

Gradient flows and proximal splitting methods: A unified view on accelerated and stochastic optimization

Guilherme França,^{1,2,*} Daniel P. Robinson,³ and René Vidal²

¹University of California, Berkeley, California 94720, USA

²Johns Hopkins University, Baltimore, Maryland 21218, USA

³Lehigh University, Bethlehem, Pennsylvania 18015, USA

 (Received 15 January 2021; revised 21 April 2021; accepted 22 April 2021; published 10 May 2021)

Optimization is at the heart of machine learning, statistics, and many applied scientific disciplines. It also has a long history in physics, ranging from the minimal action principle to finding ground states of disordered systems such as spin glasses. Proximal algorithms form a class of methods that are broadly applicable and are particularly well-suited to nonsmooth, constrained, large-scale, and distributed optimization problems. There are essentially five proximal algorithms currently known, each proposed in seminal work: Forward-backward splitting, Tseng splitting, Douglas-Rachford, alternating direction method of multipliers, and the more recent Davis-Yin. These methods sit on a higher level of abstraction compared to gradient-based ones, with deep roots in nonlinear functional analysis. In this paper we show that all of these methods are actually different discretizations of a single differential equation, namely, the simple *gradient flow* which dates back to Cauchy (1847). An important aspect behind many of the success stories in machine learning relies on “accelerating” the convergence of first-order methods. However, accelerated methods are notoriously difficult to analyze, counterintuitive, and without an underlying guiding principle. We show that similar discretization schemes applied to Newton’s equation with an additional dissipative force, which we refer to as *accelerated gradient flow*, allow us to obtain accelerated variants of all these proximal algorithms—the majority of which are new although some recover known cases in the literature. Furthermore, we extend these methods to stochastic settings, allowing us to make connections with Langevin and Fokker-Planck equations. Similar ideas apply to gradient descent, heavy ball, and Nesterov’s method which are simpler. Our results therefore provide a unified framework from which several important optimization methods are nothing but *simulations* of classical dissipative systems.

DOI: [10.1103/PhysRevE.103.053304](https://doi.org/10.1103/PhysRevE.103.053304)

I. INTRODUCTION

The simplest algorithm to solve a *smooth* optimization problem,

$$\min_{x \in \mathbb{R}^n} \varphi(x), \quad (1)$$

dates back to Cauchy [1]. It is the well-known *gradient descent* method, $x_{k+1} = x_k - h \nabla \varphi(x_k)$, where $h > 0$ is the step size and $k = 0, 1 \dots$ is the iteration number. Clearly, gradient descent corresponds to an *explicit* Euler discretization of the *gradient flow*:

$$\dot{x} = -\nabla \varphi(x), \quad (2)$$

where $x \equiv x(t)$. To minimize *nonsmooth* and *composite* functions, a series of milestone papers introduced algorithms based on *proximal operators*, which do not require explicit gradient computations. For instance, the *Douglas-Rachford* algorithm [2] was proposed in the 1950s to solve the heat equation but nowadays is a standard optimization method with important applications. Closely related is the *alternating direction method of multipliers* (ADMM), introduced independently by Glowinsky and Marrocco [3] and Gabay and

Mercier [4] in the 1970s—ADMM has been gaining significant interest in machine learning during the last decade [5]. Another method—that plays an important role in signal processing—is the *forward-backward splitting* introduced by Lions and Mercier [6] also in the 1970s. These were the only known proximal-based methods for almost 30 years, until Tseng [7] proposed a slight modification of the latter known as *forward-backward-forward splitting*. Such methods are designed to minimize composite functions $\varphi(x) = \varphi_1(x) + \varphi_2(x)$, where both φ_1, φ_2 are allowed to be nonsmooth. Finding an algorithm that minimizes $\varphi(x) = \varphi_1(x) + \varphi_2(x) + \varphi_3(x)$, where only φ_3 is assumed to be smooth, and which cannot be reduced to any of the previously known methods, was a longstanding problem that has been recently solved by Davis and Yin [8]. These five algorithms compose the list of fundamental proximal algorithms currently known—many other methods are variations of these basic themes. Such proximal methods can be derived from operator splitting techniques [9,10], which have origins in the works of Browder [11–13] and Minty [14], although nowadays form an entire field of research in convex analysis, optimization, and nonlinear functional analysis [15,16].

Perhaps surprisingly, we provide a simple yet unified perspective on these distinct methods: All of them are different discretizations of the simple gradient flow (2). More precisely,

*guifranca@gmail.com

they are first order integrators that preserve the critical points of this ODE.

“Acceleration strategies” in the context of optimization have proved to be powerful and are behind many of the empirical success stories in machine learning, such as in training neural networks. The basis of accelerated *gradient-based* methods can be traced back to Polyak [17] and Nesterov [18]—both can be seen as accelerated versions of gradient descent. Although neither are intuitive in their precise design, it has recently been shown that both can be obtained as *explicit* discretizations of a second-order ODE [19–21]. This continuous-time perspective on optimization methods is quite recent and has helped demystify the “magic” of acceleration techniques. However, the construction and analysis of accelerated methods is still obscure, without an underlying guiding principle—e.g., it is not clear how to “accelerate” a given algorithm. *Accelerated proximal-based* methods are even less known, and can play an important role since they may enjoy improved stability and be applicable in more general situations. Moreover, from a mathematical standpoint, such methods have interesting connections with nonlinear functional analysis. As we will show, several accelerated variants of each of the above mentioned proximal algorithms can be obtained as different discretizations of

$$\ddot{x} + \eta(t)\dot{x} = -\nabla\varphi(x), \quad (3)$$

with a suitable choice of damping function $\eta(t) > 0$. The resulting methods—most of which are new in the literature—are first-order integrators that preserve critical points of this ODE. Therefore, the above *classical dissipative system* has a fundamental importance to optimization. Note that this is nothing but Newton’s equation with an additional dissipative force, $-\eta(t)\dot{x}$. When $\eta(t) = \eta$ is constant and $\varphi = \omega^2 x^2/2$ it reduces to the Caldirola-Kanai oscillator [22,23], which is the classical limit of the seminal Caldeira-Leggett model [24].

Our approach makes connections between optimization and *splitting methods* for ODEs [25]. Interestingly, ADMM and its accelerated variants arise as a *rebalanced splitting*, which is a recent technique designed to preserve critical points [26]—the so-called dual variable of ADMM, originally introduced as a Lagrange multiplier, is precisely the balance coefficient of Ref. [26]. The other methods we consider also preserve critical points, but for different reasons, which in turn sheds light on the connections between ODE splitting and operator splitting ideas from convex analysis.

Stochastic optimization is an important ingredient in the machine learning toolbox to reduce the computational burden in training high-dimensional models over large datasets. By introducing stochastic gradients or stochastic proximal operators into these methods, instead of Eqs. (2) and (3), their continuum limit become an *overdamped* or *underdamped* Langevin equation, respectively. The probability distribution of such stochastic processes is described by a Fokker-Planck equation. Therefore, there is a close connection between deterministic optimization and dissipative classical mechanics, as well as stochastic optimization and nonequilibrium statistical mechanics.

This paper is organized as follows. In Sec. II, we introduce basic concepts related to proximal operators—or monotone operators and their regularizations more generally—and

illustrate how they naturally arise from *implicit* discretizations of ODEs. In Sec. III, we show relevant details about the dynamics of the gradient flows Eqs. (2) and (3). In Sec. IV, we introduce a slight variation of the balanced and rebalanced splitting schemes [26] to then show how—accelerated—ADMM arises from this approach. In Sec. V, we derive extensions of Davis-Yin, which is known to generalize both forward-backward and Douglas-Rachford. In Sec. VI, we introduce accelerated extensions of Tseng’s splitting to complete the list. The focus of this paper is on discretizations of Eq. (3) since this allows us to construct entire new families of accelerated methods that generalize the existing ones. Moreover, the known methods that are related to Eq. (2) follow as particular cases—more precisely, through a high friction limit. At this stage, we briefly summarize and interpret these results from a physics perspective in Sec. VII. Then, in Sec. VIII, we shift gears and show how one can extend—quite easily—these proximal-based methods to stochastic optimization settings. As a consequence, the connections with the continuous-time formalism are promoted to SDEs of the Langevin type, whose probability distribution are governed by Fokker-Planck equations. For completeness, in the Appendix we show that stochastic gradient descent, heavy ball, and Nesterov fit our general framework—such gradient-based methods find widespread applications in machine learning but they are actually simpler. We provide numerical experiments in Sec. IX that support our theoretical conclusions, and also illustrate the faster convergence attained by the accelerated methods. Our final remarks and potential implications of our analysis are presented in Sec. X.

II. RESOLVENT, YOSIDA REGULARIZATION, AND PROXIMAL OPERATOR

We start by introducing fundamental concepts from nonlinear functional analysis [15,16] since this is the formalism in which proximal algorithms are generally discussed. We avoid excessive formalism in the paper, but here we give a roadmap to further abstract our analysis.

The *resolvent* of an operator A can be defined as

$$J_{\lambda A} \equiv (I + \lambda A)^{-1}, \quad (4)$$

where λ is the so-called spectral parameter. Even though λ can be complex, we will only need $\lambda \in \mathbb{R}$. Another useful concept is the *Yosida regularization* of A :

$$A_{\lambda} \equiv \lambda^{-1}(I - J_{\lambda A}). \quad (5)$$

Let H be a Hilbert space with inner product $\langle \cdot | \cdot \rangle : H \times H \rightarrow \mathbb{C}$. A *multivalued* map $A : H \rightrightarrows H$, with $\text{dom}A \equiv \{x \in H \mid Ax \neq \emptyset\}$, is said to be *monotone* if and only if

$$\langle Ay - Ax | y - x \rangle \geq 0 \quad \text{for all } x, y \in \text{dom}A. \quad (6)$$

A monotone operator is said to be *maximal* if no enlargement of its graph is possible. It turns out that every monotone operator admits a maximal extension, thus we henceforth assume that all operators are maximal monotone. What matters for us is that in this case the resolvent is single-valued, i.e., $J_{\lambda A} : H \rightarrow H$ is a function. Moreover, x^* is a zero of A , namely, $x^* \in \text{zer}(A) \equiv \{x \in H \mid 0 \in Ax\}$, if and only if $J_{\lambda A}(x^*) = x^*$. Thus, zeros of A are fixed points of the resolvent, $J_{\lambda A}$, and vice

versa. Consequently, the Yosida regularization is also single-valued, and x^* is a zero of A if and only if $A_\lambda x^* = 0$, so that zeros of A are also zeros of A_λ and vice versa. The advantage of working with the Yosida regularization is that it allows us to deal with multivalued operators by considering single-valued operators. Indeed, it can be shown that $A_\lambda x \rightarrow A_0 x$ as $\lambda \downarrow 0$, where $A_0 x$ is the element of minimal norm in the set Ax .

These ideas can be made more intuitive by considering a function $\varphi : \mathbb{R}^n \rightarrow \mathbb{R}$, which for the moment we assume to be differentiable. The function φ is convex if and only if its gradient, $A = \nabla\varphi$, is (maximal) monotone. In this case, the resolvent (4) becomes the so-called *proximal operator*:¹

$$J_{\lambda\partial\varphi}(x) \equiv \text{prox}_{\lambda\varphi}(x) \equiv \arg \min_y \left(\varphi(y) + \frac{1}{2\lambda} \|y - x\|^2 \right). \quad (7)$$

Then (5) becomes $(\nabla\varphi)_\lambda(x) = \lambda^{-1}[x - \text{prox}_{\lambda\varphi}(x)]$, which is the gradient of the *Moreau envelope* φ_λ , i.e., $(\nabla\varphi)_\lambda(x) = \nabla\varphi_\lambda(x)$, where

$$\varphi_\lambda(x) \equiv \min_y \left(\varphi(y) + \frac{1}{2\lambda} \|y - x\|^2 \right). \quad (8)$$

When φ is *nonsmooth* its gradient $\nabla\varphi$ is ill-defined. However, there exists a generalization which is the notion of subdifferential set; it is defined as $\partial\varphi(x) \equiv \{g \in \mathbb{R}^n \mid \varphi(y) - \varphi(x) \geq \langle g, y - x \rangle \text{ for all } y \in \mathbb{R}^n\}$. If φ is differentiable, then $\partial\varphi(x) = \{\nabla\varphi(x)\}$. We thus see that even though φ may not be differentiable, its Moreau envelope always is, and we can thus treat the problem with standard calculus on φ_λ . We have $\lim_{\lambda \downarrow 0} \nabla\varphi_\lambda(x) \in \partial\varphi(x)$, and this limit is the vector of minimal norm in the subdifferential set $\partial\varphi(x)$.

