

Unsupervised learning of binary vectors: A Gaussian scenario

Mauro Copelli*

Department of Chemistry and Biochemistry 0340, University of California San Diego, La Jolla, California 92093-0340

Christian Van den Broeck†

Limburgs Universitair Centrum, B-3590 Diepenbeek, Belgium

(Received 26 October 1999)

We study a model of unsupervised learning where the real-valued data vectors are isotropically distributed, except for a single symmetry-breaking binary direction $\mathbf{B} \in \{-1, +1\}^N$, onto which the projections have a Gaussian distribution. We show that a candidate vector \mathbf{J} undergoing Gibbs learning in this discrete space, approaches the perfect match $\mathbf{J} = \mathbf{B}$ exponentially. In addition to the second-order “retarded learning” phase transition for unbiased distributions, we show that first-order transitions can also occur. Extending the known result that the center of mass of the Gibbs ensemble has Bayes-optimal performance, we show that taking the sign of the components of this vector (clipping) leads to the vector with optimal performance in the binary space. These upper bounds are shown generally not to be saturated with the technique of transforming the components of a special continuous vector, except in asymptotic limits and in a special linear case. Simulations are presented which are in excellent agreement with the theoretical results.

PACS number(s): 87.10.+e, 64.60.Cn, 02.50.-r

I. INTRODUCTION

Since the introduction of the Ising spin model, the study of models with discrete degrees of freedom has become a core activity in statistical mechanics. When combined with disorder, such models often have interesting connections to problems of computational complexity, to learning theory or to open problems in statistics. Discreteness and disorder introduce intrinsic difficulties, and exactly solvable models are rare. The main purpose of this paper is to present a discrete model with disorder which can be solved in full detail. The model is most naturally presented as an unsupervised learning problem, and we briefly review the connection with the existing literature.

The goal of unsupervised learning is finding structure in high-dimensional data. In one of the simplest parametric models introduced in the literature [1–9], N -dimensional independently drawn data vectors $D = \{\xi^\mu\}$, $\mu = 1, \dots, \alpha N$, are uniformly distributed, except for a single symmetry-breaking direction \mathbf{B} . If we assume that all the relevant probability distributions are known, the aim of learning is to construct an estimate vector \mathbf{J} of the true direction \mathbf{B} .

Previous studies of this model focused on the case where \mathbf{B} is constrained to have a constant size, being otherwise equiprobably sampled from the N sphere. This so-called *spherical case* is associated with a spherical *prior distribution* $P_s(\mathbf{B}) \sim \delta(\mathbf{B} \cdot \mathbf{B} - N)$. The focus of the present paper, however, is on binary (or Ising) vectors. In this case, \mathbf{B} is known to have binary components only, $B_j \in \{-1, +1\}$, $j = 1, \dots, N$. This extra knowledge is taken into account by assigning a binary prior distribution

$$P_b(\mathbf{B}) = \prod_{j=1}^N \left[\frac{1}{2} \delta(B_j - 1) + \frac{1}{2} \delta(B_j + 1) \right] \quad (1)$$

to the preferential direction.

In this framework, a Gaussian scenario was introduced in Ref. [5] as a kind of minimal model, allowing the calculations to be much simplified and the spherical case to be solved exactly. In this model, the components of ξ perpendicular to \mathbf{B} are assumed to be independent Gaussian distributed variables with zero mean and unit variance, i.e., $P(b' \equiv \mathbf{B}' \cdot \xi / \sqrt{N}) = \exp(-b'^2/2) / \sqrt{2\pi}$, where $\mathbf{B}' \cdot \mathbf{B} / N = 0$. The distribution of the component $b \equiv \mathbf{B} \cdot \xi / \sqrt{N}$ parallel to \mathbf{B} , on the other hand, can be chosen at will, and in the Gaussian scenario it is completely determined by the mean B and variance $1 - A$:

$$\mathcal{P}(b) = \frac{\mathcal{N}}{\sqrt{2\pi}} \exp\left\{-\frac{b^2}{2} - U(b)\right\}, \quad (2)$$

$$U(b) = \frac{A}{2(1-A)} b^2 - \frac{B}{1-A} b, \quad (3)$$

where $\mathcal{N} = [\int \mathcal{D}b \exp(-U(b))]^{-1}$ is a normalization constant and $\mathcal{D}b \equiv db \exp[-b^2/2] / \sqrt{2\pi}$.

In comparison with the spherical case, the binary case presents several extra difficulties, which motivates the study of this simple model. The main question to be addressed in this work is given the αN data vectors (also called patterns) and the knowledge of the probability distributions, what is the best estimate \mathbf{J} one can construct to approximate \mathbf{B} ? The answer, cast in the framework of Bayesian inference, depends on whether \mathbf{J} is allowed to have continuous components or, conversely, is required to be a binary vector. We also address the problem of whether these upper bounds can be simply attained, by first obtaining a continuous vector via minimization of a potential and then transforming its components.

The results of the replica calculation for this problem are briefly reviewed in Sec. II. Section III discusses the special

*Electronic address: mauro@hypatia.ucsd.edu

†Electronic address: christian.vandenbroeck@luc.ac.be

case of Gibbs learning, for which simulations have been performed. In Sec. IV we review the reasoning leading to the Bayesian bound in the continuous as well as the binary space, with simulations compared to the theoretical results. A simple strategy which attempts to saturate these upper bounds is studied in Sec. V, while our conclusions are presented in Sec. VI.

II. UNSUPERVISED LEARNING

In order to obtain a good candidate vector \mathbf{J} , we construct a cost function of the form $\mathcal{H} = \sum_{\mu}^{\alpha N} V(\lambda_{\mu})$, where $\lambda_{\mu} \equiv \mathbf{J} \cdot \boldsymbol{\xi}^{\mu} / \sqrt{N}$. In the Gaussian scenario, the *potential* V has a quadratic form

$$V(\lambda) = \frac{c}{2} \lambda^2 - d\lambda. \quad (4)$$

Learning is defined as sampling \mathbf{J} from the Boltzmann distribution with temperature $T = 1/\beta$

$$P(\mathbf{J}|D) = \frac{P(\mathbf{J})}{Z} \exp - \beta \sum_{\mu}^{\alpha N} V(\lambda_{\mu}), \quad (5)$$

