



ACADÉMIE  
DES SCIENCES  
INSTITUT DE FRANCE

# *Comptes Rendus*

---

## *Physique*


Guilhem Semerjian

**Some observations on the ambivalent role of symmetries in Bayesian inference problems**

Volume 26 (2025), p. 199-216

Online since: 11 February 2025

<https://doi.org/10.5802/crphys.236>

 This article is licensed under the  
CREATIVE COMMONS ATTRIBUTION 4.0 INTERNATIONAL LICENSE.  
<http://creativecommons.org/licenses/by/4.0/>



*The Comptes Rendus. Physique are a member of the  
Mersenne Center for open scientific publishing*  
[www.centre-mersenne.org](http://www.centre-mersenne.org) — e-ISSN : 1878-1535



Review article / *Article de synthèse*

# Some observations on the ambivalent role of symmetries in Bayesian inference problems

## *Quelques observations sur le rôle ambivalent des symétries dans les problèmes d'inférence Bayésienne*

Guilhem Semerjian <sup>a</sup>

<sup>a</sup> Laboratoire de physique de l'École Normale Supérieure, ENS, Université PSL, CNRS, Sorbonne Université, Université Paris Cité, F-75005 Paris, France  
*E-mail:* [guilhem.semerjian@phys.ens.fr](mailto:guilhem.semerjian@phys.ens.fr)

**Abstract.** We collect in this note some observations on the role of symmetries in Bayesian inference problems, that can be useful or detrimental depending on the way they act on the signal and on the observations. We emphasize in particular the need to gauge away unobservable invariances in the definition of a distance between a signal and its estimator, and the consequences this implies for the statistical mechanics treatment of such models, taking as a motivating example the extensive rank matrix factorization problem.

**Résumé.** Nous regroupons dans cette note quelques observations sur le rôle des symétries dans les problèmes d'inférence Bayésienne, qui peuvent être utiles ou néfastes selon la façon dont elles agissent sur le signal et les observations. Nous soulignons en particulier la nécessité de quotienter les invariances non-observables dans la définition d'une distance entre un signal et son estimateur, et les conséquences que ceci implique pour le traitement par la mécanique statistique de tels modèles, motivé par l'exemple de la factorisation de matrices de rang extensif.

**Keywords.** Bayesian inference problems, Symmetries, Statistical mechanics of disordered systems.

**Mots-clés.** Problèmes d'inférence Bayésienne, Symétries, Mécanique statistique des systèmes désordonnés.

*Manuscript received 21 November 2024, revised 14 January 2025, accepted 17 January 2025.*

## 1. Introduction

The notion of symmetry is a fascinating concept with a strong aesthetical appeal that influenced several branches of science; for instance in mathematics it led to the development of representation and invariant theories [1, 2], while modern physics, and in particular high-energy physics, used it as a guiding principle to build the standard model of particle physics [3]. It also played a role in statistical mechanics, through notably the classification of phase transitions according to the symmetry group of the order parameters, and its spontaneous breaking. In the context of disordered systems Gérard Toulouse, to whom this article is dedicated, formulated his views on glasses and spin-glasses in this group-theoretic language [4, 5]. The focus of this note is on the effects of symmetries in inference and machine learning problems, and in particular on their statistical mechanics treatment. This topic has a long history rooted in classical statistics [6–8] and in the early days of the development of neural networks [9], and has seen a more recent revival

triggering a burst of activity on the role of invariances in modern architectures of machine learning [10, 11]. These invariances can be hard-wired in the form of equivariant networks [12] (generalizing the convolutional networks that are translationally invariant to arbitrary symmetries), or exploited to perform “data augmentation” [13], namely creating artificial new samples by applying random transformations respecting the assumed invariance to observed data points. The consequences of symmetries on the loss landscape of artificial neural networks have also been studied, see for instance [14].

We shall not attempt in this note an exhaustive review of this rich and growing field of research, but instead concentrate on a specific setting, namely the Bayesian inference problems and their statistical mechanics treatment, and collect a series of simple observations on the role symmetries plays in this context. We will argue that this role can be either beneficial or detrimental depending, roughly speaking, on the relative strength of the action of the symmetry on the quantity to be inferred and on the one which is observed. A particular emphasis shall be put on the negative side of this dichotomy, where the non-trivial unobservable invariance of the signal imposes to reconsider the definition of the distance between the signal and its estimation. This point was of course noticed in previous works, but we believe its consequences may have been underappreciated, in particular for the extensive-rank matrix factorization problem whose invariance group is growing with the system size.

The rest of the manuscript is organized as follows. In Section 2 we define some notions related to Bayesian inference, to symmetries, and to the interplay between the two. Section 3 is then devoted to situations where symmetries are useful in the inference process, while Section 4 discuss the opposite case where they are detrimental. Further perspectives on possible ways to handle these difficulties are given in the conclusions of Section 5.

## 2. Definitions

### 2.1. Bayesian inference

In an inference problem a signal of interest, denoted  $S$ , cannot be measured directly but has to be determined from indirect observations, called  $Y$  in the following. In general  $S$  and  $Y$  can have a rich structure, and should be thought of as vectors of arbitrary dimensions. The Bayesian perspective on inference treats  $S$  and  $Y$  as random variables; one way to describe their distributions is as follows. The signal  $S$  is drawn from a prior law  $P_S(s) = \mathbb{P}(S = s)$  (we follow here the convention of denoting random variables with uppercase symbols, the values of their realizations with the corresponding lowercase letters, and assume for notational simplicity that the random variables take discrete values, we will later generalize to continuous ones). The observation  $Y$  is then drawn conditionally on the signal, with the probability  $P_{Y|S}(y|s)$ , which can model the addition of noise to the true signal, or the flow of information contained in  $S$  through an imperfect communication channel. From a probabilistic point of view it is however equivalent, and sometimes more convenient, to think of the pair of variables  $(S, Y)$  as drawn with the joint probability law  $P_{S,Y}(s, y) = P_S(s)P_{Y|S}(y|s)$ .

The goal of the observer is to propose an estimate  $\widehat{S}(y)$  of the true value of the signal, based on the observation of the event  $Y = y$ , and assuming the knowledge of the probability laws underlying this generation process, namely the prior distribution  $P_S$  and the noisy channel  $P_{Y|S}$ . In general the observer cannot be sure of the hidden value of the signal, and can only assign an a posteriori probability to its possible values; according to the Bayes theorem this posterior distribution  $P_{S|Y}$  reads

$$P_{S|Y}(s|y) = \frac{P_S(s)P_{Y|S}(y|s)}{P_Y(y)}, \quad (1)$$

on which the computation of  $\widehat{S}(y)$  has to be based.

The tools of statistical mechanics of disordered systems have been applied to a wide variety of such problems [15–19]. As a matter of fact the probability law (1) can be viewed as the Gibbs–Boltzmann distribution for thermal variables  $s$ , with an Hamiltonian  $H_y(s) = -\ln P_S(s) - \ln P_{Y|S}(y|s)$ , where  $y$  plays the role of a quenched disorder (with a so-called “planted” distribution), at inverse temperature  $\beta = 1$ . In this perspective the denominator in (1) is viewed as a partition function ensuring the normalization of the law,  $P_Y(y) = Z(y) = \sum_s P_S(s) P_{Y|S}(y|s)$ , and one defines the quenched free-energy as  $\mathbb{E}[\ln Z(Y)]$  (which is the opposite of the Shannon entropy of  $Y$  and can be related, up to a simple additive constant, to the mutual information between  $S$  and  $Y$ ).

Let us present a few examples that will be used for illustration in the forthcoming sections:

- The stochastic block model (SBM) is an ensemble of random graphs that was introduced to study complex networks with a community structure [20–22]: its nodes (or vertices) belong to “communities” and the probability of existence of an edge between a pair of nodes (modelling for instance a social interaction between two individuals) depends on the communities of the two nodes. In this context the signal  $S$  corresponds to the communities of each node, and the observation  $Y$  to the graph structure (i.e. the edges connecting some nodes), and the inference goal is to determine the hidden community structure from the knowledge of the graph. More precisely, denoting  $N$  the number of vertices and  $q$  the number of communities, one has  $S = (\sigma_1, \dots, \sigma_N)$  where  $\sigma_i \in \{1, \dots, q\}$  is a label specifying the community of the  $i$ -th node. The prior law  $P_S$  is a product measure, each  $\sigma_i$  is drawn independently at random from the probability law  $\bar{\eta} = (\bar{\eta}_1, \dots, \bar{\eta}_q)$ . The observed graph  $Y$  is then generated by including each of the  $N(N-1)/2$  possible edges  $\langle i, j \rangle$  between pairs of distinct vertices with probability  $c_{\sigma_i, \sigma_j} / N$ , where  $c$  is a  $q \times q$  matrix with non-negative real elements  $c_{\sigma, \tau}$ , the latter measuring the “affinity” between individuals of communities  $\sigma$  and  $\tau$ , and hence their tendency to create bonds between them. This problem was studied with statistical mechanics techniques in [23], unveiling phase transitions in the thermodynamic (large size) limit  $N \rightarrow \infty$ : depending on the parameters  $\bar{\eta}$  and  $c$  it might be possible or not to infer  $S$  from  $Y$  with an accuracy higher than a guess from the prior probability. Moreover the computational cost of this task can go from polynomial to exponential in  $N$ , opening the possibility of a phenomenon known as a statistical-computational gap when the algorithmic and information-theoretical phase transitions do not coincide. Parts of this scenario were proven rigorously later on, see [24, 25] for reviews of this research direction.
- In matrix denoising problems the signal and observation  $S$  and  $Y$  are both matrices; in the simplest version they are real and symmetric, of size  $n \times n$ , and the observation channel corresponds to the addition of an independent noise matrix  $Z$ , also symmetric, namely  $Y = S + Z$ . The matrix character of the problem shows up in the structure of the laws of  $S$  and  $Z$ , which for many natural ensembles of Random Matrix Theory (RMT), and most famously the Gaussian Orthogonal Ensemble (GOE), enjoy some invariance properties, for instance under orthogonal transformations. This inference problem was in particular studied in the seminal paper [26], and revisited more recently in [27–32].
- A related problem is the matrix factorization one, where one observes a matrix  $Y = XX^T + Z$ , with  $X$  an  $n \times r$  matrix and  $Z$  a noise, and the goal is to infer  $X$  from  $Y$  (to avoid confusion with the matrix denoising problem the signal is denoted  $X$  in this case). The matrix denoising problem can thus be seen as a simplified version of the matrix factorization one, in which one only aims at reconstructing the product of the two unknown factors (i.e. the Wishart matrix  $XX^T$ ) instead of the factor  $X$  individually. Matrix factorization appears in various problems of machine learning [33–36], and has been the subject of an intense research activity. The low-rank regime of the matrix factorization

