equation

$$Q(M_{\theta,0,\epsilon}) = 0 \tag{2.22}$$

admits $\epsilon = 0$ as a single root and has no real roots in $(-\infty, 0)$. Then for any $\gamma > 0$, (I) the smallest real root in $[-\gamma, 1]$, $\hat{\epsilon}^{(\Lambda)}(y)$ say, of

$$Q(M_{\Lambda,y,\epsilon}) = 0, \tag{2.23}$$

where

$$M_{\Lambda,y,\epsilon} = M_{\Lambda,y} A_\epsilon^{-1},$$

is such that

$$\lim_{\Lambda \uparrow \mathbb{Z}^d} \hat{\epsilon}^{(\Lambda)}(y) = \epsilon_0,$$

and (II) the vector $\hat{\theta}^{(\Lambda)}(y)$ which maximizes (in $\mathbb{R}^s$)

$$GMPL(\theta, \Lambda, y) = \prod_{(x_{\mathbf{0}} \vee \xi) \in X_{\bar{N}_{\mathbf{0}}}} [\pi_\theta(x_{\mathbf{0}} | \xi)]^{M_{\Lambda,y,\hat{\epsilon}^{(\Lambda)}(y)}(x_{\mathbf{0}} \vee \xi)} \tag{2.24}$$

converges pointwise to θ_0 (i.e., $\lim_{\Lambda \uparrow \mathbb{Z}^d} \hat{\theta}^{(\Lambda)}(y)(\eta) = \theta_0(\eta)$ for all $\eta \in X_C$ and for all $C \in \mathcal{C}$) with P_{θ_0,ϵ_0}-probability one.

3. Consistency. The Consistency Theorem is proven in this section after three introductory lemmas concerning the matrix A_ϵ; the first of these lemmas indicates that the noise transformation is well described by A_{ϵ_0}. We choose a collection of cliques, $\mathcal{C}$, and a set of interaction parameters, $\Theta \subset \mathbb{R}^s$, which are fixed throughout this section.

For $x \in X$, $z \in Z$ and $\Lambda \subset \mathbb{Z}^d$, define the matrix

$$A_{\Lambda,x,z}(\bar{\xi}^{(1)}, \bar{\xi}^{(2)}) = \begin{cases} [|\Lambda| M_{\Lambda,x}(\bar{\xi}^{(1)})]^{-1} \sum_{i \in \Lambda} \mathbf{1}[x_{\bar{N}_i}^{(\Lambda)} = \bar{\xi}^{(1)}, (x \cdot z)_{\bar{N}_i}^{(\Lambda)} = \bar{\xi}^{(2)}] & \text{if } M_{\Lambda,x}(\bar{\xi}^{(1)}) \neq 0 \\ 0 & \text{if } M_{\Lambda,x}(\bar{\xi}^{(1)}) = 0. \end{cases} \tag{3.1}$$

Note that if $y = x \cdot z$ then

$$M_{\Lambda,y} = M_{\Lambda,x} \cdot A_{\Lambda,x,z}. \tag{3.2}$$

Lemma 4. *For all $\theta_0 \in \Theta$, $\epsilon_0 \in [0, 1]$ and $\mu_{\theta_0} \in \mathcal{M}_{\theta_0}$,*

$$\lim_{\Lambda \uparrow \mathbb{Z}^d} A_{\Lambda,x,z} = A_{\epsilon_0} \tag{3.3}$$

with P_{θ_0,ϵ_0}-probability one.

Proof: Let $x \in X$, $\bar{\xi}^{(1)} \in X_{\bar{N}_0}$ and define $S_x(\bar{\xi}^{(1)}) = \{i \in \mathbb{Z}^d : x_{\bar{N}_i} = \bar{\xi}^{(1)}\}$. As ν_{ϵ_0} is a Bernoulli measure and $\mu_{\theta_0}(\bar{\xi}^{(1)}) > 0$, for any $\bar{\xi}^{(3)} \in Z_{\bar{N}_0}$, we have

$$\lim_{\Lambda\uparrow\mathbb{Z}^d} |S_x(\bar{\xi}^{(1)}) \cap \Lambda|^{-1} \sum_{i\in S_x(\bar{\xi}^{(1)})\cap\Lambda} \mathbf{1}[z_{\bar{N}_i} = \bar{\xi}^{(3)}] \quad = \quad \nu_{\epsilon_0}(\bar{\xi}^{(3)}) \tag{3.4}$$

with ν_{ϵ_0}-probability one. Since we assumed that μ_{θ_0} is ergodic,

$$\lim_{\Lambda\uparrow\mathbb{Z}^d} |\Lambda|^{-1}|S_x(\bar{\xi}^{(1)}) \cap \Lambda| = \lim_{\Lambda\uparrow\mathbb{Z}^d} M_{\Lambda,x}(\bar{\xi}^{(1)}) = M_{\theta_0}(\bar{\xi}^{(1)})$$

with μ_{θ_0}-probability one. Therefore, writing $\bar{\xi}^{(2)} = \bar{\xi}^{(1)} \cdot \bar{\xi}^{(3)}$ and letting $y_{\bar{N}_i}^{(\Lambda)} = (x \cdot z)_{\bar{N}_i}^{(\Lambda)}$ for some $z \in Z$, we have

$$\begin{aligned} \lim_{\Lambda\uparrow\mathbb{Z}^d} \frac{1}{|\Lambda|} \sum_{i\in\Lambda} \mathbf{1}[x_{\bar{N}_i}^{(\Lambda)} = \bar{\xi}^{(1)}, y_{\bar{N}_i}^{(\Lambda)} = \bar{\xi}^{(2)}] \quad &= \quad \lim_{\Lambda\uparrow\mathbb{Z}^d} \frac{1}{|\Lambda|} \sum_{i\in S_x(\bar{\xi}^{(1)})\cap\Lambda} \mathbf{1}[z_{\bar{N}_i} = \bar{\xi}^{(3)}] \\ &= \quad \nu_{\epsilon_0}(\bar{\xi}^{(3)}) M_{\theta_0}(\bar{\xi}^{(1)}) \end{aligned} \tag{3.5}$$

with P_{θ_0,ϵ_0}-probability one. It then follows that $\lim_{\Lambda\uparrow\mathbb{Z}^d} A_{\Lambda,x,z}(\bar{\xi}^{(1)}, \bar{\xi}^{(2)}) = \nu_{\epsilon_0}(\bar{\xi}^{(3)})$, again with P_{θ_0,ϵ_0}-probability one, and this concludes the proof since $\nu_\epsilon(\bar{\xi}^{(3)}) = \epsilon^{D(\bar{\xi}^{(1)},\bar{\xi}^{(2)})}(1 - \epsilon)^{|\bar{N}_0|-D(\bar{\xi}^{(1)},\bar{\xi}^{(2)})}$. ■

Lemma 4 was formulated for $\epsilon_0 \in [0, 1]$; however, A_ϵ is defined for any $\epsilon \in \mathbb{R}$ and we now proceed to describe properties of A_ϵ for generic values of ϵ.