where $Z(D) = \int d\mathbf{J} P(\mathbf{J}) \exp - \beta \mathcal{H}$ is the normalization constant and the measure $P(\mathbf{J})$ is used to enforce either a binary [$P(\mathbf{J}) = P_b(\mathbf{J})$] or a spherical [$P(\mathbf{J}) = P_s(\mathbf{J})$] constraint on \mathbf{J} . While the spherical case has been dealt with in Ref. [5], we focus now on the case where the candidate vectors have binary components. The thermodynamic properties of such a system can be read from the free energy $f \equiv - (1/\beta N) \ln Z$. In the thermodynamic limit $N \rightarrow \infty$, f becomes self-averaging, $\langle f(D) \rangle_D = f$, and can be calculated via the replica trick. This by now standard calculation will not be reproduced here, only the results are quoted: for a replica symmetric ansatz, the quadratic forms of the Gaussian scenario allow the calculations to be performed exactly, and the free energy reads

$$\begin{aligned} f = \text{Extr}_{q, R, \hat{q}, \hat{R}} & \left\{ \frac{(1-q)\hat{q}}{2\beta} + \frac{\hat{R}R}{\beta} - \frac{1}{\beta} \int \mathcal{D}z \ln \cosh(z\sqrt{\hat{q}} + \hat{R}) \right. \\ & + \frac{\alpha}{2\beta} \ln[1 + c\beta(1-q)] \\ & \left. + \alpha \left[\frac{c[q - R^2(A - B^2)] - 2BRd - d^2\beta(1-q)}{2[1 + c\beta(1-q)]} \right] \right\}. \quad (6) \end{aligned}$$

The above order parameters are also self-averaging quantities and can be interpreted as follows: $R = \mathbf{B} \cdot \mathbf{J} / N$ measures the alignment between a typical (binary) sample of Eq. (5) and the preferential direction, its absolute value as a function of α being hereafter used to account for the performance of a given potential V ; $q = \mathbf{J} \cdot \mathbf{J}' / N$ is the mutual overlap between two different samples of Eq. (5), while \hat{R} and \hat{q} are the associated conjugate parameters. The equilibrium values of the four variables are determined by the solution of the saddle point equations which arise from the extremum operator in Eq. (6).

III. GIBBS LEARNING

Gibbs learning arises as a particular but very important case in this general framework. In order to define it properly, we first recall the Bayes inversion formula

$$p(\mathbf{B}|D) \equiv \frac{P(D|\mathbf{B})P_b(\mathbf{B})}{P(D)}. \quad (7)$$

The posterior distribution $p(\mathbf{B}|D)$ expresses the knowledge about \mathbf{B} which is gained after the presentation of the data. Replacing \mathbf{B} with \mathbf{J} in this formula gives the probability density that \mathbf{J} is the ‘‘true’’ direction \mathbf{B} , given the data vectors. Note that the binary prior in Eq. (7) constrains the acceptable candidates \mathbf{J} to the corners of the N hypercube, i.e., $\mathbf{J} \in \{-1, +1\}^N$. Making use of Eq. (2), one rewrites

$$p(\mathbf{J}|D) \propto P_b(\mathbf{J}) \prod_{\mu}^p \exp - U(\mathbf{J} \cdot \boldsymbol{\xi}^{\mu} / \sqrt{N}), \quad (8)$$

apart from a normalization constant. Gibbs learning is defined as sampling from distribution (8).

A comparison with Eq. (5) shows that the thermodynamics of such a process is obtained by setting $\beta V = U$ [3,4]. Upon substitution of $\beta = 1$, $c = A/(1-A)$ and $d = B/(1-A)$ in Eq. (6), one finds that the extremum of the corresponding free-energy is reached for $q_G = R_G$ and $\hat{q}_G = \hat{R}_G$, where the subscript G will hereafter be used to denote results from Gibbs learning. The equalities reflect the symmetric role played by \mathbf{J} and \mathbf{B} in Gibbs learning, a property which has been previously noted in several publications (see, e.g., Refs. [10] and [11], among others). The four original saddle point equations are then effectively reduced to a single one

$$R_G = F_B^2(\mathcal{F}(\sqrt{R_G})), \quad (9)$$

where

$$F_B(x) \equiv \sqrt{\int \mathcal{D}z \tanh(zx + x^2)} \quad (10)$$

is a function coming from the entropic term of the free energy, while

$$\mathcal{F}^2(\sqrt{R_G}) = \alpha \left[\frac{B^2 + AR_G(A - B^2)}{1 - AR_G} \right]. \quad (11)$$

The solution of Eq. (9) also determines the value of the conjugate parameter \hat{R}_G :

$$\hat{R}_G = \mathcal{F}^2(\sqrt{R_G}). \quad (12)$$

In order to check that the replica symmetric *ansatz* is correct, we also study the entropy $s \equiv \beta^2 (\partial / \partial \beta) [f - (\ln 2)/\beta]$, which for Gibbs learning reads

$$\begin{aligned} s_G = & - \frac{(1 + R_G)\hat{R}_G}{2} + \int \mathcal{D}z \ln 2 \cosh(z\sqrt{\hat{R}_G} + \hat{R}_G) \\ & - \frac{\alpha}{2} \left[\ln \left(\frac{1 - AR_G}{1 - A} \right) + (B^2 - A)(1 - R_G) \right]. \quad (13) \end{aligned}$$

On physical grounds, this quantity should always remain positive. Additionally, by relating s_G with the mutual information i per degree of freedom between the data D and the preferential direction \mathbf{B} , Herschkowitz and Nadal [9] show that it cannot decrease faster than linearly with α . For the Gaussian scenario, the inequality reads $s_G \geq \ln 2 - (\alpha/2)[B^2 - A - \ln(1-A)]$. Before we proceed to study in detail the solution of Eq. (9), we turn to the analysis of the asymptotic behavior of the system.