More generally, our results in this paper show that all the previously mentioned algorithms correspond to discretizations of the *differential inclusion* [16],

$$\dot{x} \in -Ax, \quad (9)$$

for a monotone operator $A : H \rightrightarrows H$ that is composite, $A = A_1 + A_2 + A_3$. Similarly, the accelerated variants of these algorithms are related to the second-order differential inclusion,

$$\ddot{x} + \eta(t)\dot{x} \in -Ax. \quad (10)$$

However, dealing with differential inclusions, i.e., nonsmooth dynamical systems, involve several technicalities. A simple way to avoid the issue is to focus on their Yosida regularizations, namely,

$$\dot{x} = -A_\mu x \quad (11)$$

and

$$\ddot{x} + \eta(t)\dot{x} = -A_\mu x, \quad (12)$$

respectively, which are well-posed ODEs. (Note that $x(t)$ depends on $\mu > 0$, which we omit in the notation.) At the end

¹This can be easily seen as follows. The solution y of Eq. (7) obeys $\nabla\varphi(y) + (1/\lambda)(y - x) = 0$, i.e., $(I + \lambda\nabla\varphi)y = x$, which from Eq. (4) gives $y = J_{\lambda\nabla\varphi}x$. We replaced $\nabla\varphi \mapsto \partial\varphi$ in Eq. (7) anticipating generalization to nonsmooth cases.

of the day one can take the limit $\mu \downarrow 0$ to recover results for Eqs. (9) and (10) [16]. In the context of nonsmooth optimization this means considering the gradient of the Moreau envelope, $\nabla\varphi_\mu$, instead of the subdifferential, $\partial\varphi$ —this point will be further clarified below.

As a warmup, and also to introduce the basic building blocks of our approach, let us show a simple example on how to derive a proximal algorithm from (9), or equivalently (11). Consider an *implicit* discretization of Eq. (9):

$$\frac{x_{k+1} - x_k}{h} + \mathcal{O}(h) \in -Ax_{k+1}. \quad (13)$$

Using the resolvent (4) and neglecting $\mathcal{O}(h^2)$ terms, we can solve this recurrence as

$$x_{k+1} = J_{hA}x_k. \quad (14)$$

This algorithm finds zeros of the monotone operator A . For a nonsmooth function φ , we set $A = \partial\varphi$ to obtain

$$x_{k+1} = \text{prox}_{h\varphi}(x_k). \quad (15)$$

This is the well-known *proximal gradient* method, extensively studied in convex analysis and optimization literatures. Now, consider instead an analogous discretization of the regularized ODE (11):

$$\frac{x_{k+1} - x_k}{h} + \mathcal{O}(h) = -A_\mu x_{k+1}. \quad (16)$$

From the resolvent (4) we get

$$x_{k+1} = J_{hA_\mu}x_k. \quad (17)$$

Employing the useful formula [15]

$$J_{\lambda A_\mu} = (\lambda + \mu)^{-1}(\mu I + \lambda J_{(\lambda + \mu)A}), \quad (18)$$

we conclude that $J_{\lambda A_\mu} \rightarrow J_{\lambda A}$ as $\mu \downarrow 0$. Thus, Eq. (17) becomes precisely Eq. (14) in the limit $\mu \downarrow 0$. In the case of a nonsmooth function φ , the RHS of (11) is simply $-\nabla\varphi_\mu$, i.e., the gradient of Moreau envelope, whereas the RHS of (9) is $-\partial\varphi$, i.e., the subdifferential set. Hence, we could have considered a discretization of the gradient flow (2) with $\varphi \mapsto \varphi_\mu$ and then take the limit $\mu \downarrow 0$. Even simpler, we could have actually ignored nonsmoothness issues altogether and simply discretized Eq. (2), replacing $\nabla\varphi \mapsto \partial\varphi$ where appropriate—this results in updates in terms of the proximal operator (7) and the procedure can be formally justified by the above steps.

Next, let us consider a similar discretization approach but for the second-order ODE (12). The differential operator on the LHS can be discretized as

$$\frac{x_{k+1} - 2x_k + x_{k-1}}{h^2} + \eta_k \frac{x_k - x_{k-1}}{h} + \mathcal{O}(h). \quad (19)$$

Defining

$$\hat{x}_k \equiv x_k + \gamma_k(x_k - x_{k-1}), \quad \gamma_k \equiv (1 - h\eta_k), \quad (20)$$

where $\eta_k \equiv \eta(t_k)$ is the discretized damping coefficient, assumed to be only a function of time, we obtain

$$\ddot{\hat{x}}(t_k) + \eta(t_k)\dot{\hat{x}}(t_k) = \frac{x_{k+1} - \hat{x}_k}{h^2} + \mathcal{O}(h). \quad (21)$$

This relation will prove extremely convenient in pretty much all discretizations considered in this paper. Note that it allows us to discretize the second-order system (12) in very similar

way as the first-order system (11)—essentially, it suffices to replace $x_k \mapsto \hat{x}_k$ and $h \mapsto h^2$. Thus, an implicit discretization of Eq. (12) yields $(x_{k+1} - \hat{x}_k)/h^2 = -A_\mu x_{k+1}$, which can be readily solved with the resolvent to obtain

$$x_{k+1} = J_{h^2 A_\mu} \hat{x}_k. \quad (22)$$

Taking $\mu \downarrow 0$ yields $x_{k+1} = J_{h^2 A} \hat{x}_k$. For a nonsmooth function φ , we set $A = \partial\varphi$ and replace the resolvent by the proximal operator, hence obtaining

$$x_{k+1} = \text{prox}_{h\varphi}(\hat{x}_k), \quad (23a)$$

$$\hat{x}_{k+1} = x_{k+1} + \gamma_{k+1}(x_{k+1} - x_k). \quad (23b)$$

Note that in this last passage we redefined the step size, $h^2 \mapsto h$, and this should be also reflected in (20). The above algorithm corresponds to an “accelerated version” of the proximal gradient method (15)—in the next section we will show that this second-order system may indeed have faster convergence compared to the first-order gradient flow. Note also that the nonaccelerated method (15) can be recovered from (23) by setting $\gamma_k = 0$. Physically, this corresponds to a “large friction limit” as we will explain in more detail shortly.

Already at this stage we have several new methods encoded in Eq. (23) due to the possibility of choosing different damping functions $\eta(t)$. Reasonable choices (that will be justified in the next section) are a *constant damping*,

$$\eta(t) = \eta \quad \Rightarrow \quad \gamma_k = 1 - \sqrt{h} \eta, \quad (24)$$

which is originally related to Polyak’s heavy ball method [17], and a *decaying damping*,

$$\eta(t) = r/t \quad (r \geq 3) \quad \Rightarrow \quad \gamma_k = k/(k+r), \quad (25)$$

where usually $r = 3$ and this is originally related to Nesterov’s method [18,19].² However, other choices are possible and we allow an arbitrary $\eta(t)$ in our discretizations.

A. Large friction limit

We noted that setting $\gamma_k = 0$ into Eq. (23) yields the proximal gradient method (15), which is a discretization of the first-order gradient flow (2). We also said that this corresponds to a “large friction limit” of the second-order gradient flow (3). Since this idea will reappear later on, we provide more details.

First, from a discrete-time viewpoint, with $\gamma_k = 0$ we have $\hat{x}_k = x_k$. Dropping k for simplicity, note that the damping becomes $\eta = (1 - \gamma)/h \rightarrow 1/h$ as $\gamma \rightarrow 0$, i.e., $\eta \rightarrow \infty$ as $h \rightarrow 0$. Therefore, this corresponds indeed to an “overdamped limit.” Moreover, replacing $\eta \rightarrow 1/h$ into Eq. (21) yields

$$h\ddot{x} + \dot{x} = \frac{x_{k+1} - x_k}{h} + \mathcal{O}(h^2). \quad (26)$$

Note that \ddot{x} becomes negligible for sufficiently small step sizes, i.e., the above reduces to $\dot{x} = (x_{k+1} - x_k)/h + \mathcal{O}(h)$,

which is precisely a discretization of the LHS of Eq. (2)—or of Eq. (11) more generally.

Second, from a continuous-time viewpoint, note that Eq. (12) is in natural units where the mass $m = 1$. Restoring the mass we have

$$m\ddot{x} + m\eta\dot{x} = -A_\mu x. \quad (27)$$

The overdamped limit can be obtained with $\eta \rightarrow 1/m$ and $m \rightarrow 0$, in which case we recover Eq. (11). It is in this sense that the gradient flow is a large friction limit—or a zero-mass limit—of the accelerated gradient flow. Intuitively, a very light particle has negligible acceleration due a large frictional force, and the first-order system (11) approximates the dynamics of the second-order system (12) when η is sufficiently large. This same idea often appears in the theory of stochastic processes where a second-order Langevin equation is well-approximated by a first-order Langevin equation when the “fluid” representing the heat bath is highly viscous.

B. A note on nonsmoothness

Above, we discretized the regularized ODEs (11) and (12) and then took the limit $\mu \downarrow 0$ to reduce the fixed point iterations to the case of monotone operators. By choosing $A = \partial\varphi$ these algorithms are appropriate for minimizing a nonsmooth function φ , through the proximal operator (7). Although we were careful in taking “nonsmoothness” into account, apart from this μ limit, the discretization procedure is exactly the same as if we had considered Eqs. (2) and (3) for which φ is assumed to be differentiable. In other words, everything works fine if we replace $\nabla\varphi \mapsto \partial\varphi$ where appropriate—this point was also noted in the discussion after Eq. (18). Moreover, even when we split the operators, as we will do in the following, it is still possible to introduce some parameter μ that justifies the procedure. Therefore, to avoid unnecessary formalism, hereafter we assume that φ is differentiable for all practical purposes—one should keep in mind that this assumption can be removed by introducing Yosida regularizations or Moreau envelopes and taking the $\mu \downarrow 0$ limit.

III. CONTINUUM DYNAMICS

Here we provide some details about the dynamics of the gradient flows (2) and (3). Note that the first-order system (2) yields the simplest dynamics that follows a descent direction on φ , thus it is naturally suited for optimization purposes. The second-order system (3) corresponds to its accelerated version in a classical mechanical sense, and also follows a descent direction but can oscillate. This is an actual dissipative system with Lagrangian

$$\mathcal{L} = \frac{1}{2} e^{\theta(t)} \|\dot{x}\|^2 - e^{\theta(t)} \varphi(x), \quad (28)$$

where $\dot{\theta}(t) \equiv \eta(t)$, or equivalently with the explicit time-dependent Hamiltonian

$$\mathcal{H} = \frac{1}{2} e^{-\theta(t)} \|p\|^2 + e^{\theta(t)} \varphi(x). \quad (29)$$

The physical energy is given by $\mathcal{E} = \frac{1}{2} \|\dot{x}\|^2 + \varphi(x)$ and dissipates at a rate $\dot{\mathcal{E}} = -\eta(t) \|\dot{x}\|^2 \leq 0$, i.e., it decreases

²The particular choice of γ_k in Eq. (25) is to maintain consistency with the optimization literature but $\gamma_k = k/(k+r) \approx 1 - r/k = 1 - h\eta_k$ for large k which can equally be used.

monotonically with time, thus forcing trajectories to approach the ground state $\varphi^* \equiv \min_x \varphi(x)$. In fact, \mathcal{E} is a *Lyapunov function*, enabling us to conclude that the system is *stable*³ on a minimizer $x^* \equiv \arg \min_x \varphi(x)$. This holds for any bounded damping $\eta(t) > 0$. In addition, if $\eta(t) = \eta$ is constant, then the system is *asymptotically stable*⁴ around x^* .

