

High-dimensional optimization under nonconvex excluded volume constraints

Antonio Sclocchi¹ and Pierfrancesco Urbani²

¹*Institute of Physics, École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland*

²*Université Paris-Saclay, CNRS, CEA, Institut de physique théorique, 91191 Gif-sur-Yvette, France*



(Received 17 September 2021; accepted 1 January 2022; published 23 February 2022)

We consider high-dimensional random optimization problems where the dynamical variables are subjected to nonconvex excluded volume constraints. We focus on the case in which the cost function is a simple quadratic cost and the excluded volume constraints are modeled by a perceptron constraint satisfaction problem. We show that depending on the density of constraints, one can have different situations. If the number of constraints is small, one typically has a phase where the ground state of the cost function is unique and sits on the boundary of the island of configurations allowed by the constraints. In this case, there is a hypostatic number of marginally satisfied constraints. If the number of constraints is increased one enters a glassy phase where the cost function has many local minima sitting again on the boundary of the regions of allowed configurations. At the phase transition point, the total number of marginally satisfied constraints becomes equal to the number of degrees of freedom in the problem and therefore we say that these minima are isostatic. We conjecture that by increasing further the constraints the system stays isostatic up to the point where the volume of available phase space shrinks to zero. We derive our results using the replica method and we also analyze a dynamical algorithm, the Karush-Kuhn-Tucker algorithm, through dynamical mean-field theory and we show how to recover the results of the replica approach in the replica symmetric phase.

DOI: [10.1103/PhysRevE.105.024134](https://doi.org/10.1103/PhysRevE.105.024134)

I. INTRODUCTION

Several complex systems can be described by a large number of variables whose dynamics is trying to optimize some cost function under a set of constraints. In the most general setting, the cost function is arbitrary and the constraints may be either inequalities or equalities. In this work, we consider the generic situation in which there is a set of N real variables $\underline{x} = \{x_1, \dots, x_N\}$ and a set of $M = \alpha N$ inequality constraints

$$\mathcal{C} = \{g^\mu(\underline{x}) \geq 0, \mu = 1, \dots, M = \alpha N\}, \quad (1)$$

each characterized by a function $g^\mu(\underline{x})$ and we consider a cost function $H(\underline{x})$.

An example of this situation is found in classification problems in machine learning. In the simplest setting of binary classification, one has a set of inequalities that enforce the correct separation of the training set into two classes. The parameters of the machine (which correspond to the variables x_j here) must be adjusted to correctly satisfy these constraints. Therefore, the supervised learning task is recast into a continuous constraint satisfaction problem (CCSP). In many cases one wants to enforce that, once a solution of the problem is found, it satisfies some requirements. For example, in order to prevent numerical instabilities, one would like to have the parameters of the network be quantities that are of the good order of magnitude and do not explode to infinity. To achieve this, one typically uses regularization schemes as, for example, Ridge regularization. In other cases, one is interested in finding sparse networks in which many of the parameters are

zero. This can also be achieved by imposing an appropriate cost function on the parameters [1].

In all these cases, the machine performing the classification task is defined by the form of the functions $g_\mu(\underline{x})$ and the regularizing cost function is given by $H(\underline{x})$, which introduces some penalty to prevent that the parameters of the network behave in a way that is not the desired one. In this work, we will be rather generic and we will not focus on any practical application. We will rather focus on the characterization of the landscape of the cost function H subjected to the constraints \mathcal{C} . In particular, we want to find the ground state of the cost function defined by

$$\underline{x}^* = \underset{\underline{x}: g^\mu(\underline{x}) \geq 0 \forall \mu}{\operatorname{argmin}} H(\underline{x}). \quad (2)$$

In order to have a well-posed problem, we will assume that there exists a region of the phase space of \underline{x} such that all constraints are satisfied. We will call this region the SAT region (the terminology comes from constraint satisfaction problems in computer science [2]). Examples of optimization problems as Eq. (2) can be easily constructed and classified according to the properties of the corresponding cost function. We list here just a set of general cases in which the cost function is

(i) Separable: This means that each degree of freedom tries to optimize a cost function that is independent of the configuration of the other degrees of freedom. An example in this class is

$$H(\underline{x}) = \frac{1}{2} \sum_{i=1}^N (x_i - 1)^2 \quad (3)$$

but other nonconvex potentials with local minima may be chosen.