A. Asymptotics

The asymptotics of the solution of Eq. (9) can be immediately inferred by carrying out expansions of F_B and \mathcal{F} . In the vicinity of $R_G=0$, if we assume a smooth behavior for $R_G(\alpha)$, the predictions for the Gaussian scenario are

$$B \neq 0 \Rightarrow R_G \simeq \alpha B^2, \quad (14)$$

$$B = 0 \Rightarrow R_G \begin{cases} = 0, & \alpha \leq \alpha_G, \\ \simeq C_G(\alpha - \alpha_G), & \alpha \geq \alpha_G, \end{cases} \quad (15)$$

where

$$\alpha_G \equiv \frac{1}{A^2}, \quad (16)$$

$$C_G \equiv \frac{A^2}{1-A}.$$

We see that in the so-called biased case $B \neq 0$, it is much easier to learn. The unbiased case $B = 0$ presents much more difficulties for information about vector \mathbf{B} to be extracted, due to the intrinsic symmetry $\mathbf{B} \rightarrow -\mathbf{B}$. In this case, retarded learning occurs [2,3], meaning that a nonzero macroscopic overlap R_G will be obtained only after a critical number of examples $\alpha_G N$ is presented. For $\alpha \leq \alpha_G$, the entropy saturates its linear bound exactly [9]. This second order phase transition is identical to the one obtained in the spherical case [5,6], revealing that the binary nature of the preferential direction plays no role in the poor performance regime.

In the limit $\alpha \rightarrow \infty$, on the other hand, the differences with respect to the spherical case become pronounced: R_G approaches 1 exponentially,

$$1 - R_G \simeq \sqrt{\frac{\pi(1-A)}{2\alpha(B^2(1-A) + A^2)}} e^{-\alpha(B^2 + A^2/1-A)/2}, \quad (17)$$

as opposed to the power law observed for the spherical case. Equation (17) also implies an exponential decay to the entropy $s_G \simeq \alpha(1 - R_G)[B^2(1-A) + A^2]/[2(1-A)]$.

These qualitative asymptotic results can be shown to hold for general distributions $\mathcal{P}(b)$ [12]. In the following, we explore the Gaussian scenario in more detail, studying the behavior of $R_G(\alpha)$ away from the asymptotic regimes.

B. The biased case

The first case to be addressed is $A = 0$ with $B \neq 0$. The nonzero bias makes sure learning starts off as soon as α




FIG. 1. Overlap R_G (left axis) as a function of α' for $A=0$ [see Eq. (18)]: theory (solid line) and simulations with $N=100$ (symbols; error bars represent one standard deviation, see text for details). The dashed line represents the entropy (right axis) while the dotted line shows the linear bound (right axis).

≥ 0 , while $A = 0$ eliminates the dependence of \hat{R}_G on R_G [see Eqs. (11) and (12)], much simplifying the saddle point equations, which can be solved exactly. The behavior of R_G is seen to be simply determined by the rescaled variable

$$\alpha' \equiv \alpha B^2, \quad (18)$$

namely, $R_G = F_B^2(\sqrt{\alpha'})$. This function can be seen in Fig. 1. It shows a linear increase for small α' and an exponential behavior for $\alpha' \rightarrow \infty$. The entropy saturates its linear bound only in the limit $\alpha' \rightarrow 0$, approaching zero exponentially when $\alpha' \rightarrow \infty$ but remaining otherwise strictly positive.

Note that $A = 0$ means that b has unit variance. The patterns can thus be pictured as being distributed in an N -dimensional spherically symmetric cloud, whose displacement B from the origin conveys the information about \mathbf{B} .

Simulations. Binary disordered systems are known to be very hard to simulate due to the existence of very many local minima. A noisy dynamics with unity temperature and general cost function U will typically get stuck in one of these minima, preventing a proper sampling of the posterior distribution (8) in an acceptable time. The Gaussian scenario with $A = 0$ provides an exception to this rule, allowing Gibbs learning to be very easily implemented with a simple Metropolis algorithm [13]. Since $A = 0$ implies a linear function $U(\lambda)$, the changes in energy can be very quickly calculated because it depends only on $\mathbf{J} \cdot \sum_{\mu} \xi^{\mu}$.

Figure 1 shows the results for simulations with $N = 100$ (the smallest system size simulated) and two values of B , checking the relevance of the variable α' . For each pattern set D , ten samples of R_G and q_G were measured, after a random initialization of the system and a warming up of the dynamics (see further details below). The whole procedure was repeated for 1000 pattern sets and the standard deviation was calculated over these 10 000 samples.

The measurement of q_G during simulations is another tool to check both the property $q_G = R_G$ and the correctness of the RS ansatz. Figure 2 focus on the second simulated point of Fig. 1 ($\alpha' = 1$). It shows histograms for both R_G and q_G (measured between pairs of consecutive samples) which are




FIG. 2. Gaussian scenario with $A=0$, $B=1$, and $\alpha=1$. Histograms of q_G (thick lines) and R_G (thin lines) for $N=100$ (dashed) and $N=1000$ (solid); the vertical line is the theoretical prediction. The upper inset shows the METROPOLIS R_G dynamics for two pattern sets (same legend, see text for details). The lower inset presents the variance of the distribution for R_G (symbols) as a function of $1/N$ for $N=100, 500, 750$, and 1000 ; the dotted line is a linear fit of the three leftmost points.

virtually indistinguishable on the scale of the figure, with a mean value in excellent agreement with the theoretical prediction. The upper inset gives a glimpse of the METROPOLIS dynamics: the system is initialized randomly at $t=0$ and evolves up to $t=50$ Monte Carlo steps per site (MCS/site), at which moment a different pattern set is drawn. The system reaches thermal equilibrium after $\mathcal{O}(10)$ MCS/site, which motivated the choice of safely waiting 100 MCS/site during the simulations before any measurement was made. The system was reinitialized after every measurement of the overlaps. Note that some pattern sets yield time-averaged values of R_G which deviate from theory (notably the first one for $N=100$ and the second one for $N=1000$) and only a second average over the pattern sets gives the right results. This reflects the property of self-averaging, which only holds in the thermodynamic limit (note that deviations from theory are smaller for larger N). The lower inset shows the typical scaling with $1/\sqrt{N}$ of the width of the distribution of overlaps.

C. The unbiased case

When $B=0$, retarded learning is expected to occur, according to Eq. (15). Figure 3 shows the solution of the R_G saddle point equation for two values of A , namely, 0.6 and -0.6 . In both cases, a second-order phase transition occurs at the critical value α_G predicted by Eq. (16) and the entropy saturates exactly the linear bound before the phase transition. Based on the relation between s_G and i [9], the retarded learning phase transition can be interpreted as follows: for $\alpha \leq \alpha_G$, the system extracts maximal information from each pattern but is nonetheless unable to obtain a nonzero alignment R_G with the preferential direction \mathbf{B} . Only at $\alpha = \alpha_G$ does R_G depart from zero, which on its turn immediately gives an increasing degree of redundancy (measured by the deviation of s_G from its linear bound) to the patterns coming thereafter. Figure 3 also shows the effect of a small bias B




FIG. 3. R_G (thick lines, left axis) and s_G (thin lines, right axis) as functions of α for $A = \pm 0.6$ and $B=0$. The thin dotted lines correspond to the linear bound, which is exactly saturated up to the second order phase transition. A small bias $B=0.1$ (with $A=0.6$) breaks the symmetry and destroys the second order phase transition (thick dotted line, left axis, entropy not shown).

