



HAL
open science

Optimal errors and phase transitions in high-dimensional generalized linear models

Jean Barbier, Florent Krzakala, Nicolas Macris, Léo Miolane, Lenka
Zdeborová

► **To cite this version:**

Jean Barbier, Florent Krzakala, Nicolas Macris, Léo Miolane, Lenka Zdeborová. Optimal errors and phase transitions in high-dimensional generalized linear models. Proceedings of the National Academy of Sciences of the United States of America, 2019, 116 (12), pp.5451-5460. 10.1073/pnas.1802705116 . cea-01614258

HAL Id: cea-01614258

<https://cea.hal.science/cea-01614258v1>

Submitted on 21 Nov 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Optimal Errors and Phase Transitions in High-Dimensional Generalized Linear Models

Jean Barbier^{†◇*⊗}, Florent Krzakala^{*}, Nicolas Macris[†], Léo Miolane^{*⊗} and Lenka Zdeborová^U

Abstract

Generalized linear models (GLMs) arise in high-dimensional machine learning, statistics, communications and signal processing. In this paper we analyze GLMs when the data matrix is random, as relevant in problems such as compressed sensing, error-correcting codes or benchmark models in neural networks. We evaluate the mutual information (or “free entropy”) from which we deduce the Bayes-optimal estimation and generalization errors. Our analysis applies to the high-dimensional limit where both the number of samples and the dimension are large and their ratio is fixed. Non-rigorous predictions for the optimal errors existed for special cases of GLMs, e.g. for the perceptron, in the field of statistical physics based on the so-called replica method. Our present paper rigorously establishes those decades old conjectures and brings forward their algorithmic interpretation in terms of performance of the generalized approximate message-passing algorithm. Furthermore, we tightly characterize, for many learning problems, regions of parameters for which this algorithm achieves the optimal performance, and locate the associated sharp phase transitions separating learnable and non-learnable regions. We believe that this random version of GLMs can serve as a challenging benchmark for multi-purpose algorithms.

This paper is divided in two parts that can be read independently: The first part (main part) presents the model and main results, discusses some applications and sketches the main ideas of the proof. The second part (supplementary informations) is much more detailed and provides more examples as well as all the proofs.

Contents

I	Main part	3
1	Introduction	3
2	Main results	5
2.1	The free entropy	5
2.2	Overlap and optimal estimation error	6
2.3	Optimal generalization error	6
2.4	Optimality of approximate message-passing	7

[†] Laboratoire de Théorie des Communications, Faculté Informatique et Communications, Ecole Polytechnique Fédérale de Lausanne, Suisse.

[◇] International Center for Theoretical Physics, Trieste, Italy.

^{*} Laboratoire de Physique Statistique, CNRS & Université Pierre et Marie Curie & École Normale Supérieure & PSL Université, Paris, France.

^{*} Département d’Informatique de l’ENS, École Normale Supérieure & CNRS & PSL Research University & Inria, Paris, France.

^U Institut de Physique Théorique, CNRS & CEA & Université Paris-Saclay, Saclay, France.

[⊗] Corresponding authors: jean.barbier@lps.ens.fr, leo.miolane@gmail.com

3	Application to learning and inference	8
3.1	General observations about fixed points and terminology	9
3.2	Sensing compressively with non-linear outputs	10
3.3	Perceptron and alike	12
3.4	Empirical comparison with general purpose algorithms	13
4	Methods and proofs	13
4.1	Main theorems	14
4.2	Proof by the adaptive interpolation method	15
II	Supplementary informations	25
1	Setting	25
1.1	Generalized linear estimation: Problem statement	25
1.2	The teacher-student scenario	26
1.3	Two scalar inference channels	28
2	Main results	29
2.1	Replica-symmetric formula and mutual information	29
2.2	Optimal reconstruction (or estimation) error	31
2.3	Optimal generalization (or prediction) error	32
2.4	Optimality of the generalized approximate message-passing algorithm	33
2.5	Optimal denoising error	35
3	Application to concrete situations	35
3.1	Generic observations	36
3.2	Phase diagram of perfect learning	37
3.3	Examples of optimal generalization error	40
4	Proof of the replica formula by the adaptive interpolation method	43
4.1	Interpolating estimation problem	44
4.2	Free entropy variation along the interpolation path	46
4.3	Overlap concentration and fundamental sum rule	47
4.4	Lower and upper matching bounds	48
5	Proofs of the limits of optimal errors	51
5.1	Optimal generalization error: Proof of Theorem 4	51
5.2	Generalization error of GAMP: Proof of Proposition 2	58
5.3	Limit of the overlap: Proof of Theorem 2	59
5.4	Denoising error: Proof of Corollary 5	67
Appendix A	Some technicalities	68
A.1	The Nishimori identity	68
A.2	Unicity of the optimizer q^* of the replica formula: Proof of Proposition 1	68
A.3	Continuity properties of the mutual information	68
A.4	A simple consequence of hypotheses (h1)-(h2)-(h3)-(h4)	72
A.5	Derivative of the interpolating free entropy: Proof of Proposition 3	72
A.6	Boundedness of an overlap fluctuation	75
A.7	Proof of Proposition 6	76

Appendix B	Some properties of the scalar channels	77
B.1	The additive Gaussian scalar channel	77
B.2	The non-linear scalar channel	77
Appendix C	Approximation	84
C.1	Relaxing the hypotheses on P_0 and Φ	84
C.2	Relaxing the hypotheses on φ	85
C.3	The case of discrete channels: Removing the Gaussian noise	87
Appendix D	Some sup-inf formulas	87
Appendix E	Concentration of free entropy and overlaps	90
E.1	Concentration of the free entropy	90
E.2	Concentration of the overlap	95
Appendix F	Details on numerics	99
F.1	General purpose algorithms	99
F.2	Evaluating the replica formula	100
F.3	Breaking the symmetry in GAMP	100

Part I

Main part

1 Introduction

As datasets grow larger and more complex, modern data analysis requires solving high-dimensional estimation problems with very many parameters. Developing algorithms up to the task and understanding their limitations has become a major challenge in computer science, machine learning, statistics, signal processing, communications and related fields.

In the present contribution, we address this challenge in the case of generalized linear estimation models (GLMs) [1, 2] where data are generated as follows: Given a n -dimensional vector \mathbf{X}^* , hidden to the statistician, he/she observes instead a m -dimensional vector \mathbf{Y} where each component reads

$$Y_\mu = \varphi\left(\frac{1}{\sqrt{n}}[\Phi\mathbf{X}^*]_\mu, A_\mu\right), \quad 1 \leq \mu \leq m, \quad (1)$$

where Φ is a $m \times n$ “measurement” or “data” matrix, the random variables $(A_\mu) \stackrel{\text{iid}}{\sim} P_A$ account for noise (or randomness) of the model. The model is “linear” because the output Y_μ depends on a *linear* combination of the data $z_\mu = \frac{1}{\sqrt{n}}[\Phi\mathbf{X}^*]_\mu = \frac{1}{\sqrt{n}}\sum_{i=1}^n \Phi_{\mu i}X_i^*$. The GLM generalizes the ordinary linear regression by allowing the output function $\varphi(z, A)$ to be non-linear and/or stochastic; in the case of a deterministic model we simply write $\varphi(z)$. Explicit examples will be given below.

GLMs belong to the realm of supervised learning and arise in a wide variety of scientific fields. In signal processing one usually observes Y_μ given as a linear combination of the signal-elements \mathbf{X}^* . In a range of applications these observations are obtained via a non-linear function φ . In optics or X-ray crystallography one often measures only the amplitude of $[\Phi\mathbf{X}^*]_\mu$, leading to the phase retrieval problem [3]. A real-valued analog is the problem of sign-retrieval when we only observe $||[\Phi\mathbf{X}^*]_\mu|$ [4, 5]. Observations are sometimes quantized in order to reduce the storage, leading for instance to the problem of 1-bit compressed sensing [6]. In statistics and machine learning, classification is often described via a GLM where the output function φ

is discrete and corresponds to the labels that classify the data-points Φ_μ [1, 2, 7]. GLMs with non-linear output functions are also the basic building blocks of each layer of neural networks [8]: φ corresponds to the activation, the rows of the matrix Φ are different data samples while \mathbf{X}^* are the set of synaptic weights to be learned.

There are two main learning problems in GLMs: *i*) The *estimation* task requires, knowing the measured vector \mathbf{Y} and the matrix Φ , to infer the unknown vector \mathbf{X}^* ; *ii*) the *prediction* or *generalization* task instead requires, again knowing \mathbf{Y} and Φ , to predict accurately new values Y_{new} when new rows (i.e. data-points) are added to the matrix Φ .

In the present paper we build a rigorous theory for both these tasks for *random instances* of the GLM. In this setting each element $\Phi_{\mu i}$ of the matrix is sampled independently from a probability distribution of zero mean and unit variance, and the unknown vector \mathbf{X}^* has been also created randomly from a probability distribution P_0 , with each of its components $X_1^*, \dots, X_n^* \stackrel{\text{iid}}{\sim} P_0$. Since our main aim is to study the intrinsic information-theoretic and algorithmic limitations caused by the lack of samples and/or the amplitude of the noise, we assume throughout this paper that P_0 and φ are known to the statistician (if they are not the task can only be harder). Our results are derived in the challenging and interesting high-dimensional limit where $m, n \rightarrow \infty$ while $m/n \rightarrow \alpha$ a constant. Random instances of GLMs are both practically and theoretically relevant in many different contexts:

a) In *signal processing*, GLM estimation with a random matrix Φ has been studied with considerable attention in the context of compressed sensing [9–11] where a n -dimensional sparse signal is recovered from $m < n$ noisy measurements. While standard compressed sensing focused on the linear case –where $\varphi(z, A) = z + A$ with a Gaussian noise A – the generalized case was also widely studied [12, 13], especially for quantized output [14] and 1-bit compressed sensing [6, 15] where $\varphi(z, A) = \text{sign}(z + A)$, as well as for compressive phase retrieval when $\varphi(z, A) = |z + A|$ [16].

b) In *statistical learning*, important activity is dedicated to understand the limitation of learning with data generated by GLMs, both in the linear case, e.g. in the context of ridge regression or LASSO [17], or with non-linear probabilistic output, e.g. logistic regression. Random instances were studied in particular in the context of so-called M-estimators [18–21].

c) In studies of *artificial neural networks* there has been a large amount of works using random instances of GLMs, with φ playing the role of a non-linear activation function. In this context the random GLM was introduced as the teacher-student setting for the perceptron in the pioneering work of Gardner and Derrida [22]. Large volume of work followed and is reviewed, e.g., in [23–25]. While initial works concentrated on a simple activation functions $\varphi(z) = \text{sign}(z - K)$ (K is the threshold constant), many other functions were considered, e.g. in [26–28]. Recently, the study of random instances of neural networks have emerged as a key ingredient in understanding the performance of deep learning algorithms [29, 30]. Computing mutual informations in GLMs is also a critical issue in confirming the information bottleneck scenario of [31, 32]

d) In *communications*, error-correcting codes that use random constructions are particularly efficient, as discussed by Shannon in his seminal paper [33]. Random instances of GLMs describe both the setting of code-division multiple access –a multi-user access method used in communication technologies [34, 35]– as well as an error correction scheme called sparse superposition codes, that have been shown to achieve the Shannon capacity for *any type* of noisy channel [36–40].

Interestingly there is an important gap in the above volume of work. On the one hand there are studies that rely on the algorithmic performance of the so-called generalized approximate message-passing algorithm (GAMP) [11, 12, 41]. GAMP is remarkable in that its asymptotic ($n, m \rightarrow \infty, m/n \rightarrow \alpha$) performance can be analyzed rigorously using the so-called state evolution [42–45]. However, GAMP is not expected to be always information-theoretically optimal. On the other hand, other results are concerned with the linear case of the GLM with additive Gaussian noise for which the information-theoretically optimal performance was established in [46–48] (the methodology of these works unfortunately does not generalize straightforwardly

to the important non-linear case or to other types of additive noise). All the other works, giving information-theoretic results for the non-linear case, are based on powerful and sophisticated but *non-rigorous* techniques originating in statistical physics of disordered systems, such as the cavity and replica methods [49]. Historically, the first of these non-rigorous, yet correct, results on information-theoretic limitations of learning was for the perceptron with binary weights and was established using the replica method in [22, 50, 51], including a discontinuous phase transition to perfect learning that appears as the ratio between number of samples and the dimension exceeds $\alpha \approx 1.249$.

In the present paper we close the above gap between mathematically rigorous work and conjectures (some of them several decades old) from statistical mechanics. In particular, we prove that the results for GLMs stemming from the replica method are indeed correct and imply the optimal value of both the estimation and generalization error. These results are summarized in section “Main results”. The proof is based on the *adaptive interpolation method* recently developed in [52] and is of independent interest as it is applicable to a range of other models, see section “Methods and proofs” and the supplementary informations (SI). We compare our information-theoretic results to the performance of the GAMP algorithm and its state evolution (as reviewed briefly in section “Main results”). We determine regions of parameters where this algorithm is or is not information-theoretically optimal. Up to technical assumptions (as specified below), our results apply to all activation functions φ and priors P_0 , thus unifying a large volume of previous work where many particular functions have been analyzed on a case by case basis. This generality allows us to provide a unifying understanding of the types of phase transitions and phase diagrams that we can encounter in GLMs, which is as well of independent interest and we devote section “Application to learning and inference” to its presentation.

2 Main results

This section summarizes our main results. Their formal statement together with all technical assumptions and full proofs are provided in section “Methods and proofs” and in the SI.

For the random GLM problem as defined in the introduction, the optimal way to estimate the ground-truth signal/weights \mathbf{X}^* relies on its posterior probability distribution

$$P(\mathbf{x}|\mathbf{Y}, \Phi) = \frac{1}{\mathcal{Z}(\mathbf{Y}, \Phi)} \prod_{i=1}^n P_0(x_i) \prod_{\mu=1}^m P_{\text{out}}\left(Y_\mu \middle| \frac{[\Phi \mathbf{x}]_\mu}{\sqrt{n}}\right) \quad (2)$$

where we used the prior P_0 of \mathbf{X}^* , and introduced the likelihood P_{out} that an output Y_μ is observed given $\frac{1}{\sqrt{n}}[\Phi \mathbf{x}]_\mu$. $P_{\text{out}}(\cdot | z)$ is the probability density function of $\varphi(z, A)$ (where again the r.v. $A \sim P_A$ accounts for noise). This paper is concerned with the so-called *Bayes-optimal* setting where the prior P_0 and the likelihood P_{out} that appear in the posterior (2) were also used to generate the ground-truth signal \mathbf{X}^* and the labels \mathbf{Y} , using a known random matrix Φ .

A first quantity of interest is the *free entropy* (which is the *free energy* up to a sign) defined as $f_n(\mathbf{Y}, \Phi) \equiv \frac{1}{n} \ln \mathcal{Z}(\mathbf{Y}, \Phi)$. The expectation of the free entropy is equal to minus the conditional entropy density of the observation $-\frac{1}{n} H(\mathbf{Y}|\Phi)$, as well as (up to an additive constant) to the mutual information density between the signal and the observations $\frac{1}{n} I(\mathbf{X}^*; \mathbf{Y}|\Phi)$.

2.1 The free entropy

Our first result is the rigorous determination of the free entropy, in the high-dimensional asymptotic regime $n, m \rightarrow \infty, m/n \rightarrow \alpha$. For a random matrix Φ with independent entries of zero mean and unit variance, for output \mathbf{Y} that was generated using (1), and under appropriate technical assumptions stated precisely in

section “Methods and proofs”, the free entropy converges in probability to:

$$f_n(\mathbf{Y}, \Phi) \equiv \frac{1}{n} \ln \mathcal{Z}(\mathbf{Y}, \Phi) \xrightarrow[n \rightarrow \infty]{\mathbb{P}} \sup_{q \in [0, \rho]} \inf_{r \geq 0} f_{\text{RS}}(q, r; \rho) \quad (3)$$

where $\rho \equiv \mathbb{E}_{P_0}[(X^*)^2]$ and where the *potential* $f_{\text{RS}}(q, r; \rho)$ is

$$f_{\text{RS}}(q, r; \rho) \equiv \psi_{P_0}(r) + \alpha \Psi_{P_{\text{out}}}(q; \rho) - rq/2, \quad (4)$$

$$\psi_{P_0}(r) \equiv \mathbb{E}_{Z_0, X_0} \ln \int dP_0(x) e^{rxX_0 + \sqrt{r}xZ_0 - rx^2/2}, \quad (5)$$

$$\Psi_{P_{\text{out}}}(q; \rho) \equiv \mathbb{E}_{V, W, \tilde{Y}_0} \ln \int \mathcal{D}w P_{\text{out}}(\tilde{Y}_0 | \sqrt{q}V + \sqrt{\rho - q}W), \quad (6)$$

where $\mathcal{D}w = dw \exp(-w^2/2)/\sqrt{2\pi}$ is a standard Gaussian measure and the scalar r.v. are independently sampled from $X_0 \sim P_0$, then $V, W, Z_0 \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$ and $\tilde{Y}_0 \sim P_{\text{out}}(\cdot | \sqrt{q}V + \sqrt{\rho - q}W)$. Only the special linear case with Gaussian P_{out} was known rigorously so far [46–48]. Convergence of the *averaged* free entropy is precisely stated in Theorem 1; the one in probability follows from concentration results in the SI.

One can check by explicit comparison that for specific choices of P_0 and P_{out} the expression (4) is the replica-symmetric free entropy derived in numerous statistical physics papers (thus the RS in f_{RS}), and in particular in [22, 41, 50, 51] for $\varphi(z) = \text{sign}(z)$. The formula for general P_0 and P_{out} was conjectured based on the statistical physics derivation in [13]. Establishing (3) closes these old conjectures and yields an important step towards vindication of the cavity and replica methods for inference, alongside with e.g. [43, 53]. We now discuss the main consequences of this formula.

2.2 Overlap and optimal estimation error

Our second result concerns the overlap between a sample \mathbf{x} from the posterior (2) and the ground-truth. We obtain that as $n, m \rightarrow \infty, n/m \rightarrow \alpha$,

$$\frac{1}{n} |\mathbf{x} \cdot \mathbf{X}^*| \xrightarrow[n \rightarrow \infty]{\mathbb{P}} q^* \quad (7)$$

whenever $q^* = q^*(\alpha)$ the maximizer in formula (3) is unique. This is the case for almost every α (see the SI).

It is a simple fact of Bayesian inference that, given the measurements \mathbf{Y} and the measurement matrix Φ , the estimator $\hat{\mathbf{X}}$ that minimizes the mean-square error with the ground-truth \mathbf{X}^* is the mean of the posterior distribution (2), i.e. $\hat{\mathbf{X}} = \mathbb{E}_{P(\mathbf{x}|\mathbf{Y}, \Phi)}[\mathbf{x}]$. The minimum mean-square error (MMSE) that is achieved by such “Bayes-optimal” estimator is deduced, again in the limit $n \rightarrow \infty, m/n \rightarrow \alpha$, as follows:

$$\text{MMSE} = \frac{1}{n} \mathbb{E} \left[\|\mathbf{X}^* - \hat{\mathbf{X}}\|^2 \right] \rightarrow \rho - q^*. \quad (8)$$

We refer to Theorem 2 in section “Main theorems” for rigorous statements. Again the value of the MMSE was known rigorously so far only for the linear case with Gaussian noise [46–48] (and conjectured for the non-linear case e.g. in [13]).

2.3 Optimal generalization error

Our third result concerns the prediction error, also called generalization error. Consider again the statistical model (1). To define the *Bayes-optimal generalization error*, one is given a new row of the matrix/data point, denoted $\Phi_{\text{new}} \in \mathbb{R}^n$ (in addition to the data Φ and associated outputs \mathbf{Y} used for the learning), and is asked to estimate the corresponding output value Y_{new} . We seek for an estimator $\hat{Y}_{\text{new}} = \hat{Y}_{\text{new}}(\mathbf{Y}, \Phi, \Phi_{\text{new}})$ that achieves

$\mathcal{E}_{\text{gen}} \equiv \min_{\hat{Y}_{\text{new}}} \mathbb{E}[(Y_{\text{new}} - \hat{Y}_{\text{new}})^2]$, i.e. that minimizes the MSE with the true Y_{new} obtained using the ground-truth weights \mathbf{X}^* . Such estimator is again obtained from the posterior: $\hat{Y}_{\text{new}} = \mathbb{E}_{P_A(a)} \mathbb{E}_{P(\mathbf{x}|\mathbf{Y}, \Phi)} \varphi(\frac{1}{\sqrt{n}} \Phi_{\text{new}} \cdot \mathbf{x}, a)$. Note that this is different than the *plug-in* estimator $\tilde{Y}_{\text{new}} = \varphi(\frac{1}{\sqrt{n}} \Phi_{\text{new}} \cdot \hat{\mathbf{X}})$, which leads to a worse MSE than \hat{Y}_{new} . Yet it is often used in practice for deterministic models since most algorithms for generalized linear regression do not provide the full posterior distribution.

Our result states that the optimal generalization error follows from the I-MMSE theorem [54] applied to the free entropy (3) (see the SI for the details). The optimal generalization error reads as $n \rightarrow \infty$, $m/n \rightarrow \alpha$ (q^* is the maximizer in (3))

$$\mathcal{E}_{\text{gen}} \rightarrow \mathbb{E}_{V,a} [\varphi(\sqrt{\rho} V, a)^2] - \mathbb{E}_{V,a} [\mathbb{E}_V [\varphi(\sqrt{q^*} V + \sqrt{\rho - q^*} w, a)]^2], \quad (9)$$

where $V, w \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$ and $a \sim P_A$. See again Theorem 2 in the section “Main theorems” for the precise statement (and Theorems 3 and 4 in the SI).

Note that for labels \mathbf{Y} belonging to a discrete set the MSE might not be a suitable loss and we are more often interested in maximizing the so-called overlap, i.e. the probability of obtaining the correct label. In that case the Bayes-optimal estimator is computed as the argmax of the posterior marginals, rather than as its mean, i.e. for discrete labels $\bar{Y}_{\text{new}} = \text{argmax}_y \mathbb{P}(y = \varphi(\frac{1}{\sqrt{n}} \Phi_{\text{new}} \cdot \mathbf{x}, a))$ where again \mathbf{x} is distributed according to (2), $a \sim P_A$. The replica method has been used to compute the optimal generalization error for the perceptron where $\varphi(x) = \text{sign}(z)$ in the pioneering works of [23, 50, 55]. We note that in this special case the plug-in estimator \tilde{Y}_{new} is actually equal to the optimal one \bar{Y}_{new} .

A final note concerns the issue of overfitting. In optimization-based approaches to learning overfitting may lead to a generalization error which is too large as compared to the training error. In the Bayes-optimal setting the estimators are constructed in order not to overfit. This is related to general properties of Bayes-optimal inference and learning that are called “Nishimori conditions” in the physics literature [13] and that turn out to be crucial in our proofs.

2.4 Optimality of approximate message-passing

While the three results stated above are of an information-theoretic nature, our fourth one concerns the performance of an algorithm to solve random instances of GLMs called generalized approximate message-passing (GAMP) [11–13], which is closely related to the TAP equations developed in statistical physics [41, 56, 57].

The GAMP algorithm can be summarized as follows [11–13]: Given initial estimates $\hat{\mathbf{x}}^0, \mathbf{v}^0$ for the marginal posterior means and variances of the unknown signal vector \mathbf{X}^* entries, GAMP iterates the following equations, with $g_\mu^0 = 0$:

$$\begin{cases} V^t &= \overline{\mathbf{v}^{t-1}} \\ \boldsymbol{\omega}^t &= \Phi \hat{\mathbf{x}}^{t-1} / \sqrt{n} - V^t \mathbf{g}^{t-1} \\ g_\mu^t &= g_{P_{\text{out}}}(Y_\mu, \omega_\mu^t, V^t) & \forall \mu = 1, \dots, m \\ \lambda^t &= \alpha g_{P_{\text{out}}}^2(\mathbf{Y}, \boldsymbol{\omega}^t, V^t) \\ \mathbf{R}^t &= \hat{\mathbf{x}}^{t-1} + (\lambda^t)^{-1} \Phi^\top \mathbf{g}^t / \sqrt{n} \\ \hat{x}_i^t &= g_{P_0}(R_i^t, \lambda^t) & \forall i = 1, \dots, n \\ v_i^t &= (\lambda^t)^{-1} \partial_R g_{P_0}(R, \lambda^t)|_{R=R_i^t} & \forall i = 1, \dots, n \end{cases}$$

(here we denote by $\bar{\mathbf{u}}$ the average over all the components of a vector \mathbf{u}). The so-called thresholding function $g_{P_0}(R, \lambda)$ is defined as the mean of the normalized distribution $\propto P_0(x) \exp(-\lambda(R - x)^2/2)$ and the output function $g_{P_{\text{out}}}(Y, \omega, V)$ is similarly the mean of the normalized distribution (of x) $\propto P_{\text{out}}(Y|\omega +$

$\sqrt{V}x) \exp(-x^2/2)$.

The heuristic derivation of GAMP in statistical physics [13] suggests via the definition of the function $g_{P_{\text{out}}}$ that ω and V are the estimates of the means and average variance of the components of the variable $\mathbf{z} = \Phi \mathbf{x}$. This, in turn, suggests a GAMP prediction of labels of new data points:

$$\hat{Y}_{\text{new}}^{\text{GAMP},t} = \int y P_{\text{out}}(y|\omega_{\text{new}}^t + z\sqrt{V^t}) dy \mathcal{D}z$$

where $\omega_{\text{new}}^t \equiv \frac{1}{\sqrt{n}} \Phi_{\text{new}} \cdot \hat{\mathbf{x}}^{t-1}$. Comparing it with the test-set labels, this serves to compute GAMP's generalization error.

One of the strongest assets of GAMP is that its performance can be tracked via a closed form procedure known as state evolution (SE), again in the asymptotic limit when $n, m \rightarrow \infty, m/n \rightarrow \alpha$. For proofs of SE see [43, 44] for the linear case, and [45] for the generalized one. In our notations, SE tracks the correlation (or ‘‘overlap’’) between the true weights \mathbf{X}^* and their estimate $\hat{\mathbf{x}}^t$ defined as $q^t \equiv \lim_{n \rightarrow \infty} \frac{1}{n} \mathbf{X}^* \cdot \hat{\mathbf{x}}^t$ via:

$$q^t = 2\psi'_{P_0}(r^t), \quad r^t = 2\alpha\Psi'_{P_{\text{out}}}(q^{t-1}; \rho). \quad (10)$$

The derivatives are w.r.t. the first argument. Similarly for the evolution of GAMP's generalization error $\mathcal{E}_{\text{gen}}^{\text{GAMP},t}$ (see SI) we get that it is asymptotically, and with high probability, given by the r.h.s. of formula (9) but with q^* replaced by q^t .

It is a simple algebraic fact that the fixed points of the SE equations (10) correspond to the critical points of the potential (4). The question of GAMP achieving asymptotically optimal MMSE or generalization error therefore reduces to the study of the extrema of the two-scalar-variables potential (4). If the SE (10) converges to the same couple (q, r) as the extremizer (q^*, r^*) of (3) then GAMP is optimal, and if it does not then GAMP is sub-optimal. In the next section we illustrate this result on several examples, delimiting regions where GAMP reaches optimality. We note that optimality of AMP-based algorithms in terms of the MMSE on the ground-truth vector \mathbf{X}^* was proven for several cases where the extremizer q^* in (3) is unique, see e.g. [58], or in the linear case of GLM in [47]. Our results allow to complete the characterization of regions of parameters where the algorithm reaches optimal performance in terms of the estimation and generalization errors. While the asymptotic value of the Bayes-optimal generalization error was predicted for some cases of P_{out} and P_0 [55], and TAP-based algorithms were argued to reach this performance in [59, 60], it was not known whether this error can be achieved provably nor for what exact regions of parameters the algorithm is sub-optimal. Our present work settles this question thanks to the state evolution of the GAMP algorithm. Interestingly, heuristic arguments based on the glassy nature of the corresponding probability measure were used to argue that direct sampling or optimization-based approaches will not be able to match this performance [51]. Whether this statement is correct goes beyond the scope of the present paper.

3 Application to learning and inference

In this section, we report what our results imply for the information-theoretically optimal errors, and those reached by the GAMP algorithm for several interesting cases of output functions φ and prior distributions P_0 . We do not seek to be exhaustive in any way, we simply aim to illustrate the kind of insights about the GLM that can be obtained from our results. We focus on determination of phase transitions in performance as we vary parameters of the model, e.g. the number of samples or the sparsity of the signal. We use careful numerical procedures to compute the expectations required in the formula (4), and check that the reported results are stable towards the choice of various precision-parameters. In this section we, however, do not seek rigor in bounding formally the corresponding numerical errors. Many of the codes used in this section are

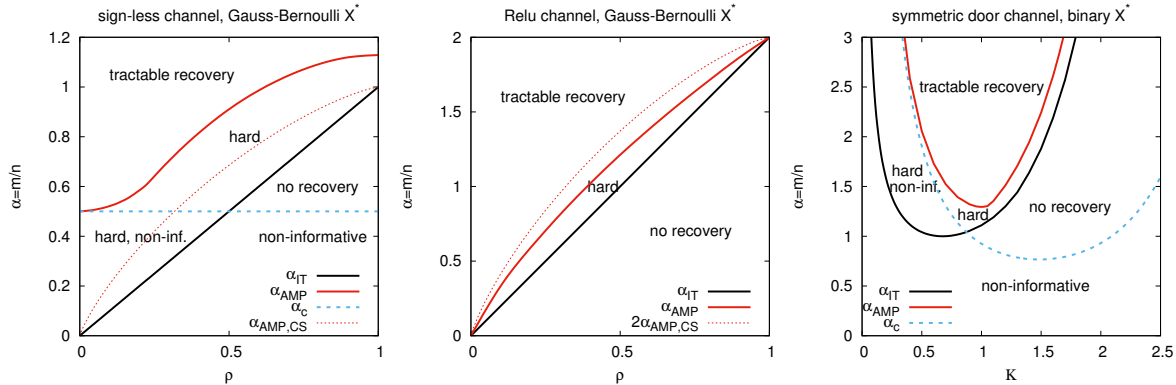


Figure 1: Phase diagrams showing boundaries of the region where almost exact recovery is possible (in absence of noise). **Left:** The case of sign-less sparse recovery, $\varphi(x) = |x|$ with a Gauss-Bernoulli signal, as a function of the ratio between number of samples/measurements and the dimension $\alpha = m/n$, and the fraction of non-zero components ρ . Evaluating (4) for this case, we find that a recovery of the signal is information-theoretically impossible for $\alpha < \alpha_{IT} = \rho$. Recovery becomes possible starting from $\alpha > \rho$, just as in the canonical compressed sensing. Algorithmically the sign-less case is much harder. Evaluating (11) we conclude that GAMP is not able to perform better than a random guess as long as $\alpha < \alpha_c = 1/2$, and the same is true for spectral algorithms, see [61]. For larger values of α , the inference using GAMP leads to better results than a purely random guess. GAMP can recover the signal and generalize perfectly only for values of α larger than α_{AMP} (full red line). The dotted red line shows for comparison the algorithmic phase transition of the canonical compressed sensing. **Center:** Analogous to the left panel, for the ReLU output function, $\varphi(x) = \max(0, x)$. Here it is always possible to perform better than random guessing using GAMP. The dotted red line shows the algorithmic phase transition when using information only about the non-zero observations. **Right:** Phase diagram for the symmetric door output function $\varphi(z) = \text{sign}(|z| - K)$ for a Rademacher signal, as a function of α and K . The stability line α_c is depicted in dashed blue, the information-theoretic phase transition to almost exact recovery α_{IT} in black, and the algorithmic one α_{AMP} in red.

given online in a github repository [62].

3.1 General observations about fixed points and terminology

• **Non-informative fixed point and its stability:** It is instrumental to analyze under what conditions $q^* = 0$ is the optimizer in (3). Our result (8) about the MMSE implies that if $q^* = 0$ then the MMSE is as large as if we had no samples/measurements at our disposition. A necessary condition for $q^* = 0$ is that it is a fixed point of the state evolution. In turn, a sufficient condition for the state evolution (10) to have such a fixed point is that *i*) the output density $P_{\text{out}}(y|z)$ is even in the argument z , and *ii*) that the prior P_0 has zero mean. A proof of this is given in the SI. In order for $q^* = 0$ to be a fixed point to which the state evolution (10) converges, it needs to be stable. We detail in the SI that under properties *i*) and *ii*) this fixed point is stable when

$$\alpha \int dy \frac{\left(\int \mathcal{D}z (z^2 - 1) P_{\text{out}}(y|\sqrt{\rho}z) \right)^2}{\int \mathcal{D}z P_{\text{out}}(y|\sqrt{\rho}z)} < 1. \quad (11)$$

In what follows we will denote α_c the largest value of α for which the above condition holds. Consequently the error reachable by the GAMP algorithm is as bad as random guessing for both the estimation and generalization errors as long as $\alpha < \alpha_c$. For $\alpha > \alpha_c$, starting with infinitesimal positive q the state evolution will move towards larger q as in [63]. Note that the condition (11) also appears in a recent work [61] as a barrier for performance of spectral algorithms.

Concerning the information-theoretically optimal error, we will call the phase where $\text{MMSE} = \rho$, i.e. $q^* =$

0 is the extremizer of (4), the *non-informative phase*. Existing literature sometimes refers to such behavior as *retarded learning* phase [64], in the sense that in that case a critical number of samples is required for the generalization error to be better than random guessing. Below we will evaluate condition (11) explicitly for several examples.

• **Almost exact recovery fixed point:** Another fixed point of (10) that is worth our particular attention is the one corresponding to almost exact recovery, meaning with average error per coordinate going to 0 as $n \rightarrow \infty$, where $q^* = \rho$. A sufficient and necessary condition for this to be a fixed point is that $\lim_{q \rightarrow \rho} \Psi'_{P_{\text{out}}}(q; \rho) = +\infty$. This means that the integral of the Fisher information of the output channel diverges:

$$\int dy d\omega \frac{e^{-\frac{\omega^2}{2\rho}}}{\sqrt{2\pi\rho}} \frac{P'_{\text{out}}(y|\omega)^2}{P_{\text{out}}(y|\omega)} = +\infty,$$

where $P'_{\text{out}}(y|\omega)$ denotes the partial derivative w.r.t. ω . This typically means that the output channel should be noiseless. For example, for the Gaussian channel with noise variance Δ , the above expression equals $1/\Delta$. For the probit channel where $P_{\text{out}}(y|z) = \text{erfc}(-yz/\sqrt{2\Delta})/2$ the above expression at small Δ is proportional to $1/\sqrt{\Delta}$.

Stability of the almost exact recovery fixed point depends non-trivially on both the properties of the output channel, and of the prior. Below we give several examples where almost exact recovery either is or is not possible. In what follow we call the region of parameters for which $\text{MMSE} = 0$, i.e. $q^* = \rho$ is the extremizer in (3), the *almost exact recovery* phase.

• **Hard phase:** As can be anticipated from the statement of our main algorithmic result, there are regions of parameters for which the error reached by GAMP is asymptotically equal to the optimal error, and regions where it is not. We will call *hard phase* the region of parameters where $\text{MMSE} < \text{MSE}_{\text{AMP}}$ with a strict inequality. Focusing on the ratio α between the number of samples and the dimensionality, we will denote α_{IT} the ratio for which the hard phase appears, and $\alpha_{\text{AMP}} > \alpha_{\text{IT}}$ the ratio for which it disappears. In other words, the hard phase is an interval $(\alpha_{\text{IT}}, \alpha_{\text{AMP}})$, and is associated to a first order phase transition in the Bayes-optimal posterior probability distribution.

It remains a formidable open question of average computational complexity whether in the setting of this paper (and for problems that are NP-complete in the worst case) there exists an efficient algorithm that achieves better performance than GAMP in the hard phase. The authors are not aware of any, and tend to conjecture that there is none.

3.2 Sensing compressively with non-linear outputs

Existing literature covers in detail the case of noiseless compressed sensing, i.e. when the output function $\varphi(z) = z$. Representative sparse prior distribution is the Gauss-Bernoulli (GB) distribution $P_0 = \rho\mathcal{N}(0, 1) + (1 - \rho)\delta_0$, where ρ is the average fraction of non-zeros, which are in this case standard Gaussians. The phase diagram of this case is well known see e.g. [67, 68]. In noiseless compressed sensing with random i.i.d. matrices and GB prior, almost exact recovery of the signal is possible for $\alpha > \alpha_{\text{IT}} = \rho$ and GAMP recovers the signal for $\alpha > \alpha_{\text{AMP,CS}}$ where $\alpha_{\text{AMP,CS}}$ is plotted in Fig. 1 (left) with a dotted red line, thus delimiting the hard phase of compressed sensing. We note that the Donoho-Tanner phase transition [9] known as the performance limit of the LASSO ℓ_1 regularization is slightly higher than $\alpha_{\text{AMP,CS}}$.

• **Sign-less output channel:** The phase diagram of noiseless compressed sensing changes intriguingly when only the absolute value of the output is measured, i.e. when $\varphi(z) = |z|$ instead of $\varphi(z) = z$. Such an output channel is reminiscent of the widely studied phase retrieval problem [3] where the signal is complex valued and only the amplitude is observed. Generalization of our results to the complex case would require extensions, as done for the algorithmic aspects in [69]. The real-valued case was studied under the name *sparse recovery from quadratic measurements* in the literature, e.g. [70] and references therein, when the number of

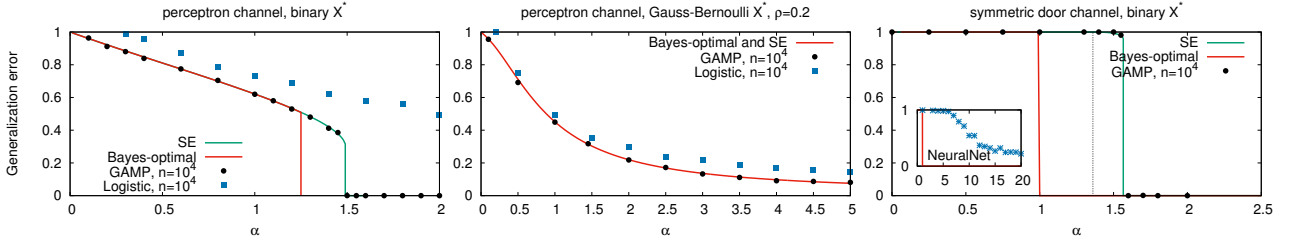


Figure 2: Optimal generalization error in three classification problems versus the sample complexity α , the size of the training set being αn . The red line is the Bayes-optimal generalization error (9) while the green line shows the (asymptotic) performances of GAMP as predicted by the state evolution (10). For comparison, we also show the results of GAMP (black dots) and, in blue, the performance of a standard out-of-the-box solver. **Left:** Perceptron, with $\varphi(x) = \text{sign}(x)$ and a binary Rademacher signal. While a perfect generalization is information-theoretically possible starting from $\alpha_{\text{IT}} \approx 1.249$, the state evolution predicts that GAMP will achieve such perfect prediction only above $\alpha_{\text{AMP}} \approx 1.493$. The results of a logistic regression with fine-tuned regularizations with the software scikit-learn [65] are shown for comparison. **Middle:** Perceptron with Gauss-Bernoulli distribution of the weights. No phase transition is observed in this case, but a smooth decrease of the error with α . The results of a logistic regression are very close to optimal. **Right:** The symmetric door activation rule with parameter K chosen in order to observe the same number of occurrence of the two classes. In this case there is a sharp phase transition from as bad as random to perfect generalization at $\alpha_{\text{IT}} = 1$. GAMP identifies the rule perfectly only starting from $\alpha_{\text{AMP}} \approx 1.566$. The non-informative fixed point is stable up to $\alpha_c = 1.36$ (dashed line). Interestingly, this non linear rule seems very hard to learn for standardly used solvers. Using Keras [66], a neural network with 2 hidden layers was able to learn only approximately the rule, only for considerably larger training set sizes and much larger number of iterations than GAMP.

non-zero variables grows slower than linearly with the dimension n . Our results give access to the phase diagram of sparse recovery from quadratic (or equivalently sign-less) measurements, that is presented in Fig. 1 (left) for the GB prior.

We observe that the information-theoretical phase transition α_{IT} is the same in the sign-less sparse recovery as in the canonical linear case, i.e. almost exact recovery is possible whenever $\alpha > \rho$. However, the algorithmic phase transition α_{AMP} above which GAMP is able to find the sparse signal¹ is strikingly larger for the sign-less case (red line in left panel of Fig. 1). We note that even for a dense signal $\rho = 1$ almost exact recovery is algorithmically possible only for $\alpha > \alpha_{\text{AMP}}(\rho = 1) \approx 1.128$. For very sparse signals, small ρ , the situation is even more striking because measurement rate of at least $\alpha > \alpha_c = 1/2$ is needed for algorithmically tractable almost exact recovery for every ρ . This is in sharp contrast with the canonical compressed sensing where $\alpha_{\text{AMP,CS}} \rightarrow 0$ as $\rho \rightarrow 0$. The nature of this algorithmic difficulty of GAMP is related to the symmetry of the output channel thanks to which the non-informative fixed point is stable for $\alpha < \alpha_c = 1/2$. Summarizing this result in one sentence, tractable compressive sensing is impossible (for $\alpha < 1/2$) if we have lost the signs. We remind that this result holds in the setting of the present paper, i.e. in particular when the sparsity ρ is of constant order. For signals where $\rho = o(1)$ the situation is expected to be different [70].

• **ReLU output channel:** Another case of output channel that attracted our interest is the rectified linear unit (ReLU), $\varphi(z) = \max(0, z)$, as widely used in multi-layer feed-forward neural networks. In the present single-layer case reconstruction with the ReLU output is interesting mathematically. With GB signals, roughly half of the measurements are given without noise, but the only information we have about the other half is its sign. A straightforward upper-bound for both information-theoretic and tractable almost exact recovery is simply twice as many measurements than needed in the canonical noiseless compressed sensing. It is interesting to ask whether this bound is tight. Results in the present paper imply that for the information-theoretic performance this bound indeed is tight. However, the phase transition α_{AMP} above which almost exact re-

¹We note that in order to break the symmetry that prevents GAMP to find the signal in constant number of iteration steps, we mismatch infinitesimally the output function φ used in the algorithm from the symmetric one used to generate the data. Another way to deal with this issue is related to a spectral initialization as discussed recently in [61].

covery is possible with the GAMP algorithm is strictly lower than twice the phase transition of compressed sensing; both are depicted in the central panel on Fig. 1. This implies that while the negative outputs are not useful information-theoretically, they do help to achieve better performance algorithmically.

3.3 Perceptron and alike

- **Binary and Gauss-Bernoulli perceptron:** One of the most studied problems that fits in the setting of the present paper is the problem of perceptron [71], where $\varphi(z) = \text{sign}(z)$, that has been analyzed for random patterns Φ in the statistical physics literature, see [23–25] for reviews. We plot in Fig. 2 the optimal generalization error (9) as follows from our results for the binary perceptron, i.e. weights taken from the Rademacher distribution $P_0 = \frac{1}{2}\delta_{+1} + \frac{1}{2}\delta_{-1}$ (left panel), and for the GB perceptron where $P_0 = \rho\mathcal{N}(0, 1) + (1-\rho)\delta_0$ (central panel). The information-theoretically optimal value of the generalization error that we report and prove agrees with existing predictions obtained by the non-rigorous replica method from [50, 51, 55]. Notably, we see that for the GB case the optimal generalization error decreases smoothly as α increases, while for the binary case the generalization error has a first order (i.e. discontinuous) phase transition towards perfect generalization at $\alpha_{\text{IT}} \approx 1.249$ as predicted already in [50]. Our results provide rigorous validation for these old conjectures.

Furthermore, our results together with recent literature on GAMP provide a refreshing clarification of the algorithmic questions. It is natural to ask for what region of parameters the optimal generalization error can be provably achieved with efficient algorithms. This question remained unanswered until now. Indeed, for the spherical perceptron the optimal generalization error was computed in [55], and argued empirically on small instances to be achievable with a TAP-like algorithm [59]. The state evolution of GAMP together with our formulas for the generalization error ((9) for the average optimal one and with q^t replacing q^* in this formula for GAMP) imply that the optimal generalization error is indeed achievable asymptotically for all α in the GB perceptron.

For the binary perceptron the optimal generalization error was computed in [50, 51]. By comparison with the state evolution of GAMP we obtain that it can also be asymptotically achieved by GAMP, but this time only outside of the hard phase $(\alpha_{\text{IT}}, \alpha_{\text{AMP}})$ with $\alpha_{\text{AMP}} \approx 1.493$. The past literature was unclear on the algorithmic question, Ref. [50] identified the spinodal of the replica-symmetric solution to be at $\alpha \approx 1.493$, but did not attribute it any algorithmic nor physical meaning. Ref. [51] argues that metastable states exist at least up to $\alpha_{\text{RSB}} \approx 1.628$ and speculates that Gibbs sampling based algorithms will not be able to reach perfect generalization before that point [23]. Taking our results into account, the main algorithmic question that remains open is whether efficient algorithms can reach perfect generalization for $\alpha_{\text{IT}} < \alpha < \alpha_{\text{AMP}}$.

- **Symmetric door:** Out of interest we explored an example of binary output channel for which $P_{\text{out}}(y|z)$ is even in the argument z , so that the non-informative fixed point $q^* = 0$ exists. Specifically we analyzed the *symmetric door* channel with $\varphi(z) = \text{sign}(|z| - K)$ and Rademacher prior P_0 . In existing literature such a perceptron was studied with the replica method in the context of lossy data compression [28]. In Fig. 1 (right panel) we report the phase diagram in terms of the stability line of the non-informative fixed point α_c (below which GAMP is not better than random guesses), the information-theoretic phase transition towards perfect generalization α_{IT} , and the phase transition of GAMP to perfect generalization α_{AMP} .

A simple counting lower bound states that for binary outputs and weights X_i^* perfect generalization is not possible for $\alpha < 1$. Thus it is interesting to notice that the symmetric door channel is able to saturate this lower-bound for $K \approx 0.6745$ for which the probability of $y_\mu = 1$ is $1/2$. This saturation was already remarked in [28]. Our results also, however, imply that in that case the perfect generalization will not be achievable with GAMP (and we conjecture no other efficient) algorithm unless $\alpha > \alpha_{\text{AMP}} \approx 1.566$. The generalization error that GAMP provides for this case is depicted in Fig. 2 (right).

3.4 Empirical comparison with general purpose algorithms

In this section we argue that many cases that fit into the setting of the present paper could serve as useful benchmarks for existing machine learning algorithms. We believe that the situation is perhaps similar to Shannon coding theorems that have driven algorithmic developments in error correcting codes, achieving the Shannon bound being the primary goal in many works in communications. In machine learning, classification is a natural task and algorithms are usually benchmarked using open access databases. In current state-of-the-art applications of machine learning we usually have very little insight about what is the sample complexity, i.e. how many samples are truly needed so that a given generalization error can be achieved. In our setting the situation is different: We can present samples (y_μ, Φ_μ) to generic out-of-the-box classification algorithms and see how their performances compare to the information-theoretic optimal performance and to the one of the GAMP algorithm that is fine-tuned to the problem.

In Fig. 2 we present examples of state-of-the-art classification algorithms compared to our results. In the left and center panels we compare the optimal and GAMP performances to a simple logistic regression, fine-tuned by manually optimizing the ridge penalty (for ℓ_2 regularization) and LASSO penalty (for a sparsity-enhancing ℓ_1 regularization) with the software scikit-learn [65]. We observe that while for the GB case the logistic regression is comparable to the performance of GAMP, for binary weights perfect generalization is not achieved close to the GAMP phase transition.

In the right panel of Fig. 2 we study classification for labels generated by the symmetric door channel. A general purpose algorithm would not know about the form of the channel. A neural network with only two hidden units is in principle able to represent the corresponding function (each of the hidden neurons can learn one of the two planes that separate data in the symmetric door function). A more intriguing question is whether a more generic multi-layer neural network is indeed able to learn this rule and how many samples does it need? In the example used in Fig. 2, using the software Keras [66] with a tensorflow backend, we show the performance of a network with two hidden layers, ReLU activation and dropout (the details for this particular run can be found on the github repository [62]). The symmetric door function thus provides a challenging benchmark that could be used to study how to improve performance of the general purpose multi-layer neural network classifiers. In the SI we provide additional examples comparing the optimal performance to general-purpose algorithms for regression.

4 Methods and proofs

In this section we give the main theorem for the free entropy and main ideas of the proof. An essential tool is the adaptive interpolation method recently introduced in [52] which is a powerful evolution of the Guerra and Toninelli interpolation method developed for spin glasses [72]. Reference [52] analyzed simpler inference problems. In particular the proof for the upper bound in [52] does not apply to GLMs and requires non-trivial new ingredients. One such new ingredient is to work with a potential $f_{\text{RS}}(q, r; \rho)$ depending on two parameters (q, r) instead of a single one as in [52]. This allows us to use convexity arguments that are crucial in order to finish the proof, see the last section “Matching bounds and end of proof”. We stress that the present analysis heavily relies on properties of Bayes-optimal inference that translate into remarkable identities between correlation functions (called Nishimori identities by physicists; see SI for their formulation) valid for *all* values of parameters. These identities are used in the derivation of (17) and (18) below, which are two essential steps of our proof. The formula from Theorem 1 relies on the Nishimori identities and does not hold out of the Bayes-optimal setting.

4.1 Main theorems

For the proof it is necessary to work with a slightly different model with an additive regularizing Gaussian noise with variance $\Delta \geq 0$:

$$Y_\mu = \varphi\left(\frac{1}{\sqrt{n}}[\Phi \mathbf{X}^*]_\mu, A_\mu\right) + \sqrt{\Delta}Z_\mu, \quad 1 \leq \mu \leq m, \quad (12)$$

where $(Z_\mu) \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$, and $(A_\mu) \stackrel{\text{iid}}{\sim} P_A$ are r.v. that represent the stochastic part of φ . It is also instrumental to think of the measurements as the outputs of a “channel” $Y_\mu \sim P_{\text{out}}(\cdot | \frac{1}{\sqrt{n}}[\Phi \mathbf{X}^*]_\mu)$ with transition density $P_{\text{out}}(y|z) = (2\pi\Delta)^{-1/2} \int dP_A(a) \exp\{-\frac{1}{2\Delta}(y - \varphi(z, a))^2\}$ if $\Delta > 0$, or $P_{\text{out}}(y|z) = \int dP_A(a) \mathbf{1}(y = \varphi(z, a))$ else, where $\mathbf{1}(\cdot)$ is the indicator function. Our main theorem holds under the following rather general hypotheses:

- (h1) The prior distribution P_0 admits a finite third moment and has at least two points in its support.
- (h2) The sequence $(\mathbb{E}[|\varphi(\frac{1}{\sqrt{n}}[\Phi \mathbf{X}^*]_1, \mathbf{A}_1)|^{2+\gamma}])_{n \geq 1}$ is bounded for some $\gamma > 0$.
- (h3) The r.v. $(\Phi_{\mu i})$ are independent with zero mean, unit variance and finite third moment bounded with n .
- (h4) For almost-all values of a (w.r.t. the distribution P_A), the function $x \mapsto \varphi(x, a)$ is continuous almost everywhere.
- (h5) $(\Delta > 0)$ or $(\Delta = 0)$ and φ takes values in \mathbb{N} .

In *general*, when φ is continuous the condition $\Delta > 0$ (but arbitrarily small) is necessary for the existence of a finite limit of the free entropy (for particular choices of (φ, P_A) this might not be needed, e.g. $\varphi(z, A) = z + A$ with $A \sim \mathcal{N}(0, \sigma^2)$). We also assume that the kernel P_{out} is *informative*, i.e. there exists y such that $P_{\text{out}}(y | \cdot)$ is not equal almost everywhere to a constant. If P_{out} is not informative, it is not difficult to show that estimation is then impossible.

We define the set of the critical points of f_{RS} , (4), also called “state evolution fixed points” (as it is clear from (10)):

$$\Gamma \equiv \left\{ (q, r) \in [0, \rho] \times (\mathbb{R}_+ \cup \{+\infty\}) \mid \begin{array}{l} q = 2\psi'_{P_0}(r) \\ r = 2\alpha\Psi'_{P_{\text{out}}}(q; \rho) \end{array} \right\}.$$

Define $f_n \equiv \mathbb{E}f_n(\mathbf{Y}, \Phi) = \frac{1}{n} \mathbb{E} \ln \mathcal{Z}(\mathbf{Y}, \Phi)$. Then the main theorem of this paper is stated as follows:

Theorem 1 (Replica-symmetric free entropy). *Suppose that (h1)-(h2)-(h3)-(h4)-(h5.a) hold. Then, for the GLM (12),*

$$\lim_{n \rightarrow \infty} f_n = \sup_{q \in [0, \rho]} \inf_{r \geq 0} f_{\text{RS}}(q, r) = \sup_{(q, r) \in \Gamma} f_{\text{RS}}(q, r).$$

Moreover, as one can see in the SI, the “sup inf” and the supremum over Γ above are achieved over the same couples. Under stronger assumptions on P_0 and P_{out} , one can show (see Theorem 6 in the SI) that $f_n(\mathbf{Y}, \Phi)$ concentrates around its mean f_n and thus obtains convergence in probability (3).

An immediate corollary of Theorem 1 is the limiting expression of the mutual information $I(\mathbf{X}^*; \mathbf{Y} | \Phi) \equiv \mathbb{E} \ln P(\mathbf{Y}, \mathbf{X}^* | \Phi) - \mathbb{E} \ln (P(\mathbf{Y} | \Phi) P(\mathbf{X}^*))$ between the observations and the unknown vector:

Corollary 1 (Mutual information). *Under the same hypotheses as in Theorem 1, the mut. info. for the GLM (12) verifies*

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n} I(\mathbf{X}^*; \mathbf{Y} | \Phi) &= \inf_{q \in [0, \rho]} \sup_{r \geq 0} i_{\text{RS}}(q, r) = \inf_{(q, r) \in \Gamma} i_{\text{RS}}(q, r), \\ i_{\text{RS}}(q, r) &\equiv \alpha\Psi_{P_{\text{out}}}(\rho; \rho) - \alpha\Psi_{P_{\text{out}}}(q; \rho) - \psi_{P_0}(r) + rq/2. \end{aligned}$$

Finally, we gather our main results related to the optimal errors in a single theorem, see the SI for more details, including results on the optimality of the GAMP algorithm:

Theorem 2 (Optimal errors). *Assume the same hypotheses as in Theorem 1. Then formula (9) for the generalization error is true as $n, m \rightarrow \infty, m/n \rightarrow \alpha$ whenever the maximizer $q^*(\alpha)$ of (3) is unique, which is the case for almost every α . If moreover all the moments of P_0 are finite, then formula (7) for the overlap as well as the following matrix-MMSE formula*

$$\frac{1}{n^2} \mathbb{E} \left[\left\| \mathbf{X}^* \mathbf{X}^{*\top} - \mathbb{E}_{P(\mathbf{x}|\mathbf{Y}, \Phi)}[\mathbf{x}\mathbf{x}^\top] \right\|_{\text{F}}^2 \right] \rightarrow \rho^2 - q^*(\alpha)^2 \quad (13)$$

are true, where $\| \cdot \|_{\text{F}}$ is the Frobenius norm.

There are cases of GLMs (e.g. the sign-less output channel $\mathbf{Y} = |\Phi \mathbf{X}^*|/\sqrt{n} + \mathbf{Z}$) where the sign of \mathbf{X}^* simply cannot be estimated (thus the absolute value in (7)). This is why our general theorem is related to an error metric (13) insensitive to this \pm symmetry. Nevertheless formula (8) for the signal MSE is formally valid when there is no such sign symmetry.

4.2 Proof by the adaptive interpolation method

We now give the main ideas behind the proof of Theorem 1. We defer to the SI the details, as well as those of Corollary 1 and Theorem 2.

A word about notation. The r.v. \mathbf{Y} (and also $\Phi, \mathbf{X}^*, \mathbf{A}, \mathbf{Z}$) are called *quenched* because once the measurements are acquired they are fixed. The expectation w.r.t. *all* quenched r.v. will be denoted by \mathbb{E} *without* subscript. In contrast, expectation of *annealed* variables w.r.t. a posterior distribution at fixed quenched variables is denoted by *Gibbs brackets* $\langle - \rangle$.

4.2.1 Two scalar inference channels

An important role in the proof is played by two simple *scalar* inference channels. The free entropy is expressed in terms of the free entropies of these channels. This “decoupling property” stands at the root of the replica approach in statistical physics.

The first scalar channel is an additive Gaussian channel. Suppose that we observe $Y_0 = \sqrt{r} X_0 + Z_0$ where $X_0 \sim P_0$ and $Z_0 \sim \mathcal{N}(0, 1)$ are independent. Consider the inference problem consisting of retrieving X_0 from the observation Y_0 . The free entropy associated with this channel is the expectation of the logarithm of the normalization factor of the associated posterior $dP(x|Y_0)$, that is given by (5) (up to a constant).

The second scalar channel that appears naturally in the problem is linked to the channel P_{out} through the following inference model. Suppose that $V, W^* \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$ where V is *known* while the inference problem is to recover the unknown W^* from the observation $\tilde{Y}_0 \sim P_{\text{out}}(\cdot | \sqrt{q} V + \sqrt{\rho - q} W^*)$ where $\rho > 0$ and $q \in [0, \rho]$. The free entropy for this model, again related to the average logarithm of the normalization factor of the posterior of w given \tilde{Y}_0 and V , is exactly (6).

4.2.2 Interpolating estimation problem

To carry on the proof, we introduce an “interpolating estimation problem” that interpolates between the original problem $Y_\mu \sim P_{\text{out}}(\cdot | \frac{1}{\sqrt{n}}[\Phi \mathbf{X}^*]_\mu)$ at $t = 0, t \in [0, 1]$ being the interpolation parameter, and the two scalar problems described above at $t = 1$. For $t \in (0, 1)$ the interpolating estimation problem is a mixture of the original and the scalar problems. This interpolation scheme is inspired by the interpolation paths used by Talagrand to study the perceptron, see [73]. Thanks to a novel ingredient specific to the adaptive interpolation method [52], it allows to obtain in a unified manner a complete proof of the replica formula for the free entropy and this in the whole phase diagram.