Lemma 5. *For $a, b \in \mathbb{R}$, let $A_{a,b}$ be the matrix defined by*

$$A_{a,b}(\bar{\xi}^{(1)}, \bar{\xi}^{(2)}) = a^{D(\bar{\xi}^{(1)},\bar{\xi}^{(2)})} b^{|\bar{N}_0|-D(\bar{\xi}^{(1)},\bar{\xi}^{(2)})}$$

for any $\bar{\xi}^{(1)}, \bar{\xi}^{(2)} \in X_{\bar{N}_0}$. Then

$$A_{\epsilon,1-\epsilon} = A_\epsilon; \tag{3.6}$$

$$A_{0,1} = I \quad \textit{(the identity matrix)}, \tag{3.7}$$

$$A_{ab} \cdot A_{cd} = A_{ad+bc,ac+bd}, \tag{3.8}$$

and

$$A_\epsilon^{-1} = A_{\epsilon,1-\epsilon}^{-1} = A_{\frac{\epsilon}{2\epsilon-1},\frac{\epsilon-1}{2\epsilon-1}} = (2\epsilon - 1)^{-|\bar{N}_0|} A_{\epsilon,\epsilon-1}, \quad \textit{for } \epsilon \neq \frac{1}{2}. \tag{3.9}$$

Proof: Properties (3.6) and (3.7) are immediate, and (3.9) follows from (3.8) and the homogeneity of $A_{a,b}$ in (a,b). Writing D_{ij} as shorthand for $D(\bar{\xi}^{(i)}, \bar{\xi}^{(j)})$, we verify property (3.8) as follows:

$$\sum_{\bar{\xi}^{(3)} \in X_{\bar{N}_0}} A_{ab}(\bar{\xi}^{(1)}, \bar{\xi}^{(3)}) A_{cd}(\bar{\xi}^{(3)}, \bar{\xi}^{(2)}) = \sum_{\bar{\xi}^{(3)} \in X_{\bar{N}_0}} a^{D_{1,3}} b^{|\bar{N}_0| - D_{1,3}} c^{D_{3,2}} d^{|\bar{N}_0| - D_{3,2}}$$

$$= \sum_{\ell=0}^{|\bar{N}_0| - D_{1,2}} \binom{|\bar{N}_0| - D_{1,2}}{\ell} (ac)^{\ell} (bd)^{|\bar{N}_0| - D_{1,2} - \ell} \sum_{m=0}^{D_{1,2}} \binom{D_{1,2}}{m} (ad)^m (bc)^{D_{1,2} - m}$$

$$= (ac + bd)^{|\bar{N}_0| - D_{1,2}} (ad + bc)^{D_{1,2}} = A_{ad+bc, ac+bd}(\bar{\xi}^{(1)}, \bar{\xi}^{(2)}).$$

■

The next lemma describes how to use the condition on the roots of (2.22) to obtain a related condition on the roots of (2.21) when $\epsilon' = \epsilon_0$.

Lemma 6. *Let $\epsilon_0 \in [0, \frac{1}{2})$ and let Q be a polynomial relation satisfying $Q(M_\theta) = 0$. Suppose that, for all $\theta \in \Theta$, the equation $Q(M_{\theta,0,\epsilon}) = 0$ admits $\epsilon = 0$ as a single root and has no real roots in $(-\infty, 0)$. Then, for all $\theta \in \Theta$, the equation*

$$Q(M_{\theta,\epsilon_0,\epsilon}) = 0 \tag{3.10}$$

admits $\epsilon = \epsilon_0$ as a single root and has no real roots in $(-\infty, \epsilon_0)$.

Proof: Fix any $\theta \in \Theta$. It follows from Lemma 5 that

$$M_{\theta,0,\epsilon'} = (2\epsilon_0 - 1)^{-|\bar{N}_0|} M_\theta A_{\epsilon', \epsilon'-1}.$$

Similarly,

$$M_{\theta,\epsilon_0,\epsilon''} = (2\epsilon_0 - 1)^{-|\bar{N}_0|} M_\theta A_{\epsilon''-\epsilon_0, \epsilon''+\epsilon_0-1} = -M_\theta A_{\frac{\epsilon''-\epsilon_0}{1-2\epsilon_0}, \frac{\epsilon''-\epsilon_0}{1-2\epsilon_0}-1},$$

where the last step uses the homogeneity of $A_{a,b}$ in (a,b). Since Q is a homogeneous polynomial, $Q(M_{\theta,0,\epsilon'})$ has a root of order n at $\epsilon' = \tilde{\epsilon}$ iff $Q(M_{\theta,\epsilon_0,\epsilon''})$ has a root of order n at $\epsilon'' = (1 - 2\epsilon_0)\tilde{\epsilon} + \epsilon_0$. The hypotheses of the lemma now guarantee that (3.10) has a single root at $\epsilon = \epsilon_0$, and it cannot have any roots in $(-\infty, \epsilon_0)$. ■

We conclude this theoretical part of the paper with the proof of the consistency of the estimators.