problem, when  $r$  remains finite in the large  $n$  limit, is now rather well understood [37–42], with a precise characterization of the information-theoretical optimal accuracy of estimation, as well as the performances of polynomial time algorithms for this task. By contrast the extensive rank regime, with  $r$  proportional to  $n$  in the large size limit, presents additional challenges that are not yet all resolved. The initial proposal of [43, 44] was later unveiled to be incorrect [45], and despite a series of recent efforts [27–30, 46–51] the current understanding of this regime is not satisfactory. We will propose in this paper some observations that may shed some light on the roots of these difficulties.

- In the machine learning context one often faces variants of the following situation [52]: from  $n$  examples of an input–output relation  $(x^{(i)}, z^{(i)})$  with  $i \in [n] \equiv \{1, \dots, n\}$  one would like to generalize and predict the output  $z^{(\text{new})}$  corresponding to an input  $x^{(\text{new})}$  not already seen. In a classical example of such supervised learning problems the inputs are high-dimensional objects, i.e.  $x \in \mathbb{R}^d$  with  $d \gg 1$ , which could be the vectorial representation of a high-resolution image, and the output is low-dimensional, e.g.  $z \in \mathbb{R}$  being a label characterizing the type of image in a classification task. Assuming the existence of a function  $f$  such that  $z^{(i)} = f(x^{(i)}, W)$ , with  $W$  some unknown parameters, and of prior distributions on the  $x$ 's and  $W$ , this generalization problem falls under the Bayesian inference setting of a correlated pair  $(S, Y)$  of random variables, the role of the signal being played by  $S = z^{(\text{new})}$ , and of the observations by  $Y = (x^{(\text{new})}, \{(x^{(i)}, z^{(i)})\}_{i \in [n]})$ .

## 2.2. Optimal estimators

We shall return now to the generic setting of a signal-observation pair of random variables  $(S, Y)$ , and discuss the constructions of estimators  $\widehat{S}$ , which should be functions from the observation space to the signal space, such that  $\widehat{S}(Y)$  is “as close as possible” to the true signal  $S$ . Of course to give a precise meaning to this notion we need to introduce a distance  $d(S, \widehat{S})$  on the signal space that measures the closeness of the signal to its estimation (we will not require  $d$  to be a distance in the mathematical sense, namely we only ask that  $d$  is non-negative and that  $d(S, \widehat{S}) = d(\widehat{S}, S)$ , without imposing the definiteness and triangular inequality for  $d$ ). The Bayesian risk of an estimator is defined as its average distance to the groundtruth signal,

$$R(\widehat{S}) = \mathbb{E}[d(S, \widehat{S}(Y))], \quad (2)$$

the average being over the joint law of  $(S, Y)$  discussed above. A Bayes-optimal (BO) estimator, for a given choice of the distance  $d$ , is then one which minimizes the risk,  $\widehat{S}^{\text{BO}} \in \operatorname{arginf} R(\widehat{S})$ . Decomposing the expression (2) of the risk as an average over  $Y$  and a conditional average over  $S|Y$  (with the law described in Equation (1)), one can perform this minimization for each point of the observation space independently, and obtain formally the expression of the optimal estimator as

$$\widehat{S}^{\text{BO}}(Y) \in \operatorname{arginf}_x \mathbb{E}[d(S, x)|Y], \quad (3)$$

with  $x$  running on the signal space. We insist on the fact, that will play an important role in the following, that the notion of optimality depends on the choice of  $d$ , i.e. on the way one measures the accuracy of estimation. Two special cases are worth spelling out more explicitly:

- suppose that the signals are vectors of discrete coordinates, i.e.  $S \in \chi^N$  with  $\chi$  finite (in the SBM model one would have  $\chi = \{1, \dots, q\}$ ). Then a natural choice for  $d$  is the Hamming distance,  $d(S, \widehat{S}) = \sum_{i=1}^N \mathbb{1}(S_i \neq \widehat{S}_i)$ , that counts the number of indices for which  $\widehat{S}$  gives a wrong answer. With this choice one can easily find from (3) that the optimal estimator is, for each coordinate  $i \in [N]$ ,  $\widehat{S}_i^{\text{BO}}(Y) \in \operatorname{argmax}_{x \in \chi} \mathbb{P}(S_i = x|Y)$ . This is called the symbol Maximal A Posteriori (MAP) estimator in the coding theory literature, since it selects the value with the maximum marginal probability in the posterior measure.

- if the signals are continuous,  $S \in \mathbb{R}^N$ , a common choice for  $d$  is the square of the  $L^2$  norm, that we shall call the Square Error (SE) distance,

$$d(S, \hat{S}) = \|S - \hat{S}\|_2^2 = \sum_{i=1}^N (S_i - \hat{S}_i)^2, \quad (4)$$

which gives the Mean Square Error (MSE) expression for the risk,  $R(\hat{S}) = \mathbb{E}[\|S - \hat{S}(Y)\|_2^2]$ . The optimal estimator is then well-known to be the posterior mean  $\hat{S}^{\text{BO}}(Y) = \mathbb{E}[S|Y]$ ; this result is easily derived from (3) by noting that the optimization can be performed independently for each  $i \in [N]$ , and that the latter can be done explicitly:

$$\hat{S}_i^{\text{BO}}(Y) \in \underset{x \in \mathbb{R}}{\operatorname{arginf}} \mathbb{E}[(S_i - x)^2|Y], \quad \mathbb{E}[(S_i - x)^2|Y] = \mathbb{E}[S_i^2|Y] - 2x\mathbb{E}[S_i|Y] + x^2, \quad (5)$$

the minimum of this quadratic function on  $x$  being easily found to be reached in  $x = \mathbb{E}[S_i|Y]$ .

In the continuous case this choice of  $d$  can be viewed as rather arbitrary: one could take more generically  $d(S, \hat{S}) = \sum_i |S_i - \hat{S}_i|^p$ , with no compelling reason to favour  $p = 2$  with respect to other values of  $p$ . The very important advantage of the choice  $p = 2$  is however the simple analytic expression it yields for the corresponding optimal estimator, namely  $\hat{S}^{\text{BO}}(Y) = \mathbb{E}[S|Y]$ . For generic values of  $p$  the optimal estimator can be seen as a Fréchet mean of the posterior distribution, but it does not admit an explicit formula in general (for  $p = 1$  the conditional median replaces the conditional mean, for generic  $p$  [53] discusses specific priors such that the optimal estimator is linear in  $Y$ ). We will come back later on this point, arguing that symmetries may force one to abandon the usual square error distance for  $d$ , and as a consequence loose the simplicity of the corresponding optimal estimator.

### 2.3. Symmetries

The notion of symmetry is formalized mathematically in terms of a group, namely a set  $G$  endowed with an associative multiplication rule admitting a neutral element  $e$  with  $eg = ge = g$  for all  $g \in G$ , and such that each element  $g$  admits an inverse  $g^{-1}$  with  $gg^{-1} = g^{-1}g = e$ . A group  $G$  is said to act on a set  $U$  if there exists an application that associates to each  $g \in G$  and  $u \in U$  an element of  $U$  denoted  $g \cdot u$ , interpreted as the image of  $u$  under the transformation  $g$ , that respects the multiplication rule of the group, in the sense that  $g \cdot (h \cdot u) = (gh) \cdot u$ , and such that the neutral element of  $G$  leaves  $U$  unchanged,  $e \cdot u = u$ . Linear representations are an important special case of such group actions, when  $U$  is a vector space and the group elements  $g$  acts on  $U$  by application of linear operators  $\rho(g)$ . A trivial action can always be defined for any group on any set, by  $g \cdot u = u$ . One defines the orbit of an element  $u$  under the action of a group  $G$  by  $G \cdot u = \{g \cdot u : g \in G\}$ , that contains all possible images of  $u$  under the transformations of  $G$ . A function  $f : U \rightarrow V$ , where  $U$  is endowed with an action of  $G$ , is said to be invariant if  $f(g \cdot u) = f(u)$  for all  $g \in G$  and  $u \in U$ . If in addition the set  $V$  also admits an action of the same group  $G$  (that we shall denote with the same symbol  $\cdot$  even if in general the underlying applications are distinct), one says that the function  $f$  is equivariant (or covariant) if  $f(g \cdot u) = g \cdot f(u)$ . The definition of invariant functions is thus a special case of the equivariant ones, when the action of  $G$  is trivial on  $V$ . The characterization of invariant and equivariant functions, and most notably polynomials, is an important branch of mathematics known as invariant theory [2], with many results in particular when  $G$  is the symmetric group of permutations acting on the coordinates of a function of several variables [54], or other classical groups like the orthogonal one [55].