Let $q_\epsilon : [0, 1] \rightarrow [0, \rho], r_\epsilon : [0, 1] \rightarrow [0, r_{\text{max}}], r_{\text{max}} \equiv 2\alpha \Psi'_{P_{\text{out}}}(\rho; \rho)$, be two continuous “interpolating functions” parametrized by $\epsilon = (\epsilon_1, \epsilon_2) \in \mathcal{B}_n \equiv [s_n, 2s_n]^2$, with $(s_n)_{n \geq 1} \in (0, 1/2]^{\mathbb{N}}$ a sequence

that tends to zero slowly enough. Set $R_1(t, \epsilon) \equiv \epsilon_1 + \int_0^t r_\epsilon(v)dv$, $R_2(t, \epsilon) \equiv \epsilon_2 + \int_0^t q_\epsilon(v)dv$ and define $S_{t,\mu} = S_{t,\mu}(\mathbf{X}^*, W_\mu^*, V_\mu, \Phi)$ as

$$S_{t,\mu} \equiv \sqrt{\frac{1-t}{n}} [\Phi \mathbf{X}^*]_\mu + \sqrt{R_2(t, \epsilon)} V_\mu + \sqrt{\rho t - R_2(t, \epsilon) + 2s_n} W_\mu^*$$

where $(V_\mu), (W_\mu^*) \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$. Consider the following observation channels, with two types of observations obtained through

$$\begin{cases} Y_{t,\mu} & \sim P_{\text{out}}(\cdot | S_{t,\mu}), & \text{for } 1 \leq \mu \leq m, \\ Y'_{t,i} & = \sqrt{R_1(t, \epsilon)} X_i^* + Z'_i, & \text{for } 1 \leq i \leq n, \end{cases} \quad (14)$$

where $(Z'_i) \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$. We assume that $\mathbf{V} = (V_\mu)_{\mu=1}^m$ is known and that the inference problem is to recover both $\mathbf{W}^* = (W_\mu^*)_{\mu=1}^m$ and $\mathbf{X}^* = (X_i^*)_{i=1}^n$ from the “ t -dependent” observations $\mathbf{Y}_t = (Y_{t,\mu})_{\mu=1}^m$ and $\mathbf{Y}'_t = (Y'_{t,i})_{i=1}^n$.

We now understand that (R_1, R_2) and $1-t$ appearing in the first and second set of measurements in (14) play the role of signal-to-noise ratios (snr) in the interpolating problem, with t giving more and more “power” to the scalar inference channels when increasing. Here is the first crucial ingredient of our interpolation scheme. In classical interpolations, these snr would all take a trivial form, i.e. be linear in t , but here, the non-trivial integral dependency in t of (R_1, R_2) allows for much more flexibility when choosing the interpolation path. This allows us to actually choose the “optimal interpolation path”. This will become clear below as well as the role of the “small perturbation” parameters (ϵ_1, ϵ_2) .

Define $u_y(x) \equiv \ln P_{\text{out}}(y|x)$ and, with a slight abuse of notations, we also define the quantity $s_{t,\mu} = s_{t,\mu}(\mathbf{x}, w_\mu, V_\mu, \Phi) \equiv S_{t,\mu}(\mathbf{x}, w_\mu, V_\mu, \Phi)$, the expression above with \mathbf{X}^*, W_μ^* replaced by \mathbf{x}, w_μ . We introduce the *interpolating Hamiltonian* $\mathcal{H}_t = \mathcal{H}_t(\mathbf{x}, \mathbf{w}; \mathbf{Y}_t, \mathbf{Y}'_t, \Phi, \mathbf{V})$

$$\mathcal{H}_t \equiv - \sum_{\mu=1}^m u_{Y_{t,\mu}}(s_{t,\mu}) + \frac{1}{2} \sum_{i=1}^n (Y'_{t,i} - \sqrt{R_1(t, \epsilon)} x_i)^2$$

and the corresponding (t -dependent) *Gibbs bracket* $\langle - \rangle_t$ which is the expectation w.r.t. the joint posterior distribution of (\mathbf{x}, \mathbf{w}) given the observations $\mathbf{Y}_t, \mathbf{Y}'_t$ (and Φ, \mathbf{V}), defined as

$$\langle L(\mathbf{x}, \mathbf{w}) \rangle_t \equiv \frac{1}{\mathcal{Z}_t(\mathbf{Y}_t, \mathbf{Y}'_t, \Phi, \mathbf{V})} \int dP_0(\mathbf{x}) \mathcal{D}\mathbf{w} L(\mathbf{x}, \mathbf{w}) e^{-\mathcal{H}_t},$$

for every continuous bounded test function L . Here $\mathcal{Z}_t \equiv \int dP_0(\mathbf{x}) \mathcal{D}\mathbf{w} \exp\{-\mathcal{H}_t(\mathbf{x}, \mathbf{w}; \mathbf{Y}_t, \mathbf{Y}'_t, \Phi, \mathbf{V})\}$ is the appropriate normalization, $\mathcal{D}\mathbf{w}$ is the standard Gaussian measure. Finally we introduce

$$f_{n,\epsilon}(t) \equiv \frac{1}{n} \mathbb{E} \ln \mathcal{Z}_t(\mathbf{Y}_t, \mathbf{Y}'_t, \Phi, \mathbf{V})$$

which is the interpolating free entropy. One verifies that

$$\begin{cases} f_{n,\epsilon}(0) = f_n - \frac{1}{2} + \mathcal{O}(s_n), \\ f_{n,\epsilon}(1) = \psi_{P_0}(\int_0^1 r_\epsilon(t) dt) - \frac{1}{2}(1 + \rho \int_0^1 r_\epsilon(t) dt) + \frac{m}{n} \Psi_{P_{\text{out}}}(\int_0^1 q_\epsilon(t) dt; \rho) + \mathcal{O}(s_n), \end{cases} \quad (15)$$

where $|\mathcal{O}(s_n)| \leq C s_n$ for a constant $C > 0$. Now comes another crucial property of the interpolating model: It is such that at $t = 0$ we recover the original problem and $f_{n,\epsilon}(0) = f_n - 1/2 + \mathcal{O}(s_n)$ (the constant $1/2$ comes from the purely noisy measurements of the second channel in (14)), while at $t = 1$ we have two scalar inference channels and thus the associated terms ψ_{P_0} and $\Psi_{P_{\text{out}}}$ appear in $f_{n,\epsilon}(1)$. These are precisely the

terms appearing in the free entropy potential (4).

4.2.3 Entropy variation along the interpolation

From the understanding of the previous section, it is natural to evaluate the variation of entropy along the interpolation, which allows to “compare” the original and purely scalar models thanks to the identity

$$f_n = f_{n,\epsilon}(1) - \int_0^1 \frac{df_{n,\epsilon}(t)}{dt} dt + \frac{1}{2} + \mathcal{O}(s_n). \quad (16)$$

Then by choosing the optimal interpolation path thanks to the non-trivial snr dependencies in t , we are able to show the equality between the replica formula and the free entropy f_n .

We thus compute the t -derivative of the free entropy (see the SI for the details of this calculation). It is given by

$$\frac{df_{n,\epsilon}(t)}{dt} = \frac{r_\epsilon(t)}{2} (q_\epsilon(t) - \rho) - \frac{1}{2} \mathbb{E} \left\langle \left(\frac{1}{n} \sum_{\mu=1}^m u'_{Y_{t,\mu}}(S_{t,\mu}) u'_{Y_{t,\mu}}(s_{t,\mu}) - r_\epsilon(t) \right) (Q - q_\epsilon(t)) \right\rangle_t + \mathcal{O}_n(1), \quad (17)$$

where $\mathcal{O}_n(1)$ is a quantity that goes to 0 in the $n, m \rightarrow \infty$ limit, uniformly in t, ϵ , and in the interpolating functions q_ϵ, r_ϵ . The *overlap* is $Q = Q_n \equiv \sum_{i=1}^n X_i^* x_i / n$.

We now state a crucial result in an informal way and refer to the SI for precise statements. Formally, the overlap concentrates around its mean (for all $t \in [0, 1]$), a behaviour called “replica-symmetric” in statistical physics. In order to make this statement rigorous, one has to include the ϵ -dependent small perturbation in (14) which effectively adds “side-information” about \mathbf{X}^* (e.g. think of $t = 0$) without affecting the asymptotic free entropy density. This perturbation forces the overlap to concentrate. We prove that: If for each t the map $R^t : (\epsilon_1, \epsilon_2) \in \mathcal{B}_n \mapsto (R_1(t, \epsilon), R_2(t, \epsilon)) \in R^t(\mathcal{B}_n)$ is a \mathcal{C}^1 diffeomorphism whose Jacobian has determinant greater or equal to 1, then we have for $s_n = \frac{1}{2} n^{-1/16}$ (see Proposition 4 of Sec. 4.3 in the SI for the precise statement)

$$\frac{1}{s_n^2} \int_{\mathcal{B}_n} d\epsilon \int_0^1 dt \mathbb{E} \langle (Q - \mathbb{E} \langle Q \rangle_t)^2 \rangle_t = \mathcal{O}(n^{-1/8}). \quad (18)$$

As will be seen below it is possible to choose interpolating functions that satisfy the required condition.

4.2.4 Canceling the remainder

Note from (15) and (4) that the first term appearing in (17) is precisely the missing one to obtain the expression of the potential on the r.h.s. of (16). Thus we would like to “cancel” the Gibbs bracket in (17). This term is called *remainder*. In order to prove the replica formula, we have to show that this remainder vanishes: Thanks to the freedom of choice of interpolation paths (r_ϵ, q_ϵ) , we are able to do so by “adapting” the interpolation. Thus we would like to choose $q_\epsilon(t) = \mathbb{E} \langle Q \rangle_t \approx Q$ because of (18). However, $\mathbb{E} \langle Q \rangle_t$ is a function of $\int_0^t q_\epsilon(v) dv$ (and of t, ϵ and $\int_0^t r_\epsilon(v) dv$). The equation $q_\epsilon(t) = \mathbb{E} \langle Q \rangle_t$ is therefore a first order differential equation over $t \mapsto \int_0^t q_\epsilon(v) dv$. Assuming for the moment that this differential equation has a solution, the Cauchy-Schwarz inequality applied to the remainder together with (18) allows to show that the absolute value of this remainder *integrated* over $(\epsilon, t) \in \mathcal{B}_n \times [0, 1]$ is $\mathcal{O}(s_n^2 n^{-1/16})$. Combining this result with (15) and (16) leads to the following *fundamental sum rule* (Proposition 5 of Sec. 4.3 in SI):

$$f_n = \frac{1}{s_n^2} \int_{\mathcal{B}_n} d\epsilon \left\{ \psi_{P_0} \left(\int_0^1 r_\epsilon(t) dt \right) + \alpha \Psi_{P_{\text{out}}} \left(\int_0^1 q_\epsilon(t) dt; \rho \right) - \frac{1}{2} \int_0^1 r_\epsilon(t) q_\epsilon(t) dt \right\} + \mathcal{O}_n(1). \quad (19)$$

4.2.5 Matching bounds and end of proof

We now possess all the necessary tools to finish the proof of Theorem 1. We first prove that $\lim_{n \rightarrow \infty} f_n = \sup_{r \geq 0} \inf_{q \in [0, \rho]} f_{\text{RS}}(q, r)$. Then in the SI, we show that *i*) this is also equal to $\sup_{q \in [0, \rho]} \inf_{r \geq 0} f_{\text{RS}}(q, r)$ which gives the first equality of the theorem; *ii*) that this sup inf is attained at the supremum of the state evolution fixed points, which gives the second equality.

• **Lower bound:** We choose the constant function $r_\epsilon(t) = r$ for $t \in [0, 1]$. In the SI we show, using the Cauchy-Lipschitz theorem and the Liouville formula, that the differential equation $q_\epsilon(t) = \mathbb{E}\langle Q \rangle_t$ possesses a unique solution and that the map $R^t : (\epsilon_1, \epsilon_2) \mapsto (\epsilon_1 + rt, \epsilon_2 + \int_0^t q_\epsilon(v) dv)$ is a \mathcal{C}^1 diffeomorphism with Jacobian greater than 1 (so (18) is valid). Identity (19) then implies $\liminf_{n \rightarrow \infty} f_n \geq \inf_{q \in [0, \rho]} f_{\text{RS}}(q, r)$ for all $r \in [0, r_{\text{max}}]$. Thus $\liminf_{n \rightarrow \infty} f_n \geq \sup_{r \in [0, r_{\text{max}}]} \inf_{q \in [0, \rho]} f_{\text{RS}}(q, r)$. In the SI an easy argument shows the r.h.s. is in fact equal to $\sup_{r \geq 0} \inf_{q \in [0, \rho]} f_{\text{RS}}(q, r)$.

• **Upper bound:** We choose the interpolating functions as solutions of the following system of 1st order differential equations: $r_\epsilon(t) = 2\alpha \Psi'_{P_{\text{out}}}(\mathbb{E}\langle Q \rangle_t)$, $q_\epsilon(t) = \mathbb{E}\langle Q \rangle_t$. Again, applying the Cauchy-Lipschitz theorem and the Liouville formula we show in the SI that this system admits a unique solution and the map $R^t : (\epsilon_1, \epsilon_2) \mapsto (R_1(t, \epsilon), R_2(t, \epsilon))$ is a \mathcal{C}^1 diffeomorphism with determinant greater or equal to 1. So with this choice of interpolating functions (18) is valid and we have (19). We show in the SI (Proposition 18) that $\Psi_{P_{\text{out}}}(q; \rho)$ is convex in q and thus $g : q \in [0, \rho] \mapsto 2\alpha \Psi_{P_{\text{out}}}(q; \rho) - r_\epsilon(t)q$ is convex too. Since by the differential equations $r_\epsilon(t) = 2\alpha \Psi'_{P_{\text{out}}}(q_\epsilon(t))$, the function g must attain its minimum at $q = q_\epsilon(t)$. By Proposition 17 in the SI $\psi_{P_0}(r)$ is convex, thus from Jensen and the last remark, the integrand $\{\cdot\cdot\}$ in (19) is bounded as

$$\begin{aligned} & \psi_{P_0}(\int_0^1 r_\epsilon(t) dt) + \alpha \Psi_{P_{\text{out}}}(\int_0^1 q_\epsilon(t) dt; \rho) - \frac{1}{2} \int_0^1 r_\epsilon(t) q_\epsilon(t) dt \\ & \leq \int_0^1 dt \left\{ \psi_{P_0}(r_\epsilon(t)) + \alpha \Psi_{P_{\text{out}}}(q_\epsilon(t); \rho) - \frac{1}{2} r_\epsilon(t) q_\epsilon(t) \right\} \\ & = \int_0^1 dt \left\{ \psi_{P_0}(r_\epsilon(t)) + \inf_{q \in [0, \rho]} (\alpha \Psi_{P_{\text{out}}}(q; \rho) - \frac{1}{2} r_\epsilon(t) q) \right\} \\ & \leq \sup_{r \geq 0} \inf_{q \in [0, \rho]} \left\{ \psi_{P_0}(r) + \alpha \Psi_{P_{\text{out}}}(q; \rho) - \frac{1}{2} r q \right\} \end{aligned}$$

which implies $\limsup_{n \rightarrow \infty} f_n \leq \sup_{r \geq 0} \inf_{q \in [0, \rho]} f_{\text{RS}}(q, r)$.

Acknowledgments

J.B. acknowledges funding from the SNSF (grant 200021-156672). F.K. and L.Z. acknowledge funding from the ERC under the European Union's 7th Framework Programme Grant Agreement 307087-SPARCS and under Horizon 2020 Research and Innovation Programme Grant Agreement 714608-SMiLe. F.K. and N.M. acknowledge support from the ANR-PAIL. Part of this work was done while L.M. was visiting EPFL.

References

- [1] J. Nelder and R. Wedderburn. Generalized linear models. *Journal of the Royal Statistical Society. Series A (General)*, 135(3):370–384, 1972.
- [2] P. McCullagh. Generalized linear models. *European Journal of Operational Research*, 16(3):285–292, 1984.
- [3] J. R. Fienup. Phase retrieval algorithms: a comparison. *Applied optics*, 21(15), 1982.

- [4] L. Demanet and P. Hand. Stable optimizationless recovery from phaseless linear measurements. *Journal of Fourier Analysis and Applications*, 20(1):199–221, 2014.
- [5] E. J. Candes, T. Strohmer, and V. Voroninski. Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming. *Communications on Pure and Applied Mathematics*, 66(8):1241–1274, 2013.
- [6] P. T. Boufounos and R. G. Baraniuk. 1-bit compressive sensing. In *42nd Annual Conference on Information Sciences and Systems (CISS)*, pages 16–21. IEEE, 2008.
- [7] P. Bühlmann and S. Van De Geer. *Statistics for high-dimensional data: methods, theory and applications*. Springer Science & Business Media, 2011.
- [8] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- [9] D. L. Donoho and J. Tanner. Sparse nonnegative solution of underdetermined linear equations by linear programming. *Proceedings of the National Academy of Sciences of the United States of America*, 102(27):9446–9451, 2005.
- [10] E. J. Candes and T. Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *IEEE Transactions on Information Theory*, 52(12):5406–5425, Dec 2006.
- [11] D. L. Donoho, A. Maleki, and A. Montanari. Message-passing algorithms for compressed sensing. *Proceedings of the National Academy of Sciences*, 106(45):18914–18919, Nov 2009.
- [12] S. Rangan. Generalized approximate message passing for estimation with random linear mixing. In *IEEE International Symposium on Information Theory Proceedings (ISIT)*, pages 2168–2172, July 2011.
- [13] L. Zdeborová and F. Krzakala. Statistical physics of inference: thresholds and algorithms. *Advances in Physics*, 65(5):453–552, 2016.
- [14] U. Kamilov, V. K. Goyal, and S. Rangan. Optimal quantization for compressive sensing under message passing reconstruction. In *IEEE International Symposium on Information Theory Proceedings (ISIT)*, pages 459–463. IEEE, 2011.
- [15] Y. Xu, Y. Kabashima, and L. Zdeborová. Bayesian signal reconstruction for 1-bit compressed sensing. *Journal of Statistical Mechanics: Theory and Experiment*, 2014(11):P11015, 2014.
- [16] P. Schniter and S. Rangan. Compressive phase retrieval via generalized approximate message passing. *IEEE Transactions on Signal Processing*, 63(4):1043–1055, 2015.
- [17] M. Bayati and A. Montanari. The lasso risk for gaussian matrices. *IEEE Transactions on Information Theory*, 58(4):1997–2017, April 2012.
- [18] N. El Karoui, D. Bean, P. J. Bickel, C. Lim, and B. Yu. On robust regression with high-dimensional predictors. *Proceedings of the National Academy of Sciences*, 110(36):14557–14562, 2013.
- [19] D. Donoho and A. Montanari. High dimensional robust m-estimation: asymptotic variance via approximate message passing. *Probability Theory and Related Fields*, 166(3):935–969, Dec 2016.
- [20] R. Gribonval and P. Machart. Reconciling "priors" & "priors" without prejudice? In *Advances in Neural Information Processing Systems*, pages 2193–2201, 2013.
- [21] M. Advani and S. Ganguli. An equivalence between high dimensional bayes optimal inference and m-estimation. In *Advances in Neural Information Processing Systems*, 2016.
- [22] E. Gardner and B. Derrida. Three unfinished works on the optimal storage capacity of networks. *Journal of Physics A: Mathematical and General*, 22(12):1983, 1989.

- [23] H. S. Seung, H. Sompolinsky, and N. Tishby. Statistical mechanics of learning from examples. *Phys. Rev. A*, 45:6056–6091, Apr 1992.
- [24] T. L. H. Watkin, A. Rau, and M. Biehl. The statistical mechanics of learning a rule. *Rev. Mod. Phys.*, 65:499–556, Apr 1993.
- [25] A. Engel and C. Van den Broeck. *Statistical mechanics of learning*. Cambridge University Press, 2001.
- [26] A. Engel and L. Reimers. Reliability of replica symmetry for the generalization problem in a toy multilayer neural network. *Europhys. Lett.*, 28:531, 1994.
- [27] G. J. Bex, R. Serneels, and C. V. den Broeck. Storage capacity and generalization error for the reversed-wedge ising perceptron. *Phys. Rev. E*, 51:6309, 1995.
- [28] T. Hosaka, Y. Kabashima, and H. Nishimori. Statistical mechanics of lossy data compression using a nonmonotonic perceptron. *Physical Review E*, 66(6):066126, 2002.
- [29] C. Baldassi, C. Borgs, J. T. Chayes, A. Ingrosso, C. Lucibello, L. Saglietti, and R. Zecchina. Unreasonable effectiveness of learning neural networks: From accessible states and robust ensembles to basic algorithmic schemes. *Proceedings of the National Academy of Sciences*, 113(48):E7655–E7662, 2016.
- [30] C. H. Martin and M. W. Mahoney. Rethinking generalization requires revisiting old ideas: statistical mechanics approaches and complex learning behavior. 2017.
- [31] N. Tishby, F. C. Pereira, and W. Bialek. The information bottleneck method. In *Proc. of the 37-th Annual Allerton Conference on Communication, Control and Computing*, pages 368–377, 1999.
- [32] R. Shwartz-Ziv and N. Tishby. Opening the black box of deep neural networks via information. *CoRR*, abs/1703.00810, 2017.
- [33] C. E. Shannon. A mathematical theory of communication, part i, part ii. *Bell Syst. Tech. J.*, 27:623–656, 1948.
- [34] T. Tanaka. A statistical-mechanics approach to large-system analysis of cdma multiuser detectors. *IEEE Transactions on Information Theory*, 48(11):2888–2910, Nov 2002.
- [35] D. Guo and S. Verdú. Randomly spread cdma: Asymptotics via statistical physics. *IEEE Transactions on Information Theory*, 51(6):1983–2010, June 2005.
- [36] A. R. Barron and A. Joseph. Toward fast reliable communication at rates near capacity with gaussian noise. In *IEEE International Symposium on Information Theory (ISIT)*, pages 315–319, June 2010.
- [37] J. Barbier and F. Krzakala. Approximate message-passing decoder and capacity-achieving sparse superposition codes. *IEEE Transactions on Information Theory*, 2017.
- [38] C. Rush, A. Greig, and R. Venkataramanan. Capacity-achieving sparse superposition codes via approximate message passing decoding. *IEEE Transactions on Information Theory*, 63(3):1476–1500, 2017.
- [39] J. Barbier, M. Dia, and N. Macris. Threshold saturation of spatially coupled sparse superposition codes for all memoryless channels. In *IEEE Information Theory Workshop (ITW)*, Sept 2016.
- [40] J. Barbier, M. Dia, and N. Macris. Universal sparse superposition codes with spatial coupling and GAMP decoding. *CoRR*, abs/1707.04203, 2017.
- [41] M. Mézard. The space of interactions in neural networks: Gardner’s computation with the cavity method. *Journal of Physics A: Mathematical and General*, 22(12):2181–2190, 1989.
- [42] E. Bolthausen. An iterative construction of solutions of the tap equations for the sherrington–kirkpatrick model. *Communications in Mathematical Physics*, 325(1):333–366, 2014.

- [43] M. Bayati and A. Montanari. The dynamics of message passing on dense graphs, with applications to compressed sensing. *IEEE Transactions on Information Theory*, 57(2):764–785, Feb 2011.
- [44] M. Bayati, M. Lelarge, and A. Montanari. Universality in polytope phase transitions and message passing algorithms. *The Annals of Applied Probability*, 25(2):753–822, 2015.
- [45] A. Javanmard and A. Montanari. State evolution for general approximate message passing algorithms, with applications to spatial coupling. *Information and Inference: A Journal of the IMA*, 2(2):115–144, 2013.
- [46] J. Barbier, M. Dia, N. Macris, and F. Krzakala. The mutual information in random linear estimation. In *54th Annual Allerton Conference on Communication, Control, and Computing*, pages 625–632. IEEE, 2016.
- [47] J. Barbier, N. Macris, M. Dia, and F. Krzakala. Mutual information and optimality of approximate message-passing in random linear estimation. *arXiv:1701.05823*, 2017.
- [48] G. Reeves and H. D. Pfister. The replica-symmetric prediction for compressed sensing with gaussian matrices is exact. In *IEEE International Symposium on Information Theory (ISIT)*, pages 665–669, 2016.
- [49] M. Mézard, G. Parisi, and M.-A. Virasoro. *Spin glass theory and beyond*. World Scientific Publishing, 1987.
- [50] G. Györgyi. First-order transition to perfect generalization in a neural network with binary synapses. *Physical Review A*, 41(12):7097, 1990.
- [51] H. Sompolinsky, N. Tishby, and H. S. Seung. Learning from examples in large neural networks. *Physical Review Letters*, 65(13):1683, 1990.
- [52] J. Barbier and N. Macris. The adaptive interpolation method: a simple scheme to prove replica formulas in bayesian inference. *Probability Theory and Related Fields*, Oct 2018.
- [53] A. Coja-Oghlan, F. Krzakala, W. Perkins, and L. Zdeborova. Information-theoretic thresholds from the cavity method. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 146–157, 2017.
- [54] D. Guo, S. Shamai, and S. Verdú. Mutual information and minimum mean-square error in gaussian channels. *IEEE Transactions on Information Theory*, 51(4):1261–1282, April 2005.
- [55] M. Opper and D. Haussler. Generalization performance of bayes optimal classification algorithm for learning a perceptron. *Physical Review Letters*, 66(20):2677, 1991.
- [56] D. J. Thouless, P. W. Anderson, and R. G. Palmer. Solution of ‘solvable model of a spin glass’. *Philosophical Magazine*, 35(3):593–601, 1977.
- [57] Y. Kabashima. Inference from correlated patterns: a unified theory for perceptron learning and linear vector channels. *Journal of Physics: Conference Series*, 95(1):012001, 2008.
- [58] D. L. Donoho, A. Javanmard, and A. Montanari. Information-theoretically optimal compressed sensing via spatial coupling and approximate message passing. *IEEE Transactions on Information Theory*, 59(11):7434–7464, Nov 2013.
- [59] M. Opper and O. Winther. Mean field approach to bayes learning in feed-forward neural networks. *Physical review letters*, 76(11):1964, 1996.
- [60] M. Opper and O. Winther. Tractable approximations for probabilistic models: The adaptive thouless-anderson-palmer mean field approach. *Physical Review Letters*, 86(17):3695, 2001.

- [61] M. Mondelli and A. Montanari. Fundamental limits of weak recovery with applications to phase retrieval. *arXiv preprint arXiv:1708.05932*, 2017.
- [62] Github repository of numerical experiments for the current paper. <https://github.com/sphinxteam/GeneralizedLinearModel2017>, 2017.
- [63] A. K. Fletcher and S. Rangan. Iterative reconstruction of rank-one matrices in noise. *Information and Inference: A Journal of the IMA*, 7(3):531–562, 2018.
- [64] D. Hansel, G. Mato, and C. Meunier. Memorization without generalization in a multilayered neural network. *EPL (Europhysics Letters)*, 20(5):471, 1992.
- [65] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [66] F. Chollet. keras. <https://github.com/fchollet/keras>, 2015.
- [67] Y. Wu and S. Verdú. Rényi information dimension: Fundamental limits of almost lossless analog compression. *IEEE Transactions on Information Theory*, 56(8):3721–3748, 2010.
- [68] F. Krzakala, M. Mézard, F. Sausset, Y. Sun, and L. Zdeborová. Statistical-physics-based reconstruction in compressed sensing. *Phys. Rev. X*, 2:021005(18), May 2012.
- [69] A. Maleki, L. Anitori, Z. Yang, and R. G. Baraniuk. Asymptotic analysis of complex lasso via complex approximate message passing (camp). *IEEE Transactions on Information Theory*, 59(7):4290–4308, 2013.
- [70] M. Soltanolkotabi. Structured signal recovery from quadratic measurements: Breaking sample complexity barriers via nonconvex optimization. *arXiv preprint arXiv:1702.06175*, 2017.
- [71] F. Rosenblatt. *The perceptron, a perceiving and recognizing automaton Project Para*. Cornell Aeronautical Laboratory, 1957.
- [72] F. Guerra and F. L. Toninelli. The thermodynamic limit in mean field spin glass models. *Communications in Mathematical Physics*, 230(1):71–79, 2002.
- [73] M. Talagrand. *Mean field models for spin glasses: Volume I: Basic examples*, volume 54. Springer Science & Business Media, 2010.
- [74] M. Mezard and A. Montanari. *Information, physics, and computation*. Oxford University Press, 2009.
- [75] C. Baldassi, A. Braunstein, N. Brunel, and R. Zecchina. Efficient supervised learning in networks with binary synapses. *Proceedings of the National Academy of Sciences*, 104(26):11079–11084, 2007.
- [76] T. Richardson and R. Urbanke. *Modern coding theory*. Cambridge university press, 2008.
- [77] M. J. Wainwright, M. I. Jordan, et al. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1–2):1–305, 2008.
- [78] J. P. Vila and P. Schniter. Expectation-maximization gaussian-mixture approximate message passing. *IEEE Transactions on Signal Processing*, 61(19):4658–4672, 2013.
- [79] J. Ziniel, P. Schniter, and P. Sederberg. Binary linear classification and feature selection via generalized approximate message passing. In *Information Sciences and Systems (CISS), 2014 48th Annual Conference on*, pages 1–6. IEEE, 2014.
- [80] F. Krzakala, M. Mézard, F. Sausset, Y. Sun, and L. Zdeborová. Probabilistic reconstruction in compressed sensing: algorithms, phase diagrams, and threshold achieving matrices. *Journal of Statistical Mechanics: Theory and Experiment*, 2012(08):P08009(57), 2012.

- [81] D. Donoho and J. Tanner. Observed universality of phase transitions in high-dimensional geometry, with implications for modern data analysis and signal processing. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 367(1906):4273–4293, 2009.
- [82] S. Oymak, A. Jalali, M. Fazel, Y. C. Eldar, and B. Hassibi. Simultaneously structured models with application to sparse and low-rank matrices. *IEEE Transactions on Information Theory*, 61(5):2886–2908, 2015.
- [83] E. B. Baum and Y.-D. Lyuu. The transition to perfect generalization in perceptrons. *Neural computation*, 3(3):386–401, 1991.
- [84] S. Diamond and S. Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 2016. To appear.
- [85] T. Goldstein and C. Studer. Phasemax: Convex phase retrieval via basis pursuit. *arXiv preprint arXiv:1610.07531*, 2016.
- [86] J. Barbier, N. Macris, and L. Miolane. The Layered Structure of Tensor Estimation and its Mutual Information. In *47th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, September 2017.
- [87] A. Montanari. Estimating random variables from random sparse observations. *European Transactions on Telecommunications*, 19(4):385–403, 2008.
- [88] S. B. Korada and N. Macris. Exact solution of the gauge symmetric p-spin glass model on a complete graph. *Journal of Statistical Physics*, 136(2):205–230, 2009.
- [89] Y. Wu and S. Verdú. Optimal phase transitions in compressed sensing. *IEEE Transactions on Information Theory*, 58(10):6241–6263, Oct 2012.
- [90] N. Macris. Griffith-kelly-sherman correlation inequalities: A useful tool in the theory of error correcting codes. *IEEE Transactions on Information Theory*, 53(2):664–683, Feb 2007.
- [91] S. B. Korada and N. Macris. Tight bounds on the capacity of binary input random cdma systems. *IEEE Transactions on Information Theory*, 56(11):5590–5613, Nov 2010.
- [92] P. Hartman. *Ordinary Differential Equations*. Society for Industrial and Applied Mathematics, 2002, 1964.
- [93] P. Milgrom and I. Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601, 2002.
- [94] D. Guo, Y. Wu, S. S. Shitz, and S. Verdú. Estimation in gaussian noise: Properties of the minimum mean-square error. *IEEE Transactions on Information Theory*, 57(4):2371–2385, 2011.
- [95] S. B. Korada and A. Montanari. Applications of the Lindeberg principle in communications and statistical learning. *IEEE Transactions on Information Theory*, 57(4):2440–2450, 2011.
- [96] Y. Wu and S. Verdú. Functional properties of minimum mean-square error and mutual information. *IEEE Transactions on Information Theory*, 58(3):1289–1301, 2012.
- [97] L. Miolane. Phase transitions in spiked matrix estimation: information-theoretic analysis. *arXiv preprint arXiv:1806.04343*, 2018.
- [98] V. I. Bogachev. *Measure theory*, volume 1. Springer Science & Business Media, 2007.
- [99] M. Lelarge and L. Miolane. Fundamental limits of symmetric low-rank matrix estimation. *ArXiv e-prints*, November 2016.

- [100] S. Chatterjee et al. A generalization of the lindeberg principle. *The Annals of Probability*, 34(6):2061–2076, 2006.
- [101] R. T. Rockafellar. *Convex analysis*. Princeton university press, 2015.
- [102] S. Boucheron, G. Lugosi, and O. Bousquet. Concentration inequalities. In *Advanced Lectures on Machine Learning*, pages 208–240. Springer, 2004.
- [103] J. Bezanon, S. Karpinski, V. Shah, and A. Edelman. Julia: A fast dynamic language for technical computing. In *Lang.NEXT*, April 2012.
- [104] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.

Part II

Supplementary informations

1 Setting

1.1 Generalized linear estimation: Problem statement

We give a formal description of the observation model to which our results apply. The generalized linear *model* covers both the *estimation* (or inference) problem and the supervised *learning* problems (see Sec. 3).

Let $n, m \in \mathbb{N}^*$. Let P_0 be a probability distribution over \mathbb{R} and let $(X_i^*)_{i=1}^n \stackrel{\text{iid}}{\sim} P_0$ be the components of a signal vector \mathbf{X}^* (this is also denoted $\mathbf{X}^* \stackrel{\text{iid}}{\sim} P_0$). We fix a function $\varphi : \mathbb{R} \times \mathbb{R}^{k_A} \rightarrow \mathbb{R}$ and consider $(\mathbf{A}_\mu)_{\mu=1}^m \stackrel{\text{iid}}{\sim} P_A$, where P_A is a probability distribution over \mathbb{R}^{k_A} , $k_A \in \mathbb{N}$. We acquire m measurements through

$$Y_\mu = \varphi\left(\frac{1}{\sqrt{n}}[\Phi\mathbf{X}^*]_\mu, \mathbf{A}_\mu\right) + \sqrt{\Delta}Z_\mu, \quad 1 \leq \mu \leq m, \quad (20)$$

where $(Z_\mu)_{\mu=1}^m \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$ is an additive Gaussian noise, $\Delta \geq 0$, and Φ is a $m \times n$ measurement matrix with independent entries that have zero mean and unit variance. The estimation problem is to recover \mathbf{X}^* from the knowledge of $\mathbf{Y} = (Y_\mu)_{\mu=1}^m$, φ , Φ , Δ , P_0 and P_A (the realization of the random stream \mathbf{A} itself, if present in the model, is unknown). We use the notation $[\Phi\mathbf{X}^*]_\mu = \sum_{i=1}^n \Phi_{\mu i} X_i^*$. When $\varphi(x, \mathbf{a}) = \varphi(x) = x$ we have a random linear estimation problem, whereas if, say, $\varphi(x) = \text{sgn}(x)$ we have a noisy single layer perceptron. Sec. 3 discusses various examples related to non-linear estimation and supervised learning.

Denote the prior over the signal as $dP_0(\mathbf{x}) = \prod_{i=1}^n dP_0(x_i)$, and similarly $dP_A(\mathbf{a}) = \prod_{\mu=1}^m dP_A(\mathbf{a}_\mu)$. It is also fruitful to think of the measurements as the outputs of a “channel”,

$$Y_\mu \sim P_{\text{out}}\left(\cdot \mid \frac{1}{\sqrt{n}}[\Phi\mathbf{X}^*]_\mu\right). \quad (21)$$

When $\Delta > 0$ the transition kernel P_{out} admits a transition density with respect to (w.r.t.) Lebesgue’s measure, given by

$$P_{\text{out}}(y|x) = \frac{1}{\sqrt{2\pi\Delta}} \int dP_A(\mathbf{a}) e^{-\frac{1}{2\Delta}(y-\varphi(x,\mathbf{a}))^2}. \quad (22)$$

When $\Delta = 0$, we will only consider discrete channels where φ takes values in \mathbb{N}^2 . In that case P_{out} admits a transition density with respect the counting measure on \mathbb{N} given by (here $\mathbf{1}(\cdot)$ is the indicator function)

$$P_{\text{out}}(y|x) = \int dP_A(\mathbf{a}) \mathbf{1}(y = \varphi(x, \mathbf{a})). \quad (23)$$

Note that for deterministic models, \mathbf{A} in (20) is absent and thus the associated $\int dP_A(\mathbf{a})$ integral in (22)-(23) simply disappears. In fact (20) is sometimes called a “random function representation” of a transition kernel P_{out} . Our analysis uses both representations (20) and (21).

Throughout this paper we often adopt the language of statistical mechanics. In particular the random variables \mathbf{Y} (and also Φ , \mathbf{X}^* , \mathbf{A} , \mathbf{Z}) are called *quenched* variables because once the measurements are acquired they have a “fixed realization”. An expectation taken w.r.t. *all* quenched random variables appearing in an

²Notice that this allows to study any channel whose outputs belong to a countable set S by applying a injection $u : S \rightarrow \mathbb{N}$ to the outputs.

expression will simply be denoted by \mathbb{E} *without* subscript. Subscripts are only used when the expectation carries over a subset of random variables appearing in an expression or when some confusion could arise.

A fundamental role is played by the posterior distribution of (the signal) \mathbf{x} given the quenched measurements \mathbf{Y} (recall that \mathbf{X}^* , \mathbf{A} and \mathbf{Z} are unknown). According to the Bayes formula this posterior is

$$dP(\mathbf{x}|\mathbf{Y}, \Phi) = \frac{1}{\mathcal{Z}(\mathbf{Y}, \Phi)} dP_0(\mathbf{x}) \prod_{\mu=1}^m P_{\text{out}}\left(Y_\mu \middle| \frac{1}{\sqrt{n}}[\Phi \mathbf{x}]_\mu\right) \quad (24)$$

$$= \frac{1}{\mathcal{Z}(\mathbf{Y}, \Phi)} dP_0(\mathbf{x}) e^{-\mathcal{H}(\mathbf{x}; \mathbf{Y}, \Phi)} \quad (25)$$

where the *Hamiltonian* is defined as

$$\mathcal{H}(\mathbf{x}; \mathbf{Y}, \Phi) := - \sum_{\mu=1}^m \ln P_{\text{out}}\left(Y_\mu \middle| \frac{1}{\sqrt{n}}[\Phi \mathbf{x}]_\mu\right) \quad (26)$$

and the *partition function* (the normalization factor) is defined as

$$\mathcal{Z}(\mathbf{Y}, \Phi) := \int dP_0(\mathbf{x}) e^{-\mathcal{H}(\mathbf{x}; \mathbf{Y}, \Phi)}. \quad (27)$$

From the point of view of statistical mechanics (25) is a Gibbs distribution and the integration over $dP_0(\mathbf{x})$ in the partition function is best thought as a “sum over annealed or fluctuating degrees of freedom”. Let us introduce a standard statistical mechanics notation for the expectation w.r.t. the posterior (24), the so called *Gibbs bracket* $\langle - \rangle$ defined as

$$\langle g(\mathbf{x}) \rangle := \mathbb{E}[g(\mathbf{X})|\mathbf{Y}, \Phi] = \int dP(\mathbf{x}|\mathbf{Y}, \Phi) g(\mathbf{x}) \quad (28)$$

for any continuous bounded function g . The main quantity of interest here is the associated averaged *free entropy* (or minus the averaged *free energy*)

$$f_n := \frac{1}{n} \mathbb{E} \ln \mathcal{Z}(\mathbf{Y}, \Phi). \quad (29)$$

It is perhaps useful to stress that $\mathcal{Z}(\mathbf{Y}, \Phi)$ is nothing else than the density of \mathbf{Y} conditioned on Φ so we have the explicit representation (used later on)

$$\begin{aligned} f_n &= \frac{1}{n} \mathbb{E}_\Phi \int d\mathbf{Y} \mathcal{Z}(\mathbf{Y}, \Phi) \ln \mathcal{Z}(\mathbf{Y}, \Phi) \\ &= \frac{1}{n} \mathbb{E}_\Phi \int d\mathbf{Y} dP_0(\mathbf{X}^*) e^{-\mathcal{H}(\mathbf{X}^*; \mathbf{Y}, \Phi)} \ln \int dP_0(\mathbf{x}) e^{-\mathcal{H}(\mathbf{x}; \mathbf{Y}, \Phi)}, \end{aligned} \quad (30)$$

where $d\mathbf{Y} = \prod_{\mu=1}^m dY_\mu$. Thus f_n is minus the conditional entropy $-H(\mathbf{Y}|\Phi)/n$ of the measurements. One of the main contributions of this paper is the derivation, thanks to the adaptive interpolation method, of the thermodynamic limit $\lim_{n \rightarrow \infty} f_n$ in the “high-dimensional regime”, namely when $n, m \rightarrow \infty$ while $m/n \rightarrow \alpha > 0$ (α is sometimes referred to as the “measurement rate” or “sampling rate”).

1.2 The teacher-student scenario

We now describe an important conceptual setting, the *teacher-student scenario* (also called planted model), that allows to then define the optimal generalization error. We voluntarily employ terms coming from machine learning instead of the signal processing terminology used until here.

First the teacher randomly generates a *classifier* $\mathbf{X}^* \in \mathbb{R}^n$ (the signal in the estimation problem) with $\mathbf{X}^* \stackrel{\text{iid}}{\sim} P_0$ and an ensemble of m *patterns* (row-vectors) $\Phi_\mu \in \mathbb{R}^n$ for $\mu = 1, \dots, m$ such that $\Phi_\mu \stackrel{\text{iid}}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$. The teacher then chooses a model (φ, P_A, Δ) or equivalently P_{out} , which are linked through (22)-(23). The teacher then output *labels* $Y_\mu \in \mathbb{R}$ through (20) or (21) for $\mu = 1, \dots, m$.

The student is given the distribution P_0 , the model (φ, P_A, Δ) or equivalently P_{out} and the training data composed of the pattern-label pairs $\{(Y_\mu; \Phi_\mu)\}_{\mu=1}^m$ generated by the teacher. His (supervised) learning task is then to predict the labels associated with new, yet unseen, patterns from all this knowledge.

How does the teacher may evaluate the student's prediction capabilities? The teacher starts by randomly generating a new line of the matrix, or pattern, Φ_{new} . Then, still using the same \mathbf{X}^* , he generates the associated new label $Y_{\text{new}} \sim P_{\text{out}}(\cdot | \Phi_{\text{new}} \cdot \mathbf{X}^* / \sqrt{n})$. He is now ready to evaluate the student generalization performance. For that purpose, an important quantity is the *generalization error* (or prediction error). If we denote $\hat{Y}(\Phi_{\text{new}}, \Phi, \mathbf{Y})$ the estimator used by the student (which is thus a measurable function of the observations), the generalization error is defined as

$$\mathcal{E}_{\text{gen}}(\hat{Y}) := \mathbb{E}[(Y_{\text{new}} - \hat{Y}(\Phi_{\text{new}}, \Phi, \mathbf{Y}))^2]. \quad (31)$$

The optimal generalization error is then defined as the minimum of \mathcal{E}_{gen} over all estimators $\hat{Y}(\Phi_{\text{new}}, \Phi, \mathbf{Y})$:

$$\mathcal{E}_{\text{gen}}^{\text{opt}} := \min_{\hat{Y}} \mathcal{E}_{\text{gen}}(\hat{Y}) = \text{MMSE}(Y_{\text{new}} | \Phi_{\text{new}}, \Phi, \mathbf{Y}) = \mathbb{E}[(Y_{\text{new}} - \mathbb{E}[Y_{\text{new}} | \Phi_{\text{new}}, \Phi, \mathbf{Y}])^2]. \quad (32)$$

Here, and for the rest of the paper, we define the minimum mean-square error (MMSE) function as follows: Given two random variables \mathbf{A}, \mathbf{B} , the MMSE in estimating \mathbf{A} given \mathbf{B} is defined as

$$\text{MMSE}(\mathbf{A} | \mathbf{B}) := \mathbb{E}[\|\mathbf{A} - \mathbb{E}[\mathbf{A} | \mathbf{B}]\|^2], \quad (33)$$

where $\mathbb{E}[\mathbf{A} | \mathbf{B}]$ is the expectation of \mathbf{A} with respect to its posterior given \mathbf{B} .

A word about notations: Let us emphasize on the link between the different notations that we use in the present supplementary material and in the main text. E.g., the expectation w.r.t. to the posterior of Y_{new} appearing in (32) can be written equivalently as:

$$\mathbb{E}[Y_{\text{new}} | \Phi_{\text{new}}, \Phi, \mathbf{Y}] = \mathbb{E}_{P_A(\mathbf{a})} \mathbb{E}_{P(\mathbf{x} | \Phi, \mathbf{Y})} \left(\frac{\Phi_{\text{new}} \cdot \mathbf{x}}{\sqrt{n}}, \mathbf{a} \right) = \mathbb{E}_{P_A(\mathbf{a})} \left\langle \varphi \left(\frac{\Phi_{\text{new}} \cdot \mathbf{x}}{\sqrt{n}}, \mathbf{a} \right) \right\rangle. \quad (34)$$

To see that, just write:

$$\begin{aligned} \mathbb{E}[Y_{\text{new}} | \Phi_{\text{new}}, \Phi, \mathbf{Y}] &:= \int dY_{\text{new}} Y_{\text{new}} P(Y_{\text{new}} | \Phi_{\text{new}}, \Phi, \mathbf{Y}) \\ &= \int dY_{\text{new}} Y_{\text{new}} P_{\text{out}} \left(Y_{\text{new}} \left| \frac{1}{\sqrt{n}} \Phi_{\text{new}} \cdot \mathbf{x} \right. \right) dP(\mathbf{x} | \mathbf{Y}, \Phi) \\ &= \left\langle \int dY_{\text{new}} Y_{\text{new}} P_{\text{out}} \left(Y_{\text{new}} \left| \frac{1}{\sqrt{n}} \Phi_{\text{new}} \cdot \mathbf{x} \right. \right) \right\rangle \\ &= \left\langle \int dP_A(\mathbf{a}) dY_{\text{new}} Y_{\text{new}} \frac{1}{\sqrt{2\pi\Delta}} e^{-\frac{1}{2\Delta} \left\{ Y_{\text{new}} - \varphi \left(\frac{1}{\sqrt{n}} \Phi_{\text{new}} \cdot \mathbf{x}, \mathbf{a} \right) \right\}^2} \right\rangle \\ &= \mathbb{E}_{P_A(\mathbf{a})} \left\langle \varphi \left(\frac{1}{\sqrt{n}} \Phi_{\text{new}} \cdot \mathbf{x}, \mathbf{a} \right) \right\rangle. \end{aligned} \quad (35)$$

Here we used definition (22) for the transition kernel, but using instead (23) would lead to the same identity.

1.3 Two scalar inference channels

An important role in our proof of the asymptotic expression of the free entropy is played by simple *scalar* inference channels. As we will see, the free entropy is expressed in terms of the free entropy of these channels. This “decoupling property” results from the mean-field approach in statistical physics, used through in the replica method to perform a formal calculation of the free entropy of the model [49, 74]. Let us now introduce these two scalar denoising models.

The first one is an additive Gaussian channel. Let $r \geq 0$, which plays the role of a signal-to-noise ratio (snr). Suppose that $X_0 \sim P_0$ and that we observe

$$Y_0 = \sqrt{r} X_0 + Z_0, \quad (36)$$

where $Z_0 \sim \mathcal{N}(0, 1)$ independently of X_0 . Consider the inference problem consisting of retrieving X_0 from the observations Y_0 . The associated posterior distribution is

$$dP(x|Y_0) = \frac{dP_0(x)e^{\sqrt{r}Y_0x - rx^2/2}}{\int dP_0(x)e^{\sqrt{r}Y_0x - rx^2/2}}. \quad (37)$$

In this expression all the x -independent terms have been simplified between the numerator and the normalization. The free entropy associated with this channel is just the expectation of the logarithm of the normalization factor

$$\psi_{P_0}(r) := \mathbb{E} \ln \int dP_0(x)e^{\sqrt{r}Y_0x - rx^2/2}. \quad (38)$$

The basic properties of ψ_{P_0} are presented in Appendix B.1.

The second scalar channel that appears naturally in the problem is linked to the transition kernel P_{out} through the following inference model. Suppose that $V, W^* \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$ where V is *known* while the inference problem is to recover the unknown W^* from the following observation

$$\tilde{Y}_0 \sim P_{\text{out}}(\cdot | \sqrt{q}V + \sqrt{\rho - q}W^*), \quad (39)$$

where $\rho > 0, q \in [0, \rho]$. Notice that the channel (39) is equivalent to $\tilde{Y}_0 = \varphi(\sqrt{q}V + \sqrt{\rho - q}W^*, \mathbf{A}) + \sqrt{\Delta}Z$ with $\Delta \geq 0$ and where $(\mathbf{A}, Z) \sim P_A \otimes \mathcal{N}(0, 1)$, independently of V, W^* . The free entropy for this model, again related to the normalization of the posterior $dP(w|\tilde{Y}_0, V)$, is

$$\Psi_{P_{\text{out}}}(q; \rho) = \Psi_{P_{\text{out}}}(q) := \mathbb{E} \ln \int \mathcal{D}w P_{\text{out}}(\tilde{Y}_0 | \sqrt{q}V + \sqrt{\rho - q}w), \quad (40)$$

where $\mathcal{D}w := dw(2\pi)^{-1/2}e^{-w^2/2}$ is the standard Gaussian measure. In (40) above, P_{out} denotes either the transition density with respect to Lebesgue’s measure (given by (22)) in the case $\Delta > 0$, or the density with respect to the counting measure over \mathbb{N} (given by (23)), in the case of a “discrete” channel (φ takes values in \mathbb{N} and $\Delta = 0$). We prove in Appendix B.2 that this function is convex, differentiable and non-decreasing w.r.t. its first argument.

2 Main results

2.1 Replica-symmetric formula and mutual information

Let us now introduce our first main result, a single-letter *replica-symmetric formula* for the asymptotic free entropy of model (20), (21). The result holds under the following rather general hypotheses. We will consider two cases, that is when there is some Gaussian noise ($\Delta > 0$, see (h5.a) below) and the case without Gaussian noise ($\Delta = 0$, see (h5.b) below):

- (h1) The prior distribution P_0 admits a finite third moment and has at least two points in its support.
- (h2) There exists $\gamma > 0$ such that the sequence $(\mathbb{E}[|\varphi(\frac{1}{\sqrt{n}}[\Phi\mathbf{X}^*]_1, \mathbf{A}_1)|^{2+\gamma}])_{n \geq 1}$ is bounded.
- (h3) The random variables (Φ_{μ_i}) are independent with zero mean, unit variance and finite third moment that is bounded with n .
- (h4) For almost-all values of $\mathbf{a} \in \mathbb{R}^{k_A}$ (w.r.t. P_A), the function $x \mapsto \varphi(x, \mathbf{a})$ is continuous almost everywhere.

We will also assume that one of the two following hypotheses hold:

- (h5.a) $\Delta > 0$.
- (h5.b) $\Delta = 0$ and φ takes values in \mathbb{N} .

Remark 1. *The above hypotheses are here stated using the “random function” representation of (20). In many cases, it can be useful to state them using the “transition kernel” representation of (21). The hypotheses (h2) and (h4) are respectively equivalent³ to:*

- (h2') *There exists $\gamma > 0$ such that $\mathbb{E}[|Y_1|^{2+\gamma}]$ remains bounded with n .*
- (h4') *$x \in \mathbb{R} \mapsto P_{\text{out}}(\cdot|x)$ is continuous almost everywhere for the weak convergence.*

Under the above hypothesis (h5.a) (respectively (h5.b)), the transition kernel P_{out} admits a density with respect to Lebesgue’s measure on \mathbb{R} (resp. the counting measure on \mathbb{N}) that will be denoted by $P_{\text{out}}(\cdot|x)$. We will call the kernel P_{out} *informative* if there exists $y \in \mathbb{R}$ (resp. $y \in \mathbb{N}$) such that $P_{\text{out}}(y|\cdot)$ is not equal almost everywhere to a constant. If P_{out} is not informative, it is not difficult to show that estimation is then impossible.

Let us define the *replica-symmetric potential* (or just potential). Call $\rho := \mathbb{E}[(X^*)^2]$ where $X^* \sim P_0$. Then the potential is

$$f_{\text{RS}}(q, r; \rho) = f_{\text{RS}}(q, r) := \psi_{P_0}(r) + \alpha \Psi_{P_{\text{out}}}(q; \rho) - \frac{rq}{2}. \quad (41)$$

We define also $f_{\text{RS}}(\rho, +\infty) = \lim_{r \rightarrow \infty} f_{\text{RS}}(\rho, r)$. From now on denote $\psi'_{P_0}(r)$ and $\Psi'_{P_{\text{out}}}(q) = \Psi'_{P_{\text{out}}}(q; \rho)$ the derivatives of $\psi_{P_0}(r)$ and $\Psi_{P_{\text{out}}}(q; \rho)$ w.r.t. their first argument. We need also to define the set of the critical points of f_{RS} :

$$\Gamma := \left\{ (q, r) \in [0, \rho] \times (\mathbb{R}_+ \cup \{+\infty\}) \left| \begin{array}{l} q = 2\psi'_{P_0}(r) \\ r = 2\alpha\Psi'_{P_{\text{out}}}(q; \rho) \end{array} \right. \right\}, \quad (42)$$

where, with a slight abuse of notation, we define $\psi'_{P_0}(+\infty) = \lim_{r \rightarrow \infty} \psi'_{P_0}(r)$ and $\Psi'_{P_{\text{out}}}(\rho) = \lim_{q \rightarrow \rho} \Psi'_{P_{\text{out}}}(q)$. These limits are well defined by convexity of ψ_{P_0} and $\Psi_{P_{\text{out}}}$. The elements of Γ are called “fixed points of the state evolution”. Our first main result is

³The implications (h2) \Leftrightarrow (h2') and (h4) \Rightarrow (h4') are obvious. If (h4') holds one can show, by inverting cumulative distribution functions, that there exists a function $\varphi : \mathbb{R} \times [0, 1] \rightarrow \mathbb{R}$ such that (20) holds for $A_\mu \stackrel{\text{iid}}{\sim} P_A = \text{Unif}([0, 1])$ and that (h4) is verified.

Theorem 1 (Replica-symmetric formula for the free entropy). *Suppose that hypotheses (h1)-(h2)-(h3)-(h4) hold. Suppose that either hypothesis (h5.a) or (h5.b) holds. Then, for the generalized linear estimation model (20), (21) the thermodynamic limit of the free entropy (29) verifies*

$$f_\infty := \lim_{n \rightarrow \infty} f_n = \sup_{q \in [0, \rho]} \inf_{r \geq 0} f_{\text{RS}}(q, r) = \sup_{(q, r) \in \Gamma} f_{\text{RS}}(q, r). \quad (43)$$

Moreover, if P_{out} is informative, then the “sup inf” and the supremum over Γ in (43) are achieved over the same couples (q, r) .

An immediate corollary of Theorem 1 is the limiting expression of the mutual information between the signal and the observations. To state the result, we need to introduce two mutual informations associated to the two scalar channels presented in Sec. 1.3, namely

$$I_{P_0}(r) := I(X_0; \sqrt{r} X_0 + Z_0) = \frac{r\rho}{2} - \psi_{P_0}(r) \quad (44)$$

for the channel (36) and

$$\mathcal{I}_{P_{\text{out}}}(q) := I(W^*; \tilde{Y}_0 | V) = \Psi_{P_{\text{out}}}(\rho) - \Psi_{P_{\text{out}}}(q) \quad (45)$$

for the channel (39).

Corollary 2 (Single-letter formula for the mutual information). *The thermodynamic limit of the mutual information for model (20), (21) between the observations and the hidden variables verifies*

$$i_\infty := \lim_{n \rightarrow \infty} \frac{1}{n} I(\mathbf{X}^*; \mathbf{Y} | \Phi) = \inf_{q \in [0, \rho]} \sup_{r \geq 0} i_{\text{RS}}(q, r) = \inf_{(q, r) \in \Gamma} i_{\text{RS}}(q, r), \quad (46)$$

where

$$i_{\text{RS}}(q, r) := I_{P_0}(r) + \alpha \mathcal{I}_{P_{\text{out}}}(q) - \frac{r}{2}(\rho - q). \quad (47)$$

Proof. This follows from a simple calculation:

$$\begin{aligned} \frac{1}{n} I(\mathbf{X}^*; \mathbf{Y} | \Phi) &= \frac{1}{n} H(\mathbf{Y} | \Phi) - \frac{1}{n} H(\mathbf{Y} | \mathbf{X}^*, \Phi) = -f_n + \frac{1}{n} \mathbb{E} \ln P(\mathbf{Y} | \mathbf{X}^*, \Phi) \\ &= -f_n + \frac{m}{n} \mathbb{E} \ln P_{\text{out}}(Y_1 | \Phi_1 \cdot \mathbf{X}^* / \sqrt{n}). \end{aligned} \quad (48)$$

By the central limit theorem (that we can apply under hypotheses (h1)-(h3)) we have

$$S_n := \frac{1}{\sqrt{n}} \Phi_1 \cdot \mathbf{X}^* = \frac{1}{\sqrt{n}} \sum_{i=1}^n \Phi_{1,i} X_i^* \xrightarrow[n \rightarrow \infty]{(d)} \mathcal{N}(0, \rho).$$

Now, under the hypotheses (h2)-(h4) and either (h5.a) or (h5.b) it is not difficult to verify that

$$\begin{aligned} \mathbb{E} \ln P_{\text{out}}(Y_1 | \Phi_1 \cdot \mathbf{X}^* / \sqrt{n}) &= \mathbb{E} \int dY P_{\text{out}}(Y | S_n) \ln P_{\text{out}}(Y | S_n) \\ &\xrightarrow[n \rightarrow \infty]{} \mathbb{E} \int dY P_{\text{out}}(Y | \sqrt{\rho} V) \ln P_{\text{out}}(Y | \sqrt{\rho} V) = \Psi_{P_{\text{out}}}(\rho) \end{aligned}$$

where $V \sim \mathcal{N}(0, 1)$. We conclude, using (48):

$$\frac{1}{n} I(\mathbf{X}^*; \mathbf{Y} | \Phi) = -f_n + \alpha \Psi_{P_{\text{out}}}(\rho) + o_n(1) \quad (49)$$

where $\lim_{n \rightarrow \infty} o_n(1) = 0$. □

The next proposition, proved in Appendix A.2, states that for almost every $\alpha > 0$ there is one unique optimizer q^* in (43) (or equivalently in (46)):

Proposition 1. *Assume that the assumptions of Theorem 1 hold and that P_{out} is informative. Define*

$$D^* := \{ \alpha > 0 \mid (43) \text{ (or equivalently (46)) admits a unique optimizer } q^*(\alpha) \}. \quad (50)$$

the set D^* is equal to \mathbb{R}_+^* minus some countable set. Moreover $\alpha \mapsto q^*(\alpha)$ is continuous on D^* .