(ii) Nonseparable and convex: The cost function of each degree of freedom depends on the configuration of part or the whole system and the global cost is a convex function when it is considered in the whole phase space (in absence of the excluded volume constraints). As an example one can consider:

$$H(\underline{x}) = \frac{1}{P} \sum_{\mu=1}^P x_i \eta_i^\mu \eta_j^\mu x_j \quad (4)$$

being η^μ i.i.d. random vectors. Another example is the following. Consider the vector \underline{x} to be unconstrained. Then one wants to minimize the volume of the phase space of the vector \underline{x} , for example by minimizing $|\underline{x}|^2$, such that there is a configuration that satisfies all constraints. This problem is related to the packing problem and has been analyzed in Refs. [2–9], see also Ref. [10] for a review with the corresponding connection to deep learning.

(iii) Nonseparable and nonconvex: We have again a nonseparable cost, which is also nonconvex. As an example, one can think of high-dimensional random Gaussian functions [11]

$$H(\underline{x}) = \frac{1}{N} \sum_{i < j < k} J_{ijk} x_i x_j x_k \quad (5)$$

and the coupling constants are Gaussian variables of the order of $1/N$. In other words, one can superimpose a spin glass landscape on top of a continuous constraint satisfaction problem (CCSP) identified by the set of constraints.

This class of optimization problems is quite vast. In the following, we will consider the simplest case of having a simple separable convex optimization cost and the simplest set of nonconvex excluded volume constraints. We will show that the nonconvexity in the constraints brings about glassiness and marginal stability even when the cost function is convex by itself.

II. A TOY MODEL

To be concrete we consider a simple toy model. We assume that each variable x_i wants to minimize a simple convex cost function but it is constrained to be in a possibly nonconvex region of phase space. We consider the simplest cost function that is

$$H(\underline{x}) = \frac{1}{2} \sum_{i=1}^N (x_i - 1)^2 \quad (6)$$

and we want to minimize it under the constraints that

$$g^\mu(\underline{x}) \geq 0, \quad \forall \mu = 1, \dots, M = \alpha N \quad (7)$$

with

$$g^\mu(\underline{x}) = \frac{1}{\sqrt{N}} \underline{\xi}^\mu \cdot \underline{x} - \sigma, \quad (8)$$

where $\underline{\xi}^\mu = \{\xi_1^\mu, \dots, \xi_N^\mu\}$ is a random vector whose components are Gaussian random variables with zero mean and unit variance. Furthermore, we constrain \underline{x} to be on the N -dimensional sphere $|\underline{x}|^2 = N$. The constraints in Eqs. (7) and

(8) define the perceptron continuous constraint satisfaction problem (CCSP). Equations (6), (7), and (8) define a quadratic programming problem with inequality constraints [12]. However, as we will see, the topology of the phase space and the nature of the constraints are such that the global optimization problem can become nonconvex.

In Refs. [2,3,13–15] the properties of the space of solutions of this CCSP have been extensively studied. As a function of σ , one has a critical value $\alpha_J(\sigma)$ of the density of constraints below which the CCSP is SAT, meaning that with probability one (for $N \rightarrow \infty$) the phase space defined by the constraints has a finite positive volume, and above which the problem is UNSAT, meaning that there is no configuration of \underline{x} that satisfies all constraints.

In Ref. [2] the UNSAT phase of the perceptron has been studied by associating a harmonic energy cost to the unsatisfied constraints. In Ref. [3] the SAT phase of the perceptron CCSP has been characterized. The present setting, instead, is quite different: we still consider the SAT phase of the perceptron CCSP, corresponding to $\alpha < \alpha_J(\sigma)$, but on top of it we want to optimize the cost function $H(\underline{x})$ of Eq. (6). It is therefore a problem of constrained optimization.

The nature of the constraints changes completely when σ passes from being positive, where the CCSP is convex [3], to a negative value, where the solution of the CCSP lies at the intersection of nonconvex domains, which may be in general nonconvex.

In order to study the properties of the landscape of this optimization problem, we consider the partition function defined as

$$Z = \int d\underline{x} \exp[-\beta H(\underline{x})] \left[\prod_{\mu=1}^M \theta(g^\mu(\underline{x})) \right], \quad (9)$$

where $\beta = 1/T$ is the inverse of the temperature T and $\theta(x)$ is the Heaviside step function. Given the partition function, one constructs the average free energy as

$$f = -\frac{1}{\beta N} \overline{\ln Z}, \quad (10)$$

where the overline denotes the average over the random vectors ξ s. To compute the properties of the landscape of the cost function, we need to study the behavior of the free energy in the zero-temperature limit $\beta \rightarrow \infty$.