Proof (of the Consistency Theorem). For $\epsilon \in \mathbb{R}$, $x \in X$, $z \in Z$ and $\Lambda \subset \mathbb{Z}^d$, let $M_{\Lambda,x,z,\epsilon} = M_{\Lambda,x} A_{\Lambda,x,z} A_\epsilon^{-1}$, and for $\gamma, a, b > 0$, let $R_{\gamma,a,b} = \{t \in \mathbf{C} : -\gamma \leq \mathrm{Re}(t) \leq a, |\mathrm{Im}(t)| \leq b\}$. For any $\gamma, b > 0$, $a \in (\epsilon_0, \frac{1}{2})$, and $\bar{\xi} \in X_{\bar{N}_0}$, we claim that

$$\lim_{\Lambda \uparrow \mathbb{Z}^d} \sup_{\epsilon \in R_{\gamma,a,b}} |M_{\Lambda,x,z,\epsilon}(\bar{\xi}) - M_{\theta_0,\epsilon_0,\epsilon}(\bar{\xi})| = 0 \tag{3.11}$$

with P_{θ_0,ϵ_0}-probability one. In fact, for any $\bar{\xi} \in X_{\bar{N}_0}$

$$\begin{aligned}
&\sup_{\epsilon \in R_{\gamma,a,b}} |M_{\Lambda,x,z,\epsilon}(\bar{\xi}) - M_{\theta_0,\epsilon_0,\epsilon}(\bar{\xi})| \\
&\leq 2^{|\bar{N}_0|} \max_{\bar{\xi}^{(1)} \in X_{\bar{N}_0}} |M_{\Lambda,x,z}(\bar{\xi}^{(1)}) - M_{\theta_0,\epsilon_0}(\bar{\xi}^{(1)})| \cdot \sup_{\substack{\epsilon \in R_{\gamma,a,b} \\ \bar{\xi}^{(1)} \in X_{\bar{N}_0}}} |A_\epsilon^{-1}(\bar{\xi}^{(1)}, \bar{\xi})| \\
&\leq 2^{2|\bar{N}_0|} \{ \max_{\bar{\xi}^{(1)}, \bar{\xi}^{(2)} \in X_{\bar{N}_0}} |A_{\Lambda,x,z}(\bar{\xi}^{(2)}, \bar{\xi}^{(1)}) - A_{\epsilon_0}(\bar{\xi}^{(2)}, \bar{\xi}^{(1)})| \\
&\quad + \max_{\bar{\xi}^{(2)} \in X_{\bar{N}_0}} |M_{\Lambda,x}(\bar{\xi}^{(2)}) - M_{\theta_0}(\bar{\xi}^{(2)})| \} \cdot \sup_{\substack{\epsilon \in R_{\gamma,a,b} \\ \bar{\xi}^{(1)} \in X_{\bar{N}_0}}} |A_\epsilon^{-1}(\bar{\xi}^{(1)}, \bar{\xi})|,
\end{aligned}$$

so that (3.11) follows from the ergodic theorem, Lemma 4 and the continuity of the mapping $\epsilon \mapsto A_\epsilon{-1}$.

For all $\epsilon_0 \in [0, \frac{1}{2})$, it follows from the assumptions about the roots of (2.22) and from Lemma 6 that $\epsilon = \epsilon_0$ is a single root of $Q(M_{\theta_0,\epsilon_0,\epsilon}) = 0$ and that this equation has no other real root in $(-\infty, \epsilon_0)$. Since $Q(M_{\theta_0,\epsilon_0,\epsilon})$ is a polynomial in ϵ, for every sufficiently small $\delta > 0$

$$\operatorname{sgn}[Q(M_{\theta_0,\epsilon_0,\epsilon_0-\delta})] = -\operatorname{sgn}[Q(M_{\theta_0,\epsilon_0,\epsilon_0+\delta})], \tag{3.12}$$

and there exists $c = c(\delta) > 0$ such that

$$|Q(M_{\theta_0,\epsilon_0,\epsilon_0-\delta})|, |Q(M_{\theta_0,\epsilon_0,\epsilon_0+\delta})| \geq c \tag{3.13}$$

and

$$\inf_{\epsilon \in R_{\gamma,\epsilon_0-\delta,c}} |Q(M_{\theta_0,\epsilon_0,\epsilon})| \geq c. \tag{3.14}$$

It follows from (3.11) and the continuity of Q in its arguments that, with P_{θ_0,ϵ_0}-probability one, there is a Λ large enough that $Q(M_{\Lambda,x,z,\epsilon})$ satisfies (3.12)-(3.14) with c replaced by $c/2$ in (3.13) and (3.14). As $Q(M_{\Lambda,x,z,\epsilon})$ also is a polynomial in ϵ, this implies that $Q(M_{\Lambda,x,z,\epsilon}) = 0$ has at least one root in $[\epsilon_0 - \delta, \epsilon_0 + \delta]$, and no other roots in $R_{\gamma,\epsilon_0-\delta,c/2}$. Therefore, $\hat{\epsilon}^{(\Lambda)}(y)$ can be taken to be the smallest (real) root in $[-\gamma, \epsilon_0 + \delta]$ of $Q(M_{\Lambda,y,\epsilon}) = 0$, and letting $\delta \downarrow 0$ shows that

$$\hat{\epsilon}^{(\Lambda)}(y) \to \epsilon_0 \text{ as } \Lambda \uparrow \mathbb{Z}^d, \tag{3.15}$$

with P_{θ_0,ϵ_0}-probability one. This estimator for ϵ_0 allows us to "clean," i.e. reconstruct, the probabilities of complete local patterns in μ_{θ_0}: for all $\bar{\xi} \in X_{\bar{N}_0}$

$$\begin{aligned}
&|M_{\Lambda,y,\hat{\epsilon}^{(\Lambda)}(y)}(\bar{\xi}) - M_{\theta_0}(\bar{\xi})| \\
&\leq 2^{|\bar{N}_0|} \{ \max_{\bar{\xi}^{(1)} \in X_{\bar{N}_0}} |A^{-1}_{\hat{\epsilon}^{(\Lambda)}(y)}(\bar{\xi}^{(1)}, \bar{\xi}) - A^{-1}_{\epsilon_0}(\bar{\xi}^{(1)}, \bar{\xi})| \\
&\quad + \max_{\bar{\xi}^{(1)} \in X_{\bar{N}_0}} |M_{\Lambda,y}(\bar{\xi}^{(1)}) - M_{\theta_0,\epsilon_0}(\bar{\xi}^{(1)})| \cdot \max_{\bar{\xi}^{(1)} \in X_{\bar{N}_0}} |A^{-1}_{\epsilon_0}(\bar{\xi}^{(1)}, \bar{\xi})| \},
\end{aligned}$$

so that

$$M_{\Lambda,y,\hat{\epsilon}(\Lambda)(y)} \to M_{\theta_0} \tag{3.16}$$

with P_{θ_0,ϵ_0}-probability one, by the continuity of the mapping $\epsilon \mapsto A_\epsilon^{-1}$, (3.15), and (3.11).