As an application of Theorem 1 we can compute the free entropy of the “planted perceptron” on the hypercube and the sphere. This perceptron model has already been studied in physics [22] and more recently in statistics, where it is known as “one-bit compressed sensing” [6, 15]. The limit of the free entropy follows from an application of Theorem 1 with $\varphi(x) = \text{sgn}(x)$ and $P_0 = \frac{1}{2}\delta_{-1} + \frac{1}{2}\delta_1$ (for the hypercube) or $P_0 = \mathcal{N}(0, 1)$ (for the sphere). For $\mu \in \{1, \dots, m\}$ we define

$$S_\mu := \left\{ \mathbf{x} \in \mathbb{R}^n \mid \text{sgn}(\mathbf{x} \cdot \Phi_\mu) = \text{sgn}(\mathbf{X}^* \cdot \Phi_\mu) \right\}. \quad (51)$$

We will use the notation $\mathcal{N}(x) = \mathbb{P}(Z \leq x)$ for $Z \sim \mathcal{N}(0, 1)$. Let \mathbb{S}_n be the unit sphere in \mathbb{R}^n and μ_n the uniform probability measure on \mathbb{S}_n .

Corollary 3 (Free entropy of the planted perceptron). *Let $Z, V \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$. We have*

$$\begin{aligned} & \frac{1}{n} \mathbb{E} \ln \left(\# \bigcap_{\mu=1}^m S_\mu \cap \{-1, 1\}^n \right) \\ & \xrightarrow{n \rightarrow \infty} \ln(2) + \sup_{q \in [0, 1]} \inf_{r \geq 0} \left\{ \mathbb{E} \ln \cosh(\sqrt{r}Z + r) + 2\alpha \mathbb{E} \left[\mathcal{N}\left(\frac{\sqrt{q}V}{\sqrt{1-q}}\right) \ln \mathcal{N}\left(\frac{\sqrt{q}V}{\sqrt{1-q}}\right) \right] - \frac{r(q+1)}{2} \right\}, \end{aligned} \quad (52)$$

$$\begin{aligned} & \frac{1}{n} \mathbb{E} \ln \mu_n \left(\bigcap_{\mu=1}^m S_\mu \cap \mathbb{S}_n \right) \\ & \xrightarrow{n \rightarrow \infty} \sup_{q \in [0, 1]} \left\{ \frac{1}{2} \ln(1-q) + 2\alpha \mathbb{E} \left[\mathcal{N}\left(\frac{\sqrt{q}V}{\sqrt{1-q}}\right) \ln \mathcal{N}\left(\frac{\sqrt{q}V}{\sqrt{1-q}}\right) \right] + \frac{q}{2} \right\}. \end{aligned} \quad (53)$$

2.2 Optimal reconstruction (or estimation) error

We first consider the problem of estimating \mathbf{X}^* given \mathbf{Y} and Φ . The following theorem states that the optimizer $q^*(\alpha)$ of the replica-symmetric formula (43) gives the asymptotic correlation between the planted solution \mathbf{X}^* and a typical sample from the posterior distribution $P(\cdot \mid \mathbf{Y}, \Phi)$:

Theorem 2 (Limit of the overlap). *Assume that all the moments of P_0 are finite and that P_{out} is informative. Assume that (h1)-(h2)-(h3)-(h4) hold and that either (h5.a) or (h5.b) holds. Then for all $\alpha \in D^*$,*

$$\frac{1}{n} |\mathbf{x} \cdot \mathbf{X}^*| = \frac{1}{n} \left| \sum_{i=1}^n x_i X_i^* \right| \xrightarrow{n \rightarrow \infty} q^*(\alpha), \quad \text{in probability,} \quad (54)$$

where $\mathbf{x} = (x_1, \dots, x_n)$ is sampled from the posterior distribution of the signal $P(\cdot \mid \mathbf{Y}, \Phi)$ given by (24), independently of everything else.

Theorem 2 is proved in Sec. 5.3. Notice that in all generality it is only possible to estimate \mathbf{X}^* up to its sign (think for instance to $\mathbf{Y} = |\Phi \mathbf{X}^*| / \sqrt{n} + \sqrt{\Delta} \mathbf{Z}$), this is why the absolute values in (54) are needed. For this reason, the usual MSE on \mathbf{X}^*

$$\text{mse}(\widehat{\mathbf{X}}) := \frac{1}{n} \mathbb{E} \left[\|\mathbf{X}^* - \widehat{\mathbf{X}}(\mathbf{Y}, \Phi)\|^2 \right]$$

is not (in all generality) an appropriate error metric. Indeed, in the case where $\mathbf{Y} = |\Phi \mathbf{X}^*|/\sqrt{n} + \sqrt{\Delta} \mathbf{Z}$, where $\Phi, \mathbf{X}^*, \mathbf{Z}$ have all independent $\mathcal{N}(0, 1)$ entries, then $\mathbb{E}[\mathbf{X}^* | \mathbf{Y}, \Phi] = 0$ and $\min_{\hat{\mathbf{X}}} \text{mse}(\hat{\mathbf{X}}) = 1$. This means that the minimum mean-square error is always equal to the variance and thus, in this sense, it is never possible to estimate the signal better than trivial estimators. For this reason, the appropriate error metric for the reconstruction problem is the MSE on $\mathbf{X}^* \mathbf{X}^{*\top}$. From Theorem 2 one deduces the limit of the MMSE in estimating $\mathbf{X}^* \mathbf{X}^{*\top}$:

Corollary 4 (Matrix minimum mean-square error). *Under the same conditions as in Theorem 2, for all $\alpha \in D^*$ we have*

$$\text{MMSE}_n := \frac{1}{n^2} \mathbb{E} \left[\left\| \mathbf{X}^* \mathbf{X}^{*\top} - \mathbb{E}[\mathbf{X}^* \mathbf{X}^{*\top} | \mathbf{Y}, \Phi] \right\|_{\text{F}}^2 \right] \xrightarrow{n \rightarrow \infty} \rho^2 - q^*(\alpha)^2, \quad (55)$$

where $\| \cdot \|_{\text{F}}$ denotes the Frobenius norm.

2.3 Optimal generalization (or prediction) error

In order to express the optimal generalization error we introduce the following function (recall that \tilde{Y}_0 is drawn from the channel (39)):

$$\mathcal{E}(q) := \text{MMSE}(\tilde{Y}_0 | V) = \mathbb{E}[(\tilde{Y}_0 - \mathbb{E}[\tilde{Y}_0 | V])^2] \quad (56)$$

$$= \mathbb{E}_V \int dY Y^2 P_{\text{out}}(Y | \sqrt{\rho} V) - \mathbb{E}_V \left[\mathbb{E}_W \left[\int dY Y P_{\text{out}}(Y | \sqrt{q} V + \sqrt{\rho - q} W) \right]^2 \right] \quad (57)$$

$$= \mathbb{E}[\varphi(\sqrt{\rho} V, \mathbf{A})^2] - \mathbb{E}_V [\mathbb{E}_{W, \mathbf{A}}[\varphi(\sqrt{q} V + \sqrt{\rho - q} W, \mathbf{A})]^2] + \Delta \quad (58)$$

where $V, W \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$, $\mathbf{A} \sim P_{\mathbf{A}}$ are independent random variables, $\mathbb{E}_{W, \mathbf{A}}$ denotes the expectation w.r.t. W and \mathbf{A} only and $\mathbb{E}_W[-]^2 = (\mathbb{E}_W[-])^2$. We recall $\rho := \mathbb{E}[(X^*)^2]$ with $X^* \sim P_0$. With a slight abuse of notation $\int dY$ denotes in (57) either the integration w.r.t. Lebesgue's measure on \mathbb{R} in the case $\Delta > 0$ or the integration w.r.t. the counting measure on \mathbb{N} (in the case $\Delta = 0$).

Recall the teacher-student setting of Sec. 1.2: The generalization error is related to the estimation of a new output $Y_{\text{new}} \sim P_{\text{out}}(\cdot | \Phi_{\text{new}} \cdot \mathbf{X}^* / \sqrt{n})$ where Φ_{new} is a new row of the matrix, and is defined by (32).

Theorem 3 (Optimal generalization error). *Assume that P_{out} is informative, that (h1)-(h2)-(h3)-(h4) hold and that either (h5.a) or (h5.b) hold. Then for all $\alpha \in D^*$ we have*

$$\mathcal{E}_{\text{gen}}^{\text{opt}}(\alpha) \xrightarrow{n \rightarrow \infty} \mathcal{E}(q^*(\alpha)) \quad (59)$$

where $q^*(\alpha)$ is the optimizer of the replica-symmetric formula (43), see Proposition 1.

Theorem 3 follows from a more general result, that we state now. Let $f : \mathbb{R} \rightarrow \mathbb{R}$ and consider the generalized optimal generalization error

$$\mathcal{E}_{f,n}(\alpha) := \text{MMSE}(f(Y_{\text{new}}) | \Phi_{\text{new}}, \mathbf{Y}, \Phi) = \mathbb{E}[(f(Y_{\text{new}}) - \mathbb{E}[f(Y_{\text{new}}) | \Phi_{\text{new}}, \mathbf{Y}, \Phi])^2] \quad (60)$$

which is the minimum mean-square error on $f(Y_{\text{new}})$. In particular $\mathcal{E}_{\text{gen}}^{\text{opt}}(\alpha) = \mathcal{E}_{f,n}(\alpha)$ for $f : x \mapsto x$. We define also

$$\mathcal{E}_f(q) := \text{MMSE}(f(\tilde{Y}_0) | V) = \mathbb{E}[(f(\tilde{Y}_0) - \mathbb{E}[f(\tilde{Y}_0) | V])^2] \quad (61)$$

$$= \mathbb{E}[f(\varphi(\sqrt{\rho} V, \mathbf{A}) + \sqrt{\Delta} Z)^2] - \mathbb{E}_V [\mathbb{E}_{W, Z, \mathbf{A}}[f(\varphi(\sqrt{q} V + \sqrt{\rho - q} W, \mathbf{A}) + \sqrt{\Delta} Z)^2]], \quad (62)$$

where \tilde{Y}_0 is the output of the second scalar channel (39).

Theorem 4 (Generalized optimal generalization error). *Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a measurable function such that $\mathbb{E}[|f(Y_1)|^{2+\gamma}]$ remains bounded as n grows, for some $\gamma > 0$. Assume that P_{out} is informative and that (h1)-(h2)-(h3)-(h4) hold and that either (h5.a) or (h5.b) holds. Then for all $\alpha \in D^*$ we have*

$$\mathcal{E}_{f,n}(\alpha) \xrightarrow[n \rightarrow \infty]{} \mathcal{E}_f(q^*(\alpha)) \quad (63)$$

where $q^*(\alpha)$ is the optimizer of the replica-symmetric formula (43), see Proposition 1.

Theorem 4 is proved in Sec. 5.1.

2.4 Optimality of the generalized approximate message-passing algorithm

2.4.1 The generalized approximate message-passing algorithm

While the main results presented until now are information-theoretic, our next one concerns the performance of a popular algorithm to solve random instances of generalized linear problems, called generalized approximate message-passing (GAMP). We shall not re-derive its properties here, and instead refer to the original papers for details. This approach has a long history, especially in statistical physics [41, 56, 57, 75], error correcting codes [76], and graphical models [77]. For a modern derivation in the context of linear models, see [11, 68, 78]. The case of generalized linear models was discussed by Rangan in [12], and has been used for classification purpose in [79].

We first need to define two so-called threshold functions that are associated to the two scalar channels (36) and (39). The first one is the posterior mean of the signal in channel (36) with signal-to-noise ratio r :

$$g_{P_0}(y, r) := \mathbb{E}[X_0 | Y_0 = y]. \quad (64)$$

The second one is the posterior mean of W^* in channel (39) with “noise level” $\eta = \rho - q$:

$$g_{P_{\text{out}}}(\tilde{y}, v, \eta) := \mathbb{E}[W^* | \tilde{Y}_0 = \tilde{y}, \sqrt{q} V = v]. \quad (65)$$

These functions act componentwise when applied to vectors.

Given initial estimates $(\hat{\mathbf{x}}^0, \mathbf{v}^0)$ for the means and variances of the elements of the signal vector \mathbf{X}^* , GAMP takes as input the observation vector \mathbf{Y} and then iterates the following equations with initialization $g_\mu^0 = 0$ for all $\mu = 1, \dots, m$ (we denote by $\bar{\mathbf{u}}$ the average over all the components of the vector \mathbf{u} and Φ^\top is the transpose of the matrix Φ): From $t = 1$ until convergence,

$$\begin{cases} V^t &= \overline{\mathbf{v}^{t-1}} \\ \boldsymbol{\omega}^t &= \Phi \hat{\mathbf{x}}^{t-1} / \sqrt{n} - V^t \mathbf{g}^{t-1} \\ g_\mu^t &= g_{P_{\text{out}}}(Y_\mu, \omega_\mu^t, V^t) \\ \lambda^t &= \alpha g_{P_{\text{out}}}^2(\mathbf{Y}, \boldsymbol{\omega}^t, V^t) \\ \mathbf{R}^t &= \hat{\mathbf{x}}^{t-1} + (\lambda^t)^{-1} \Phi^\top \mathbf{g}^t / \sqrt{n} \\ \hat{x}_i^t &= g_{P_0}(R_i^t, \lambda^t) \\ \mathbf{v}_i^t &= (\lambda^t)^{-1} \partial_R g_{P_0}(R, \lambda^t)|_{R=R_i^t} \end{cases} \quad \forall \mu = 1, \dots, m \quad \forall i = 1, \dots, n \quad (66)$$

One of the strongest asset of GAMP is that its performance can be tracked rigorously in the limit $n, m \rightarrow \infty$ while $m/n \rightarrow \alpha$ via a procedure known as state evolution (SE), see [43, 44] for the linear case, and [12, 45] for the generalized one. In our notations, state evolution tracks the asymptotic value of the overlap between the true hidden value \mathbf{X}^* and its estimate by GAMP $\hat{\mathbf{x}}^t$ defined as $q^t := \lim_{n \rightarrow \infty} \mathbf{X}^* \cdot \hat{\mathbf{x}}^t / n$ (that is related to the asymptotic mean-square error (MSE) E^t between \mathbf{X}^* and its estimate $\hat{\mathbf{x}}^t$ by $E^t = \rho - q^t$, where recall that

$\rho := \mathbb{E}[(X^*)^2]$ with $X^* \sim P_0$) via:

$$\begin{cases} q^{t+1} &= 2\psi'_{P_0}(r^t), \\ r^t &= 2\alpha\Psi'_{P_{\text{out}}}(q^t; \rho). \end{cases} \quad (67)$$

From Theorem 1 we realize that the fixed points of these equations correspond to the critical points of the asymptotic free entropy in (43). In fact, in the replica heuristic, the optimizer q^* of the potential is conjectured to give the optimal value of the overlap, a fact that was proven for the linear channel [46–48]. We will see in Sec. 3 that $q^t \xrightarrow[t \rightarrow \infty]{} q^*$ for a large set of parameters.

2.4.2 Estimation and generalization error of GAMP

Perhaps more surprisingly, one can use GAMP in the teacher-student scenario described in Sec. 1.2 in order to provide an estimation of a new output $Y_{\text{new}} \sim P_{\text{out}}(\cdot | \Phi_{\text{new}} \cdot \mathbf{X}^* / \sqrt{n})$ where Φ_{new} is a new row of the matrix. As $\hat{\mathbf{x}}^t$ is the GAMP estimate of the posterior expectation of \mathbf{X}^* , with estimated variance \mathbf{v}^t , the natural heuristic is to consider for the posterior probability distribution of the random variable $\Phi_{\text{new}} \cdot \mathbf{X}^* / \sqrt{n}$ a Gaussian with mean $\Phi_{\text{new}} \cdot \hat{\mathbf{x}}^{t-1} / \sqrt{n}$ and variance $V^t = E^t = \rho - q^t$ (the fact that the variance and MSE are equal follows from the Nishimori identity of Proposition 12 but applied to GAMP instead of the Gibbs measure, see e.g. [13] where this is shown). This allows to estimate the posterior mean of the output, which leads to the GAMP prediction (recall the P_{out} definition (22)-(23)):

$$\hat{Y}^{\text{GAMP},t} := \int y P_{\text{out}}\left(y \mid \frac{1}{\sqrt{n}} \Phi_{\text{new}} \cdot \hat{\mathbf{x}}^{t-1} + \sqrt{\rho - q^t} w\right) \mathcal{D}w dy, \quad (68)$$

where $\mathcal{D}w$ denotes the standard Gaussian measure. The following claim, from [12], gives the precise estimation error of GAMP. It is stated there as a claim because some steps of the proof are missing. The paper [45] affirms in its abstract to prove the claim of [12], but without further details. For these reasons, we believe that the claim holds, however we prefer to state it here as a claim (instead of a theorem).

Claim 1 (GAMP estimation error, [12]). *We have almost surely for all $t \in \mathbb{N}$,*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \hat{\mathbf{x}}^t \cdot \mathbf{X}^* = \lim_{n \rightarrow \infty} \frac{1}{n} \|\hat{\mathbf{x}}^t\|^2 = q^t, \quad (69)$$

as well as

$$\lim_{n \rightarrow \infty} \frac{1}{n^2} \mathbb{E} \left[\|\mathbf{X}^* \mathbf{X}^{*\top} - \hat{\mathbf{x}}^t (\hat{\mathbf{x}}^t)^\top\|^2 \right] = \rho^2 - (q^t)^2. \quad (70)$$

Comparing (70) with the MMSE given by Corollary 5, we see that if $\lim_{t \rightarrow \infty} q^t = q^*(\alpha)$, then GAMP achieves the MMSE. Provided that Claim 1 holds we can deduce the generalization error of GAMP:

Proposition 2 (GAMP generalization error). *Suppose that hypotheses (h1)-(h2)-(h4) hold. Moreover suppose that either (h5.a) or (h5.b) holds. Assume that $(\Phi_{\mu i}) \stackrel{iid}{\sim} \mathcal{N}(0, 1)$, and that $x \mapsto P_{\text{out}}(\cdot | x)$ is continuous almost everywhere for the Wasserstein distance of order 2. Let $t \in \mathbb{N}$. Assume that the limit (69) holds in probability and that there exists $\eta > 0$ such that $\mathbb{E}[\|\hat{Y}^{\text{GAMP},t}\|^{2+\eta}]$ remains bounded (as n grows). Then we have*

$$\lim_{n \rightarrow \infty} \mathcal{E}_{\text{gen}}^{\text{GAMP},t} := \lim_{n \rightarrow \infty} \mathbb{E}[(Y_{\text{new}} - \hat{Y}^{\text{GAMP},t})^2] = \mathcal{E}(q^t). \quad (71)$$

Remark 2. *If we modify slightly the GAMP estimator of (68) by changing the first y into $f(y)$, it is not difficult to show (following the steps of Proposition 2) that this new estimator achieves an asymptotic error of $\mathcal{E}_f(q^t)$, given by (61), for estimating $f(Y_{\text{new}})$.*

Proposition 2 is proved in Sec. 5.2. We see that this formula matches the one for the Bayes-optimal generalization error, see Theorem 3, up to the fact that instead of $q^*(\alpha)$ (the optimizer of the replica formula (43)) appearing in the optimal error formula, here it is q^t which appears. Thus clearly, when q^t converges to $q^*(\alpha)$ (we shall see that this is the case in many situations in the examples of Sec. 3) this yields a very interesting and non trivial result: *GAMP achieves the Bayes-optimal generalization error* in a plethora of models (a task again often believed to be intractable) and this for large sets of parameters.

2.5 Optimal denoising error

Another interesting error measure to study is the following “denoising error”. Assume that the observations are noisy, i.e. $\Delta > 0$ in (20). The goal here is to denoise the observations Y_μ and estimate the signal which in this case is $\varphi\left(\frac{1}{\sqrt{n}}[\Phi \mathbf{X}^*]_\mu, \mathbf{A}_\mu\right)$.

The minimum denoising error (in L^2 sense) is actually a simple corollary from the replica-symmetric formula of Theorem 1 and follows from a so-called “I-MMSE relation”, see Proposition 13. We will need the joint posterior distribution of (W^*, \mathbf{A}) given (V, \tilde{Y}_0) for the scalar channel (39). So we define the Gibbs bracket for the scalar channel by (here $\mathbf{a} \in \mathbb{R}^{k_A}$):

$$\langle g(w, \mathbf{a}) \rangle_{\text{sc}} := \mathbb{E}[g(W^*, \mathbf{A}) | \tilde{Y}_0, V] = \frac{\int \mathcal{D}w dP_A(\mathbf{a}) g(w, \mathbf{a}) e^{-\frac{1}{2\Delta} \left\{ \tilde{Y}_0 - \varphi(\sqrt{q}V + \sqrt{\rho - q}w, \mathbf{a}) \right\}^2}}{\int \mathcal{D}w dP_A(\mathbf{a}) e^{-\frac{1}{2\Delta} \left\{ \tilde{Y}_0 - \varphi(\sqrt{q}V + \sqrt{\rho - q}w, \mathbf{a}) \right\}^2}}, \quad (72)$$

for any continuous bounded function g . When the function depends only on w it may be re-written as

$$\langle g(w) \rangle_{\text{sc}} = \frac{\int \mathcal{D}w g(w) P_{\text{out}}(\tilde{Y}_0 | \sqrt{q}V + \sqrt{\rho - q}w)}{\int \mathcal{D}w P_{\text{out}}(\tilde{Y}_0 | \sqrt{q}V + \sqrt{\rho - q}w)}. \quad (73)$$

Corollary 5 (Optimal denoising error). *Suppose that hypotheses (h1)-(h2)-(h3)-(h4) hold. Suppose that either hypothesis (h5.a) or (h5.b) holds. Then for almost every $\Delta > 0$, for any optimal couple (q^*, r^*) of (43),*

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{m} \text{MMSE} \left(\varphi \left(\frac{1}{\sqrt{n}} \Phi \mathbf{X}^*, \mathbf{A} \right) \middle| \Phi, \mathbf{Y} \right) &= \text{MMSE} \left(\varphi(\sqrt{q^*}V + \sqrt{\rho - q^*}W^*, \mathbf{A}) \middle| \tilde{Y}_0, V \right) \\ &= \mathbb{E}[\varphi(\sqrt{\rho}V, \mathbf{A})^2] - \mathbb{E}[\langle \varphi(\sqrt{q^*}V + \sqrt{\rho - q^*}w, \mathbf{a}) \rangle_{\text{sc}}^2], \end{aligned} \quad (74)$$

where $\langle - \rangle_{\text{sc}}$ acts jointly on (w, \mathbf{a}) and is defined by (72), and $V, W^* \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$.

Note that the joint posterior over both the signal \mathbf{X}^* and the random stream \mathbf{A} is simply expressed as

$$dP(\mathbf{x}, \mathbf{a} | \mathbf{Y}, \Phi) \propto dP_0(\mathbf{x}) dP_A(\mathbf{a}) \prod_{\mu=1}^m e^{-\frac{1}{2\Delta} \left\{ Y_\mu - \varphi \left(\frac{1}{\sqrt{n}} [\Phi \mathbf{x}]_\mu, \mathbf{a}_\mu \right) \right\}^2}. \quad (75)$$

The proof of Corollary 5 is presented in Sec. 5.4.

3 Application to concrete situations

In this section, we show how our main results can be applied to several models of interest in fields ranging from machine learning to signal processing, and unveil several interesting new phenomena in learning of generalized linear models. For various specific cases of prior P_0 and output P_{out} , we evaluate numerically the free entropy potential (41), its stationary points Γ and identify which of them gives the information-theoretic results, i.e. is the optimizer in (43). We also identify which of the stationary points corresponds to the result

obtained asymptotically by the GAMP algorithm, i.e. the fixed point of the state evolution (67). Finally we compute the corresponding generalization error (71). We stress that in this section the results are based on numerical investigation of the resulting formulas: We do not aim at rigor that would involve precise bounds and more detailed analytical control for the corresponding integrals.

3.1 Generic observations

Using the functions $g_{P_{\text{out}}}$ and g_{P_0} introduced in Sec. 2.4 we can rewrite the fixed point equations (42) as

$$q = 2\psi'_{P_0}(r) = \mathbb{E}[g_{P_0}(Y_0, r)^2], \quad (76)$$

$$r = 2\alpha\Psi'_{P_{\text{out}}}(q) = \frac{\alpha}{\rho - q}\mathbb{E}[g_{P_{\text{out}}}(\tilde{Y}_0, \sqrt{q}V, \rho - q)^2], \quad (77)$$

where the expectation in (76) corresponds to the scalar channel (36) and the expectation in (77) corresponds to the second scalar channel (39).

Non-informative fixed point and its stability: It is interesting to analyze under what conditions $q^* = 0$ is the optimizer of (43). Notice that $q^* = 0$ corresponds to the error on the recovery of the signal as large as it would be if we had no observations at our disposition. Theorem 1 gives that any optimal couple (q^*, r^*) of (43) should be a fixed point of the state evolution equations (76)–(77). A sufficient condition for $(q, r) = (0, 0)$ to be a fixed point of (76)–(77) is that:

- (a) The transition density $P_{\text{out}}(y|z)$ is even in the argument z .
- (b) The prior P_0 has zero mean.

In order to see this, notice that if $P_{\text{out}}(y|z)$ is even in z then from the definition (65) of the function $g_{P_{\text{out}}}$ we have $g_{P_{\text{out}}}(y, 0, \rho) = 0$ and consequently from (77) we have $\Psi'_{P_{\text{out}}}(0) = 0$. From the second point, notice that we have $\psi'_{P_0}(0) = \frac{1}{2}\mathbb{E}_{P_0}[X_0]^2 = 0$.

We assume now that the transition density $P_{\text{out}}(y|z)$ is even in the argument z and that the prior P_0 has zero mean. In order for $q = 0$ to be the global maximizer q^* of (43) or to be a relevant fixed point of the state evolution (67) (relevant in the sense that GAMP might indeed converge to it in a practical setting) we need $q = 0$ to be a stable fixed point of the above equations (76)–(77). We therefore need to expand (76)–(77) around $q = 0$, and doing so, we obtain that $q = 0$ is stable if

$$2\alpha\Psi''_{P_{\text{out}}}(0) \times 2\psi''_{P_0}(0) = \alpha\mathbb{E}[(\langle w^2 \rangle_{\text{sc}} - \langle w \rangle_{\text{sc}}^2 - 1)^2] < 1, \quad (78)$$

where the expectation corresponds to the scalar channel (39) with $q = 0$ and the Gibbs bracket $\langle - \rangle_{\text{sc}}$ is given by (72). The expectation quantifies how the observation of \tilde{Y}_0 in the scalar channel (39) modifies the variance of W^* (which is 1 without any observation). Rewriting this condition more explicitly into a form that is convenient for numerical evaluation we get (recalling (73), $q = 0$ and condition (a))

$$\alpha \int dy \frac{(\int \mathcal{D}z (z^2 - 1) P_{\text{out}}(y|\sqrt{\rho}z))^2}{\int \mathcal{D}z P_{\text{out}}(y|\sqrt{\rho}z)} < 1, \quad (79)$$

where recall that $\mathcal{D}z$ is a standard Gaussian measure. We conjecture that the condition (79) delimits precisely the region where polynomial-time algorithms do not perform better than “random guessing” (see the discussion below, where we will make this stability condition explicit for several examples of symmetric output channels). Note that the condition (79) also appears in a recent work [61] as a barrier for performance of spectral algorithms.

Exact recovery fixed point: Another particular fixed point of (76)–(77) that we observe is the one corresponding to exact recovery $q^* = \rho$. A sufficient and necessary condition for this to be a fixed point is that $\lim_{q \rightarrow \rho} \Psi'_{P_{\text{out}}}(q) = +\infty$. Heuristically, this means that the integral of the Fisher information of the output channel should diverge:

$$\int dy d\omega \frac{e^{-\frac{\omega^2}{2\rho}} P'_{\text{out}}(y|\omega)^2}{\sqrt{2\pi\rho} P_{\text{out}}(y|\omega)} = +\infty, \quad (80)$$

where $P'_{\text{out}}(y|\omega)$ denotes the partial derivative w.r.t. ω . This typically means that the channel should be noiseless. For the Gaussian channel with noise variance Δ , the above expression equals $1/\Delta$. For the probit channel where $P_{\text{out}}(y|z) = \text{erfc}(-yz/\sqrt{2\Delta})/2$ the above expression at small Δ is proportional to $1/\sqrt{\Delta}$.

Stability of the exact recovery fixed point was also investigated, but we did not obtain any unified expression. The stability depends non-trivially on both the properties of the output channel, but also on the properties of the prior. Below we give several examples where exact recovery either is or is not possible, or where there is a phase transition between the two regimes.

3.2 Phase diagram of perfect learning

In this section we consider *deterministic (noiseless) output channels* and ask: How many measurements are needed in order to perfectly recover the signal?

Our crucial point is to compare with the well explored phase diagram of Bayesian (noiseless) compressed sensing in the case of the linear channel [68, 80]. As the number of samples (measurements) varies we encounter five different regimes of parameters:

- The *tractable recovery* phase: This is the region in the parameter space where GAMP achieves perfect reconstruction.
- The *non-informative* phase: Region where perfect reconstruction is information-theoretically impossible and moreover even the Bayes-optimal estimator is as bad as a random guess based on the prior information and on the knowledge of the output function.
- The *no recovery* phase: Region where perfect reconstruction is information-theoretically impossible, but an estimator positively correlated with the ground truth exists.
- The *hard* phase: Region where the perfect reconstruction is information-theoretically possible, but where GAMP is unable to achieve it. At the same time, in this region GAMP leads to a better generalization error than the one corresponding to the non-informative fixed point. It remains a challenging open question whether polynomial-time algorithms can achieve perfect reconstruction in this regime.
- The *hard non-informative* phase: This phase corresponds to the region where perfect reconstruction is information-theoretically possible but where GAMP only achieves an error as bad as randomly guessing, given by the trivial fixed point. In this phase as well, the existence of polynomial-time exact recovery algorithms is an open question. This phase does not exist for the linear channel.

Some of the codes used in this section can be consulted online on the github repository [62].

3.2.1 The linear channel

The case of exact recovery of a sparse signal after it passed through a noiseless linear channel, i.e. $\varphi(x) = x$, is studied in the literature in great details, especially in the context of compressed sensing [10, 11]. For a signal with a fraction ρ of non-zero entries it is found that as soon as $\alpha > \rho$, perfect reconstruction is theoretically possible, although it may remain computationally difficult. The whole field of compressed sensing builds on

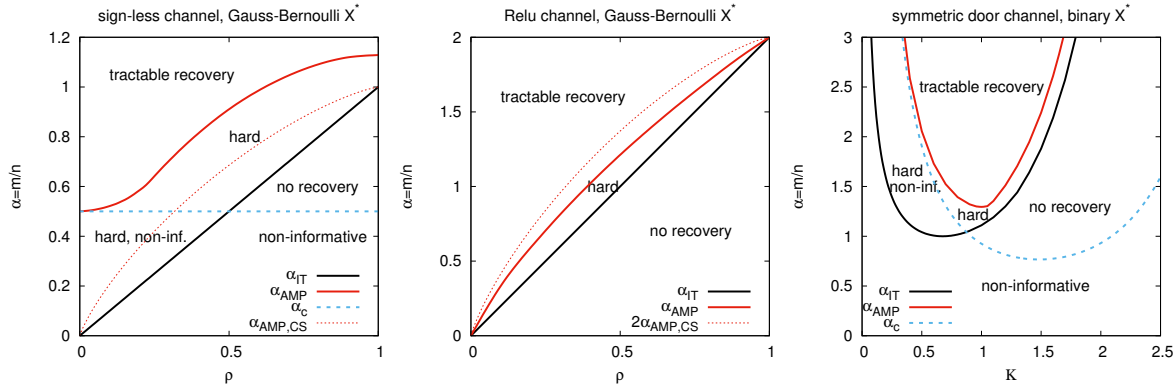


Figure 3: Phase diagrams showing boundaries of the region where exact recovery is possible (in absence of noise). **Left:** The case of sign-less sparse recovery, $\varphi(x) = |x|$ with a Gauss-Bernoulli signal, as a function of the ratio between number of samples/measurements and the dimension $\alpha = m/n$, and the fraction of non-zero components ρ . Evaluating the free entropy for this case, we find that a recovery of the signal is information-theoretically impossible for $\alpha < \alpha_{IT} = \rho$. Recovery becomes possible starting from $\alpha > \rho$, just as in the canonical compressed sensing. Algorithmically the sign-less case is much harder. Evaluating (79) we conclude that GAMP is not able to perform better than a random guess as long as $\alpha < \alpha_c = 1/2$. For larger values of α , the inference using GAMP leads to better results than a purely random guess. GAMP can exactly recover the signal and generalize perfectly only for values of α larger than α_{AMP} (full red line). The dotted red line shows for comparison the algorithmic phase transition of the canonical compressed sensing. **Center:** Analogous to the left panel, for the ReLU output function, $\varphi(x) = \max(0, x)$. Here it is always possible to perform better than random guessing using GAMP. The dotted red line shows the algorithmic phase transition when using information only about the non-zero observations. **Right:** Phase diagram for the symmetric door output function $\varphi(z) = \text{sign}(|z| - K)$ for a Rademacher signal, as a function of α and K . The stability line α_c is depicted in dashed blue, the information-theoretic phase transition to exact recovery α_{IT} in black, and the algorithmic one α_{AMP} in red.

the realization that, using the ℓ_1 norm minimization technique, one can efficiently recover the signal for larger α values than the so-called Donoho-Tanner transition [11, 81].

In the context of the present paper, when the empirical distribution of the signal is known, one can fairly easily beat the ℓ_1 transition and reconstruct the signal up to lower values of α using the Bayesian GAMP algorithm [11, 12, 68, 80]. In this case, three different phases are present [68, 80]: *i*) For $\alpha < \rho$, perfect reconstruction is impossible; *ii*) for $\rho < \alpha < \alpha_s$ reconstruction is possible, but not with any known polynomial-complexity algorithm; *iii*) for $\alpha > \alpha_s$, the so-called *spinodal transition* computed with state evolution, GAMP provides a polynomial-complexity algorithm able to reach perfect reconstruction. The line $\alpha_s(\rho)$ depends on the distribution of the signal. For a Gauss-Bernoulli signal with a fraction ρ of non-zero (Gaussian) values we compare the GAMP performance to the optimal one in Fig. 3 (left and right). This is the same figure as in the main text. We copy it here so that the SI is self-contained.

3.2.2 The rectified linear unit (ReLU) channel

Let us start by discussing the case of a generalized linear model with the ReLU output channel, i.e. $\varphi(x) = \max(0, x)$, with a signal coming from a Gauss-Bernoulli distribution $P_0 = \rho\mathcal{N}(0, 1) + (1 - \rho)\delta_0$ with a fraction ρ of non-zero (Gaussian) values. We are motivated by the omnipresent use of the ReLU activation function in deep learning, and explore its properties for GLMs that can be seen as a simple single layer neural network.

Our analysis shows that a perfect generalization (and thus a perfect reconstruction of the signal as well) is possible whenever the number of samples per dimension (measurement rate) $\alpha > 2\rho$, and impossible when $\alpha < 2\rho$. This is very intuitive, since half of the measurements (those non-zero) are giving as much information as in the linear case, thus the factor 2.

How hard is it to actually solve the problem with an efficient algorithm? The answer is given by applying the state evolution analysis to GAMP, which tells us that only for even larger values of α , beyond the spinodal transition, does GAMP reach a perfect recovery. Notice, however, that this spinodal transition occurs at a significantly lower measurement rate α than one would reach just keeping the non-zero measurements. This shows that, actually, these zero measurements contain a useful information for the algorithm. The situation is shown in the center panel of Fig. 3: The zero measurements do not help information-theoretically but they, however, do help algorithmically.

3.2.3 The sign-less channel

We now discuss the sign-less channel where only the absolute value of the linear mixture is observed, i.e. $\varphi(x) = |x|$. This case can be seen as the real-valued analog of the famous phase retrieval problem. We again consider the signal to come from a Gauss-Bernoulli distribution with a fraction ρ of non-zero (Gaussian) values.

Sparse phase retrieval has been well explored in the literature in the regime where the number s of non-zeros is sub-leading in the dimension, $s = \mathcal{O}(n)$. This case is known to present a large algorithmic gap. While analogously to compressed sensing exact recovery is information-theoretically possible for a number of measurement $\Omega(s \ln(n/s))$, best known algorithms achieve it only with $\Omega(s^2/\ln n)$ measurements [82], see also [70] and references therein for a good discussion of other related literature. This is sometimes referred to as the s^2 barrier. We are not aware of a study where, as in our setting, the sparsity is $s = \rho n$ and the number of measurements is αn with $\alpha = \Omega(1)$. Our analysis in this regime hence sheds a new light on the hardness of the problem of recovering a sparse signal from sign-less measurements.

Our analysis of the mutual information shows that a perfect reconstruction is information-theoretically possible as soon as $\alpha > \rho$: In other words, the problem is –information-theoretically– as easy, or as hard as the compressed sensing one. This is maybe less surprising when one thinks of the following algorithm: Try all 2^m choices of the possible signs for the m outputs, and solve a compressed sensing problem for each of them. Clearly, this should yields a perfect solution only in the case of the actual combination of signs.

Algorithmically, however, the problem is much harder than for the linear output channel. As shown in the left side of Fig. 3, for small ρ one requires a much larger fraction α of measurements in order for GAMP to recover the signal. For the linear channel the algorithmic transition $\alpha_s(\rho) \rightarrow 0$ as $\rho \rightarrow 0$, while for the sign-less channel we get $\alpha_s(\rho) \rightarrow 1/2$ as $\rho \rightarrow 0$. In other words if one loses the signs one cannot perform recovery in compressed sensing with less than $n/2$ measurements.

What we observe in this example for $\alpha < 1/2$ is in the statistical physics literature on neural networks known as retarded learning [64]. This appears in problems where the $\varphi(x)$ function is symmetric, as seen at the beginning of this section: There is always a critical point of the mutual information with an overlap value $q = 0$. For this problem, this critical point is actually “stable” (meaning that it is actually a local minimum in q in the mutual information (46)) for all $\alpha < 1/2$ independently of ρ . To see that, we have to go back to (78). For the absolute value channel, the posterior distribution of W^* given $\tilde{Y}_0 = \sqrt{\rho} |W^*|$ in the second scalar channel (39) is a mean of two Dirac masses on $-W^*$ and W^* . Thus, the posterior variance is $\langle w^2 \rangle_{sc} - \langle w \rangle_{sc}^2 = (W^*)^2$. Consequently, (78) leads to the stability condition $\alpha < 1/2$.

This has the two following implications: *i*) In the non-informative phase, when $\alpha < 1/2$ and $\rho > \alpha$, the minimum at $q = 0$ is actually the global one. In this case, the MMSE on \mathbf{X}^* and the generalization error are the ones given by using 0 as a guess for each element of \mathbf{X}^* ; in other words, there is no useful information that one can exploit and no algorithmic approach can be better than a random guess. *ii*) In the hard non-informative phase when $\alpha < 1/2$, GAMP initialized at random, i.e. close to the $q = 0$ fixed point, will remain there. This suggests that in this region, even if a perfect reconstruction is information-theoretically possible, it will still be very hard to beat a random guess with a tractable algorithm.

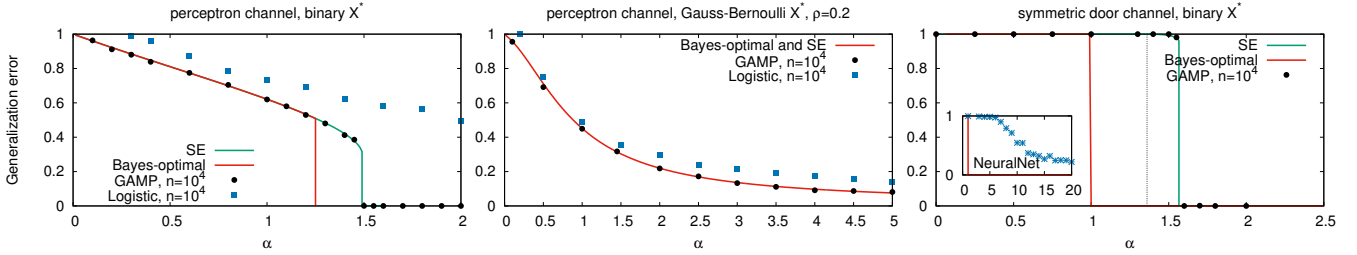


Figure 4: Generalization error in three classification problems as a function of the number of data-samples per dimension α . The red line is the Bayes-optimal generalization error, while the green one shows the (asymptotic) performances of GAMP as predicted by the state evolution (SE), when different. For comparison, we also show the result of GAMP (black dots) and, in blue, the performance of a standard out-of-the-box solver, both tested on a single randomly generated instance. **Left:** Perceptron, with $\varphi(x) = \text{sgn}(x)$ and a Rademacher (± 1) signal. While a perfect generalization is information-theoretically possible starting from $\alpha = 1.249(1)$, the state evolution predicts that GAMP will allow such perfect prediction only from $\alpha = 1.493(1)$. The results of a logistic regression with fine-tuned ridge penalty with the software scikit-learn [65] are shown for comparison. **Middle:** Perceptron with Gauss-Bernoulli coefficients for the signal. No phase transition is observed in this case, but a smooth decrease of the error with α . The results of a logistic regression with fine-tuned ℓ_1 sparsity-enhancing penalty (again with [65]) are very close to optimal. **Right:** The symmetric door activation rule with parameter $K = 0.67449$ chosen in order to observe the same number of occurrence of the two classes. In this case there is a sharp phase transition at $\alpha = 1$ from a situation where it is impossible to learn the rule, so that the generalization is not better than a random guess, to a situation where the optimal generalization error drops to zero. However, GAMP identifies the rule perfectly only starting from $\alpha_s = 1.566(1)$ (GAMP error stays 1 up to $\alpha_{\text{stab}} = 1.36$, see the black dashed curve). Interestingly, this non-linear rule seems very hard to learn for other existing algorithms. Using Keras [66], a neural network with two hidden layers was able to learn approximately the rule, but only for much larger training set sizes (shown in inset, the Keras/tensorflow code for this particular run can be found on the github repository [62]).

3.2.4 The symmetric door channel

The third output channel we study in detail is the symmetric door channel, where $\varphi(x) = \text{sgn}(|x| - K)$. In case of channels with discrete set of outputs exact recovery is only possible when the prior is also discrete. In the present case we consider the signal to be Rademacher, where each element is chosen at random between 1 and -1 , i.e. $P_0 = \frac{1}{2}\delta_{+1} + \frac{1}{2}\delta_{-1}$. This channel was studied previously using the replica method in the context of optimal data compression [28].

This output channel is in the class of symmetric channels for which overlap $q = 0$ is a fixed point. This fixed point is stable for $\alpha < \alpha_c(K)$. Exact recovery is information-theoretically possible above $\alpha_{\text{IT}}(K)$ and tractable with the GAMP algorithm above the spinodal transition $\alpha_s(K)$. The values of these three transition lines are depicted in the right panel of Fig. 3.

We note that $\alpha_{\text{IT}} \geq 1$ is a generic bound on exact recovery for every K , required by a simple counting argument. While a-priori it is not clear whether this bound is saturated for some K , we observe that it is for $K = 0.67449$ such that half of the observed measurements are negative and the rest positive. This is, however, not efficiently achievable with GAMP. The saturation of the $\alpha_{\text{IT}} \geq 1$ bound was remarked previously in the context of the work [28] on optimal data compression. Our work predicts that this information-theoretic result will not be achievable with known efficient algorithms.

3.3 Examples of optimal generalization error

Besides the formula for the mutual information, the main result of this paper is the Theorem 3 for the optimal generalization error, and formula (71) for the generalization error achieved by the GAMP algorithm. In this section we evaluate both these generalization errors for several cases of priors and output functions. We study both regression problems, where the output is real-valued, and classification problems, where the output is

discrete.

While in realistic regression and classification problems the matrix Φ corresponds to the data, and is thus not i.i.d. random, we view the practical interest of our theory as a benchmark for state-of-the-art algorithms. Our work provides an exact asymptotic analysis of optimal generalization error and sample complexity for a range of simple rules where a teacher uses random data to generate labels. The challenge for state-of-the-art multi-purpose algorithms is to try to match as closely as possible the performance that can be obtained with GAMP that is fine-tuned to the specific form of the output and prior.

3.3.1 Threshold output: The perceptron

The example of non-linear output that is the most widely explored in the literature is the threshold output, where the deterministic output (or “activation”) function is $\varphi(x) = \text{sgn}(x)$. This output in the teacher-student setting of the present paper is known as the perceptron problem [22], or equivalently, the one-bit compressed sensing in signal processing [6]. Its solution has been discussed in details within the replica formalism (see for instance [15, 55, 57, 83]) and we confirm all of these heuristic computations within our approach. Let $V \sim \mathcal{N}(0, 1)$. The formula (57), (63) for the generalization error then reduces to (recall $q^* = q^*(\alpha)$ is an optimizer of (43))

$$\lim_{n \rightarrow \infty} \mathcal{E}_{\text{gen}}^{\text{opt}} = 1 - \int \mathcal{D}V \left(\frac{2}{\sqrt{\pi}} \int_0^V \sqrt{\frac{q^*}{2(\rho - q^*)}} dt e^{-t^2/2} \right)^2 = 1 - \mathbb{E} \left[\text{erf} \left(V \sqrt{\frac{q^*}{2(\rho - q^*)}} \right)^2 \right]. \quad (81)$$

In Fig. 4 (left) we plot the optimal generalization error of the perceptron with a Rademacher signal, the state evolution prediction of the generalization error of the GAMP algorithm, together with the error actually achieved by GAMP on one randomly generated instance of the problem. We also compare these to the performance of a standard logistic regression. As expected from existing literature [22, 23] we confirm that in this case the information-theoretic transition appears at a number of samples per dimension $\alpha_{\text{IT}} = 1.249(1)$, while the algorithmic transition is at $\alpha_s = 1.493(1)$. Logistic regression does not seem to be able to match the performance on GAMP in this case.

In Fig. 4 (center) we plot the generalization error for a Gauss-Bernoulli signal with density $\rho = 0.2$. Cases as this one were studied in detail in the context of one-bit compressed sensing [15] and GAMP was found to match the optimal generalization performance with no phase transitions observed, which is confirmed by our analysis. In this case the logistic regression is rather close to the performance of GAMP.

3.3.2 Symmetric Door

The next classification problem, i.e. discrete output rule, we study is the symmetric door function $\varphi(x) = \text{sgn}(|x| - K)$. In this case the generalization error (57) becomes (here again $V \sim \mathcal{N}(0, 1)$)

$$\lim_{n \rightarrow \infty} \mathcal{E}_{\text{gen}}^{\text{opt}} = 1 - \mathbb{E}_V \left[\left\{ \text{erf} \left(\frac{K - \sqrt{q^*} V}{\sqrt{2(\rho - q^*)}} \right) - \text{erf} \left(- \frac{K + \sqrt{q^*} V}{\sqrt{2(\rho - q^*)}} \right) - 1 \right\}^2 \right]. \quad (82)$$

In Fig. 4 (right) we plot the generalization error for $K = 0, 67449$ such that $1/2$ of the outputs are 1 and $1/2$ are -1 . The symmetric door output is an example of function for which the optimal generalization error for $\alpha < \alpha_{\text{IT}} = 1$ (for that specific value of K , see phase diagram in the right panel of Fig. 3) is as bad as if we were guessing randomly. The GAMP algorithm still achieves such a bad generalization until $\alpha_{\text{stab}} = 1.36$, and achieves perfect generalization only for $\alpha > \alpha_s = 1.566(1)$.

Interestingly, labels created from this very simple symmetric door rule seem to be very challenging to learn for general purpose algorithms. We tried to optimize parameters of a two-layers neural network and

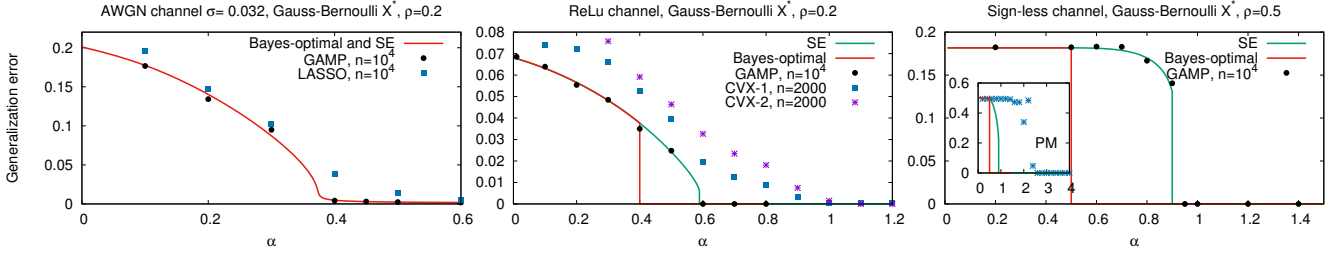


Figure 5: The generalization error for three regression problems is plotted as a function of the number of samples per dimension α . The red line is again the Bayes-optimal generalization error, while the green one shows the (asymptotic) performances of GAMP as predicted by the state evolution (SE), when different. Again, we also show the result of GAMP on a particular instance (black dots) and, in blue, the performance of an out-of-the-box solver. **Left:** White Gaussian noise output and a Gauss-Bernoulli signal. For this choice of noise, there is no sharp transition (as opposed to what happens at smaller noises). The results of a LASSO with fine-tuned ℓ_1 sparsity-enhancing penalty (with [65]) are very close to optimal. **Middle:** Here we analyze a ReLU output function $\varphi(x) = \max(0, x)$, still with a Gauss-Bernoulli signal. Now there is an information-theoretic phase transition at $\alpha = 2\rho = 0.4$, but GAMP requires $\alpha_s = 0.589(1)$ to reach perfect recovery. We show for comparison the results of maximum likelihood estimation performed with CVXPY—a powerful python-embedded language for convex optimization [84]—using two methods that are both amenable to convex optimization: In CVX-1 we use only the non-zero values of \mathbf{Y} , and perform a minimization of the ℓ_1 norm of \mathbf{x} subject to $Y_\mu = \Phi_\mu \cdot \mathbf{x}$ for $\mu \in \{1, \dots, m\}$ such that $Y_\mu \neq 0$, while in CVX-2, we use all the dataset, with the constraint that $Y_\mu = \Phi_\mu \cdot \mathbf{x}$ for $\mu \in \{1, \dots, m\}$ such that $Y_\mu \neq 0$ (as before) and the additional restriction $\Phi_\mu \cdot \mathbf{x} \leq 0$ for $\mu \in \{1, \dots, m\}$ such that $Y_\mu = 0$. In both case, a perfect generalization is obtained only for $\alpha \gtrsim 1$. **Right:** The sign-less output function $\varphi(x) = |x|$. The information-theoretic perfect recovery starts at $\alpha = \rho = 0.5$, but the problem is again harder algorithmically for GAMP that succeeds only above $\alpha_s = 0.90(1)$. Again, the problem appears to be hard for other solvers. In inset, we show the performance for the estimation problem using PhaseMax [85], which is able to learn the rule only using about four times as many measurements than needed information-theoretically.

only managed to get the performances shown in the inset of Fig. 4 (right). It is an interesting theoretical challenge whether a deeper neural network can learn this simple rule from fewer samples.

3.3.3 Linear regression

The additive white Gaussian noise (AWGN) channel, or linear regression, is defined by $\varphi(x, A) = x + \sigma A$ with $A \sim \mathcal{N}(0, 1)$. This models the (noisy) linear regression problem, as well as noisy random linear estimation and compressed sensing. In this case (57) leads to

$$\lim_{n \rightarrow \infty} \mathcal{E}_{\text{gen}}^{\text{opt}} = \rho - q^* + \sigma^2. \quad (83)$$

This result agrees with the generalization error analyzed heuristically in [23] in the limit $\sigma \rightarrow 0$. Fig. 5 (left) depicts the generalization error for this example. The performance of GAMP in this case is very close to the one of LASSO.

3.3.4 Rectified linear unit (ReLU)

In Fig. 5 (center) we analyze the generalization error for the ReLU output function, $\varphi(x) = \max(0, x)$. This channel models the behavior of a single neuron with the rectified linear unit activation [8] widely used in multilayer neural networks. In this case (57) becomes after simple algebra and Gaussian integration by parts

(again $V \sim \mathcal{N}(0, 1)$),

$$\lim_{n \rightarrow \infty} \mathcal{E}_{\text{gen}}^{\text{opt}} = \frac{\rho}{2} - \frac{q^*}{4} \left(1 + \mathbb{E}_V \left[V^2 \text{erf} \left(V \sqrt{\frac{q^*}{2(\rho - q^*)}} \right)^2 \right] \right) - \frac{(\rho - q^*)^{3/2}}{\sqrt{\rho + q^*}} \left(\frac{1}{2\pi} + \frac{q^*}{\rho\pi} \sqrt{\frac{\rho + q^*}{\rho - q^*}} \right). \quad (84)$$

For sparse Gauss-Bernoulli signals in Fig. 5 (center) we observe again the information-theoretic transition to perfect generalization to be distinct from the algorithmic one. At the same time our test with existing algorithms were not able to closely match the performance of GAMP. This hence also remains an interesting benchmark.

3.3.5 Sign-less channel

In Fig. 5 (right) we analyze the generalization error for the sign-less output function where $\varphi(x) = |x|$. This models a situation similar to compressed sensing, except that the sign of the output has been lost. This is a real-valued analog of the phase retrieval problem as discussed in Sec. 3.2.3. In this case the generalization error (57) becomes (again $V \sim \mathcal{N}(0, 1)$)

$$\lim_{n \rightarrow \infty} \mathcal{E}_{\text{gen}}^{\text{opt}} = \rho - \mathbb{E}_V [b(V\sqrt{q^*}, \rho - q^*)^2], \quad (85)$$

where

$$b(x, y) = \sqrt{\frac{2y}{\pi}} e^{-\frac{x^2}{2y}} + \frac{x}{2} \text{erfc} \left(-\frac{x}{\sqrt{2y}} \right) - \frac{x}{2} \left\{ 1 + \text{erf} \left(-\frac{x}{\sqrt{2y}} \right) \right\}. \quad (86)$$

Our comparison with the performance of a state-of-the-art algorithm PhaseMax [85] suggests that also for this simple benchmark there is room for improvement in term of matching the performance of GAMP.

3.3.6 Sigmoid, or logistic regression

Let us also consider an output function with auxiliary randomization. After having generated the classifier \mathbf{X}^* , the teacher randomly associates the label $+1$ to the pattern Φ_μ with probability $f_\lambda(n^{-1/2}\Phi_\mu \cdot \mathbf{X}^*)$, where $f_\lambda(x) = (1 + \exp(-\lambda x))^{-1} \in [0, 1]$ is the sigmoid of parameter $\lambda > 0$, and the label -1 with probability $1 - f_\lambda(n^{-1/2}\Phi_\mu \cdot \mathbf{X}^*)$. One of the (many) possible ways for the teacher to do so is by selecting $\varphi(x, A) = \mathbf{1}(A \leq f_\lambda(x)) - \mathbf{1}(A > f_\lambda(x))$, where $\mathbf{1}(E)$ is the indicator function of the event E . He then generates a stream of uniform random numbers $\mathbf{A} \stackrel{\text{iid}}{\sim} \mathcal{U}_{[0,1]}$ and obtains the labels through (20) (with $\Delta = 0$). Let $V, w \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$. In this setting the error (57) becomes

$$\lim_{n \rightarrow \infty} \mathcal{E}_{\text{gen}}^{\text{opt}} = 2 - 4 \mathbb{E}_V [\{ \mathbb{E}_w f_\lambda(\sqrt{q^*} V + \sqrt{\rho - q^*} w) \}^2]. \quad (87)$$

This formula reduces to (81) when $\lambda \rightarrow \infty$ as it should.

4 Proof of the replica formula by the adaptive interpolation method

We now prove Theorem 1. Our main tool will be an interpolation method recently introduced in [52] and called “adaptive interpolation method”. Here we formulate the method as a direct evolution of the Guerra and Toninelli interpolation method developed in the context of spin glasses [72]. In contrast with the discrete and more pedestrian version of the adaptive interpolation method presented in [52], here we employ a continuous

approach which is more straightforward (see [52] for the links between the discrete and continuous versions of the method) and that has also been recently used in [86] for studying non-symmetric tensor estimation.

We will prove Theorem 1 under the following hypotheses:

- (H1) The support of the prior distribution P_0 is included in $[-S, S]$, for some $S > 0$.
- (H2) φ is a bounded \mathcal{C}^2 function with bounded first and second derivatives w.r.t. its first argument.
- (H3) $(\Phi_{\mu i}) \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$.