$=0.1$ in an otherwise symmetric distribution: for sufficiently large α (say, $\alpha \gg \alpha_G$), the effect is negligible, but for small α the broken symmetry destroys the second order phase transition.

It is interesting to note in Fig. 3 that even though the phase transition for $A = -0.6$ and $A = 0.6$ occurs at the same critical value, the overlap increases much slower in the former case than in the latter. Recalling the definition of A [Eqs. (2),(3)], this means that prolate Gaussian distributions (N -dimensional ‘‘cigars’’ [2]) convey less information about the preferential direction than oblate distributions (N -dimensional ‘‘pancakes’’ [2]) for the same absolute value of A .

However, the second order phase transition at $\alpha_G = A^{-2}$ is not the only interesting phenomenon for this model. First order phase transitions are also possible, depending on the value of A . They can occur in two situations: either for $\alpha > \alpha_G$, in which case two consecutive phase transitions take place during learning (a second-order one followed by a first-order one) or $\alpha < \alpha_G$, in which case the asymptotic result Eq. (15) is overridden. The first-order phase transition appears when there is more than one solution to the saddle point equation. In such cases the solution with minimal free energy has maximal probability of occurrence, being thus the thermodynamically stable state. Such first-order phase transitions have been found for the spherical case with a two-peaked distribution [7], but not with the Gaussian scenario [5], which shows that they are due to the discrete nature of the search space in this case.

An overview of this phenomenology is presented in Fig. 4. It shows the three typical behaviors that occur for $B=0$. For comparison, the case $A=0.6$ plotted in Fig. 3 is shown again, as an example of a parameter region where there is only a second order phase transition (at $\alpha_G = 2.78$). For $A = 0.78$ the second-order phase transition at $\alpha_G = 1.643$ is followed by a first-order phase transition at $\alpha_G^{(f)} = 1.704$ (upper inset, lower axis), while for $A = 0.85$ only a first-order phase transition takes place at $\alpha = 1.27$, overriding the second-order phase transition at $\alpha = 1.38$ (lower inset) which was




FIG. 4. Solutions R_G of the saddle point Eqs. (9) and (11) as a function of α for $B=0$ and three values of A . s_G is plotted with dashed lines and the thermodynamically stable solutions are plotted with thick lines. $A=0.78$ (upper inset): R_G (left axis) vs α (bottom axis) and s_G (right axis) vs α (top axis—note the different α scale, which zooms in the first-order phase transition); $A=0.85$ (lower inset): R_G (left axis) and s_G (right axis) vs α .

predicted on asymptotics and smoothness grounds. Note that none of these first-order phase transitions can be predicted by the asymptotic expansion, Eq. (15). It is also interesting to observe that some solutions of the saddle point equation may violate the linear bound and/or the positivity of the entropy (notably $A=0.85$ in Fig. 4). However, it turns out that these branches are always thermodynamically unstable, while the stable solutions satisfy all the requirements.

The whole phase diagram for $B=0$ is shown in Fig. 5. For $A > A_1 \approx 0.773$, a first-order phase transition takes place at the line $\alpha_G^{(f)}(A)$, after the second-order one has already occurred. For increasing A , $\alpha_G^{(f)}(A)$ gets closer and closer to $\alpha_G(A)$, until there is finally a collapse at $A = A_2 \approx 0.808$. For larger values of A , only the first-order phase transition occurs.

IV. OPTIMAL LEARNING: THE BAYESIAN PERSPECTIVE

We now switch to the following question: given the αN data vectors and the prior information about \mathbf{B} , what is the




FIG. 5. Phase diagram for the unbiased case ($B=0$). Second order phase transitions occur at $\alpha = \alpha_G = A^{-2}$ (solid line), while first order phase transitions take place at $\alpha = \alpha_G^{(f)}$ (dashed line). See text for details.

best performance R one could possibly attain with a vector \mathbf{J} ? Watkin and Nadal [3] answered this question in a Bayesian framework by defining *optimal learning* (see also Refs. [14,15]). We briefly review their reasoning here and extend it to take into account the binary nature of the vectors.

We define the quality measure $Q(\mathbf{B}, \mathbf{J}) \equiv \mathbf{B} \cdot \mathbf{J} / N$, which quantifies how well \mathbf{B} is approximated by any candidate vector \mathbf{J} satisfying $\mathbf{J} \cdot \mathbf{J} = N$. Since \mathbf{B} is unknown, Q is formally inaccessible. But one can take its average with respect to the posterior distribution (7), leading to $\tilde{Q}(\mathbf{J}, D) \equiv \int d\mathbf{B} Q(\mathbf{B}, \mathbf{J}) p(\mathbf{B}|D)$. \tilde{Q} is then a formally accessible *bona fide* quantity which can be used to measure the performance of \mathbf{J} .

Optimal learning is *defined* as constructing a vector \mathbf{J}_B which maximizes \tilde{Q} . The linearity of Q in \mathbf{J} immediately implies

$$\tilde{Q}(\mathbf{J}, D) = N^{-1} \mathbf{J} \cdot \int d\mathbf{B} \mathbf{B} p(\mathbf{B}|D), \quad (19)$$

leading on its turn to

$$\mathbf{J}_B = \frac{1}{\sqrt{R_G}} \int d\mathbf{B} \mathbf{B} p(\mathbf{B}|D), \quad (20)$$

where the $\sqrt{R_G}$ factor guarantees the proper normalization of \mathbf{J}_B . This is the so-called Bayesian vector, which is the center of mass of the Gibbs ensemble. In the thermodynamic limit, its performance $R_B \equiv \mathbf{B} \cdot \mathbf{J}_B / N$ is shown to be simply related to that of Gibbs learning [15,3]: $R_B = \sqrt{R_G}$.