The average of the logarithm of the partition function can be computed using the replica method [16]

$$f = -\frac{1}{\beta N} \lim_{n \rightarrow 0} \partial_n \overline{Z^n}. \quad (11)$$

The average over the n th power of the partition function can be obtained by considering n to be integer and then taking the analytic continuation down to $n \rightarrow 0$. Performing the integral over the random vectors ξ s (the details are essentially very close to what is reported in Ref. [2] and we will omit them) we obtain

$$\overline{Z^n} = \int d\hat{Q} \exp[NA(\hat{Q})], \quad (12)$$

where $A(\hat{Q})$ is an action and \hat{Q} is an $(n+1) \times (n+1)$ integration matrix. For $N \rightarrow \infty$ we can evaluate the integral by saddle point. The expression of the function $A(\hat{Q})$ is given by

$$A(\hat{Q}) = \frac{1}{2} \ln \det \hat{Q} - \beta(1-m) + \alpha \ln \mathcal{Z}, \quad (13)$$

where we have assumed that $\hat{Q}_{1a} = \hat{Q}_{a1} = m$ independently on a .¹ At the saddle point we have the following interpretation for the overlap matrix \hat{Q} :

$$\begin{aligned} \hat{Q}_{ab} &= \frac{1}{N} \underline{x}_a \cdot \underline{x}_b \quad a, b > 1 \\ \hat{Q}_{1a} &= \hat{Q}_{a1} = \frac{1}{N} \sum_{i=1}^N [x_a]_i = m \quad a > 1. \end{aligned} \quad (14)$$

Therefore $\hat{Q}_{1a} = \hat{Q}_{a1} = m$ is an order parameter that quantifies the overlap between the configuration \underline{x}_a and the minimum $[1, 1, \dots, 1]$ of $H(\underline{x})$, while $\hat{Q}_{ab} = \frac{1}{N} \underline{x}_a \cdot \underline{x}_b$, for $a, b > 1$, is the overlap between the configurations of different replicas of the system. Note that the spherical constraint on \underline{x} imposes that $\hat{Q}_{aa} = 1$. The local partition function \mathcal{Z} is given by

$$\mathcal{Z} = \exp \left[\frac{1}{2} \sum_{a,b=1}^n \hat{Q}_{ab} \frac{\partial^2}{\partial h_a \partial h_b} \right] \prod_{c=1}^n \theta[h_c - \sigma] \Big|_{h=0}, \quad (15)$$

where we have defined \hat{Q} as the matrix obtained from \hat{Q} by removing the first row and column. The saddle point equations for the matrix \hat{Q} can be obtained in full generality following the same steps as in Ref. [2]. Here we will not repeat the computation and we will consider the replica symmetric solution and its stability.

A. Replica symmetric solution and its stability

The replica symmetric ansatz to the solution of the saddle point equations amounts to have $\hat{Q}_{a \neq b} = q$ for all $a \neq b$. The order parameters m and q are then given by the following saddle point equations:

$$\begin{aligned} \frac{q-m^2}{(1-q)^2} &= \alpha \int dh P(q, h) [f'(q, h)]^2 \\ \frac{m}{1-q} &= \beta, \end{aligned} \quad (16)$$

where

$$\begin{aligned} f(q, h) &= \ln \gamma_{1-q} \star \theta(h) \\ P(q, h) &= \gamma_q(h + \sigma) \end{aligned} \quad (17)$$

and we used the notation according to which the prime is the derivative with respect to (w.r.t.) h , \star denotes a convolution operation while

$$\gamma_A(h) = \frac{1}{\sqrt{2\pi A}} e^{-h^2/(2A)}. \quad (18)$$

¹The fact that \hat{Q}_{1a} is independent on the replica index is a property that holds at saddle point. Indeed one can show that at the saddle point $\hat{Q}_{1a} = \frac{1}{N} \sum_{i=1}^N \langle x_i \rangle$, which is therefore independent of a .

The stability condition of the replica symmetric solution is given by studying the Hessian $\frac{\partial^2 A(\hat{Q})}{\partial \hat{Q}_{ab} \partial \hat{Q}_{cd}}$ of the function $A(\hat{Q})$ [2,16,17]. The replica symmetric solution is stable if the saddle point solution satisfies the inequality

$$\frac{1}{(1-q)^2} \geq \alpha \int dh P(q, h) [f''(q, h)]^2. \quad (19)$$

When Eq. (19) becomes a strict equality, the model enters in a glassy regime where multiple minima appear, and the model is characterized by replica symmetry breaking and (possibly) marginal stability [16]. The equations (16)–(19) have been derived in the finite-temperature regime. However, we are interested in looking at the zero-temperature limit, which corresponds to the ground state of the cost function. In order to take the $T \rightarrow 0$ limit, we assume that in this regime

$$q = 1 - \chi T \quad T \rightarrow 0^+, \quad (20)$$

which gives $m = \chi$. Substituting this relation in the first equation of Eqs. (16) we get that χ satisfies

$$\frac{1-\chi^2}{\chi^2} = T^2 \alpha \int dh P(1, h) [f'(1, h)]^2. \quad (21)$$