As explained in Sec. II A, the gradient flow (2) corresponds to a large friction limit of the accelerated gradient flow (3). It is straightforward to show that the gradient flow is *asymptotically stable* on x^* .⁵ Thus, its trajectories converge to x^* , and so do trajectories of (3) with a constant damping. However, when $\eta(t)$ is a decreasing function of time, such as in Eq. (25), we can only conclude stability in general. This means that trajectories can oscillate around x^* without ever converging. This is intuitive since $\eta(t)$ becomes very small for large t and the system becomes almost conservative—we refer to Ref. [27] for a more thorough stability analysis of these systems.

Besides stability, it is possible to estimate *how fast* trajectories approach a minimum of φ . This can also be done via a Lyapunov analysis, under certain convexity assumptions on φ . A function φ is said to be *convex* if its gradient $\nabla\varphi$ is maximal monotone, i.e., it obeys the inequality (6). A function φ is said to be *strongly convex* with parameter $\mu > 0$ if it obeys a stronger condition: $\langle \nabla\varphi(y) - \nabla\varphi(x) | y - x \rangle \geq \mu \|y - x\|^2$. Let us mention some known rates of convergence which follow from our results in Ref. [28]. For the gradient flow (2) we have

$$\varphi(x(t)) - \varphi^* = \mathcal{O}(t^{-1}) \quad (\text{convex}), \quad (30a)$$

$$\|x(t) - x^*\|^2 = \mathcal{O}(e^{-\mu t}) \quad (\text{strongly convex}). \quad (30b)$$

For the accelerated gradient flow (3) with constant damping, $\eta(t) = \eta = \text{const.}$, we have

$$\varphi(x(t)) - \varphi^* = \mathcal{O}(t^{-1}) \quad (\text{convex}), \quad (31a)$$

$$\|x(t) - x^*\|^2 = \mathcal{O}(e^{-\sqrt{\mu}t}) \quad (\text{strongly convex}). \quad (31b)$$

For the accelerated gradient flow (3) with decaying damping, $\eta(t) = r/t$ where $r \geq 3$, we have

$$\varphi[x(t)] - \varphi^* = \mathcal{O}(t^{-2}) \quad (\text{convex}), \quad (32a)$$

$$\|x(t) - x^*\|^2 = \mathcal{O}(t^{-2r/3}) \quad (\text{strongly convex}). \quad (32b)$$

We thus see that the accelerated gradient flow (3) may converge faster than the gradient flow (2) in some situations. For instance, the bound (31b) has a \sqrt{t} improvement in the exponential compared to (30b), while (32a) is an order of magnitude faster compared to (30a). Besides these rates, the stability of the system also plays a role. However, one should note that these rates are upper bounds, and thus may not always reflect the actual behavior of the system which may be faster for a particular φ .

³Stability means that the trajectories stay nearby x^* for all times.

⁴This is stronger, i.e., $x(t) \rightarrow x^*$ as $t \rightarrow \infty$. This result can be derived from LaSalle’s invariance principle.

⁵Indeed, consider the Lyapunov function $\mathcal{E} \equiv \varphi(x) - \varphi^*$. We have $\mathcal{E} \geq 0$ and $\dot{\mathcal{E}} = -\|\nabla\varphi\|^2 \leq 0$, and note that such inequalities are strict off a critical point.

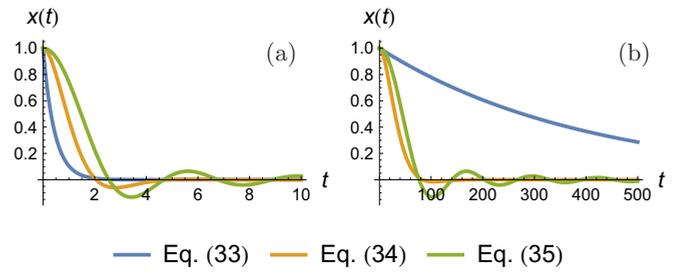


FIG. 1. Solutions of Eqs. (2) and (3) with $\varphi = (1/2)\omega^2 x^2$. For constant damping (24) we choose η slightly below the critical value, and for decaying damping (25) we choose $r = 3$. (a) $\omega = 1.5$. (b) $\omega = 0.05$.

When φ is quadratic we can solve Eqs. (2) and (3) exactly for some choices of damping $\eta(t)$. Recall that this captures the behavior close to an isolated minimum: $\varphi(x) - \varphi(x^*) \approx \frac{1}{2}(x - x^*)^T \nabla^2\varphi(x^*)(x - x^*)$. We can change coordinates to a basis where the Hessian $\nabla^2\varphi(x^*)$ is diagonal so that the components of the ODE become decoupled. It is thus sufficient to consider the one-dimensional case $\varphi(x) = \omega^2 x^2/2$. Thus, the gradient flow (2) has solution

$$x(t) = x_0 e^{-\omega^2 t}, \quad (33)$$

which agrees with the rate (30b). The accelerated gradient flow (3) with constant damping has solution

$$x(t) = x_0 e^{-\eta t/2} [\cosh(\xi t/2) + (\eta/\xi) \sinh(\xi t/2)], \quad (34)$$

where $\xi \equiv \sqrt{\eta^2 - 4\omega^2}$ (we assumed $\dot{x}_0 = 0$). This solution shows an exponential decay that agrees with the bound (31b). Similarly, the accelerated gradient flow (3) with a decaying damping has solution

$$x(t) = x_0 2^\nu \omega^{-\nu} \Gamma(\nu + 1) t^{-\nu} J_\nu(\omega t), \quad \nu \equiv \frac{r-1}{2}, \quad (35)$$

where J_ν is the Bessel function of the first kind (again, $\dot{x}_0 = 0$). A series expansion of J_ν for large t reveals that $J_\nu(t) \sim 1/\sqrt{t}$, which implies $x(t) \sim t^{-r/2}$. This is a power law, faster than the general upper bound (32b), however slower than the exponential decay in Eq. (34).

We illustrate the above solutions in Fig. 1. Note that when $\omega > 1$ the gradient flow tends to converge faster to equilibrium, however, when $\omega < 1$ there is a significant faster convergence of the accelerated gradient flow in both cases.⁶ In particular, the oscillations can be better controlled with a constant damping.

The continuous-time dynamics is often easier to analyze compared to the potentially complicated recurrence relations of a discrete-time algorithm. Thus, knowing the behavior of systems (2) and (3) provide useful insights to understand and design “good” optimization methods. We expect that *reasonable discretizations* of these systems will reproduce their behavior, at least for sufficiently small choices of the step size.

⁶Most challenging optimization problems in machine learning tend to have $\omega < 1$ since this comes from a poor condition number of the objective function.

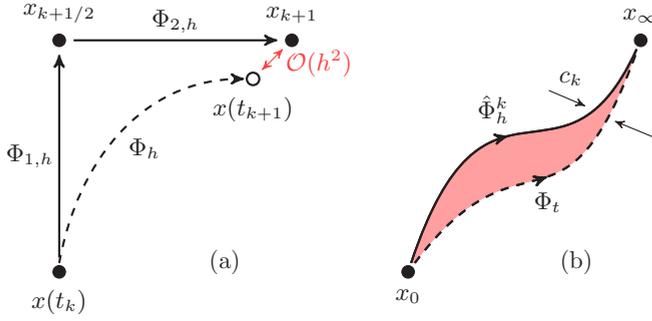


FIG. 2. (a) Splitting approach in the flow composition (38). In one time step the error between the continuous and discrete trajectory is $\|x(t_{k+1}) - x_{k+1}\| = \mathcal{O}(h^2)$. (b) The (re)balanced splitting introduces c_k which forces the discrete-time evolution to converge to a critical point of the ODE: $\hat{\Phi}_h^k(x_0) = \Phi_t(x_0) \rightarrow x_\infty$ as $k, t \rightarrow \infty$.

More precisely, it is a general result in numerical analysis [29] that if $\nabla\varphi$ is Lipschitz continuous with constant L , a numerical integrator of order $p \geq 1$ has *global error* $e_k \equiv x(t_k) - x_k$ such that $\|e_k\| \leq C_t h^p$, where $t_k = kh$ and $C_t = C(e^{Lt} - 1)$ for some constant $C > 0$. Thus, given a fixed simulation time t , one can control e_k by making the step size h sufficiently small. For our purposes, the dissipative system presumably converges to a minimum of φ fast enough, thus we expect that the simulation time will not be excessively large which helps controlling C_t —this is a better scenario compared to long-time simulations of conservative systems which are common, e.g., in molecular dynamics and astrophysics. Moreover, this estimate on C_t is general and pessimistic; a particular method may have a much smaller C_t . In the following sections we will show that several well-known—but also new—optimization methods are actually discretizations of the gradient flow (2) or the accelerated gradient flow (3).

IV. ACCELERATED EXTENSIONS OF ADMM

A. Balanced and rebalanced splitting

Before making connections with ADMM we need to introduce some ideas about splitting methods for ODEs. Thus, consider the dynamical system

$$\dot{x} = A(x), \quad A \equiv A_1 + A_2, \quad (36)$$

where A_1, A_2 represent smooth and single-valued vector fields. Suppose this is an intractable problem, i.e., the structure of φ makes the problem not amenable to a numerical procedure. We denote the flow of Eq. (36) by Φ_t . The idea is to split the vector field A such that each individual system

$$\dot{x} = A_1(x), \quad \dot{x} = A_2(x), \quad (37)$$

is integrable or has a feasible numerical approximation. We denote their respective flows by $\Phi_{1,t}$ and $\Phi_{2,t}$. For a step size $h > 0$, it can be shown [29] that the simplest composition

$$\hat{\Phi}_h = \Phi_{2,h} \circ \Phi_{1,h} \quad (38)$$

yields a first-order approximation, namely, the *local error* is $\|\Phi_h(x) - \hat{\Phi}_h(x)\| = \mathcal{O}(h^2)$ —see Fig. 2(a) for an illustration. However, in general, splittings such as (37) *do not* preserve critical points of the original ODE. The proposal of Ref. [26]

is to introduce a *balance coefficient*, $c = c(t)$, and replace Eq. (37) by

$$\dot{x} = A_1(x) + c, \quad \dot{x} = A_2(x) - c. \quad (39)$$

By appropriately choosing c we can then preserve critical points. To see this, first suppose that x_∞ is a critical point of Eq. (36), i.e., $A_1(x_\infty) + A_2(x_\infty) = 0$. If c_∞ obeys

$$c_\infty = \frac{1}{2}(A_2(x_\infty) - A_1(x_\infty)), \quad (40)$$

then x_∞ is also a critical point of both individual Eqs. (39). Conversely, suppose x_∞ is a critical point of both individual Eqs. (39). We then have

$$c_\infty = A_2(x_\infty) = -A_1(x_\infty) = \frac{1}{2}(A_2(x_\infty) - A_1(x_\infty)), \quad (41)$$

implying that x_∞ is also a critical point of the original system (36). This can be implemented with $c_{k+1} = \frac{1}{2}(A_2(x_k) - A_1(x_k))$, together with suitable discretizations of Eqs. (39)—see Fig. 2(b) for an illustration. However, this approach requires explicit computation of the vector fields A_i 's. In optimization this means computing gradients, which may not be available. To address this issue we consider a related approach.

The *rebalanced splitting* [26] is particularly well-suited for the *implicit* discretizations we have in mind. We thus integrate $\dot{x} = A_1(x) + c_k$ over the interval $[t_k, t_k + h]$, with initial condition $x(t_k) = x_k$, to obtain the intermediate point $x_{k+1/2}$. Then we integrate $\dot{x} = A_2(x) - c_k$ over the same interval, with initial condition $x(t_k) = x_{k+1/2}$, to obtain the endpoint x_{k+1} . Note that c_k is kept fixed during this procedure. Thus,

$$x_{k+1/2} = x_k + \int_{t_k}^{t_k+h} (A_1[x(t)] + c_k) dt, \quad (42a)$$

$$x_{k+1} = x_{k+1/2} + \int_{t_k}^{t_k+h} (A_2[x(t)] - c_k) dt. \quad (42b)$$

In light of Eq. (41), two reasonable ways of computing c_k are given by the *average* of either $\frac{1}{2}(A_2 - A_1)$ or A_2 . We choose the latter—as we will see this allow us to derive ADMM—which with Eq. (42) yields

$$\begin{aligned} c_{k+1} &= \frac{1}{h} \int_{t_k}^{t_k+h} A_2[x(t)] dt \\ &= c_k + h^{-1}(x_{k+1} - x_{k+1/2}). \end{aligned} \quad (43)$$

In contrast to the previous case, we now need not compute A_i explicitly. Our derivation above is slightly different than the one in Ref. [26].