A. The best binary

Up to now the reasoning is fairly general. The whole procedure can actually be carried out without explicitly mentioning what the prior distribution $P(\mathbf{B})$ is. For clarity, in the following \mathbf{J}_B will specifically denote the Bayesian vector for a *binary* prior. Note, however, that despite being the center of mass of the ensemble of *binary* vectors sampled from the posterior distribution, \mathbf{J}_B has real-valued components [12], in general.

One would therefore like to address the next question: what is the *best binary* vector one can construct? In other words, what is the binary vector—inferable from the data—that outperforms—on average—any other binary vector in approximating \mathbf{B} ? The answer is again straightforward [16]: the vector \mathbf{J}_{bb} which maximizes \tilde{Q} among the binary vectors is simply obtained by the clipping prescription, namely $[\mathbf{J}_{bb}]_j = \text{sgn}([\mathbf{J}_B]_j)$, $j = 1, \dots, N$ or, in shorthand notation,

$$\mathbf{J}_{bb} = \Phi_{\text{clip}}(\mathbf{J}_B). \quad (21)$$

This can be easily checked by noting that the quantity to be maximized [the right-hand side of Eq. (19)] is proportional to $\sum_{j=1}^N J_j [\mathbf{J}_B]_j$. In what follows, \mathbf{J}_{bb} is called the best binary vector. Summarizing, if \mathbf{B} is known to be binary, \mathbf{J}_B is the best estimator one can provide. But if the estimator is required to be binary as well, then \mathbf{J}_{bb} is the optimal choice.

The proof that \mathbf{J}_B and \mathbf{J}_{bb} are optimal estimators in their respective spaces, is relatively simple [4]. What we show below is that maximal \tilde{Q} implies maximal alignment R with

\mathbf{B} , in the thermodynamic limit. For the best binary, one departs from the inequality $\tilde{Q}(\mathbf{J}_{bb}) - \tilde{Q}(\mathbf{J}) \geq 0$, $\forall \mathbf{J} \in \{-1, +1\}^N$ and takes the average with respect to the data distribution

$$\int dD P(D) [\tilde{Q}(\mathbf{J}_{bb}) - \tilde{Q}(\mathbf{J})] = \int d\mathbf{B} P_b(\mathbf{B}) \int dD P(D|\mathbf{B}) \times \left[\frac{\mathbf{B} \cdot \mathbf{J}_{bb}}{N} - \frac{\mathbf{B} \cdot \mathbf{J}}{N} \right] \geq 0. \quad (22)$$

If we now assume that $R_{bb} \equiv \mathbf{B} \cdot \mathbf{J}_{bb}/N$ and $R \equiv \mathbf{B} \cdot \mathbf{J}/N$ are self-averaging, the average over the data can be by-passed and one obtains

$$\int d\mathbf{B} P_b(\mathbf{B}) [R_{bb} - R] \stackrel{N \rightarrow \infty}{\geq} 0. \quad (23)$$

Finally, one notices that $P_b(\mathbf{B})$ is a uniform prior, making no distinction between any particular binary vector (this is reflected, for instance, in the free energy ϕ being independent of the particular choice of \mathbf{B}): this allows the last average in Eq. (23) to be by-passed as well, leading to the stronger upper bound $R_{bb} \geq R$.

B. Performance and simulations

The performance R_{bb} of the best binary vector can be explicitly calculated by extending previously obtained results for the clipping prescription [17,18]. In Ref. [18], Schietse *et al.* study the effect of a general transformation $\tilde{J}_j = \sqrt{N} \phi(J_j) / \sqrt{\sum_i \phi^2(J_i)}$ on the components of a properly normalized continuous vector \mathbf{J} satisfying $\mathbf{B} \cdot \mathbf{J}/N = R$. If \mathbf{B} is binary, ϕ is odd and $\tilde{\mathbf{R}} \equiv \mathbf{B} \cdot \tilde{\mathbf{J}}/N$ is self-averaging, then the following relation follows:

$$\tilde{R} = \frac{\int P(x) \phi(x) dx}{\left[\int P(x) \phi^2(x) dx \right]^{1/2}}, \quad (24)$$

where the variable $x \equiv B_1 J_1$ is expected to be distributed independently of the index, because of the permutation symmetry among the axes.

We are left then with the problem of calculating $P(x)$, after which Eq. (24) can be applied for $\phi(x) = \text{sgn}(x)$, providing R_{bb} as a function of R_G . If \mathbf{J} is uniformly distributed on the cone $\mathbf{B} \cdot \mathbf{J} = NR$, then $P(x)$ is just a Gaussian with mean R and variance $1 - R^2$ [18]. However, this can hardly be expected to hold for the Bayesian vector, since it is a sum of Ising vectors. One would naively expect \mathbf{J}_B to be closer to the corners of the N -hypercube instead. To obtain the relevant $P_{\text{c.m.}}(x) \equiv P(x = B_1 [\mathbf{J}_B]_1)$, we calculate the m th quenched moment of $[\mathbf{J}_B]_1$,

$$\langle ([\mathbf{J}_B]_1)^m \rangle_D = \frac{1}{R_G^{m/2}} \left\langle Z^{-m} \left(\int d\mathbf{J} P_b(\mathbf{J}) J_1 e^{-\beta \sum_{\mu} U(\lambda_{\mu})} \right)^m \right\rangle_D. \quad (25)$$




FIG. 6. Probability distribution of $y = B_1 \langle J_1 \rangle_J$ for different values of R_G .

The above expression can be evaluated again with the use of the replica trick. The replica symmetric result is

$$\langle ([\mathbf{J}_B]_1)^m \rangle_D = \frac{1}{[R_G(\alpha)]^{m/2}} \int \mathcal{D}z \{ \tanh[z \sqrt{\hat{R}_G(\alpha)} + \hat{R}_G(\alpha) B_1] \}^m \quad (26)$$

or, equivalently,

$$\langle x^m \rangle_D = \frac{1}{[R_G(\alpha)]^{m/2}} \int \mathcal{D}z \{ \tanh[z \sqrt{\hat{R}_G(\alpha)} + \hat{R}_G(\alpha)] \}^m, \quad (27)$$

where the values of $R_G(\alpha)$ and $\hat{R}_G(\alpha)$ should be taken at the solution of the saddle point Eqs. (9) and (12) for Gibbs learning. Notice that the passage from Eq. (26) to Eq. (27) is valid *only* if \mathbf{B} is binary. From Eq. (27) one immediately rewrites the probability distribution $P_{\text{c.m.}}(x)$, by identifying a change of stochastic variables $x = R_G^{-1/2} \tanh(z \sqrt{\hat{R}_G} + \hat{R}_G)$, with z normally distributed:

$$P_{\text{c.m.}}(x) = \frac{\sqrt{R_G}}{\sqrt{2\pi \hat{R}_G} (1 - R_G x^2)} \times \exp \left[\frac{-1}{2\hat{R}_G} \left[\frac{1}{2} \ln \left(\frac{1 + \sqrt{R_G} x}{1 - \sqrt{R_G} x} \right) - \hat{R}_G \right]^2 \right]. \quad (28)$$

Note that since R_G and \hat{R}_G are simply related to each other [Eqs. (9) and (12)], $P_{\text{c.m.}}(x)$ can always be parametrized in function of R_G only. In Fig. 6, the probability distribution of $y \equiv x \sqrt{R_G}$ is plotted for different values of R_G , illustrating the fact that $|y| \leq 1$.