We now will use a maximum pseudo-likelihood method to estimate θ_0. In the case of no noise (Besag (1977)), the method uses the function

$$MPL(\theta,\Lambda,x) = \prod_{i\in\Lambda} \pi_\theta(x_i|x_{N_i}^{(\Lambda)}) = \prod_{(x_{\mathbf{0}}\vee\xi)\in X_{\bar{N}_{\mathbf{0}}}} \pi_\theta(x_{\mathbf{0}}|\xi)^{M_{\Lambda,x}(x_{\mathbf{0}}\vee\xi)} \tag{3.17}$$

which converges to

$$MPL(\theta,\theta_0) = \prod_{(x_{\mathbf{0}}\vee\xi)\in X_{\bar{N}_{\mathbf{0}}}} \pi_\theta(x_{\mathbf{0}}|\xi)^{M_{\theta_0}(x_{\mathbf{0}}\vee\xi)} \tag{3.18}$$

with μ_{θ_0}-probability one, as $\Lambda \uparrow \mathbb{Z}^d$. Our method uses (2.24) instead of (3.17). We first prove the convergence of $GMPL(\theta,\Lambda,y)$ to $MPL(\theta,\theta_0)$ when θ belongs to some compact set $\Theta_c \subset \mathbb{R}^s$; note that here θ is not restricted to Θ. The functions $\theta \mapsto \pi_\theta(x_{\mathbf{0}}|\xi)$ are continuous and strictly positive; therefore,

$$\begin{aligned}
&\sup_{\theta\in\Theta_c} |\log GMPL(\theta,\Lambda,y) - \log MPL(\theta,\theta_0)| \\
\le\ & 2^{|\bar{N}_{\mathbf{0}}|} \sup_{\bar{\xi}\in X_{\bar{N}_{\mathbf{0}}}} |M_{\Lambda,y,\hat{\epsilon}(\Lambda)(y)}(\bar{\xi}) - M_{\theta_0}(\bar{\xi})| \sup_{\substack{\theta\in\Theta_c \\ (x_{\mathbf{0}}\vee\xi)\in X_{\bar{N}_{\mathbf{0}}}}} |\log(\pi_\theta(x_{\mathbf{0}}|\xi))| \\
\le\ & K(\Theta_c) \sup_{\bar{\xi}\in X_{\bar{N}_{\mathbf{0}}}} |M_{\Lambda,y,\hat{\epsilon}(\Lambda)(y)}(\bar{\xi}) - M_{\theta_0}(\bar{\xi})| \ ,
\end{aligned}$$

with $K(\Theta_c)$ a suitable constant depending on Θ_c, so that

$$\sup_{\theta\in\Theta_c} |\log GMPL(\theta,\Lambda,y) - \log MPL(\theta,\theta_0)| \to 0 \tag{3.19}$$

by (3.16), with P_{θ_0,ϵ_0}-probability one.

Next, we show that $\log GMPL(\theta,\Lambda,y)$ is a strictly concave function of θ. Let θ, $\theta' \in \mathbb{R}^s$, with $\theta' \neq 0$. Then

$$\begin{aligned}
&\frac{d}{dt}\log GMPL(\theta+t\theta',\Lambda,y) \qquad (3.20)\\
&= \sum_{(x_{\mathbf{0}}\vee\xi)\in X_{\bar{N}_{\mathbf{0}}}} M_{\Lambda,y,\hat{\epsilon}(\Lambda)(y)}(x_{\mathbf{0}}\vee\xi)\ \pi_{\theta+t\theta'}(-x_{\mathbf{0}}|\xi)\ \theta'\cdot[U_{\mathbf{0}}(-x_{\mathbf{0}}\vee\xi) - U_{\mathbf{0}}(x_{\mathbf{0}}\vee\xi)]
\end{aligned}$$

and

$$\begin{aligned}
&\frac{d^2}{dt^2}\log GMPL(\theta+t\theta',\Lambda,y) \qquad (3.21)\\
&= -\sum_{(x_{\mathbf{0}}\vee\xi)\in X_{\bar{N}_{\mathbf{0}}}} M_{\Lambda,y,\hat{\epsilon}(\Lambda)(y)}(x_{\mathbf{0}}\vee\xi)\pi_{\theta+t\theta'}(-x_{\mathbf{0}}|\xi)\pi_{\theta+t\theta'}(x_{\mathbf{0}}|\xi)(\theta'\cdot[U_{\mathbf{0}}(-x_{\mathbf{0}}\vee\xi) - U_{\mathbf{0}}(x_{\mathbf{0}}\vee\xi)])^2.
\end{aligned}$$

Since $\theta' \neq 0$ and $M_{\theta_0}(\bar{\xi}) > 0$, it follows from (2.7), (3.16) and (3.21) that for all $\theta \in \mathbb{R}^s$, $\log GMPL(\theta, \Lambda, y)$ is strictly concave with P_{θ_0,ϵ_0}-probability one if Λ is sufficiently large.

Finally, note that (3.20) and (3.21) also hold (with M_{θ_0} replacing $M_{\Lambda,y,\hat{\epsilon}^{(\Lambda)}(y)}$) if $MPL(\theta, \theta_0)$ replaces $GMPL(\theta, \Lambda, y)$. In this case, we may set $\theta = \theta_0$ in (3.20) and use (2.7) to see that $\log MPL(\theta, \theta_0)$ has a unique maximum, and that it achieves this maximum at $\theta = \theta_0$. Therefore, by taking Θ_c such that $\theta_0 \in \Theta_c$, it follows from (3.19) that with P_{θ_0,ϵ_0}-probability one, if Λ is large enough, $\log GMPL(\theta, \Lambda, y)$ has a unique maximum in $\mathbb{R}^s$, at $\hat{\theta}^{(\Lambda)}(y) \in \Theta_c$, and that $\hat{\theta}^{(\Lambda)}(y) \to \theta_0$ as $\Lambda \uparrow \mathbb{Z}^d$. ■

The Consistency Theorem was formulated for the (homogeneous) polynomial relations. Lemma 3 indicates how to produce at least one polynomial equation of the type (2.17); however it may happen that all of the polynomial relations which are obtained in this fashion are *null-relations* in that each is identically zero in ϵ for certain values if $\theta \in \Theta$. In such circumstances, it may be possible to use further insight into the model to obtain additional non-null (or *effective*) relations — an example is given in section 4 (for the general nearest-neighbor Markov random field on $\mathbb{Z}$). Loosely speaking, the polynomial relations which we have shown how to construct in Lemma 3 can fail to be effective if either Θ is taken to be an unreasonably large subset of $\mathbb{R}^s$, or there are some special symmetries in the process. In the latter case, it is not unreasonable to expect that the presence of the many symmetries necessary to make all of the these polynomial relations null should lead to (enough of) an exact solution of the model to enable one to construct some additional relations which are not null. It is also conceivable that some these new relations may not be polynomial relations in the sense defined above — in that they may be nonhomogeneous. Actually, in the example of section 4 mentioned above, we are able to find two effective relations: one is homogeneous (and thus is covered by the Consistency Theorem), the other is not.