B. Deriving extensions of ADMM

We are now in a position to show how generalizations of ADMM emerge from such an approach. We focus on problem

$$\min_{x \in \mathbb{R}^n} \varphi(x), \quad \varphi = \varphi_1 + \varphi_2 + \varphi_3, \quad (44)$$

and moreover on discretizations of the accelerated gradient flow (3) since discretizations of the gradient flow (2) can be recovered as particular cases; recall the discussion in Sec. II A.

Algorithm 1: Family of accelerated extensions of ADMM for solving the problem (44).

Choose step size h and damping function γ_k
Initialize x_0, \hat{x}_0 and $c_0 = 0$
for $k = 0, 1, \dots$ **do**
 $x_{k+1/2} = \text{prox}_{h\varphi_1}(\hat{x}_k - h\nabla\varphi_3(\hat{x}_k) + hc_k)$
 $x_{k+1} = \text{prox}_{h\varphi_2}(x_{k+1/2} - hc_k)$
 $c_{k+1} = c_k + h^{-1}(x_{k+1} - x_{k+1/2})$
 $\hat{x}_{k+1} = x_{k+1} + \gamma_{k+1}(x_{k+1} - x_k)$
end

With a balance coefficient $c = c(t)$ we write Eq. (3) as

$$\dot{x} = p, \quad (45a)$$

$$\dot{p} = \underbrace{-\eta(t)p - \nabla\varphi_1(x) - \nabla\varphi_3(x) + c}_{A_1} - \underbrace{\nabla\varphi_2(x)}_{A_2} - c. \quad (45b)$$

Splitting as indicated, and further combining the resulting equations, we obtain two independent ODEs:

$$\ddot{x} + \eta(t)\dot{x} = -\nabla\varphi_1(x) - \nabla\varphi_3(x) + c, \quad (46)$$

$$\ddot{x} = -\nabla\varphi_2(x) - c. \quad (47)$$

Using Eq. (21), and redefining the step size $h^2 \mapsto h$ as discussed in Eq. (23), a semi-implicit discretization of Eq. (46) is

$$\frac{x_{k+1/2} - \hat{x}_k}{h} = -\nabla\varphi_1(x_{k+1/2}) - \nabla\varphi_3(\hat{x}_k) + c_k. \quad (48)$$

This can be solved with the resolvent (4) to obtain

$$x_{k+1/2} = J_{h\partial\varphi_1}(\hat{x}_k - h\nabla\varphi_3(\hat{x}_k) + hc_k). \quad (49)$$

We now discretize Eq. (47) as

$$\frac{\tilde{x}_{k+1} - 2x_{k+1/2} + \hat{x}_k}{h} = -\nabla\varphi_2(x_{k+1}) - c_k, \quad (50)$$

where

$$\tilde{x}_{k+1} \equiv x_{k+1} + (x_{k+1/2} - \hat{x}_k). \quad (51)$$

Note that \tilde{x}_{k+1} is related to x_{k+1} via the ‘‘momentum’’ term ($x_{k+1/2} - \hat{x}_k$) based on the first splitting.⁷ With this and the resolvent we obtain

$$x_{k+1} = J_{h\partial\varphi_2}(x_{k+1/2} - hc_k). \quad (52)$$

The balance coefficient follows readily from Eq. (43):⁸

$$c_{k+1} = c_k + h^{-1}(x_{k+1} - x_{k+1/2}). \quad (53)$$

⁷ \tilde{x}_{k+1} is slightly further away from x_{k+1} which makes the algorithm ‘‘look ahead’’ and implicitly introduces dependency on the curvature of φ_2 . Although the introduction of \tilde{x}_{k+1} may seem artificial, it will be justified below when we compute the error in approximating the continuous trajectory.

⁸To justify that $h^2 \mapsto h$ does not change this, note that with $A_2 = -\nabla\varphi_2$ in Eq. (43) an implicit discretization corresponds to approximating the integral by its upper limit, thus $c_{k+1} = \frac{1}{h} \int_{t_k}^{t_k+h} B dt \approx -\nabla\varphi_2(x_{k+1})$. Using Eqs. (50) and (51) yields Eq. (53).

Combining the above steps results into a family of accelerated extensions of ADMM summarized in Algorithm 1.

Let us stress some important aspects of Algorithm 1. The standard ADMM [3–5] corresponds to the particular case where $\varphi_3 = 0$ and no acceleration is used, i.e., $\gamma_k = 0$. Thus, Algorithm 1 not only generalizes ADMM to handle problems in the form of Eq. (44) but also includes acceleration with arbitrary damping functions $\eta(t)$. The so-called dual vector in ADMM, originally obtained as a Lagrange multiplier [5], is here represented by the balance coefficient c_k and thus acquires a new meaning: its role is to preserve critical points of the ODE.

When decaying damping (25) is chosen and $\varphi_3 = 0$, Algorithm 1 is similar to the so-called ‘‘fast ADMM’’ [30]. They differ in that the latter also ‘‘accelerates’’ the dual variable c_k . Connections between fast ADMM and ODEs was considered recently by us in Refs. [27,28] and also corresponds to Eq. (3), however the discretization is not a rebalanced splitting within the above framework.

In light of the discussion in Sec. II, it is clear that Algorithm 1 can be generalized to account for monotone operators by the replacements $\text{prox}_{h\varphi_1} \rightarrow J_{hA_1}$ and $\text{prox}_{h\varphi_2} \rightarrow J_{hA_2}$.

Finally, although we focused on the accelerated gradient flow (3), analogous procedure applies to the gradient flow (2) and leads to Algorithm 1 with $\gamma_k = 0$.⁹

C. Order of accuracy

Next, we show that the above discretization is justified since it is a first-order approximation to the continuous trajectory, i.e., $\|\Phi_h(x) - \hat{\Phi}_h(x)\| = \mathcal{O}(h^2)$. From the definition of the resolvent (4) we have that $y = J_{h\nabla\varphi}(x)$ if and only if $y = x - h\nabla\varphi(y)$, thus

$$y = x - h\nabla\varphi(x - h\nabla\varphi(y)) = x - h\nabla\varphi(x) + \mathcal{O}(h^2). \quad (54)$$

This relation implies the following approximations, valid up to $\mathcal{O}(h^2)$, for the updates in Algorithm 1:

$$x_{k+1/2} \approx \hat{x}_k - h\nabla\varphi_3(\hat{x}_k) + hc_k - h\nabla\varphi_1(\hat{x}_k), \quad (55)$$

$$x_{k+1} \approx \hat{x}_k - h\nabla\varphi(\hat{x}_k). \quad (56)$$

Recall Eq. (20), namely, $\gamma_k = 1 - \sqrt{h}\eta(t_k)$, since we redefined $h^2 \mapsto h$. Thus,

$$\hat{x}_k = x_k + (1 - \eta(t_k)\sqrt{h})\sqrt{h}p_k, \quad (57)$$

where

$$p_k \equiv \frac{x_k - x_{k-1}}{\sqrt{h}}. \quad (58)$$

From Eqs. (57) and (56), and now restoring the original step size ($h \mapsto h^2$), we conclude that

$$p_{k+1} \approx p_k - h\eta(t_k)p_k - h\nabla\varphi(x_k), \quad (59)$$

$$x_{k+1} = x_k + hp_{k+1} \approx x_k + hp_k. \quad (60)$$

⁹We have $\dot{x} = -\nabla(\varphi_1 + \varphi_3)(x) + c$ and $\dot{x} = -\nabla\varphi_2(x) - c$. For the former, $x_{k+1/2} - x_k = -h\nabla\varphi_1(x_{k+1/2}) - h\nabla\varphi_3(x_k) + hc_k$, whereas for latter, $x_{k+1} - x_{k+1/2} = -h\nabla\varphi_2(x_{k+1}) - hc_k$. Using the resolvent and Eq. (43) yield Algorithm 1 with $\gamma_k = 0$.

Finally, the evolution of Eq. (3) in one time step gives

$$\begin{aligned} p(t+h) &= p(t) + h\dot{p}(t) + \mathcal{O}(h^2) \\ &\approx p(t) - h\eta(t)p(t) - h\nabla\varphi[x(t)], \end{aligned} \quad (61)$$

$$\begin{aligned} x(t+h) &= x(t) + h\dot{x}(t) + \mathcal{O}(h^2) \\ &\approx x(t) + hp(t). \end{aligned} \quad (62)$$

Comparing with Eqs. (59) and (60) implies that the algorithm's state simulates the continuous trajectory up to an error $\mathcal{O}(h^2)$, therefore the discretization is first-order accurate. Similar conclusion holds for the nonaccelerated algorithm in relation to the gradient flow (2).

We mention a subtlety when φ is nonsmooth, or when considering monotone operators more generally. A crucial step in the above derivation was the Taylor approximation of the resolvent (54). For a maximal monotone operator A , in the most general case only a slightly weaker approximation is available [15, Remark 23.47]:

$$J_{hA} = I - hA_0 + \mathcal{O}(h), \quad (63)$$

where $A_0x = \lim_{\mu \downarrow 0} A_\mu x$ (see Sec. II). The previous arguments still hold true, however, due to Eq. (63), and assuming that we can expand $A_0[x + \mathcal{O}(h)] = A_0(x) + \mathcal{O}(h)$, the local error is now $\mathcal{O}(h)$ instead of $\mathcal{O}(h^2)$. It is important to note that this is a consequence of the nonsmoothness of φ , or the multivaluedness of A , and not of the discretization procedure. This comment applies to all cases considered in this paper.

V. ACCELERATED EXTENSIONS OF DAVIS-YIN

A. Discretization

We now introduce accelerated extensions of Davis-Yin [8]. This time we split the system (3) without a balance coefficient, namely, we choose vector fields

$$A_1(x) = -\eta(t)\dot{x} - \nabla\varphi_1(x), \quad (64)$$

$$A_2(x) = -\nabla\varphi_2(x) - \nabla\varphi_3(x). \quad (65)$$

Instead of Eqs. (46) and (47) we now obtain

$$\ddot{x} + \eta(t)\dot{x} = -\nabla\varphi_1(x), \quad (66)$$

$$\ddot{x} = -\nabla\varphi_2(x) - \nabla\varphi_3(x). \quad (67)$$

Using Eq. (21) (and redefining $h^2 \mapsto h$), an implicit discretization of Eq. (66) is

$$\frac{x_{k+1/4} - \hat{x}_k}{h} = -\nabla\varphi_1(x_{k+1/4}), \quad (68)$$

which with the resolvent (4) gives

$$x_{k+1/4} \equiv \Phi_{1,h}(\hat{x}_k) = J_{h\partial\varphi_1}(\hat{x}_k). \quad (69)$$

Next, to “inject momentum” in the direction of $\nabla\varphi_1$, we define the “translation operator”

$$\mathcal{T}_h(z) \equiv z - h\nabla\varphi_1(x_{k+1/4}). \quad (70)$$

The next point is thus obtained as

$$x_{k+1/2} \equiv \mathcal{T}_h(x_{k+1/4}) = 2x_{k+1/4} - \hat{x}_k. \quad (71)$$

Algorithm 2: Family of accelerated extensions of Davis-Yin (DY) for problem (44).