These stronger assumptions will then be relaxed in Appendix C to the weaker assumptions (h1)-(h2)-(h3)-(h4) and (h5.a) or (h5.b). Since the observations (20) are equivalent to the rescaled observations

$$\tilde{Y}_\mu := \Delta^{-1/2} Y_\mu = \Delta^{-1/2} \varphi\left(\frac{1}{\sqrt{n}}[\Phi \mathbf{X}^*]_\mu, \mathbf{A}_\mu\right) + Z_\mu, \quad 1 \leq \mu \leq m, \quad (88)$$

the variance Δ of the Gaussian noise can be “incorporated” inside the function φ . Thus, it suffices to prove Theorem 1 for $\Delta = 1$ and we suppose, for the rest of the proof, that we are in this equivalent case.

4.1 Interpolating estimation problem

We introduce an “interpolating estimation problem” that interpolates between the original problem (21) at $t = 0$, $t \in [0, 1]$ being the interpolation parameter, and the two scalar problems described in Sec. 1.3 at $t = 1$ which are analytically tractable. For $t \in (0, 1)$ the interpolating estimation problem is a mixture of the original and scalar problems. This interpolation scheme is inspired from the interpolation paths used by Talagrand to study the perceptron, see [73]. There are two major differences between the “non-planted perceptron” studied by Talagrand, and the “planted perceptron” that we are investigating:

- In the planted case, the presence of a planted solution forces (under small perturbations) the correlations to vanish for all values of the parameters, see [87, 88]. In the non-planted case, proving such decorrelation is much more involved, and is proved only in a limited region of the parameter space, see [73].
- However, in the planted case, there can be arbitrarily many solutions to the state evolution equations (42) (see Remark 21 in [89]), whereas in the region studied by [73], there is only one solution. For this reason, our interpolation method needs to be more sophisticated in order to interpolate with the “right fixed point”.

We fix a sequence $(s_n)_{n \geq 1} \in (0, 1/2]^{\mathbb{N}}$ that converges to 0 as n goes to infinity (s_n will be chosen in Sec. 4.3 below to be equal to $\frac{1}{2}n^{-1/16}$). We define $\mathcal{B}_n := [s_n, 2s_n]^2$. For all $\epsilon = (\epsilon_1, \epsilon_2) \in \mathcal{B}_n$, we consider two continuous “interpolation functions” $q_\epsilon : [0, 1] \rightarrow [0, \rho]$ and $r_\epsilon : [0, 1] \rightarrow [0, r_{\max}]$, where $r_{\max} := 2\alpha \sup_{q \in [0, \rho]} \Psi'_{P_{\text{out}}}(q; \rho) = 2\alpha \Psi'_{P_{\text{out}}}(\rho; \rho)$ (recall that by Proposition 18, $\Psi'_{P_{\text{out}}}$ is non-decreasing). We define also for all $t \in [0, 1]$ and all $\epsilon \in \mathcal{B}_n$

$$R_1(t, \epsilon) := \epsilon_1 + \int_0^t r_\epsilon(v) dv, \quad R_2(t, \epsilon) := \epsilon_2 + \int_0^t q_\epsilon(v) dv. \quad (89)$$

We will be mainly interested in functions r_ϵ, q_ϵ that satisfy some regularity properties. We will use the following definition:

Definition 1 (Regularity). *We say that the families of functions $(q_\epsilon)_{\epsilon \in \mathcal{B}_n}$ and $(r_\epsilon)_{\epsilon \in \mathcal{B}_n}$, taking values respectively in $[0, \rho]$ and $[0, r_{\max}]$, are regular if for all $t \in [0, 1]$ the mapping*

$$R^t : \begin{cases} (s_n, 2s_n)^2 & \rightarrow & R^t((s_n, 2s_n)^2) \\ \epsilon & \mapsto & (R_1(t, \epsilon), R_2(t, \epsilon)) \end{cases} \quad (90)$$

is a C^1 diffeomorphism, whose Jacobian is greater or equal to 1.

Define

$$S_{t,\mu} := \sqrt{\frac{1-t}{n}} [\Phi \mathbf{X}^*]_{\mu} + \sqrt{R_2(t, \epsilon)} V_{\mu} + \sqrt{\rho t - R_2(t, \epsilon) + 2s_n} W_{\mu}^* \quad (91)$$

where $V_{\mu}, W_{\mu}^* \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$. Consider the following observation channels, with two types of observations obtained through

$$\begin{cases} Y_{t,\mu} \sim P_{\text{out}}(\cdot | S_{t,\mu}), & 1 \leq \mu \leq m, \\ Y'_{t,i} = \sqrt{R_1(t, \epsilon)} X_i^* + Z'_i, & 1 \leq i \leq n, \end{cases} \quad (92)$$

where $(Z'_i)_{i=1}^n \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$. We assume that $\mathbf{V} = (V_{\mu})_{\mu=1}^m$ is known. Then the inference problem is to recover both unknowns $\mathbf{W}^* = (W_{\mu}^*)_{\mu=1}^m$ and $\mathbf{X}^* = (X_i^*)_{i=1}^n$ from the knowledge of \mathbf{V} , Φ and the “time-dependent” observations $\mathbf{Y}_t = (Y_{t,\mu})_{\mu=1}^m$ and $\mathbf{Y}'_t = (Y'_{t,i})_{i=1}^n$.

We now understand that $R_1(t, \epsilon)$ appearing in the second set of measurements in (92), and the terms $1-t$, $R_2(t, \epsilon)$ and $\rho t - R_2(t, \epsilon) + 2s_n$ appearing in the first set all play the role of signal-to-noise ratios in the interpolating model, with t giving more and more “power” (or weight) to the scalar inference channels when increasing. Here is the first crucial and novel ingredient of our interpolation scheme. In the classical interpolation method, these signal intensities would all take a trivial form (i.e. would be linear in t) but here, the non-trivial (integral) dependency in t of the intensities through the use of the interpolation functions q and r allows for much more flexibility when choosing the interpolation path. This will allow us to actually choose the “optimal interpolation path” (this will become clear soon).

Define $u_y(x) := \ln P_{\text{out}}(y|x)$ and, with a slight abuse of notations,

$$s_{t,\mu} = s_{t,\mu}(\mathbf{x}, w_{\mu}) := \sqrt{\frac{1-t}{n}} [\Phi \mathbf{x}]_{\mu} + \sqrt{R_2(t, \epsilon)} V_{\mu} + \sqrt{\rho t - R_2(t, \epsilon) + 2s_n} w_{\mu}. \quad (93)$$

We introduce the *interpolating Hamiltonian*

$$\mathcal{H}_{t,\epsilon}(\mathbf{x}, \mathbf{w}; \mathbf{Y}_t, \mathbf{Y}'_t, \Phi, \mathbf{V}) := - \sum_{\mu=1}^m \ln P_{\text{out}}(Y_{t,\mu} | s_{t,\mu}) + \frac{1}{2} \sum_{i=1}^n (Y'_{t,i} - \sqrt{R_1(t, \epsilon)} x_i)^2. \quad (94)$$

The dependence in Φ and \mathbf{V} of the Hamiltonian is through the $(s_{t,\mu})_{\mu=1}^m$. It becomes, when the observations are replaced by their expression (92),

$$\mathcal{H}_{t,\epsilon}(\mathbf{x}, \mathbf{w}; \mathbf{Y}_t, \mathbf{Y}'_t, \Phi, \mathbf{V}) = - \sum_{\mu=1}^m u_{Y_{t,\mu}}(s_{t,\mu}) + \frac{1}{2} \sum_{i=1}^n (\sqrt{R_1(t, \epsilon)} (X_i^* - x_i) + Z'_i)^2. \quad (95)$$

We also introduce the corresponding Gibbs bracket $\langle - \rangle_{n,t,\epsilon}$ which is the expectation operator w.r.t. the (t, ϵ) -dependent posterior distribution of $(\mathbf{X}^*, \mathbf{W}^*)$ given $(\mathbf{Y}_t, \mathbf{Y}'_t, \Phi, \mathbf{V})$. It is defined as

$$\langle g(\mathbf{x}, \mathbf{w}) \rangle_{n,t,\epsilon} := \frac{1}{\mathcal{Z}_{t,\epsilon}(\mathbf{Y}_t, \mathbf{Y}'_t, \Phi, \mathbf{V})} \int dP_0(\mathbf{x}) \mathcal{D}\mathbf{w} g(\mathbf{x}, \mathbf{w}) e^{-\mathcal{H}_{t,\epsilon}(\mathbf{x}, \mathbf{w}; \mathbf{Y}_t, \mathbf{Y}'_t, \Phi, \mathbf{V})}, \quad (96)$$

for every continuous bounded function g on $\mathbb{R}^n \times \mathbb{R}^m$. In (96) $\mathcal{D}\mathbf{w} = (2\pi)^{-m/2} \prod_{\mu=1}^m dw_{\mu} e^{-w_{\mu}^2/2}$ is the m -dimensional standard Gaussian distribution and $\mathcal{Z}_{t,\epsilon}(\mathbf{Y}_t, \mathbf{Y}'_t, \Phi, \mathbf{V})$ is the appropriate normalization (or

partition function):

$$\mathcal{Z}_{t,\epsilon}(\mathbf{Y}_t, \mathbf{Y}'_t, \Phi, \mathbf{V}) := \int dP_0(\mathbf{x}) D\mathbf{w} e^{-\mathcal{H}_{t,\epsilon}(\mathbf{x}, \mathbf{w}; \mathbf{Y}_t, \mathbf{Y}'_t, \Phi, \mathbf{V})}. \quad (97)$$

Finally the *interpolating free entropy* is

$$f_{n,\epsilon}(t) := \frac{1}{n} \mathbb{E} \ln \mathcal{Z}_{t,\epsilon}(\mathbf{Y}_t, \mathbf{Y}'_t, \Phi, \mathbf{V}). \quad (98)$$

Note that the presence of the *perturbation* $\epsilon = (\epsilon_1, \epsilon_2)$ induces only a small change in the free entropy, namely of the order of s_n :

Lemma 1 (Small free entropy variation under perturbation). *For all $\epsilon_1, \epsilon_2 \in [s_n, 2s_n]$,*

$$|f_{n,\epsilon}(0) - f_{n,\epsilon=(0,0)}(0)| \leq C s_n \quad (99)$$

for some constant C that only depends on S, α and φ .

Proof. Let us compute

$$\left| \frac{df_{n,\epsilon}(0)}{d\epsilon_1} \right| = \frac{1}{2} |\mathbb{E} \langle Q \rangle_{n,0,\epsilon}| \leq \frac{S^2}{2}, \quad (100)$$

by hypothesis (H1). Next we compute

$$\left| \frac{df_{n,\epsilon}(0)}{d\epsilon_2} \right| = \frac{1}{2n} \sum_{\mu=1}^m |\mathbb{E} [u'_{Y_0,\mu}(S_{0,\mu}) \langle u'_{Y_0,\mu}(s_{0,\mu}) \rangle_{n,0,\epsilon}]|. \quad (101)$$

This identity is obtained using very similar steps as in Sec. A.5 to which we refer. Under hypothesis (H2) this quantity is bounded by a constant that only depends on α and φ . Then by the mean value theorem we obtain $|f_{n,\epsilon}(0) - f_{n,(0,0)}(0)| \leq C \|\epsilon\| \leq 2\sqrt{2}s_n C$ for some constant C that only depends on S, α and φ . \square

One verifies easily, using the Lemma 1, that for all $\epsilon \in \mathcal{B}_n$

$$\begin{cases} f_{n,\epsilon}(0) &= f_{n,(0,0)}(0) + \mathcal{O}(s_n) \\ &= f_n - \frac{1}{2} + \mathcal{O}(s_n), \\ f_{n,\epsilon}(1) &= \psi_{P_0}(R_1(1, \epsilon)) - \frac{1}{2}(1 + \rho R_1(1, \epsilon)) + \frac{m}{n} \Psi_{P_{\text{out}}}(R_2(1, \epsilon); \rho + 2s_n) \\ &= \psi_{P_0}(\int_0^1 r_\epsilon(t) dt) - \frac{1}{2}(1 + \rho \int_0^1 r_\epsilon(t) dt) + \frac{m}{n} \Psi_{P_{\text{out}}}(\int_0^1 q_\epsilon(t) dt; \rho) + \mathcal{O}(s_n). \end{cases} \quad (102)$$

where f_n is given by (30) and where $\mathcal{O}(s_n)$ denotes a quantity that is bounded by $C s_n$ for some constant $C > 0$ that only depends on S, φ and α . For the last equality we used Proposition 17 in Appendix B.1 which says that ψ_{P_0} is $\frac{\rho}{2}$ -Lipschitz and, similarly to Proposition 18, it is not difficult to verify that $(q_1, q_2) \mapsto \Psi_{P_{\text{out}}}(q_1; q_2)$ is \mathcal{C}^1 on the compact set $\{(q_1, q_2) \mid 0 \leq q_1 \leq q_2 \leq \rho + 1\}$ and is thus Lipschitz. We emphasize a crucial property of the interpolating model: It is such that at $t = 0$ we recover the original model and thus $f_{n,\epsilon}(0) \approx f_n - 1/2$ (the trivial constant comes from the purely noisy measurements of the second channel in (92)), while at $t = 1$ we have the two scalar inference channels and thus the associated terms ψ_{P_0} and $\Psi_{P_{\text{out}}}$ discussed in Sec. 1.3 appear in $f_{n,\epsilon}(1)$. These are precisely the terms appearing in the potential (41).

4.2 Free entropy variation along the interpolation path

From the understanding of the previous section, it is at this stage very natural to evaluate the variation of free entropy along the interpolation path, which allows to “compare” the original and purely scalar models

thanks to the identity

$$f_n = f_{n,\epsilon}(0) + \frac{1}{2} + \mathcal{O}(s_n) = f_{n,\epsilon}(1) - \int_0^1 \frac{df_{n,\epsilon}(t)}{dt} dt + \frac{1}{2} + \mathcal{O}(s_n), \quad (103)$$

where the first equality follows from (102). As discussed above, part of the potential (41) appears in $f_{n,\epsilon}(1)$. If the interpolation is properly done, the missing terms required to obtain the potential on the r.h.s. of (103) should naturally appear. Then by choosing the optimal interpolation path thanks to the non-trivial snr dependencies in t (i.e. by selecting the proper interpolating functions q and r), we will be able to show the equality between the replica formula and the free entropy $\lim_{n \rightarrow \infty} f_n$.

We thus now compute the t -derivative of the free entropy along the interpolation path (see Appendix A.5 for the proof). Let $u'_y(x)$ be the derivative (w.r.t. x) of $u_y(x)$. Then we have the following.

Proposition 3 (Free entropy variation). *The derivative of the free entropy (98) verifies, for all $\epsilon \in \mathcal{B}_n$ and all $t \in (0, 1)$*

$$\frac{df_{n,\epsilon}(t)}{dt} = -\frac{1}{2} \mathbb{E} \left\langle \left(\frac{1}{n} \sum_{\mu=1}^m u'_{Y_{t,\mu}}(S_{t,\mu}) u'_{Y_{t,\mu}}(s_{t,\mu}) - r_\epsilon(t) \right) (Q - q_\epsilon(t)) \right\rangle_{n,t,\epsilon} + \frac{r_\epsilon(t)}{2} (q_\epsilon(t) - \rho) + \mathcal{O}_n(1), \quad (104)$$

where $\mathcal{O}_n(1)$ is a quantity that goes to 0 in the $n, m \rightarrow \infty$ limit, uniformly in $t \in (0, 1)$, $\epsilon \in \mathcal{B}_n$ and uniformly in the choice of the functions q_ϵ and r_ϵ . The overlap is

$$Q_n = Q := \frac{1}{n} \mathbf{X}^* \cdot \mathbf{x} = \frac{1}{n} \sum_{i=1}^n X_i^* x_i \quad (105)$$

where \mathbf{x} is a sample from the posterior of model (92) associated with the Gibbs bracket $\langle - \rangle_{n,t,\epsilon}$, see (96).

4.3 Overlap concentration and fundamental sum rule

The next lemma plays a key role in our proof. Essentially it states that the overlap concentrates around its mean, a behavior called “replica symmetric” in statistical physics. Similar results have been obtained in the context of the analysis of spin glasses [53, 73]. Here we use a formulation tailored to Bayesian inference problems as developed in the context of LDPC codes, random linear estimation [47] and Nishimori symmetric spin glasses [88, 90, 91].

Proposition 4 (Overlap concentration). *Assume that the interpolation functions $(q_\epsilon), (r_\epsilon)$ are regular, see Definition 1. Let $s_n = \frac{1}{2} n^{-1/16}$ for all $n \geq 1$. Under assumptions (H1), (H2) and (H3) there exists a constant $C(\varphi, S, \alpha)$ that depends only on S, φ and α such that*

$$\frac{1}{s_n^2} \int_{\mathcal{B}_n} d\epsilon \int_0^1 dt \mathbb{E} \langle (Q - \mathbb{E} \langle Q \rangle_{n,t,\epsilon})^2 \rangle_{n,t,\epsilon} \leq \frac{C(\varphi, S, \alpha)}{n^{1/8}}. \quad (106)$$

Proposition 4 follows from Proposition 29 proved in Appendix E.2, combined with (240) and Fubini’s theorem. Note from (102) and (41) that the second term appearing in (104) is precisely the missing one that is required in order to obtain the expression of the potential on the r.h.s. of (103). Thus in order to prove Theorem 1 we would like to “cancel” the Gibbs bracket in (104), which is the so called *remainder* (once integrated over t). This is made possible thanks to the adaptive interpolating functions.

One possible way to cancel the remainder is to choose $q_\epsilon(t) = \mathbb{E} \langle Q \rangle_{n,t,\epsilon}$, which is approximately equal to Q because it concentrates by Proposition 4. However, $\mathbb{E} \langle Q \rangle_{n,t,\epsilon}$ depends on $\int_0^t q_\epsilon(v) dv$ (and on $t, \int_0^t r_\epsilon(v) dv$ and ϵ too). The equation $q_\epsilon(t) = \mathbb{E} \langle Q \rangle_{n,t,\epsilon}$ is therefore a first order differential equation over $t \mapsto \int_0^t q_\epsilon(v) dv$.

We will see in details in Sec. 4.4 that it possesses a solution, but for the moment we just assume it exists in order to derive the following *fundamental sum rule*, which is a core identity in the proof scheme:

Proposition 5 (Fundamental sum rule). *Assume that the interpolation functions (q_ϵ) and (r_ϵ) are regular (see Definition 1). Assume that for all $t \in [0, 1]$ and $\epsilon \in \mathcal{B}_n$ we have $q_\epsilon(t) = \mathbb{E}\langle Q \rangle_{n,t,\epsilon}$. Then*

$$f_n = \frac{1}{s_n^2} \int_{\mathcal{B}_n} \left\{ \psi_{P_0} \left(\int_0^1 r_\epsilon(t) dt \right) + \alpha \Psi_{P_{\text{out}}} \left(\int_0^1 q_\epsilon(t) dt; \rho \right) - \frac{1}{2} \int_0^1 q_\epsilon(t) r_\epsilon(t) dt \right\} d\epsilon + \mathcal{O}_n(1), \quad (107)$$

where $\mathcal{O}_n(1)$ denotes a quantity that goes to 0 as $n \rightarrow \infty$ uniformly w.r.t. the choice of the interpolation functions.

Proof. By the Cauchy-Schwarz inequality

$$\begin{aligned} & \left(\frac{1}{s_n^2} \int_{\mathcal{B}_n} d\epsilon \int_0^1 dt \mathbb{E} \left\langle \left(\frac{1}{n} \sum_{\mu=1}^m u'_{Y_{t,\mu}}(S_{t,\mu}) u'_{Y_{t,\mu}}(s_{t,\mu}) - r_\epsilon(t) \right) (Q - q_\epsilon(t)) \right\rangle_{n,t,\epsilon} \right)^2 \\ & \leq \frac{1}{s_n^2} \int_{\mathcal{B}_n} d\epsilon \int_0^1 dt \mathbb{E} \left\langle \left(\frac{1}{n} \sum_{\mu=1}^m u'_{Y_{t,\mu}}(S_{t,\mu}) u'_{Y_{t,\mu}}(s_{t,\mu}) - r_\epsilon(t) \right)^2 \right\rangle_{n,t,\epsilon} \times \frac{1}{s_n^2} \int_{\mathcal{B}_n} d\epsilon \int_0^1 dt \mathbb{E} \langle (Q - q_\epsilon(t))^2 \rangle_{n,t,\epsilon}. \end{aligned}$$

The first term of this product is bounded by some constant $C(\varphi, \alpha)$ that only depend on φ and α , see Appendix A.6. The second term is bounded by $C(\varphi, S, \alpha) n^{-1/8}$ by Proposition 4, since we assumed that for all $\epsilon \in \mathcal{B}_n$ and all $t \in [0, 1]$ we have $q_\epsilon(t) = \mathbb{E}\langle Q \rangle_{n,t,\epsilon}$. We have therefore

$$\left| \frac{1}{s_n^2} \int_{\mathcal{B}_n} d\epsilon \int_0^1 dt \mathbb{E} \left\langle \left(\frac{1}{n} \sum_{\mu=1}^m u'_{Y_{t,\mu}}(S_{t,\mu}) u'_{Y_{t,\mu}}(s_{t,\mu}) - r_\epsilon(t) \right) (Q - q_\epsilon(t)) \right\rangle_{n,t,\epsilon} \right| \leq \frac{C(\varphi, S, \alpha)}{n^{1/16}}.$$

Therefore from (104)

$$\frac{1}{s_n^2} \int_{\mathcal{B}_n} d\epsilon \int_0^1 dt \frac{df_{n,\epsilon}(t)}{dt} = \frac{1}{2s_n^2} \int_{\mathcal{B}_n} d\epsilon \int_0^1 dt \{ q_\epsilon(t) r_\epsilon(t) - r_\epsilon(t) \rho \} + \mathcal{O}_n(1) + \mathcal{O}(n^{-1/16}). \quad (108)$$

Here the small terms are going to 0 both uniformly w.r.t. to the choice of q_ϵ and r_ϵ . When replacing (108) in (103) and combining it with (102) we reach the claimed identity (107), but up to the fact that $\Psi_{P_{\text{out}}}(\int_0^1 q_\epsilon(t) dt; \rho)$ is multiplied by m/n instead of α . Recalling that $m/n \rightarrow \alpha$ as $m, n \rightarrow \infty$ allows to finish the argument (notice that $\Psi_{P_{\text{out}}}$ is continuous and hence bounded on $[0, \rho]$, see Proposition 18). \square

We are now ready to prove matching bounds.

4.4 Lower and upper matching bounds

We now possess all the necessary tools to prove Theorem 1 in three steps.

- (i) We prove that, under assumptions (H1), (H2) and (H3), $\lim_{n \rightarrow \infty} f_n = \sup_{r \geq 0} \inf_{q \in [0, \rho]} f_{\text{RS}}(q, r)$.
- (ii) Under hypothesis (H2), the function $\Psi_{P_{\text{out}}}$ is convex, Lipschitz and non-decreasing (Proposition 19). We thus apply Corollary 8 of Appendix D to get $\sup_{r \geq 0} \inf_{q \in [0, \rho]} f_{\text{RS}}(q, r) = \sup_{q \in [0, \rho]} \inf_{r \geq 0} f_{\text{RS}}(q, r)$. We then deduce from (i) that $\lim_{n \rightarrow \infty} f_n = \sup_{q \in [0, \rho]} \inf_{r \geq 0} f_{\text{RS}}(q, r)$ under (H1)-(H2)-(H3).
- (iii) Finally, the approximation arguments given in Appendix C permit to relax (H1)-(H2) to the weaker hypotheses (h1)-(h2) and allow to replace the Gaussian assumption (H3) on Φ by (h1)-(h3)-(h4). The fact that for discrete channels the Gaussian noise can then be removed, allowing to replace (h5.a) (i.e. $\Delta > 0$ treated until here) to (h5.b) (i.e. $\Delta = 0$ and φ takes values in \mathbb{N}), is proven in Sec. C.3. This

proves the first equality of Theorem 1. The last equality in (43) and the remaining part of Theorem 1 follow then from Lemma 23.

It thus remains to tackle (i), but before that we need a definition. For $t \in [0, 1]$ and $\epsilon \in \mathcal{B}_n$, we write $R^t(\epsilon) = (R_1(t, \epsilon), R_2(t, \epsilon))$. The quantity $\mathbb{E}\langle Q \rangle_{n,t,\epsilon}$ is a function of $n, t, R^t(\epsilon)$ that we write $\mathbb{E}\langle Q \rangle_{n,t,\epsilon} = F_n(t, R^t(\epsilon))$, where F_n is a function defined on

$$D_n := \left\{ (t, r_1, r_2) \in [0, 1] \times \mathbb{R}_+ \times \mathbb{R}_+ \mid r_2 \leq \rho t + 2s_n \right\}. \quad (109)$$

The following proposition, proven in Appendix A.7, will be useful.

Proposition 6. F_n is a continuous function from D_n to $[0, \rho]$. Let D_n° denotes the interior of D_n . F_n admits partial derivatives with respect to its second and third argument on D_n° . These partial derivatives are both continuous and non-negative on D_n° .

Let us now start with the lower bound.

4.4.1 Lower bound

Proposition 7 (Lower bound). *The free entropy (29) verifies*

$$\liminf_{n \rightarrow \infty} f_n \geq \sup_{r \geq 0} \inf_{q \in [0, \rho]} f_{RS}(q, r). \quad (110)$$

Proof. We consider, for $(\epsilon_1, \epsilon_2) \in \mathcal{B}_n$ and a fixed value $r \in [0, r_{\max}]$, the following 1st order differential equation:

$$y(0) = (\epsilon_1, \epsilon_2) \quad \text{and} \quad \forall t \in [0, 1], \quad y'(t) = (r, F_n(t, y(t))). \quad (111)$$

By the Cauchy-Lipschitz Theorem (see for instance Theorem 3.1 in Chapter V from [92]) this equation admits a (unique) solution that we write $y(\cdot, \epsilon) = (y_1(\cdot, \epsilon), y_2(\cdot, \epsilon))$. The hypotheses of the Cauchy-Lipschitz Theorem are verified, because of Proposition 6. We define then, for all $t \in [0, 1]$,

$$r_\epsilon(t) = y'_1(t, \epsilon) = r \quad \text{and} \quad q_\epsilon(t) = y'_2(t, \epsilon) = F_n(t, y(t, \epsilon)) \in [0, \rho].$$

We have therefore $R_1(t, \epsilon) = \epsilon_1 + \int_0^t y'_1(s, \epsilon) ds = y_1(t, \epsilon)$ and similarly $R_2(t, \epsilon) = y_2(t, \epsilon)$. We obtain that for all $t \in [0, 1]$,

$$q_\epsilon(t) = F_n(t, y(t, \epsilon)) = F_n(t, (R_1(t, \epsilon), R_2(t, \epsilon))) = \mathbb{E}\langle Q \rangle_{n,t,\epsilon}.$$

Let us show now that the functions (q_ϵ) and (r_ϵ) are regular (see Definition 1). Let $t \in [0, 1]$. The function $R^t : \epsilon \mapsto (R_1(t, \epsilon), R_2(t, \epsilon)) = y(t, \epsilon)$ is the flow of (111) and is thus injective (by unicity of the solution) and \mathcal{C}^1 because of the regularity properties (see Proposition 6) of F_n . The Jacobian of the flow is given by the Liouville formula (see Corollary 3.1 in Chapter V from [92]):

$$\det \left(\frac{\partial R^t}{\partial \epsilon}(\epsilon) \right) = \exp \left(\int_0^t dv \frac{\partial F_n}{\partial y_2}(v, y(v, \epsilon)) \right) \geq 1,$$

because by Proposition 6 we have $\partial_{y_2} F_n \geq 0$. We obtain (by the local inversion Theorem) that R^t is a \mathcal{C}^1 diffeomorphism, and since its Jacobian is greater or equal to 1 the functions (q_ϵ) and (r_ϵ) are regular.

We have seen that for all $\epsilon \in \mathcal{B}_n$ and all $t \in [0, 1]$, $q_\epsilon(t) = \mathbb{E}\langle Q \rangle_{t,n,\epsilon}$, so we can apply Proposition 5 to get

$$\begin{aligned} f_n &= \frac{1}{s_n^2} \int_{\mathcal{B}_n} \left\{ \psi_{P_0}(r) + \alpha \Psi_{P_{\text{out}}} \left(\int_0^1 q_\epsilon(t) dt; \rho \right) - \frac{r}{2} \int_0^1 q_\epsilon(t) dt \right\} d\epsilon + \mathcal{O}_n(1) \\ &= \frac{1}{s_n^2} \int_{\mathcal{B}_n} f_{\text{RS}} \left(\int_0^1 q_\epsilon(t) dt, r \right) d\epsilon + \mathcal{O}_n(1) \\ &\geq \inf_{q \in [0, \rho]} f_{\text{RS}}(q, r) + \mathcal{O}_n(1) \end{aligned}$$

and thus $\liminf_{n \rightarrow \infty} f_n \geq \inf_{q \in [0, \rho]} f_{\text{RS}}(q, r)$. This is true for all $r \in [0, r_{\text{max}}]$ so we get

$$\liminf_{n \rightarrow \infty} f_n \geq \sup_{r \in [0, r_{\text{max}}]} \inf_{q \in [0, \rho]} f_{\text{RS}}(q, r). \quad (112)$$

Let $r \geq r_{\text{max}}$. We have for all $q \in [0, \rho]$, $\partial_q f_{\text{RS}}(q, r) = \alpha \Psi'_{P_{\text{out}}}(q) - \frac{r}{2} \leq 0$, because $r \geq r_{\text{max}} \geq 2\alpha \Psi'_{P_{\text{out}}}(q)$. Therefore for all $r \geq r_{\text{max}}$, $\inf_{q \in [0, \rho]} f_{\text{RS}}(q, r) = f_{\text{RS}}(\rho, r)$ and

$$\frac{\partial}{\partial r} \inf_{q \in [0, \rho]} f_{\text{RS}}(q, r) = \frac{\partial}{\partial r} f_{\text{RS}}(\rho, r) = \psi'_{P_0}(r) - \frac{\rho}{2} \leq 0,$$

because by Proposition 17, ψ_{P_0} is $\frac{\rho}{2}$ -Lipschitz. The function $r \mapsto \inf_{q \in [0, \rho]} f_{\text{RS}}(q, r)$ is therefore non-increasing on $[r_{\text{max}}, +\infty)$. Going back to (112), we conclude

$$\liminf_{n \rightarrow \infty} f_n \geq \sup_{r \in [0, r_{\text{max}}]} \inf_{q \in [0, \rho]} f_{\text{RS}}(q, r) = \sup_{r \geq 0} \inf_{q \in [0, \rho]} f_{\text{RS}}(q, r). \quad (113)$$

□

4.4.2 Upper bound

Proposition 8 (Upper bound). *The free entropy (29) verifies*

$$\limsup_{n \rightarrow \infty} f_n \leq \sup_{r \geq 0} \inf_{q \in [0, \rho]} f_{\text{RS}}(q, r). \quad (114)$$

Proof. We consider, for $(\epsilon_1, \epsilon_2) \in \mathcal{B}_n$, the following order-1 system of differential equations:

$$y(0) = (\epsilon_1, \epsilon_2) \quad \text{and} \quad \forall t \in [0, 1], \quad y'(t) = \begin{pmatrix} 2\alpha \Psi'_{P_{\text{out}}}(F_n(t, y(t))) \\ F_n(t, y(t)) \end{pmatrix}. \quad (115)$$

By Proposition 18 the function $\Psi'_{P_{\text{out}}}$ is \mathcal{C}^1 and takes values in $[0, r_{\text{max}}]$. By Proposition 6, the function F_n is continuous, bounded and admits partial derivatives w.r.t. its second and third arguments, that are continuous. We can therefore apply the Cauchy-Lipschitz Theorem as in the proof of Proposition 7: The equation (115) admits a (unique) solution that we write $y(\cdot, \epsilon) = (y_1(\cdot, \epsilon), y_2(\cdot, \epsilon))$. We define then, for all $t \in [0, 1]$,

$$r_\epsilon(t) = y'_1(t, \epsilon) = 2\alpha \Psi'_{P_{\text{out}}}(F_n(t, y(t, \epsilon))) \in [0, r_{\text{max}}] \quad \text{and} \quad q_\epsilon(t) = y'_2(t, \epsilon) = F_n(t, y(t, \epsilon)) \in [0, \rho].$$

We have therefore $R_1(t, \epsilon) = \epsilon_1 + \int_0^t y'_1(s, \epsilon) ds = y_1(t, \epsilon)$ and similarly $R_2(t, \epsilon) = y_2(t, \epsilon)$. We obtain that for all $t \in [0, 1]$,

$$q_\epsilon(t) = F_n(t, y(t, \epsilon)) = F_n(t, (R_1(t, \epsilon), R_2(t, \epsilon))) = \mathbb{E}\langle Q \rangle_{n,t,\epsilon}.$$

Let us show now that the functions (q_ϵ) and (r_ϵ) are regular (see Definition 1). Let $t \in [0, 1]$. The function $R^t : \epsilon \mapsto (R_1(t, \epsilon), R_2(t, \epsilon)) = y(t, \epsilon)$ is the flow of (115) and is thus injective and \mathcal{C}^1 because of the regularity

properties (see Proposition 6) of F_n . The Jacobian of the flow is again given by the Liouville formula:

$$\det\left(\frac{\partial R^t}{\partial \epsilon}(\epsilon)\right) = \exp\left(\int_0^t dv 2\alpha \frac{\partial F_n}{\partial y_1}(v, y(v, \epsilon)) \Psi''_{P_{\text{out}}}(F_n(v, y(v, \epsilon))) + \int_0^t dv \frac{\partial F_n}{\partial y_2}(v, y(v, \epsilon))\right) \geq 1,$$

because by Proposition 6, $\frac{\partial F_n}{\partial y_1}$ and $\frac{\partial F_n}{\partial y_2}$ are both non negative and since $\Psi_{P_{\text{out}}}$ is convex (see Proposition 18), we have also $\Psi''_{P_{\text{out}}} \geq 0$. We obtain (by the local inversion Theorem) that R^t is a \mathcal{C}^1 diffeomorphism. Its Jacobian is greater or equal to 1, and the functions (q_ϵ) and (r_ϵ) are therefore regular.

We have seen that for all $\epsilon \in \mathcal{B}_n$ and all $t \in [0, 1]$, $q_\epsilon(t) = \mathbb{E}\langle Q \rangle_{t, n, \epsilon}$, so we can apply Proposition 5 to get

$$\begin{aligned} f_n &= \frac{1}{s_n^2} \int_{\mathcal{B}_n} \left\{ \psi_{P_0}\left(\int_0^1 r_\epsilon(t) dt\right) + \alpha \Psi_{P_{\text{out}}}\left(\int_0^1 q_\epsilon(t) dt; \rho\right) - \frac{1}{2} \int_0^1 q_\epsilon(t) r_\epsilon(t) dt \right\} d\epsilon + \mathcal{O}_n(1) \\ &\leq \frac{1}{s_n^2} \int_{\mathcal{B}_n} \int_0^1 \left\{ \psi_{P_0}(r_\epsilon(t)) + \alpha \Psi_{P_{\text{out}}}(q_\epsilon(t); \rho) - \frac{1}{2} q_\epsilon(t) r_\epsilon(t) \right\} dt d\epsilon + \mathcal{O}_n(1) \end{aligned} \quad (116)$$

by Jensen's inequality, because by Propositions 17 and 18 the functions ψ_{P_0} and $\Psi_{P_{\text{out}}}$ are convex.

Let us fix $\epsilon \in \mathcal{B}_n$ and $t \in [0, 1]$. By definition of r_ϵ and q_ϵ , we have

$$r_\epsilon(t) = 2\alpha \Psi'_{P_{\text{out}}}(F_n(t, y(t, \epsilon))) = 2\alpha \Psi'_{P_{\text{out}}}(q_\epsilon(t)). \quad (117)$$

The function $g : q \in [0, \rho] \mapsto 2\alpha \Psi_{P_{\text{out}}}(q; \rho) - r_\epsilon(t)q$ is convex by Proposition 18. By equation (117) above, we see that $g'(q_\epsilon(t)) = 0$ and therefore:

$$\alpha \Psi_{P_{\text{out}}}(q_\epsilon(t); \rho) - \frac{1}{2} q_\epsilon(t) r_\epsilon(t) = \inf_{q \in [0, \rho]} \left\{ \alpha \Psi_{P_{\text{out}}}(q; \rho) - \frac{1}{2} q r_\epsilon(t) \right\}.$$

This holds for all $\epsilon \in \mathcal{B}_n$ and all $t \in [0, 1]$. Plugging this back in (116), we get:

$$\begin{aligned} f_n &\leq \frac{1}{s_n^2} \int_{\mathcal{B}_n} \int_0^1 \inf_{q \in [0, \rho]} \left\{ \psi_{P_0}(r_\epsilon(t)) + \alpha \Psi_{P_{\text{out}}}(q; \rho) - \frac{1}{2} q r_\epsilon(t) \right\} dt d\epsilon + \mathcal{O}_n(1) \\ &\leq \sup_{r \geq 0} \inf_{q \in [0, \rho]} \left\{ \psi_{P_0}(r) + \alpha \Psi_{P_{\text{out}}}(q; \rho) - \frac{1}{2} q r \right\} + \mathcal{O}_n(1). \end{aligned}$$

This proves Proposition 8. □

From the arguments given at the beginning of the section, this ends the proof of Theorem 1.

5 Proofs of the limits of optimal errors

5.1 Optimal generalization error: Proof of Theorem 4

5.1.1 Formal derivation and proof idea: A teacher-student scenario with side information

Before proving Theorem 4 rigorously, we find useful to provide a conceptual framework allowing to formally derive the generalization error (a framework that will actually serve as a basis for the rigorous derivation presented in the next section). In order to obtain the (generalized) optimal generalization error, we need first to assume that the new “test labels” are also observed by the student in the teacher-student scenario of Sec. 1.2 but with a very low signal-to-noise ratio. The presence of this side information will allow us to use the I-MMSE relation (Proposition 13) to obtain the generalization error when small, but non-zero, information about the test labels is known by the student. Then, by formally taking the limit of vanishing side information

on the resulting expression (and *assuming* that the large n and vanishing side information limits commute), we will recover the generalization error. We thus now introduce the following “train-test” observation model.

The set of patterns and labels are divided into two sets by the teacher: The *training set* \mathcal{S}^{tr} of size m that will be used as the main source of information by the student in order to then generalize, and the *test set* \mathcal{S}^{te} of size $m' = \epsilon n$ that will be used by the teacher in order to evaluate the performance of the student, but also by the student as small additional side information. Let us be more precise: The teacher gives to the student both the patterns and associated labels of the training set, namely $\mathcal{S}^{\text{tr}} := \{(Y_\mu; \Phi_\mu)\}_{\mu=1}^m$ (recall the labels are given by (20), (21)). For the test set, the test patterns to classify are given to the student but the associated labels are (almost) not: Let $\epsilon, \lambda \geq 0$. Instead of the test labels $\{\tilde{Y}_\mu\}_{\mu=1}^{m'}$ (that should be totally unknown to the student in the ideal setting), what is given to the student is

$$U_\mu = \sqrt{\lambda} Y'_\mu + Z'_\mu, \quad \text{for } 1 \leq \mu \leq m' = \epsilon n, \quad (118)$$

where $Z'_\mu \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$, and Y'_μ is given by

$$Y'_\mu = f(\tilde{Y}_\mu), \quad \tilde{Y}_\mu \sim P_{\text{out}}\left(\cdot \mid \frac{\Phi'_\mu \cdot \mathbf{X}^*}{\sqrt{n}}\right), \quad (119)$$

where $f : \mathbb{R} \rightarrow \mathbb{R}$ is a continuous bounded function and $\Phi'_\mu \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \mathbf{I}_n)$ independently of everything else. We will first prove Theorem 4 for continuous bounded functions f , and then relax this at the end of the proof. The test set given to the student, in addition of the training set, is $\mathcal{S}^{\text{te}} = \mathcal{S}^{\text{te}}(\lambda, \epsilon) := \{(U_\mu = \sqrt{\lambda} Y'_\mu + Z'_\mu; \Phi'_\mu)\}_{\mu=1}^{\epsilon n}$ where λ is *typically very small*. Indeed, we are particularly interested in the case $\lambda, \epsilon \rightarrow 0$ when the student has *no* information about the test labels, which is the ideal setting we want to study. But in order to employ the I-MMSE relation we consider instead very small $\lambda > 0$.

The learning of the classifier \mathbf{X}^* given \mathcal{S}^{tr} and \mathcal{S}^{te} is a slight extension of model (20). Define $\mathbf{Y}' = (Y'_\mu)_{\mu=1}^{m'}$ as the vector of test labels (before they are corrupted by additional noise through (118)). Then the (generalized) optimal generalization error *with side information* (i.e. at $\lambda, \epsilon > 0$) in this “train-test” observation model is

$$\mathcal{E}_{f,n}^{\text{side}}(\lambda, \epsilon) := \min_{\hat{\mathbf{Y}}'} \frac{1}{\epsilon n} \mathbb{E}[\|\mathbf{Y}' - \hat{\mathbf{Y}}'(\mathcal{S}^{\text{te}}, \mathcal{S}^{\text{tr}})\|^2] = \frac{1}{\epsilon n} \mathbb{E}[\|\mathbf{Y}' - \mathbb{E}[\mathbf{Y}' | \mathcal{S}^{\text{te}}, \mathcal{S}^{\text{tr}}]\|^2]. \quad (120)$$

The “true” generalization error (60) is recovered by defining instead $\mathcal{S}^{\text{te}} = \Phi'$ or equivalently letting $\lambda, \epsilon \rightarrow 0$, i.e. when only \mathcal{S}^{tr} and the test patterns are given to the student: $\lim_{\lambda, \epsilon \rightarrow 0} \mathcal{E}_{f,n}^{\text{side}}(\lambda, \epsilon) = \mathcal{E}_{f,n}$. Note that the f function only plays a role in the test set, while the labels of the training data are generated through the “pure model” (20), (21).

From there one can use the I-MMSE relation of Proposition 13 in order to formally compute the limiting $n \rightarrow \infty$ expression of (120). Indeed,

$$\frac{\partial}{\partial \lambda} \frac{1}{n} I(\mathbf{Y}'; \sqrt{\lambda} \mathbf{Y}' + \mathbf{Z}' | \mathbf{Y}, \Phi, \Phi') = \frac{\epsilon}{2} \mathcal{E}_{f,n}^{\text{side}}(\lambda, \epsilon). \quad (121)$$

Fortunately, by a straightforward extension of the interpolation method presented in Sec. 4 one can generalize Theorem 1 to take into account this additional side information and access this mutual information (see the end of the section for the proof):

Lemma 2. *For all $\epsilon, \lambda \geq 0$ we have*

$$\frac{1}{n} I(\mathbf{Y}'; \sqrt{\lambda} \mathbf{Y}' + \mathbf{Z}' | \mathbf{Y}, \Phi, \Phi') \xrightarrow{n \rightarrow \infty} \inf_{q \in [0, \rho]} \sup_{r \geq 0} \tilde{i}_{\text{RS}}(q, r, \lambda) - i_\infty, \quad (122)$$

where i_∞ is given by Corollary 2 and

$$\tilde{i}_{\text{RS}}(q, r, \lambda) := i_{\text{RS}}(q, r) + \epsilon I(f(Y^{(q)}); \sqrt{\lambda} f(Y^{(q)}) + Z' | V) \quad (123)$$

$$= I_{P_0}(r) + \alpha \mathcal{I}_{P_{\text{out}}}(q; \rho) + \epsilon I(f(Y^{(q)}); \sqrt{\lambda} f(Y^{(q)}) + Z' | V) - \frac{r}{2}(\rho - q). \quad (124)$$

Recall that $Y^{(q)}$ is sampled from the “second scalar channel” (39): $Y^{(q)} \sim P_{\text{out}}(\cdot | \sqrt{q} V + \sqrt{\rho - q} W^*)$, where $V, W^* \stackrel{iid}{\sim} \mathcal{N}(0, 1)$.

Define the following MMSE function:

$$M_f : (\lambda, q) \mapsto \text{MMSE}(f(Y^{(q)}) | \sqrt{\lambda} f(Y^{(q)}) + Z', V). \quad (125)$$

By concavity arguments detailed in the next section, we have almost everywhere

$$\lim_{n \rightarrow \infty} \frac{\partial}{\partial \lambda} \frac{1}{n} I(\mathbf{Y}'; \sqrt{\lambda} \mathbf{Y}' + \mathbf{Z}' | \mathbf{Y}, \Phi, \Phi') = \frac{\partial}{\partial \lambda} \lim_{n \rightarrow \infty} \frac{1}{n} I(\mathbf{Y}'; \sqrt{\lambda} \mathbf{Y}' + \mathbf{Z}' | \mathbf{Y}, \Phi, \Phi') = \frac{\partial}{\partial \lambda} \inf_{q \in [0, \rho]} \sup_{r \geq 0} \tilde{i}_{\text{RS}}(q, r, \lambda)$$

using Lemma 2 for the last equality. Assuming $\partial_\lambda \inf_{q \in [0, \rho]} \sup_{r \geq 0} \tilde{i}_{\text{RS}}(q, r, \lambda) = \partial_\lambda \tilde{i}_{\text{RS}}(q, r, \lambda)|_{(q_\lambda^*, r_\lambda^*)}$, where $(q_\lambda^*, r_\lambda^*)$ is an optimal couple, (121) and the last identity combined lead to

$$\frac{\epsilon}{2} \lim_{n \rightarrow \infty} \mathcal{E}_{f,n}^{\text{side}}(\lambda, \epsilon) = \epsilon \frac{\partial}{\partial \lambda} I(f(Y^{(q)}); \sqrt{\lambda} f(Y^{(q)}) + Z' | V) \Big|_{q_\lambda^*} = \frac{\epsilon}{2} M_f(\lambda, q_\lambda^*), \quad (126)$$

using again the I-MMSE relation for the last equality. Thus $\lim_{n \rightarrow \infty} \mathcal{E}_{f,n}^{\text{side}}(\lambda, \epsilon) = M_f(\lambda, q_\lambda^*)$.

A formal calculation of the vanishing side information limit of $M_f(\lambda, q_\lambda^*)$ gives back $\mathcal{E}_f(q^*(\alpha))$ (recall (61) and $q^*(\alpha)$ is the optimizer of the replica-symmetric formula (43)), so that $\lim_{\lambda, \epsilon \rightarrow 0} \lim_{n \rightarrow \infty} \mathcal{E}_{f,n}^{\text{side}}(\lambda, \epsilon) = \mathcal{E}_f(q^*(\alpha))$. It is very natural to believe that the vanishing side information limit of $\lim_{n \rightarrow \infty} \mathcal{E}_{f,n}^{\text{side}}(\lambda, \epsilon)$ should give back the true asymptotic generalization error. So if one could justify the commutation of limits

$$\lim_{\lambda, \epsilon \rightarrow 0} \lim_{n \rightarrow \infty} \mathcal{E}_{f,n}^{\text{side}}(\lambda, \epsilon) = \lim_{n \rightarrow \infty} \lim_{\lambda, \epsilon \rightarrow 0} \mathcal{E}_{f,n}^{\text{side}}(\lambda, \epsilon) = \lim_{n \rightarrow \infty} \mathcal{E}_{f,n}$$

this would end the proof. We prove this point in the next section.

Proof of Lemma 2: Extending the interpolation method presented in Sec. 4, one can generalize Theorem 1 to take into account this additional side information. This gives directly

$$\frac{1}{n} I(\mathbf{X}^*; \mathbf{Y}, \sqrt{\lambda} \mathbf{Y}' + \mathbf{Z}' | \Phi, \Phi') \xrightarrow{n \rightarrow \infty} \tilde{I}_\infty(\alpha, \epsilon, \lambda) := \inf_{q \in [0, \rho]} \sup_{r \geq 0} \tilde{I}_{\text{RS}}(q, r, \lambda) \quad (127)$$

where $\tilde{I}_{\text{RS}}(q, r, \lambda)$ is given by

$$\tilde{I}_{\text{RS}}(q, r, \lambda) := I_{P_0}(r) + \alpha \mathcal{I}_{P_{\text{out}}}(q; \rho) + \epsilon I(W^*; \sqrt{\lambda} f(Y^{(q)}) + Z' | V) - \frac{r}{2}(\rho - q). \quad (128)$$

Conditionally on $(V, f(Y^{(q)}))$, the random variables W^* and $\sqrt{\lambda} f(Y^{(q)}) + Z'$ are independent, therefore

$$I(f(Y^{(q)}); \sqrt{\lambda} f(Y^{(q)}) + Z' | V) = I(W^*, f(Y^{(q)}); \sqrt{\lambda} f(Y^{(q)}) + Z' | V).$$

Now, by the chain rule of the mutual information we have

$$I(W^*, f(Y^{(q)}); \sqrt{\lambda} f(Y^{(q)}) + Z' | V) = I(W^*; \sqrt{\lambda} f(Y^{(q)}) + Z' | V) + I(f(Y^{(q)}); \sqrt{\lambda} f(Y^{(q)}) + Z' | V, W^*).$$

We obtain that

$$I(W^*; \sqrt{\lambda}f(Y^{(q)}) + Z'|V) = I(f(Y^{(q)}); \sqrt{\lambda}f(Y^{(q)}) + Z'|V) - I(f(Y^{(q)}); \sqrt{\lambda}f(Y^{(q)}) + Z'|V, W^*). \quad (129)$$

Notice that the last mutual information in the above equation does not depend on q nor r . Therefore we have:

$$\inf_{q \in [0, \rho]} \sup_{r \geq 0} \tilde{I}_{\text{RS}}(q, r, \lambda) = -\epsilon I(f(Y^{(q)}); \sqrt{\lambda}f(Y^{(q)}) + Z'|V, W^*) + \inf_{q \in [0, \rho]} \sup_{r \geq 0} \tilde{i}_{\text{RS}}(q, r, \lambda). \quad (130)$$

Now, by the chain rule, we have

$$\frac{1}{n} I(\mathbf{X}^*; \mathbf{Y}, \sqrt{\lambda} \mathbf{Y}' + \mathbf{Z}' | \Phi, \Phi') = \frac{1}{n} I(\mathbf{X}^*; \mathbf{Y} | \Phi) + \frac{1}{n} I(\mathbf{X}^*; \sqrt{\lambda} \mathbf{Y}' + \mathbf{Z}' | \mathbf{Y}, \Phi, \Phi'). \quad (131)$$

The limit of the left-hand side is given by (127). By Corollary 2, we have $\lim_{n \rightarrow \infty} I(\mathbf{X}^*; \mathbf{Y} | \Phi)/n = i_\infty$. It remains to investigate the last term of the equation above. By the arguments used to prove (129), we have

$$\begin{aligned} I(\mathbf{X}^*; \sqrt{\lambda} \mathbf{Y}' + \mathbf{Z}' | \mathbf{Y}, \Phi, \Phi') &= I(\mathbf{Y}'; \sqrt{\lambda} \mathbf{Y}' + \mathbf{Z}' | \mathbf{Y}, \Phi, \Phi') - I(\mathbf{Y}'; \sqrt{\lambda} \mathbf{Y}' + \mathbf{Z}' | \mathbf{Y}, \Phi, \mathbf{X}^*, \Phi') \\ &= I(\mathbf{Y}'; \sqrt{\lambda} \mathbf{Y}' + \mathbf{Z}' | \mathbf{Y}, \Phi, \Phi') - I(\mathbf{Y}'; \sqrt{\lambda} \mathbf{Y}' + \mathbf{Z}' | \mathbf{X}^*, \Phi'). \end{aligned} \quad (132)$$

We have $I(\mathbf{Y}'; \sqrt{\lambda} \mathbf{Y}' + \mathbf{Z}' | \mathbf{X}^*, \Phi')/n = \epsilon I(Y'_1; \sqrt{\lambda} Y'_1 + Z'_1 | \mathbf{X}^*, \Phi'_1)$ and it is not difficult to show, using similar computations as in the proof of Corollary 2, that

$$I(Y'_1; \sqrt{\lambda} Y'_1 + Z'_1 | \mathbf{X}^*, \Phi'_1) \xrightarrow[n \rightarrow \infty]{} I(f(Y^{(q)}); \sqrt{\lambda} f(Y^{(q)}) + Z'|V, W^*),$$

(recall that the right-hand side does not depend on q). Combining this with (132), (131), (127), Corollary 2 and (130), we obtain the desired result. \square

5.1.2 Proof of Theorem 4

In order to compute the limit of the (generalized) generalization error, we work in the teacher-student scenario with side-information discussed in the previous section.

Lemma 3. For all $\alpha, \lambda > 0$ the set

$$D_{\alpha, \lambda} := \{ \epsilon \geq 0 \mid \text{the infimum in (122) is achieved at a unique } q_{\alpha, \epsilon, \lambda}^* \} \quad (133)$$

is equal to $[0, +\infty)$ minus some countable set. Moreover, $\epsilon \mapsto q_{\alpha, \lambda, \epsilon}^*$ is continuous on $D_{\alpha, \epsilon}$.

Proof. This follows from the same arguments than the proof of Proposition 1. \square

Lemma 4. For all $\alpha, \lambda > 0$, we have for all $\epsilon \in D_{\alpha, \lambda} \setminus \{0\}$

$$\lim_{n \rightarrow \infty} \text{MMSE}(Y'_1 | \mathbf{Y}, \mathbf{U}, \Phi, \Phi') = M_f(\lambda, q_{\alpha, \epsilon, \lambda}^*),$$

where $q_{\alpha, \epsilon, \lambda}^*$ is the unique minimizer of (122).

Proof. Let us fix $\alpha, \epsilon > 0$. Consider the function

$$h_{\alpha, \epsilon} : \lambda \mapsto \inf_{q \in [0, \rho]} \sup_{r \geq 0} \tilde{i}_{\text{RS}}(q, r, \lambda). \quad (134)$$

Corollary 4 from [93] gives that $h_{\alpha,\epsilon}$ is differentiable at λ if and only if

$$\left\{ \epsilon \frac{\partial}{\partial \lambda} I(f(Y^{(q)}); \sqrt{\lambda}f(Y^{(q)}) + Z'|V) = \frac{\epsilon}{2} M_f(\lambda, q) \mid q \text{ minimizer of (122) (or equivalently of (134))} \right\}$$

is a singleton (the equality comes from the I-MMSE relation from Proposition 13). In such case, Corollary 4 from [93] also gives that

$$h'_{\alpha,\epsilon}(\lambda) = \frac{\epsilon}{2} M_f(\lambda, q), \quad (135)$$

for all q minimizer of (134). So if now $\epsilon \in D_{\alpha,\lambda} \setminus \{0\}$, then the minimizer is unique and thus $h_{\alpha,\epsilon}$ is differentiable at λ , with derivative $h'_{\alpha,\lambda}(\lambda) = \epsilon M_f(\lambda, q_{\alpha,\epsilon,\lambda}^*)/2$. However, by (122) in Lemma 2, $h_{\alpha,\epsilon}$ is the pointwise limit on \mathbb{R}_+ of the sequence of concave functions

$$(h_n)_{n \geq 1} = \left(\lambda \mapsto \frac{1}{n} I(\mathbf{Y}'; \sqrt{\lambda} \mathbf{Y}' + \mathbf{Z}' | \mathbf{Y}, \Phi, \Phi') + i_\infty \right)_{n \geq 1}.$$

Consequently, a standard convex analysis result gives that $h'_n(\lambda) \xrightarrow{n \rightarrow \infty} h'_{\alpha,\epsilon}(\lambda)$. By the I-MMSE relation (Proposition 13) we have $h'_n(\lambda) = \epsilon \text{MMSE}(Y'_1 | \mathbf{Y}, \mathbf{U}, \Phi, \Phi')/2$ and we conclude using the fact that $\epsilon \neq 0$. \square

Lemma 5. For all $\alpha \in D^*$ given by (50),

$$\lim_{\lambda \rightarrow 0} \lim_{\epsilon \rightarrow 0} M_f(\lambda, q_{\alpha,\epsilon,\lambda}^*) = \mathcal{E}_f(q^*(\alpha)).$$

Proof. Let $\alpha \in D^*$ and $\lambda > 0$. We have by definition of $D_{\alpha,\lambda}$, of D^* and using the link between \tilde{i}_{RS} and i_{RS} given by (123), that $0 \in D_{\alpha,\lambda}$. By Lemma 3 above, we have

$$q_{\alpha,\epsilon,\lambda}^* \xrightarrow{\epsilon \rightarrow 0, \epsilon \in D_{\alpha,\lambda}} q_{\alpha,0,\lambda}^* = q^*(\alpha).$$

Analogously to Proposition 22, $M_f(\lambda, \cdot)$ is continuous on $[0, \rho]$, thus $\lim_{\epsilon \rightarrow 0} M_f(\lambda, q_{\alpha,\epsilon,\lambda}^*) = M_f(\lambda, q^*(\alpha))$. And we obtain the result by taking $\lim_{\lambda \rightarrow 0} M_f(\lambda, q^*(\alpha)) = \mathcal{E}_f(q^*(\alpha))$, using that $M_f(\cdot, q)$ is continuous for $q \in [0, \rho]$ fixed (by Proposition 13) and by comparing (125) and (61). \square

In order to simplify the proof, we assume that $m = \alpha n$. By definition of the generalization error (60) and of the labels \mathbf{Y}' given by (119),

$$\mathcal{E}_{f,n}(\alpha) := \text{MMSE}(Y'_1 | \mathbf{Y}, \Phi, \Phi').$$

Lemma 6 (Lower bound on the generalization error). For all $\alpha \in D^*$,

$$\liminf_{n \rightarrow \infty} \mathcal{E}_{f,n}(\alpha) \geq \mathcal{E}_f(q^*(\alpha)).$$

Proof. Let $\alpha \in D^*$, $\lambda > 0$ and $\epsilon \in D_{\alpha,\lambda} \setminus \{0\}$. Obviously,

$$\mathcal{E}_{f,n}(\alpha) \geq \text{MMSE}(Y'_1 | \mathbf{Y}, \mathbf{U}, \Phi, \Phi') \xrightarrow{n \rightarrow \infty} M_f(\lambda, q_{\alpha,\epsilon,\lambda}^*),$$

where we used Lemma 4. Consequently $\liminf_{n \rightarrow \infty} \mathcal{E}_{f,n}(\alpha) \geq M_f(\lambda, q_{\alpha,\epsilon,\lambda}^*)$ and we obtain the lower bound by letting $\epsilon, \lambda \rightarrow 0$ and using Lemma 5. \square

Let us now prove the converse upper bound.