To be able to handle situations in which we wish to consider nonhomogeneous polynomials, we point out that if no use is made of Lemma 6 and the homogeneity of Q, then the proof of the Consistency Theorem yields the following result.

Corollary. *Let C and Θ be given, and let $\epsilon_0 \in [0, \frac{1}{2})$ and $\theta_0 \in \Theta$. Suppose that $Q = Q(M)$ is a polynomial in the entries of M such that for every $\theta \in \Theta$, $Q(M_\theta) = 0$ and the only solution $\epsilon \in (-\infty, \epsilon_0]$ of*

$$Q(M_{\theta,\epsilon_0,\epsilon}) = 0$$

is a single root at ϵ_0. Then for any $\gamma > 0$,
(I) the smallest real root in $[-\gamma, 1]$, call it $\hat{\epsilon}^{(\Lambda)}(y)$, of

$$Q(M_{\Lambda,y,\epsilon}) = 0,$$

is such that

$$\lim_{\Lambda \uparrow \mathbb{Z}^d} \hat{\epsilon}^{(\Lambda)}(y) = \epsilon_0,$$

and (II) the vector $\hat{\theta}^{(\Lambda)}(y)$ which maximizes

$$GMPL(\theta,\Lambda,y) = \prod_{(x_{\mathbf{0}}\vee\xi)\in X_{\tilde{N}_{\mathbf{0}}}} [\pi_\theta(x_{\mathbf{0}}|\xi)]^{M_{\Lambda,y,\hat{\epsilon}^{(\Lambda)}(y)}(x_{\mathbf{0}}\vee\xi)}$$

in $\mathbb{R}^s$ converges pointwise to θ_0 with P_{θ_0,ϵ_0}-probability one.

4. Examples in one dimension. The reader should be able to read this section right after the introduction [comments between brackets are for the benefit of those who have also read sections 2 and 3], occasionally looking at some parts of sections 2 and 3 when indicated.

In this section we consider the one-dimensional nearest neighbor Markov random field on $\{-1,1\}^{\mathbb{Z}}$. Stationary one-dimensional nearest neighbor Markov random fields are in one-to-one correspondence with a collection of stationary one-step Markov chains, the latter defined by means of a two-parameter transition matrix $\begin{pmatrix} 1-a & a \\ b & 1-b \end{pmatrix}$, with $a,b \in (0,1)$. The Markov chains with $a=0$, $a=1$, $b=0$ or $b=1$ do not correspond to Markov random fields because each of these chains prohibits some pattern of a pair of colors — which corresponds to having an interaction which is infinite on that pattern. We shall further assume that the $a \neq 1-b$, which is the same as requiring that the Markov chain not be a Bernoulli process — since in that case the interaction is not identifiable. We can use the correspondence between Markov random fields and Markov chains as a method for writing down the probabilities of the specification of colors in some finite array as explicit functions of the interaction, ϕ_0. This explicit functional dependence is called an *exact solution* of the model.

The existence of an exact solution should mean that the estimation of the parameters θ_0 and ϵ_0 is not overly difficult, and this presentation serves mainly to illustrate our method in a simple case, and also to offer an explicit estimator that the reader might want to compare with his/her preferred one. It would now be easy to obtain our polynomial relations from the exact solution — eventually we will do this, and the impatient reader can go directly to the subsection on the general nearest-neighbor Markov random fields below — but we choose instead to begin working from the general setup of Markov random fields, using interactions between pixel colors [along the lines of section 2]. In this way we shall also illustrate what one does when the exact solution is not available.

Before we go on, let us mention two possible alternative estimators which could be used in this situation. First, as the necessary probabilities can be easily obtained from the exact solution, one can try to use the Maximum Likelihood Estimator to estimate a, b and ϵ simultaneously. Second, if one has the a-priori information that the Markov chain is symmetric, i.e. $a=b$, then the method of moments, as in Frigessi and Piccioni (1990), turns out to be an computationally easier task. However, it is not clear how to extend the method of moments to the asymmetric case, $a \neq b$, and we thus make this our starting point, providing a computationally fast method for estimating the parameter in an asymmetric one-step Markov chain with independent symmetric noise. The general case, with $a,b \in (0,1)$ and $a \neq 1-b$ (but without the hypothesis that $a \neq b$) will be discussed afterwards.

Asymmetric nearest-neighbor Markov random fields. As mentioned above, we begin our treatment by studying the *interactions*, which are functions of the colors in cliques. It is suggested that the reader who has not already done so, should read the definitions in section 2 up to (2.3). Here cliques are all the elements of $\mathbb{Z}$, and all the pairs (which we can take to be ordered, for simplicity) of nearest neighbor elements of $\mathbb{Z}$. Modulo translations, there are two cliques in the present model, for a total of six specifications of colors:

$$(+,+),(+,-),(-,+),(-,-),(+),(-); \tag{4.1}$$

here and in the following we abbreviate by dropping the 1 in $+1$ and -1. It is assumed that a noise level $\epsilon_0 \in [0, \frac{1}{2})$ and an asymmetric interaction, ϕ_0, are fixed but unknown; an interaction is a real-valued function of the six specifications above, and it is asymmetric if it is not invariant under the exchange of $+$ and $-$ in the specification of colors above. Our goal is to estimate the parameters ϕ_0 and ϵ_0, which characterize the Markov random field and the noise process, respectively, by observing the specification of colors in finite portions of the observable image $y \in \{-1,1\}^{\mathbb{Z}}$; this image is distributed according to the product measure $P_{\phi_0,\epsilon_0} = \mu_{\phi_0} \times \nu_{\epsilon_0}$, where μ_{ϕ_0} is a Markov random field of ϕ_0 and ν_{ϵ_0} is the Bernoulli measure with density ϵ_0. The interaction can only be estimated up to an equivalence class, where two interactions are said to be equivalent if they generate the same Markov random field [see Lemma 1].