Choose step size h and damping function γ_k

Initialize x_0 and \hat{x}_0

for $k = 0, 1, \dots$ **do**

$$x_{k+1/4} = \text{prox}_{h\varphi_1}(\hat{x}_k)$$

$$x_{k+1/2} = 2x_{k+1/4} - \hat{x}_k$$

$$x_{k+3/4} = \text{prox}_{h\varphi_2}(x_{k+1/2} - h\nabla\varphi_3(x_{k+1/4}))$$

$$x_{k+1} = \hat{x}_k + x_{k+3/4} - x_{k+1/4}$$

$$\hat{x}_{k+1} = x_{k+1} + \gamma_{k+1}(x_{k+1} - x_k)$$

end

A semi-implicit discretization of Eq. (67) is

$$\frac{x_{k+3/4} - 2x_{k+1/4} + \hat{x}_k}{h} = -\nabla\varphi_2(x_{k+3/4}) - \nabla\varphi_3(x_{k+1/4}), \quad (72)$$

which can again be solved with the resolvent (4) to obtain

$$\begin{aligned} x_{k+3/4} &\equiv \Phi_{2,h}(\hat{x}_k) \\ &= J_{h\partial\varphi_2}(x_{k+1/2} - h\nabla\varphi_3(x_{k+1/4})). \end{aligned} \quad (73)$$

Finally, applying the inverse $\mathcal{T}_h^{-1}(z) \equiv z + h\nabla\varphi_1(x_{k+1/4})$ and using Eq. (68) we obtain

$$x_{k+1} \equiv \mathcal{T}_h^{-1}(x_{k+3/4}) = x_{k+3/4} - (x_{k+1/4} - \hat{x}_k). \quad (74)$$

We collect the above steps into Algorithm 2. The original Davis-Yin method [8] is recovered by setting $\gamma_k = 0$. Such a case corresponds to an overdamped limit—recall the discussion of Sec. II A—which is indeed a discretization of the gradient flow (2), as can also be easily verified by repeating the above procedure to this simpler case.

Algorithm 2 is equivalent to the fixed point iteration $x_{k+1} = \hat{\Phi}_h(\hat{x}_k)$ with

$$\hat{\Phi}_h \equiv \mathcal{T}_h^{-1} \circ \Phi_{2,h} \circ \mathcal{T}_h \circ \Phi_{1,h}, \quad (75)$$

where these individual maps are defined in Eqs. (69), (70), and (73). Thus, the translation operator \mathcal{T}_h is actually a “pre-processor map,” which is a common technique in numerical analysis [29]. The discretization associated to Davis-Yin can be summarized diagrammatically:

$$\begin{array}{ccccc} \hat{x}_k & \xrightarrow{\Phi_{1,h}} & x_{k+1/4} & \xrightarrow{\mathcal{T}_h} & x_{k+1/2} \\ & \searrow & & & \downarrow \Phi_{2,h} \\ & & & & x_{k+3/4} \\ & & & \xleftarrow{\mathcal{T}_h^{-1}} & x_{k+1} \end{array} \quad (76)$$

B. Order of accuracy

Using the expansion (54) we can approximate the updates of Algorithm 2 up to $\mathcal{O}(h^2)$:

$$x_{k+1/4} \approx \hat{x}_k - h\nabla\varphi_1(\hat{x}_k), \quad (77a)$$

$$x_{k+1/2} \approx \hat{x}_k - 2h\nabla\varphi_1(\hat{x}_k), \quad (77b)$$

$$x_{k+3/4} \approx \hat{x}_k - 2h\nabla\varphi(\hat{x}_k), \quad (77c)$$

$$x_{k+1} \approx \hat{x}_k - h\nabla\varphi(\hat{x}_k). \quad (77d)$$

But Eq. (77d) is exactly the same as Eq. (56), thus the remaining steps of the argument follow as before, implying that Algorithm 2 is also a first-order integrator to the accelerated gradient flow Eq. (3). The same holds true for standard Davis-Yin ($\gamma_k = 0$) in relation to the gradient flow (2).

C. Preserving critical points

Since Algorithm 2 arises from a splitting that is not balanced, it is not a priori obvious if critical points of the underlying ODE are preserved. We now show that this is indeed the case. We can write the operator (75) as

$$\hat{\Phi}_h = I + J_{h\partial\varphi_2} \circ (2J_{h\partial\varphi_1} - I - h\nabla\varphi_3 \circ J_{h\partial\varphi_1}) - J_{h\partial\varphi_1}. \quad (78)$$

Assuming the algorithm converges, we must have a fixed point equation $x_\infty = \hat{\Phi}_h(x_\infty)$. We thus need to show that this equation generates critical points of the ODE. To this end, let x^* be such a critical point, i.e.,

$$(\nabla\varphi_1 + \nabla\varphi_2 + \nabla\varphi_3)(x^*) = 0. \quad (79)$$

This is equivalent to $(I + h\nabla\varphi_2)(x^*) = (I - h\nabla\varphi_1 - h\nabla\varphi_3)(x^*)$, and with the aid of the resolvent (4) it can be written as

$$x^* = J_{h\partial\varphi_2} \circ (I - h\nabla\varphi_1 - h\nabla\varphi_3)(x^*). \quad (80)$$

Using the identity

$$(2J_{h\partial\varphi} - I) \circ (I + h\nabla\varphi) = I - h\nabla\varphi, \quad (81)$$

we thus have

$$x^* = J_{h\partial\varphi_2} \circ [(2J_{h\partial\varphi_1} - I) \circ (I + h\nabla\varphi_1) - h\nabla\varphi_3](x^*). \quad (82)$$

Define $x_\infty \equiv (I + h\nabla\varphi_1)(x^*)$, i.e., $x^* = J_{h\partial\varphi_1}(x_\infty)$. The above equation then yields

$$J_{h\partial\varphi_1}(x_\infty) = J_{h\partial\varphi_2} \circ [2J_{h\partial\varphi_1} - I - h\nabla\varphi_3 \circ J_{h\partial\varphi_1}](x_\infty), \quad (83)$$

which is equivalent to $x_\infty = \hat{\Phi}_h(x_\infty)$ according to Eq. (78). Therefore, critical points of Eq. (79) yield fixed points of the operator (78) and vice versa. This shows that Algorithm 2 preserves critical points of the underlying ODE.

D. Accelerated extensions of Douglas-Rachford

Douglas-Rachford (DR) [2,6] is recovered from Algorithm 2 in the particular case where $\gamma_k = 0$ and $\varphi_3 = 0$. Therefore, in the case where $\varphi_3 = 0$ but $\gamma_k \neq 0$, Algorithm 2 contains accelerated extensions of Douglas-Rachford—the case with decaying damping was considered in Ref. [31]. From the previous arguments, we know that such methods are discretizations of the accelerated gradient flow (3), whereas the standard Douglas-Rachford is a discretization of the gradient flow (2). Moreover, such discretizations preserve critical points and are first-order integrators.

E. Accelerated extensions of forward-backward

The forward-backward method [6] is recovered from Algorithm 2 when $\gamma_k = 0$ and $\varphi_1 = 0$. Thus, when $\varphi_1 = 0$ but $\gamma_k \neq 0$, Algorithm 2 reduces to

$$\begin{aligned} x_{k+1} &= \text{prox}_{h\varphi_2}(\hat{x}_k - h\nabla\varphi_3(\hat{x}_k)), \\ \hat{x}_{k+1} &= x_{k+1} + \gamma_{k+1}(x_{k+1} - x_k). \end{aligned} \quad (84)$$

Algorithm 3: Family of accelerated extensions of Tseng’s method.

```

Choose step size  $h$  and damping function  $\gamma_k$ 
Initialize  $x_0$  and  $\hat{x}_0$ 
for  $k = 0, 1, \dots$  do
   $x_{k+1/2} = \text{prox}_{h\varphi_2}(\hat{x}_k - h\nabla\varphi_3(\hat{x}_k))$ 
   $x_{k+1} = x_{k+1/2} - h(\nabla\varphi_3(x_{k+1/2}) - \nabla\varphi_3(\hat{x}_k))$ 
   $\hat{x}_{k+1} = x_{k+1} + \gamma_{k+1}(x_{k+1} - x_k)$ 
end

```

The decaying damping case (25) was considered in Ref. [32]. From an ODE perspective, the above discretization is not a splitting method but rather a direct semi-implicit discretization of Eq. (3). Anyhow, our previous arguments show that such accelerated variants of forward-backward are first-order integrators of this ODE and preserve critical points—the same holds true for the standard forward-backward ($\gamma_k = 0$) in relation to the gradient flow (2).

VI. ACCELERATED EXTENSIONS OF TSENG’S SPLITTING

The last proximal method remaining to be considered is the forward-backward-forward or Tseng’s splitting [7]. Thus, consider Eq. (3) with $\varphi_1 = 0$ and written as

$$\dot{p} = \underbrace{-\eta(t)p - \nabla\varphi_2(x) - \nabla\varphi_3(x)}_{A_1} + \underbrace{\nabla\varphi_3(x) - \nabla\varphi_3(x)}_{A_2}. \quad (85)$$

Note that $A_2 = 0$, however in a discretization there might be numerical inaccuracies which introduces a kind of “perturbation” on top of the forward-backward method, which arises from the first component of this system. Splitting as indicated yields

$$\ddot{x} + \eta(t)\dot{x} = -\nabla\varphi_2(x) - \nabla\varphi_3(x), \quad (86)$$

$$\ddot{x} = \nabla\varphi_3(x) - \nabla\varphi_3(x). \quad (87)$$

A semi-implicit discretization of the first equation yields

$$x_{k+1/2} = J_{h\partial\varphi_2}[\hat{x}_k - h\nabla\varphi_3(\hat{x}_k)], \quad (88)$$

which is the forward-backward method (84). Eq. (87) can be discretized as $\tilde{x}_{k+1} - 2x_{k+1/2} + \hat{x}_k = h\nabla\varphi_3(\hat{x}_k) - h\nabla\varphi_3(x_{k+1/2})$, where \tilde{x}_{k+1} is given by (51). Thus,

$$x_{k+1} = x_{k+1/2} - h(\nabla\varphi_3(x_{k+1/2}) - \nabla\varphi_3(\hat{x}_k)). \quad (89)$$

Therefore, we derived Algorithm 3.

The original Tseng’s splitting is recovered with $\gamma_k = 0$, in which case it is a discretization of the gradient flow (2). Due to Eq. (87) we expect to have a “contraction” on the acceleration which indicates that Algorithm 3 may be slower than (84) (this was actually observed in our experiments and this method tends to be slower than forward-backward).

In a similar way as already done in Secs. IV and V, through Taylor expansions it is straightforward to show that the above discretization is first-order accurate.

We can also show that the above discretization preserves critical points. Indeed, Algorithm 3 is equivalent to iterations

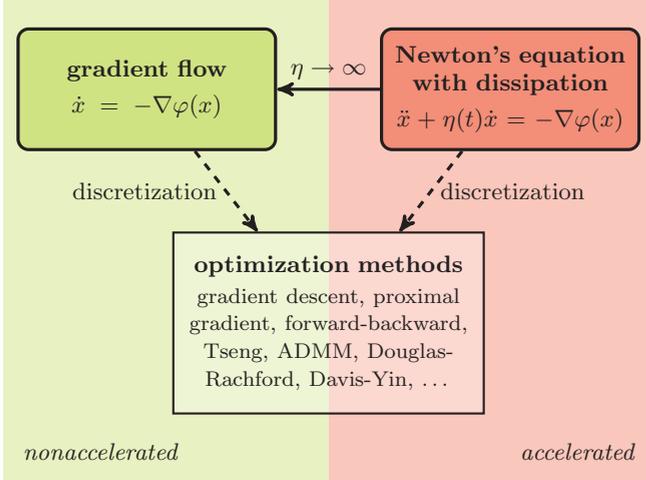


FIG. 3. Several well-known optimization methods arise as discretizations of the gradient flow. Moreover, “accelerated extensions” of these methods arise as discretizations of Newton’s equation with a dissipative term; the gradient flow corresponds to a large friction limit of the latter. Therefore, these optimization methods consist of actual *simulations* of a classical dissipative physical system.

$x_{k+1} = \hat{\Phi}_h(\hat{x}_k)$ with

$$\hat{\Phi}_h = (I - h\nabla\varphi_3) \circ J_{h\partial\varphi_2} \circ (I - h\nabla\varphi_3) + h\nabla\varphi_3. \quad (90)$$

Assuming the algorithm converges, $x_\infty = \hat{\Phi}_h(x_\infty)$, i.e.,

$$(I - h\nabla\varphi_3)(x_\infty) = (I - h\nabla\varphi_3) \circ J_{h\partial\varphi_2} \circ (I - h\nabla\varphi_3)(x_\infty). \quad (91)$$

Moreover, assuming that h is sufficiently small so that the inverse $(I - h\nabla\varphi_3)^{-1}$ exists, the above yields $x_\infty = J_{h\partial\varphi_2} \circ (I - h\nabla\varphi_3)(x_\infty)$. By the definition of resolvent in Eq. (4) this is equivalent to $\nabla(\varphi_2 + \varphi_3)(x_\infty) = 0$. Thus, the iterates of this algorithm generates critical points of the underlying ODE.