Furthermore, in the zero-temperature limit, we have

$$f(1, h) \simeq -\frac{h^2}{2\chi T} \theta(-h), \quad (22)$$

which gives

$$1 - \chi^2 = \alpha \int dh P(1, h) h^2 \theta(-h). \quad (23)$$

The marginal stability of the RS solution is given by

$$1 = \alpha \int dh P(1, h) \theta(-h). \quad (24)$$

This relation defines the region in which replica symmetry breaking (RSB) appears. Note that Eq. (20) requires χ to be positive since q is constrained to be $q \leq 1$. The phase diagram of the model is plotted in Fig. 1 where we considered only the case of $\sigma < 0$, which gives rise to nonconvex excluded volume constraints. The red line in the phase diagram represents the replica symmetric approximation for the SAT/UNSAT (jamming) transition line. Above this line, there is no configuration of \underline{x} that satisfies all constraints and therefore the optimization problem makes sense only in the region below this line. This line is simply obtained by taking the $\chi \rightarrow 0$ limit of the saddle point equations. Indeed, above this line no physically meaningful solution should exist. Therefore this line coincides with the one obtained in Ref. [3]. Another way to see that the $\chi \rightarrow 0$ is the right limit to get the SAT/UNSAT line is by noting that at saddle point $\chi = m$. We expect that, on this line, only one solution of the CCSP is left and, by rotational invariance of the disorder, it is orthogonal to the global minimum of the cost function. This implies that $m = 0$.

The dashed green line instead corresponds to the line that separates a strictly stable phase below it, where the ground state of the cost function is unique and not critical, from a marginally stable phase where the landscape of the cost

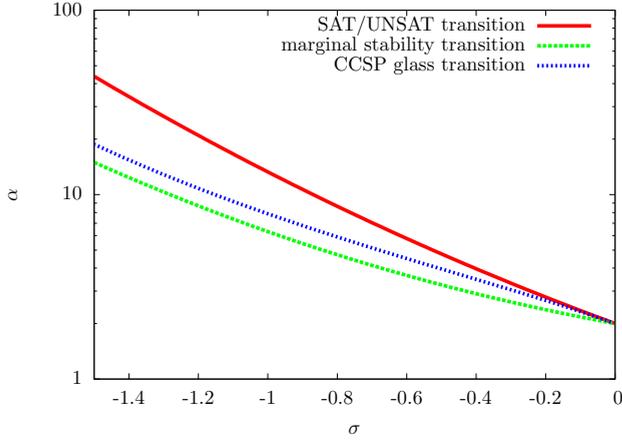


FIG. 1. The phase diagram of the toy optimization problem we are considering. The model is defined only below the SAT/UNSAT transition line for which we plotted in red the replica symmetric approximation. The SAT/UNSAT transition line represents where the volume of configurations allowed by the constraints shrinks to zero. Below the green dashed line, the minimum of the cost function is unique and the replica symmetric solution is stable, meaning that the stability condition Eq. (19) is satisfied. The dashed green line corresponds to Eq. (19) becoming a strict equality. Above this line, instead, the system enters in a glassy phase characterized by multiple minima. This phase is described by replica symmetry breaking. In the replica symmetric phase, the global minimum of the cost function is characterized by a hypostatic number of marginally satisfied constraints. Upon reaching the marginal stability line the number of those constraints becomes equal to the number of degrees of freedom so that the minimum becomes isostatic. We plot in blue also the glass transition line of the perceptron CCSP, whose equation is given in Refs. [2,3] and is independent of the choice of the optimization function $H(\underline{x})$ since it only depends on the perceptron constraints. Below this line, the phase space defined by the constraints is ergodic while right above ergodicity is broken and one has replica symmetry breaking.

function is glassy and local minima are marginally stable. Note that the CCSP defined by the perceptron constraints also undergoes a replica symmetry breaking transition where the space of solutions disconnects before reaching the red line. However, this line, plotted also in Fig. 1, is different from the one signaling the onset of glassiness of the cost function and, in particular, it is closer to the SAT/UNSAT transition line. Therefore we get that even if the CCSP is replica symmetric, the landscape of the cost function may be glassy.