The neighborhood of the origin is $\{-1,1\} \subset \mathbb{Z}$, the complete neighborhood of the origin is $\{-1,\mathbf{0},1\} \subset \mathbb{Z}$, and interactions are vectors in $\mathbb{R}^6$. We list the $2^{|N_0|}$ local patterns:

$$(++),(+-),(-+),(--), \tag{4.2}$$

and the $2^{|\bar{N}_0|}$ complete local patterns:

$$(+++),(++-),(-++),(-+-),(+-+),(+--),(--+),(---), \tag{4.3}$$

and we fix these orderings.

The energy function (at the origin) is

$$U^{\phi}(\bar{\xi}_{(-1,\mathbf{0},1)}) = \phi(\bar{\xi}_{(-1,\mathbf{0})}) + \phi(\bar{\xi}_{(\mathbf{0},1)}) + \phi(\bar{\xi}_{\mathbf{0}})$$

for any interaction ϕ and any $\bar{\xi}_{(-1,\mathbf{0},1)} \in X_{\bar{N}_0}$. An examination of this energy function leads to the construction of some polynomial relations [following Lemmas 1 and 3].

(I) Let $\mathbf{U}(\phi)$ be the vector indexed by $\xi \in X_{N_0}$ given by $\mathbf{U}(\phi)(\xi) = U^{\phi}(+ \vee \xi) - U^{\phi}(- \vee \xi)$, where we define $(\pm \vee \xi) = (\xi_{(-1)}, \pm, \xi_{(+1)}) \in X_{\bar{N}_0}$ for $\xi = \xi_{(-1,1)} \in X_{N_0}$. [For special choices of ϕ, $\mathbf{U}(\phi)$ is a column of the related matrix $\mathbf{U}$ in (2.19).] This defines a map $\phi = (\phi(++), \phi(+-), \phi(-+), \phi(--), \phi(+), \phi(-)) \mapsto \mathbf{U}(\phi) = \mathbf{U}\phi$ from $\mathbb{R}^6$ to $\mathbb{R}^4$ described by the matrix

$$\mathbf{U} = \begin{pmatrix} 2 & -1 & -1 & 0 & 1 & -1 \\ 1 & 0 & 0 & -1 & 1 & -1 \\ 1 & 0 & 0 & -1 & 1 & -1 \\ 0 & 1 & 1 & -2 & 1 & -1 \end{pmatrix} \tag{4.4}$$

when using the order given in (4.1) and (4.2).

(II) The model can be parametrized [see Lemma 1] by first finding a basis of a linear space $\mathcal{S}$ independent of the null space of the matrix $\mathbf{U}$, and then writing the interactions as linear combinations of these basis elements, which can be chosen to be vectors in $\mathbb{Z}^6$. A customary choice in the present case is to form the basis from $\phi_1 = (1,-1,-1,1,0,0)$ and $\phi_2 = (0,0,0,0,1,-1)$, in which case the two parameters are called β and h. Regardless of which basis we choose, we have that $\dim(\mathcal{S}) = 2 < 4 = 2^{|N_0|}$ [as predicted by Lemma 2].

(III) We find polynomial relations [see Lemma 3] by solving the equation

$$\alpha\mathbf{U} = 0,$$

i.e., by finding vectors α which are orthogonal to the columns of $\mathbf{U}$. Since $\mathbf{U}$ is a rank 2 $[2 = \dim(\mathcal{S})]$ linear transformation into $\mathbb{R}^4$ $[4 = 2^{|N_0|}]$, it is possible to find two such independent vectors, for instance

$$\alpha^{(1)} = (0,1,-1,0) \tag{4.5}$$

and

$$\alpha^{(2)} = (1,-2,0,1). \tag{4.6}$$

(IV) Next, for any $\theta = (\theta(1),\theta(2)) \in \mathbb{R}^2$, let M_θ be the vector indexed by $\bar{\xi} \in X_{\bar{N}_0}$ of the probabilities $M_\theta(\bar{\xi}) = \mu_{\theta(1)\phi_1+\theta(2)\phi_2}(\bar{\xi})$, where ϕ_1 and ϕ_2 are a basis for $\mathcal{S}$. A pair of polynomial equations can be obtained from (2.16), using the α's in (4.5) and (4.6):

$$Q_{\alpha^{(1)}}(M) = M(++-)M(--+) - M(+--)M(-++) = 0 \tag{4.7}$$

and

$$Q_{\alpha^{(2)}}(M) = M(+++)(M(+--))^2M(-+-) - M(+-+)(M(++-))^2M(---) = 0. \tag{4.8}$$

(V) Let $M_{\theta,\epsilon_0,\epsilon} = M_\theta A_{\epsilon_0}A_\epsilon^{-1}$, where A_ϵ is the matrix defined in (2.14). It is clear that $\epsilon = \epsilon_0$ is a solution of $Q_{\alpha^{(i)}}(M_{\theta,\epsilon_0,\epsilon}) = 0$ for $i = 1,2$, but it it is not easy to tell whether the relations $Q_{\alpha^{(i)}}$ are null in the sense that for some value of θ, $Q_{\alpha^{(i)}}(M_{\theta,\epsilon_0,\epsilon}) = 0$ for all ϵ. It is to decide whether these relations are null or effective that we now turn to the exact solution, as proceeding without it would prove quite arduous.