#### 2.4. Symmetries of inference problems

Putting together the two ingredients introduced above, we shall say that a group  $G$  is a symmetry of an inference problem if it acts on both the signal and observations spaces in such a way that  $(g \cdot S, g \cdot Y) \stackrel{d}{=} (S, Y)$  for all  $g \in G$ , with  $\stackrel{d}{=}$  denoting equality in distribution of random variables. Let us insist on the facts that in this definition the same  $g$  acts simultaneously on the signal and on the observations, and that even if one uses the same symbol  $\cdot$  the two actions are in general distinct (even the two spaces of signal and observations are generically different). In terms of the two laws  $P_S$  and  $P_{Y|S}$  we used above to define the random variables, this definition amounts to  $P_S(g \cdot s) = P_S(s)$  and  $P_{Y|S}(g \cdot y | g \cdot s) = P_{Y|S}(y | s)$  for all  $g \in G$ .

In the next two Sections we shall explore the consequences of such symmetries in inference problems, arguing that they can be helpful in some cases, or detrimental in others. In colloquial terms, the fuzzy boundary between these two sides depends on the “relative strength” of the actions on  $G$  on the signal and observations spaces: as an oversimplified rule of thumb, the more  $G$  acts on the observations (i.e. the bigger the orbits it induces in this space), the better, while on the contrary the more  $G$  acts on the signal, the worst. We will give more flesh to these imprecise statements in the rest of the manuscript, at this point an intuitive justification can be provided as follows: the action of  $G$  on  $S$  can be interpreted as a blurring of the sought-for signal, hence as a loss of information, if it does not come along with an action on the observations  $Y$ , that are the only directly accessible sources of knowledge on  $S$ .

### 3. Symmetries as an asset

In all this Section we assume that the signal is continuous,  $S \in \mathbb{R}^N$  (even if in some cases  $S$  can have an internal structure, like matrix indices, it can always be seen as a vector), and that the quality of estimation is measured in terms of the Square Error distance  $d$  defined in Equation (4). We further assume that a group  $G$  is a symmetry of the inference problem, in the sense defined above, and that its action on the signal space is a linear isometric representation (i.e.  $g \cdot S = \rho(g)S$ , with  $\rho(G)$  an  $N \times N$  orthogonal matrix).

As discussed previously this choice of  $d$  implies that the Bayes-optimal estimator is the posterior mean,  $\hat{S}^{\text{BO}}(Y) = \mathbb{E}[S|Y]$ . It is a simple exercise to check that the symmetry hypothesis leads to the equivariance of this optimal estimator, i.e.  $\hat{S}^{\text{BO}}(g \cdot Y) = g \cdot \hat{S}^{\text{BO}}(Y)$ . If the problem is sufficiently simple for the conditional mean to be computable exactly (or at least asymptotically exactly, for instance in a large dimension limit), this equivariance property is merely a sanity check of the correctness of the result. There are however problems where a direct computation of  $\hat{S}^{\text{BO}}$  is challenging; we shall describe now an alternative approach, known as the low-degree polynomial strategy [31, 42, 56], that can be useful in such cases, and where the symmetries play a very helpful role.

The idea of this method is to look for an approximation of  $\hat{S}^{\text{BO}}$  in a variational way, considering the estimators of the form

$$\hat{S}(Y) = \sum_{\beta \in \mathcal{A}} c_\beta b_\beta(Y), \quad (6)$$

where the  $b_\beta$  are simple functions from the observation space to  $\mathbb{R}^N$ ,  $\mathcal{A}$  is a finite index set, and the  $c_\beta$  are free parameters in this linear combination. The latter are chosen to minimize the risk of  $\hat{S}$ , which corresponds here to the MSE  $R(\hat{S}) = \mathbb{E}[\|S - \hat{S}(Y)\|_2^2]$ . Inserting the variational expression (6) one obtains  $R(\hat{S}) = \mathbb{E}[\|S\|_2^2] + c^T \mathcal{M} c - 2\mathcal{R}^T c$ , where  $c = \{c_\beta\}_{\beta \in \mathcal{A}}$  and  $\mathcal{R}$  are vectors of size  $|\mathcal{A}|$ ,  $\mathcal{M}$  is a square matrix of the same size, the elements of  $\mathcal{R}$  and  $\mathcal{M}$  being given by

$$\mathcal{M}_{\beta, \beta'} = \mathbb{E}[\langle b_\beta(Y), b_{\beta'}(Y) \rangle], \quad \mathcal{R}_\beta = \mathbb{E}[\langle S, b_\beta(Y) \rangle], \quad (7)$$

where  $\langle S, T \rangle = \sum_{i=1}^N S_i T_i$  is the canonical scalar product in  $\mathbb{R}^N$ . The optimal estimator, within this restricted variational space of linear combinations of trial functions indexed by  $\mathcal{A}$ , is thus obtained by minimizing the quadratic function of the free parameters  $c$ , yielding the linear system of equations  $\mathcal{M}c = \mathcal{R}$ . In general it is not too difficult to compute the elements of  $\mathcal{M}$  and  $\mathcal{R}$ , the complexity of this approach can thus be much lower than the direct computation of the conditional mean. In practice the determination of the elements of  $\mathcal{M}$  and  $\mathcal{R}$  is reduced to the computation of (joint) moments of  $S$  and  $Y$  when the basic functions  $b_\beta$  of the decomposition are taken as polynomials in the observations (assuming  $Y$  is a vector of real random variables). This explains the name “low-degree polynomial” of the method (which initially emerged in the context of hypothesis testing problems, see [57, 58]).

It should be clear that the quality of the approximate estimator  $\widehat{S}$  obtained in this way improves (i.e. that its MSE decreases) if one includes more and more functions in the trial set  $\{b_\beta\}_{\beta \in \mathcal{A}}$ ; at the same time the complexity of the computation, and in particular the size of the linear system of equations  $\mathcal{M}c = \mathcal{R}$ , also increases along the way, requiring to settle for a compromise between these two considerations. It turns out that symmetries have a very beneficial consequence here: one does not pay any price in terms of quality, and may gain a lot in terms of complexity, by focusing on equivariant basic functions  $b_\beta$ . This statement is quite intuitive: since our objective was to emulate  $\widehat{S}^{\text{BO}}$ , which is itself equivariant, it is not so surprising that its best approximation is generated by linear combinations of equivariant functions. One can give a formal proof of this fact (see Appendix B.3 in [42], where it is called the Hunt-Stein lemma, and Section III.C of [31]), by first showing that for any estimator the risk cannot increase under symmetrization:  $R(\widehat{S}^{\text{eq}}) \leq R(\widehat{S})$ , where  $\widehat{S}^{\text{eq}}$  is the equivariant function defined by

$$\widehat{S}^{\text{eq}}(Y) = \mathbb{E}_g[g^{-1} \cdot \widehat{S}(g \cdot Y)], \tag{8}$$

the expectation being on the invariant (Haar) measure on  $G$ . The mapping  $\widehat{S} \rightarrow \widehat{S}^{\text{eq}}$  defined in (8) is a linear transformation in the space of functions from observations to signals, that projects onto the subspace of equivariant functions. Since  $R(\widehat{S}^{\text{eq}}) \leq R(\widehat{S})$  the MSE reached by a variational family  $\{b_\beta\}_{\beta \in \mathcal{A}}$  cannot be strictly better than the one of  $\{b_\beta^{\text{eq}}\}_{\beta \in \mathcal{A}}$ ; the latter can nevertheless be of a strictly smaller cardinality, if originally distinct  $b_\beta$  are projected to the same element by the symmetrization (8), hence a reduction of the computational cost with no loss in estimation accuracy.

Let us give a few examples of this phenomenon:

- The first one is almost trivial but has some pedagogical virtue. Consider indeed a scalar denoising problem: both signal and observation are real random variables, and the observation channel is  $Y = S + Z \in \mathbb{R}$  with the noise  $Z$  independent of  $S$ . The Bayes-optimal estimator  $\widehat{S}^{\text{BO}}(Y) = \mathbb{E}[S|Y]$  has no closed form expression in general (an exception arising when  $S$  and  $Z$  are both Gaussian), one can thus try a polynomial approximation of degree  $D$  with  $\widehat{S}(Y) = c_0 + c_1 Y + \dots + c_D Y^D$ . According to the generic formulation above the optimal value of the coefficients  $c$  will be obtained from the solution of  $\mathcal{M}c = \mathcal{R}$ , where  $\mathcal{M}$  is a  $(D + 1) \times (D + 1)$  matrix with elements  $\mathcal{M}_{p,p'} = \mathbb{E}[Y^{p+p'}]$ , and  $\mathcal{R}$  the vector with  $\mathcal{R}_p = \mathbb{E}[SY^p]$ . Suppose now that the signal and noise have even distributions, in the sense that  $S \stackrel{d}{=} -S$ ,  $Z \stackrel{d}{=} -Z$ . As a consequence  $(S, Y) \stackrel{d}{=} (-S, -Y)$ , which shows that the multiplicative group  $G = \{-1, 1\}$ , acting on the signal and observation space by multiplication, is a symmetry of the inference problem. The equivariant functions in this toy model are odd in the usual sense (they are such that  $f(-Y) = -f(Y)$ ), the symmetry considerations allows us to reduce a priori the variational functions to the odd monomials and work with  $\widehat{S}(Y) = c_1 Y + c_3 Y^3 + \dots$ ; equivalently, the even monomials vanish under the symmetrization by Equation (8). In other words the use of the symmetry divides the number of the trial functions by 2, for the same


maximum degree  $D$ , and hence the same quality of approximation. Of course in this trivial case one could have considered all polynomials, and discovered very easily that  $c_2 = c_4 = \dots = 0$  in the minimizers, but in more complicated situations the identification of the relevant terms is not so obvious, and the gain in the size of the trial set can be much more substantial.