B. Distribution of gaps

To characterize the phase diagram, we compute the distribution of the gaps defined as the values of g^μ computed in the configuration that represents the minimum of the cost function. Using the same approach of Ref. [2], we can show that the distribution of g^μ is given by

$$\rho(h) = \frac{1}{N} \sum_{\mu=1}^M \langle \delta(h - g^\mu(\underline{x})) \rangle = c\delta(h) + \rho^+(h). \quad (25)$$

In the replica symmetric region, we have that $\rho^+(h)$ is given by

$$\rho^+(h) = \alpha P(1, h)\theta(h) \quad (26)$$

and

$$c = \alpha - \alpha \int_0^\infty dh P(1, h) = \alpha \int_{-\infty}^0 dP(1, h). \quad (27)$$

Therefore, when the replica symmetric solution becomes marginally stable, the system becomes isostatic, corresponding to $c \rightarrow 1$, meaning that the number of marginally satisfied constraints equals the number of degrees of freedom in the system. At this point we may ask what happens in the glassy phase, beyond the instability line where the replica symmetric solution breaks down. We do not attempt the RSB solution of the model but we give a simple argument. At the jamming point we know that the allowed phase space of configurations shrinks to a point. This point is jamming critical and isostatic [2]. Therefore we get that, at jamming as well as at the instability transition point, the system is isostatic. We, therefore, conjecture that the RSB phase is again isostatic and jamming critical. This would imply that we have another optimization problem for which the landscape is critical everywhere far from jamming as it happens for other cases [18–20].

III. KARUSH-KUHN-TUCKER DYNAMICS

In order to solve constrained optimization problems, one can rely on the so-called Karush-Kuhn-Tucker (KKT) conditions. We consider a dynamical version of such conditions given by the following dynamical system:

$$\begin{aligned} \dot{x}_i(t) &= -\zeta(t)x_i(t) - [x_i(t) - 1] + \frac{1}{\sqrt{N}} \sum_{\mu=1}^M \xi_i^\mu f_\mu(t) \\ \dot{f}_\mu(t) &= -f_\mu(t)g_\mu(t), \end{aligned} \quad (28)$$

where $\zeta(t)$ is a Lagrange multiplier needed to enforce the spherical constraint on the variables x_i . It is simple to show that a fixed point of these equations is a solution of the optimization problem. Indeed if at long times $g_\mu > 0$, we have $f_\mu = 0$. Instead if we have $g_\mu = 0$ we have $f_\mu > 0$. The variables f_μ , which are the Lagrange multipliers of the KKT conditions, are nothing but the forces that the constraints are exerting in order to fix the corresponding variables g_μ to zero. Therefore this system of equations, when it reaches a fixed point, directly gives the statistics of forces corresponding to the g_μ that are identically equal to zero.

This set of dynamical equations can be analyzed exactly in the $N \rightarrow \infty$ limit through dynamical mean-field theory. We first start by considering the case in which the dynamical system is initialized at random for the x_i while we will fix $f_\mu(0) = 1$, which is a convenient choice. Then, to analyze the dynamics in the $N \rightarrow \infty$ limit, we consider the Martin-Siggia-Rose-Jenssen-De-Dominicis dynamical path integral [11] defined as

$$Z_{\text{dyn}} = \int \mathcal{D}(\underline{f}, \underline{\hat{f}}) \mathcal{D}(\underline{x}, \underline{\hat{x}}) \mathcal{D}(\underline{t}, \underline{\hat{t}}) \exp[S_{\text{dyn}}], \quad (29)$$

where the dynamical action is given by

$$S_{\text{dyn}} = i \sum_{\mu=1}^M \int dt \{ \hat{f}_{\mu}(t) [\dot{f}_{\mu}(t) + f_{\mu}(t)(r_{\mu}(t) - \sigma)] + \hat{r}_{\mu}(t) r_{\mu}(t) \} \\ + i \sum_{i=1}^N \int dt \hat{x}_i(t) \left[\dot{x}_i + \zeta(t)x_i(t) + (x_i(t) - 1) - \frac{1}{\sqrt{N}} \sum_{\mu=1}^M \xi_i^{\mu} f_{\mu}(t) \right] - \frac{i}{\sqrt{N}} \sum_{i=1}^N \int dt \sum_{\mu=1}^M \hat{r}_{\mu}(t) \xi_i^{\mu} x_i(t) \quad (30)$$

and the overline stands for the average over the random patterns. Taking the average over these vectors gives

$$\overline{\exp \left[-\frac{i}{\sqrt{N}} \sum_{\mu=1}^M \sum_{i=1}^N \xi_i^{\mu} \int dt [x_i(t) f_{\mu}(t) + \hat{r}_{\mu}(t) x_i(t)] \right]} \\ = \exp \left[-\frac{1}{2} \sum_{\mu=1}^M \int dt dt' [f_{\mu}(t) f_{\mu}(t') D(t, t') \right. \\ \left. + \hat{r}_{\mu}(t) \hat{r}_{\mu}(t') C(t, t') + 2i \hat{r}_{\mu}(t) f_{\mu}(t') R(t, t') \right], \quad (31)$$

where we have introduced the dynamical order parameters

$$C(t, t') = \frac{1}{N} \underline{x}(t) \cdot \underline{x}(t') \\ D(t, t') = \frac{1}{N} \hat{x}_i(t) \cdot \hat{x}_i(t') \\ R(t, t') = -\frac{1}{N} \underline{x}(t) \cdot i \hat{x}_i(t'). \quad (32)$$