(VI) It is most convenient to change from the original parameters $\theta_0 = (\theta_1,\theta_2)$ to the parameters a and b appearing in the transition matrix for the Markov chain corresponding to the Markov random field. Let us indicate this change of parameters by the transformation $(a,b) \mapsto \theta(a,b)$. We then have

$$\begin{aligned} M_{a,b} &= M_{\theta(a,b)} \\ &= (a+b)^{-1}[b(1-a)^2, ab(1-a), ab(1-a), a^2b, ab^2, ab(1-b), ab(1-b), a(1-b)^2], \end{aligned} \tag{4.9}$$

where we use the ordering of the complete local patterns given in (4.3). A direct calculation (left to the reader) shows that $Q_{\alpha^{(1)}}$ is a null relation. On the other hand, temporarily assuming that $\epsilon_0 = 0$, we note that $Q_{\alpha^{(2)}}(M_{\theta(a,b),0,\epsilon}) = 0$ reduces to

$$\epsilon(\epsilon-1)ab(b-a)[a-(1-b)]^2(\epsilon^2-\epsilon+ab)[(a+b)\epsilon^2-(a+b)\epsilon+ab(2-a-b)] = 0, \quad (4.10)$$

where we have ignored various factors of $a+b$ and $2\epsilon-1$. Our assumptions on the parameters a and b ($a, b \neq 0$, $a \neq 1-b$, and $a \neq b$) guarantee that (4.10) is not identically zero. Therefore we only need to study the roots of the last two factors in order to verify that they cannot be confused with the root at $\epsilon = \epsilon_0 (= 0)$ when estimating ϵ_0 from the data. It is easily seen that the real roots (in ϵ) of last two factors in (4.10), are in $(0,1)$ for all $a, b \in (0,1)$. The homogeneity of the polynomial $Q_{\alpha^{(2)}}$ allows us to conclude from the observation that $\epsilon = 0$ is (a single root and) the smallest real root of (4.10), that for any $\epsilon_0 \in [0, \frac{1}{2})$, the root at $\epsilon = \epsilon_0$ of $Q_{\alpha^{(2)}}(M_{\theta(a,b),\epsilon_0,\epsilon}) = 0$ is (a single root and) the smallest real root [see Lemma 6].

(VIII) We are now ready to define our estimators. Recall that for any given interaction $\phi_0 \in \mathbb{R}^6$, there is a vector $\theta_0 \in \mathbb{R}^2$ with $\theta_0(1)\phi_1 + \theta_0(2)\phi_2 = \phi_0$, where ϕ_1 and ϕ_2 the elements of some basis for $\mathcal{S}$, as described in (II) above. We will estimate ϵ_0 and θ_0. First, let $y^{(\Lambda)}$ be the periodic extension to $\mathbb{Z}$ of the observed image y_Λ (which is the restriction of the observable image y to Λ), and define $M_{\Lambda,y}$ to be the vector, indexed by $\bar{\xi} \in X_{\bar{N}_0}$, of the empirical frequencies of the specifications of colors $\bar{\xi}$ in $y^{(\Lambda)}$ in Λ [see (2.12]. Then, form the equation in ϵ,

$$Q_{\alpha^{(2)}}(M_{\Lambda,y}A_\epsilon^{-1}) = 0, \quad (4.11)$$

[as in (2.23)]. Next, find the smallest real root $\hat{\epsilon}^{(\Lambda)}(y)$ of (4.1); [as we proved in the Consistency Theorem]

$$\hat{\epsilon}^{(\Lambda)}(y) \to \epsilon_0$$

with P_{ϕ_0,ϵ_0}-probability one as $\Lambda \uparrow \mathbb{Z}^d$. Finally, consider the function $GMPL : \mathbb{R}^2 \to \mathbb{R}$ given [as in (2.24)] by

$$\theta \mapsto GMPL(\theta, \Lambda, y) = \prod_{(x_0 \vee \xi) \in X_{\bar{N}_0}} [\pi_\theta(x_0|\xi)]^{M_{\Lambda,y,\hat{\epsilon}^{(\Lambda)}(y)}(x_0 \vee \xi)} \quad (4.12)$$

where $\pi_\theta(x_0|\xi) = M_\theta(x_0 \vee \xi)/(M_\theta(+\vee\xi) + M_\theta(-\vee\xi))$, for $\xi \in X_{N_0}$. The vector $\hat{\theta}^{(\Lambda)}(y) \in \mathbb{R}^2$ which maximizes $GMPL(\theta, \Lambda, y)$ satisfies [see the Consistency Theorem (and its proof)]

$$\hat{\theta}^{(\Lambda)}(y) \to \theta_0$$

with P_{ϕ_0,ϵ_0}-probability one as $\Lambda \uparrow \mathbb{Z}^d$, and this concludes the discussion of the asymmetric case.

General nearest-neighbor Markov random fields. We retain our 'finite energy' ($a, b \neq 0$) and non-Bernoulli ($a \neq 1-b$) assumptions on the parameters a and b, but now relax the asymmetry ($a \neq b$) hypothesis. Note that it was essential in (4.10) that $a \neq b$; if we are not given this a-priori information, then $Q_{\alpha^{(2)}}$ is also a null-relation, and unfortunately no effective relations are derived from our theory. However, one can now go further, and add more relations by using the exact solution (4.9) explicitly. An examination of (4.9) yields a pair of equations which are independent of the two null-relations (4.7) and (4.8):

$$Q_3(M) = M(++-) - M(-++) = 0, \tag{4.13}$$

and

$$Q_4(M) = M(+++)M(-+-) - (M(++-))^2 = 0. \tag{4.14}$$

Another direct calculation (left to the reader) shows that Q_3 is a null relation. However, it is also readily seen that $Q_4(M_{\theta(a,b),0,\epsilon}) = 0$ is equivalent (neglecting various factors of $a+b$ and $2\epsilon - 1$) to

$$\epsilon(\epsilon-1)ab[a-(1-b)]^2 = 0. \tag{4.15}$$

Therefore Q_4 is an effective relation, and as above, for any $\epsilon_0 \in [0, \frac{1}{2})$ and $\phi_0 \in \mathbb{R}^6$, we can define the estimators $\hat{\epsilon}^{(\Lambda)}$ and $\hat{\theta}^{(\Lambda)}$ to be the smallest (real) root of $\epsilon \mapsto Q_4(M_y A_\epsilon^{-1}) = 0$ and the vector in $\mathbb{R}^2$ maximizing (4.12) with this current value of $\hat{\epsilon}^{(\Lambda)}$, respectively. These estimators converge (with P_{θ_0,ϵ_0}-probability one) to ϵ_0 and θ_0, respectively, where θ_0 is related to the interaction ϕ_0 by $\theta_0(1)\phi_1 + \theta_0(2)\phi_2 = \phi_0$. [Although Q_4 was not obtained via the construction described in Lemma 3, it is a polynomial relation, and so the convergence of the estimators is according to the Consistency Theorem.]