Equation (28) should be compared to the Gaussian distribution obtained in Ref. [18]. It shows that \mathbf{J}_B is indeed closer to the corners of the N hypercube, optimally incorporating the information that \mathbf{B} is binary.

We have run simulations for $A=0$ and $B=1$ as described in Sec. III B. For a system size $N=500$, the center of mass was constructed with $n=50$ samplers, being normalized afterwards. Each component of \mathbf{B} and \mathbf{J}_B was used to measure x , the procedure being repeated 100 times for each of the 100




FIG. 7. Gaussian scenario with $A=0$ and $B=1$. Probability distribution of $x=B_1[\mathbf{J}_B]_1$ for $\alpha=0.25$, $\alpha=0.6$ (left plot), and $\alpha=1.17$ (right plot) or, correspondingly, $R_G=0.2$, $R_G=0.4$ and $R_G=0.6$. Solid line: simulations. Dashed line: theory [Eq. (28)].

pattern sets. A comparison between the resulting histogram and the theoretical prediction can be seen in Fig. 7. The good agreement shows that Eq. (28) correctly describes the statistical properties of the Bayesian vector.

We can finally proceed to calculate the performance R_{bb} of the best binary vector. Making use of Eqs. (24) and (28), we make a change of variables to obtain

$$R_{bb} = \int P_{\text{c.m.}}(x) \text{sgn}(x) dx = 1 - 2H[\sqrt{\hat{R}_G(\alpha)}], \quad (29)$$

where $H(x) \equiv \int_x^\infty \mathcal{D}z$. Making use of the relation between R_G and \hat{R}_G , we finally write

$$R_{bb} = 1 - 2H[F_B^{-1}(R_B)] = 1 - 2H[F_B^{-1}(\sqrt{R_G})] \quad (30)$$

$$= 1 - 2H[\mathcal{F}(R_B)] = 1 - 2H(\sqrt{\hat{R}_G}), \quad (31)$$

where F_B^{-1} is the inverse of F_B .

Equation (30) expresses an upper bound for binary candidate vectors \mathbf{J} in approximating \mathbf{B} , satisfying two obvious inequalities $R_G \leq R_{bb} \leq R_B$. The asymptotic behavior of R_{bb} can always be written in terms of R_G : in the poor performance regime ($R_G \rightarrow 0$), one recovers previous results for clipping a spherical vector [18,19], $R_{bb} \approx \sqrt{2R_G/\pi}$, while in the large α regime ($R_G \rightarrow 1$) a faster exponential decay is achieved than with Gibbs learning $1 - R_{bb} \approx (2/\pi)(1 - R_G)$ [12].

As a spin-off of the calculation, we have also obtained the overlap between the Bayesian vector and its clipped counterpart, $\Gamma \equiv \mathbf{J}_B \cdot \mathbf{J}_{bb} / N = N^{-1} \sum_j^N |[\mathbf{J}_B]_j|$. This quantity can be easily computed:

$$\Gamma = \frac{1}{R_B} \int \mathcal{D}z |\tanh[z\sqrt{\hat{R}_G(\alpha)} + \hat{R}_G(\alpha)]| = \frac{R_{bb}}{R_B}, \quad (32)$$

if one notices the counterintuitive identity $\int \mathcal{D}z |\tanh(za + a^2)| = \int \mathcal{D}z \text{sgn}(za + a^2)$, $\forall a$, which is proved in the Appendix. The simple result $\Gamma = R_{bb}/R_B$ immediately implies the equality $(\mathbf{J}_{bb} - \Gamma \mathbf{J}_B) \cdot (\mathbf{J}_B - R_B \mathbf{B}) = 0$, for which we still have not found a deeper interpretation.




FIG. 8. R_{bb} , R_B , and Γ : comparison between simulations and theory, for different values of α . Error bars represent one standard deviation. The diagonal is plotted for comparison.

The curves R_{bb} , R_B , and Γ as functions of R_G are plotted on Fig. 8, together with results for simulations with the same parameters as those of Fig. 7. The data is in excellent agreement with the theoretical results, errors generally remaining below the margin of one standard deviation. Note that $\Gamma \geq \sqrt{2/\pi}$, with equality holding only for $R_G \rightarrow 0$. This result is another confirmation of the picture that \mathbf{J}_B lies closer to the binary vectors, since $\sqrt{2/\pi}$ is the overlap between a continuous vector isotropically sampled from the N hypersphere and its clipped counterpart. Figure 9 zooms in the fourth column of points of Fig. 8 ($\alpha=1.17$), showing the histograms of the overlaps. One observes that the distribution of Γ is much sharper than the other ones, while the statistics of R_G is better because it has $n=50$ times more samples.

V. TRANSFORMING THE COMPONENTS

A. General results

Since sampling the binary Gibbsian vectors is usually a very difficult task, the construction of the Bayesian vector according to the center of mass recipe is not always possible, in practice. Alternative methods should therefore be developed for approximating the R_B and R_{bb} performances. One




FIG. 9. Histograms of the overlaps R_G , R_B , R_{bb} , and Γ for $\alpha=1.17$. The vertical lines show the theoretical predictions.

such method is the technique of transforming the components of a previously obtained *spherical* vector, as described in Sec. IV B. In the following, we first derive general results (for any function U) and subsequently look at the Gaussian scenario in detail.