Lemma 7. There exists a constant $C > 0$ (that only depend on f) such that for all $\alpha, \lambda > 0$ and all $\epsilon \in D_{\alpha,\lambda} \setminus \{0\}$

$$\limsup_{n \rightarrow \infty} \mathcal{E}_{f,n}(\alpha + \epsilon) \leq M_f(\lambda, q_{\alpha,\epsilon,\lambda}^*) + C\lambda.$$

Proof. We will let the signal-to-noise ratio (snr) of the observation of Y_1' go to zero. Let us denote by λ_1 this snr: $U_1 = \sqrt{\lambda_1} Y_1' + Z_1'$. We will let λ_1 go from λ to 0 while the other snr for the observations of U_μ for $\mu = 2, \dots, \epsilon n$ will remain equal to λ . Recall that we denote $\mathbf{U} = (U_\mu)_{\mu=1}^{\epsilon n}$. Using Proposition 9 from [94],

$$\left| \frac{\partial}{\partial \lambda_1} \text{MMSE}(Y_1' | \mathbf{Y}, \mathbf{U}, \Phi, \Phi') \right| = \mathbb{E}[\text{Var}(Y_1' | \mathbf{Y}, \mathbf{U}, \Phi, \Phi')^2] \leq \mathbb{E}[(Y_1')^4] \leq \|f\|_\infty^4.$$

We define $C := \|f\|_\infty^4$. Consequently, by the mean value theorem,

$$|\text{MMSE}(Y_1' | \mathbf{Y}, \mathbf{U}, \Phi, \Phi') - \text{MMSE}(Y_1' | \mathbf{Y}, (U_\mu)_{\mu=2}^{\epsilon n}, \Phi, \Phi')| \leq C\lambda. \quad (136)$$

Since $(U_\mu)_{\mu=2}^{\epsilon n}$ contains less information than $(\tilde{Y}_\mu)_{\mu=2}^{\epsilon n}$ because of the additional Gaussian noise and the application of the function f , we have

$$\text{MMSE}(Y_1' | \mathbf{Y}, (U_\mu)_{\mu=2}^{\epsilon n}, \Phi, \Phi') \geq \text{MMSE}(Y_1' | \mathbf{Y}, (\tilde{Y}_\mu)_{\mu=2}^{\epsilon n}, \Phi, \Phi') = \mathcal{E}_{f,n}(\alpha + \epsilon - 1/n) \geq \mathcal{E}_{f,n}(\alpha + \epsilon). \quad (137)$$

The last identity combined with (136) leads to

$$\text{MMSE}(Y_1' | \mathbf{Y}, \mathbf{U}, \Phi, \Phi') + C\lambda \geq \mathcal{E}_{f,n}(\alpha + \epsilon). \quad (138)$$

By Lemma 4 we know that $\lim_{n \rightarrow \infty} \text{MMSE}(Y_1' | \mathbf{Y}, \mathbf{U}, \Phi, \Phi') = M_f(\lambda, q_{\alpha, \epsilon, \lambda}^*)$. Thus we conclude by taking the limsup in the inequality above. \square

Corollary 6 (Upper bound on the generalization error). *For all $\alpha \in D^*$,*

$$\limsup_{n \rightarrow \infty} \mathcal{E}_{f,n}(\alpha) \leq \mathcal{E}_f(q^*(\alpha)).$$

Proof. Let $\alpha \in D^*$, $\lambda > 0$ and $\epsilon_1 > 0$ such that $\alpha - \epsilon_1 \in D^*$. Since by Lemma 3 the set $D_{\alpha - \epsilon_1, \lambda}$ is dense in \mathbb{R}_+ , we can find $\epsilon_2 \in D_{\alpha - \epsilon_1, \lambda}$ such that $0 < \epsilon_2 \leq \epsilon_1$. Using Lemma 7 above, we have

$$\limsup_{n \rightarrow \infty} \mathcal{E}_{f,n}(\alpha - \epsilon_1 + \epsilon_2) \leq M_f(\lambda, q_{\alpha - \epsilon_1, \epsilon_2, \lambda}^*) + C\lambda.$$

Now, using the fact that $\epsilon_2 \leq \epsilon_1$ we have

$$\limsup_{n \rightarrow \infty} \mathcal{E}_{f,n}(\alpha) \leq \limsup_{n \rightarrow \infty} \mathcal{E}_{f,n}(\alpha - \epsilon_1 + \epsilon_2) \leq M_f(\lambda, q_{\alpha - \epsilon_1, \epsilon_2, \lambda}^*) + C\lambda.$$

Now, by Lemma 5 we have

$$\lim_{\lambda \rightarrow 0} \lim_{\epsilon_2 \rightarrow 0} M_f(\lambda, q_{\alpha - \epsilon_1, \epsilon_2, \lambda}^*) + C\lambda = \mathcal{E}_f(q^*(\alpha - \epsilon_1))$$

which leads to $\limsup_{n \rightarrow \infty} \mathcal{E}_{f,n}(\alpha) \leq \mathcal{E}_f(q^*(\alpha - \epsilon_1))$. We conclude by letting $\epsilon_1 \rightarrow 0$ (recall that by Proposition 1 D^* is dense in \mathbb{R}_+ so it is possible to find $\epsilon_1 > 0$ arbitrary small such that $\alpha - \epsilon_1 \in D^*$), using the continuity of \mathcal{E}_f (by Proposition 22) and the continuity of q^* (by Proposition 1). \square

Proof of Theorem 4: For the moment we have proven Theorem 4 when f is continuous and bounded. We are going to relax this assumption by approximation. Let $f : \mathbb{R} \rightarrow \mathbb{R}$ such that $\mathbb{E}[|f(Y_{\text{new}})|^{2+\gamma}]$ remains bounded as n goes to infinity, for some $\gamma > 0$. Let $\epsilon > 0$. By density of the continuous and bounded functions in the space $L^2(\mathbb{R})$ equipped with the law of $Y^{(q)} \sim P_{\text{out}}(\cdot | \sqrt{q}V + \sqrt{\rho - q}W)$ ($V, W \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$), we can find a continuous bounded function $\tilde{f} : \mathbb{R} \rightarrow \mathbb{R}$ such that $\mathbb{E}[(f(Y^{(q)}) - \tilde{f}(Y^{(q)}))^2] \leq \epsilon$.

Lemma 8. For all $q \in [0, \rho]$ (because the law of $Y^{(q)}$ does not depend on q), we have

$$f(Y_{\text{new}}) - \tilde{f}(Y_{\text{new}}) \xrightarrow[n \rightarrow \infty]{(d)} f(Y^{(q)}) - \tilde{f}(Y^{(q)}). \quad (139)$$

Proof. Let $(\mathbf{A}_{\text{new}}, Z_{\text{new}}) \sim P_A \otimes \mathcal{N}(0, 1)$ such that $Y_{\text{new}} = \varphi(\Phi_{\text{new}} \cdot \mathbf{X}^*/\sqrt{n}, \mathbf{A}_{\text{new}}) + \sqrt{\Delta}Z_{\text{new}}$. By the central limit theorem (that we apply under (h1)-(h3) and using (h4))

$$\varphi\left(\frac{\Phi_{\text{new}} \cdot \mathbf{X}^*}{\sqrt{n}}, \mathbf{A}_{\text{new}}\right) \xrightarrow[n \rightarrow \infty]{(d)} \varphi(\sqrt{\rho}Z, \mathbf{A}_{\text{new}}), \quad (140)$$

where $Z \sim \mathcal{N}(0, 1)$ is independent from \mathbf{A}_{new} . Under (h5.b) this proves Lemma 8, because in that case $Y_{\text{new}} = \varphi(\Phi_{\text{new}} \cdot \mathbf{X}^*/\sqrt{n}, \mathbf{A}_{\text{new}})$ takes values in \mathbb{N} . Under (h5.a) we let $g : \mathbb{R} \rightarrow \mathbb{R}$ be a continuous bounded function and we write $h := f - \tilde{f}$. Then

$$\begin{aligned} \mathbb{E}[g \circ h(Y_{\text{new}})] &= \mathbb{E}\left[g \circ h\left(\varphi\left(\frac{\Phi_{\text{new}} \cdot \mathbf{X}^*}{\sqrt{n}}, \mathbf{A}_{\text{new}}\right) + \sqrt{\Delta}Z_{\text{new}}\right)\right] \\ &= \mathbb{E}\left[\frac{1}{\sqrt{2\pi\Delta}} \int g \circ h(z) \exp\left\{-\frac{1}{2\Delta}\left(z - \varphi\left(\frac{\Phi_{\text{new}} \cdot \mathbf{X}^*}{\sqrt{n}}, \mathbf{A}_{\text{new}}\right)\right)^2\right\} dz\right]. \end{aligned}$$

The function $x \mapsto \int g \circ h(z) \frac{e^{-\frac{1}{2\Delta}(z-x)^2}}{\sqrt{2\pi\Delta}} dz$ is continuous and bounded: (140) then gives that

$$\mathbb{E}[g \circ h(Y_{\text{new}})] \xrightarrow[n \rightarrow \infty]{} \mathbb{E}\left[\frac{1}{\sqrt{2\pi\Delta}} \int g \circ h(z) \exp\left\{-\frac{1}{2\Delta}\left(z - \varphi(\sqrt{\rho}Z, \mathbf{A}_{\text{new}})\right)^2\right\} dz\right] = \mathbb{E}[g \circ h(Y^{(q)})],$$

which concludes the proof by the Portemanteau Theorem. \square

The sequence $((f(Y_{\text{new}}) - \tilde{f}(Y_{\text{new}}))^2)_{n \geq 0}$ is uniformly integrable because bounded in $L^{1+\gamma}$ with $\gamma > 0$. Consequently, Lemma 8 above implies

$$\mathbb{E}[(f(Y_{\text{new}}) - \tilde{f}(Y_{\text{new}}))^2] \xrightarrow[n \rightarrow \infty]{} \|f(Y^{(q)}) - \tilde{f}(Y^{(q)})\|_{L^2}^2 = \mathbb{E}[(f(Y^{(q)}) - \tilde{f}(Y^{(q)}))^2] \leq \epsilon.$$

Therefore, we can find $n_0 \in \mathbb{N}$ such that for all $n \geq n_0$, $\|f(Y_{\text{new}}) - \tilde{f}(Y_{\text{new}})\|_{L^2}^2 = \mathbb{E}[(f(Y_{\text{new}}) - \tilde{f}(Y_{\text{new}}))^2] \leq 2\epsilon$. If we now apply Theorem 4 for \tilde{f} , we can find $n_1 \geq n_0$ such that for all $n \geq n_1$, $|\mathcal{E}_{f,n}^{1/2} - \mathcal{E}_{\tilde{f},n}(q^*(\alpha))^{1/2}| \leq \sqrt{\epsilon}$. Let $n \geq 1$, and compute

$$\begin{aligned} \left|\mathcal{E}_{f,n}^{1/2} - \mathcal{E}_{\tilde{f},n}^{1/2}\right| &= \left|\|f(Y_{\text{new}}) - \mathbb{E}[f(Y_{\text{new}})|\mathbf{Y}, \Phi, \Phi_{\text{new}}]\|_{L^2} - \|\tilde{f}(Y_{\text{new}}) - \mathbb{E}[\tilde{f}(Y_{\text{new}})|\mathbf{Y}, \Phi, \Phi_{\text{new}}]\|_{L^2}\right| \\ &\leq \|f(Y_{\text{new}}) - \tilde{f}(Y_{\text{new}})\|_{L^2} + \|\mathbb{E}[f(Y_{\text{new}}) - \tilde{f}(Y_{\text{new}})|\mathbf{Y}, \Phi, \Phi_{\text{new}}]\|_{L^2} \\ &\leq 2\|f(Y_{\text{new}}) - \tilde{f}(Y_{\text{new}})\|_{L^2} \leq 2\sqrt{2}\epsilon, \end{aligned}$$

where we successively used the triangular inequality twice for the first inequality ($\|a - b\| - \|x - y\| \leq \|a - x + y - b\| \leq \|a - x\| + \|y - b\|$) and Jensen's inequality for the second. By the same arguments we have also $|\mathcal{E}_f(q)^{1/2} - \mathcal{E}_{\tilde{f}}(q)^{1/2}| \leq 2\sqrt{\epsilon}$ for all $q \in [0, \rho]$. We conclude that for all $n \geq n_1$,

$$\begin{aligned} \left|\mathcal{E}_{f,n}^{1/2} - \mathcal{E}_f(q^*(\alpha))^{1/2}\right| &\leq \left|\mathcal{E}_{f,n}^{1/2} - \mathcal{E}_{\tilde{f},n}^{1/2}\right| + \left|\mathcal{E}_{\tilde{f},n}(q^*(\alpha))^{1/2} - \mathcal{E}_f(q^*(\alpha))^{1/2}\right| + \left|\mathcal{E}_{\tilde{f},n}^{1/2} - \mathcal{E}_{\tilde{f}}(q^*(\alpha))^{1/2}\right| \\ &\leq (2\sqrt{2} + 3)\sqrt{\epsilon}, \end{aligned} \quad (141)$$

which proves Theorem 4. \square

5.2 Generalization error of GAMP: Proof of Proposition 2

Let us decompose:

$$\mathcal{E}_{\text{gen}}^{\text{GAMP},t} := \mathbb{E}[(Y_{\text{new}} - \widehat{Y}^{\text{GAMP},t})^2] = \mathbb{E}[Y_{\text{new}}^2] + \mathbb{E}[(\widehat{Y}^{\text{GAMP},t})^2] - 2\mathbb{E}[Y_{\text{new}}\widehat{Y}^{\text{GAMP},t}]. \quad (142)$$

Lemma 9. *We have*

$$\mathbb{E}[Y_{\text{new}}\widehat{Y}^{\text{GAMP},t}] \xrightarrow{n \rightarrow \infty} \mathbb{E}_V \left[\mathbb{E}_W \left[\int dY Y P_{\text{out}}(Y | \sqrt{q^t} V + \sqrt{\rho - q^t} W) \right]^2 \right]. \quad (143)$$

Proof. Start by writing

$$\mathbb{E}[Y_{\text{new}}\widehat{Y}^{\text{GAMP},t}] = \mathbb{E} \int y y' P_{\text{out}}\left(y \mid \frac{\Phi_{\text{new}} \cdot \mathbf{X}^*}{\sqrt{n}}\right) P_{\text{out}}\left(y' \mid \frac{\Phi_{\text{new}} \cdot \widehat{\mathbf{x}}^t}{\sqrt{n}} + \sqrt{\rho - q^t} W\right) dy dy'$$

where $W \sim \mathcal{N}(0, 1)$ is independent of everything else. $\Phi_{\text{new}} \sim \mathcal{N}(0, \mathbf{I}_n)$ is independent of \mathbf{X}^* and $\widehat{\mathbf{x}}^t$, so, conditionally on $\mathbf{X}^*, \widehat{\mathbf{x}}^t$ we have

$$\left(\frac{\Phi_{\text{new}} \cdot \mathbf{X}^*}{\sqrt{n}}, \frac{\Phi_{\text{new}} \cdot \widehat{\mathbf{x}}^t}{\sqrt{n}} \right) \sim \mathcal{N}\left(0, \frac{1}{n} \begin{pmatrix} \|\mathbf{X}^*\|^2 & \widehat{\mathbf{x}}^t \cdot \mathbf{X}^* \\ \widehat{\mathbf{x}}^t \cdot \mathbf{X}^* & \|\widehat{\mathbf{x}}^t\|^2 \end{pmatrix}\right).$$

We assumed that (69) holds, i.e. $\mathbf{X}^* \cdot \widehat{\mathbf{x}}^t/n \rightarrow q^t$ and $\|\widehat{\mathbf{x}}^t\|^2/n \rightarrow q^t$, in probability. By the law of large numbers $\|\mathbf{X}^*\|^2/n \rightarrow \rho$ in probability. Consequently,

$$\left(\frac{\Phi_{\text{new}} \cdot \mathbf{X}^*}{\sqrt{n}}, \frac{\Phi_{\text{new}} \cdot \widehat{\mathbf{x}}^t}{\sqrt{n}} \right) \xrightarrow[n \rightarrow \infty]{(d)} \mathcal{N}\left(0, \begin{pmatrix} \rho & q^t \\ q^t & q^t \end{pmatrix}\right).$$

Since $x \mapsto P_{\text{out}}(\cdot|x)$ is continuous almost everywhere for the Wasserstein distance of order 2, the function $h : (a, b) \mapsto \mathbb{E}_W \int y y' P_{\text{out}}(y|a) P_{\text{out}}(y'|b + \sqrt{\rho - q^t} W) dy dy'$ with $W \sim \mathcal{N}(0, 1)$ is continuous almost everywhere. Therefore

$$H_n := h\left(\frac{\Phi_{\text{new}} \cdot \mathbf{X}^*}{\sqrt{n}}, \frac{\Phi_{\text{new}} \cdot \widehat{\mathbf{x}}^t}{\sqrt{n}}\right) \xrightarrow[n \rightarrow \infty]{(d)} h(\sqrt{q^t} Z_0 + \sqrt{\rho - q^t} Z_1, \sqrt{q^t} Z_0), \quad (144)$$

where $Z_0, Z_1 \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$. We have by Jensen's inequality

$$\begin{aligned} \mathbb{E}[|H_n|^{1+\eta}] &\leq \mathbb{E}[|Y_{\text{new}}\widehat{Y}^{\text{GAMP},t}|^{1+\eta}] \leq \mathbb{E}\left[\left(\frac{1}{2}Y_{\text{new}}^2 + \frac{1}{2}(\widehat{Y}^{\text{GAMP},t})^2\right)^{1+\eta}\right] \\ &\leq \frac{1}{2}\mathbb{E}|Y_1|^{2+2\eta} + \frac{1}{2}\mathbb{E}|\widehat{Y}^{\text{GAMP},t}|^{2+2\eta}. \end{aligned}$$

By assumption, there exists $\eta > 0$ such that the two last terms above remain bounded with n : H_n is therefore bounded in $L^{1+\eta}$ and is therefore uniformly integrable. From (144) we thus get

$$\begin{aligned} \mathbb{E}[Y_{\text{new}}\widehat{Y}^{\text{GAMP},t}] &= \mathbb{E}[H_n] \xrightarrow{n \rightarrow \infty} \mathbb{E}[h(\sqrt{q^t} Z_0 + \sqrt{\rho - q^t} Z_1, \sqrt{q^t} Z_0)] \\ &= \mathbb{E}_V \left[\mathbb{E}_W \left[\int dY Y P_{\text{out}}(Y | \sqrt{q^t} V + \sqrt{\rho - q^t} W) \right]^2 \right]. \end{aligned}$$

□

Following the arguments of Lemma 9 one can also show that

$$\begin{aligned}\mathbb{E}[(\widehat{Y}^{\text{GAMP},t})^2] &\xrightarrow{n \rightarrow \infty} \mathbb{E}_V \left[\mathbb{E}_W \left[\int dY Y P_{\text{out}}(Y | \sqrt{q^t} V + \sqrt{\rho - q^t} W) \right]^2 \right], \\ \mathbb{E}[Y_{\text{new}}^2] &\xrightarrow{n \rightarrow \infty} \mathbb{E}_V \int dY Y^2 P_{\text{out}}(Y | \sqrt{\rho} V).\end{aligned}$$

This proves (together with (142) and Lemma 9) Proposition 2.

5.3 Limit of the overlap: Proof of Theorem 2

Recall the definition of the overlap (105): $Q_n := \mathbf{X}^* \cdot \mathbf{x}/n$, where $\mathbf{x} = (x_1, \dots, x_n)$ is a sample from the posterior distribution $P(\mathbf{X}^* | \mathbf{Y}, \Phi)$, independently of everything else. In this section we will show that $|Q_n|$ converges in probability to $q^*(\alpha)$, when $\alpha \in D^*$ given by (50). We will first show an upper-bound in Sec. 5.3.1 below, before proving the converse lower-bound in Sec. 5.3.2.

5.3.1 Upper bound on the overlap

Proposition 9 (Upper bound on the overlap). *For all $\alpha \in D^*$ and for all $\epsilon > 0$,*

$$\mathbb{P}(|Q_n| \geq q^*(\alpha) + \epsilon) \xrightarrow{n \rightarrow \infty} 0.$$

Let us fix $\alpha \in D^*$ and let $p \geq 1$. In order to obtain an upper bound on the overlap, we consider an observation model with some (small) extra information (that takes the form of a tensor of order $2p$) in addition of the original model (21), i.e. we observe

$$\begin{cases} \mathbf{Y} \sim P_{\text{out}}(\cdot | \Phi \mathbf{X}^* / \sqrt{n}), \\ \mathbf{Y}' = \sqrt{\frac{\lambda}{n^{2p-1}}} (\mathbf{X}^*)^{\otimes 2p} + \mathbf{Z}', \end{cases} \quad (145)$$

where $\lambda \geq 0$, $\mathbf{Z}' = (Z'_{i_1 \dots i_{2p}})_{1 \leq i_1, \dots, i_{2p} \leq n} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$ and $(\mathbf{X}^*)^{\otimes 2p} = (X_{i_1} \dots X_{i_{2p}})_{1 \leq i_1, \dots, i_{2p} \leq n}$. In order to prove Proposition 9 we need the two results below, which are proven after the proof of Proposition 9.

Proposition 10 (Mutual information of the perturbed model). *For all $\lambda \geq 0$, the mutual information for model (145) verifies*

$$\lim_{n \rightarrow \infty} \frac{1}{n} I(\mathbf{X}^*; \mathbf{Y}, \mathbf{Y}' | \Phi) = I(\lambda), \quad (146)$$

where the right-hand-side is

$$I(\lambda) := \inf_{q \in [0, \rho]} \sup_{r \geq 0} \left\{ I_{P_0}(r + 2p\lambda q^{2p-1}) + \alpha \mathcal{I}_{P_{\text{out}}}(q) - \frac{r}{2}(\rho - q) + \frac{2p-1}{2} \lambda q^{2p} - \rho p \lambda q^{2p-1} + \frac{\lambda}{2} \rho^{2p} \right\}. \quad (147)$$

Lemma 10. *The function I defined above by (147) is concave on \mathbb{R}_+ . Its left- and right-derivatives are given by*

$$\begin{aligned} I'(\lambda^+) &= \min \left\{ \frac{1}{2} (\rho^{2p} - q_*(\lambda)^{2p}) \mid q_*(\lambda) \text{ achieves the infimum in (147)} \right\}, \\ I'(\lambda^-) &= \max \left\{ \frac{1}{2} (\rho^{2p} - q_*(\lambda)^{2p}) \mid q_*(\lambda) \text{ achieves the infimum in (147)} \right\}.\end{aligned}$$

We are now in position to prove Proposition 9.

Proof of Proposition 9: By the I-MMSE relation of Proposition 13,

$$\frac{1}{n} \frac{\partial}{\partial \lambda} I(\mathbf{X}^*; \mathbf{Y}, \mathbf{Y}' | \Phi) = \frac{1}{n} \frac{\partial}{\partial \lambda} I((\mathbf{X}^*)^{\otimes 2p}; \mathbf{Y}, \mathbf{Y}' | \Phi) = \frac{1}{2n^{2p}} \text{MMSE}((\mathbf{X}^*)^{\otimes 2p} | \mathbf{Y}, \mathbf{Y}', \Phi).$$

Using Proposition 10 and Lemma 10 above we obtain by concavity that

$$\frac{1}{2n^{2p}} \text{MMSE}((\mathbf{X}^*)^{\otimes 2p} | \mathbf{Y}, \mathbf{Y}', \Phi) = \frac{1}{n} \frac{\partial}{\partial \lambda} I(\mathbf{X}^*; \mathbf{Y}, \mathbf{Y}' | \Phi) \xrightarrow{n \rightarrow \infty} I'(\lambda) = \frac{1}{2} (\rho^{2p} - q_*(\lambda)^{2p}),$$

for all $\lambda > 0$ for which the infimum of (147) is achieved at a unique $q_*(\lambda)$. Consequently,

$$\liminf_{n \rightarrow \infty} \frac{1}{n^{2p}} \text{MMSE}((\mathbf{X}^*)^{\otimes 2p} | \mathbf{Y}, \Phi) \geq \liminf_{n \rightarrow \infty} \frac{1}{n^{2p}} \text{MMSE}((\mathbf{X}^*)^{\otimes 2p} | \mathbf{Y}, \mathbf{Y}', \Phi) = \rho^{2p} - q_*(\lambda)^{2p}. \quad (148)$$

Let us now suppose that $\alpha \in D^*$. In that case, there exists a unique $q_*(\lambda = 0) = q^*(\alpha)$ that achieves the infimum in (147). Consequently, $I'(0^+) = \frac{1}{2} (\rho^{2p} - q^*(\alpha)^{2p})$. By concavity, $I'(\lambda) \rightarrow I'(0^+)$ as $\lambda \rightarrow 0$, which gives $q_*(\lambda) \rightarrow q^*(\alpha)$. By taking the $\lambda \rightarrow 0$ limit in (148) above we get

$$\liminf_{n \rightarrow \infty} \frac{1}{n^{2p}} \text{MMSE}((\mathbf{X}^*)^{\otimes 2p} | \mathbf{Y}, \Phi) \geq \rho^{2p} - q^*(\alpha)^{2p}.$$

One verifies easily that

$$\frac{1}{n^{2p}} \text{MMSE}((\mathbf{X}^*)^{\otimes 2p} | \mathbf{Y}, \Phi) = \rho^{2p} - \mathbb{E}[Q_n^{2p}] + o_n(1), \quad (149)$$

so we deduce that

$$\limsup_{n \rightarrow \infty} \mathbb{E}[Q_n^{2p}] \leq q^*(\alpha)^{2p}.$$

Let $\epsilon > 0$. By Markov's inequality we have

$$\mathbb{P}(|Q_n| \geq q_*(\alpha) + \epsilon) \leq \frac{\mathbb{E}[Q_n^{2p}]}{(q_*(\alpha) + \epsilon)^{2p}}.$$

By taking the lim sup in n on both sides we obtain

$$\limsup_{n \rightarrow \infty} \mathbb{P}(|Q_n| \geq q_*(\alpha) + \epsilon) \leq \frac{q^*(\alpha)^{2p}}{(q_*(\alpha) + \epsilon)^{2p}},$$

and Proposition 9 follows by taking the $p \rightarrow \infty$ limit in the inequality above. \square

We now prove the two preliminary results used in the proof of Proposition 9.

Proof of Proposition 10: The proof is very similar to the one of Theorem 1 (and Corollary 2), by the adaptive interpolation method (see Sec. 4), so we provide only the main arguments and omit to write the small perturbation (i.e. the ϵ_1, ϵ_2 present in Sec. 4) for simplicity.

In order to tackle model (145) we need first to study a simpler one, namely when we have access to the simultaneous observations $\mathbf{Y} \sim P_{\text{out}}(\cdot | \Phi \mathbf{X}^* / \sqrt{n})$ and $\mathbf{Y}'' = \sqrt{\gamma} \mathbf{X}^* + \mathbf{Z}''$. Define, for $\gamma \geq 0$, the free entropy (expected log-partition function) of this model:

$$F_n(\gamma) := \frac{1}{n} \mathbb{E} \ln \int dP_0(\mathbf{x}) \exp \left(\sum_{i=1}^n \sqrt{\gamma} Z_i'' x_i + \gamma x_i X_i^* - \frac{\gamma}{2} x_i^2 \right) \prod_{\mu=1}^m P_{\text{out}} \left(Y_\mu \middle| \frac{1}{\sqrt{n}} \Phi_\mu \cdot \mathbf{x} \right), \quad (150)$$

where $Z_i'' \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$ are independent of everything else. Let us define

$$F_{\text{RS}}(\gamma) := \sup_{q \in [0, \rho]} \inf_{r \geq 0} \left\{ \psi_{P_0}(r + \gamma) + \alpha \Psi_{P_{\text{out}}}(q) - \frac{rq}{2} \right\}. \quad (151)$$

A slight and easy modification of the Theorem 1 gives that for all $\gamma \geq 0$

$$F_n(\gamma) \xrightarrow{n \rightarrow \infty} F_{\text{RS}}(\gamma). \quad (152)$$

F_n is a convex function of γ (this can be checked by relating it to the mutual information like in Corollary 2 and then using the I-MMSE relation of Proposition 13), thus F_{RS} is too. The function F_{RS} is therefore continuous on \mathbb{R}_+ . F_n is also a non-decreasing function of γ (this is again checked using the I-MMSE relation). By Dini's second theorem we obtain that the convergence of (152) is uniform over all compact subsets of \mathbb{R}_+ .

Now that we have studied this simpler model, we come back to the analysis of (145). We proceed by interpolation as in Sec. 4.1. Let $q : [0, 1] \rightarrow [0, \rho]$ be a continuous interpolating function. For $t \in [0, 1]$, consider the following ‘‘interpolating estimation model’’:

$$\begin{cases} Y_\mu & \sim P_{\text{out}}(\cdot | \Phi \mathbf{X}^* / \sqrt{n}), & 1 \leq \mu \leq m, \\ \mathbf{Y}'_t & = \sqrt{\frac{\lambda(1-t)}{n^{2p-1}}} (\mathbf{X}^*)^{\otimes 2p} + \mathbf{Z}', \\ Y''_{t,i} & = \sqrt{2p\lambda \int_0^t q(v)^{2p-1} dv} X_i^* + Z''_i, & 1 \leq i \leq n. \end{cases} \quad (153)$$

Define the corresponding interpolating free entropy:

$$f_n(t) := \frac{1}{n} \mathbb{E} \ln \int dP_0(\mathbf{x}) e^{H_{n,t}(\mathbf{x})} \prod_{\mu=1}^m P_{\text{out}}\left(Y_\mu \middle| \frac{1}{\sqrt{n}} \Phi_\mu \cdot \mathbf{x}\right),$$

where the Hamiltonian of the model is

$$\begin{aligned} H_{n,t}(\mathbf{x}) &= \sum_{i=1}^n \left\{ \sqrt{2p\lambda \int_0^t q(v)^{2p-1} dv} Z''_i x_i + 2p\lambda \int_0^t q(v)^{2p-1} dv x_i X_i^* - p\lambda \int_0^t q(v)^{2p-1} dv x_i^2 \right\} \\ &+ \sum_{i_1, \dots, i_{2p}} \left\{ \sqrt{\frac{\lambda(1-t)}{n^{2p-1}}} Z'_{i_1 \dots i_{2p}} x_{i_1} \dots x_{i_{2p}} + \frac{\lambda(1-t)}{n^{2p-1}} x_{i_1} \dots x_{i_{2p}} X_{i_1}^* \dots X_{i_{2p}}^* - \frac{\lambda(1-t)}{2n^{2p-1}} x_{i_1}^2 \dots x_{i_{2p}}^2 \right\}. \end{aligned}$$

We aim at computing $f_n := f_n(0)$. We have $f_n(1) = F_n(2p\lambda \int_0^1 q(t)^{2p-1} dt)$. Similarly to Proposition 3 one can compute (see [52] where this computation is done):

$$f'_n(t) = -\frac{\lambda}{2} \mathbb{E} \langle Q_t^{2p} - 2pq(t)^{2p-1} Q_t \rangle_t \quad (154)$$

where $Q_t = \sum_{i=1}^n X_i^* x_i / n$ is the overlap between the planted solution \mathbf{X}^* and $\mathbf{x} = (x_1, \dots, x_n)$, a sample from the posterior distribution $P(\mathbf{X}^* | \mathbf{Y}, \mathbf{Y}'_t, \mathbf{Y}''_t)$. Similarly as in Sec. 4.1 the Gibbs bracket $\langle - \rangle_t$ denotes the expectation w.r.t. this t -dependent posterior acting on \mathbf{x} , \mathbb{E} is w.r.t. the quenched variables $\mathbf{Y}, \mathbf{Y}'_t, \mathbf{Y}''_t$. By convexity of the function $x \mapsto x^{2p}$, we have for all $a, b \in \mathbb{R}$, $a^{2p} - 2pab^{2p-1} \geq (1-2p)b^{2p}$. Consequently, if we choose q to be a constant function, i.e. $q(t) = q$ for all $t \in [0, 1]$, we have

$$f'_n(t) \leq \frac{\lambda}{2} (2p-1) q^{2p}.$$

This gives

$$f_n = f_n(0) = f_n(1) - \int_0^1 f'_n(t) dt \geq F_n(2p\lambda q^{2p-1}) - \frac{\lambda}{2}(2p-1)q^{2p}.$$

By taking the \liminf in n on both sides, we obtain $\liminf_{n \rightarrow \infty} f_n \geq F_{\text{RS}}(2p\lambda q^{2p-1}) - \frac{\lambda}{2}(2p-1)q^{2p}$ using (152) and since this holds for all $q \in [0, \rho]$ we get

$$\liminf_{n \rightarrow \infty} f_n \geq \sup_{q \in [0, \rho]} \left\{ F_{\text{RS}}(2p\lambda q^{2p-1}) - \frac{\lambda}{2}(2p-1)q^{2p} \right\}.$$

Let us now prove the converse upper-bound. One can show as in Sec. 4.3 that the overlap Q_t concentrates around its expectation: Proposition 4 applies. This perturbation does not change the free entropy in the limit $n \rightarrow \infty$ nor the following derivation, so we do not track it explicitly for the sake of simplicity. Let us go back to (154). Therefore, using this concentration and then choosing $q(t) = \tilde{q}(t) = \mathbb{E}\langle Q_t \rangle_t$ as done in Sec. 4.4, we obtain that

$$f'_n(t) = \frac{\lambda}{2}(2p-1)\tilde{q}(t)^{2p} + o_n(1).$$

Consequently,

$$\begin{aligned} f_n &= f_n(1) - \int_0^1 f'_n(t) dt = F_n \left(2p\lambda \int_0^1 \tilde{q}(t)^{2p-1} dt \right) - \frac{\lambda}{2}(2p-1) \int_0^1 \tilde{q}(t)^{2p} dt + o_n(1) \\ &\leq F_n \left(2p\lambda \int_0^1 \tilde{q}(t)^{2p-1} dt \right) - \frac{\lambda}{2}(2p-1) \left(\int_0^1 \tilde{q}(t)^{2p-1} dt \right)^{\frac{2p}{2p-1}} + o_n(1) \\ &\leq \sup_{q \in [0, \rho]} \left\{ F_n(2p\lambda q^{2p-1}) - \frac{\lambda}{2}(2p-1)q^{2p} \right\} + o_n(1). \end{aligned}$$

We use now the fact that the convergence in (152) is uniform over all compact sets to get the upper-bound: $\limsup_{n \rightarrow \infty} f_n \leq \sup_{q \in [0, \rho]} \left\{ F_{\text{RS}}(2p\lambda q^{2p-1}) - \frac{\lambda}{2}(2p-1)q^{2p} \right\}$. We conclude that

$$\lim_{n \rightarrow \infty} f_n = \sup_{q \in [0, \rho]} \left\{ F_{\text{RS}}(2p\lambda q^{2p-1}) - \frac{\lambda}{2}(2p-1)q^{2p} \right\}. \quad (155)$$

We are now going to simplify the right-hand side of the above equation.

Lemma 11. F_{RS} is a convex function on \mathbb{R}_+ , whose left- and right-derivatives at $\gamma \geq 0$ are:

$$\begin{aligned} F'_{\text{RS}}(\gamma^+) &= \max \left\{ \frac{1}{2}q_*(\gamma) \mid q_*(\gamma) \text{ achieves the supremum in (151)} \right\}, \\ F'_{\text{RS}}(\gamma^-) &= \min \left\{ \frac{1}{2}q_*(\gamma) \mid q_*(\gamma) \text{ achieves the supremum in (151)} \right\}. \end{aligned}$$

In particular, F_{RS} is differentiable at $\gamma \geq 0$ if and only if the supremum in (151) is achieved at a unique $q_*(\gamma)$.

Proof. We already know that F_{RS} is convex (as a limit of convex functions, see (152)). We have

$$F_{\text{RS}}(\gamma) = \sup_{q \in [0, \rho]} \inf_{r \geq 0} \left\{ \psi_{P_0}(r + \gamma) + \alpha \Psi_{P_{\text{out}}}(q) - \frac{rq}{2} \right\} = \sup_{q \in [0, \rho]} \left\{ \alpha \Psi_{P_{\text{out}}}(q) - g(\gamma, q/2) \right\} \quad (156)$$

where $g(\gamma, x) = \sup_{r \geq 0} \{xr - \psi_{P_0}(\gamma + r)\}$ is the Legendre transform of $r \mapsto \psi_{P_0}(\gamma + r)$. Let us now compute $\frac{\partial g}{\partial \gamma}(\gamma, x)$. If $x \leq \psi'_{P_0}(\gamma)$, then the supremum in r is achieved at $r = 0$, $g(x, \gamma) = -\psi_{P_0}(\gamma)$. If now $x > \psi'_{P_0}(\gamma)$ then

$$g(\gamma, x) = \sup_{r \geq -\gamma} \{xr - \psi_{P_0}(\gamma + r)\} = -\gamma x + \sup_{r \geq 0} \{xr - \psi_{P_0}(r)\}.$$

The first equality comes from the fact that the supremum can not be achieved on $[-\gamma, 0]$ because for all $r \in [-\gamma, 0]$, $x > \psi'_{P_0}(\gamma) \geq \psi'_{P_0}(\gamma + r)$. We obtain

$$g(\gamma, x) = \begin{cases} -\psi_{P_0}(\gamma) & \text{if } x \leq \psi'_{P_0}(\gamma), \\ -x\gamma + g(0, x) & \text{if } x > \psi'_{P_0}(\gamma). \end{cases}$$

From there, we conclude that $\frac{\partial g}{\partial \gamma}(\gamma, x) = -\max(\psi'_{P_0}(\gamma), x)$. By Lemma 23, every optimal couple $(q_*(\gamma), r_*(\gamma))$ satisfy $q_*(\gamma) = 2\psi'_{P_0}(\gamma + r_*(\gamma))$. This implies (by convexity of ψ_{P_0}) that $q_*(\gamma)/2 \geq \psi'_{P_0}(\gamma)$. Using Corollary 4 from [93] F_{RS} we get that

$$F'_{\text{RS}}(\gamma^+) = \max \left\{ -\frac{\partial g}{\partial \gamma}(\gamma, q_*(\gamma)) \mid q_*(\gamma) \text{ maximizer of (151)} \right\} = \max \left\{ \frac{1}{2}q_*(\gamma) \mid q_*(\gamma) \text{ maximizer of (151)} \right\}$$

and analogously for $F'_{\text{RS}}(\gamma^-)$. \square

Lemma 12. *We have*

$$\sup_{\tilde{q} \in [0, \rho]} \left\{ F_{\text{RS}}(2p\lambda\tilde{q}^{2p-1}) - (2p-1)\frac{\lambda}{2}\tilde{q}^{2p} \right\} = \sup_{q \in [0, \rho]} \inf_{r \geq 0} \left\{ \psi_{P_0}(r + 2p\lambda q^{2p-1}) + \alpha \Psi_{P_{\text{out}}}(q) - \frac{rq}{2} - (2p-1)\frac{\lambda}{2}q^{2p} \right\}.$$

Proof. Consider the equality above. The inequality l.h.s \geq r.h.s. is obvious because it suffices to restrict the supremum over $(q, \tilde{q}) \in [0, \rho]^2$ to the supremum over the couples (q, q) for $q \in [0, \rho]$.

Let us prove now the converse inequality. Let us do the change of variable $x = \tilde{q}^{2p-1}$ and define $H(x) := F_{\text{RS}}(2p\lambda x) - (2p-1)\frac{\lambda}{2}x^{2p/(2p-1)}$. F_{RS} is left- and right-differentiable everywhere, so is H . We have

$$H'(x) = 2p\lambda F'_{\text{RS}}(2p\lambda x) - \frac{\lambda}{2}2px^{1/(2p-1)} = p\lambda(2F'_{\text{RS}}(2p\lambda x) - x^{1/(2p-1)}) \quad (157)$$

at the points at which H is differentiable, and analogously for the left- and right-derivatives of H . Let $x \in [0, \rho^{2p-1}]$ be a point at which H achieves its supremum over $[0, \rho^{2p-1}]$. Let us distinguish 3 cases:

- Case 1: $x = 0$. In that case, we have $H'(0^+) \leq 0$ and thus $F'_{\text{RS}}(0^+) \leq 0$. Using Lemma 11, we obtain that the only $q \in [0, \rho]$ that achieves the supremum in (151) is $q = 0 = x^{1/(2p-1)}$.
- Case 2: $0 < x < \rho^{2p-1}$. We have then $H'(x^-) \geq 0$ and $H'(x^+) \leq 0$. Using (157), we deduce that

$$2F'_{\text{RS}}((2p\lambda x)^+) \leq x^{1/(2p-1)} \leq 2F'_{\text{RS}}((2p\lambda x)^-).$$

F_{RS} is convex, so the above inequalities collapses into equalities and we get that F_{RS} is differentiable at $2p\lambda x$ with derivative given by $F'_{\text{RS}}(2p\lambda x) = x^{1/(2p-1)}/2$. Lemma 11 above gives then that the supremum in (151) is achieved uniquely at $q = x^{1/(2p-1)}$.

- Case 3: $x = \rho^{2p-1}$. Using the same arguments than in Case 1, we obtain also $q = x^{1/(2p-1)}$.

Conclusion: In all 3 cases above, $q = x^{1/(2p-1)}$ achieves the supremum in (151). Recall that we used the change of variable $x = \tilde{q}^{2p-1}$. Consequently, if $\tilde{q} \in [0, \rho]$ achieves the supremum of $\tilde{q} \mapsto F_{\text{RS}}(2p\lambda\tilde{q}^{2p-1}) - (2p-1)\lambda\tilde{q}^{2p}$, then \tilde{q} achieves also the supremum in (151). This proves the converse bound. \square

By Lemma 12 and (155) above, we get that

$$f_n \xrightarrow{n \rightarrow \infty} \sup_{q \in [0, \rho]} \inf_{r \geq 0} \left\{ \psi_{P_0}(r + 2p\lambda q^{2p-1}) + \alpha \Psi_{P_{\text{out}}}(q) - \frac{rq}{2} - (2p-1)\frac{\lambda}{2}q^{2p} \right\}.$$

Proposition 10 follows then by rewriting the above limit in terms of mutual information, as we did to deduce Corollary 2 from Theorem 1. \square

Proof of Lemma 10: The proof follows exactly the same steps than the one of Lemma 11, so we omit it for the sake brevity. \square

5.3.2 Lower bound using the generalization error

Let us fix $\alpha \in D^*$. The sequence of the overlaps $(Q_n)_{n \geq 1}$ is tight (because bounded in L^1). By Prokhorov's Theorem we know that the sequence of the laws of $(Q_n)_{n \geq 1}$ is relatively compact. We can thus consider a subsequence along which it converges in law, to some random variable Q . In order to simplify the notations (and because working with an extraction does not change the proof) we will assume in the sequel that

$$Q_n \xrightarrow[n \rightarrow \infty]{(d)} Q,$$

for some random variable Q . We aim now at showing that $|Q| = q^*(\alpha)$ almost-surely.

Lemma 13 (Upper bound on the overlap). $|Q| \leq q^*(\alpha)$ almost-surely.

Proof. Let $\epsilon > 0$. The set $[0, q^*(\alpha) + \epsilon]$ is closed, so by Portemanteau's Theorem

$$\mathbb{P}(|Q| \leq q^*(\alpha) + \epsilon) \geq \limsup_{n \rightarrow \infty} \mathbb{P}(|Q_n| \leq q^*(\alpha) + \epsilon) = 1,$$

by Proposition 9. So $\mathbb{P}(|Q| \leq q^*(\alpha) + \epsilon) = 1$ for all $\epsilon > 0$ which gives $\mathbb{P}(|Q| \leq q^*(\alpha)) = 1$. \square

We are going to prove the converse lower bound using Theorem 4. Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a continuous bounded function. Theorem 4 gives $\mathcal{E}_{f,n}(\alpha) \xrightarrow[n \rightarrow \infty]{} \mathcal{E}_f(q^*(\alpha))$. The function \mathcal{E}_f can be written as

$$\mathcal{E}_f(q) = \frac{1}{2} \mathbb{E} \left[h_f(\sqrt{q}Z_0 + \sqrt{\rho - q}Z_1, \sqrt{q}Z_0 + \sqrt{\rho - q}Z'_1) \right]$$

where $Z_0, Z_1, Z'_1 \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$ and $h_f : (a, b) \in \mathbb{R}^2 \mapsto \int (f(y_1) - f(y_2))^2 P_{\text{out}}(y_1|a) P_{\text{out}}(y_2|b) dy_1 dy_2$. By a central limit argument, we have:

Lemma 14.

$$\left(\frac{\mathbf{x} \cdot \Phi_{\text{new}}}{\sqrt{n}}, \frac{\mathbf{X}^* \cdot \Phi_{\text{new}}}{\sqrt{n}} \right) \xrightarrow[n \rightarrow \infty]{(d)} (Z_1, Z_2),$$

where (Z_1, Z_2) is sampled, conditionally on Q , from $\mathcal{N}\left(0, \begin{pmatrix} \rho & Q \\ Q & \rho \end{pmatrix}\right)$.

Proof. Notice that \mathbf{x} and \mathbf{X} are independent of Φ_{new} . If $(\Phi_{\text{new},1}, \dots, \Phi_{\text{new},n}) \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$, then Lemma 14 is obvious because in that case

$$\left(\frac{\mathbf{x} \cdot \Phi_{\text{new}}}{\sqrt{n}}, \frac{\mathbf{X}^* \cdot \Phi_{\text{new}}}{\sqrt{n}} \right) \sim \mathcal{N}\left(0, \frac{1}{n} \begin{pmatrix} \|\mathbf{x}\|^2 & \mathbf{x} \cdot \mathbf{X}^* \\ \mathbf{x} \cdot \mathbf{X}^* & \|\mathbf{X}^*\|^2 \end{pmatrix}\right) \quad \text{and} \quad \frac{1}{n} \begin{pmatrix} \|\mathbf{x}\|^2 & \mathbf{x} \cdot \mathbf{X}^* \\ \mathbf{x} \cdot \mathbf{X}^* & \|\mathbf{X}^*\|^2 \end{pmatrix} \xrightarrow[n \rightarrow \infty]{(d)} \begin{pmatrix} \rho & Q \\ Q & \rho \end{pmatrix}.$$

Let us now suppose that the entries of Φ_{new} are not i.i.d. standard Gaussian (but still verify hypothesis (h3)). Let $g_1, \dots, g_n \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$. Let $L : \mathbb{R}^2 \rightarrow \mathbb{R}$ be a bounded \mathcal{C}^3 function, with bounded partial derivatives. We have to show that

$$\mathbb{E} \left[L \left(\frac{\mathbf{x} \cdot \Phi_{\text{new}}}{\sqrt{n}}, \frac{\mathbf{X}^* \cdot \Phi_{\text{new}}}{\sqrt{n}} \right) \right] \xrightarrow[n \rightarrow \infty]{} \mathbb{E} [L(Z_1, Z_2)]. \quad (158)$$

We have seen above that $\mathbb{E}\left[L\left(\frac{\mathbf{x}\cdot\mathbf{g}}{\sqrt{n}}, \frac{\mathbf{X}^*\cdot\mathbf{g}}{\sqrt{n}}\right)\right] \xrightarrow[n\rightarrow\infty]{} \mathbb{E}[L(Z_1, Z_2)]$. We now apply Theorem 5 (Theorem 2 from [95]) conditionally on \mathbf{x}, \mathbf{X}^* to obtain

$$\mathbb{E}\left[L\left(\frac{\mathbf{x}\cdot\Phi_{\text{new}}}{\sqrt{n}}, \frac{\mathbf{X}^*\cdot\Phi_{\text{new}}}{\sqrt{n}}\right)\right] = \mathbb{E}\left[L\left(\frac{\mathbf{x}\cdot\mathbf{g}}{\sqrt{n}}, \frac{\mathbf{X}^*\cdot\mathbf{g}}{\sqrt{n}}\right)\right] + \mathcal{O}_n(n^{-1/2}),$$

which proves (158) and therefore Lemma 14. \square

Proposition 11. *We have*

$$\mathcal{E}_{f,n}(\alpha) \xrightarrow[n\rightarrow\infty]{} \frac{1}{2}\mathbb{E}[h_f(Z_1, Z_2)],$$

where (Z_1, Z_2) is defined in Lemma 14 above.

Proof. We have

$$\begin{aligned} \mathcal{E}_{f,n} &= \mathbb{E}\left[\left(f(Y_{\text{new}}) - \mathbb{E}[f(Y_{\text{new}})|\Phi_{\text{new}}, \Phi, \mathbf{Y}]\right)^2\right] \\ &= \frac{1}{2}\mathbb{E}\left[\int (f(y_{\text{new}}) - f(y))^2 P_{\text{out}}(y_{\text{new}}|\Phi_{\text{new}} \cdot \mathbf{X}^*/\sqrt{n}) P_{\text{out}}(y|\Phi_{\text{new}} \cdot \mathbf{x}/\sqrt{n}) dy_{\text{new}} dy\right] \\ &= \frac{1}{2}\mathbb{E}\left[h_f\left(\frac{\mathbf{x}\cdot\Phi_{\text{new}}}{\sqrt{n}}, \frac{\mathbf{X}^*\cdot\Phi_{\text{new}}}{\sqrt{n}}\right)\right]. \end{aligned}$$

By Lemma 14 above, we have $\left(\frac{\mathbf{x}\cdot\Phi_{\text{new}}}{\sqrt{n}}, \frac{\mathbf{X}^*\cdot\Phi_{\text{new}}}{\sqrt{n}}\right) \xrightarrow[n\rightarrow\infty]{(d)} (Z_1, Z_2)$. Using (h4) (and the fact that either (h5.a) or (h5.b) hold), we can find a Bore set $S \subset \mathbb{R}$ of full Lebesgue's measure such that $x \mapsto P_{\text{out}}(y|x)$ is continuous on S , for all $y \in \mathbb{R}$. By dominated convergence (recall that f is assumed to be bounded), we obtain that h_f is continuous on $S \times S$. The set of discontinuity points of h_f has thus zero measure for the law of (Z_1, Z_2) . Indeed if we condition on Q :

- if $|Q| < \rho$, then (Z_1, Z_2) has a density over \mathbb{R}^2 .
- if $Q = \rho$, then $Z_1 = Z_2$ almost surely, but h_f is continuous on $\{(s, s) \mid s \in S\}$ that has full Lebesgue's measure on the diagonal $\{(x, x) \mid x \in \mathbb{R}\}$.
- if $Q = -\rho$, then $Z_1 = -Z_2$ almost surely and we use then similar arguments as for the previous point.

We have therefore:

$$h_f\left(\frac{\mathbf{x}\cdot\Phi_{\text{new}}}{\sqrt{n}}, \frac{\mathbf{X}^*\cdot\Phi_{\text{new}}}{\sqrt{n}}\right) \xrightarrow[n\rightarrow\infty]{(d)} h_f(Z_1, Z_2),$$

and Lemma 11 follows from the fact that h_f is bounded. \square

Let us now define:

$$H_f : \begin{cases} [-\rho, \rho] & \rightarrow \mathbb{R} \\ q & \mapsto \frac{1}{2}\mathbb{E}[h_f(G^{(q)})] \end{cases} \quad (159)$$

where $G^{(q)} \sim \mathcal{N}(0, \begin{pmatrix} \rho & q \\ q & \rho \end{pmatrix})$. Notice that H_f is equal to the function \mathcal{E}_f on $[0, \rho]$. By Proposition 11 above and Theorem 4, we have:

$$H_f(q^*(\alpha)) = \lim_{n\rightarrow\infty} \mathcal{E}_{f,n}(\alpha) = \mathbb{E}[H_f(Q)]. \quad (160)$$

Lemma 15. *For all $q \in [-\rho, \rho]$, $H_f(q) \geq H_f(|q|)$.*

Proof. Let $q \in [0, \rho]$ and $Z_0, Z_1, Z'_1 \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$.

$$H_f(-q) = \frac{1}{2}\mathbb{E}\left[h_f(\sqrt{q}Z_0 + \sqrt{\rho-q}Z_1, -\sqrt{q}Z_0 + \sqrt{\rho-q}Z'_1)\right].$$

Let us denote by \mathbb{E}_{Z_1} and $\mathbb{E}_{Z'_1}$ the expectations with respect to Z_1 and Z'_1 . By replacing h_f by its expression, we have

$$\begin{aligned} H_f(-q) &= \frac{1}{2} \mathbb{E} \int (f(y) - f(y'))^2 P_{\text{out}}(y | \sqrt{q}Z_0 + \sqrt{\rho - q}Z_1) P_{\text{out}}(y' | -\sqrt{q}Z_0 + \sqrt{\rho - q}Z'_1) dy dy' \\ &= \frac{1}{2} \mathbb{E} \int (f(y) - f(y'))^2 \mathbb{E}_{Z_1} P_{\text{out}}(y | \sqrt{q}Z_0 + \sqrt{\rho - q}Z_1) \mathbb{E}_{Z'_1} P_{\text{out}}(y' | -\sqrt{q}Z_0 + \sqrt{\rho - q}Z'_1) dy dy' \\ &= \frac{1}{2} \mathbb{E} \int (f(y) - f(y'))^2 \tilde{P}_{\text{out}}(y | Z_0) \tilde{P}_{\text{out}}(y' | -Z_0) dy dy', \end{aligned}$$

where $\tilde{P}_{\text{out}}(y | z) = \mathbb{E}_{Z_1} P_{\text{out}}(y | \sqrt{q}z + \sqrt{\rho - q}Z_1)$. Let now Y and Y' be two random variables that are independent conditionally on Z_0 and distributed as

$$Y \sim \tilde{P}_{\text{out}}(\cdot | Z_0) \quad \text{and} \quad Y' \sim \tilde{P}_{\text{out}}(\cdot | -Z_0).$$

Then we have

$$\begin{aligned} H_f(-q) &= \frac{1}{2} \mathbb{E} \left[(f(Y) - f(Y'))^2 \right] = \frac{1}{2} \mathbb{E} \left[(f(Y) - \mathbb{E}[f(Y) | Z_0] + \mathbb{E}[f(Y) | Z_0] - f(Y'))^2 \right] \\ &= \frac{1}{2} \mathbb{E} \left[(f(Y) - \mathbb{E}[f(Y) | Z_0])^2 \right] + \frac{1}{2} \mathbb{E} \left[(\mathbb{E}[f(Y) | Z_0] - f(Y'))^2 \right], \end{aligned}$$

because Y and Y' are independent conditionally on Z_0 . The conditional expectation $\mathbb{E}[f(Y) | Z_0]$ is Z_0 -measurable, therefore $\mathbb{E} \left[(\mathbb{E}[f(Y) | Z_0] - f(Y'))^2 \right] \geq \mathbb{E} \left[(\mathbb{E}[f(Y') | Z_0] - f(Y'))^2 \right] = \mathbb{E} \left[(\mathbb{E}[f(Y) | Z_0] - f(Y))^2 \right]$. We conclude

$$H_f(-q) \geq \mathbb{E} \left[(\mathbb{E}[f(Y) | Z_0] - f(Y))^2 \right] = H_f(q).$$

□

We have now all the tools needed to prove Theorem 2. Using Lemma 15 and (160) above, we get that $\mathbb{E}H_f(|Q|) \leq \mathbb{E}H_f(Q) = H_f(q^*(\alpha))$. Since H_f is equal to \mathcal{E}_f on $[0, \rho]$ this gives

$$\mathbb{E}[\mathcal{E}_f(|Q|)] \leq \mathcal{E}_f(q^*(\alpha)). \tag{161}$$

If $q^*(\alpha) = 0$, then Theorem 2 follows simply from Proposition 9. We suppose now that $q^*(\alpha) > 0$ and consider $\epsilon \in (0, q^*(\alpha))$. We define $p(\epsilon) = \mathbb{P}(|Q| \leq q^*(\alpha) - \epsilon)$. We are going to show that $p(\epsilon) = 0$. We assumed that P_{out} is informative, so by Proposition 23 and Proposition 24 in Appendix B.2, there exists a continuous bounded function $f : \mathbb{R} \mapsto \mathbb{R}$ such that \mathcal{E}_f is strictly decreasing on $[0, \rho]$. In the following, f is assumed to be such a function. We have

$$\begin{aligned} \mathbb{E}[\mathcal{E}_f(|Q|)] &= \mathbb{E} \left[\mathbf{1}(|Q| \leq q^*(\alpha) - \epsilon) \mathcal{E}_f(|Q|) + \mathbf{1}(|Q| > q^*(\alpha) - \epsilon) \mathcal{E}_f(|Q|) \right] \\ &\geq p(\epsilon) \mathcal{E}_f(q^*(\alpha) - \epsilon) + (1 - p(\epsilon)) \mathcal{E}_f(q^*(\alpha)). \end{aligned}$$

because \mathcal{E}_f is non-increasing and because $|Q| \leq q^*(\alpha)$ almost-surely (Lemma 13). Combining this with (161) leads to

$$p(\epsilon) \mathcal{E}_f(q^*(\alpha)) \geq p(\epsilon) \mathcal{E}_f(q^*(\alpha) - \epsilon).$$

Since \mathcal{E}_f is strictly decreasing: $\mathcal{E}_f(q^*(\alpha)) < \mathcal{E}_f(q^*(\alpha) - \epsilon)$, which implies $p(\epsilon) = 0$. This is true for all $\epsilon > 0$, consequently $|Q| \geq q^*(\alpha)$ almost-surely. We get (using Lemma 13) that

$$|Q| = q^*(\alpha), \quad \text{almost-surely.}$$

We conclude that the only possible limit in law of the tight sequence $(|Q_n|)_{n \geq 1}$ is $q^*(\alpha)$. Therefore $|Q_n| \rightarrow q^*(\alpha)$ in law and in probability because $q^*(\alpha)$ is a constant.