- Such non-trivial simplifying aspects of symmetries have been exploited in the low-rank matrix factorization problem in [42], and in the matrix denoising one in [31]. Recall that in the latter case one has also  $Y = S + Z$ , but now these objects are  $n \times n$  matrices, and the objective is to send the size  $n$  to infinity. In this context a polynomial estimator  $\widehat{S}(Y)$  is a matrix whose elements  $\widehat{S}(Y)_{i,j}$  are multivariate polynomials in the entries of  $Y$ , i.e. in  $Y_{1,1}, Y_{1,2}, \dots, Y_{n,n}$ . One thus faces the difficulty that, even for a fixed maximal degree  $D$ , the number of such monomials increases with  $n$ , and as a consequence the size  $|\mathcal{A}|$  of the matrix  $\mathcal{M}$  and vector  $\mathcal{R}$  of the variational approach diverges in the large  $n$  limit, making the whole procedure quite impractical. Fortunately symmetry considerations allow to escape this deadend: the probability laws of several classical ensembles of RMT are invariant under the conjugation by the orthogonal group (i.e. under  $M \rightarrow OMO^T$  for all  $O$  with  $OO^T = O^T O = \mathbb{1}$ ). If  $S$  and  $Z$  enjoy this invariance in law then the orthogonal group becomes a symmetry of the inference problem. One can then restrict the polynomial estimators  $\widehat{S}(Y)$  to the equivariant ones, without loosing on the accuracy according to the generic discussion above. This is actually a strong requirement (in colloquial terms, the action of this group induces “big enough” orbits) and the number of equivariant polynomials is drastically reduced with respect to the number of arbitrary ones. One can indeed show that the equivariant polynomials of degree at most  $D$  are linear combinations of

$$Y^p (\text{Tr } Y)^{q_1} (\text{Tr } Y^2)^{q_2} \dots (\text{Tr } Y^D)^{q_D}, \quad (9)$$

for non-negative integers  $p, \{q_i\}$ , the total degree of the polynomial being  $p + \sum_i i q_i \leq D$ . Note that the dimension of this subspace of equivariant polynomials depends on  $D$  but not on  $n$ , which allows to perform the large size limit in the variational approach with  $\mathcal{M}$  and  $\mathcal{R}$  of fixed sizes. The hypothesis of invariance under orthogonal conjugation is violated for some ensembles of random matrices, like the arbitrary Wigner ones with i.i.d. non-Gaussian matrix elements, that however preserves a weaker subgroup of the orthogonal one, namely their laws are invariant by permutation of the row and column indices. The equivariant polynomials under the permutation group are more numerous than those under the orthogonal group written in (9), but their number is still independent of  $n$  for a fixed degree, and they can be indexed by suitable finite graphs; see [31, 42] for more details on this point, and [59] for an extension to the tensor case.

- Symmetries can also play a simplifying role in the machine learning setting defined in Section 2, where  $S = z^{(\text{new})}$  is to be inferred from  $Y = (x^{(\text{new})}, \{(x^{(i)}, z^{(i)})\}_{i \in [n]})$ . In this context it is indeed natural to assume that the function  $z = f(x)$  is invariant under the action of a group  $G$  on the inputs  $x$ : for image classification the label  $z$  should not change if the input image  $x$  is translated, or slightly deformed. In the language of the present paper this means that  $G$  is a symmetry of the inference problem with a trivial action on the signal and a non-trivial action on the observations. This is somehow the ideal situation, in the sense that the reduction of the number of candidate variational functions is the strongest possible and collapse to a projection on the invariant functions. There is in addition a permutation symmetry on the indices  $i$  of the sample pairs  $\{(x^{(i)}, z^{(i)})\}_{i \in [n]}$  that can be further exploited. This phenomenon was studied and put on a quantitative basis in [60], through the definition of the *degeneracy* of a group in terms of



**Figure 1.** A sketch illustrating the definition of the quotiented distance of Equation (10).

the asymptotic reduction in the number of invariant polynomials compared to arbitrary ones. To be precise the study of [60] does not treat the Bayes-optimal estimation, but shows the reduction of sample complexity or network width induced by the degeneracy for some specific architectures (random features and ridge kernel regression).

#### 4. Symmetries as a nuisance

We will now describe somehow opposite situations, in which the presence of a symmetry induces difficulties in the inference problems. According to the informal discussion above this is to be expected when the action of  $G$  is stronger on the signal than on the observation. We shall in fact make in this section the drastic assumption that  $G$  acts trivially on the observations, i.e. that  $(g \cdot S, Y) \stackrel{d}{=} (S, Y)$  for all  $g \in G$ , with a non-trivial action on  $S$ . In terms of the laws defining the process this means that the prior is invariant,  $P_S(g \cdot s) = P_S(s)$ , and that  $P_{Y|S}(y|g \cdot s) = P_{Y|S}(y|s)$ . It should be clear that the “signal” that can be recovered in such a situation is not  $S$  itself, but rather the orbit  $G \cdot S = \{g \cdot S : g \in G\}$ , since all the points of this orbit generate observations with the same law, and are thus perfectly indistinguishable from the knowledge available to the observer. An example that is visually easy to grasp is provided in Figure 1: the signal space is here  $S = (S_1, S_2) \in \mathbb{R}^2$ , and  $G = \text{SO}(2)$  acts by rotating this plane, in such a way that the orbits are circles centered on the origin.

We insisted in Section 2.2 on the crucial role played by the choice of the distance  $d(S, \hat{S})$ , and in particular on the dependency of the optimal estimator on  $d$ . This point is particularly relevant here: if for instance one uses the SE distance in the problem sketched in Figure 1 one obtains as an optimal estimator  $\hat{S}^{\text{BO}}(Y) = \mathbb{E}[S|Y] = 0$ , since the posterior is uniform on the circles and then averages out to the origin of the plane. This apparent contradiction that the “optimal” estimator is a trivial constant even if  $Y$  contains a non-trivial information arises from the incorrect choice of the distance  $d$  used. In this very simple example the way out is easy to find: one can use polar coordinates  $(r, \theta)$  in the plane instead of the cartesian ones  $(S_1, S_2)$ , notice that the orbits are in bijection with the radius  $r$ , hence take for  $d(S, \hat{S})$  any distance between  $r$  and  $\hat{r}$ , independently of  $\theta$  and  $\hat{\theta}$ . However for high-dimensional signal spaces, and complicated group actions on them, it is not obvious to find an explicit parametrization of the space that allows to separate the coordinates of the orbit from the coordinates inside an orbit. In this case a generic prescription is to start from a distance  $d$  (possibly too fine, i.e. that separate points of the same orbit), and to quotient it by the action of the group, defining a new distance  $d_G$  with

$$d_G(S, \hat{S}) = \inf_{g \in G} d(g \cdot S, \hat{S}). \tag{10}$$

We assume that  $d$  and the action of  $G$  are such that  $d(g \cdot S, \widehat{S}) = d(S, g^{-1} \cdot \widehat{S})$ , which yields the symmetry  $d_G(S, \widehat{S}) = d_G(\widehat{S}, S)$ , and the independence of  $d_G$  on the orbit representant, namely  $d_G(g \cdot S, \widehat{S}) = d_G(S, \widehat{S})$ . In the example of Figure 1 the quotient of the SE distance  $d(S, \widehat{S}) = (S_1 - \widehat{S}_1)^2 + (S_2 - \widehat{S}_2)^2$  by the rotation group according to Equation (10) yields  $d_G(S, \widehat{S}) = (r - \widehat{r})^2$ , a valid distance between orbits in agreement with the intuitive reasoning above.

We emphasized in Section 2.2 that very few choices of the distance  $d$  yield a closed-form expression for the corresponding Bayes-optimal estimator  $\widehat{S}^{\text{BO}}$ . As a consequence the presence of a symmetry, and hence the need to trade the distance  $d$  by its coarse-grained version  $d_G$ , will in general complicate the resolution of such inference problems (even when the initial distance measure was the SE or Hamming ones). We will present later on a series of explicit examples, but before that we point out some previous works where this phenomenon was already discussed.