A natural choice for the vector to be transformed is $\mathbf{J}_{\text{opt}}^s$, which can be obtained by minimizing—in the N hypersphere—an optimally constructed cost function [5–8] $\mathcal{H} = \sum_{\mu} V_{\text{opt}}^s(\lambda_{\mu})$. It attains the Bayes-optimal performance $R_{\text{opt}}^s \equiv \mathbf{B} \cdot \mathbf{J}_{\text{opt}}^s / N$ for the spherical case, which satisfies

$$R_{\text{opt}}^s = F_s(\mathcal{F}(R_{\text{opt}}^s)), \quad (33)$$

where $F_s(x) \equiv x/\sqrt{1+x^2}$. Note that $\mathbf{J}_{\text{opt}}^s$ saturates the performance of the center of mass of the Gibbs ensemble *for a spherical prior*. Equation (33) should be compared to the performance of \mathbf{J}_B , which obeys

$$R_B = F_B(\mathcal{F}(R_B)). \quad (34)$$

While $R_{\text{opt}}^s \simeq R_B$ for small α , the differences between \mathbf{J}_B and $\mathbf{J}_{\text{opt}}^s$ are clearly manifested in the asymptotic behavior for large α , with R_{opt}^s approaching unity with a power law [5,6] instead of the exponential decay of Eq. (17).

We would like to depart from $\mathbf{J}_{\text{opt}}^s$ and obtain approximations to both \mathbf{J}_{bb} and \mathbf{J}_B . The first one is obtained by clipping, $\mathbf{J}_{\text{clip}}^s \equiv \Phi_{\text{clip}}(\mathbf{J}_{\text{opt}}^s)$. The second one relies on an optimal transformation [18] $\phi^*(x) \equiv [P(x) - P(-x)]/[P(x) + P(-x)]$, which maximizes the transformed overlap.¹ The vector obtained by such a transformation on $\mathbf{J}_{\text{opt}}^s$ is denoted by \mathbf{J}_*^s , thus $[\mathbf{J}_*^s]_j = \phi^*([\mathbf{J}_{\text{opt}}^s]_j)/R_*^s$, $j = 1, \dots, N$, where $R_*^s \equiv \mathbf{B} \cdot \mathbf{J}_*^s / N$.

Since $\mathbf{J}_{\text{opt}}^s$ contains no information about the binary nature of \mathbf{B} , the results of Schietse *et al.* can be directly applied to render $R_{\text{clip}}^s \equiv \mathbf{B} \cdot \mathbf{J}_{\text{clip}}^s / N$ and R_*^s . In this case, $P(x)$ is Gaussian and one obtains [18]

$$R_{\text{clip}}^s = 1 - 2H(\mathcal{F}(R_{\text{opt}}^s)) \quad (35)$$

$$R_*^s = F_B(\mathcal{F}(R_{\text{opt}}^s)). \quad (36)$$

We would like to compare Eqs. (30) and (34) with Eqs. (35) and (36), respectively. Despite their resemblance in form, one notices that the former should be solved, while the latter just map the solution of Eq. (33). In order to compare the equations, one should first note that $F_B(x) \geq F_s(x)$, $\forall x \geq 0$. Since $\partial \mathcal{F} / \partial R \geq 0$, in general $R_B \geq R_{\text{opt}}^s$. This result in turn immediately implies the inequalities

$$R_{\text{clip}}^s \leq R_{bb}, \quad (37)$$

$$R_*^s \leq R_B, \quad (38)$$

with equality holding for both equations in the asymptotic limits $\alpha \rightarrow \infty$ and $R_G \rightarrow 0$. This general behavior is confirmed in Fig. 10, which shows the results for the Gaussian scenario in the two relevant cases: zero and nonzero bias.

¹Consistently, the optimal transformation for $P(x) = P_{\text{c.m.}}(x)$ becomes $\phi^*(x) \propto x$, that is, no improvement is possible.




FIG. 10. Overlaps as functions of α for two choices of parameters in the Gaussian scenario: $A = 1/3$ with $B = 0.1$ (upper curves) and $A = -1/3$ with $B = 0$ (lower curves). The upper bounds R_B (solid) and R_{bb} (dashed) are depicted with thick lines, while the approximations R_*^s (solid) and R_{clip}^s (dashed) are plotted with thin lines.

Another available measure of the success of the optimal transformation ϕ^* in rendering a good approximation for \mathbf{J}_B , is the probability distribution $P(x_*)$, where $x_* \equiv \phi^*(x)/R_*^s$. In order to obtain $P(x_*)$, one just has to recall that $P(x)$ is Gaussian with mean R_{opt}^s and variance $1 - (R_{\text{opt}}^s)^2$. The optimal transformation is then $x_* = \phi^*(x)/R_*^s = \tanh\{R_{\text{opt}}^s x / [1 - (R_{\text{opt}}^s)^2]\} / R_*^s$, and can be regarded as an attempt to attach some structure to the distribution of the transformed x_* . With a simple change of variables, $P(x_*)$ is readily seen to be

$$P(x_*) = \frac{R_*^s \sqrt{1 - (R_{\text{opt}}^s)^2}}{\sqrt{2\pi} R_{\text{opt}}^s [1 - (R_*^s x_*)^2]} \times \exp \left\{ \frac{-[1 - (R_{\text{opt}}^s)^2]}{2(R_{\text{opt}}^s)^2} \left[\frac{1}{2} \ln \left(\frac{1 + R_*^s x_*}{1 - R_*^s x_*} \right) - \frac{(R_{\text{opt}}^s)^2}{1 - (R_{\text{opt}}^s)^2} \right]^2 \right\}. \quad (39)$$

A comparison with Eq. (28) shows that the two equations are very similar, but not identical. Some similarity in shape should indeed be expected, mainly because $P(x_*)$, just like $P_{\text{c.m.}}(x)$, must be such that $[P(x) - P(-x)]/[P(x) + P(-x)] \propto x$, in order to consistently prevent any further improvement by a similar transformation. One can verify in Fig. 11 that the resemblance between the probability distributions is closely associated with the success of R_*^s in saturating the upper bound R_B . The curves correspond to the Gaussian scenario with $A = 1/3$ and $B = 0.1$ for two values of α (one can thus refer to the upper solid curves of Fig. 10). Note that for $\alpha = 8$, the difference between R_B and R_*^s is very small in Fig. 10, which is reflected in the solid curves of Fig. 11 being very close to each other. Accordingly, the dashed curves in Fig. 11 get further apart for $\alpha = 10$ as the mismatch between the overlaps increase in Fig. 10.