5.4 Denoising error: Proof of Corollary 5

Start by noticing that the denoising error, i.e. the right hand side of (74), is obtained through the I-MMSE theorem, see Proposition 13, applied to $\tilde{i}_n := I(\mathbf{X}^*, \mathbf{A}; \mathbf{Y} | \Phi) / n$:

$$\frac{\partial \tilde{i}_n}{\partial \Delta^{-1}} = \frac{1}{2n} \text{MMSE} \left(\varphi \left(\frac{1}{\sqrt{n}} \Phi \mathbf{X}^*, \mathbf{A} \right) \middle| \Phi, \mathbf{Y} \right). \quad (162)$$

This mutual information is simply computed using our main theorem. Indeed,

$$\tilde{i}_\infty := \lim_{n \rightarrow \infty} \tilde{i}_n = \lim_{n \rightarrow \infty} \frac{1}{n} H(\mathbf{Y} | \Phi) - \lim_{n \rightarrow \infty} \frac{1}{n} H(\mathbf{Y} | \Phi, \mathbf{X}^*, \mathbf{A}) = -f_\infty - \lim_{n \rightarrow \infty} \frac{1}{n} H(\mathbf{Y} | \Phi, \mathbf{X}^*, \mathbf{A}).$$

One can simply check that $\lim_{n \rightarrow \infty} H(\mathbf{Y} | \Phi, \mathbf{X}^*, \mathbf{A}) / n = \alpha \ln(2\pi\Delta e) / 2$ by similar computations as in the proof of Corollary 2. Therefore, defining

$$\tilde{i}_{\text{RS}}(q, \Delta) := -\frac{\alpha}{2} \ln(2\pi\Delta e) - \alpha \Psi_{P_{\text{out}}}(q) - \inf_{r \geq 0} \left\{ \psi_{P_0}(r) - \frac{qr}{2} \right\},$$

we have $\tilde{i}_\infty = \sup_{q \in [0, \rho]} \tilde{i}_{\text{RS}}(q, \Delta)$ from Theorem 1.

One can verify easily that \tilde{i}_n is a concave differentiable function of Δ^{-1} (this is again related to the I-MMSE theorem). Thus its limit \tilde{i}_∞ is also a concave function of Δ^{-1} . Therefore, a standard analysis lemma gives that the derivative of \tilde{i}_n w.r.t. Δ^{-1} converges to the derivative of \tilde{i}_∞ at every point at which \tilde{i}_∞ is differentiable (i.e. almost every points, by concavity): $\lim_{n \rightarrow \infty} \partial_{\Delta^{-1}} \tilde{i}_n = \partial_{\Delta^{-1}} \tilde{i}_\infty = \partial_{\Delta^{-1}} \sup_{q \in [0, \rho]} \tilde{i}_{\text{RS}}(q, \Delta)$. The first limit is given by the limit of the right hand side of (162). It thus remains to compute $\partial_{\Delta^{-1}} \sup_{q \in [0, \rho]} \tilde{i}_{\text{RS}}(q, \Delta)$.

Assume for a moment that the $\partial_{\Delta^{-1}}$ and $\sup_{q \in [0, \rho]}$ operations commute. Then we need to compute $\partial_{\Delta^{-1}} \Psi_{P_{\text{out}}}(q)$; this follows from the I-MMSE theorem. Indeed, if we denote $S = \varphi(\sqrt{q}V + \sqrt{\rho - q}W^*, \mathbf{A})$, notice that $\Psi_{P_{\text{out}}}(q) = -I(S; S + \sqrt{\Delta}Z | V) - \ln(2\pi e\Delta) / 2$, because $\Psi_{P_{\text{out}}}(q) = -H(S + \sqrt{\Delta}Z | V)$ and $I(S + \sqrt{\Delta}Z; S | V) = H(S + \sqrt{\Delta}Z | V) - H(S + \sqrt{\Delta}Z | V, S) = -\Psi_{P_{\text{out}}}(q) - \ln(2\pi e\Delta) / 2$. Therefore

$$\begin{aligned} \frac{\partial \Psi_{P_{\text{out}}}(q)}{\partial \Delta^{-1}} &= \frac{\Delta}{2} - \frac{\partial}{\partial \Delta^{-1}} I(S + \sqrt{\Delta}Z; S | V) = \frac{\Delta}{2} - \frac{1}{2} \text{MMSE}(S | V, S + \sqrt{\Delta}Z) \\ &= \frac{\Delta}{2} - \frac{1}{2} \left(\mathbb{E}[\varphi(\sqrt{\rho}V, \mathbf{A})^2] - \mathbb{E}[\langle \varphi(\sqrt{q}V + \sqrt{\rho - q}w, \mathbf{a}) \rangle_{\text{sc}}^2] \right). \end{aligned}$$

Consequently,

$$\frac{\partial \tilde{i}_{\text{RS}}(q, \Delta)}{\partial \Delta^{-1}} = \frac{\alpha}{2} \left(\mathbb{E}[\varphi(\sqrt{\rho}V, \mathbf{A})^2] - \mathbb{E}[\langle \varphi(\sqrt{q}V + \sqrt{\rho - q}w, \mathbf{a}) \rangle_{\text{sc}}^2] \right).$$

Now, Theorem 1 from [93] gives that at every Δ^{-1} at which \tilde{i}_∞ is differentiable

$$\frac{\partial \tilde{i}_\infty}{\partial \Delta^{-1}} = \frac{\partial}{\partial \Delta^{-1}} \sup_{q \in [0, \rho]} \tilde{i}_{\text{RS}}(q, \Delta) = \frac{\alpha}{2} \left(\mathbb{E}[\varphi(\sqrt{\rho}V, \mathbf{A})^2] - \mathbb{E}[\langle \varphi(\sqrt{q^*}V + \sqrt{\rho - q^*}w, \mathbf{a}) \rangle_{\text{sc}}^2] \right)$$

where $q^* \in [0, \rho]$ is a point where the supremum above is achieved, and thus corresponds to an optimal couple in (43). As explained above, $\lim_{n \rightarrow \infty} \partial_{\Delta^{-1}} \tilde{i}_n = \partial_{\Delta^{-1}} \tilde{i}_\infty$ at every Δ^{-1} at which \tilde{i}_∞ is differentiable, which concludes the proof.

Appendix A: Some technicalities

A.1 The Nishimori identity

Proposition 12 (Nishimori identity). *Let $(\mathbf{X}, \mathbf{Y}) \in \mathbb{R}^{n_1} \times \mathbb{R}^{n_2}$ be a couple of random variables. Let $k \geq 1$ and let $\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(k)}$ be k i.i.d. samples (given \mathbf{Y}) from the conditional distribution $P(\mathbf{X} = \cdot | \mathbf{Y})$, independently of every other random variables. Let us denote $\langle - \rangle$ the expectation operator w.r.t. $P(\mathbf{X} = \cdot | \mathbf{Y})$ and \mathbb{E} the expectation w.r.t. (\mathbf{X}, \mathbf{Y}) . Then, for all continuous bounded function g we have*

$$\mathbb{E}\langle g(\mathbf{Y}, \mathbf{X}^{(1)}, \dots, \mathbf{X}^{(k)}) \rangle = \mathbb{E}\langle g(\mathbf{Y}, \mathbf{X}^{(1)}, \dots, \mathbf{X}^{(k-1)}, \mathbf{X}) \rangle. \quad (163)$$

Proof. This is a simple consequence of Bayes formula. It is equivalent to sample the couple (\mathbf{X}, \mathbf{Y}) according to its joint distribution or to sample first \mathbf{Y} according to its marginal distribution and then to sample \mathbf{X} conditionally to \mathbf{Y} from its conditional distribution $P(\mathbf{X} = \cdot | \mathbf{Y})$. Thus the $(k+1)$ -tuple $(\mathbf{Y}, \mathbf{X}^{(1)}, \dots, \mathbf{X}^{(k)})$ is equal in law to $(\mathbf{Y}, \mathbf{X}^{(1)}, \dots, \mathbf{X}^{(k-1)}, \mathbf{X})$. \square

A.2 Unicity of the optimizer q^* of the replica formula: Proof of Proposition 1

The function

$$h : \alpha \mapsto \inf_{q \in [0, \rho]} \left\{ \alpha \mathcal{I}_{P_{\text{out}}}(q) + \sup_{r \geq 0} \left\{ I_{P_0}(r) - \frac{r}{2}(\rho - q) \right\} \right\} \quad (164)$$

is concave (as an infimum of linear functions). An “envelope” theorem (Corollary 4 from [93]) gives that h is differentiable at α if and only if

$$\left\{ \mathcal{I}_{P_{\text{out}}}(q) \mid q \text{ minimizer of (164)} \right\}$$

is a singleton. We assumed that P_{out} is informative, so Proposition 21 gives that $\mathcal{I}_{P_{\text{out}}}$ is strictly decreasing. We obtain thus that the set of points at which h is differentiable is exactly D^* . Since h is concave, D^* is equal to \mathbb{R}_+^* minus a countable set. Corollary 4 from [93] gives also that $h'(\alpha) = \mathcal{I}_{P_{\text{out}}}(q^*(\alpha))$, for all $\alpha \in D^*$. The function h is concave, so its derivative h' is non-increasing. Since $\mathcal{I}_{P_{\text{out}}}$ is strictly decreasing, we obtain that $\alpha \in D^* \mapsto q^*(\alpha)$ is non-decreasing.

Let now $\alpha_0 \in D^*$. By concavity of h , $h'(\alpha) \rightarrow h'(\alpha_0)$ when $\alpha \in D^* \rightarrow \alpha_0$. Therefore:

$$\mathcal{I}_{P_{\text{out}}}(q^*(\alpha)) \xrightarrow{\alpha \in D^* \rightarrow \alpha_0} \mathcal{I}_{P_{\text{out}}}(q^*(\alpha_0))$$

which implies $q^*(\alpha) \rightarrow q^*(\alpha_0)$ by strict monotonicity of $\mathcal{I}_{P_{\text{out}}}$.

A.3 Continuity properties of the mutual information

We establish in this section two continuity properties of the mutual information, namely Proposition 14 and Corollary 7. Recall definition (33) of the MMSE function. The following proposition comes from [54] and will be repeatedly used in the sequel.

Proposition 13 (I-MMSE theorem, [54, 94]). *Let P_X be a probability distribution over \mathbb{R}^n that admits a finite second moment. Let $\mathbf{X} \sim P_X$ and $\mathbf{Z} \sim \mathcal{N}(0, \mathbf{I}_n)$ be independent random variables. Then the function*

$$I_{P_X} : \begin{cases} \mathbb{R}_+ & \rightarrow & \mathbb{R} \\ \lambda & \mapsto & I(\mathbf{X}; \sqrt{\lambda}\mathbf{X} + \mathbf{Z}) \end{cases}$$

is concave, continuously differentiable over \mathbb{R}_+ , with derivative given by

$$I'_{P_X}(\lambda) = \frac{1}{2} \text{MMSE}(\mathbf{X} | \sqrt{\lambda}\mathbf{X} + \mathbf{Z}) = \frac{1}{2} \mathbb{E} \left[\|\mathbf{X} - \mathbb{E}[\mathbf{X} | \sqrt{\lambda}\mathbf{X} + \mathbf{Z}]\|^2 \right].$$

Remark: We will often apply Proposition 13 in a “conditional fashion”. Let \mathbf{U} be some random variable independent from \mathbf{Z} , then

$$\frac{\partial}{\partial \lambda} I(\mathbf{X}; \sqrt{\lambda}\mathbf{X} + \mathbf{Z} | \mathbf{U}) = \frac{1}{2} \text{MMSE}(\mathbf{X} | \sqrt{\lambda}\mathbf{X} + \mathbf{Z}, \mathbf{U}) = \frac{1}{2} \mathbb{E} \left[\|\mathbf{X} - \mathbb{E}[\mathbf{X} | \sqrt{\lambda}\mathbf{X} + \mathbf{Z}, \mathbf{U}]\|^2 \right].$$

Proposition 14. Let P_1 and P_2 be two probability distributions on \mathbb{R}^n , that admits a finite second moment. We denote by $W_2(P_1, P_2)$ the Wasserstein distance of order 2 between P_1 and P_2 .

$$|I(\mathbf{X}_1; \mathbf{X}_1 + \mathbf{Z}) - I(\mathbf{X}_2; \mathbf{X}_2 + \mathbf{Z})| \leq (\sqrt{\mathbb{E}\|\mathbf{X}_1\|^2} + \sqrt{\mathbb{E}\|\mathbf{X}_2\|^2}) W_2(P_1, P_2).$$

A similar result was proved in [96] but with a weaker bound for the W_2 distance.

Proof. Let $\epsilon > 0$. Let us fix a coupling of $\mathbf{X}_1 \sim P_1$ and $\mathbf{X}_2 \sim P_2$ such that

$$(\mathbb{E}\|\mathbf{X}_1 - \mathbf{X}_2\|^2)^{1/2} \leq W_2(P_1, P_2) + \epsilon.$$

Let us consider for $t_1, t_2 \in [0, 1]$ the observation model

$$\begin{cases} \mathbf{Y}_1^{(t_1)} &= \sqrt{t_1}\mathbf{X}_1 + \mathbf{Z}_1, \\ \mathbf{Y}_2^{(t_2)} &= \sqrt{1-t_2}\mathbf{X}_2 + \mathbf{Z}_2, \end{cases}$$

where $\mathbf{Z}_1, \mathbf{Z}_2 \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \mathbf{I}_n)$ are independent from $(\mathbf{X}_1, \mathbf{X}_2)$. Define $J(t_1, t_2) = I(\mathbf{X}_1, \mathbf{X}_2; \mathbf{Y}_1^{(t_1)}, \mathbf{Y}_2^{(t_2)})$ and $I(t) = J(t, t)$. Let us now differentiate J with respect to t_1 . Using the chain rule for the mutual information,

$$\begin{aligned} J(t_1, t_2) &= I(\mathbf{X}_1, \mathbf{X}_2; \mathbf{Y}_2^{(t_2)}) + I(\mathbf{X}_1, \mathbf{X}_2; \mathbf{Y}_1^{(t_1)} | \mathbf{Y}_2^{(t_2)}) = I(\mathbf{X}_1, \mathbf{X}_2; \mathbf{Y}_2^{(t_2)}) + I(\mathbf{X}_1; \mathbf{Y}_1^{(t_1)} | \mathbf{Y}_2^{(t_2)}) + I(\mathbf{X}_2; \mathbf{Y}_1^{(t_1)} | \mathbf{X}_1, \mathbf{Y}_2^{(t_2)}) \\ &= I(\mathbf{X}_1, \mathbf{X}_2; \mathbf{Y}_2^{(t_2)}) + I(\mathbf{X}_1; \mathbf{Y}_1^{(t_1)} | \mathbf{Y}_2^{(t_2)}) \end{aligned}$$

because, conditionally on $\mathbf{X}_1, \mathbf{X}_2$ and $\mathbf{Y}_1^{(t_1)}$ are independent. The quantity $I(\mathbf{X}_1, \mathbf{X}_2; \mathbf{Y}_2^{(t_2)})$ does not depend on t_1 , therefore by the “I-MMSE relation” from Proposition 13:

$$\frac{\partial J}{\partial t_1}(t_1, t_2) = \frac{1}{2} \text{MMSE}(\mathbf{X}_1 | \mathbf{Y}_1^{(t_1)}, \mathbf{Y}_2^{(t_2)})$$

and similarly

$$\frac{\partial J}{\partial t_2}(t_1, t_2) = -\frac{1}{2} \text{MMSE}(\mathbf{X}_2 | \mathbf{Y}_1^{(t_1)}, \mathbf{Y}_2^{(t_2)}).$$

Let us write $\mathbf{E}_i = \mathbb{E}[\mathbf{X}_i | \mathbf{Y}_1^{(t)}, \mathbf{Y}_2^{(t)}]$ for $i = 1, 2$, then

$$I'(t) = \frac{1}{2} \text{MMSE}(\mathbf{X}_1 | \mathbf{Y}_1^{(t)}, \mathbf{Y}_2^{(t)}) - \frac{1}{2} \text{MMSE}(\mathbf{X}_2 | \mathbf{Y}_1^{(t)}, \mathbf{Y}_2^{(t)}) = \frac{1}{2} \mathbb{E} \left[\|\mathbf{X}_1 - \mathbf{E}_1\|^2 - \|\mathbf{X}_2 - \mathbf{E}_2\|^2 \right]$$

so that

$$\begin{aligned}
|I'(t)| &= \frac{1}{2} \mathbb{E} [(\|\mathbf{X}_1 - \mathbf{E}_1\| + \|\mathbf{X}_2 - \mathbf{E}_2\|)(\|\mathbf{X}_1 - \mathbf{E}_1\| - \|\mathbf{X}_2 - \mathbf{E}_2\|)] \\
&\leq \frac{1}{2} \mathbb{E} [\|\mathbf{X}_1\|^2 + \|\mathbf{X}_2\|^2]^{1/2} \mathbb{E} [(\|\mathbf{X}_1 - \mathbf{E}_1\| - \|\mathbf{X}_2 - \mathbf{E}_2\|)^2]^{1/2} \\
&\leq \frac{1}{2} \mathbb{E} [\|\mathbf{X}_1\|^2 + \|\mathbf{X}_2\|^2]^{1/2} \mathbb{E} [\|\mathbf{X}_1 - \mathbf{X}_2 + \mathbf{E}_2 - \mathbf{E}_1\|^2]^{1/2} \\
&\leq \frac{1}{2} (\sqrt{\mathbb{E}\|\mathbf{X}_1\|^2} + \sqrt{\mathbb{E}\|\mathbf{X}_2\|^2}) \mathbb{E} [2\|\mathbf{X}_1 - \mathbf{X}_2\|^2 + 2\|\mathbf{E}_2 - \mathbf{E}_1\|^2]^{1/2} \\
&\leq (\sqrt{\mathbb{E}\|\mathbf{X}_1\|^2} + \sqrt{\mathbb{E}\|\mathbf{X}_2\|^2}) (W_2(P_1, P_2) + \epsilon).
\end{aligned}$$

We obtain the result by letting $\epsilon \rightarrow 0$. \square

Proposition 15. *Let P_U be a probability distribution over \mathbb{N}^m that admits a finite second moment. Let $\mathbf{U} \sim P_U$ and $\mathbf{Z} \sim \mathcal{N}(0, \mathbf{I}_m)$ be two independent random variables. Then $H(\mathbf{U}) = -\sum_{\mathbf{n} \in \mathbb{N}^m} P_U(\mathbf{n}) \ln P_U(\mathbf{n})$ is finite and for all $\Delta \in (0, 1]$,*

$$|I(\mathbf{U}; \mathbf{U} + \sqrt{\Delta}\mathbf{Z}) - H(\mathbf{U})| \leq 48me^{-1/(16\Delta)}.$$

Proof. Let us define for $\Delta > 0$, $h(\Delta) = I(\mathbf{U}; \mathbf{U} + \sqrt{\Delta}\mathbf{Z}) = I_{P_U}(\Delta^{-1})$. By Proposition 13 we have for all $\Delta > 0$,

$$h'(\Delta) = -\frac{1}{2\Delta^2} \text{MMSE}(\mathbf{U} | \mathbf{U} + \sqrt{\Delta}\mathbf{Z}). \quad (165)$$

We are now going to upper bound $\text{MMSE}(\mathbf{U} | \mathbf{U} + \sqrt{\Delta}\mathbf{Z})$ by considering the following estimator:

$$\hat{\theta}_i = \arg \min_{u \in \mathbb{N}} |u - U_i + \sqrt{\Delta}Z_i|,$$

for all $i \in \{1, \dots, m\}$. Note that $\hat{\theta}_i$ is well-defined almost-surely since there is a.s. a unique minimizer above. We have

$$\mathbb{P}(\hat{\theta}_i \neq U_i) \leq \mathbb{P}(\sqrt{\Delta}|Z_i| \geq 1/2) = 2\mathbb{P}\left(\mathcal{N}(0, 1) \geq \frac{1}{2\sqrt{\Delta}}\right) \leq 2\frac{1}{\sqrt{2\pi}}2\sqrt{\Delta}e^{-1/(8\Delta)} \leq 2\sqrt{\Delta}e^{-1/(8\Delta)},$$

by usual bounds on the Gaussian cumulative distribution function. We have then

$$\begin{aligned}
\text{MMSE}(\mathbf{U} | \mathbf{U} + \sqrt{\Delta}\mathbf{Z}) &\leq \mathbb{E}\|\mathbf{U} - \hat{\theta}\|^2 = \sum_{i=1}^m \mathbb{E}(U_i - \hat{\theta}_i)^2 = \sum_{i=1}^m \mathbb{E}[\mathbf{1}(\hat{\theta}_i \neq U_i)(U_i - \hat{\theta}_i)^2] \\
&\leq \sum_{i=1}^m 2\mathbb{E}[\mathbf{1}(\hat{\theta}_i \neq U_i)(U_i - (U_i + \sqrt{\Delta}Z_i))^2] + 2\mathbb{E}[\mathbf{1}(\hat{\theta}_i \neq U_i)(U_i + \sqrt{\Delta}Z_i - \hat{\theta}_i)^2] \\
&\leq \sum_{i=1}^m 2\mathbb{E}[\mathbf{1}(\hat{\theta}_i \neq U_i)\Delta Z_i^2] + \frac{1}{2}\mathbb{E}[\mathbf{1}(\hat{\theta}_i \neq U_i)] \\
&\leq \sum_{i=1}^m 2\Delta\mathbb{P}(\hat{\theta}_i \neq U_i)^{1/2}\mathbb{E}[Z_i^4]^{1/2} + \frac{1}{2}\mathbb{P}(\hat{\theta}_i \neq U_i) \\
&\leq me^{-1/(16\Delta)}(2\sqrt{6}\Delta^{5/4} + \sqrt{\Delta}) \leq 6me^{-1/(16\Delta)}
\end{aligned}$$

for $\Delta \leq 1$. Plugging this inequality in (165), we obtain for all $\Delta \in (0, 1]$,

$$|h'(\Delta)| \leq \frac{3m}{\Delta^2}e^{-1/(16\Delta)}. \quad (166)$$

Since $h(1)$ is finite and $\int_0^1 \frac{e^{-1/(16\Delta)}}{\Delta^2} d\Delta < +\infty$ we obtain that

$$\sup_{\Delta \in (0,1]} |h(\Delta)| < +\infty. \quad (167)$$

By definition of h :

$$h(\Delta) = I(\mathbf{U}; \mathbf{U} + \sqrt{\Delta}\mathbf{Z}) = -\frac{m}{2} - \mathbb{E} \ln \sum_{\mathbf{U} \in \mathbb{N}^m} P_U(\mathbf{U}) \exp\left(-\frac{1}{2\Delta} \|\mathbf{U} + \sqrt{\Delta}\mathbf{Z} - \mathbf{U}\|^2\right). \quad (168)$$

By the previous equality and (167), the family of (non-negative) random variables

$$\left(-\ln \sum_{\mathbf{U} \in \mathbb{N}^m} P_U(\mathbf{U}) \exp\left(-\frac{1}{2\Delta} \|\mathbf{U} + \sqrt{\Delta}\mathbf{Z} - \mathbf{U}\|^2\right)\right)_{\Delta \in (0,1]}$$

is bounded in L^1 . Notice that (by dominated convergence)

$$-\ln \sum_{\mathbf{U} \in \mathbb{N}^m} P_U(\mathbf{U}) \exp\left(-\frac{1}{2\Delta} \|\mathbf{U} + \sqrt{\Delta}\mathbf{Z} - \mathbf{U}\|^2\right) \xrightarrow{\Delta \rightarrow 0} -\ln\left(P_U(\mathbf{U})e^{-\frac{1}{2}\|\mathbf{Z}\|^2}\right) = \frac{1}{2}\|\mathbf{Z}\|^2 - \ln P_U(\mathbf{U})$$

almost-surely. This gives (by Fatou's Lemma) that this almost-sure limit is integrable and thus that $H(\mathbf{U}) = -\mathbb{E} \ln P_U(\mathbf{U})$ is finite. Let us now show that $h(\Delta) \xrightarrow{\Delta \rightarrow 0} H(\mathbf{U})$. We have almost-surely

$$\ln\left(P_U(\mathbf{U})e^{-\frac{1}{2}\|\mathbf{Z}\|^2}\right) \leq \ln \sum_{\mathbf{U} \in \mathbb{N}^m} P_U(\mathbf{U}) \exp\left(-\frac{1}{2\Delta} \|\mathbf{U} + \sqrt{\Delta}\mathbf{Z} - \mathbf{U}\|^2\right) \leq 0.$$

Since we now know that the left-hand side is integrable (because $H(\mathbf{U})$ is finite), we can apply the dominated convergence theorem to obtain that

$$\mathbb{E} \ln \sum_{\mathbf{U} \in \mathbb{N}^m} P_U(\mathbf{U}) \exp\left(-\frac{1}{2\Delta} \|\mathbf{U} + \sqrt{\Delta}\mathbf{Z} - \mathbf{U}\|^2\right) \xrightarrow{\Delta \rightarrow 0} \mathbb{E} \ln\left(P_U(\mathbf{U})e^{-\frac{1}{2}\|\mathbf{Z}\|^2}\right) = H(\mathbf{U}) - \frac{m}{2},$$

which combined with (168) gives $h(\Delta) \xrightarrow{\Delta \rightarrow 0} H(\mathbf{U})$. Now, using the bound on the derivative of h (166) we conclude that for all $\Delta \in (0, 1]$,

$$|h(\Delta) - H(\mathbf{U})| \leq 3m \int_0^\Delta \frac{e^{-1/(16t)}}{t^2} dt = 3m \left[16e^{-1/(16t)}\right]_0^\Delta = 48me^{-1/(16\Delta)}.$$

□

Corollary 7. *Let \mathbf{U} be a random variable over \mathbb{N}^m with finite second moment, let \mathbf{X} be a random variable over \mathbb{R}^n and let $\mathbf{Z} \sim \mathcal{N}(0, \mathbf{I}_m)$. We assume (\mathbf{U}, \mathbf{X}) to be independent from \mathbf{Z} . Then, for all $\Delta \in (0, 1]$,*

$$|I(\mathbf{X}; \mathbf{U} + \sqrt{\Delta}\mathbf{Z}) - I(\mathbf{X}; \mathbf{U})| \leq 100me^{-1/(16\Delta)}.$$

Proof. We have by the chain rule of the mutual information:

$$I(\mathbf{U}; \mathbf{U} + \sqrt{\Delta}\mathbf{Z}) = I(\mathbf{U}, \mathbf{X}; \mathbf{U} + \sqrt{\Delta}\mathbf{Z}) = I(\mathbf{X}; \mathbf{U} + \sqrt{\Delta}\mathbf{Z}) + I(\mathbf{U}; \mathbf{U} + \sqrt{\Delta}\mathbf{Z}|\mathbf{X}).$$

By applying Proposition 15 twice, we get

$$|I(\mathbf{U}; \mathbf{U} + \sqrt{\Delta}\mathbf{Z}) - H(\mathbf{U})|, |I(\mathbf{U}; \mathbf{U} + \sqrt{\Delta}\mathbf{Z}|\mathbf{X}) - H(\mathbf{U}|\mathbf{X})| \leq 48me^{-1/(16\Delta)}.$$

Since $I(\mathbf{X}; \mathbf{U}) = H(\mathbf{U}) - H(\mathbf{U}|\mathbf{X})$ we obtain the desired inequality. \square

A.4 A simple consequence of hypotheses (h1)-(h2)-(h3)-(h4)

Proposition 16. *Assume that hypotheses (h1)-(h2)-(h3)-(h4) hold. Then there exists $\eta > 0$ such that*

$$\mathbb{E}[\varphi(\sqrt{\rho}Z, \mathbf{A})^{2+\eta}] < \infty,$$

where the expectation above is with respect to $(Z, \mathbf{A}) \sim \mathcal{N}(0, 1) \otimes P_A$.

Proof. By the Central Limit Theorem (using the fact that the third moments of $(X_i^* \Phi_{1,i})$ are bounded with n , because of hypotheses (h1) and (h3)) we have $([\Phi \mathbf{X}^*]_1 / \sqrt{n}, \mathbf{A}_1) \xrightarrow[n \rightarrow \infty]{(d)} (\sqrt{\rho}G, \mathbf{A}_1)$. This implies that

$$\varphi\left(\frac{[\Phi \mathbf{X}^*]_1}{\sqrt{n}}, \mathbf{A}_1\right) \xrightarrow[n \rightarrow \infty]{(d)} \varphi(\sqrt{\rho}G, \mathbf{A}_1), \quad (169)$$

because $\varphi(\cdot, \mathbf{A}_1)$ is almost-surely continuous almost-everywhere, by assumption (h4). The sequence of random variables $(\varphi([\Phi \mathbf{X}^*]_1 / \sqrt{n}, \mathbf{A}_1))_n$ is by assumption (h2) bounded in $L^{2+\eta}$ for some $\eta > 0$. By (169) we conclude that $\mathbb{E}[\varphi(\sqrt{\rho}G, \mathbf{A}_1)^{2+\eta}] < \infty$. \square

A.5 Derivative of the interpolating free entropy: Proof of Proposition 3

Recall $u'_y(x)$ is the x -derivative of $u_y(x) = \ln P_{\text{out}}(y|x)$. Moreover denote $P'_{\text{out}}(y|x)$ and $P''_{\text{out}}(y|x)$ the first and second x -derivatives, respectively, of $P_{\text{out}}(y|x)$. We will first prove that for all $t \in (0, 1)$

$$\frac{df_{n,\epsilon}(t)}{dt} = -\frac{1}{2} \mathbb{E} \left\langle \left(\frac{1}{n} \sum_{\mu=1}^m u'_{Y_{t,\mu}}(S_{t,\mu}) u'_{Y_{t,\mu}}(s_{t,\mu}) - r(t) \right) (Q - q(t)) \right\rangle_{n,t,\epsilon} + \frac{r(t)}{2} (q(t) - \rho) - \frac{A_n}{2}, \quad (170)$$

where recall $Q := \sum_{i=1}^n X_i^* x_i / n$ and

$$A_{n,\epsilon} := \mathbb{E} \left[\frac{1}{\sqrt{n}} \sum_{\mu=1}^m \frac{P''_{\text{out}}(Y_{t,\mu}|S_{t,\mu})}{P_{\text{out}}(Y_{t,\mu}|S_{t,\mu})} \left(\frac{1}{\sqrt{n}} \sum_{i=1}^n ((X_i^*)^2 - \rho) \right) \frac{1}{n} \ln Z_{t,\epsilon} \right]. \quad (171)$$

Once this is done, we will prove that $A_{n,\epsilon}$ goes to 0 as $n \rightarrow \infty$ uniformly in $t \in [0, 1]$, in order to obtain Proposition 3.

A.5.1 Proof of (170)

Recall definition (98) which becomes, when written as a function of the interpolating Hamiltonian (95),

$$f_{n,\epsilon}(t) = \frac{1}{n} \mathbb{E}_{\Phi, \mathbf{V}} \int d\mathbf{Y}_t d\mathbf{Y}'_t dP_0(\mathbf{X}^*) \mathcal{D}\mathbf{W}^* e^{-\mathcal{H}_{t,\epsilon}(\mathbf{X}^*, \mathbf{W}^*; \mathbf{Y}_t, \mathbf{Y}'_t, \Phi, \mathbf{V})} \ln \int dP_0(\mathbf{x}) \mathcal{D}\mathbf{w} e^{-\mathcal{H}_{t,\epsilon}(\mathbf{x}, \mathbf{w}; \mathbf{Y}_t, \mathbf{Y}'_t, \Phi, \mathbf{V})}. \quad (172)$$

We will need the Hamiltonian t -derivative $\mathcal{H}'_{t,\epsilon}$ given by

$$\mathcal{H}'_{t,\epsilon}(\mathbf{X}^*, \mathbf{W}^*; \mathbf{Y}_t, \mathbf{Y}'_t, \Phi, \mathbf{V}) = -\sum_{\mu=1}^m \frac{dS_{t,\mu}}{dt} u'_{Y_{t,\mu}}(S_{t,\mu}) - \frac{r(t)}{2\sqrt{R_1(t)}} \sum_{i=1}^n X_i^* (Y'_{t,i} - \sqrt{R_1(t)} X_i^*). \quad (173)$$

The derivative of the interpolating free entropy thus reads, for $0 < t < 1$,

$$\frac{df_{n,\epsilon}(t)}{dt} = - \underbrace{\frac{1}{n} \mathbb{E}[\mathcal{H}'_{t,\epsilon}(\mathbf{X}^*, \mathbf{W}^*; \mathbf{Y}_t, \mathbf{Y}'_t, \Phi, \mathbf{V}) \ln \mathcal{Z}_{t,\epsilon}]}_{T_1} - \underbrace{\frac{1}{n} \mathbb{E}\langle \mathcal{H}'_{t,\epsilon}(\mathbf{x}, \mathbf{w}; \mathbf{Y}_t, \mathbf{Y}'_t, \Phi, \mathbf{V}) \rangle_{n,t,\epsilon}}_{T_2} \quad (174)$$

where recall the definition of $\mathcal{Z}_{t,\epsilon} = \mathcal{Z}_{t,\epsilon}(\mathbf{Y}_t, \mathbf{Y}'_t, \Phi, \mathbf{V})$ given by (97).

Let us compute T_1 . Let $1 \leq \mu \leq m$. Let us start with the following term

$$\begin{aligned} & \mathbb{E} \left[\frac{dS_{t,\mu}}{dt} u'_{Y_{t,\mu}}(S_{t,\mu}) \ln \mathcal{Z}_{t,\epsilon} \right] \\ &= \frac{1}{2} \mathbb{E} \left[\left(- \frac{[\Phi \mathbf{X}^*]_\mu}{\sqrt{n(1-t)}} + \frac{q(t)}{\sqrt{R_2(t)}} V_\mu + \frac{\rho - q(t)}{\sqrt{\rho t - R_2(t)} + 2s_n} W_\mu^* \right) u'_{Y_{t,\mu}}(S_{t,\mu}) \ln \mathcal{Z}_{t,\epsilon} \right]. \end{aligned} \quad (175)$$

Let us compute the first term of the right-hand side of the last identity. By Gaussian integration by parts w.r.t $\Phi_{\mu i}$ we obtain

$$\begin{aligned} & \frac{1}{\sqrt{n(1-t)}} \mathbb{E} [[\Phi \mathbf{X}^*]_\mu u'_{Y_{t,\mu}}(S_{t,\mu}) \ln \mathcal{Z}_{t,\epsilon}] \\ &= \frac{1}{\sqrt{n(1-t)}} \sum_{i=1}^n \mathbb{E} \left[\int d\mathbf{Y}_t d\mathbf{Y}'_t e^{-\mathcal{H}_{t,\epsilon}(\mathbf{X}^*, \mathbf{W}^*; \mathbf{Y}_t, \mathbf{Y}'_t, \Phi, \mathbf{V})} \Phi_{\mu i} X_i^* u'_{Y_{t,\mu}}(S_{t,\mu}) \ln \mathcal{Z}_{t,\epsilon} \right] \\ &= \frac{1}{n} \sum_{i=1}^n \left(\mathbb{E} [(X_i^*)^2 (u''_{Y_{t,\mu}}(S_{t,\mu}) + u'_{Y_{t,\mu}}(S_{t,\mu})^2) \ln \mathcal{Z}_{t,\epsilon}] + \mathbb{E} \langle X_i^* x_i u'_{Y_{t,\mu}}(S_{t,\mu}) u'_{Y_{t,\mu}}(s_{t,\mu}) \rangle_{n,t,\epsilon} \right) \\ &= \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n (X_i^*)^2 \frac{P''_{\text{out}}(Y_{t,\mu}|S_{t,\mu})}{P_{\text{out}}(Y_{t,\mu}|S_{t,\mu})} \ln \mathcal{Z}_{t,\epsilon} \right] + \mathbb{E} \left\langle \frac{1}{n} \sum_{i=1}^n X_i^* x_i u'_{Y_{t,\mu}}(S_{t,\mu}) u'_{Y_{t,\mu}}(s_{t,\mu}) \right\rangle_{n,t,\epsilon}, \end{aligned} \quad (176)$$

where we used the identity

$$u''_{Y_{t,\mu}}(x) + u'_{Y_{t,\mu}}(x)^2 = \frac{P''_{\text{out}}(Y_{t,\mu}|x)}{P_{\text{out}}(Y_{t,\mu}|x)}. \quad (177)$$

We now compute the second term of the right hand side of (175). Using again Gaussian integrations by parts but this time w.r.t $V_\mu, W_\mu^* \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$ as well as the previous formula, we obtain similarly

$$\begin{aligned} & \mathbb{E} \left[\left(\frac{q(t)}{\sqrt{R_2(t)}} V_\mu + \frac{\rho - q(t)}{\sqrt{\rho t - R_2(t)} + 2s_n} W_\mu^* \right) u'_{Y_{t,\mu}}(S_{t,\mu}) \ln \mathcal{Z}_{t,\epsilon} \right] \\ &= \mathbb{E} \left[\int d\mathbf{Y}_t d\mathbf{Y}'_t e^{-\mathcal{H}_{t,\epsilon}(\mathbf{X}^*, \mathbf{W}^*; \mathbf{Y}_t, \mathbf{Y}'_t, \Phi, \mathbf{V})} \left(\frac{q(t)}{\sqrt{R_2(t)}} V_\mu + \frac{\rho - q(t)}{\sqrt{\rho t - R_2(t)} + 2s_n} W_\mu^* \right) u'_{Y_{t,\mu}}(S_{t,\mu}) \ln \mathcal{Z}_{t,\epsilon} \right] \\ &= \mathbb{E} \left[\rho \frac{P''_{\text{out}}(Y_{t,\mu}|S_{t,\mu})}{P_{\text{out}}(Y_{t,\mu}|S_{t,\mu})} \ln \mathcal{Z}_{t,\epsilon} \right] + \mathbb{E} \langle q(t) u'_{Y_{t,\mu}}(S_{t,\mu}) u'_{Y_{t,\mu}}(s_{t,\mu}) \rangle_{n,t,\epsilon}. \end{aligned} \quad (178)$$

Combining equations (175), (176) and (178) together, we have

$$\begin{aligned} & - \mathbb{E} \left[\frac{dS_{t,\mu}}{dt} u'_{Y_{t,\mu}}(S_{t,\mu}) \ln \mathcal{Z}_{t,\epsilon} \right] \\ &= \frac{1}{2} \mathbb{E} \left[\frac{P''_{\text{out}}(Y_{t,\mu}|S_{t,\mu})}{P_{\text{out}}(Y_{t,\mu}|S_{t,\mu})} \left(\frac{1}{n} \sum_{i=1}^n (X_i^*)^2 - \rho \right) \ln \mathcal{Z}_{t,\epsilon} \right] + \frac{1}{2} \mathbb{E} \left\langle \left(\frac{1}{n} \sum_{i=1}^n X_i^* x_i - q(t) \right) u'_{Y_{t,\mu}}(S_{t,\mu}) u'_{Y_{t,\mu}}(s_{t,\mu}) \right\rangle_{n,t,\epsilon}. \end{aligned}$$

As seen from (173), (174) it remains to compute $\mathbb{E}[X_j^*(Y'_{t,j} - \sqrt{R_1(t)} X_j^*) \ln \mathcal{Z}_{t,\epsilon}]$. Recalling that for $1 \leq j \leq n$,

$Y'_{t,j} - \sqrt{R_1(t)}X_j^* = Z'_j$ and then using again a Gaussian integration by parts w.r.t $Z'_j \sim \mathcal{N}(0, 1)$ we obtain

$$\begin{aligned}
\mathbb{E}[X_j^*(Y'_{t,j} - \sqrt{R_1(t)}X_j^*) \ln \mathcal{Z}_{t,\epsilon}] &= \mathbb{E}[X_j^* Z'_j \ln \mathcal{Z}_{t,\epsilon}] = \mathbb{E}\left[X_j^* Z'_j \ln \int dP_0(\mathbf{x}) \mathcal{D}\mathbf{w} e^{-\mathcal{H}_{t,\epsilon}(\mathbf{x}, \mathbf{w}; \mathbf{Y}_t, \mathbf{Y}'_t, \Phi, \mathbf{V})}\right] \\
&= \mathbb{E}\left[X_j^* Z'_j \ln \int dP_0(\mathbf{x}) \mathcal{D}\mathbf{w} \exp\left\{\sum_{\mu=1}^m u_{Y_{t,\mu}}(s_{t,\mu}) - \frac{1}{2} \sum_{i=1}^n (\sqrt{R_1(t)}X_i^* + Z'_i - \sqrt{R_1(t)}x_i)^2\right\}\right] \\
&= -\mathbb{E}[X_j^* \langle \sqrt{R_1(t)}(X_j^* - x_j) + Z'_j \rangle_{n,t,\epsilon}] \\
&= -\sqrt{R_1(t)}(\rho - \mathbb{E}\langle X_j^* x_j \rangle_{n,t,\epsilon}). \tag{179}
\end{aligned}$$

Thus, by taking the sum,

$$-\frac{r(t)}{2\sqrt{R_1(t)}} \mathbb{E}\left[\frac{1}{n} \sum_{i=1}^n X_i^*(Y'_{t,i} - \sqrt{R_1(t)}X_i^*) \ln \mathcal{Z}_{t,\epsilon}\right] = \frac{r(t)\rho}{2} - \frac{r(t)}{2} \mathbb{E}\left\langle \frac{1}{n} \sum_{i=1}^n X_i^* x_i \right\rangle_{n,t,\epsilon}. \tag{180}$$

Therefore, for all $t \in (0, 1)$,

$$\begin{aligned}
T_1 &= \frac{1}{2} \mathbb{E}\left[\frac{1}{\sqrt{n}} \sum_{\mu=1}^m \frac{P''_{\text{out}}(Y_{t,\mu}|S_{t,\mu})}{P_{\text{out}}(Y_{t,\mu}|S_{t,\mu})} \left(\frac{1}{\sqrt{n}} \sum_{i=1}^n ((X_i^*)^2 - \rho)\right) \frac{1}{n} \ln \mathcal{Z}_{t,\epsilon}\right] + \frac{r(t)\rho}{2} - \frac{r(t)q(t)}{2} \\
&\quad + \frac{1}{2} \mathbb{E}\left\langle \left(\frac{1}{n} \sum_{\mu=1}^m u'_{Y_{t,\mu}}(S_{t,\mu}) u'_{Y_{t,\mu}}(s_{t,\mu}) - r(t)\right) \left(\frac{1}{n} \sum_{i=1}^n X_i^* x_i - q(t)\right) \right\rangle_{n,t,\epsilon}. \tag{181}
\end{aligned}$$

To obtain (170), it remains to show that $T_2 = 0$. This is a direct consequence of the Nishimori identity (see Appendix A.1):

$$T_2 = \frac{1}{n} \mathbb{E}\langle \mathcal{H}'_{t,\epsilon}(\mathbf{x}, \mathbf{w}; \mathbf{Y}_t, \mathbf{Y}'_t, \Phi) \rangle_{n,t,\epsilon} = \frac{1}{n} \mathbb{E}\langle \mathcal{H}'_{t,\epsilon}(\mathbf{X}^*, \mathbf{W}^*; \mathbf{Y}_t, \mathbf{Y}'_t, \Phi) \rangle = 0. \tag{182}$$

For obtaining the Lemma, it remains to show that $A_{n,\epsilon}$ goes to 0 uniformly in $t \in [0, 1]$.

A.5.2 Proof that $A_{n,\epsilon}$ vanishes as $n \rightarrow \infty$

We now consider the final step, that is showing that $A_{n,\epsilon}$ given by (171) vanishes in the $n \rightarrow \infty$ limit uniformly in $t \in [0, 1]$ under conditions (H1)-(H2)-(H3). First we show that

$$\mathbb{E}\left[\frac{1}{\sqrt{n}} \sum_{\mu=1}^m \frac{P''_{\text{out}}(Y_{t,\mu}|S_{t,\mu})}{P_{\text{out}}(Y_{t,\mu}|S_{t,\mu})} \left(\frac{1}{\sqrt{n}} \sum_{i=1}^n ((X_i^*)^2 - \rho)\right)\right] = 0. \tag{183}$$

Once this is done, we use the fact that $\frac{1}{n} \ln \mathcal{Z}_{t,\epsilon}$ concentrates around $f_{n,\epsilon}(t)$ to prove that $A_{n,\epsilon}$ converges to 0 as $n \rightarrow \infty$. We start by noticing the simple fact that for all $s \in \mathbb{R}$, $\int P''_{\text{out}}(y|s) dy = 0$. Consequently, for $\mu \in \{1, \dots, m\}$,

$$\mathbb{E}\left[\frac{P''_{\text{out}}(Y_{t,\mu}|S_{t,\mu})}{P_{\text{out}}(Y_{t,\mu}|S_{t,\mu})} \middle| \mathbf{X}^*, \mathbf{S}_t\right] = \int dY_{t,\mu} P''_{\text{out}}(Y_{t,\mu}|S_{t,\mu}) = 0. \tag{184}$$

Thus, using the ‘‘tower property’’ of the conditionnal expectation:

$$\mathbb{E}\left[\left(\sum_{i=1}^n ((X_i^*)^2 - \rho)\right) \sum_{\mu=1}^m \frac{P''_{\text{out}}(Y_{t,\mu}|S_{t,\mu})}{P_{\text{out}}(Y_{t,\mu}|S_{t,\mu})}\right] = \mathbb{E}\left[\left(\sum_{i=1}^n ((X_i^*)^2 - \rho)\right) \mathbb{E}\left[\sum_{\mu=1}^m \frac{P''_{\text{out}}(Y_{t,\mu}|S_{t,\mu})}{P_{\text{out}}(Y_{t,\mu}|S_{t,\mu})} \middle| \mathbf{X}^*, \mathbf{S}_t\right]\right] = 0$$

which gives (183). We now show that $A_{n,\epsilon}$ goes to 0 uniformly in $t \in [0, 1]$ as $n \rightarrow \infty$. Using successively (183) and the Cauchy-Schwarz inequality, we have

$$\begin{aligned} |A_{n,\epsilon}| &= \left| \mathbb{E} \left[\frac{1}{\sqrt{n}} \sum_{\mu=1}^m \frac{P''_{\text{out}}(Y_{t,\mu}|S_{t,\mu})}{P_{\text{out}}(Y_{t,\mu}|S_{t,\mu})} \left(\frac{1}{\sqrt{n}} \sum_{i=1}^n ((X_i^*)^2 - \rho) \right) \left(\frac{1}{n} \ln \mathcal{Z}_{t,\epsilon} - f_{n,\epsilon}(t) \right) \right] \right| \\ &\leq \mathbb{E} \left[\left(\frac{1}{\sqrt{n}} \sum_{\mu=1}^m \frac{P''_{\text{out}}(Y_{t,\mu}|S_{t,\mu})}{P_{\text{out}}(Y_{t,\mu}|S_{t,\mu})} \right)^2 \left(\frac{1}{\sqrt{n}} \sum_{i=1}^n ((X_i^*)^2 - \rho) \right)^2 \right]^{1/2} \mathbb{E} \left[\left(\frac{1}{n} \ln \mathcal{Z}_{t,\epsilon} - f_{n,\epsilon}(t) \right)^2 \right]^{1/2}. \end{aligned} \quad (185)$$

Using again the ‘‘tower property’’ of conditional expectations

$$\begin{aligned} &\mathbb{E} \left[\left(\sum_{\mu=1}^m \frac{P''_{\text{out}}(Y_{t,\mu}|S_{t,\mu})}{P_{\text{out}}(Y_{t,\mu}|S_{t,\mu})} \right)^2 \left(\sum_{i=1}^n ((X_i^*)^2 - \rho) \right)^2 \right] \\ &= \mathbb{E} \left[\left(\sum_{i=1}^n ((X_i^*)^2 - \rho) \right)^2 \mathbb{E} \left[\left(\sum_{\mu=1}^m \frac{P''_{\text{out}}(Y_{t,\mu}|S_{t,\mu})}{P_{\text{out}}(Y_{t,\mu}|S_{t,\mu})} \right)^2 \middle| \mathbf{X}^*, \mathbf{S}_t \right] \right]. \end{aligned} \quad (186)$$

Now, using the fact that conditionally on \mathbf{S}_t , the random variables $\left(\frac{P''_{\text{out}}(Y_{t,\mu}|S_{t,\mu})}{P_{\text{out}}(Y_{t,\mu}|S_{t,\mu})} \right)_{1 \leq \mu \leq m}$ are i.i.d. and centered, we have

$$\mathbb{E} \left[\left(\sum_{\mu=1}^m \frac{P''_{\text{out}}(Y_{t,\mu}|S_{t,\mu})}{P_{\text{out}}(Y_{t,\mu}|S_{t,\mu})} \right)^2 \middle| \mathbf{X}^*, \mathbf{S}_t \right] = \mathbb{E} \left[\left(\sum_{\mu=1}^m \frac{P''_{\text{out}}(Y_{t,\mu}|S_{t,\mu})}{P_{\text{out}}(Y_{t,\mu}|S_{t,\mu})} \right)^2 \middle| \mathbf{S}_t \right] = m \mathbb{E} \left[\left(\frac{P''_{\text{out}}(Y_1|S_{t,1})}{P_{\text{out}}(Y_1|S_{t,1})} \right)^2 \middle| \mathbf{S}_t \right]. \quad (187)$$

Under condition (H2), it is not difficult to show that there exists a constant $C > 0$ such that

$$\mathbb{E} \left[\left(\frac{P''_{\text{out}}(Y_{t,1}|S_{t,1})}{P_{\text{out}}(Y_{t,1}|S_{t,1})} \right)^2 \middle| \mathbf{S}_t \right] \leq C. \quad (188)$$

Combining now (188), (187) and (186) we obtain that

$$\mathbb{E} \left[\left(\sum_{\mu=1}^m \frac{P''_{\text{out}}(Y_{t,\mu}|S_{t,\mu})}{P_{\text{out}}(Y_{t,\mu}|S_{t,\mu})} \right)^2 \left(\sum_{i=1}^n ((X_i^*)^2 - \rho) \right)^2 \right] \leq mC \mathbb{E} \left[\left(\sum_{i=1}^n ((X_i^*)^2 - \rho) \right)^2 \right] = mnC \text{Var}((X_1^*)^2).$$

Going back to (185), therefore there exists a constant $C' > 0$ such that

$$|A_{n,\epsilon}| \leq C' \mathbb{E} \left[\left(\frac{1}{n} \ln \mathcal{Z}_{t,\epsilon} - f_{n,\epsilon}(t) \right)^2 \right]^{1/2}. \quad (189)$$

By Theorem 6 we have $\mathbb{E}[(n^{-1} \ln \mathcal{Z}_{t,\epsilon} - f_{n,\epsilon}(t))^2] \rightarrow 0$ as $n \rightarrow \infty$ uniformly in $t \in [0, 1]$. Thus $A_{n,\epsilon}$ goes to 0 as $n \rightarrow \infty$ uniformly in $t \in [0, 1]$, ϵ and w.r.t. the choice of the interpolation functions. This ends the proof of Proposition 3.

A.6 Boundedness of an overlap fluctuation

In this appendix we show that the ‘‘overlap fluctuation’’

$$\mathbb{E} \left\langle \left(\frac{1}{n} \sum_{\mu=1}^m u'_{Y_{t,\mu}}(S_{t,\mu}) u'_{Y_{t,\mu}}(s_{t,\mu}) - r_\epsilon(t) \right)^2 \right\rangle_{n,t,\epsilon} \leq 2r_{\max}^2 + 2 \mathbb{E} \left\langle \left(\frac{1}{n} \sum_{\mu=1}^m u'_{Y_{t,\mu}}(S_{t,\mu}) u'_{Y_{t,\mu}}(s_{t,\mu}) \right)^2 \right\rangle_{n,t,\epsilon} \quad (190)$$

is bounded uniformly in t under hypothesis (H2) on φ . From the representation (22) (recall we consider $\Delta = 1$ in Sec. 4)

$$u_{Y_{t,\mu}}(s) = \ln P_{\text{out}}(Y_{t,\mu}|s) = \ln \int dP_A(\mathbf{a}_\mu) \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}(Y_{t,\mu} - \varphi(s, \mathbf{a}_\mu))^2} \quad (191)$$

and thus

$$u'_{Y_{t,\mu}}(s) = \frac{\int dP_A(\mathbf{a}_\mu) (Y_{t,\mu} - \varphi(s, \mathbf{a}_\mu)) \varphi'(s, \mathbf{a}_\mu) e^{-\frac{1}{2}(Y_{t,\mu} - \varphi(s, \mathbf{a}_\mu))^2}}{\int dP_A(\mathbf{a}_\mu) e^{-\frac{1}{2}(Y_{t,\mu} - \varphi(s, \mathbf{a}_\mu))^2}} \quad (192)$$

where φ' is the derivative w.r.t. the first argument. From (20) at $\Delta = 1$ we get $|Y_{t,\mu}| \leq \sup |\varphi| + |Z_\mu|$, where the supremum is taken over both arguments of φ , and thus immediately obtain for all $s \in \mathbb{R}$

$$|u'_{Y_{t,\mu}}(s)| \leq (2 \sup |\varphi| + |Z_\mu|) \sup |\varphi'|. \quad (193)$$

From (193) and (190) we see that it suffices to check that

$$\frac{m^2}{n^2} \mathbb{E} \left[((2 \sup |\varphi| + |Z_\mu|)^2 (\sup |\varphi'|)^2) \right] \leq C(\varphi, \alpha)$$

where $C(\varphi, \alpha)$ is a constant depending only on φ and α . This is easily seen by expanding all squares and using that m/n is bounded.

A.7 Proof of Proposition 6

The continuity and differentiability properties of F_n follow from the standard theorems of continuity and derivation under the integral sign. The domination hypotheses are easily verified because we are working under hypotheses (H1)-(H2).

The overlap $\mathbb{E}\langle Q \rangle_{n,t,\epsilon}$ is related to the minimum mean-square error by

$$\frac{1}{n} \text{MMSE}(\mathbf{X}^* | \mathbf{Y}_t, \mathbf{Y}'_t, \mathbf{V}, \Phi) = \frac{1}{n} \mathbb{E} \left[\|\mathbf{X}^* - \mathbb{E}[\mathbf{X}^* | \mathbf{Y}_t, \mathbf{Y}'_t, \mathbf{V}, \Phi]\|^2 \right] = \frac{1}{n} \mathbb{E} \left[\|\mathbf{X}^* - \langle \mathbf{x} \rangle_{n,t,\epsilon}\|^2 \right] = \rho - \mathbb{E}\langle Q \rangle_{n,t,\epsilon}.$$

Since the left hand side belongs to $[0, \rho]$, we obtain that $\mathbb{E}\langle Q \rangle_{n,t,\epsilon} \in [0, \rho]$.

It remains therefore to prove that $\text{MMSE}(\mathbf{X}^* | \mathbf{Y}_t, \mathbf{Y}'_t, \mathbf{V}, \Phi)$ is separately non-increasing in R_1 and R_2 . R_1 only appears in the definition of \mathbf{Y}'_t . Recall that $\mathbf{Y}'_t = \sqrt{R_1} \mathbf{X}^* + \mathbf{Z}'$, where \mathbf{Z}' is a standard Gaussian vector, so $\text{MMSE}(\mathbf{X}^* | \mathbf{Y}_t, \mathbf{Y}'_t, \mathbf{V}, \Phi)$ is obviously a non-increasing function of R_1 .

R_2 only plays a role in $\mathbf{Y}_t \sim P_{\text{out}}(\cdot | \sqrt{(1-t)/n} \Phi \mathbf{X}^* + \sqrt{R_2} \mathbf{V} + \sqrt{\rho t - R_2 + 2s_n} \mathbf{W}^*)$. Let $0 < r_2 \leq r'_2 < \rho t$. Let $\mathbf{V}' \sim \mathcal{N}(0, \mathbf{I}_m)$, independently of everything else. Define

$$\tilde{\mathbf{Y}}_t \sim P_{\text{out}} \left(\cdot \mid \sqrt{\frac{1-t}{n}} \Phi \mathbf{X}^* + \sqrt{r_2} \mathbf{V} + \sqrt{r'_2 - r_2} \mathbf{V}' + \sqrt{\rho t - r'_2 + 2s_n} \mathbf{W}^* \right),$$

independently of everything else. Now notice that

$$\begin{aligned} \text{MMSE}(\mathbf{X}^* | \mathbf{Y}_t, \mathbf{Y}'_t, \mathbf{V}, \Phi) |_{R_2=r_2} &= \text{MMSE}(\mathbf{X}^* | \tilde{\mathbf{Y}}_t, \mathbf{Y}'_t, \mathbf{V}, \Phi), \\ \text{MMSE}(\mathbf{X}^* | \mathbf{Y}_t, \mathbf{Y}'_t, \mathbf{V}, \Phi) |_{R_2=r'_2} &= \text{MMSE}(\mathbf{X}^* | \tilde{\mathbf{Y}}_t, \mathbf{Y}'_t, \mathbf{V}, \mathbf{V}', \Phi), \end{aligned}$$

which implies of course that $\text{MMSE}(\mathbf{X}^* | \mathbf{Y}_t, \mathbf{Y}'_t, \mathbf{V}, \Phi) |_{R_2=r_2} \geq \text{MMSE}(\mathbf{X}^* | \mathbf{Y}_t, \mathbf{Y}'_t, \mathbf{V}, \Phi) |_{R_2=r'_2}$. We have proved that $\text{MMSE}(\mathbf{X}^* | \mathbf{Y}_t, \mathbf{Y}'_t, \mathbf{V}, \Phi)$ is a non-increasing function of R_2 .

Appendix B: Some properties of the scalar channels

B.1 The additive Gaussian scalar channel

We recall some properties (see [54] and [97] for proofs) of the free entropy of the first scalar channel (36).