A well-known case where this kind of symmetry occurs is the SBM introduced in Section 2.1, in particular in its symmetric version where the prior distribution of the labels is uniform,  $\bar{\eta} = (1/q, \dots, 1/q)$ , and where the elements of the affinity matrix  $c_{\sigma, \tau}$  only depend on whether  $\sigma = \tau$  or not (i.e.  $c$  is constant on the diagonal, and off-diagonal). In that case it is easy to check that  $G = \mathcal{S}_q$ , the symmetric group of permutations of  $[q]$ , acting on the labels  $S = (\sigma_1, \dots, \sigma_N)$  by  $\pi \cdot S = (\pi(\sigma_1), \dots, \pi(\sigma_N))$ , and leaving the observed graph untouched, is a symmetry of the problem. This was of course noticed in [23], which used as a measure of the distance between the signal and its estimation the quotiented version of the Hamming distance, according to the definition (10) with  $G = \mathcal{S}_q$ . A non-trivial estimation, in the large size limit, is obtained if this (normalized) quotiented distance is strictly smaller than what can be achieved by a dummy estimator that does not access  $Y$ , and only uses the prior information on  $S$ . There has been some discussions, see in particular [25], on the appropriate generalization of this definition when the parameters  $\bar{\eta}$  and  $c$  are only partially symmetric. In [25] it was emphasized that the definition (10) with  $G = \mathcal{S}_q$  was not satisfactory in this case, and proposed an alternative definition based on the possibility to partition the vertices in two sets which have a different empirical content of at least two communities. We believe that (10) could also be used in the partially symmetric case, provided one takes in this definition  $G$  as the largest subgroup of  $\mathcal{S}_q$  with  $(g \cdot S, Y) \stackrel{\text{d}}{=} (S, Y)$  for all  $g \in G$  (which will be a proper subgroup of  $\mathcal{S}_q$  if there is only a partial symmetry in the parameters of the SBM). In the statistical mechanics treatment of [23] the symmetry difficulty was bypassed thanks to the spontaneous breaking of the posterior measure into various “pure states” related one to the other by the action of  $G$ , i.e. by the permutations of the labels; in technical terms the cavity (or Belief Propagation) equations admit a trivial symmetric solution, that would yield the Bayes optimal estimator if the naive Hamming distance were to be used, and several symmetry breaking solutions (corresponding to each pure states). Picking one of the latter allows to solve the Bayes-optimal estimation with the quotiented distance (10). Notice that here the size of the symmetry group is at most  $q!$ , and thus remains finite in the large graph limit (we will see later on an example where this will not be the case).

We would also like to mention here the research field called orbit recovery, or estimation under a group action, see for instance [61, 62] and reference therein, that is tightly related to the present discussion (see also [63] for computational aspects of related problems). In this setting one generates an observation  $Y = (Y_1, \dots, Y_n)$  through

$$Y_i = g_i \cdot S + \xi_i, \quad (11)$$

with  $g_i$  i.i.d. copies of an uniformly chosen member of a group  $G$  (i.e. the  $g_i$  are drawn from the Haar measure on  $G$ ), and  $\xi_i$  some i.i.d. noise variables, with intensity controlled by a parameter  $\Delta$ . The motivation for these studies comes from applications to cryo-electron microscopy (cryoEM) [64, 65], an experimental method where successive projections of randomly rotated molecules are observed, or its abstraction called multi-reference alignment (MRA) [66].

As discussed above the signal  $S$  is here indistinguishable from any member of the orbit  $G \cdot S$ , since the  $g_i$ 's are uniform on  $G$ . There is however an important difference in the perspective taken by these works, which concentrate on the dependence of the sample complexity on the noise intensity  $\Delta$ , i.e. on how large  $n$  should be when  $\Delta$  diverges in order to recover the orbit of  $S$  within a fixed accuracy. On the contrary here we consider a single sample, and focus on the Bayes-optimal estimation one can extract from it.

The procedure to achieve this task is in principle clear: inserting the definition (10) of the quotiented distance in the expression (3) of the Bayes-optimal estimator yields

$$\widehat{S}^{\text{BO}}(Y) \in \underset{x}{\operatorname{arg\,inf}} \mathbb{E} \left[ \inf_{g \in G} d(g \cdot S, x) \mid Y \right]; \quad (12)$$

note that by definition this optimal estimator is not unique, if  $\widehat{S}(Y)$  minimizes the average quotiented distance (for a given  $Y$ ) this will also be the case of all members of the orbit  $G \cdot \widehat{S}(Y)$ . The actual computation of (12) seems daunting in general, as it relies on two extremizations (over  $g$  and  $x$ ) intertwined with an average (over  $S|Y$ ). We shall now consider a few concrete examples where one can give more explicit expressions of  $d_G$  and of the corresponding  $\widehat{S}^{\text{BO}}$ .

- Our first example will be elementary; we consider a scalar signal  $S \in \mathbb{R}$ , and the SE distance  $d(S, \widehat{S}) = (S - \widehat{S})^2$  as the naive measure for the accuracy of estimation. The multiplicative group  $G = \{-1, 1\}$ , acting by multiplication on  $S$  (i.e.  $g \cdot S = gS$  for  $g = \pm 1$ ), induces the quotiented distance

$$d_G(S, \widehat{S}) = \min((S - \widehat{S})^2, (S + \widehat{S})^2) = S^2 + \widehat{S}^2 - 2 \max(S\widehat{S}, -S\widehat{S}) = S^2 + \widehat{S}^2 - 2|S||\widehat{S}| = (|S| - |\widehat{S}|)^2, \quad (13)$$

which one recognizes as a valid distance between the orbits  $G \cdot S = \{|S|, -|S|\}$  and  $G \cdot \widehat{S}$ . To obtain the Bayes-optimal estimators for this quotiented distance we compute

$$\mathbb{E}[(|x| - |S|)^2 | Y] = x^2 - 2|x|\mathbb{E}[|S| | Y] + \mathbb{E}[S^2 | Y], \quad (14)$$

which is minimized in  $x = \mathbb{E}[|S| | Y]$  and  $x = -\mathbb{E}[|S| | Y]$ . This yields the two optimal estimators,

$$\widehat{S}^{\text{BO}}(Y) = \mathbb{E}[|S| | Y] \quad \text{and} \quad \widehat{S}^{\text{BO}}(Y) = -\mathbb{E}[|S| | Y], \quad (15)$$

as could be expected intuitively. As an example of a pair  $(S, Y)$  admitting this symmetry one can take for  $S$  a symmetric real random variable (i.e. with  $S \stackrel{d}{=} -S$ ), and observe it through a multiplicative noise with symmetric distribution, namely defining the observation as  $Y = ZS$  with  $Z$  independent of  $S$ , and such that  $Z \stackrel{d}{=} -Z$ . As in the example of Figure 1 the Bayes-optimal estimator for the naive SE distance  $d$  is trivial,  $\mathbb{E}[S|Y] = 0$ , even when  $Y$  contains a lot of information on the orbit of  $S$  (and potentially allows to determine it exactly, as for example when  $Z = \pm 1$  with equal probabilities 1/2).

- The example above was very simple, and led to an optimal estimator which could be expressed in terms of a posterior mean, not of the signal itself, but of a simple function of  $S$  (here the absolute value). It is however instructive as it shows that such a situation should be considered as a fortunate coincidence. Consider indeed the following slight variation of the previous example: now  $S \in \mathbb{R}^N$  with  $N \geq 2$ , with  $d(S, \widehat{S}) = \|S - \widehat{S}\|_2^2$  and again  $G = \{-1, 1\}$  acting multiplicatively on  $S$ . The quotiented distance thus reads

$$d_G(S, \widehat{S}) = \min(\|S - \widehat{S}\|_2^2, \|S + \widehat{S}\|_2^2) = \|S\|_2^2 + \|\widehat{S}\|_2^2 - 2|\langle S, \widehat{S} \rangle|. \quad (16)$$

At variance with the  $N = 1$  case the absolute value of the scalar product  $S_1 \widehat{S}_1 + \dots + S_N \widehat{S}_N$  cannot be factored as the product of a function of  $S$  and one of  $\widehat{S}$ . As a consequence in the expression of the optimal estimator

$$\widehat{S}^{\text{BO}}(Y) \in \underset{x \in \mathbb{R}^N}{\text{arginf}} \left( \|x\|_2^2 - 2 \mathbb{E} \left[ |\langle S, x \rangle| \mid Y \right] \right), \quad (17)$$

the minimization on  $x$  cannot be performed explicitly (for a generic posterior law of  $S|Y$ ), and  $\widehat{S}^{\text{BO}}(Y)$  is not expressible as the posterior mean of some simple function of  $S$ .

- We will now present another fortunate case in which analytical simplifications are possible even though the action of the group is much less trivial than above. We consider vectorial signals,  $S \in \mathbb{R}^N$ , with the usual SE distance  $d(S, \widehat{S}) = \|S - \widehat{S}\|_2^2$ , and assume now that the symmetry group of the inference problem is  $G = \mathcal{S}_N$ , the group of permutations of  $[N]$ , acting on vectors by permutation of the coordinates: for  $\pi \in \mathcal{S}_N$ ,  $(\pi \cdot S)_i = S_{\pi^{-1}(i)}$ . This is indeed an action, isometric with respect to  $d$  in the sense that  $d(\pi \cdot S, \widehat{S}) = d(S, \pi^{-1} \cdot \widehat{S})$ . We consider the quotiented distance  $d_G$  from Equation (10), that reads

$$d_G(S, \widehat{S}) = \min_{\pi \in \mathcal{S}_N} \sum_{i=1}^N [(S_{\pi^{-1}(i)})^2 + (\widehat{S}_i)^2 - 2S_{\pi^{-1}(i)} \widehat{S}_i] \quad (18)$$

$$= \|S\|_2^2 + \|\widehat{S}\|_2^2 - 2 \max_{\pi \in \mathcal{S}_N} \sum_{i=1}^N S_i \widehat{S}_{\pi(i)}. \quad (19)$$

It turns out that the last maximization can be performed rather explicitly, thanks to the so-called rearrangement inequality, see for instance theorem 368 in [67]. To state this result let us denote  $S_{[1]} \leq S_{[2]} \leq \dots \leq S_{[N]}$  a permutation of  $S_1, \dots, S_N$  that puts the coordinates in non-decreasing order, and  $S^\uparrow$  the corresponding vector, with  $(S^\uparrow)_i = S_{[i]}$ . The same notation will be used for  $\widehat{S}$  as well. The rearrangement inequality states that the maximum in the last term of (19) is reached when  $\pi$  orders the elements of  $\widehat{S}$  in the same way as the one of  $S$ . More explicitly, assuming without loss of generality that  $S_1 \leq \dots \leq S_N$ , the maximum is reached for  $\pi$  such that  $\widehat{S}_{\pi(i)} = \widehat{S}_{[i]}$ . We can thus write

$$d_G(S, \widehat{S}) = \sum_{i=1}^N [(S_i)^2 + (\widehat{S}_i)^2 - 2S_{[i]} \widehat{S}_{[i]}] = d(S^\uparrow, \widehat{S}^\uparrow), \quad (20)$$

in other words the quotiented distance coincides with the usual SE distance on the ordered version of the vectors.