As usual, at the saddle point level $D(t, t') = 0$ by causality. The single gap effective process then gives the equation

$$r(t) = \int_0^t dt' R(t, t') f(t') + \eta(t), \quad (33)$$

where the statistics of the noise $\eta(t)$ is Gaussian and given by

$$\overline{\eta(t)} = 0 \quad \overline{\eta(t)\eta(t')} = C(t, t'). \quad (34)$$

Finally we have

$$f(t) = \exp \left[-\int_0^t ds (r(s) - \sigma) \right]. \quad (35)$$

Note that Eq. (33) must be understood as a distributional equation. Furthermore, we have used that at the initial time $t = 0$ forces are initialized to $f_{\mu}(0) = 1$. This initialization does not affect the endpoint of the dynamics and provides a way to enforce that all forces are either zero or positive. The single-site effective process is instead

$$\dot{x}(t) = -\zeta(t)x(t) - (x(t) - 1) + \Xi(t) - \int_0^t ds K(t, s)x(s), \quad (36)$$

where the statistics of the Gaussian noise $\Xi(t)$ is given by

$$\overline{\Xi(t)} = 0 \quad \overline{\Xi(t)\Xi(t')} = \alpha \langle f(t)f(t') \rangle = M(t, t') \quad (37)$$

and the kernel $K(t, s)$ is defined as

$$K(t, s) = \alpha \left. \frac{\delta \langle f(t) \rangle_{\eta}}{\delta P(s)} \right|_{P=0}, \quad (38)$$

where the perturbation $P(t)$ is an additive perturbation on the right-hand side of Eq. (33). Therefore we obtain the following equations for correlation, response, and magnetization $m(t) = \sum_i x_i(t)/N$

$$\partial_t m(t) = -\zeta(t)m(t) - [m(t) - 1] \\ - \int_0^t ds K(t, s)m(s) \\ \partial_t C(t, t') = -\zeta(t)C(t, t') - [C(t, t') - m(t')] \\ - \int_0^t ds K(t, s)C(t', s) + \int_0^{t'} ds M(t, s)R(t', s) \\ \partial_t R(t, t') = -\zeta(t)R(t, t') - R(t, t') + \delta(t - t') \\ - \int_{t'}^t ds K(t, s)R(s, t'). \quad (39)$$

Note that the Lagrange multiplier $\zeta(t)$ can be obtained directly from the equation for $C(t, t')$.

From these dynamical equations it follows that

$$m(t) = \int_0^t ds R(t, s), \quad (40)$$

which is a relation connecting the response function $R(t, s)$ to the magnetization $m(t)$.

Replica symmetric solution

Let us consider what happens in the region where replica symmetry is unbroken. If we define

$$\chi = \lim_{t \rightarrow \infty} \int_0^t ds R(t, s) \quad (41)$$

we immediately have

$$m_{\infty} \equiv \lim_{t \rightarrow \infty} m(t) = \chi. \quad (42)$$

Considering Eq. (33) and taking the $t \rightarrow \infty$ limit we get that

$$r_{\infty} = \chi f_{\infty} + \eta \quad (43)$$

being η a normal Gaussian random variable. This equation tells us that if $f_{\infty} = 0$ then r_{∞} is a Gaussian random variable constrained to be such that $r_{\infty} - \sigma$ is positive. Therefore the gap distribution is

$$P(h) = \gamma_1(h + \sigma). \quad (44)$$

On the other hand, if $r_{\infty} = \sigma$ then $f_{\infty} > 0$ and its distribution coincides with the one of the random variable $(\sigma - \eta)/\chi$ with the constraint that $\sigma - \eta > 0$. Finally we need to establish an equation for χ and show that it coincides with the one coming