For the sake of completeness we point out that there are eight components to $M_{a,b}$. We know that these eight components are functions of two parameters, they sum to 1, and they satisfy $Q_{\alpha^{(1)}}(M_{a,b}) = Q_{\alpha^{(2)}}(M_{a,b}) = Q_3(M_{a,b}) = Q_4(M_{a,b}) = 0$. Thus we should be able to extract one more relation from the complete solution, and indeed we can. One choice for a fifth independent equation is

$$Q_5(M) = M(-+-)[M(-+-)+M(+-+)][M(+++)+M(++-)]^2 - M^2(++-)M(+-+) = 0. \tag{4.16}$$

This fifth equation is different from the first four in that it is not homogeneous. On the other hand it is still possible to base the estimation on this relation, and we conclude this paper by showing that Q_5 is effective as an example of how to proceed in the case of a nonhomogeneous polynomial.

A somewhat lengthy calculation shows that $Q_5(M_{\theta(a,b),\epsilon_0,\epsilon}) = 0$ reduces to

$$\begin{aligned}
&(\epsilon-\epsilon_0)[\epsilon-(1-\epsilon_0)](2\epsilon-1)^2 ab[a-(1-b)]^2 \\
&\quad\cdot\ [\epsilon^2 - \epsilon + ab + (1-4b)\epsilon_0(1-\epsilon_0)] \\
&\quad\cdot\ \{3(a+b)\epsilon^3 - [2a+7b+5(a-b)\epsilon_0]\epsilon^2 \\
&\qquad +[b(5-a^2-ab) + (4a-6b+4a^2b+4ab^2)\epsilon_0 + (a+b-4a^2b-4ab^2)\epsilon_0^3]\epsilon \\
&\qquad +[b(a^2-1) + b(1-5a^2+ab)\epsilon_0 + (b-2a+8a^2b-4ab^2)\epsilon_0^2 \\
&\qquad\quad +(a-b-4a^2b+4ab^2)\epsilon_0^3]\} = 0.
\end{aligned} \tag{4.17}$$

We emphasize that because of the nonhomogeneity of Q_5 it is no longer sufficient to only verify that the smallest root in ϵ of $Q_5(M_{\theta(a,b),0,\epsilon}) = 0$ is $\epsilon = 0$; we must show more generally that the smallest root in ϵ of $Q_5(M_{\theta(a,b),\epsilon_0,\epsilon}) = 0$ is $\epsilon = \epsilon_0$. Another feature of the lack of homogeneity is that it now becomes critical to remember the denominators $a+b$ and $2\epsilon-1$ in $M_{a,b}$ and A_ϵ^{-1} when computing (4.17), whereas they could have been neglected in the similar computations of (4.10) and (4.15). To see that the penultimate factor in (4.17) does not vanish for any $\epsilon \leq \epsilon_0$, we note that at $\epsilon = \epsilon_0$, this term, its derivative with respect to ϵ, and its second derivative with respect to ϵ are $ab(2\epsilon_0 - 1)^2$ (positive), $(2\epsilon_0 - 1)$ (negative), and 2 (positive), respectively. Hence this factor must be strictly positive for all $\epsilon \leq \epsilon_0$; a similar line of reasoning shows that the last term in (4.17) is strictly negative for all $\epsilon \leq \epsilon_0$. Therefore Q_5 is an effective (albeit nonhomogeneous) polynomial. We can now produce the estimators $\hat{\epsilon}^{(\Lambda)}$ and $\hat{\theta}^{(\Lambda)}$, as described above for the homogeneous cases, by using $Q_5(M_y A_\epsilon^{-1}) = 0$ to find $\hat{\epsilon}^{(\Lambda)}$ from the observed image y_Λ. [The convergence of the estimators is guaranteed by the corollary at the end of section 3.]

References

Barsky, D.J. (1993): Some combinatorial aspects of the Ehrenfest diffusion process. In preparation.

Barsky, D.J., Gandolfi, A. (1993): A generalized maximum pseudo-likelihood estimator for Markov random fields with noise, II: Two-dimensional examples. In preparation.

Baxter, R.J. (1982): Exactly solved models in statistical mechanics. Academic Press.

Besag, J. (1977): Efficiency of pseudo-likelihood estimation for simple Gaussian fields. *Biometrika*, **64**, 616-618.

Chalmond, B. (1987): Image restoration using an estimated Markov model. *Signal Processing*, **15**, 115-129.

Comets, F., Gidas, B. (1992): Parameter estimation for Gibbs distributions from partially observed data. *Ann. Appl. Prob.*, **2**, No. 1, 142-170.

Dempster, A.P., Laird, N.M., Rubin, D.B. (1977): Maximum likelihood for incomplete data via the EM algorithm. *J. R. Stat. Soc.*, Series B, **39**, 1-38.

Denny, J.L., Wright, A.L., (1991): Inference about the shape of neighboring points in fields. Spatial Statistics and Imaging, IMS Lecture Notes-Monograph Series, **20**, 46-54.

Frigessi, A., Piccioni, M. (1990): Parameter estimation for the two-dimensional Ising fields corrupted by noise. *Stoch. Proc. and Appl.*, **34**, 297-311

Geman, S., McClure, D.E. (1985): Bayesian image analysis: an application to single photon emission tomography. Proceedings of the Amer. Stat. Assoc., Statistical Computing Section, 12-18

Gi, C., Seymour, L. (1991): On the selection of Markov random fields texture models. Nimeo Series, **2062**, Dept. of Stat., Univ. of North Carolina at Chapel Hill.

Gidas, B. (1991): Parameter estimation for Gibbs distributions, I: Fully observed data. Markov Random Fields: Theory and Applications. (eds. R. Chellappa and A. Jain) Academic Press.

Meloche, J., Ruben, M.Z. (1992): Black and white image restoration. Preprint, Dept. of Stat., University of British Columbia, Tech Report #15.

Newman, C.M. (1987): Decomposition of binary random fields and zeros of partition functions. *Ann. Prob.*, **15**, No. 3, 1126-1130.

Ruelle, D. (1978): Thermodynamic Formalism. Addison-Wesley.

Younes, L. (1989): Parameter inference for imperfectly observed Gibbsian fields. *Prob. Th. and Rel. Fields*, **82**, 625-645.

Younes, L. (1991): Parameter estimation for imperfectly observed Gibbs fields and some comments on Chalmond's EM Gibbsian algorithm. Preprint.