As a consequence the Bayes-optimal estimator for this  $d_G$  can be computed explicitly,  $\widehat{S}^{\text{BO}}(Y) = \mathbb{E}[S^\uparrow | Y]$ , or in components  $(\widehat{S}^{\text{BO}}(Y))_i = \mathbb{E}[S_{[i]} | Y]$  (of course any permutation of this estimator achieves the same risk and is also optimal). This provides another example, less trivial than the one above, of a quotiented distance whose optimal estimator can be expressed as a posterior mean of some function of the signal (here the order statistics of  $S$ ). The practical implementation of such a formula might still be complicated, in particular in the large  $N$  limit, but at least one has an explicit expression to start with.

To complete the discussion of this example let us now define a pair of random variable  $(S, Y)$  that fulfils this symmetry assumption. We shall take  $S_1, \dots, S_N$  as independent identically distributed copies of an absolutely continuous random variable (hence with probability 1 the  $S_i$  are distinct, which simplifies the discussion below). The law of  $S$  is clearly invariant under  $\mathcal{S}_N$ . Consider in addition the observation channel  $Y = (Y_1, \dots, Y_N)$ , where for  $p \in [N]$ :

$$Y_p = \sum_{i=1}^N (S_i)^p. \quad (21)$$

The  $Y_p$ 's are deterministic functions of  $S$ , invariant under the permutations of coordinates of  $S$ , hence one has indeed  $(\pi \cdot S, Y) \stackrel{d}{=} (S, Y)$ . The theory of symmetric functions [54] (in which the  $Y_p$ 's are known as the power sum symmetric polynomials) demonstrates that the knowledge of  $Y$  allows to determine the set  $\{S_1, \dots, S_N\}$  (but of course not the ordered vector). The posterior law  $P(\cdot|Y)$  is thus the linear combination of  $N!$  atoms, with equal weights, on the permutations of the vector  $S$  corresponding to this realization of  $Y$ . The risk of the Bayes-optimal estimator for the quotiented distance thus vanishes, since the orbit of  $S$  is perfectly reconstructed from  $Y$ . On the contrary if one had used the naive SE distance the optimal estimator would have been  $(\widehat{S}^{\text{BO}}(Y))_1 = \dots = (\widehat{S}^{\text{BO}}(Y))_N = (1/N)Y_1$ , giving a risk proportional to the variance of the law of the  $S_i$ 's.

- The last example of a detrimental symmetry we shall consider now is motivated by the extensive rank matrix factorization problem [27–30, 43–51] we mentioned in Section 2.1. We recall that in this setting the signal, denoted  $X$  to avoid confusion with the matrix denoising problem, is an  $n \times r$  matrix observed through the channel  $Y = XX^T + Z$ , the observation  $Y$  being an  $n \times n$  matrix, and  $Z$  modelling an additive noise independent of  $X$ . Denoting  $X_{i,\mu}$  the matrix elements of  $X$ , one has more explicitly

$$Y_{i,j} = \sum_{\mu=1}^r X_{i,\mu}X_{j,\mu} + Z_{i,j}, \quad (22)$$

for  $i, j \in [n]$ . Suppose now that in the prior law on  $X$  the elements  $X_{i,\mu}$  are i.i.d. random variables (one could actually make the weaker assumption that the vectors  $X_\mu = \{X_{i,\mu}\}_{i \in [n]}$  are i.i.d. vectors in  $\mathbb{R}^n$ ); then the permutation group  $G = \mathcal{S}_r$ , acting on  $X$  as  $(\pi \cdot X)_{i,\mu} = X_{i,\pi^{-1}(\mu)}$  for  $\pi \in G$ , is a symmetry of the inference problem in the sense that  $(\pi \cdot X, Y) \stackrel{d}{=} (X, Y)$ . Since its action is trivial on the observation this is an example of a detrimental symmetry in which the distance measuring the accuracy of estimation has to be amended in order to gauge the non-trivial unobservable orbits of the signal (this symmetry acting on the “internal” indices  $\mu$  of  $X$  should not be confused with the one of the “external” indices  $i$ , that was shown to be beneficial in the matrix denoising problem in Section 3).

The natural naive distance is here the squared Frobenius norm of the difference between  $X$  and its estimator  $\widehat{X}$ , or in other words the sum of the SE for each matrix element,

$$d(X, \widehat{X}) = \sum_{i=1}^n \sum_{\mu=1}^r (X_{i,\mu} - \widehat{X}_{i,\mu})^2. \quad (23)$$

The optimal estimator with respect to  $d$  is as usual the posterior mean  $\widehat{X}(Y) = \mathbb{E}[X|Y]$ ; its components  $\widehat{X}_{i,\mu}(Y)$  are independent on  $\mu$  because of the invariance under  $G$ , which would yield a rather unsatisfactory estimate  $\widehat{X}(Y)$  of rank 1 (that would actually vanish identically if in addition the law of  $X$  is assumed to be even, i.e. if  $X \stackrel{d}{=} -X$ , since  $Y$  is invariant under a change of sign in  $X$ ). As explained in generic terms above one should use instead the quotiented version of  $d$ , according to the definition (10) for the group  $G = \mathcal{S}_r$ , that reads (see also Section V of [32] for a discussion of this distance)

$$d_G(X, \widehat{X}) = \sum_{i=1}^n \sum_{\mu=1}^r [(X_{i,\mu})^2 + (\widehat{X}_{i,\mu})^2] - 2 \max_{\pi \in \mathcal{S}_r} \sum_{i=1}^n \sum_{\mu=1}^r X_{i,\mu} \widehat{X}_{i,\pi(\mu)}. \quad (24)$$

The last term is of the form

$$\max_{\pi \in \mathcal{S}_r} \sum_{\mu=1}^r T_{\mu,\pi(\mu)}, \quad \text{with } T_{\mu,\nu} = \sum_{i=1}^n X_{i,\mu} \widehat{X}_{i,\nu} = \langle X_\mu, \widehat{X}_\nu \rangle, \quad (25)$$

where in the last equality we used  $\langle \cdot, \cdot \rangle$  to denote the standard scalar product in  $\mathbb{R}^n$ . One can recognize here a famous optimization problem, known as the bipartite matching, or

linear assignment problem. This can be solved in a time polynomial in  $r$  [68] (in [32] a greedy procedure was used to approximate this optimization), but in general does not admit an explicit analytical solution (one cannot use the rearrangement inequality anymore, apart from the trivial case  $n = 1$ ). As a consequence the Bayes-optimal estimator for the quotiented distance takes its generic form (12), which reads here

$$\widehat{X}^{\text{BO}}(Y) \in \underset{x \in \mathbb{R}^{n \times r}}{\operatorname{arginf}} \left\{ \sum_{i=1}^n \sum_{\mu=1}^r (x_{i,\mu})^2 - 2 \mathbb{E} \left[ \max_{\pi \in \mathcal{S}_r} \sum_{i=1}^n \sum_{\mu=1}^r X_{i,\mu} x_{i,\pi(\mu)} \mid Y \right] \right\}, \quad (26)$$

with a priori no simplification in terms of the conditional expectation of some explicit function of  $X$ . As we shall discuss further in the conclusions this seems a particularly challenging problem in the extensive rank case with  $r$  proportional to  $n$  in the large size limit, that is not expressible in terms of local marginal laws of the posterior probability  $P_{X|Y}$ .

## 5. Conclusions

We have collected in this paper a series of simple observations on the effects of invariances in Bayesian inference problems. We have argued that their consequences depend on the relative strength of the actions of the symmetry group on the hidden signal and on the observed quantities. When the latter are sufficiently affected by the symmetry one can exploit the constraints induced on the candidate estimators to reduce the complexity of the low-degree polynomial approach; on the contrary an invariance of the signal without observable effects imposes a revision of the notion of distance between an estimator and the groundtruth, that should be invariant on the orbits of the group.

We believe that this last point may have been overlooked and would deserve deeper investigations, in particular in the context of the statistical mechanics approach to inference problems. As briefly sketched in Section 2.1 the latter envisions the posterior probability (1) as a graphical model with dynamical variables in the signal space, the observations corresponding to a planted quenched disorder. The typical properties of such a model, and in particular the quenched free-energy or mutual information  $I(S; Y)$ , are then obtained via the replica or cavity method, while single-sample quantities (i.e. for one given realization of the disorder  $Y$ ) like the marginal probabilities of a finite number of variables are accessible by their algorithmic counterparts, namely the Belief Propagation or Approximate Message Passing techniques (we refer the reader to [15–19] for reviews of these methods). The latter allows to compute Bayes-optimal estimators for the Hamming or SE distance, that are indeed expressible in terms of the local marginals of the posterior, the former giving access to the corresponding average risk (in particular thanks to the I-MMSE relation [69] for Gaussian noise channels). Unfortunately these connections break down when other distances between the signal and its estimation are used, and in particular when a non-trivial group action on the signal space imposes to quotient the distance according to Equation (10): the mutual information  $I(S; Y)$  can no longer be related to the average risk, and the optimal estimator is not anymore expressible in terms of local posterior averages.