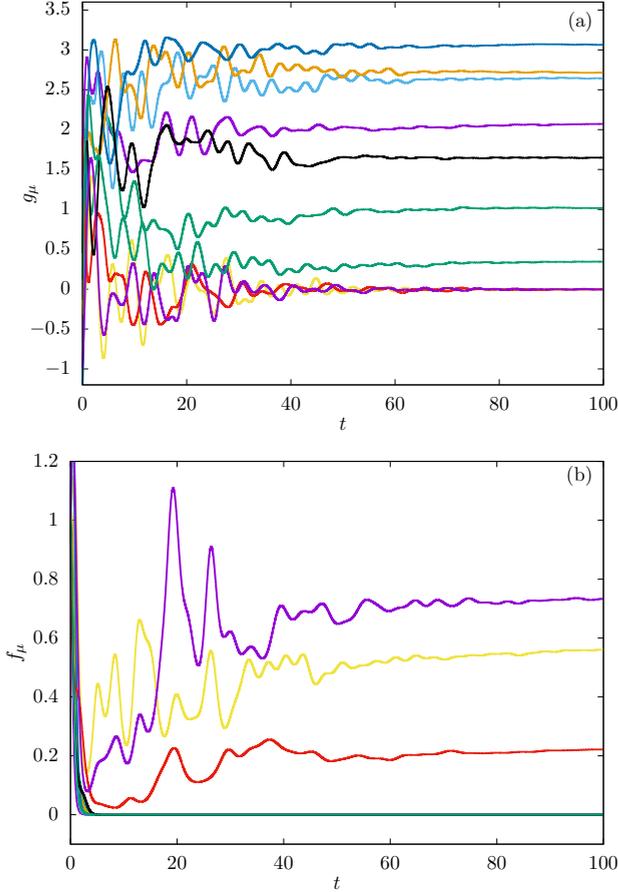


FIG. 2. Time evolution of (a) g_μ and (b) f_μ in the RS phase ($\alpha = 5$, $\sigma = -1$ in this example), each color corresponding to a different $\mu = 1, \dots, 10$, obtained by numerically integrating Eq. (28) with an explicit Runge-Kutta method of order 5 [21]. In the RS phase, the system quickly converges to the steady state. We notice that in the steady state of the RS phase, for a given μ , $g_\mu > 0$, and $f_\mu = 0$ or $g_\mu = 0$ and $f_\mu > 0$. In this experiment, the f_μ have been initialized to 1 while the x_i have been initialized as standard Gaussian variables. We used $N = 100$.

from the replica approach. We consider the long time limit of the equation for $m(t)$, which gives

$$0 = -\zeta_\infty \chi - (\chi - 1) - K_\infty \chi, \quad (45)$$

where

$$K_\infty = \lim_{t \rightarrow \infty} \int_0^\infty ds K(t, s). \quad (46)$$

On the other hand the long time limit of the equation for M is obtained by considering the long time limit of the equation for C . This gives

$$0 = -\zeta_\infty + 1 - \chi - K_\infty + M_\infty \chi, \quad (47)$$

which can be combined with Eq. (45) to get

$$M_\infty \chi^2 = 1 - \chi^2. \quad (48)$$

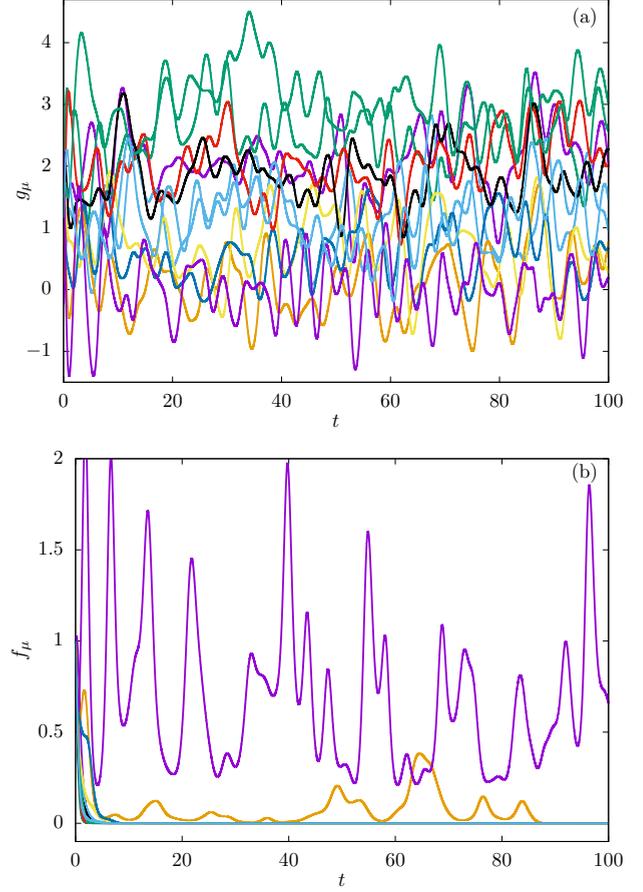


FIG. 3. Time evolution of (a) g_μ and (b) f_μ in the RSB phase ($\alpha = 8$, $\sigma = -1$ in this example), each color corresponding to a different $\mu = 1, \dots, 10$, obtained by numerically integrating Eq. (28) with an explicit Runge-Kutta method of order 5 [21]. Since the RSB phase is marginally stable, the time evolution is chaotic and does not converge to a stationary state. In this experiment, the f_μ have been initialized to 1 while the x_i have been initialized as standard Gaussian variables. We used $N = 100$.