FIG. 11. Distributions $P_{c.m.}(x)$ (thick) and $P(x_*)$ (thin) according to Eqs. (28) and (39), respectively. The values $\alpha=8$ (solid) and $\alpha=10$ (dashed) refer to the Gaussian scenario with $A=1/3$ and $B=0.1$ (see Fig. 10).

B. The biased case

The simple biased case with $A=0$ and $B \neq 0$ provides an interesting exception to the performances of \mathbf{J}_*^s and \mathbf{J}_{clip}^s . The fact that $U(b)$ is linear implies $\partial F / \partial R = 0$, as can be readily verified in Eq. (11). This, on the other hand, implies the equalities in Eqs. (37) and (38), that is,

$$R_{clip}^s = R_{bb}, \quad (40)$$

$$R_*^s = R_B. \quad (41)$$

Therefore the strategy described in the previous section is successful in attaining the upper bounds of Sec. IV, and not only asymptotically. It should be noted that for a linear U , the vector \mathbf{J}_{opt}^s can be simply constructed with the Hebbian rule, $\mathbf{J}_{opt}^s \propto \sum_{\mu} \alpha^N \xi^{\mu}$, $\forall \alpha$. Therefore the best binary performance is attainable by the clipped Hebbian vector \mathbf{J}_{clip}^s , in this case. The second equality, however, seems to us more remarkable, because it establishes a result which we could not find elsewhere in the literature: the optimal transformation manages to completely incorporate the information about the binary nature of \mathbf{B} , leading to the Bayes-optimal performance R_B without the need of explicitly constructing the center of mass of the Gibbs ensemble. In other words, the technique of non-linearly transforming the components of the vectors, introduced in Ref. [17] and extended in Ref. [18], is able to give a definitive answer to the problem it aims to solve.

VI. CONCLUSIONS

We have presented results on learning a binary preferential direction from disordered data. Constraining the candidate vectors to the binary space as well, we first showed that

Gibbs learning presents not only an exponential asymptotic decay and second-order “retarded learning” phase transitions, but also first-order phase transitions for a simple Gaussian scenario.

On the question of what is the optimal estimator, given the data and the knowledge that the preferential direction is binary, we have shown that the answer depends on which space the estimator \mathbf{J} is allowed to lie in. The best continuous estimator is the Bayesian vector, which is the center of mass of binary Gibbsian vectors. The best binary vector is obtained by clipping the Bayesian vector. We have calculated its properties in detail, providing an upper bound to the performance of binary vectors.

Finally, we have also studied one possible way of constructing approximations to these two optimal estimators. By transforming the components of a previously obtained continuous vector, we show that the upper bounds cannot be saturated, in general. Exceptions to this rule are the asymptotic limits (both $R_G \rightarrow 0$ and $\alpha \rightarrow \infty$) and the special case of a linear function U . Interestingly, the linear case also seems to be the only one in which Gibbs sampling can be performed without computational difficulties. We are therefore left with a situation where the approximations work perfectly only in the case where they are actually not needed. We believe this kind of result reinforces the need of investigating the connection between results in statistical mechanics and computational complexity theory.

ACKNOWLEDGMENTS

We would like to thank Manfred Opper for very fruitful discussions. We also acknowledge financial support from FWO Vlaanderen and the Belgian IUAP program (Prime Minister’s Office). This work was supported in part by the Engineering Research Program of the Office of Basic Energy Sciences at the U.S. Department of Energy under Grant No. DE-FG03-86ER13606.

APPENDIX

In order to show that Eq. (32) is correct, one just has to show that the integral below vanishes identically, $\forall a$:

$$\begin{aligned} & \int \mathcal{D}z [\text{sgn}(az+a^2) - |\tanh(az+a^2)|] \\ &= \int \mathcal{D}z \text{sgn}(az+a^2) [1 - \tanh(az+a^2)] \\ &= e^{-a^2/2} \int \frac{dy}{\sqrt{2\pi}} e^{-(y-a)^2/2} \text{sgn}(ay) [1 - \tanh(ay)] \\ &= e^{-a^2/2} \int \mathcal{D}y \text{sgn}(ay) e^{ay} \left[\frac{e^{-ay}}{\cosh(ay)} \right] = 0. \quad (\text{A1}) \end{aligned}$$

- [1] M. Biehl and A. Mietzner, Europhys. Lett. **24**, 421 (1993).
 [2] M. Biehl and A. Mietzner, J. Phys. A **27**, 1885 (1994).
 [3] T. L. H. Watkin and J.-P. Nadal, J. Phys. A **27**, 1899 (1994).
 [4] P. Reimann and C. Van den Broeck, Phys. Rev. E **53**, 3989

- (1996).
 [5] P. Reimann, C. Van den Broeck, and G. J. Bex, J. Phys. A **29**, 3521 (1996).
 [6] C. Van den Broeck and P. Reimann, Phys. Rev. Lett. **76**, 2188

- (1996).
- [7] A. Buhot and M. B. Gordon, *Phys. Rev. E* **57**, 3326 (1998).
- [8] M. B. Gordon and A. Buhot, *Physica A* **257**, 85 (1998).
- [9] D. Herschkowitz and J.-P. Nadal, *Phys. Rev. E* **59**, 3344 (1999).
- [10] G. Györgyi, *Phys. Rev. A* **41**, 7097 (1990).
- [11] M. Opper, *Phys. Rev. Lett.* **72**, 2113 (1994).
- [12] M. Copelli, C. Van den Broeck, and M. Opper, *J. Phys. A* **32**, L555 (1999).
- [13] K. Binder and D. W. Heermann, *Monte Carlo Simulation in Statistical Physics: An Introduction* (Springer-Verlag, Berlin, 1988).
- [14] M. Opper and D. Haussler, *Phys. Rev. Lett.* **66**, 2677 (1991).
- [15] T. L. H. Watkin, *Europhys. Lett.* **21**, 871 (1993).
- [16] T. L. H. Watkin, A. Rau, and M. Biehl, *Rev. Mod. Phys.* **65**, 499 (1993).
- [17] D. Bollé and G. M. Shim, *Network Comput. Neural Syst.* **6**, 619 (1995).
- [18] J. Schietse, M. Bouten, and C. Van den Broeck, *Europhys. Lett.* **32**, 279 (1995).
- [19] C. Van den Broeck and M. Bouten, *Europhys. Lett.* **22**, 223 (1993).