When the order of the group remains finite in the thermodynamic limit, as is the case in the symmetric SBM, one can rely on the spontaneous breaking of the posterior measure in pure states that shows up in the statistical mechanics treatment as symmetry breaking solutions of the replica or cavity equations and allows to compute (10) from (1). It is much less clear that such a phenomenon can occur when the symmetry group grows in the thermodynamic limit, in particular in the case of extensive rank matrix factorization, an observation that could constitute the starting point for future research. We believe that in such a case the computation of Bayes-optimal estimators should instead start from the formula in (12) (or (26) in the special case of

matrix estimation problems invariant under column permutation), and not from the graphical model interpretation of the posterior (1). This is a much more challenging “three-level problem”: for a given value of  $x$  and  $S$  in the signal space one has to solve a minimization problem over  $g \in G$ , for instance by considering  $d(g \cdot S, x)$  as an energy function and introducing a fictitious temperature ultimately sent to 0. The solution of this “inner problem” would then provide a function of  $(S, x)$ , which should be averaged over  $S$  according to the posterior law  $S|Y$ , and finally minimized over  $x$ . Multi-level problems of a similar kind were studied with statistical mechanics methods in [70, 71]: the inner problem is solved by a message passing algorithm (for the assignment problem of (26) this could be done with the BP approach of [72]), these messages constituting the degrees of freedom of the next-level problem. In principle this approach could be followed for the quotiented distance of (26) in the extensive rank matrix factorization problem, with the hope that incorporating explicitly the unobservable symmetry in such a way may alleviate the difficulties encountered by the formalism of [43, 44]. We leave this specific question, as well as the more generic problem of the possibility of spontaneous breaking of an extensive symmetry group of a posterior measure into pure states, as open issues for future works.

Let us finally underline that even if we concentrated in this work on a static, Bayes-optimal approach to inference problems, symmetries also play a role in the dynamical studies of learning algorithms, see for instance [73–78] and references therein. In this context they can again be beneficial, allowing in particular to project a dynamical evolution in a high-dimensional space down to a finite number of order parameters, or detrimental whenever the non-trivial recovery of the groundtruth signal requires the breaking of symmetries present in a random, uninformative initial condition.

## Declaration of interests

The authors do not work for, advise, own shares in, or receive funds from any organization that could benefit from this article, and have declared no affiliations other than their research organizations.

## References

- [1] W. Fulton and J. Harris, *Representation Theory, a First Course*, Springer: New York, 1991.
- [2] R. Goodman and N. R. Wallach, *Symmetry, Representations, and Invariants*, Springer: New York, 2009.
- [3] A. Zee, *Group Theory in a Nutshell for Physicists*, Princeton University Press: Princeton, 2016.
- [4] G. Toulouse, “Theory of frustration effect in spin-glasses: I”, *Commun. Phys.* **2** (1977), pp. 115–119.
- [5] G. Toulouse, “Symmetry and topology concepts for spin glasses and other glasses”, *Phys. Rep.* **49** (1979), pp. 267–272.
- [6] M. L. Eaton, *Group invariance applications in statistics*, Regional Conference Series in Probability and Statistics, Institute of Mathematical Statistics: Hayward, 1989, pp. i–133.
- [7] R. A. Wijsman, *Invariant measures on groups and their use in statistics*, Lecture Notes-Monograph Series, Institute of Mathematical Statistics: Hayward, 1990, pp. i–218.
- [8] E. L. Lehmann and G. Casella, *Theory of Point Estimation*, 2nd edition, Springer: New York, 1998.
- [9] J. Shawe-Taylor, “Building symmetries into feedforward networks”, in *First IEE International Conference on Artificial Neural Networks*, (Conf. Publ. No. 313), IET: London, 1989, pp. 158–162.
- [10] M. M. Bronstein, J. Bruna, T. Cohen and P. Velickovic, *Geometric deep learning: grids, groups, graphs, geodesics, and gauges*, preprint, 2021, 2104.13478.
- [11] S. Villar, D. W. Hogg, W. Yao, G. A. Kevrekidis and B. Schölkopf, “Towards fully covariant machine learning”, *Trans. Mach. Learn. Res.* (2024). (submitted).
- [12] T. Cohen and M. Welling, “Group equivariant convolutional networks”, in *Proceedings of The 33rd International Conference on Machine Learning* (M. F. Balcan and K. Q. Weinberger, eds.), Proceedings of Machine Learning Research, PMLR, PMLR: New York, 2016, pp. 2990–2999.
- [13] S. Chen, E. Dobriban and J. H. Lee, “A group-theoretic framework for data augmentation”, *J. Mach. Learn. Res.* **21** (2020), no. 245, pp. 1–71.

- [14] B. Simsek, F. Ged, A. Jacot, F. Spadaro, C. Hongler, W. Gerstner and J. Brea, “Geometry of the loss landscape in overparameterized neural networks: symmetries and invariances”, in *Proceedings of the 38th International Conference on Machine Learning* (M. Meila and T. Zhang, eds.), Proceedings of Machine Learning Research, PMLR, PMLR: New York, 2021, pp. 9722–9732.
- [15] H. Nishimori, *Statistical Physics of Spin Glasses and Information Processing: An Introduction*, Oxford University Press: Oxford, 2001.
- [16] A. Engel and C. Van den Broeck, *Statistical Mechanics of Learning*, Cambridge University Press: Cambridge, 2001.
- [17] M. Mézard and A. Montanari, *Physics, Information, Computation*, Oxford University Press: Oxford, 2009.
- [18] L. Zdeborová and F. Krzakala, “Statistical physics of inference: thresholds and algorithms”, *Adv. Phys.* **65** (2016), no. 5, pp. 453–552.
- [19] *Spin Glass Theory and Far Beyond*, (P. Charbonneau, E. Marinari, M. Mézard, G. Parisi, F. Ricci-Tersenghi, G. Sicuro and F. Zamponi, eds.), World Scientific: Singapore, 2023.
- [20] S. E. Fienberg and S. S. Wasserman, “Categorical data analysis of single sociometric relations”, *Sociol. Methodol.* **12** (1981), pp. 156–192.
- [21] P. W. Holland, K. B. Laskey and S. Leinhardt, “Stochastic blockmodels: first steps”, *Social Networks* **5** (1983), no. 2, pp. 109–137.
- [22] B. Bollobás, S. Janson and O. Riordan, “The phase transition in inhomogeneous random graphs”, *Random Struct. Algorithms* **31** (2007), no. 1, pp. 3–122.
- [23] A. Decelle, F. Krzakala, C. Moore and L. Zdeborová, “Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications”, *Phys. Rev. E* **84** (2011), article no. 066106.
- [24] C. Moore, “The computer science and physics of community detection: landscapes, phase transitions, and hardness”, *Bull. EATCS* **1** (2017), no. 121.
- [25] E. Abbe, “Community detection and stochastic block models: recent developments”, *J. Mach. Learn. Res.* **18** (2018), no. 177, pp. 1–86.
- [26] J. Bun, R. Allez, J.-P. Bouchaud and M. Potters, “Rotational invariant estimator for general noisy matrices”, *IEEE Trans. Inf. Theory* **62** (2016), no. 12, pp. 7475–7490.
- [27] A. Maillard, F. Krzakala, M. Mézard and L. Zdeborová, “Perturbative construction of mean-field equations in extensive-rank matrix factorization and denoising”, *J. Stat. Mech.: Theory Exp.* **2022** (2022), no. 8, article no. 083301.
- [28] J. Barbier and N. Macris, “Statistical limits of dictionary learning: Random matrix theory and the spectral replica method”, *Phys. Rev. E* **106** (2022), article no. 024136.
- [29] E. Troiani, V. Erba, F. Krzakala, A. Maillard and L. Zdeborová, “Optimal denoising of rotationally invariant rectangular matrices”, *Proc. Math. Sci. Mach. Learn. (MSML)* **190** (2022), pp. 97–112.
- [30] F. Pourkamali, J. Barbier and N. Macris, “Matrix inference in growing rank regimes”, *IEEE Trans. Inform. Theory* **70** (2024), no. 11, pp. 8133–8163.
- [31] G. Semerjian, “Matrix denoising: bayes-optimal estimators via low-degree polynomials”, *J. Stat. Phys.* **191** (2024), article no. 139.
- [32] J. Barbier, F. Camilli, J. Ko and K. Okajima, *On the phase diagram of extensive-rank symmetric matrix denoising beyond rotational invariance*, preprint, 2024, 2411.01974.
- [33] B. Olshausen and D. Field, “Emergence of simple-cell receptive field properties by learning a sparse code for natural images”, *Nature* **381** (1996), pp. 607–609.
- [34] M. Zibulevsky and B. A. Pearlmutter, “Blind source separation by sparse decomposition in a signal dictionary”, *Neural Comput.* **13** (2001), no. 4, pp. 863–882.
- [35] K. Kreutz-Delgado, J. F. Murray, B. D. Rao, K. Engan, T.-W. Lee and T. J. Sejnowski, “Dictionary learning algorithms for sparse representation”, *Neural Comput.* **15** (2003), no. 2, pp. 349–396.
- [36] D. A. Spielman, H. Wang and J. Wright, “Exact recovery of sparsely-used dictionaries”, in *Proceedings of the 25th Annual Conference on Learning Theory*, PMLR: New York, 2012, pp. 37.1–37.18.
- [37] S. Rangan and A. K. Fletcher, “Iterative estimation of constrained rank-one matrices in noise”, in *2012 IEEE International Symposium on Information Theory Proceedings*, IEEE: Cambridge, 2012, pp. 1246–1250.
- [38] J. Barbier, M. Dia, N. Macris, F. Krzakala, T. Lesieur and L. Zdeborová, “Mutual information for symmetric rank-one matrix estimation: A proof of the replica formula”, in *Advances in Neural Information Processing Systems* (D. Lee et al., ed.), Curran Associates, Inc.: Red Hook, 2016, pp. 424–432. Online at [https://proceedings.neurips.cc/paper\\_files/paper/2016/file/621bf66ddb7c962aa0d22ac97d69b793-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2016/file/621bf66ddb7c962aa0d22ac97d69b793-Paper.pdf). (accessed February 10, 2025).
- [39] T. Lesieur, F. Krzakala and L. Zdeborová, “Constrained low-rank matrix estimation: phase transitions, approximate message passing and applications”, *J. Stat. Mech.: Theory Exp.* **2017** (2017), no. 7, article no. 073403.
- [40] M. Lelarge and L. Miolane, “Fundamental limits of symmetric low-rank matrix estimation”, *Probab. Theory Relat. Fields* **173** (2019), pp. 859–929.
- [41] J. Barbier and N. Macris, “The adaptive interpolation method: a simple scheme to prove replica formulas in Bayesian inference”, *Probab. Theory Relat. Fields* **174** (2019), pp. 1133–1185.