Proposition 17. *Let $X_0 \sim P_0$ be a real random variable with finite second moment. Let $r \geq 0$ and $Y_0 = \sqrt{r}X_0 + Z_0$, where $Z_0 \sim \mathcal{N}(0, 1)$ is independent from X_0 . Then the function*

$$\psi_{P_0} : r \mapsto \mathbb{E} \ln \int dP_0(x) e^{\sqrt{r}Y_0x - rx^2/2}$$

is convex, differentiable, non-decreasing and $\frac{1}{2}\mathbb{E}[X_0^2]$ -Lipschitz on \mathbb{R}_+ . Moreover, ψ_{P_0} is strictly convex, if P_0 is not a Dirac measure.

B.2 The non-linear scalar channel

We prove here some properties of the free entropy of the second scalar channel (39), where $V, W^* \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$ and

$$Y^{(q)} \sim P_{\text{out}}(\cdot \mid \sqrt{q}V + \sqrt{\rho - q}W^*). \quad (194)$$

In this channel, the statistician observes V and $Y^{(q)}$ and wants to recover W^* . Recall that by definition $\mathcal{I}_{P_{\text{out}}}(q) = I(W^*; Y^{(q)} \mid V) = \Psi_{P_{\text{out}}}(\rho) - \Psi_{P_{\text{out}}}(q)$ so the properties we will prove on $\Psi_{P_{\text{out}}}$ can be directly translated for $\mathcal{I}_{P_{\text{out}}}$, and vice-versa.

Proposition 18. *Suppose that for all $x \in \mathbb{R}$, $P_{\text{out}}(\cdot \mid x)$ is the law of $\varphi(x, A) + \sqrt{\Delta}Z$ where $\Delta > 0$, $\varphi : \mathbb{R} \times \mathbb{R}^{k_A} \rightarrow \mathbb{R}$ is a measurable function and $(Z, A) \sim \mathcal{N}(0, 1) \otimes P_A$, for some probability distribution P_A over \mathbb{R}^{k_A} . In that case P_{out} admits a density given by*

$$P_{\text{out}}(y \mid x) = \frac{1}{\sqrt{2\pi\Delta}} \int dP_A(\mathbf{a}) e^{-\frac{1}{2\Delta}(y - \varphi(x, \mathbf{a}))^2}.$$

Assume that φ is bounded and \mathcal{C}^2 with respect to its first coordinate, with bounded first and second derivatives. Then $q \mapsto \Psi_{P_{\text{out}}}(q)$ is convex, \mathcal{C}^2 and non-decreasing on $[0, \rho]$.

Proof. Let $V, W^* \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$ and $Y^{(q)}$ be the output of the scalar channel given by (194). Then for all $q \in [0, \rho]$,

$$\Psi_{P_{\text{out}}}(q) = \mathbb{E} \ln \int dw \frac{e^{-\frac{w^2}{2}}}{\sqrt{2\pi}} P_{\text{out}}(Y^{(q)} \mid \sqrt{q}V + \sqrt{\rho - q}w)$$

Under the hypotheses we made on φ , we will be able to use continuity and differentiation under the expectation, because all the domination hypotheses will be easily verified. It is thus easy to check that $\Psi_{P_{\text{out}}}$ is continuous on $[0, \rho]$.

We compute now the first derivative. Recall that $\langle - \rangle_{\text{sc}}$, defined in (72), denotes the posterior distribution of W^* given $Y^{(q)}$. We will use the notation $u_y(x) = \ln P_{\text{out}}(y \mid x)$. For $q \in (0, \rho)$ we have

$$\begin{aligned} \Psi'_{P_{\text{out}}}(q) &= \frac{1}{2} \mathbb{E} \left\langle u'_{Y^{(q)}}(\sqrt{q}V + \sqrt{\rho - q}w) u'_{Y^{(q)}}(\sqrt{q}V + \sqrt{\rho - q}W^*) \right\rangle_{\text{sc}} \\ &= \frac{1}{2} \mathbb{E} \left\langle u'_{Y^{(q)}}(\sqrt{q}V + \sqrt{\rho - q}w) \right\rangle_{\text{sc}}^2 \geq 0, \end{aligned}$$

where $w \sim \langle - \rangle_{\text{sc}}$, independently of everything else. $\Psi_{P_{\text{out}}}$ is therefore non-decreasing. Using the boundedness assumption on φ and its derivatives, it is not difficult to check that $\Psi'_{P_{\text{out}}}$ is indeed bounded.

We will now compute $\Psi''_{P_{\text{out}}}$. To lighten the notations, we write $u'(w)$ for $u'_{Y(q)}(\sqrt{q}V + \sqrt{\rho - q}w)$. We compute

$$\partial_q \mathbb{E} \left\langle u'(w)u'(W^*) \right\rangle_{\text{sc}} = \mathbb{E} \left[\left\langle \left(\frac{1}{2\sqrt{q}}V - \frac{1}{2\sqrt{\rho - q}}W^* \right) u'(W^*) \right\rangle_{\text{sc}} \left\langle u'(w)u'(W^*) \right\rangle_{\text{sc}} \right] \quad (A)$$

$$+ 2\mathbb{E} \left\langle \left(\frac{1}{2\sqrt{q}}V - \frac{1}{2\sqrt{\rho - q}}W^* \right) u''(W^*)u'(w) \right\rangle_{\text{sc}} \quad (B)$$

$$+ \mathbb{E} \left\langle \left(\frac{1}{2\sqrt{q}}V - \frac{1}{2\sqrt{\rho - q}}W^* \right) u'(W^*)^2 u'(w) \right\rangle_{\text{sc}} \quad (C)$$

$$- \mathbb{E} \left\langle u'(W^*)u'(w) \right\rangle_{\text{sc}} \left\langle \left(\frac{1}{2\sqrt{q}}V - \frac{1}{2\sqrt{\rho - q}}w \right) u'(w) \right\rangle_{\text{sc}} \quad (D) \quad (195)$$

Notice that (A) = (C). We compute, using Gaussian integration by parts and the Nishimori identity (Proposition 12)

$$(A) = \frac{1}{2} \mathbb{E} \left[u'(W^*) \left\langle u'(W^*)u''(w) \right\rangle_{\text{sc}} \right] + \frac{1}{2} \mathbb{E} \left[u'(W^*) \left\langle u'(W^*)u'(w)^2 \right\rangle_{\text{sc}} \right] \\ - \frac{1}{2} \mathbb{E} \left[u'(W^*) \left\langle u'(W^*)u'(w) \right\rangle_{\text{sc}} \left\langle u'(w) \right\rangle_{\text{sc}} \right] \quad (196)$$

$$(B) = \mathbb{E} \left\langle u''(W^*)u''(w) \right\rangle_{\text{sc}} + \mathbb{E} \left\langle u''(W^*)u'(w)^2 \right\rangle_{\text{sc}} - \mathbb{E} \left\langle u''(W^*)u'(w) \right\rangle_{\text{sc}} \left\langle u'(w) \right\rangle_{\text{sc}} \quad (197)$$

$$(D) = -\mathbb{E} \left\langle \left(\frac{1}{2\sqrt{q}}V - \frac{1}{2\sqrt{\rho - q}}W^* \right) u'(W^*)u'(w^{(1)})u'(w^{(2)}) \right\rangle_{\text{sc}} \\ = -\mathbb{E} \left\langle u'(W^*)u''(w^{(1)})u'(w^{(2)}) \right\rangle_{\text{sc}} - \mathbb{E} \left\langle u'(W^*)u'(w^{(1)})^2 u'(w^{(2)}) \right\rangle_{\text{sc}} \\ + \mathbb{E} \left\langle u'(W^*)u'(w^{(1)})u'(w^{(2)}) \right\rangle_{\text{sc}} \left\langle u'(w) \right\rangle_{\text{sc}} \quad (198)$$

We now replace (196), (197) and (198) in (195):

$$2\Psi''_{P_{\text{out}}}(q) = \mathbb{E} \left\langle u'(W^*)^2 u''(w) \right\rangle_{\text{sc}} + \mathbb{E} \left\langle u'(W^*)^2 u'(w)^2 \right\rangle_{\text{sc}} - \mathbb{E} \left\langle u'(W^*)^2 u'(w^{(1)})u'(w^{(2)}) \right\rangle_{\text{sc}} \\ + \mathbb{E} \left\langle u''(W^*)u''(w) \right\rangle_{\text{sc}} + \mathbb{E} \left\langle u''(W^*)u'(w)^2 \right\rangle_{\text{sc}} - \mathbb{E} \left\langle u''(W^*)u'(w^{(1)})u'(w^{(2)}) \right\rangle_{\text{sc}} \\ - \mathbb{E} \left\langle u'(W^*)u''(w^{(1)})u'(w^{(2)}) \right\rangle_{\text{sc}} - \mathbb{E} \left\langle u'(W^*)u'(w^{(1)})^2 u'(w^{(2)}) \right\rangle_{\text{sc}} + \mathbb{E} \left\langle u'(w) \right\rangle_{\text{sc}}^4.$$

Using the identity $u''_Y(x) + u'_Y(x)^2 = \frac{P''_{\text{out}}(Y|x)}{P_{\text{out}}(Y|x)}$, this factorizes and gives

$$\Psi''_{P_{\text{out}}}(q) = \frac{1}{2} \mathbb{E} \left[\left(\left\langle \frac{P''_{\text{out}}(Y|\sqrt{q}V + \sqrt{\rho - q}w)}{P_{\text{out}}(Y|\sqrt{q}V + \sqrt{\rho - q}w)} \right\rangle_{\text{sc}} - \left\langle u'_{Y(q)}(\sqrt{q}V + \sqrt{\rho - q}w) \right\rangle_{\text{sc}} \right)^2 \right] \geq 0. \quad (199)$$

$\Psi_{P_{\text{out}}}$ is thus convex on $[0, \rho]$. It is not difficult to verify (by standard arguments of continuity under the integral) that $\Psi''_{P_{\text{out}}}$ is continuous on $[0, \rho]$, which gives that $\Psi_{P_{\text{out}}}$ is \mathcal{C}^2 on its domain. \square

Proposition 19. *Suppose that for all $x \in \mathbb{R}$, $P_{\text{out}}(\cdot | x)$ is the law of $\varphi(x, A) + \sqrt{\Delta}Z$ where $\varphi : \mathbb{R} \times \mathbb{R}^{k_A} \rightarrow \mathbb{R}$ is a measurable function and $(Z, A) \sim \mathcal{N}(0, 1) \otimes P_A$, for some probability distribution P_A over \mathbb{R}^{k_A} . Assume also that*

$$\mathbb{E}[\varphi(\sqrt{\rho}Z, A)^2] < \infty, \quad (200)$$

and that we are in one of the following cases:

- (i) $\Delta > 0$.
- (ii) $\Delta = 0$ and φ takes values in \mathbb{N} .

Then $q \mapsto \Psi_{P_{\text{out}}}(q)$ is continuous, convex and non-decreasing over $[0, \rho]$.

Notice that (200) is for instance verified under hypotheses (h1)-(h2)-(h3)-(h4), see Proposition 16.

Proof. We deduce Proposition 19 from Proposition 18 above by an approximation procedure. Since $\Psi_{P_{\text{out}}} = \Psi_{P_{\text{out}}}(\rho) - \mathcal{I}_{P_{\text{out}}}$, we will work with the mutual information $\mathcal{I}_{P_{\text{out}}}$. Let us define $U^{(q)} = \varphi(\sqrt{q}V + \sqrt{\rho - q}W^*, A)$ and $Y^{(q)} = U^{(q)} + \sqrt{\Delta}Z$.

We start by proving Proposition 19 under the assumption (i). Let $\epsilon > 0$. By density of the \mathcal{C}^∞ functions with compact support in L^2 (see for instance Corollary 4.2.2 from [98]), one can find a \mathcal{C}^∞ function $\widehat{\varphi}$ with compact support, such that

$$\mathbb{E}\left[\left(\varphi(\sqrt{\rho}Z, A) - \widehat{\varphi}(\sqrt{\rho}Z, A)\right)^2\right] \leq \epsilon^2.$$

Let us write $\widehat{U}^{(q)} = \widehat{\varphi}(\sqrt{q}V + \sqrt{\rho - q}W^*, A)$ and $\widehat{Y}^{(q)} = \widehat{U}^{(q)} + \sqrt{\Delta}Z$. We have by the chain rule for the mutual information

$$\begin{aligned} I(U^{(q)}; Y^{(q)}|V) &= I(W^*, U^{(q)}; Y^{(q)}|V) \\ &= I(U^{(q)}; Y^{(q)}|V, W^*) + I(W^*; Y^{(q)}|V) = I(U^{(q)}; Y^{(q)}|V, W^*) + \mathcal{I}_{P_{\text{out}}}(q) \end{aligned} \quad (201)$$

and similarly, $\mathcal{I}_{\widehat{P}_{\text{out}}}(q) = I(\widehat{U}^{(q)}; \widehat{Y}^{(q)}|V) - I(\widehat{U}^{(q)}; \widehat{Y}^{(q)}|V, W^*)$. By Proposition 14, there exists a constant $C > 0$ such that

$$|I(\widehat{U}^{(q)}; \widehat{Y}^{(q)}|V) - I(U^{(q)}; Y^{(q)}|V)| \leq C\epsilon \quad \text{and} \quad |I(\widehat{U}^{(q)}; \widehat{Y}^{(q)}|V, W^*) - I(U^{(q)}; Y^{(q)}|V, W^*)| \leq C\epsilon.$$

We get that for all $q \in [0, \rho]$, $|\mathcal{I}_{P_{\text{out}}}(q) - \mathcal{I}_{\widehat{P}_{\text{out}}}(q)| \leq C\epsilon$. The function $\mathcal{I}_{P_{\text{out}}}$ can therefore be uniformly approximated by continuous, concave, non-increasing functions on $[0, \rho]$: $\mathcal{I}_{P_{\text{out}}}$ is therefore continuous, concave and non-increasing.

Let us now prove Proposition 19 under the assumption (ii). Under this assumption we have $\mathcal{I}_{P_{\text{out}}}(q) = I(W^*; U^{(q)}|V)$ and by the case (i) we know that the function $i_\Delta(q) = I(W^*; U^{(q)} + \sqrt{\Delta}Z|V)$ is concave and non-increasing for all $\Delta > 0$. By Corollary 7 we obtain that for all $q \in [0, \rho]$ and all $\Delta \in (0, 1]$ we have

$$|\mathcal{I}_{P_{\text{out}}}(q) - i_\Delta(q)| \leq 100e^{-1/(16\Delta)},$$

which proves (by uniform approximation) that $\mathcal{I}_{P_{\text{out}}}$ is continuous, concave and non-increasing. \square

Proposition 20. *Under the same hypotheses than Proposition 19 above, Ψ_{out} is differentiable over $[0, \rho]$ and for all $q \in [0, \rho]$*

$$\Psi'_{P_{\text{out}}}(q) = \frac{1}{2(\rho - q)} \mathbb{E}\langle w \rangle_{\text{sc}}^2,$$

where we recall that $\langle - \rangle_{\text{sc}}$ is defined by (72).

Proof. The fact that $\Psi_{P_{\text{out}}}$ is differentiable on $[0, \rho]$ follows from differentiation under the expectation sign. In order to see it, we define $X = \sqrt{q}V + \sqrt{\rho - q}W^*$. Then, for all $q \in [0, \rho]$:

$$\Psi_{P_{\text{out}}}(q) = \mathbb{E} \int dX \frac{1}{\sqrt{2\pi(\rho - q)}} e^{-\frac{(X - \sqrt{q}V)^2}{2(\rho - q)}} \int dY P_{\text{out}}(Y|X) \ln \int dx \frac{1}{\sqrt{2\pi(\rho - q)}} e^{-\frac{(x - \sqrt{q}V)^2}{2(\rho - q)}} P_{\text{out}}(Y|x). \quad (202)$$

We are now in a good setting to differentiate under the expectation sign. We have for all $q \in (0, \rho)$,

$$\frac{\partial}{\partial q} \left[\frac{1}{\sqrt{\rho - q}} e^{-\frac{(X - \sqrt{q}V)^2}{2(\rho - q)}} \right] = \frac{1}{2\sqrt{\rho - q}} \left(\frac{1}{\rho - q} - \frac{(X - \sqrt{q}V)^2}{(\rho - q)^2} + \frac{V(X - \sqrt{q}V)}{\sqrt{q}(\rho - q)} \right) e^{-\frac{(X - \sqrt{q}V)^2}{2(\rho - q)}}. \quad (203)$$

Thus

$$\begin{aligned} \Psi'_{P_{\text{out}}}(q) &= \frac{1}{2} \mathbb{E} \left[\left(\frac{1}{\rho - q} - \frac{(X - \sqrt{q}V)^2}{(\rho - q)^2} + \frac{V(X - \sqrt{q}V)}{\sqrt{q}(\rho - q)} \right) \ln \int dx \frac{1}{\sqrt{2\pi(\rho - q)}} e^{-\frac{(x - \sqrt{q}V)^2}{2(\rho - q)}} P_{\text{out}}(Y|x) \right] \\ &\quad + \frac{1}{2} \mathbb{E} \left\langle \frac{1}{\rho - q} - \frac{(x - \sqrt{q}V)^2}{(\rho - q)^2} + \frac{V(x - \sqrt{q}V)}{\sqrt{q}(\rho - q)} \right\rangle_{\text{sc}} \end{aligned}$$

where the Gibbs brackets $\langle - \rangle_{\text{sc}}$ denotes the expectation with respect to $x \sim P(X|Y^{(q)}, V)$. The second term of the sum above is equal to zero. Indeed by the Nishimori identity (Proposition 12):

$$\begin{aligned} \mathbb{E} \left\langle \frac{1}{\rho - q} - \frac{(x - \sqrt{q}V)^2}{(\rho - q)^2} + \frac{V(x - \sqrt{q}V)}{\sqrt{q}(\rho - q)} \right\rangle_{\text{sc}} &= \mathbb{E} \left[\frac{1}{\rho - q} - \frac{(X - \sqrt{q}V)^2}{(\rho - q)^2} + \frac{V(X - \sqrt{q}V)}{\sqrt{q}(\rho - q)} \right] \\ &= \frac{1}{\rho - q} \mathbb{E} [1 - (W^*)^2] = 0. \end{aligned}$$

We now compute, by Gaussian integration by parts with respect to $V \sim \mathcal{N}(0, 1)$:

$$\begin{aligned} &\mathbb{E} \left[\frac{V(X - \sqrt{q}V)}{\sqrt{q}(\rho - q)} \ln \int dx \frac{1}{\sqrt{2\pi(\rho - q)}} e^{-\frac{(x - \sqrt{q}V)^2}{2(\rho - q)}} P_{\text{out}}(Y^{(q)}|x) \right] \\ &= \mathbb{E} \left[\frac{-1}{\rho - q} \ln \int dx \frac{e^{-\frac{(x - \sqrt{q}V)^2}{2(\rho - q)}}}{\sqrt{2\pi(\rho - q)}} P_{\text{out}}(Y^{(q)}|x) \right] + \mathbb{E} \left[\frac{(X - \sqrt{q}V)^2}{(\rho - q)^2} \ln \int dx \frac{e^{-\frac{(x - \sqrt{q}V)^2}{2(\rho - q)}}}{\sqrt{2\pi(\rho - q)}} P_{\text{out}}(Y^{(q)}|x) \right] \\ &\quad + \mathbb{E} \left\langle \frac{(X - \sqrt{q}V)(x - \sqrt{q}V)}{(\rho - q)^2} \right\rangle_{\text{sc}}. \end{aligned}$$

Bringing all together, we conclude:

$$\Psi'_{P_{\text{out}}}(q) = \frac{1}{2} \mathbb{E} \left\langle \frac{(X - \sqrt{q}V)(x - \sqrt{q}V)}{(\rho - q)^2} \right\rangle_{\text{sc}} = \frac{1}{2(\rho - q)} \mathbb{E} \langle w \rangle_{\text{sc}}^2.$$

This derivative is continuous at $q = 0$ thus $\Psi_{P_{\text{out}}}$ is differentiable at $q = 0$ with derivative given by the same expression. \square

Proposition 21. *Assume that the hypotheses of Proposition 19 hold and suppose also that the kernel P_{out} is informative. Then $\Psi_{P_{\text{out}}}$ is strictly increasing on $[0, \rho]$.*

Proof. Let us suppose that $\Psi_{P_{\text{out}}}$ is not strictly increasing on $[0, \rho]$. There exists thus $q \in (0, \rho)$ such that $\Psi'_{P_{\text{out}}}(q) = 0$. This means that $\langle w \rangle_{\text{sc}} = 0$ almost surely and therefore that

$$\int_{\mathbb{R}} P_{\text{out}}(Y^{(q)} | \sqrt{q}V + \sqrt{\rho - q}w) w e^{-w^2/2} dw = 0$$

almost-surely. Let us write $\sigma = \sqrt{\rho - q}$. Consequently,

$$\int_{\mathbb{R}} P_{\text{out}}(y | v + \sigma w) w e^{-w^2/2} dw = 0 \tag{204}$$

for almost all y in \mathbb{R} (if we are under assumption (i)) or all $y \in \mathbb{N}$ (under assumption (ii)) and almost all $v \in \mathbb{R}$. We will now use the following lemma:

Lemma 16. Let $Z \sim \mathcal{N}(0, 1)$ and let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a bounded function. Suppose that for almost all $v \in \mathbb{R}$,

$$\mathbb{E}[Zf(v + Z)] = 0.$$

Then, there exists a constant $C \in \mathbb{R}$ such that $f(v) = C$ for almost every v .

Proof. Let us define the function

$$h : t \mapsto \mathbb{E}[f(Z - t)] = \frac{1}{\sqrt{2\pi}} \int f(x)e^{-(x+t)^2/2} dx.$$

We have $h'(t) = \frac{-1}{\sqrt{2\pi}} \int f(x)(x+t)e^{-(x+t)^2/2} dx = -\mathbb{E}[Zf(Z - t)] = 0$ and therefore h is equal to some constant $C \in \mathbb{R}$. We are going to show that $f = C$ almost everywhere. Without loss of generality we can assume that $C = 0$, otherwise it suffices to consider the function $\tilde{f} = f - C$. Now we have for all $n \geq 0$, $t \in \mathbb{R}$

$$0 = h^{(n)}(t) = \frac{1}{\sqrt{2\pi}} \int f(x) \frac{\partial^n}{\partial t^n} e^{-(x+t)^2/2} dx = \frac{1}{\sqrt{2\pi}} \int f(x) (-1)^n H_n(x+t) e^{-(x+t)^2/2} dx,$$

where H_n is n^{th} Hermite polynomial, defined as $H_n(x) = (-1)^n e^{x^2/2} \frac{d^n}{dx^n} e^{-x^2/2}$. Therefore, for all $n \geq 0$,

$$\int f(x) H_n(x) e^{-x^2/2} dx = 0,$$

which implies that $f = 0$ almost everywhere since the Hermite functions form an orthonormal basis of $L^2(\mathbb{R})$. \square

We apply now Lemma 16 to (204) where the function f is given by $f(x) = P_{\text{out}}(y | \sigma x)$. We thus obtain that for almost every y , $P_{\text{out}}(y | \cdot)$ is almost everywhere equal to a constant. Under assumption (ii), we get that for all $y \in \mathbb{N}$, $P_{\text{out}}(y | \cdot)$ is almost everywhere equal to a constant: this contradicts the hypothesis that P_{out} is informative.

If now assumption (i) holds, then by (22) the density function $P_{\text{out}}(\cdot | x)$ is continuous on \mathbb{R} for all $x \in \mathbb{R}$. Let us fix $y \in \mathbb{R}$. We are going to show that $P_{\text{out}}(y | \cdot)$ is almost everywhere equal to a constant C_y . Given what we just showed, we can construct a sequence $(y_n)_n \in \mathbb{R}^{\mathbb{N}}$ that converges to y such that for all $n \geq 0$, there exists $E_n \subset \mathbb{R}$ with full Lebesgue's measure and $C_n \in \mathbb{R}$ such that for all $x \in E_n$,

$$P_{\text{out}}(y_n | x) = C_n.$$

Let us define $E = \bigcap_{n \geq 0} E_n$. E has therefore full Lebesgue's measure. Let now $x_1, x_2 \in E$. By continuity of $P_{\text{out}}(\cdot | x_i)$, we get

$$P_{\text{out}}(y_n | x_i) \xrightarrow[n \rightarrow \infty]{} P_{\text{out}}(y | x_i), \quad \text{for } i = 1, 2.$$

Since we know that for all $n \geq 0$ that $P_{\text{out}}(y_n | x_1) = C_n = P_{\text{out}}(y_n | x_2)$, we deduce that $P_{\text{out}}(y | x_1) = P_{\text{out}}(y | x_2)$. This proves that $P_{\text{out}}(y | \cdot)$ is almost everywhere equal to a constant C_y and contradicts the fact that P_{out} is informative. \square

We turn now our attention to the study of the function:

$$\mathcal{E}_f : \begin{cases} [0, \rho] & \rightarrow \mathbb{R}_+ \\ q & \mapsto \mathbb{E}[(f(Y^{(q)}) - \mathbb{E}[f(Y^{(q)}) | V])^2] \end{cases} \quad (205)$$

where $f : \mathbb{R} \rightarrow \mathbb{R}$ is a continuous bounded function. We will prove that \mathcal{E}_f is continuous (Proposition 22) and strictly decreasing (Proposition 23) under the following hypotheses.

- (a) For all $x \in \mathbb{R}$, $P_{\text{out}}(\cdot | x)$ is the law of $\varphi(x, \mathbf{A}) + \sqrt{\Delta} Z$ where $\varphi : \mathbb{R} \times \mathbb{R}^{k_A} \rightarrow \mathbb{R}$ is a measurable function and $(Z, \mathbf{A}) \sim \mathcal{N}(0, 1) \otimes P_A$, for some probability distribution P_A over \mathbb{R}^{k_A} .
- (b) For almost all $a \in \mathbb{R}^{k_A}$ (w.r.t. P_A), $\varphi(\cdot, a)$ is continuous almost everywhere.

We suppose also that we are in one of the following cases:

- (i) $\Delta > 0$.
- (ii) $\Delta = 0$ and φ takes values in \mathbb{N} .

Proposition 22. *Under the hypotheses presented above, \mathcal{E}_f is continuous on $[0, \rho]$.*

Proof. Consider expression (62): The first term does not depend on q and the second one is continuous by Lebesgue's convergence theorem. \square

Proposition 23. *Assume that the hypotheses of Proposition 22 hold. Suppose that $x \mapsto \int f(y)P_{\text{out}}(y | x)dy$ is not almost-everywhere equal to a constant. Then \mathcal{E}_f is strictly decreasing on $[0, \rho]$.*

Proof. $\mathcal{E}_f(q) = \mathbb{E}[f(Y^{(q)})^2] - \mathbb{E}[\mathbb{E}[f(Y^{(q)})|V]^2]$. Since the first term does not depend on q , it suffices to show that $H : q \mapsto \mathbb{E}[\mathbb{E}[f(Y^{(q)})|V]^2]$ is strictly increasing on $[0, \rho]$. We have for $q \in (0, \rho)$:

$$\mathbb{E}[f(Y^{(q)})|V] = \int \int f(y) \frac{e^{-w^2/2}}{\sqrt{2\pi}} P_{\text{out}}(y|\sqrt{q}V + \sqrt{\rho-q}w) dy dw = \int \int f(y) \frac{e^{-\frac{(x-\sqrt{q}V)^2}{2(\rho-q)}}}{\sqrt{2\pi(\rho-q)}} P_{\text{out}}(y|x) dy dx.$$

So we have, using (203):

$$\begin{aligned} \frac{\partial}{\partial q} \mathbb{E}[f(Y^{(q)})|V] &= \int \int \frac{f(y)}{2} \left(\frac{1}{\rho-q} - \frac{(x-\sqrt{q}V)^2}{(\rho-q)^2} + \frac{V(x-\sqrt{q}V)}{\sqrt{q}(\rho-q)} \right) \frac{e^{-\frac{(x-\sqrt{q}V)^2}{2(\rho-q)}}}{\sqrt{2\pi(\rho-q)}} P_{\text{out}}(y|x) dy dw \\ &= \frac{1}{2(\rho-q)} \mathbb{E} \left[f(Y^{(q)}) \left(1 - W^{*2} + \frac{\sqrt{\rho-q} V W^*}{\sqrt{q}} \right) \middle| V \right]. \end{aligned}$$

We obtain

$$H'(q) = \frac{1}{\rho-q} \mathbb{E} \left[\mathbb{E}[f(Y^{(q)})|V] \mathbb{E} \left[f(Y^{(q)}) \left(1 - W^{*2} + \frac{\sqrt{\rho-q} V W^*}{\sqrt{q}} \right) \middle| V \right] \right]. \quad (206)$$

We compute by Gaussian integration by parts:

$$\begin{aligned} \mathbb{E} \left[\mathbb{E}[f(Y^{(q)})|V] \mathbb{E} \left[f(Y^{(q)}) V W^* \middle| V \right] \right] &= \mathbb{E} \left[V \mathbb{E}[f(Y^{(q)})|V] \mathbb{E} \left[f(Y^{(q)}) W^* \middle| V \right] \right] \\ &= \mathbb{E} \left[\frac{\partial}{\partial V} \mathbb{E}[f(Y^{(q)})|V] \mathbb{E} \left[f(Y^{(q)}) W^* \middle| V \right] \right] + \mathbb{E} \left[\mathbb{E}[f(Y^{(q)})|V] \frac{\partial}{\partial V} \mathbb{E} \left[f(Y^{(q)}) W^* \middle| V \right] \right]. \end{aligned} \quad (207)$$

We compute successively

$$\begin{aligned}
\frac{\partial}{\partial V} \mathbb{E}[f(Y^{(q)})|V] &= \frac{\partial}{\partial V} \int \int f(y) \frac{e^{-\frac{(x-\sqrt{q}V)^2}{2(\rho-q)}}}{\sqrt{2\pi(\rho-q)}} P_{\text{out}}(y|x) dy dx \\
&= \int \int f(y) \frac{\sqrt{q}(x-\sqrt{q}V)}{\rho-q} \frac{e^{-\frac{(x-\sqrt{q}V)^2}{2(\rho-q)}}}{\sqrt{2\pi(\rho-q)}} P_{\text{out}}(y|x) dy dx = \frac{\sqrt{q}}{\sqrt{\rho-q}} \mathbb{E}\left[f(Y^{(q)})W^*|V\right].
\end{aligned} \tag{208}$$

$$\begin{aligned}
\frac{\partial}{\partial V} \mathbb{E}[f(Y^{(q)})W^*|V] &= \frac{\partial}{\partial V} \int \int f(y) \frac{x-\sqrt{q}V}{\sqrt{\rho-q}} \frac{e^{-\frac{(x-\sqrt{q}V)^2}{2(\rho-q)}}}{\sqrt{2\pi(\rho-q)}} P_{\text{out}}(y|x) dy dx \\
&= \int \int f(y) \left(\frac{-\sqrt{q}}{\sqrt{\rho-q}} + \frac{\sqrt{q}(x-\sqrt{q}V)^2}{(\rho-q)^{3/2}} \right) \frac{e^{-\frac{(x-\sqrt{q}V)^2}{2(\rho-q)}}}{\sqrt{2\pi(\rho-q)}} P_{\text{out}}(y|x) dy dx \\
&= \frac{\sqrt{q}}{\sqrt{\rho-q}} \mathbb{E}\left[f(Y^{(q)})(-1+W^{*2})|V\right].
\end{aligned} \tag{209}$$

By plugging (207)-(208)-(209) back in (206) we get:

$$H'(q) = \frac{1}{\rho-q} \mathbb{E}\left[\mathbb{E}[f(Y^{(q)})W^*|V]^2\right] \geq 0.$$

Let us suppose now that H is not strictly increasing on $[0, \rho]$. This means that we can find $q \in (0, \rho)$ such that $H'(q) = 0$ and therefore $\mathbb{E}[f(Y^{(q)})W^*|V] = 0$ almost-surely. This gives that for almost all $v \in \mathbb{R}$,

$$\mathbb{E}\left[W \int f(y) P_{\text{out}}(y|\sqrt{q}v + \sqrt{\rho-q}W) dy\right] = 0,$$

where \mathbb{E} is the expectation with respect to $W \sim \mathcal{N}(0, 1)$. Lemma 16 gives then that the function $x \mapsto \int f(y) P_{\text{out}}(y|x) dy$ is almost everywhere equal to a constant: we obtain a contradiction. We conclude that H is strictly increasing on $[0, \rho]$ and thus \mathcal{E}_f is strictly decreasing on $[0, \rho]$. \square

Proposition 24. *Assume that the hypotheses of Proposition 22 hold. If the channel P_{out} is informative, then there exists a continuous bounded function $f : \mathbb{R} \rightarrow \mathbb{R}$ such that $x \mapsto \int f(y) P_{\text{out}}(y|x) dy$ is not almost everywhere equal to a constant.*

Proof. Let us suppose that for all continuous bounded function $f : \mathbb{R} \rightarrow \mathbb{R}$ we have

$$\int f(y) P_{\text{out}}(y|x) dy = C_f$$

for almost all $x \in \mathbb{R}$, for some constant $C_f \in \mathbb{R}$. Let $X \sim \mathcal{N}(0, 1)$ and $Y \sim P_{\text{out}}(\cdot|X)$. We have then $\mathbb{E}[f(Y)|X] = C_f = \mathbb{E}[f(Y)]$ almost surely. Let $g : \mathbb{R} \rightarrow \mathbb{R}$ be another continuous bounded function and compute:

$$\mathbb{E}[g(X)f(Y)] = \mathbb{E}[g(X)\mathbb{E}[f(Y)|X]] = \mathbb{E}[g(X)]\mathbb{E}[f(Y)].$$

It follows that X and Y are independent: The measures $P_{\text{out}}(y|x) \frac{e^{-x^2/2}}{\sqrt{2\pi}} dy dx$ and $\mathbb{E}[P_{\text{out}}(y|X)] \frac{e^{-x^2/2}}{\sqrt{2\pi}} dy dx$ are therefore equal. Consequently, for almost every x, y we have

$$P_{\text{out}}(y|x) = \mathbb{E}[P_{\text{out}}(y|X)].$$

This gives that for almost every y , $P_{\text{out}}(y|\cdot)$ is almost everywhere equal to a constant. We conclude by

the arguments presented at the end of the proof of Proposition 21 that P_{out} is not informative, which is a contradiction. \square

Appendix C: Approximation

Let us recall the various hypotheses considered in this paper, starting with the stronger set:

- (H1) The prior distribution P_0 has a bounded support.
- (H2) φ is a bounded \mathcal{C}^2 function with bounded first and second derivatives w.r.t. its first argument.
- (H3) $(\Phi_{\mu i}) \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$.

The aim of this section is to relax them to the weaker ones:

- (h1) The prior distribution P_0 admits a finite third moment and has at least two points in its support.
- (h2) There exists $\gamma > 0$ such that the sequence $(\mathbb{E}[|\varphi(\frac{1}{\sqrt{n}}[\Phi \mathbf{X}^*]_1, \mathbf{A}_1)|^{2+\gamma}])_{n \geq 1}$ is bounded.
- (h3) The random variables $(\Phi_{\mu i})$ are independent with zero mean, unit variance and finite third moment that is bounded with n .
- (h4) For almost-all values of $\mathbf{a} \in \mathbb{R}^{k_A}$ (w.r.t. P_A), the function $x \mapsto \varphi(x, \mathbf{a})$ is continuous almost everywhere.

The hypotheses on the presence or not of the Gaussian noise in (20) are:

- (h5.a) $\Delta > 0$.
- (h5.b) $\Delta = 0$ and φ takes values in \mathbb{N} .

In this section, we suppose that Theorem 1 holds for channels of the form (20) (with $\Delta > 0$) under the hypotheses (H1), (H2) and (H3), as proven in Section 4.

We show in this section that this imply that Theorem 1 holds under the weaker hypotheses (h1)-(h2)-(h3)-(h4), and either (h5.a) or (h5.b). This section is organized as follows: We first prove Theorem 1 under (h1)-(h2)-(h3)-(h4) and (h5.a) (i.e. $\Delta > 0$). This is done by first relaxing the hypotheses on P_0 and Φ (Sec. C.1) and then the hypotheses on φ (Sec. C.2). Finally, in Sec. C.3, we let $\Delta \rightarrow 0$ in order to prove Theorem 1 under (h1)-(h2)-(h3)-(h4) and (h5.b).

Note that the statement of Theorem 1 is equivalent to the statement of Corollary 2, which simply express the result in terms of mutual information. This formulation will be slightly more convenient to relax the hypotheses. We will therefore prove in this section that (46) holds under the hypotheses (h1)-(h2)-(h3)-(h4), and either (h5.a) or (h5.b). The statement of Theorem 1 can then be directly obtained by using the expressions of I_{P_0} , $\mathcal{I}_{P_{\text{out}}}$ in terms of ψ_{P_0} , $\Psi_{P_{\text{out}}}$, the relation (49) and Lemma 23.

C.1 Relaxing the hypotheses on P_0 and Φ

As explained at the beginning of Sec. 4, it suffices to consider the case $\Delta = 1$. We start by relaxing the hypothesis (H1).

Lemma 17 (Relaxing P_0). *Suppose that (h1)-(H2)-(H3) and (h5.a) hold. Then Theorem 1 holds.*

Proof. The ideas are basically the same that in [99] (Sec. 6.2.2). We omit the details here for the sake of brevity. \square

We now relax the Gaussian assumption on the ‘‘measurement matrix’’ Φ .

Lemma 18 (Relaxing Φ). *Suppose that $\varphi : \mathbb{R} \times \mathbb{R}^{k_A} \rightarrow \mathbb{R}$ is \mathcal{C}^∞ with compact support and that (h1)-(h3)-(h5.a) hold. Then Theorem 1 holds.*

Proof. The proof is based on the Lindeberg generalization theorem (Theorem 2 from [95]) which is a variant of the generalized ‘‘Lindeberg principle’’ from [100]:

Theorem 5 (Lindeberg generalization theorem). *Let $(U_i)_{1 \leq i \leq n}$ and $(V_i)_{1 \leq i \leq n}$ be two collections of random variables with independent components and $f : \mathbb{R}^n \rightarrow \mathbb{R}$ a \mathcal{C}^3 function. Denote $a_i = |\mathbb{E}U_i - \mathbb{E}V_i|$ and $b_i = |\mathbb{E}[U_i^2] - \mathbb{E}[V_i^2]|$. Then*

$$|\mathbb{E}f(\mathbf{U}) - \mathbb{E}f(\mathbf{V})| \leq \sum_{i=1}^n \left\{ a_i \mathbb{E} |\partial_i f(U_{1:i-1}, 0, V_{i+1:n})| + \frac{b_i}{2} \mathbb{E} |\partial_i^2 f(U_{1:i-1}, 0, V_{i+1:n})| \right. \\ \left. + \frac{1}{2} \mathbb{E} \int_0^{U_i} |\partial_i^3 f(U_{1:i-1}, 0, V_{i+1:n})| (U_i - s)^2 ds + \frac{1}{2} \mathbb{E} \int_0^{V_i} |\partial_i^3 f(U_{1:i-1}, 0, V_{i+1:n})| (V_i - s)^2 ds \right\}.$$

Let $(\Phi'_{\mu,i}) \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$ and let $(\Phi_{\mu,i})$ be a family of independent random variables, with zero mean and unit variance. Let f'_n be the free entropy (29) with design matrix Φ' and f_n be the free entropy (29) with design matrix Φ .

We will apply Theorem 5 to the function

$$F : \mathbf{U} \in \mathbb{R}^{m \times n} \mapsto \frac{1}{n} \mathbb{E} \ln \int_{\mathbf{x}, \mathbf{a}} dP_A(\mathbf{a}) dP_0(\mathbf{x}) e^{-\frac{1}{2} \sum_{\mu=1}^m \left(\varphi \left(\frac{1}{\sqrt{n}} [\mathbf{U}\mathbf{X}^*]_{\mu}, \mathbf{A}_{\mu} \right) - \varphi \left(\frac{1}{\sqrt{n}} [\mathbf{U}\mathbf{x}]_{\mu}, \mathbf{a}_{\mu} \right) + Z_{\mu} \right)^2}$$

where the expectation \mathbb{E} is taken w.r.t. \mathbf{X}^* , \mathbf{A} and \mathbf{Z} . We have

$$f_n = \mathbb{E}F(\Phi) \quad \text{and} \quad f'_n = \mathbb{E}F(\Phi').$$

It is not difficult to verify that F is a \mathcal{C}^3 function and that for all $1 \leq \mu \leq m$ and $1 \leq i \leq n$:

$$\left\| \frac{\partial^3 F}{\partial U_{\mu,i}^3} \right\|_{\infty} \leq \frac{C}{n^{5/2}},$$

for some constant C that only depends on φ and the first three moments of P_0 . Thus, an application of Theorem 5 gives $|f_n - f'_n| \leq \frac{C}{\sqrt{n}}$. By Proposition 17, we know that Theorem 1 holds for f'_n , thus it holds for f_n . \square

C.2 Relaxing the hypotheses on φ

It remains to relax the hypotheses on φ . This section is dedicated to the proof of the following proposition, which is of course exactly the statement of Theorem 1.

Proposition 25 (Relaxing φ). *Suppose that (h1)-(h2)-(h3)-(h4) and (h5.a) hold. Then, Theorem 1 holds for the output channel (20).*

To prove Proposition 25 we will approximate the function φ with a function $\widehat{\varphi}$ which is \mathcal{C}^{∞} with compact support. In the following, G is a standard Gaussian random variable, independent of everything else.

Proposition 26. *Suppose that (h1)-(h2)-(h3)-(h4) hold. Then, for all $\epsilon > 0$, there exist $\widehat{\varphi} \in \mathcal{C}^{\infty}(\mathbb{R} \times \mathbb{R}^{k_A})$ with compact support, such that*

$$\mathbb{E} [(\varphi(\sqrt{\rho}G, \mathbf{A}) - \widehat{\varphi}(\sqrt{\rho}G, \mathbf{A}))^2] \leq \epsilon,$$

and for n large enough, we have

$$\mathbb{E} \left[\left(\varphi \left(\frac{1}{\sqrt{n}} [\Phi \mathbf{X}^*]_1, \mathbf{A}_1 \right) - \widehat{\varphi} \left(\frac{1}{\sqrt{n}} [\Phi \mathbf{X}^*]_1, \mathbf{A}_1 \right) \right)^2 \right] \leq \epsilon.$$

Proof. By the Central Limit Theorem (using the fact that the third moments of $(X_i^* \Phi_{1,i})$ are bounded with n , because of hypotheses (h1) and (h3))

$$\left(\frac{[\Phi \mathbf{X}^*]_1}{\sqrt{n}}, \mathbf{A}_1 \right) \xrightarrow[n \rightarrow \infty]{(d)} (\sqrt{\rho} G, \mathbf{A}_1). \quad (210)$$

This implies that

$$\varphi \left(\frac{[\Phi \mathbf{X}^*]_1}{\sqrt{n}}, \mathbf{A}_1 \right) \xrightarrow[n \rightarrow \infty]{(d)} \varphi(\sqrt{\rho} G, \mathbf{A}_1), \quad (211)$$

because $\varphi(\cdot, \mathbf{A}_1)$ is almost-surely continuous almost-everywhere, by assumption (h4). The following sequence $(\varphi(\frac{[\Phi \mathbf{X}^*]_1}{\sqrt{n}}, \mathbf{A}_1))_n$ is by assumption (h2) bounded in L^2 , thus by (211) we have that $\mathbb{E}[\varphi(\sqrt{\rho} G, \mathbf{A}_1)^2] < \infty$. Let $\epsilon > 0$. We have just proved that $\varphi \in L^2(\mathbb{R} \times \mathbb{R}^{k_A})$ with the measure induced by $(\sqrt{\rho} G, \mathbf{A}_1)$. There exists (see for instance Corollary 4.2.2 in [98]) a C^∞ function with compact support $\hat{\varphi}$ such that $\mathbb{E}[(\varphi(\sqrt{\rho} G, \mathbf{A}) - \hat{\varphi}(\sqrt{\rho} G, \mathbf{A}))^2] \leq \epsilon$.

One deduce from (210) and (211) that

$$\left(\varphi \left(\frac{1}{\sqrt{n}} [\Phi \mathbf{X}^*]_1, \mathbf{A}_1 \right) - \hat{\varphi} \left(\frac{1}{\sqrt{n}} [\Phi \mathbf{X}^*]_1, \mathbf{A}_1 \right) \right)^2 \xrightarrow[n \rightarrow \infty]{(d)} (\varphi(\sqrt{\rho} G, \mathbf{A}) - \hat{\varphi}(\sqrt{\rho} G, \mathbf{A}))^2.$$

Now, hypothesis (h2) gives that the sequence above is uniformly integrable. This gives that

$$\mathbb{E} \left(\varphi \left(\frac{1}{\sqrt{n}} [\Phi \mathbf{X}^*]_1, \mathbf{A}_1 \right) - \hat{\varphi} \left(\frac{1}{\sqrt{n}} [\Phi \mathbf{X}^*]_1, \mathbf{A}_1 \right) \right)^2 \xrightarrow[n \rightarrow \infty]{} \mathbb{E}(\varphi(\sqrt{\rho} G, \mathbf{A}) - \hat{\varphi}(\sqrt{\rho} G, \mathbf{A}))^2 \leq \epsilon.$$

Consequently, the left-hand side is smaller than 2ϵ for n large enough. This concludes the proof. \square

In the remaining of this section, we prove Proposition 25. Let $\epsilon > 0$. Let φ and $\hat{\varphi}$ as in Proposition 26. Let us define $\mathbf{Y} = \varphi(n^{-1/2} \Phi \mathbf{X}^*, \mathbf{A}) + \sqrt{\Delta} Z$ and $\hat{\mathbf{Y}} = \hat{\varphi}(n^{-1/2} \Phi \mathbf{X}^*, \mathbf{A}) + \sqrt{\Delta} Z$.

Lemma 19. *Suppose that (h1)-(h2)-(h3)-(h4) and (h5.a) hold. There exists a constant $C > 0$ such that for n large enough*

$$\left| \frac{1}{n} I(\mathbf{X}^*; \mathbf{Y} | \Phi) - \frac{1}{n} I(\mathbf{X}^*; \hat{\mathbf{Y}} | \Phi) \right| \leq C \sqrt{\epsilon}.$$

Proof. We have, for n large enough

$$\mathbb{E} \|\mathbf{Y} - \hat{\mathbf{Y}}\|^2 = m \mathbb{E} \left[\left(\varphi \left(\frac{1}{\sqrt{n}} [\Phi \mathbf{X}^*]_1, \mathbf{A}_1 \right) - \hat{\varphi} \left(\frac{1}{\sqrt{n}} [\Phi \mathbf{X}^*]_1, \mathbf{A}_1 \right) \right)^2 \right] \leq m \epsilon.$$

By Proposition 14, we obtain that there exists a constant $C > 0$ (that depends only on Δ and φ) such that

$$|I(\mathbf{X}^*; \mathbf{Y} | \Phi) - I(\mathbf{X}^*; \hat{\mathbf{Y}} | \Phi)| \leq C m \sqrt{\epsilon},$$

which gives the result. \square

Let P_{out} denote the transition kernel associated to φ and \hat{P}_{out} the one associated to $\hat{\varphi}$. Analogously to the previous Lemma, one can show:

Lemma 20. *There exists a constant $C' > 0$ such that for all $q \in [0, \rho]$, $|\mathcal{I}_{P_{\text{out}}}(q) - \mathcal{I}_{\hat{P}_{\text{out}}}(q)| \leq C' \sqrt{\epsilon}$.*

From there we obtain that

$$\left| \inf_{q \in [0, \rho]} \sup_{r \geq 0} i_{\text{RS}}(q, r) - \inf_{q \in [0, \rho]} \sup_{r \geq 0} \hat{i}_{\text{RS}}(q, r) \right| \leq C' \sqrt{\epsilon}. \quad (212)$$

Applying Theorem 1 for \widehat{P}_{out} , we obtain that for n large enough $|\frac{1}{n}I(\mathbf{X}^*; \widehat{\mathbf{Y}}|\Phi) - \inf_{q \in [0, \rho]} \sup_{r \geq 0} \widehat{i}_{\text{RS}}(q, r)| \leq \sqrt{\epsilon}$. We now combine this with (212) and Lemma 19 we obtain that for n large enough

$$\left| \frac{1}{n}I(\mathbf{X}^*; \mathbf{Y}|\Phi) - \inf_{q \in [0, \rho]} \sup_{r \geq 0} i_{\text{RS}}(q, r) \right| \leq \left| \frac{1}{n}I(\mathbf{X}^*; \widehat{\mathbf{Y}}|\Phi) - \inf_{q \in [0, \rho]} \sup_{r \geq 0} \widehat{i}_{\text{RS}}(q, r) \right| + (C+C')\sqrt{\epsilon} \leq (C+C'+1)\sqrt{\epsilon},$$

which concludes the proof of Proposition 25, because of (49) and the definition of the functions I_{P_0} and $\mathcal{I}_{P_{\text{out}}}$ in Corollary 2.

C.3 The case of discrete channels: Removing the Gaussian noise

Now that we proved (Proposition 25) that Theorem 1 holds under hypotheses (h1)-(h2)-(h3)-(h4) and (h5.a), we are going to show that it holds under (h1)-(h2)-(h3)-(h4) and (h5.b) by letting $\Delta \rightarrow 0$. We suppose in this section that φ takes values in \mathbb{N} and write $\mathbf{Y} = \varphi(\Phi \mathbf{X}^* / \sqrt{n}, \mathbf{A})$. By Proposition 25 we know that for all $\Delta > 0$,

$$\frac{1}{n}I(\mathbf{X}^*; \mathbf{Y} + \sqrt{\Delta} \mathbf{Z}|\Phi) \xrightarrow{n \rightarrow \infty} \inf_{q \in [0, \rho]} \sup_{r \geq 0} \left\{ I_{P_0}(r) + \alpha I(W^*; \varphi(\sqrt{q}V + \sqrt{\rho - q}W^*, A) + \sqrt{\Delta}Z|V) - \frac{r}{2}(\rho - q) \right\},$$

where $\mathbf{Z} \sim \mathcal{N}(0, \mathbf{I}_m)$ and $(V, W^*, Z, A) \sim \mathcal{N}(0, 1)^{\otimes 3} \otimes P_A$. Since \mathbf{Y} takes values in \mathbb{N}^m and φ takes values in \mathbb{N} , we can apply Corollary 7 twice to obtain that for all $\Delta \in (0, 1]$,

$$\left| I(\mathbf{X}^*; \mathbf{Y} + \sqrt{\Delta} \mathbf{Z}|\Phi) - I(\mathbf{X}^*; \mathbf{Y}|\Phi) \right| \leq 100m e^{-1/(16\Delta)}$$

and (recall that by definition $\mathcal{I}_{P_{\text{out}}}(q) = I(W^*; \varphi(\sqrt{q}V + \sqrt{\rho - q}W^*, A)|V)$):

$$\left| I(W^*; \varphi(\sqrt{q}V + \sqrt{\rho - q}W^*, A) + \sqrt{\Delta}Z|V) - \mathcal{I}_{P_{\text{out}}}(q) \right| \leq 100e^{-1/(16\Delta)}.$$

Since our control over Δ is uniform in n , we can permute the $n \rightarrow \infty$ limit with the $\Delta \rightarrow 0$ limit to get:

$$\frac{1}{n}I(\mathbf{X}^*; \mathbf{Y}|\Phi) \xrightarrow{n \rightarrow \infty} \inf_{q \in [0, \rho]} \sup_{r \geq 0} \left\{ I_{P_0}(r) + \alpha \mathcal{I}_{P_{\text{out}}}(q) - \frac{r}{2}(\rho - q) \right\}.$$

Appendix D: Some sup-inf formulas

This appendix gathers some useful lemmas for the manipulations of “sup-inf” formulas like (43).

Lemma 21. *Let $f, g : \mathbb{R}_+ \rightarrow \mathbb{R}$ be two non-decreasing convex functions. We have*

$$\sup_{x \geq 0} \inf_{y \geq 0} \{f(x) + g(y) - xy\} = \sup_{y \geq 0} \inf_{x \geq 0} \{f(x) + g(y) - xy\}.$$

Proof. Let us define the monotone conjugate (see the end of §12 of [101]) of f and g :

$$f^*(y) = \sup_{x \geq 0} \{xy - f(x)\} \quad \text{and} \quad g^*(y) = \sup_{x \geq 0} \{xy - g(x)\}.$$

These conjugates satisfy an analog of the Fenchel-Moreau Theorem: $f(x) = \sup_{y \geq 0} \{xy - f^*(y)\}$ and $g(y) =$

$\sup_{x \geq 0} \{xy - g^*(x)\}$, see Theorem 12.4 from [101]. We have then

$$\begin{aligned} \sup_{x \geq 0} \inf_{y \geq 0} \{f(x) + g(y) - xy\} &= \sup_{x \geq 0} \{f(x) - g^*(x)\} = \sup_{x \geq 0} \sup_{y \geq 0} \{xy - f^*(y) - g^*(x)\} \\ &= \sup_{y \geq 0} \left\{ -f^*(y) + \sup_{x \geq 0} \{xy - g^*(x)\} \right\} = \sup_{y \geq 0} \{ -f^*(y) + g(y) \} \\ &= \sup_{y \geq 0} \inf_{x \geq 0} \{f(x) + g(y) - xy\}. \end{aligned}$$

□

The next Lemma on the Legendre transform will be useful.

Lemma 22. *Let $V \subset \mathbb{R}$ be a non-empty, closed interval and let $g : V \rightarrow \mathbb{R}$ be a continuous convex function. Define*

$$g^* : x \in \mathbb{R} \mapsto \sup_{y \in V} \{xy - g(y)\} \in \mathbb{R} \cup \{+\infty\}. \quad (213)$$

Let $\text{dom } g^ = \{x \in \mathbb{R} \mid g^*(x) < \infty\}$. Then g^* is a closed convex function and $\text{dom } g^*$ is a non-empty interval. Moreover, for all $x \in \text{dom } g^*$,*

$$\partial g^*(x) = \arg \max_{y \in V} \{xy - g(y)\}. \quad (214)$$

In particular, if g is strictly convex then g^ is differentiable around every point in the interior of $\text{dom } g^*$.*

Proof. We first extend the function g on \mathbb{R} by setting $g(x) = +\infty$ for all $x \notin V$. Notice that this does not change the definition of the function g^* . g is then a proper, closed convex function (see for instance [101] for the definitions of these properties). By Theorem 12.2 in [101], g^* is also a proper closed convex function on \mathbb{R} , which gives that $\text{dom } g^*$ is a non-empty interval. We now apply Corollary 23.5.1 from [101] to obtain

$$y \in \partial g^*(x) \iff x \in \partial g(y) \iff y \text{ maximizes } \{xy - g(y)\},$$

for all $x \in \text{dom } g^*$, which concludes the proof. □

Corollary 8. *Let $f : \mathbb{R}_+ \rightarrow \mathbb{R}$ be a convex, Lipschitz, non-decreasing function. Define $\rho = \sup_{x \geq 0} f'(x^+)$. Let $g : [0, \rho] \rightarrow \mathbb{R}$ be a convex, Lipschitz, non-decreasing function. For $q_1 \in \mathbb{R}_+$ and $q_2 \in [0, \rho]$ we define $\psi(q_1, q_2) = f(q_1) + g(q_2) - q_1 q_2$. Then*

$$\sup_{q_1 \geq 0} \inf_{q_2 \in [0, \rho]} \psi(q_1, q_2) = \sup_{q_2 \in [0, \rho]} \inf_{q_1 \geq 0} \psi(q_1, q_2).$$

Proof. In order to apply Lemma 21 we need to extend g on \mathbb{R}_+ . We thus define for $x \geq 0$

$$g(x) = \begin{cases} g(x) & \text{if } x \leq \rho, \\ g(\rho) + (x - \rho)g'(\rho^-) & \text{if } x \geq \rho. \end{cases}$$

Obviously g is a convex, Lipschitz, non-decreasing function on \mathbb{R}_+ . One can thus apply Lemma 21:

$$\sup_{q_1 \geq 0} \inf_{q_2 \geq 0} \{f(q_1) + g(q_2) - q_1 q_2\} = \sup_{q_2 \geq 0} \inf_{q_1 \geq 0} \{f(q_1) + g(q_2) - q_1 q_2\}. \quad (215)$$

We will show now that $\sup_{q_1 \geq 0} \inf_{q_2 \geq 0} \psi(q_1, q_2) = \sup_{q_1 \geq 0} \inf_{q_2 \in [0, \rho]} \psi(q_1, q_2)$. Let us define for $q_1 \geq 0$

$$g^*(q_1) = \sup_{q_2 \in [0, \rho]} \{q_1 q_2 - g(q_2)\} \quad \text{and} \quad h(q_1) = \inf_{q_2 \in [0, \rho]} \{f(q_1) + g(q_2) - q_1 q_2\} = f(q_1) - g^*(q_1).$$

For $q_1 \geq g'(\rho^-)$ we have $g^*(q_1) = q_1\rho - g(\rho)$. The function h is therefore non-increasing on $[g'(\rho^-), +\infty)$, because f is ρ -Lipschitz. We get that $h = f - g^*$ achieves its supremum on $[0, g'(\rho^-)]$. Let q_1^* be the smallest point at which this supremum is achieved. Let us show that $\inf_{q_2 \in [0, \rho]} \{g(q_2) - q_1^*q_2\} = \inf_{q_2 \geq 0} \{g(q_2) - q_1^*q_2\}$.

- If $q_1^* = 0$, then the minimum over $[0, \rho]$ is achieved at $q_2 = 0$, because g is non-decreasing. By convexity, q_2 is also the minimizer over \mathbb{R}_+ : both infimum are equal.
- If $q_1^* > 0$, the optimality condition of q_1^* gives $f'(q_1^{*-}) - (g^*)'(q_1^{*-}) \geq 0$. By (214) we obtain that there exists $q_2^* \in \arg \min_{q_2 \in [0, \rho]} \{g(q_2) - q_1^*q_2\}$ such that $f'(q_1^{*-}) \geq q_2^*$. If $q_2^* < \rho$ we conclude, as above, that both infimum are equal. Suppose now that $q_2^* = \rho$ and define $q_1' = g'(\rho^-)$. By the optimality condition of $q_2^* = \rho$ we have $q_1' = g'(\rho^-) \leq q_1^*$. Compute

$$h(q_1^*) - h(q_1') = f(q_1^*) - f(q_1') - \rho(q_1^* - q_1') \leq 0$$

because f is ρ -Lipschitz. Since $q_1' \leq q_1^*$ and q_1^* is defined as the smallest maximizer of h , we get that $q_1^* = q_1'$. The left-hand derivative of $q_2 \mapsto g(q_2) - q_1^*q_2$ at $q = \rho$ is therefore equal to 0: ρ minimizes $q_2 \mapsto g(q_2) - q_1^*q_2$ over \mathbb{R}_+ : both infimum are equal.