VII. A CLASSICAL VIEW ON DETERMINISTIC OPTIMIZATION

At this stage, let us summarize and interpret the previous results from a physics perspective. They can be summarized by Fig. 3. Several well-known optimization methods—including gradient- or proximal-based—are simply different discretizations of the gradient flow. This includes gradient descent (see the Appendix), forward-backward, Tseng’s splitting, ADMM, Douglas-Rachford, Davis-Yin, and potentially many other methods. However, accelerated extensions of these methods arise as discretizations of Newton’s equation with a dissipative term. Besides the accelerated proximal methods we introduced, we also have Nesterov and heavy ball (see the Appendix and Ref. [21]), and potentially many others as well. All these methods are provably first-order integrators, i.e., they have a local error $\mathcal{O}(h^2)$, and preserve critical points of the associated ODE. In addition, let us mention that we have recently generalized symplectic integrators to general dissipative Hamiltonian systems [33], which may offer a systematic approach to construct gradient-based optimization methods. Therefore, all these optimization methods can be seen as actual *simulations* of a classical dissipa-

tive system; when someone solves an optimization problem by implementing one of these methods in a—presumably classical—computer, there is a simulation of one classical system, i.e., dissipative Newton’s equation, by another classical system, i.e., the computer.

There are many interesting open questions even at this classical level. For instance, given a potential $\varphi : \mathbb{R}^n \rightarrow \mathbb{R}$ where certain properties of its landscape are known, is there a lower bound on how fast a dissipative system can approach the ground state? If so, which specific system would achieve such rate of convergence? Given φ , what is the best way to dissipate energy, i.e., the best damping $\eta(t)$? What is the tradeoff between energy dissipation and stability? Answering these type of questions would allow us to design “optimal” optimization algorithms through suitable discretizations of dissipative physical systems.

A dissipative Newtonian system is deterministic, and so are the optimization methods obtained as its discretizations. Such a deterministic approach is suited to convex problems which have a unique global minimum, or to obtain local minima in the neighborhood of the initial state. However, to escape poor local minima in more complex landscapes, i.e., to solve more challenging nonconvex problems, we need to introduce some sort of perturbation or noise. Thus, in the next section, we will generalize these methods to stochastic optimization settings. This will allow us to make connections with Langevin and Fokker-Planck equations which are ubiquitous in nonequilibrium statistical mechanics. It is worth noticing that such an approach is in some sense more closely related to sampling than to pure or deterministic optimization.

VIII. STOCHASTIC OPTIMIZATION

A. Stochastic gradient

One of the motivations behind stochastic optimization is to lighten the computational burden in computing full gradients over entire datasets, which is a bottleneck for high-dimensional problems with large data. The basic idea dates back to Robbins and Monro [34] and nowadays is widely used in machine learning, especially in training neural networks. Consider replacing the deterministic problem (1) by its stochastic counterpart

$$\min_{x \in \mathbb{R}^n} \mathbb{E}_\omega[\ell(x; \omega)], \quad (92)$$

where ω is a random variable from a sample space Ω . Specifically, suppose we have training data $\{\omega_1, \dots, \omega_N\}$ so that $\ell_i(x) \equiv \ell(x; \omega_i)$ is a random variable. Numerically, the above expectation is approximated by the empirical mean,

$$\bar{\ell}(x) \equiv \frac{1}{N} \sum_{i=1}^N \ell_i(x), \quad (93)$$

which is exact when $N \rightarrow \infty$. Thus, instead of computing $\nabla \bar{\ell}(x) = \frac{1}{N} \sum_{i=1}^N \nabla \ell_i(x)$ that may not be feasible, at each iteration of the algorithm we sample a “minibatch” \mathcal{B} , of size S , drawn uniformly at random—without replacement—from an index set $\{1, \dots, N\}$ and compute the so-called *stochastic*

gradient

$$\widetilde{\nabla} \ell(x) \equiv \frac{1}{S} \sum_{i \in \mathcal{B}} \nabla \ell_i(x). \tag{94}$$

Note that when $S = N$ this becomes the true gradient of the empirical loss (93). Importantly, when the dataset is very large, i.e., $S \ll N$ and $N \rightarrow \infty$, the *central limit theorem* comes into play and

$$\widetilde{\nabla} \ell(x) = \overline{\nabla} \ell(x) + \xi(x), \tag{95}$$

where $\xi(x) \sim \mathcal{N}(0, \Sigma(x))$. Thus, the stochastic gradient is an unbiased estimator of the true gradient of the empirical loss. It is reasonable to assume that the covariance matrix takes the form

$$\Sigma(x) = \frac{1}{S} C(x) C^T(x) \tag{96}$$

for some matrix $C(x)$. We do not know the specific form of $C(x)$, which is data dependent; however, in principle, it can be estimated. When $N, S \rightarrow \infty$, but with the ratio $S/N \ll 1$ kept fixed, we have $\Sigma \rightarrow 0$ and the stochastic gradient becomes the gradient of the expectation in Eq. (92)—this can be seen as a thermodynamic limit.

The stochastic gradient (94) can be implemented into the previous algorithms quite easily. Since φ_3 is the only function assumed to be differentiable in problem (44), we consider

$$\min_{x \in \mathbb{R}^n} \varphi_1(x) + \varphi_2(x) + \mathbb{E}_\omega[\varphi_3(x; \omega)]. \tag{97}$$

The entire family of Algorithms 1, 2, and 3 can be adapted to such case by adding two simple steps at each iteration, i.e., in the very first line of the “for” loop:

- (1) Sample a minibatch $\mathcal{B} \subset \{1, \dots, N\}$ of size S uniformly at random and without replacement;
- (2) Replace $\nabla \varphi_3 \rightarrow \widetilde{\nabla} \varphi_3$ in the subsequent updates.

B. Stochastic proximal operator

We can use similar ideas for proximal operators. As before, in each iteration of the algorithm we sample a minibatch \mathcal{B} and define

$$\widetilde{\ell}(x) \equiv \frac{1}{S} \sum_{i \in \mathcal{B}} \ell_i(x), \quad \ell_i \equiv \ell(x, \omega_i). \tag{98}$$

Thus, at each iteration, the algorithm has access to a random function $\widetilde{\ell}(x)$ that presumably “mimics” $\mathbb{E}_\omega[\ell(x; \omega)]$. We replace the proximal operator of the empirical loss, $\text{prox}_{h\bar{\ell}}(x)$, by its stochastic counterpart

$$\begin{aligned} \widetilde{\text{prox}}_{h\widetilde{\ell}}(x) &\equiv \text{prox}_{h\widetilde{\ell}}(x) \\ &= \arg \min_y \left(\widetilde{\ell}(y) - \frac{1}{2h} \|y - x\|^2 \right). \end{aligned} \tag{99}$$

Suppose we introduce stochasticity through φ_2 in problem (44), i.e.,

$$\min_{x \in \mathbb{R}^n} \varphi_1(x) + \mathbb{E}_\omega[\varphi_2(x; \omega)] + \varphi_3(x). \tag{100}$$

Then the family of Algorithms 1, 2, and 3 are adapted by adding the following instructions at each iteration:

- (1) Sample a minibatch $\mathcal{B} \subset \{1, \dots, S\}$ of size S uniformly at random and without replacement;
- (2) Replace $\text{prox}_{h\varphi_2} \rightarrow \widetilde{\text{prox}}_{h\varphi_2}$ in the next updates.

Note that, also in this case, a similar relation to (95) holds. Indeed, from Eq. (54) we get

$$\begin{aligned} \widetilde{\text{prox}}_{h\widetilde{\ell}}(x) &\approx x - h \widetilde{\nabla} \ell(x) \\ &= x - h \overline{\nabla} \ell(x) + h \xi(x) \\ &\approx \text{prox}_{h\bar{\ell}}(x) + h \xi(x). \end{aligned} \tag{101}$$

C. Langevin and Fokker-Planck equations

In the deterministic case we have the situation depicted in Fig. 3. In light of the discussion above, introducing a stochastic gradient or a stochastic proximal operator into these methods is equivalent to introducing a random perturbation in the associated ODEs. Thus, the only difference compared to the deterministic case is that $\nabla \varphi(x)$ is replaced by a “stochastic gradient,” $\widetilde{\nabla} \varphi(x)$, during a time interval of one step size h . As a consequence of Eqs. (95) and (101), we can describe this process by a Brownian motion provided we account for the correct power of the step size h when discretizing the noise term. Therefore, we must choose

$$\sqrt{\frac{h}{S}} C(x) dW \rightarrow \sqrt{\frac{h}{S}} C(x_k) \sqrt{h} \epsilon_k \equiv h \xi(x_k), \tag{102}$$

where W is a standard Wiener process, $\epsilon_k \sim \mathcal{N}(0, I)$, and $\xi(x)$ is the noise of the stochastic gradient (95), or the noise of the stochastic proximal operator (101). The gradient flow (2) is thus replaced by the *overdamped Langevin equation*

$$dx = -\nabla \varphi(x) dt + \sqrt{\frac{h}{S}} C(x) dW. \tag{103}$$

Similarly, the accelerated gradient flow (3) is replaced by the *underdamped Langevin equation*

$$dx = pdt, \tag{104a}$$

$$dp = -\nabla \varphi(x) dt - \eta(t) pdt + \sqrt{\frac{h}{S}} C(x) dW. \tag{104b}$$

There is one subtle point about these SDEs. They have a *multiplicative* white noise which is often ambiguous, e.g., the Itô-Stratonovich dilemma. In our context, it should be noted that the stochastic versions of the proximal methods previously discussed can be obtained from these SDEs as long as we discretize the noise consistently with the gradient of either $\nabla \varphi_3$ or $\nabla \varphi_2$, i.e., the previous splitting schemes must be followed carefully by discretizing $C(x)$ appropriately so that we can combine the noise into the stochastic gradient $\widetilde{\nabla} \varphi_3$ —in the case of problem (97)—or the stochastic proximal operator $\widetilde{\text{prox}}_{h\varphi_2}$ —in the case of problem (100). Naturally, there is no ambiguity if we assume *additive* noise, i.e., constant C , which should already provide insights into these methods, at least qualitatively.