From the definition of M_∞ we finally have

$$M_\infty = \alpha \langle f_\infty^2 \rangle = \alpha \int \frac{d\eta}{\sqrt{2\pi}} e^{-\eta^2/2} \frac{(\eta + \sigma)^2}{\chi^2} \theta(\eta + \sigma) \quad (49)$$

and therefore using Eqs. (48) and (49) we get Eq. (16). This concludes the derivation of the replica symmetric equation from the KKT dynamics.

We implemented numerically the algorithm of Eq. (28) and we found that in the RS phase the algorithm goes very quickly to the solution of the optimization problem, as it can be observed in Fig. 2. In the RSB phase, instead, it seems that the algorithm does not converge and therefore we believe that in this region it makes a chaotic surf on the isostatic landscape (Fig. 3).

IV. CONCLUSIONS

We have considered the generic setting of optimization problems under nonconvex excluded volume constraints. We

have analyzed a simple example of this kind in which the cost function is separable and convex and in which the nonconvex excluded volume constraints are modeled by a negative perceptron.

We have found that, when the number of constraints is low enough, the cost function has a simple ground state, which is captured by the replica symmetric solution of the model. At high density of constraints, instead, the optimization landscape undergoes an RSB transition where minima become marginally stable. Remarkably, we find that the RSB transition point happens before the point in which the accessible phase space defined by the constraints splits into glassy

regions. We have also shown how to recover these results from the dynamical mean-field theory of the KKT algorithm. We leave the analysis of more complex cost functions for future work.

ACKNOWLEDGMENTS

This work was supported by “Investissements d’Avenir” LabExPALM (ANR-10-LABX-0039-PALM). A.S. thanks LPTMS where part of this work has been done and the Simons Foundation Grant (No. 454941 Silvio Franz and No. 454953 Matthieu Wyart) for support.

-
- [1] B. Neyshabur, in *Advances in Neural Information Processing Systems*, edited by H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (NeurIPS, San Diego, 2020), Vol. 33, pp. 8078–8088.
 - [2] S. Franz, G. Parisi, M. Sevelev, P. Urbani, and F. Zamponi, *SciPost Phys.* **2**, 019 (2017).
 - [3] S. Franz and G. Parisi, *J. Phys. A: Math. Theor.* **49**, 145001 (2016).
 - [4] S. Franz, G. Parisi, P. Urbani, and F. Zamponi, *Proc. Natl. Acad. Sci. USA* **112**, 14539 (2015).
 - [5] H. Yoshino, *SciPost Phys* **4**, 040 (2018).
 - [6] C. Brito, H. Ikeda, P. Urbani, M. Wyart, and F. Zamponi, *Proc. Natl. Acad. Sci. USA* **115**, 11736 (2018).
 - [7] S. Franz, S. Hwang, and P. Urbani, *Phys. Rev. Lett.* **123**, 160602 (2019).
 - [8] H. Ikeda, P. Urbani, and F. Zamponi, *J. Phys. A: Math. Theor.* **52**, 344001 (2019).
 - [9] H. Ikeda, *Phys. Rev. Research* **2**, 033220 (2020).
 - [10] M. Geiger, L. Petrini, and M. Wyart, *Phys. Rep.* **924**, 1 (2021).
 - [11] T. Castellani and A. Cavagna, *J. Stat. Mech.* (2005) **P05012**.
 - [12] S. Boyd and L. Vandenberghe, *Convex Optimization* (Cambridge University Press, Cambridge, 2004).
 - [13] E. Gardner, *J. Phys. A: Math. Gen.* **21**, 257 (1988).
 - [14] E. Gardner and B. Derrida, *J. Phys. A: Math. Gen.* **21**, 271 (1988).
 - [15] M. Stojnic, [arXiv:1306.3980](https://arxiv.org/abs/1306.3980).
 - [16] M. Mézard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987).
 - [17] J. De Almeida and D. J. Thouless, *J. Phys. A: Math. Gen.* **11**, 983 (1978).
 - [18] S. Franz, A. Sclocchi, and P. Urbani, *Phys. Rev. Lett.* **123**, 115702 (2019).
 - [19] S. Franz, A. Sclocchi, and P. Urbani, *SciPost Phys.* **9**, 012 (2020).
 - [20] A. Sclocchi and P. Urbani, *SciPost Phys.* **10**, 13 (2021).
 - [21] J. R. Dormand and P. J. Prince, *J. Comput. Appl. Math.* **6**, 19 (1980).