We have proved that $\inf_{q_2 \in [0, \rho]} \{g(q_2) - q_1^*q_2\} = \inf_{q_2 \geq 0} \{g(q_2) - q_1^*q_2\}$. Therefore

$$\sup_{q_1 \geq 0} \inf_{q_2 \in [0, \rho]} \psi(q_1, q_2) = \inf_{q_2 \in [0, \rho]} \psi(q_1^*, q_2) = \inf_{q_2 \geq 0} \psi(q_1^*, q_2) \leq \sup_{q_1 \geq 0} \inf_{q_2 \geq 0} \psi(q_1, q_2).$$

We conclude that $\sup_{q_1 \geq 0} \inf_{q_2 \geq 0} \psi(q_1, q_2) = \sup_{q_1 \geq 0} \inf_{q_2 \in [0, \rho]} \psi(q_1, q_2)$ because the converse inequality is trivial. It remains to show now that $\sup_{q_2 \geq 0} \inf_{q_1 \geq 0} \psi(q_1, q_2) = \sup_{q_2 \in [0, \rho]} \inf_{q_1 \geq 0} \psi(q_1, q_2)$ to prove the Lemma, because of (215). The inequality “ \geq ” is obvious and the inequality “ \leq ” follows from the fact that $\inf_{q_1 \geq 0} \psi(q_1, q_2) = -\infty$ if $q_2 > \rho$. \square

Lemma 23. *Let g be a strictly convex, differentiable, Lipschitz non-decreasing function on \mathbb{R}_+ . Define $\rho = \sup_{x \geq 0} g'(x)$. Let f be a convex, continuous, strictly increasing function on $[0, \rho]$, differentiable on $[0, \rho)$. For $(q_1, q_2) \in [0, \rho] \times \mathbb{R}_+$ we define $\psi(q_1, q_2) = f(q_1) + g(q_2) - q_1q_2$. Then*

$$\sup_{q_1 \in [0, \rho]} \inf_{q_2 \geq 0} \psi(q_1, q_2) = \sup_{(q_1, q_2) \in \Gamma} \psi(q_1, q_2), \quad (216)$$

where

$$\Gamma = \left\{ (q_1, q_2) \in [0, \rho] \times (\mathbb{R}_+ \cup \{+\infty\}) \mid \begin{array}{l} q_1 = g'(q_2) \\ q_2 = f'(q_1) \end{array} \right\},$$

where all the function are extended by their limits at the points at which they may not be defined (for instance $g'(+\infty) = \lim_{q \rightarrow \infty} g'(q)$, $f'(\rho) = \lim_{q \rightarrow \rho} f'(q)$). Moreover, the above extremas are achieved precisely on the same couples.

Proof. Let q_1^* be a maximizer of $f - g^*$ over $[0, \rho]$. q_1^* is well defined because f is continuous and g^* is continuous over $[0, \rho)$ and is either continuous at ρ or goes to $+\infty$ at ρ (this comes from the fact that g^* is a closed convex function, see Lemma 22).

Case 1: $0 < q_1^* < \rho$. By strict convexity of g , $\psi(q_1, \cdot)$ admits a unique minimizer q_2^* and $(g^*)'(q_1^*) = q_2^*$ by Lemma 22. Thus, the optimality condition at q_1^* gives

$$0 = f'(q_1^*) - (g^*)'(q_1^*) = f'(q_1^*) - q_2^*.$$

The optimality of q_2^* gives then $q_1^* \leq g'(q_2^*)$. Suppose that $q_1^* < g'(q_2^*)$. This is only possible when $q_2^* = 0$. Define $q_1' = g'(q_2^*) = g'(0)$. Remark that $g^*(q_1') = -g(0) = g^*(q_1^*)$. We supposed that $q_1' > q_1^*$ thus, by

strict monotonicity of f , $f(q'_1) - g^*(q'_1) > f(q_1^*) - g^*(q_1^*)$ which contradict the optimality of q_1^* . We obtain therefore that $q_1^* = g'(q_2^*)$.

Case 2: $q_1^* = 0$. The optimality condition gives now

$$0 \leq f'(q_1^* = 0) \leq q_2^*, \quad (217)$$

where q_2^* is again the unique minimizer of $\psi(q_1^* = 0, \cdot) = f(0) + g$. g is strictly increasing, so $q_2^* = 0$. Therefore $q_2^* = 0 = f'(q_1^* = 0)$, by (217). As before we have necessarily, by optimality of q_2^* that $q_1^* = g'(q_2^*)$.

Case 3: $q_1^* = \rho$. In that case $\arg \min_{q_2 \geq 0} \{g(q_2) - q_1^* q_2\} = \emptyset$ because g is strictly convex and ρ -Lipschitz. Lemma 22 gives then that $\partial g^*(\rho) = \emptyset$ which implies (see Theorem 23.3 from [101]) that $(g^*)'(\rho^-) = +\infty$. Since $q_1^* = \rho$ maximizes $f - g^*$, we necessarily have then $f'(\rho^-) = +\infty$.

Using the slight abuse of notation explained in the Lemma, we have $f'(q_1^*) = +\infty = q_2^*$, where $q_2^* = +\infty$ is the unique “minimizer” of $\psi(q_1^*, \cdot)$, by strict convexity of g . By definition of ρ we have also $g'(q_2^*) = g'(+\infty) = \rho = q_1^*$.

We conclude from the three cases above that the “sup-inf” in (216) is achieved, and that all the couples (q_1^*, q_2^*) that achieve this “sup-inf” belong to Γ . Thus

$$\sup_{q_1 \in [0, \rho]} \inf_{q_2 \geq 0} \psi(q_1, q_2) \leq \sup_{(q_1, q_2) \in \Gamma} \psi(q_1, q_2).$$

Let now be $(q_1, q_2) \in \Gamma$. By convexity of g we see easily that $\psi(q_1, q_2) = \inf_{q'_2} \psi(q_1, q'_2)$. Thus, $\psi(q_1, q_2) \leq \sup_{q'_1} \inf_{q'_2} \psi(q'_1, q'_2)$. Therefore

$$\sup_{(q_1, q_2) \in \Gamma} \psi(q_1, q_2) \leq \sup_{q_1 \in [0, \rho]} \inf_{q_2 \geq 0} \psi(q_1, q_2).$$

This concludes the proof of (216). It remains to see that a couple $(q_1^*, q_2^*) \in \Gamma$ that achieves the supremum in (216) also achieves the “sup-inf”. This simply follows from the fact that $\psi(q_1^*, q_2^*) = \inf_{q_2} \psi(q_1^*, q_2)$ and (216). \square

Appendix E: Concentration of free entropy and overlaps

E.1 Concentration of the free entropy

The goal of this appendix is to prove that the free entropy of the interpolating model studied in Sec. 4.1 concentrates around its expectation. To simplify the notations we use $C(\varphi, S, \alpha)$ for a generic non-negative constant depending *only* on φ , S and α (S is the supremum over the signal values). We will also use the notation $K = 1 + \max(\rho, r_{\max})$ for a constant (depending only on $\rho \leq S^2$ and φ) that upper bounds both R_1 and R_2 given by (89). It is also understood that n and m are large enough and $m/n \rightarrow \alpha$.

Theorem 6 (Free entropy concentration). *Under assumptions (H1), (H2) and (H3) there exists a non-negative constant $C(\varphi, S, \alpha)$ such that the partition function (97) concentrates as*

$$\text{Var} \left(\frac{1}{n} \ln \mathcal{Z}_{t, \epsilon} \right) = \mathbb{E} \left[\left(\frac{1}{n} \ln \mathcal{Z}_{t, \epsilon} - \frac{1}{n} \mathbb{E} \ln \mathcal{Z}_{t, \epsilon} \right)^2 \right] \leq \frac{C(\varphi, S, \alpha)}{n}. \quad (218)$$

The remaining of this appendix is dedicated to the proof of Theorem 6. We first recall some set-up and

notation for the convenience of the reader. Recall that the interpolating Hamiltonian (95)–(93) is

$$\mathcal{H}_{t,\epsilon}(\mathbf{x}, \mathbf{w}; \mathbf{Y}_t, \mathbf{Y}'_t, \Phi, \mathbf{V}) = - \sum_{\mu=1}^m \ln P_{\text{out}}(Y_{t,\mu} | s_{t,\mu}(\mathbf{x}, w_\mu)) + \frac{1}{2} \sum_{i=1}^n (Y'_{t,i} - \sqrt{R_1(t)} x_i)^2 \quad (219)$$

where

$$s_{t,\mu}(\mathbf{x}, w_\mu) := \sqrt{\frac{1-t}{n}} [\Phi \mathbf{x}]_\mu + k_1(t) V_\mu + k_2(t) w_\mu, \quad k_1(t) := \sqrt{R_2(t)}, \quad k_2(t) := \sqrt{\rho t - R_2(t) + 2s_n}.$$

We find it convenient to use the random function representation (20) for the interpolating model, namely

$$\begin{cases} Y_{t,\mu} = \varphi \left(\sqrt{\frac{1-t}{n}} [\Phi \mathbf{X}^*]_\mu + k_1(t) V_\mu + k_2(t) W_{\mu}^*, \mathbf{A}_\mu \right) + Z_\mu, \\ Y'_{t,i} = \sqrt{R_1(t)} X_i^* + Z'_i. \end{cases}$$

In this representation the random variables $(\mathbf{A}_\mu)_{1 \leq \mu \leq m} \stackrel{\text{iid}}{\sim} P_A$ are arbitrary, and $(Z_\mu)_{1 \leq \mu \leq m} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$, $(Z'_i)_{1 \leq i \leq n} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$. We have

$$\begin{aligned} P_{\text{out}}(Y_{t,\mu} | s_{t,\mu}(\mathbf{x}, w_\mu)) &= \int dP_A(\mathbf{a}_\mu) \frac{1}{\sqrt{2\pi}} \exp \left\{ -\frac{1}{2} \left(Y_{t,\mu} - \varphi(s_{t,\mu}(\mathbf{x}, w_\mu), \mathbf{a}_\mu) \right)^2 \right\} \\ &= \int dP_A(\mathbf{a}_\mu) \frac{1}{\sqrt{2\pi}} \exp \left\{ -\frac{1}{2} \left(\Gamma_{t,\mu}(\mathbf{x}, w_\mu, \mathbf{a}_\mu) + Z_\mu \right)^2 \right\} \end{aligned} \quad (220)$$

where, using the random function representation,

$$\begin{aligned} \Gamma_{t,\mu}(\mathbf{x}, w_\mu, \mathbf{a}_\mu) & \quad (221) \\ &= \varphi \left(\sqrt{\frac{1-t}{n}} [\Phi \mathbf{X}^*]_\mu + k_1(t) V_\mu + k_2(t) W_{\mu}^*, \mathbf{A}_\mu \right) - \varphi \left(\sqrt{\frac{1-t}{n}} [\Phi \mathbf{x}]_\mu + k_1(t) V_\mu + k_2(t) w_\mu, \mathbf{a}_\mu \right). \end{aligned}$$

From (219), (220), (221) we can express the free entropy of the interpolating model as

$$\frac{1}{n} \ln \mathcal{Z}_{t,\epsilon} = \frac{1}{n} \ln \int dP_0(\mathbf{x}) dP_A(\mathbf{a}) \mathcal{D}\mathbf{w} e^{-\mathcal{H}_{t,\epsilon}(\mathbf{x}, \mathbf{w}, \mathbf{a})} - \frac{m}{2n} \ln(2\pi) \quad (222)$$

where $\mathcal{D}\mathbf{w}$ denote the standard m -dimensional Gaussian measure and where the Hamiltonian $\mathcal{H}_{t,\epsilon}$ is re-expressed as

$$\mathcal{H}_{t,\epsilon}(\mathbf{x}, \mathbf{w}, \mathbf{a}) = \frac{1}{2} \sum_{\mu=1}^m (\Gamma_{t,\mu}(\mathbf{x}, w_\mu, \mathbf{a}_\mu) + Z_\mu)^2 + \frac{1}{2} \sum_{i=1}^n (\sqrt{R_1(t)} (X_i^* - x_i) + Z'_i)^2. \quad (223)$$

The interpretation here is that $\mathbf{x}, \mathbf{w}, \mathbf{a}$ are *annealed* variables and $\Phi, \mathbf{V}, \mathbf{A}, \mathbf{Y}_t, \mathbf{Y}'_t, \mathbf{X}^*, \mathbf{W}^*$, or equivalently $\Phi, \mathbf{V}, \mathbf{A}, \mathbf{Z}, \mathbf{Z}', \mathbf{X}^*, \mathbf{W}^*$ are *quenched*. The inference problem is to recover $\mathbf{X}^*, \mathbf{W}^*$ given $\Phi, \mathbf{V}, \mathbf{Y}_t, \mathbf{Y}'_t$. The free entropy can be further re-expressed as

$$\frac{1}{n} \ln \mathcal{Z}_{t,\epsilon} = \frac{1}{n} \ln \hat{\mathcal{Z}}_{t,\epsilon} - \frac{1}{2n} \sum_{\mu=1}^m Z_\mu^2 - \frac{1}{2n} \sum_{i=1}^n Z_i'^2 - \frac{m}{2n} \ln(2\pi) \quad (224)$$

where

$$\frac{1}{n} \ln \hat{Z}_{t,\epsilon} = \frac{1}{n} \ln \int dP_0(\mathbf{x}) dP_A(\mathbf{a}) \mathcal{D}\mathbf{w} e^{-\hat{\mathcal{H}}_{t,\epsilon}(\mathbf{x}, \mathbf{w}, \mathbf{a})}, \quad (225)$$

$$\begin{aligned} \hat{\mathcal{H}}_{t,\epsilon}(\mathbf{x}, \mathbf{w}, \mathbf{a}) &= \frac{1}{2} \sum_{\mu=1}^m \left\{ \Gamma_{t,\mu}(\mathbf{x}, w_\mu, \mathbf{a}_\mu)^2 + 2Z_\mu \Gamma_{t,\mu}(\mathbf{x}, w_\mu, \mathbf{a}_\mu) \right\} \\ &\quad + \frac{1}{2} \sum_{i=1}^n \left\{ R_1(t)(X_i^* - x_i)^2 + 2Z'_i \sqrt{R_1(t)}(X_i^* - x_i) \right\}. \end{aligned} \quad (226)$$

In order to prove Theorem 6 it remains to show that there exists a constant $C(\varphi, S, \alpha) > 0$ such that $\text{Var}(\ln \hat{Z}_{t,\epsilon}/n) \leq C(\varphi, S, \alpha)/n$. This concentration property together with (224) implies (218).

We will first show concentration w.r.t. all Gaussian variables $\Phi, \mathbf{V}, \mathbf{Z}, \mathbf{Z}', \mathbf{W}^*$ thanks to the classical Gaussian Poincaré inequality, then the concentration w.r.t. \mathbf{A} and finally the one w.r.t. \mathbf{X}^* using classical bounded differences arguments. The order in which we prove the concentrations matters. We recall here these two variances bounds. The reader can refer to [102] (Chapter 3) for detailed proofs of these statements.

Proposition 27 (Gaussian Poincaré inequality). *Let $\mathbf{U} = (U_1, \dots, U_N)$ be a vector of N independent standard normal random variables. Let $g : \mathbb{R}^N \rightarrow \mathbb{R}$ be a continuously differentiable function. Then*

$$\text{Var}(g(\mathbf{U})) \leq \mathbb{E}[\|\nabla g(\mathbf{U})\|^2]. \quad (227)$$

Proposition 28 (Bounded difference). *Let $\mathcal{U} \subset \mathbb{R}$. Let $g : \mathcal{U}^N \rightarrow \mathbb{R}$ a function that satisfies the bounded difference property, i.e., there exists some constants $c_1, \dots, c_N \geq 0$ such that*

$$\sup_{\substack{u_1, \dots, u_N \in \mathcal{U}^N \\ u'_i \in \mathcal{U}}} |g(u_1, \dots, u_i, \dots, u_N) - g(u_1, \dots, u'_i, \dots, u_N)| \leq c_i \quad \text{for all } 1 \leq i \leq N.$$

Let $\mathbf{U} = (U_1, \dots, U_N)$ be a vector of N independent random variables that take values in \mathcal{U} . Then

$$\text{Var}(g(\mathbf{U})) \leq \frac{1}{4} \sum_{i=1}^N c_i^2. \quad (228)$$

E.1.1 Concentration with respect to the Gaussian random variables $\mathbf{Z}, \mathbf{Z}', \mathbf{V}, \mathbf{W}^*, \Phi$

Lemma 24. *Let \mathbb{E}_G denotes the joint expectation w.r.t. $\mathbf{Z}, \mathbf{Z}', \mathbf{V}, \mathbf{W}^*, \Phi$ only. There exists a constant $C(\varphi, S, \alpha) > 0$ such that*

$$\mathbb{E} \left[\left(\frac{1}{n} \ln \hat{Z}_{t,\epsilon} - \frac{1}{n} \mathbb{E}_G \ln \hat{Z}_{t,\epsilon} \right)^2 \right] \leq \frac{C(\varphi, S, \alpha)}{n}. \quad (229)$$

Lemma 24 follows directly from Lemmas 25 and 26 below.

Lemma 25. *Let $\mathbb{E}_{\mathbf{Z}, \mathbf{Z}'}$ denotes the expectation w.r.t. \mathbf{Z}, \mathbf{Z}' only. There exists a constant $C(\varphi, S, \alpha) > 0$ such that*

$$\mathbb{E} \left[\left(\frac{1}{n} \ln \hat{Z}_{t,\epsilon} - \frac{1}{n} \mathbb{E}_{\mathbf{Z}, \mathbf{Z}'} \ln \hat{Z}_{t,\epsilon} \right)^2 \right] \leq \frac{C(\varphi, S, \alpha)}{n}. \quad (230)$$

Proof. We consider here $g = \ln \hat{Z}_{t,\epsilon}/n$ only as a function of \mathbf{Z} and \mathbf{Z}' and work conditionally on all other random variables. We have

$$\|\nabla g\|^2 = \sum_{\mu=1}^m \left| \frac{\partial g}{\partial Z_\mu} \right|^2 + \sum_{i=1}^n \left| \frac{\partial g}{\partial Z'_i} \right|^2. \quad (231)$$

Each of these partial derivatives are of the form $|\partial_u g| = |n^{-1} \langle \partial_u \hat{\mathcal{H}}_{t,\epsilon} \rangle_{\hat{\mathcal{H}}_{t,\epsilon}}|$ where the Gibbs bracket $\langle - \rangle_{\hat{\mathcal{H}}_{t,\epsilon}}$ pertains to the effective Hamiltonian (226). We find

$$\begin{aligned} \left| \frac{\partial g}{\partial Z_\mu} \right| &= n^{-1} |\langle \Gamma_{t,\mu} \rangle_{\hat{\mathcal{H}}_{t,\epsilon}}| \leq 2n^{-1} \sup |\varphi|, \\ \left| \frac{\partial g}{\partial Z'_i} \right| &= n^{-1} \sqrt{R_1(t)} |X_i^* - \langle x_i \rangle_{\hat{\mathcal{H}}_{t,\epsilon}}| \leq 2n^{-1} \sqrt{K} S, \end{aligned}$$

and replacing in (231) we get $\|\nabla g\|^2 \leq 4n^{-1} (\frac{m}{n} (\sup |\varphi|)^2 + K S^2)$. Applying Proposition 27 we have

$$\mathbb{E}_{\mathbf{Z}, \mathbf{Z}'} \left[\left(\frac{1}{n} \ln \hat{\mathcal{Z}}_{t,\epsilon} - \frac{1}{n} \mathbb{E}_{\mathbf{Z}, \mathbf{Z}'} \ln \hat{\mathcal{Z}}_{t,\epsilon} \right)^2 \right] \leq \frac{C(\varphi, S, \alpha)}{n}. \quad (232)$$

Taking the expectation in (232) gives the lemma. \square

Lemma 26. *There exists a constant $C(\varphi, S, \alpha) > 0$ such that*

$$\mathbb{E} \left[\left(\frac{1}{n} \mathbb{E}_{\mathbf{Z}, \mathbf{Z}'} \ln \hat{\mathcal{Z}}_{t,\epsilon} - \frac{1}{n} \mathbb{E}_G \ln \hat{\mathcal{Z}}_{t,\epsilon} \right)^2 \right] \leq \frac{C(\varphi, S, \alpha)}{n}. \quad (233)$$

Proof. We consider here $g = \mathbb{E}_{\mathbf{Z}, \mathbf{Z}'} \ln \hat{\mathcal{Z}}_{t,\epsilon} / n$ as a function of \mathbf{V} , \mathbf{W}^* , Φ and we work conditionally on the other random variables. Let $\partial_x \varphi$ be the derivative of φ w.r.t. its first argument. We compute

$$\begin{aligned} \left| \frac{\partial g}{\partial V_\mu} \right| &= n^{-1} \left| \mathbb{E}_{\mathbf{Z}, \mathbf{Z}'} \left\langle (\Gamma_{t,\mu} + Z_\mu) \frac{\partial \Gamma_{t,\mu}}{\partial V_\mu} \right\rangle_{\hat{\mathcal{H}}_{t,\epsilon}} \right| \\ &\leq n^{-1} \mathbb{E}_{\mathbf{Z}, \mathbf{Z}'} \left[(2 \sup |\varphi| + |Z_\mu|) 2\sqrt{K} \sup |\partial_x \varphi| \right] = n^{-1} \left(2 \sup |\varphi| + \sqrt{\frac{2}{\pi}} \right) 2\sqrt{K} \sup |\partial_x \varphi|. \end{aligned}$$

The same inequality holds for $|\frac{\partial g}{\partial W_\mu^*}|$. To compute the derivative w.r.t. $\Phi_{\mu i}$ we first remark

$$\begin{aligned} \frac{\partial \Gamma_{t,\mu}}{\partial \Phi_{\mu i}} &= \sqrt{\frac{1-t}{n}} \left\{ X_i^* \partial_x \varphi \left(\sqrt{\frac{1-t}{n}} [\Phi \mathbf{X}^*]_\mu + k_1(t) V_\mu + k_2(t) W_\mu^*, \mathbf{A}_\mu \right) \right. \\ &\quad \left. - x_i \partial_x \varphi \left(\sqrt{\frac{1-t}{n}} [\Phi \mathbf{x}]_\mu + k_1(t) V_\mu + k_2(t) w_\mu, \mathbf{a}_\mu \right) \right\}. \end{aligned}$$

Therefore, as $t \in [0, 1]$,

$$\begin{aligned} \left| \frac{\partial g}{\partial \Phi_{\mu i}} \right| &= n^{-1} \left| \mathbb{E}_{\mathbf{Z}, \mathbf{Z}'} \left\langle (\Gamma_{t,\mu} + Z_\mu) \frac{\partial \Gamma_{t,\mu}}{\partial \Phi_{\mu i}} \right\rangle_{\hat{\mathcal{H}}_{t,\epsilon}} \right| \\ &\leq n^{-3/2} \mathbb{E}_{\mathbf{Z}, \mathbf{Z}'} \left[(2 \sup |\varphi| + |Z_\mu|) 2S \sup |\partial_x \varphi| \right] = n^{-3/2} \left(2 \sup |\varphi| + \sqrt{\frac{2}{\pi}} \right) 2S \sup |\partial_x \varphi|. \end{aligned}$$

Putting these inequalities together we find

$$\begin{aligned} \|\nabla g\|^2 &= \sum_{\mu=1}^m \left| \frac{\partial g}{\partial V_\mu} \right|^2 + \sum_{\mu=1}^m \left| \frac{\partial g}{\partial W_\mu^*} \right|^2 + \sum_{\mu=1}^m \sum_{i=1}^n \left| \frac{\partial g}{\partial \Phi_{\mu i}} \right|^2 \\ &\leq 2 \frac{m}{n^2} \left(2 \sup |\varphi| + \sqrt{\frac{2}{\pi}} \right)^2 4K (\sup |\partial_x \varphi|)^2 + \frac{mn}{n^3} \left(2 \sup |\varphi| + \sqrt{\frac{2}{\pi}} \right)^2 4S^2 (\sup |\partial_x \varphi|)^2. \end{aligned}$$

The lemma follows again from Proposition 27. \square

E.1.2 Bounded difference with respect to A_μ

The next step is an application of the variance bound of Lemma 28 to show that $\mathbb{E}_G \ln \hat{\mathcal{Z}}_{t,\epsilon}/n$ concentrates w.r.t. \mathbf{A} (we still keep \mathbf{X}^* fixed for the moment).

Lemma 27. *Let \mathbb{E}_A denotes the expectation w.r.t. \mathbf{A} only. There exists a constant $C(\varphi, \alpha) > 0$ such that*

$$\mathbb{E} \left[\left(\frac{1}{n} \mathbb{E}_G \ln \hat{\mathcal{Z}}_{t,\epsilon} - \frac{1}{n} \mathbb{E}_{G,\mathbf{A}} \ln \hat{\mathcal{Z}}_{t,\epsilon} \right)^2 \right] \leq \frac{C(\varphi, \alpha)}{n}. \quad (234)$$

Proof. Let us consider $g = \mathbb{E}_G \ln \hat{\mathcal{Z}}_{t,\epsilon}/n$ as a function of \mathbf{A} only. Let $\nu \in \{1, \dots, m\}$. We must estimate variations $g(\mathbf{A}) - g(\mathbf{A}^{(\nu)})$ corresponding to two configurations \mathbf{A} and $\mathbf{A}^{(\nu)}$ with $A_\mu^{(\nu)} = A_\mu$ for $\mu \neq \nu$ and $A_\nu^{(\nu)} = \tilde{A}_\nu$. We will use the notations $\hat{\mathcal{H}}_{t,\epsilon}^{(\nu)}$ and $\Gamma_{t,\mu}^{(\nu)}$ to denote respectively the quantities $\hat{\mathcal{H}}_{t,\epsilon}$ and $\Gamma_{t,\mu}$ where \mathbf{A} is replaced by $\mathbf{A}^{(\nu)}$. By an application of Jensen's inequality one finds

$$\frac{1}{n} \mathbb{E}_G \langle \hat{\mathcal{H}}_{t,\epsilon}^{(\nu)} - \hat{\mathcal{H}}_{t,\epsilon} \rangle_{\hat{\mathcal{H}}_{t,\epsilon}^{(\nu)}} \leq g(\mathbf{A}) - g(\mathbf{A}^{(\nu)}) \leq \frac{1}{n} \mathbb{E}_G \langle \hat{\mathcal{H}}_{t,\epsilon}^{(\nu)} - \hat{\mathcal{H}}_{t,\epsilon} \rangle_{\hat{\mathcal{H}}_{t,\epsilon}} \quad (235)$$

where the Gibbs brackets pertain to the effective Hamiltonians (226). From (226) we obtain

$$\hat{\mathcal{H}}_{t,\epsilon}^{(\nu)} - \hat{\mathcal{H}}_{t,\epsilon} = \frac{1}{2} \sum_{\mu=1}^m \left(\Gamma_{t,\mu}^{(\nu)2} - \Gamma_{t,\mu}^2 + 2Z_\mu(\Gamma_{t,\mu}^{(\nu)} - \Gamma_{t,\mu}) \right) = \frac{1}{2} \left(\Gamma_{t,\nu}^{(\nu)2} - \Gamma_{t,\nu}^2 + 2Z_\nu(\Gamma_{t,\nu}^{(\nu)} - \Gamma_{t,\nu}) \right).$$

Consequently

$$\begin{aligned} \frac{1}{2n} \mathbb{E}_G \left\langle \Gamma_{t,\nu}^{(\nu)2} - \Gamma_{t,\nu}^2 + 2Z_\nu(\Gamma_{t,\nu}^{(\nu)} - \Gamma_{t,\nu}) \right\rangle_{\hat{\mathcal{H}}_{t,\epsilon}^{(\nu)}} &\leq g(\mathbf{A}) - g(\mathbf{A}^{(\nu)}) \\ &\leq \frac{1}{2n} \mathbb{E}_G \left\langle \Gamma_{t,\nu}^{(\nu)2} - \Gamma_{t,\nu}^2 + 2Z_\nu(\Gamma_{t,\nu}^{(\nu)} - \Gamma_{t,\nu}) \right\rangle_{\hat{\mathcal{H}}_{t,\epsilon}}. \end{aligned} \quad (236)$$

Notice that $|\Gamma_{t,\nu}^{(\nu)2} - \Gamma_{t,\nu}^2 + 2Z_\nu(\Gamma_{t,\nu}^{(\nu)} - \Gamma_{t,\nu})| \leq 8(\sup |\varphi|)^2 + 8|Z_\nu| \sup |\varphi|$. Thus we conclude by (236) that g satisfies a bounded difference property:

$$|g(\mathbf{A}) - g(\mathbf{A}^{(\nu)})| \leq \frac{4}{n} \sup |\varphi| \left(\sup |\varphi| + \sqrt{\frac{2}{\pi}} \right). \quad (237)$$

Lemma 27 follows then by an application of Proposition 28. \square

E.1.3 Bounded difference with respect to X_i^*

Let $\mathbb{E}_\Theta = \mathbb{E}_{A,G}$ denote the expectation w.r.t. all quenched variables except \mathbf{X}^* . It remains to bound the variance of $\mathbb{E}_\Theta \ln \hat{\mathcal{Z}}_{t,\epsilon}/n$ (which only depends on \mathbf{X}^*).

Lemma 28. *There exists a constant $C(\varphi, S, \alpha) > 0$ such that*

$$\mathbb{E} \left[\left(\frac{1}{n} \mathbb{E}_\Theta \ln \hat{\mathcal{Z}}_{t,\epsilon} - \frac{1}{n} \mathbb{E} \ln \hat{\mathcal{Z}}_{t,\epsilon} \right)^2 \right] \leq \frac{C(\varphi, S, \alpha)}{n}. \quad (238)$$

Proof. The lemma is proved using again a bounded difference argument. Let $g = \mathbb{E}_\Theta \ln \hat{\mathcal{Z}}_{t,\epsilon}/n$ a function of \mathbf{X}^* . Let $j \in \{1, \dots, n\}$. Let $\mathbf{X}^*, \mathbf{X}^{*(j)} \in [-S, S]^n$ be two input signals such that $X_i^{*(j)} = X_i^*$ for $i \neq j$.

We are going to interpolate between $g(\mathbf{X}^*)$ and $g(\mathbf{X}^{*(j)})$. For $s \in [0, 1]$ we define

$$\psi(s) = g(s\mathbf{X}^* + (1-s)\mathbf{X}^{*(j)}).$$

Obviously $\psi(1) = g(\mathbf{X}^*)$ and $\psi(0) = g(\mathbf{X}^{*(j)})$. Using Gaussian integration by parts, it is not difficult to verify that for $s \in [0, 1]$

$$|\psi'(s)| \leq \frac{C(\varphi, S, \alpha)}{n}.$$

This implies the bounded difference property $|g(\mathbf{X}^*) - g(\mathbf{X}^{*(j)})| \leq C(\varphi, S, \alpha)/n$ and using Proposition 28 we obtain the lemma. \square

E.1.4 Proof of Theorem 6

From Lemmas 24, 27 and 28 above, we obtain directly that $\text{Var}(\ln \hat{\mathcal{Z}}_{t,\epsilon}/n) \leq C(\varphi, S, \alpha)/n$ for some constant $C(\varphi, S, \alpha) > 0$. As mentioned before this implies, thanks to (224), the Theorem 6.

E.2 Concentration of the overlap

In this appendix we provide the proof of Proposition 4. Recall the notation $\langle - \rangle_{n,t,\epsilon}$ for the Gibbs bracket associated to the Hamiltonian (95). It is crucial that it preserves the Nishimori identity of Appendix A.1, i.e. it must come from an inference problem with known parameters. Consider the corresponding average free entropy $f_{n,\epsilon}(t)$. In this section we think of it as a function of $R_1 = R_1(t, \epsilon)$ and $R_2 = R_2(t, \epsilon)$ given by (89), i.e. $(R_1, R_2) \mapsto f_{n,\epsilon}(t)$. Similarly the free entropy for a realization of the quenched variables is also viewed here as a function $(R_1, R_2) \mapsto F_{n,\epsilon}(t) := \ln \mathcal{Z}_{t,\epsilon}(\mathbf{Y}_t, \mathbf{Y}'_t, \Phi, \mathbf{V})/n$. For this section, we drop the indices in the Gibbs bracket $\langle - \rangle_{n,t,\epsilon}$ and simply write $\langle - \rangle$.

Let

$$\mathcal{L} := \frac{1}{n} \sum_{i=1}^n \left(\frac{x_i^2}{2} - x_i X_i^* - \frac{x_i Z'_i}{2\sqrt{R_1}} \right).$$

The fluctuations of the overlap $Q := n^{-1} \sum_{i=1}^n X_i^* x_i$ and those of \mathcal{L} are related through the remarkable identity

$$\mathbb{E} \langle (\mathcal{L} - \mathbb{E} \langle \mathcal{L} \rangle)^2 \rangle = \frac{1}{4} \mathbb{E} \langle (Q - \mathbb{E} \langle Q \rangle)^2 \rangle + \frac{1}{2} \mathbb{E} [\langle Q^2 \rangle - \langle Q \rangle^2] + \frac{1}{4nR_1} \mathbb{E} [(X_1^*)^2]. \quad (239)$$

In particular

$$\mathbb{E} \langle (\mathcal{L} - \mathbb{E} \langle \mathcal{L} \rangle)^2 \rangle \geq \frac{1}{4} \mathbb{E} \langle (Q - \mathbb{E} \langle Q \rangle)^2 \rangle. \quad (240)$$

A detailed derivation of (239) involves only lengthy but straightforward algebra, using the Nishimori identity and integrations by parts w.r.t. the Gaussian noise Z'_i , and can be found in Sec. 6 of [52]. Proposition 4 is then a direct consequence of the following:

Proposition 29 (Concentration of \mathcal{L} on $\mathbb{E} \langle \mathcal{L} \rangle$). *Let $\mathcal{B}_n := [s_n, 2s_n]^2$, where the sequence $(s_n) \in (0, 1/2]^{\mathbb{N}}$. Assume that the interpolation functions (r_ϵ) and (q_ϵ) are regular (recall Definition 1). Under assumptions (H1), (H2) and (H3) there exists a constant $C(\varphi, S, \alpha)$ such that*

$$\int_{\mathcal{B}_n} d\epsilon \mathbb{E} \langle (\mathcal{L} - \mathbb{E} \langle \mathcal{L} \rangle_{n,t,\epsilon})^2 \rangle_{n,t,\epsilon} \leq \frac{C(\varphi, S, \alpha)}{n^{1/4}}. \quad (241)$$

The proof of this proposition is broken in two parts. Notice that

$$\mathbb{E} \langle (\mathcal{L} - \mathbb{E} \langle \mathcal{L} \rangle)^2 \rangle = \mathbb{E} \langle (\mathcal{L} - \langle \mathcal{L} \rangle)^2 \rangle + \mathbb{E} [(\langle \mathcal{L} \rangle - \mathbb{E} \langle \mathcal{L} \rangle)^2]. \quad (242)$$

Thus it suffices to prove the two following lemmas (see the proofs below). The first lemma expresses concentration w.r.t. the posterior distribution (or “thermal fluctuations”) and is an elementary consequence of concavity properties of the free entropy and the Nishimori identity.

Lemma 29 (Concentration of \mathcal{L} on $\langle \mathcal{L} \rangle$). *Under the same hypotheses as in Proposition 29 we have*

$$\int_{\mathcal{B}_n} d\epsilon \mathbb{E} \langle (\mathcal{L} - \langle \mathcal{L} \rangle_{n,t,\epsilon})^2 \rangle_{n,t,\epsilon} \leq \frac{\rho(1+\rho)}{n}. \quad (243)$$

The second lemma expresses the concentration of the average overlap w.r.t. the realizations of quenched disorder variables and is a consequence of the concentration of the free entropy (more precisely Theorem 6 in Appendix E.1).

Lemma 30 (Concentration of $\langle \mathcal{L} \rangle$ on $\mathbb{E} \langle \mathcal{L} \rangle$). *Under the same hypotheses as in Proposition 29 there exists a constant $C(\varphi, S, \alpha)$ such that*

$$\int_{\mathcal{B}_n} d\epsilon \mathbb{E} [(\langle \mathcal{L} \rangle_{n,t,\epsilon} - \mathbb{E} \langle \mathcal{L} \rangle_{n,t,\epsilon})^2] \leq \frac{C(\varphi, S, \alpha)}{n^{1/4}}. \quad (244)$$

We now turn to the proof of Lemmas 29 and 30. The main ingredient is a set of formulas for the first two derivatives of the free entropy w.r.t. $R_1 = R_1(t, \epsilon)$. For any given realisation of the quenched disorder,

$$\frac{dF_{n,\epsilon}(t)}{dR_1} = -\langle \mathcal{L} \rangle - \frac{1}{2n} \sum_{i=1}^n \left((X_i^*)^2 + \frac{1}{\sqrt{R_1}} X_i^* Z_i' \right), \quad (245)$$

$$\frac{1}{n} \frac{d^2 F_{n,\epsilon}(t)}{dR_1^2} = \langle \mathcal{L}^2 \rangle - \langle \mathcal{L} \rangle^2 - \frac{1}{4n^2 R_1^{3/2}} \sum_{i=1}^n \langle x_i \rangle Z_i'. \quad (246)$$

Averaging (245) and (246), using a Gaussian integration by parts w.r.t. Z_i' and the (Nishimori) identity $\mathbb{E} \langle x_i X_i^* \rangle = \mathbb{E} [\langle x_i \rangle^2]$ we find

$$\frac{df_{n,\epsilon}(t)}{dR_1} = -\mathbb{E} \langle \mathcal{L} \rangle - \frac{\rho}{2} = \frac{1}{2n} \sum_{i=1}^n \mathbb{E} [\langle x_i \rangle^2] - \frac{\rho}{2}, \quad (247)$$

$$\frac{1}{n} \frac{d^2 f_{n,\epsilon}(t)}{dR_1^2} = \mathbb{E} [\langle \mathcal{L}^2 \rangle - \langle \mathcal{L} \rangle^2] - \frac{1}{4n^2 R_1} \sum_{i=1}^n \mathbb{E} [\langle x_i^2 \rangle - \langle x_i \rangle^2]. \quad (248)$$

Proof of Lemma 29

From (248) we have

$$\begin{aligned} \mathbb{E} \langle (\mathcal{L} - \langle \mathcal{L} \rangle)^2 \rangle &= \frac{1}{n} \frac{d^2 f_{n,\epsilon}(t)}{dR_1^2} + \frac{1}{4n^2 R_1} \sum_{i=1}^n \mathbb{E} [\langle x_i^2 \rangle - \langle x_i \rangle^2] \\ &\leq \frac{1}{n} \frac{d^2 f_{n,\epsilon}(t)}{dR_1^2} + \frac{\rho}{4n\epsilon_1}, \end{aligned} \quad (249)$$

where we used $\mathbb{E} \langle x_i^2 \rangle = \mathbb{E}_{P_0} [(X^*)^2] = \rho$ by the Nishimori identity, and $R_1 \geq \epsilon_1$. Recall $\mathcal{B}_n := [s_n, 2s_n]^2$. By assumption q and r are regular. Therefore $R^t : (\epsilon_1, \epsilon_2) \mapsto (R_1(t, \epsilon), R_2(t, \epsilon))$ is a diffeomorphism whose

Jacobian $J(R^t)$ verifies $J(R^t)(\epsilon) \geq 1$ for all $\epsilon \in \mathcal{B}_n$. Integrating over $\epsilon \in \mathcal{B}_n$ we obtain

$$\begin{aligned} \int_{\mathcal{B}_n} d\epsilon \mathbb{E} \langle (\mathcal{L} - \langle \mathcal{L} \rangle)^2 \rangle &\leq \frac{1}{n} \int_{R^t(\mathcal{B}_n)} \frac{dR_1 dR_2}{J(R^t)} \frac{d^2 f_{n,\epsilon}(t)}{dR_1^2} + \frac{\rho s_n}{4n} \int_{s_n}^{2s_n} \frac{d\epsilon_1}{\epsilon_1} \\ &\leq \frac{1}{n} \int_{R^t(\mathcal{B}_n)} dR_1 dR_2 \frac{d^2 f_{n,\epsilon}(t)}{dR_1^2} + \frac{\rho s_n}{4n} \ln 2, \end{aligned} \quad (250)$$

where in the integral above $J(R^t)$ is a function of $(R^t)^{-1}(R_1, R_2)$. Note that from (89) we have $R^t(\mathcal{B}_n) \subset [s_n, 2s_n + r_{\max}] \times [s_n, 2s_n + \rho]$ and therefore

$$\begin{aligned} \int_{\mathcal{B}_n} d\epsilon \mathbb{E} \langle (\mathcal{L} - \langle \mathcal{L} \rangle)^2 \rangle &\leq \frac{1}{n} \int_{s_n}^{2s_n + \rho} dR_2 \left\{ \frac{df_{n,\epsilon}(t)}{dR_1} \Big|_{R_1=s_n} - \frac{df_{n,\epsilon}(t)}{dR_1} \Big|_{R_1=2s_n+r_{\max}} \right\} + \frac{\rho s_n}{4n} \ln 2 \\ &\leq \frac{\rho(s_n + \rho)}{n} + \frac{\rho s_n}{4n} \ln 2 \end{aligned} \quad (251)$$

using (247) combined with $\mathbb{E} \langle x_i^2 \rangle = \rho$ to assert that the derivative of the free entropy is bounded in absolute value by $\rho/2$. This concludes the proof of Lemma 29 using $s_n \leq 1/2$ and $(\ln 2)/4 < 1$. \square

Proof of Lemma 30

Consider the two functions

$$\tilde{F}(R_1) := F_{n,\epsilon}(t) - \frac{\sqrt{R_1}}{n} S \sum_{i=1}^n |Z'_i|, \quad \tilde{f}(R_1) := \mathbb{E} \tilde{F}(R_1) = f_{n,\epsilon}(t) - \frac{\sqrt{R_1}}{n} S \sum_{i=1}^n \mathbb{E} |Z'_i|. \quad (252)$$

Because of (246) we see that the second derivative of $\tilde{F}(R_1)$ is positive so that it is convex (without this extra term $F_{n,\epsilon}(t)$ is not necessarily convex in R_1 , although $f_{n,\epsilon}(t)$ is, which can be shown easily). Note that $\tilde{f}(R_1)$ is convex too. Convexity allows us to use the following lemma (proved at the end of this section):

Lemma 31 (A bound on differences of derivatives due to convexity). *Let $G(x)$ and $g(x)$ be convex functions. Let $\delta > 0$ and define $C_\delta^+(x) := g'(x + \delta) - g'(x) \geq 0$ and $C_\delta^-(x) := g'(x) - g'(x - \delta) \geq 0$. Then*

$$|G'(x) - g'(x)| \leq \delta^{-1} \sum_{u \in \{x-\delta, x, x+\delta\}} |G(u) - g(u)| + C_\delta^+(x) + C_\delta^-(x).$$

From (252)

$$\tilde{F}(R_1) - \tilde{f}(R_1) = F_{n,\epsilon}(t) - f_{n,\epsilon}(t) - \sqrt{R_1} S A, \quad \text{with} \quad A = \frac{1}{n} \sum_{i=1}^n (|Z'_i| - \mathbb{E} |Z'_i|). \quad (253)$$

and from (245), (247) we obtain for the difference of derivatives (w.r.t. R_1)

$$\tilde{F}'(R_1) - \tilde{f}'(R_1) = \mathbb{E} \langle \mathcal{L} \rangle - \langle \mathcal{L} \rangle + \frac{\rho}{2} - \frac{1}{2n} \sum_{i=1}^n \left((X_i^*)^2 + \frac{1}{\sqrt{R_1}} X_i^* Z'_i \right) - \frac{SA}{2\sqrt{R_1}}. \quad (254)$$

From (253), (254) it is easy to show that Lemma 31 implies

$$\begin{aligned} |\langle \mathcal{L} \rangle - \mathbb{E}\langle \mathcal{L} \rangle| &\leq \delta^{-1} \sum_{u \in \{R_1 - \delta, R_1, R_1 + \delta\}} (|F_{n,\epsilon}(t, R_1 = u) - f_{n,\epsilon}(t, R_1 = u)| + S|A|\sqrt{u}) \\ &\quad + C_\delta^+(R_1) + C_\delta^-(R_1) + \frac{S|A|}{2\sqrt{\epsilon_1}} + \left| \frac{\rho}{2} - \frac{1}{2n} \sum_{i=1}^n \left((X_i^*)^2 + \frac{1}{\sqrt{R_1}} X_i^* Z_i' \right) \right| \end{aligned} \quad (255)$$

where $C_\delta^+(R_1) := \tilde{f}'(R_1 + \delta) - \tilde{f}'(R_1) \geq 0$ and $C_\delta^-(R_1) := \tilde{f}'(R_1) - \tilde{f}'(R_1 - \delta) \geq 0$. We used $R_1 \geq \epsilon_1$ for the term $S|A|/(2\sqrt{\epsilon_1})$. Note that δ will be chosen later on strictly smaller than s_n (namely $\delta = s_n n^{-1/4}$) so that $R_1 - \delta \geq \epsilon_1 - \delta \geq s_n - \delta$ remains positive. Remark that by independence of the noise variables $\mathbb{E}[A^2] \leq an^{-1}$ for some constant $a > 0$; and that by independence between signal and noise, the last term in the absolute value in (255), call it B , satisfies $\mathbb{E}[B^2] \leq bn^{-1}$ for some constant $b > 0$. We now square the identity (255) and take its expectation. Then using $(\sum_{i=1}^p v_i)^2 \leq p \sum_{i=1}^p v_i^2$ (by convexity), and that $R_1 \leq K$ ($K = 1 + \max(\rho, r_{\max})$ upper bounds both R_1 and R_2 given by (89)), as well as the free entropy concentration Theorem 6,

$$\frac{1}{10} \mathbb{E}[(\langle \mathcal{L} \rangle - \mathbb{E}\langle \mathcal{L} \rangle)^2] \leq 3\delta^{-2} (C + aS^2(K + \delta)) \frac{1}{n} + C_\delta^+(R_1)^2 + C_\delta^-(R_1)^2 + \frac{S^2 a}{4\epsilon_1 n} + \frac{b}{n}. \quad (256)$$

where $C = C(\varphi, S, \alpha)$ is a positive constant depending only on φ, S and α that comes from the use of Theorem 6. Recall $|C_\delta^\pm(R_1)| = |\tilde{f}'(R_1 \pm \delta) - \tilde{f}'(R_1)|$. We have

$$|\tilde{f}'(R_1)| \leq \frac{1}{2} \left(\rho + \frac{S}{\sqrt{R_1}} \right) \leq \frac{1}{2} \left(\rho + \frac{S}{\sqrt{\epsilon_1}} \right) \quad (257)$$

from (247), (252) and $R_1 \geq \epsilon_1$. This implies $|C_\delta^\pm(R_1)| \leq \rho + S/\sqrt{\epsilon_1} \leq \rho + S/\sqrt{s_n}$ as $\epsilon_1 \geq s_n$. Recall also that $\mathcal{B}_n := [s_n, 2s_n]^2$. Then

$$\begin{aligned} &\int_{\mathcal{B}_n} d\epsilon \{ C_\delta^+(R_1(t, \epsilon))^2 + C_\delta^-(R_1(t, \epsilon))^2 \} \\ &\leq \left(\rho + \frac{S}{\sqrt{s_n}} \right) \int_{\mathcal{B}_n} d\epsilon \{ C_\delta^+(R_1(t, \epsilon)) + C_\delta^-(R_1(t, \epsilon)) \} \\ &= \left(\rho + \frac{S}{\sqrt{s_n}} \right) \int_{R^t(\mathcal{B}_n)} \frac{dR_1 dR_2}{J(R^t)} \{ C_\delta^+(R_1) + C_\delta^-(R_1) \} \\ &\leq \left(\rho + \frac{S}{\sqrt{s_n}} \right) \int_{R^t(\mathcal{B}_n)} dR_1 dR_2 \{ C_\delta^+(R_1) + C_\delta^-(R_1) \} \\ &\leq \left(\rho + \frac{S}{\sqrt{s_n}} \right) \int_{s_n}^{2s_n + \rho} dR_2 \left[\left(\tilde{f}(2s_n + r_{\max} + \delta) - \tilde{f}(2s_n + r_{\max} - \delta) \right) \right. \\ &\quad \left. + \left(\tilde{f}(s_n - \delta) - \tilde{f}(s_n + \delta) \right) \right]. \end{aligned} \quad (258)$$

where we used that the Jacobian $J(R^t)$ of the \mathcal{C}^1 diffeomorphism $R^t : (\epsilon_1, \epsilon_2) \mapsto (R_1(t, \epsilon), R_2(t, \epsilon))$ is greater or equal to 1 (by regularity of the interpolation functions q and r) and $R^t(\mathcal{B}_n) \subset [s_n, 2s_n + r_{\max}] \times [s_n, 2s_n + \rho]$. The mean value theorem and (257) imply $|\tilde{f}(R_1 - \delta) - \tilde{f}(R_1 + \delta)| \leq \delta(\rho + S/\sqrt{s_n})$ uniformly in R_2 . Therefore

$$\int_{\mathcal{B}_n} d\epsilon \{ C_\delta^+(R_1(t, \epsilon))^2 + C_\delta^-(R_1(t, \epsilon))^2 \} \leq 2\delta(s_n + \rho) \left(\rho + \frac{S}{\sqrt{s_n}} \right)^2. \quad (259)$$

Thus, integrating (256) over $\epsilon \in \mathcal{B}_n$ yields (using $\text{Vol}(\mathcal{B}_n) = s_n^2$ and $s_n \leq 1/2$)

$$\begin{aligned} & \int_{\mathcal{B}_n} d\epsilon \mathbb{E}[(\langle \mathcal{L} \rangle - \mathbb{E}[\langle \mathcal{L} \rangle])^2] \\ & \leq 30(C + aS^2\sqrt{K + \delta})\delta^{-2}s_n^2n^{-1} + 20\delta(1/2 + \rho)\left(\rho + \frac{S}{\sqrt{s_n}}\right)^2 + 5S^2a\frac{\ln 2}{2}\frac{s_n}{n} + \frac{bs_n^2}{n}. \end{aligned}$$

Finally we choose $\delta = s_n n^{-1/4}$ and obtain the desired result. \square

Proof of Lemma 31: Convexity implies that for any $\delta > 0$ we have

$$\begin{aligned} G'(x) - g'(x) & \leq \frac{G(x + \delta) - G(x)}{\delta} - g'(x) \\ & \leq \frac{G(x + \delta) - G(x)}{\delta} - g'(x) + g'(x + \delta) - \frac{g(x + \delta) - g(x)}{\delta} \\ & = \frac{G(x + \delta) - g(x + \delta)}{\delta} - \frac{G(x) - g(x)}{\delta} + C_\delta^+(x), \\ G'(x) - g'(x) & \geq \frac{G(x) - G(x - \delta)}{\delta} - g'(x) + g'(x - \delta) - \frac{g(x) - g(x - \delta)}{\delta} \\ & = \frac{G(x) - g(x)}{\delta} - \frac{G(x - \delta) - g(x - \delta)}{\delta} - C_\delta^-(x). \end{aligned}$$

Combining these two inequalities ends the proof. \square

Appendix F: Details on numerics

Most of our experiments and codes are provided on the associated GitHub repository [62], with codes in the Julia programming language [103] (with a Jupyter notebook interface) and in matlab. In this appendix, we shall give additional details on how the plots have been obtained.

F.1 General purpose algorithms

We have been using free available softwares in our experiments: Standard machine learning tasks such as LASSO or logistic regression were done using scikit-learn [65]. Keras [66], with a tensorflow backend [104], was used for neural networks. We also used CVXPY, a python-embedded language for performing convex optimization [84], as well as PhaseMax for phase retrieval [85] experiments.

Figure 4 (in both the main text and SI) contrasts results of these general purpose algorithms with the optimal generalization error in three classification problems. In the left pannel of Fig. 4 (for the binary perceptron), we used logistic regression with hand-tuned ℓ_2 regularization (basically, we have hand-selected the regularization parameter in order to obtain the best results) with the function Logistic Regression in the software scikit-learn [65]. In the center pannel of Fig. 4 (this time for a sparse signal), we used the same software but this time with a sparsity enhancing ℓ_1 regularization, again fine-tuned by hand. In the right pannel of Fig. 4, we show, in the inset, how a neural network with 2 hidden layers was able to learn only approximately the “symmetric door” rule. In this experiment, we used Keras with a tensorflow backend.

The data are created with a signal-vector of dimension $n = 2500$ and are sent into a network made with a first layer of dimension 2500×64 followed by a rectified linear unit (ReLU), and a dropout layer with fraction 0.2 for regularization. This is followed by a second layer of dimension 64×64 , this time with a sigmoid activation and again a dropout layer. Finally, we classify with a final output layer with a softmax, using the categorical cross-entropy as the loss function. The minimization is done using the RMSprop optimizer for 1000 epochs. The code is shown in the GitHub repository [62].

We have tried many variations around this network. Interestingly, the dropout layers have a strong effect on the regularization, and help significantly in improving the generalization error. Also interestingly, the number of epochs used for fitting was an important parameter. Indeed, the quality of the fit improves drastically as the number of epochs is increased: It seems that it actually takes a lot of time to escape the initial point, where prediction is just as bad as random. Finally, we also tried to increase the depth of the neural nets. Interestingly this did not affect the performance and the network was fitting the data and generalizing just as well with deeper and deeper networks. We believe that it should be very instructive to further study empirically this problem.

In Fig. 5 in this SI, similar plots are shown for three regression problems. In the left panel, we used the LASSO function in the software scikit-learn [65] with its sparsity enhancing ℓ_1 regularization, again fine-tuned by hand.

For the middle panel, we had to turn to a different software. In this case, the idea was to solve $\mathbf{Y} = \text{Relu}(\Phi \mathbf{x}) = \max(0, \Phi \mathbf{x})$ (componentwise) subject to a sparse penalty on \mathbf{x} . Luckily, this can be turned into a linear programming framework: Minimize the ℓ_1 norm of \mathbf{x} subject to the constraint that $\mathbf{Y} = \text{Relu}(\Phi \mathbf{x})$ which can be implemented by enforcing $Y_\mu = \Phi_\mu \cdot \mathbf{x}$ for μ 's such that $Y_\mu > 0$, and $\Phi_\mu \cdot \mathbf{x} < 0$ for μ 's such that $Y_\mu = 0$. This linear program is solved with CVXPY [84] (CVX-2). We also show the results when only the indices μ associated with positive $Y_\mu > 0$ are used for comparison (CVX-1). Finally, for the right figure, we used PhaseMax [85] out-of-the-box to solve the problem.

F.2 Evaluating the replica formula

In order to evaluate numerically the replica formula (41) we proceed as it is common in the statistical physics literature since the early papers on spin glasses. First, we found the critical points (42) by iterating the state evolution equations (67) —also called the replica self-consistent equations— starting from two different initial conditions ($q^{t=0} = 0$ and $q^{t=0} = \rho$). Next we computed the associated value of the free entropy and then selected, if two different fixed points were found, the correct one following the prescription given by Theorem 1. We also took special care in checking that we could not identify other fixed points. An example of such a procedure is shown in the GitHub repository [62], with codes in the Julia programming language [103], for the perceptron problem.

F.3 Breaking the symmetry in GAMP

A last notable point concerns the symmetry issue in GAMP. Indeed, when $q = 0$ is a fixed point of the state evolution (which is the case if the prior has zero mean and the channel is symmetric), then GAMP should stay in this fixed point forever. This is the case, for instance, for two problems considered in the present paper: The symmetric door output function $\varphi(z) = \text{sgn}(|z| - K)$ with a Rademacher prior ± 1 , and for the sign-less channel $\varphi(z) = |z|$. In both cases, both z and $-z$ are giving the same output, and therefore so does both \mathbf{X}^* and $-\mathbf{X}^*$. Notice, however, that this is not a problem in the computation of the free entropy. Here, one has to compute all the fixed points anyway. It is also not a problem if the prior is breaking the symmetry (for instance if one is working with a binary signal where $X_i^* = 1$ with probability $p_+ = 1/2 + \epsilon$ and $X_i^* = -1$ with probability $p_- = 1/2 - \epsilon$ with, say, $\epsilon = 10^{-7}$). In this case the symmetry is broken, GAMP works, and the state evolution predicts its behavior correctly. Even though this problem is thus restricted to a very small class of channels and priors, and even though perturbations solve it, it is still an interesting mathematical challenge, especially from the rigorous point of view. Indeed, this problem has attracted attention recently where initializations based on spectral algorithms were analyzed [61].

In the present paper, we adopted a pragmatic point of view. We did not break the symmetry in the data generative model (as that would make the problem slightly easier), instead we broke the symmetry in the GAMP solver, thus making it slightly but un-noticeably suboptimal. An example of our code is given in the

associated GitHub repository [62]. We created the data in the symmetric manner, but when we ran GAMP to solve the problem, we broke the channel symmetry slightly. For instance, instead of solving with a door function that returns 1 only for $-0.674489 < z < 0.674489$, we use a function that instead returned 1 for $-0.674489 < z < 0.6745$. The same strategy was used for the absolute value function, which can be replaced by $\varphi(z) = z$ for $x > -\epsilon$ and $-z$ otherwise. Again, when ϵ is small, this makes only an unnoticeable difference in the figures. This trick allowed GAMP to solve symmetric problems without trouble in practice, and to reach perfect recovery even in the symmetric problems as close to the theoretical threshold as numerically desired.