- [42] A. Montanari and A. S. Wein, *Equivalence of approximate message passing and low-degree polynomials in rank-one matrix estimation*, preprint, 2022, 2212.06996.
- [43] A. Sakata and Y. Kabashima, “Statistical mechanics of dictionary learning”, *Europhys. Lett.* **103** (2013), no. 2, article no. 28008.
- [44] Y. Kabashima, F. Krzakala, M. Mézard, A. Sakata and L. Zdeborová, “Phase transitions and sample complexity in bayes-optimal matrix factorization”, *IEEE Trans. Inf. Theory* **62** (2016), no. 7, pp. 4228–4265.
- [45] H. C. Schmidt, *Statistical physics of sparse and dense models in optimization and inference*, PhD Thesis, Université Paris-Saclay, 2018.
- [46] J. Barbier, J. Ko and A. A. Rahman, *A multiscale cavity method for sublinear-rank symmetric matrix factorization*, preprint, 2024, 2403.07189.
- [47] F. Pourkamali and N. Macris, “Bayesian extensive-rank matrix factorization with rotational invariant priors”, in *Advances in Neural Information Processing Systems* (A. Oh et al., ed.), Curran Associates Inc.: Red Hook, 2023, pp. 24025–24073. Online at [https://proceedings.neurips.cc/paper\\_files/paper/2023/file/4b8afc47273c746662a96dfdf562f87f-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/4b8afc47273c746662a96dfdf562f87f-Paper-Conference.pdf). (accessed February 10, 2025).
- [48] F. Pourkamali and N. Macris, “Rectangular rotational invariant estimator for general additive noise matrices”, in *2023 IEEE International Symposium on Information Theory (ISIT), Taipei, Taiwan*, IEEE: Cambridge, 2023, pp. 2081–2086.
- [49] I. D. Landau, G. C. Mel and S. Ganguli, “Singular vectors of sums of rectangular random matrices and optimal estimation of high-rank signals: The extensive spike model”, *Phys. Rev. E* **108** (2023), article no. 054129.
- [50] F. Camilli and M. Mézard, “Matrix factorization with neural networks”, *Phys. Rev. E* **107** (2023), article no. 064308.
- [51] F. Camilli and M. Mézard, “The decimation scheme for symmetric matrix factorization”, *J. Phys. A: Math. Theor.* **57** (2024), no. 8, article no. 085002.
- [52] T. Hastie, R. Tibshirani and J. Friedman, *The Elements of Statistical Learning*, 2nd edition, Springer: New York, 2009.
- [53] L. P. Barnes, A. Dytso, J. Liu and H. V. Poor, “Multivariate priors and the linearity of optimal Bayesian estimators under Gaussian noise”, in *2024 IEEE International Symposium on Information Theory (ISIT), Athens, Greece*, IEEE: Cambridge, 2024, pp. 987–992.
- [54] I. G. Macdonald, *Symmetric Functions and Hall Polynomials*, Oxford University Press: Oxford, 1998.
- [55] H. Weyl, *The Classical Groups: Their Invariants and Representations*, Princeton University Press: Princeton, 1966.
- [56] T. Schramm and A. S. Wein, “Computational barriers to estimation from low-degree polynomials”, *Ann. Stat.* **50** (2022), no. 3, pp. 1833–1858.
- [57] S. B. Hopkins and D. Steurer, “Efficient bayesian estimation from few samples: community detection and related problems”, in *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, IEEE: Cambridge, 2017, pp. 379–390.
- [58] D. Kunisky, A. S. Wein and A. S. Bandeira, “Notes on computational hardness of hypothesis testing: predictions using the low-degree likelihood ratio”, in *Mathematical Analysis, its Applications and Computation*, Springer International Publishing: Cham, 2022, pp. 1–50.
- [59] D. Kunisky, C. Moore and A. S. Wein, “Tensor cumulants for statistical inference on invariant distributions”, in *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS), Chicago, IL, USA*, IEEE: Cambridge, 2024, pp. 1007–1026.
- [60] S. Mei, T. Misiakiewicz and A. Montanari, “Learning with invariances in random features and kernel models”, in *Proceedings of Thirty Fourth Conference on Learning Theory* (M. Belkin and S. Kpotufe, eds.), Proceedings of Machine Learning Research, PMLR, PMLR: New York, 2021, pp. 3351–3418.
- [61] E. Abbe, J. M. Pereira and A. Singer, “Estimation in the group action channel”, in *2018 IEEE International Symposium on Information Theory (ISIT)*, IEEE: Cambridge, 2018, pp. 561–565.
- [62] A. S. Bandeira, B. Blum-Smith, J. Kileel, J. Niles-Weed, A. Perry and A. S. Wein, “Estimation under group actions: recovering orbits from invariants”, *Appl. Comput. Harmon. Anal.* **66** (2023), pp. 236–319.
- [63] P. Bürgisser, M. L. Doğan, V. Makam, M. Walter and A. Wigderson, “Complexity of robust orbit problems for torus actions and the abc-conjecture”, in *39th Computational Complexity Conference (CCC 2024)* (R. Santhanam, ed.), Leibniz International Proceedings in Informatics (LIPIcs), Schloss Dagstuhl – Leibniz-Zentrum für Informatik: Dagstuhl, Germany, 2024, 14:1–14:48.
- [64] M. Adrian, J. Dubochet, J. Lepault and A. W. McDowell, “Cryo-electron microscopy of viruses”, *Nature* **308** (1984), pp. 32–36.
- [65] A. Singer and Y. Shkolnisky, “Three-dimensional structure determination from common lines in Cryo-EM by eigenvectors and semidefinite programming”, *SIAM J. Imaging Sci.* **4** (2011), no. 2, pp. 543–572.
- [66] F. Sigworth, “A maximum-likelihood approach to single-particle image refinement”, *J. Struct. Biol.* **122** (1998), no. 3, pp. 328–339.
- [67] G. Hardy, J. Littlewood and G. Pólya, in *Inequalities*, Cambridge Mathematical Library, Cambridge University Press: Cambridge, 1934.

- [68] J. Edmonds and R. M. Karp, “Theoretical improvements in algorithmic efficiency for network flow problems”, *J. ACM* **19** (1972), no. 2, pp. 248–264.
- [69] D. Guo, S. Shamai and S. Verdu, “Mutual information and minimum mean-square error in Gaussian channels”, *IEEE Trans. Inf. Theory* **51** (2005), no. 4, pp. 1261–1282.
- [70] F. Altarelli, A. Braunstein, A. Ramezanzpour and R. Zecchina, “Stochastic optimization by message passing”, *J. Stat. Mech.: Theory Exp.* **2011** (2011), no. 11, article no. P11009.
- [71] M. Castellana and L. Zdeborová, “Adversarial satisfiability problem”, *J. Stat. Mech.: Theory Exp.* **2011** (2011), no. 03, article no. P03023.
- [72] M. Bayati, D. Shah and M. Sharma, “Max-product for maximum weight matching: convergence, correctness, and LP duality”, *IEEE Trans. Inform. Theory* **54** (2008), no. 3, pp. 1241–1251.
- [73] S. Mei, A. Montanari and P.-M. Nguyen, “A mean field view of the landscape of two-layer neural networks”, *Proc. Natl. Acad. Sci.* **115** (2018), no. 33, E7665–E7671.
- [74] S. Sarao Mannelli, G. Biroli, C. Cammarota, F. Krzakala, P. Urbani and L. Zdeborová, “Marvels and pitfalls of the langevin algorithm in noisy high-dimensional inference”, *Phys. Rev. X* **10** (2020), article no. 011057.
- [75] G. Ben Arous, R. Gheissari and A. Jagannath, “Online stochastic gradient descent on non-convex losses from high-dimensional inference”, *J. Mach. Learn. Res.* **22** (2021), no. 106, pp. 1–51.
- [76] G. Gluch and R. Urbanke, *Noether: the more things change, the more stay the same*, preprint, 2021, 2104.05508.
- [77] K. Hajjar and L. Chizat, “On the symmetries in the dynamics of wide two-layer neural networks”, *Electron. Res. Arch.* **31** (2023), no. 4, pp. 2175–2212.
- [78] Y. Dandi, E. Troiani, L. Arnaboldi, L. Pesce, L. Zdeborová and F. Krzakala, “The benefits of reusing batches for gradient descent in two-layer networks: breaking the curse of information and leap exponents”, in *Proceedings of the 41st International Conference on Machine Learning*, ICML’24, JMLR.org: Vienna, 2025.