We should also point out that stochastic versions of gradient descent, Polyak’s heavy ball, and Nesterov—which are gradient-based methods and widely used in machine learning—also follow from this approach. In these cases the discretization is explicit, i.e., one has a single gradient, no splitting, and no proximal operators. Moreover, the noise term is discretized in the Itô sense. For instance, stochastic gradient descent is simply an Euler-Maruyama discretization of Eq. (103). The stochastic version of Nesterov arises similarly

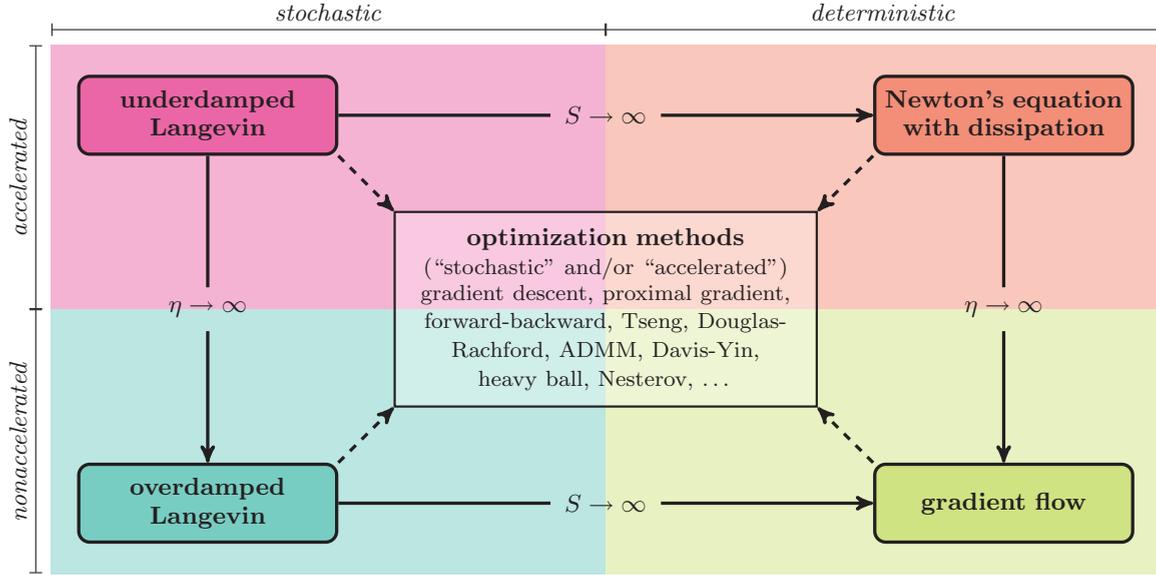


FIG. 4. Different optimization algorithms arise from different discretizations (dashed lines) of the *same physical system*. The left column of the diagram represents *stochastic processes* described by a Langevin equation, while the right column represents *deterministic processes* from (dissipative) classical mechanics—the transition between these phases is controlled by the “temperature” $T \sim 1/S$ where S is the batch size. The upper row of the diagram corresponds to an underdamped or *accelerated regime*, while the lower row corresponds to an *overdamped regime* where acceleration is negligible—the transition is controlled by the damping coefficient η . Discretizations from each colored quadrant yield (variants of) optimization algorithms in these “different phases.”

from Eq. (104) by using Eq. (21). For heavy ball, one should use a “conformal symplectic integrator” for the deterministic part of Eq. (104)—see Ref. [21]—and compose with an Itô discretization of the noise; we provide these derivations in the Appendix for completeness.

An interesting aspect of the above SDEs is that the ratio h/S plays the role of an “effective temperature” T . This is intuitive: small S means more noise in the stochastic gradient approximation, which is equivalent to raising the temperature of the “heat bath,” while increasing the step size simply amplifies the noise. The limit $S \rightarrow \infty$ —which presumes $N \rightarrow \infty$ and $S/N \ll 1$ —corresponds to removing the heat bath so that the previous SDEs together with their discretizations become deterministic. In similar vein as discussed in Sec. II A, the Langevin Eq. (103) can be recovered from Eq. (104) in the large friction limit. Therefore, the overall picture relating all possible variants of these optimization methods is depicted in Fig. 4; depending from which “quadrant” one chooses to discretize, and depending which discretization scheme is chosen, one can obtain an optimization algorithm with qualifiers such as “accelerated,” “stochastic,” or both, or none. The underdamped Langevin (top left quadrant) is the most general model that unifies all methods.

Now, we can readily write down the *Fokker-Planck* equation associated to the above SDEs, which describe the probability density $P(x, t)$ —or $P(x, p, t)$ for the accelerated methods—of the stochastic process. From the Chapman-Kolmogorov equation, through a standard derivation, we find that in the case of Eq. (103) we have

$$\frac{\partial P}{\partial t} = \nabla \cdot [P \nabla \varphi(x)] + \Delta_D P, \quad (105)$$

where we defined the diffusion matrix and the “stochastic Laplacian” as

$$D \equiv \frac{h}{2S} C(x) C(x)^T, \quad \Delta_D P \equiv \sum_{i,j} \frac{\partial^2}{\partial x_i \partial x_j} (D_{ij} P). \quad (106)$$

Similarly, the Fokker-Planck equation associated to the underdamped Langevin Eq. (104) is given by

$$\frac{\partial P}{\partial t} = -\nabla_x \cdot (pP) + \nabla_p \cdot [P \nabla \varphi(x) + \eta(t)pP] + \Delta_D P. \quad (107)$$

However, here $\Delta_D P = \sum_{i,j} D_{ij}(x) \partial_{p_i} \partial_{p_j} P$ since the noise is coupled only to the momenta.

In Eqs. (103) and (104) the Brownian motion arises from a stochastic gradient or stochastic proximal operator approximation, hence the diffusion dependence on the minibatch and step sizes. Alternatively, one can directly perturb the deterministic optimization methods with a Gaussian noise. Both situations are conceptually similar except that in the latter case we have discretizations of the usual Langevin dynamics where the Brownian motion is controlled independently by a heat bath. The noise term of Eq. (104) is thus replaced by the standard $\sqrt{2D}dW$, which obeys Einstein’s fluctuation-dissipation relation $D = \eta k_B T$, assuming $\eta = \text{const}$. The overdamped limit yields Eq. (103) with the same term but $D = k_B T$. In this setting, we have “sampling methods” where the dynamics first relax to equilibrium and then perform random excursions around a local minimum of φ . From this perspective, the stochastic optimization methods are conceptually closer to a sampling approach, however with smaller and limited control over the noise.

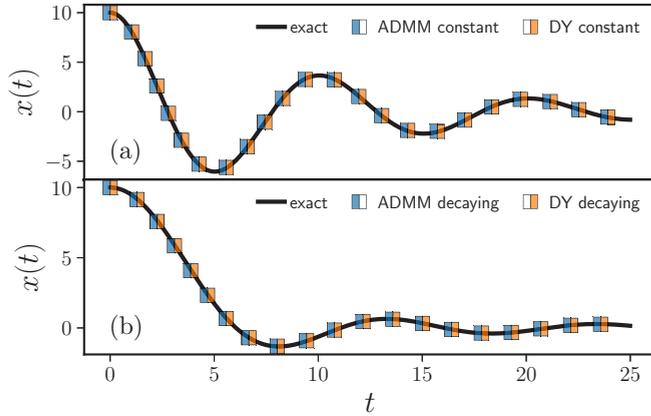


FIG. 5. Comparison of the discretization provided by Algorithms 1 and 2 versus the exact solution of the second-order gradient flow (3) with Eq. (108). (a) Constant damping with $\eta = 0.2$. (b) Decaying damping with $r = 3$. In both cases we chose step size $h = 0.01$.

IX. NUMERICAL EXPERIMENTS

A. Order of accuracy

We start by illustrating how accurate the previous optimization methods approximate the gradient flow ODEs. We consider Algorithms 1 and 2—recall that the latter reduces to (accelerated versions of) Douglas-Rachford and forward-backward as particular cases, so we are indirectly testing them as well. We compare these methods with the exact solutions (34) and (35). To this end, we choose a composite function:

$$\varphi = \varphi_1 + \varphi_2 + \varphi_3, \quad \varphi_i(x) = \omega_i^2 x^2 / 2, \quad (108)$$

with $\omega_1 = 1/2$, $\omega_2 = 1/3$, and $\omega_3 = 1/5$ —recall that there is *splitting* of these individual functions in the algorithms. As we can see in Fig. 5, these methods closely match the exact solutions. Naturally, for large step sizes there might be significant deviations. However, since these methods were proven to be first-order integrators, they are accurate up to a global error $\mathcal{O}(h)$. To verify this, for a given step size h we compute

$$\max_{k \in [0, K]} |x(t_k) - x_k| \quad (109)$$

for $K = t_{\max}/h$, where t_{\max} is a fixed simulation time (we choose $t_{\max} = 25$ as in Fig. 5). We thus compare this maximum error over the entire history of the system against a range of step sizes. Results for the constant damping case are shown in Fig. 6—we performed the same simulation with decaying damping but the plot is nearly indistinguishable from this one.

In these experiments there is no visible difference between—accelerated—Davis-Yin and ADMM. This is also explained by our theoretical results since both methods consist of numerical integrators of the same order of accuracy and to the same physical system.

B. Langevin approximation

We now wish to verify the Langevin approximation to the stochastic variants of the optimization methods introduced in the previous section. Consider thus a quadratic function as in Eq. (108), again with $\omega_1 = 1/2$ and $\omega_2 = 1/3$, but now φ_3 is

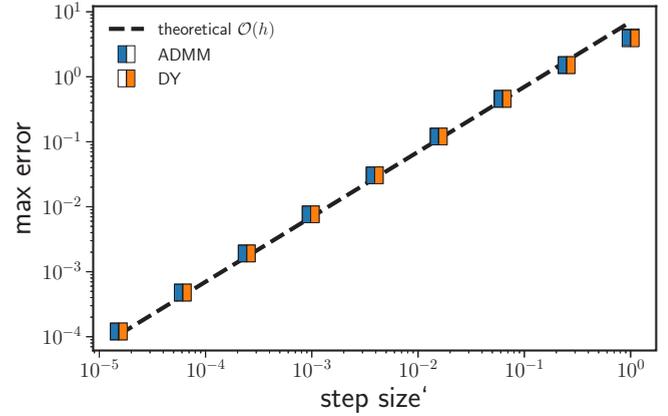


FIG. 6. Maximum error (109) over a range of step sizes. We consider constant damping under the setting of Fig. 5(a); similar results hold for decaying damping.

generated randomly so as to obtain a problem in the form of Eq. (97). We set

$$\varphi_3(x) = \frac{1}{2N} \sum_{i=1}^N \theta_i^2 x^2, \quad (110)$$

with θ_i sampled uniformly from the interval $[0,1]$. We set $N = 1000$ and φ_3 is kept fixed in all simulations, i.e., we have an “empirical mean” as in Eq. (93). To obtain stochastic gradients (94), at each iteration of Algorithms 1 and 2 we sample a minibatch of θ_i ’s of size $S = 1$ —this is the only source of stochasticity. For simplicity, we focus on the nonaccelerated methods ($\gamma_k = 0$) which are modelled by the overdamped Langevin Eq. (103). Since the function is quadratic, if we assume a constant D into Eq. (105)—which we know is not the case but it should already capture the qualitative behavior—we then have an Ornstein-Uhlenbeck process whose probability density, $P(x, t|x_0, t_0 = 0)$, is given by

$$\sqrt{\frac{\lambda}{2\pi D(1 - e^{-2\lambda t})}} \exp \left\{ -\frac{\lambda(x - x_0 e^{-\lambda t})^2}{2D(1 - e^{-2\lambda t})} \right\}, \quad (111)$$

with $\lambda = \omega_1^2 + \omega_2^2 + \frac{1}{2N} \sum_{i=1}^N \theta_i^2$. In our simulations we choose a step size $h = 0.1$, maximum time $t_{\max} = 10$, and initial position $x_0 = 10$. We consider 2000 Monte Carlo runs, with φ_3 fixed so that only the minibatch changes in each simulation. In Fig. 7 we show histograms for stochastic ADMM and Davis-Yin for a single time instant ($t = 2$). We can see that both methods give essentially the same results and agree with the out-of-equilibrium Gaussian prediction (111). Similar results hold for other time instants, as illustrated in Fig. 8. Here we set $D = 4$ to obtain the shaded gray area—this value was estimated with the simulation for $t = 1$, and we picked this time since the variance was sufficiently large. We also indicate the standard deviation of both methods (vertical red lines) which actually vary with x , however everything is well-described by Eq. (111), in agreement with the Langevin approximation. We mention that simulations with other values of the batch and step sizes were also considered, verifying similar results

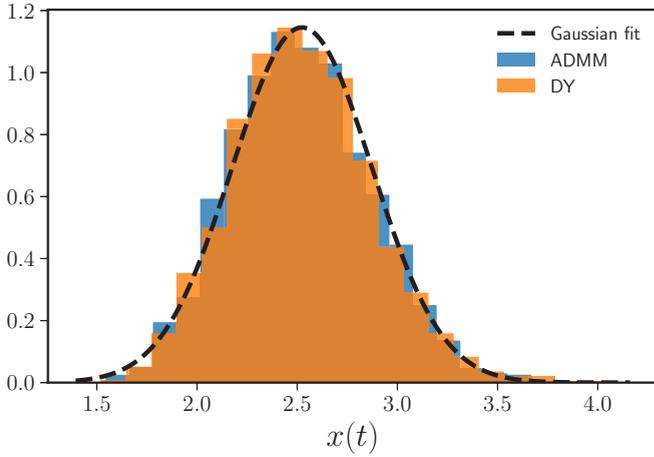


FIG. 7. Histogram of stochastic variants of Algorithms 1 and 2 (with $\gamma_k = 0$) for a quadratic problem which is modelled by an Ornstein-Uhlenbeck process; we considered Eq. (110) with $N = 1000$, minibatch of size $S = 1$, step size $h = 0.1$, initial position $x_0 = 10$, and 2000 Monte Carlo runs. We show the simulation for $t = 2$. The Gaussian fit has mean $\mu = 2.52$ and variance $\sigma = 0.35$. The theoretical prediction from Eq. (111) is $\mu = 2.51$, in close agreement.

but where the Gaussian approximation is more or less peaked according to the the scaling h/S from Eq. (106).

C. Machine learning experiments

We now wish to verify whether the accelerated methods we introduced are able to achieve faster convergence compared to the base methods, which are the actual known methods in the literature. According to the discussion in Sec. III, we expect that this might be the case. We focus on two types of damping: constant (24) and decaying (25). When nothing is specified it means that no acceleration is used, i.e., $\gamma_k = 0$ in

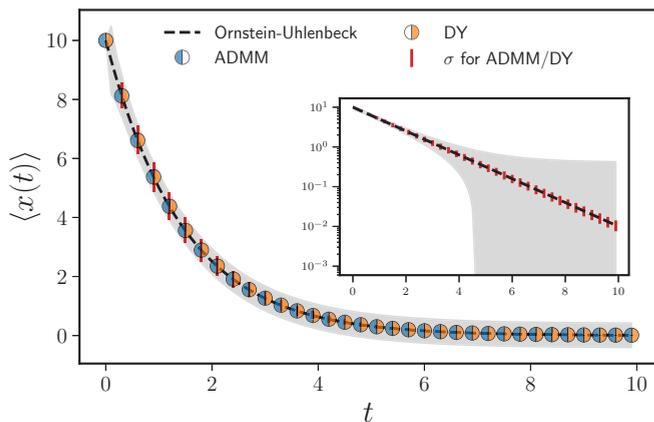


FIG. 8. Same simulation as in Fig. 7 for $t \in [0, 10]$. The dashed line is the mean of the Ornstein-Uhlenbeck process Eq. (111), and the shaded area indicates the standard deviation with $D = 4$. The markers are the means of ADMM and Davis-Yin, and the vertical red lines are standard deviations for both methods, which are very close. The inset is exactly the same plot but with markers omitted and the y axis on log scale.

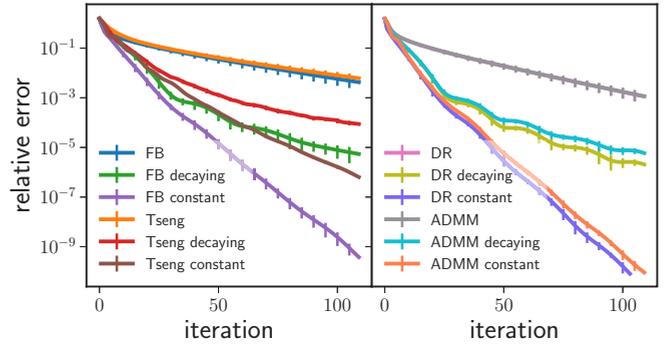


FIG. 9. Performance of several methods on problem (112). We perform 10 Monte Carlo runs and show the mean and standard deviation (error bars) of the error $|\varphi(x_k) - \varphi^*|/\varphi^*$ versus the iteration k . The accelerated variants have faster convergence.

Eq. (20), which yields the base methods. We consider ADMM (Algorithm 1), DY (Algorithm 2), and Tseng (Algorithm 3). Forward-backward (FB) is Algorithm 2 with $\varphi_1 = 0$ and Douglas-Rachford (DR) with $\varphi_3 = 0$.

Consider a *LASSO regression* problem that is of fundamental importance to machine learning and statistics:

$$\min_{x \in \mathbb{R}^n} \varphi(x), \quad \varphi(x) \equiv \frac{1}{2} \|Ax - b\|^2 + \alpha \|x\|_1, \quad (112)$$

where $A \in \mathbb{R}^{m \times n}$ is a given matrix, $b \in \mathbb{R}^m$ is a given signal, and $\alpha > 0$ is a coupling constant. Here $\|\cdot\|_1$ denotes the ℓ_1 -norm, known to induce “sparsity”—this function is not differentiable although its proximal operator has a well-known closed form solution called *soft thresholding*. Following Ref. [10, Sec. 7.1.3], we generate data by sampling $A_{ij} \sim \mathcal{N}(0, 1)$ and then normalize its columns to have unit norm. We sample $x_\bullet \in \mathbb{R}^n \sim \mathcal{N}(0, 1)$ with sparsity level 95% (only 5% of its entries are nonzero) and then add noise to obtain the observed signal $b = Ax_\bullet + e$, where $e \sim \mathcal{N}(0, 10^{-3})$. We choose dimensions $m = 500$ and $n = 2500$ —the signal-to-noise ratio is ≈ 250 , and x_\bullet has 125 nonzero entries. We set $\alpha = 0.1\alpha_{\max}$ where $\alpha_{\max} = \|A^T b\|_\infty$ is the largest value such that problem (112) admits a nontrivial solution—the factor 0.1 was verified to yield good results after few trials. We evaluate the algorithms by computing the relative error $|\varphi(x_k) - \varphi^*|/\varphi^*$ where φ^* is the solution obtained with an independent and reliable solver—we use the default implementation of CVXPY which is a standard optimization library in the Python language. For all algorithms, we choose step size $h = 0.08$, which is the largest choice such that all algorithms converge. For decaying damping we choose $r = 3$ in Eq. (25)—this choice is standard but we verified that other values did not improve—and for constant damping we choose $\eta = 0.5$ in Eq. (24)—this value was tuned with a rough grid search. In Fig. 9 we report the mean and standard deviation (small error bars) across 10 randomly generated instances of problem (112). The figure shows that the accelerated variants of each method improve over the base method. In particular, the constant damping variants are the fastest.

Next, we consider a matrix completion problem which is also of fundamental importance in machine learning. The goal is to reconstruct a *low-rank* matrix where we are only allowed to observe a few of its entries. Moreover, we assume these

entries are constrained to specified range. More precisely, suppose that for a low-rank matrix $M \in \mathbb{R}^{n \times m}$ we observe entries (i, j) that are collected in a set Ω : let $\mathcal{P}_\Omega : \mathbb{R}^{n \times m} \rightarrow \mathbb{R}^{n \times m}$ be the projection onto the support of observed entries. The observable data matrix is thus $M_{\text{obs}} = \mathcal{P}_\Omega(M)$, where $\mathcal{P}_\Omega(M)_{ij} = M_{ij}$ if $(i, j) \in \Omega$ and $\mathcal{P}_\Omega(M)_{ij} = 0$ otherwise. The goal is to estimate the missing entries of M . This can be done [35] by solving the convex problem $\min_X \|X\|_*$ such that $\mathcal{P}_\Omega(X) = \mathcal{P}_\Omega(M)$ —here $\|X\|_*$ denotes the nuclear norm. We consider a modification of this approach by imposing constraints $a \leq X_{ij} \leq b$ for given constants a and b . Specifically,

$$\min_{X \in \mathbb{R}^{n \times m}} \underbrace{\alpha \|X\|_*}_{\varphi_1} + \underbrace{\mathcal{I}_{[a,b]}(X)}_{\varphi_2} + \underbrace{\frac{1}{2} \|\mathcal{P}_\Omega(X) - \mathcal{P}_\Omega(M)\|_F^2}_{\varphi_3}, \quad (113)$$

where $\|\cdot\|_F$ denotes the Frobenius norm, $\mathcal{I}_{[a,b]}(X) = 0$ if $a \leq X_{ij} \leq b$ and ∞ otherwise. A higher $\alpha > 0$ induces lower rank solutions. This problem can be solved with Algorithms 1 and 2 with the proximal operator

$$\text{prox}_{h\|\cdot\|_*}(X) = UD_h(\Sigma)V^T, \quad (114)$$

where $X = U\Sigma V^T$ is the singular value decomposition of X and $D_h(\Sigma)_{ii} = \max\{\Sigma_{ii} - h, 0\}$ (see [35] for details). The proximal operator of φ_2 is just the projection

$$\text{prox}_{h\mathcal{I}_{[a,b]}}(X)_{ij} = \max\{a, \min\{X_{ij}, b\}\}. \quad (115)$$

Moreover, $\nabla\varphi_3(X) = \mathcal{P}_\Omega(X - M)$. In our methods we choose the following stopping criterion:

$$\|X_{k+1} - X_k\|_F / \|X_k\|_F \leq \epsilon \quad (116)$$

where ϵ is a small tolerance. To evaluate performance we report the relative error

$$\|X_k - M\|_F / \|M\|_F. \quad (117)$$

Following Ref. [35], we generate a low-rank matrix as $M = L_1 L_2^T$ where $L_1, L_2 \in \mathbb{R}^{100 \times 5}$ with entries sampled i.i.d. from $\mathcal{N}(3, 1)$. Thus, M has rank 5 (with probability one) and each entry is positive with high probability (each test instance was checked to have positive entries). We sample sn^2 entries of M uniformly, with a sampling ratio $s = 0.4$ (i.e., only 40% of the matrix M is observed). We choose

$$\begin{aligned} a &= \min\{(M_{\text{obs}})_{ij}\} - \sigma/2, \\ b &= \max\{(M_{\text{obs}})_{ij}\} + \sigma/2, \end{aligned} \quad (118)$$

where σ is the standard deviation of all entries of M_{obs} . In terms of algorithm's parameters, we choose a step size $h = 1$ for all methods, $r = 3$ for decaying damping (25), and $\eta = 0.1$ for constant damping (24)—the choice of step size is standard for this type of problem with proximal methods, while the choice of damping was obtained with a rough grid search and provided good results. In the stopping criterion (116) we choose $\epsilon = 10^{-10}$. In Fig. 10 we report the mean and standard deviation (error bars) across 10 randomly generated instances of problem (113) with $\alpha = 3.5$; all methods terminate successfully and recover a matrix with the correct rank (for this choice of α) and a final relative error of $\approx 5 \times 10^{-3}$. The number of iterations of each method to achieve the tolerance error are reported in Fig. 11.

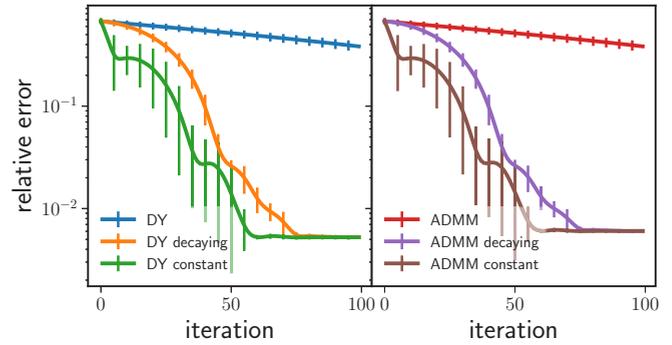


FIG. 10. Convergence of different algorithms on problem (113). We perform 10 Monte Carlo runs and indicate the mean and standard deviation (error bars) for the relative error (117) versus iteration k . Note the improvement of the accelerated methods.

To obtain more accurate solutions we consider “annealing” the parameter α . We wish to verify if the accelerated methods can still speedup convergence in this scenario. We follow the procedure of Ref. [36] which is as follows. Given a sequence $\alpha_1 > \alpha_2 > \dots > \alpha_L = \bar{\alpha} > 0$ for some $\bar{\alpha}$, we run each algorithm with α_j and then use its solution as a warm start for the solution to the next run with α_{j+1} (all other parameters are kept fixed). Starting with $\alpha_0 = \delta \|M_{\text{obs}}\|_F$ for some $\delta \in (0, 1)$ we use the schedule

$$\alpha_{j+1} = \max\{\delta\alpha_j, \bar{\alpha}\} \quad (119)$$

until reaching $\bar{\alpha}$. In our tests we choose $\delta = 0.25$ and $\bar{\alpha} = 10^{-8}$ —these parameters were tuned with a few trials and proved to yield good results. The remaining parameters are the same as those used in creating Fig. 10, except that for the constant damping variants we now use $\eta = 0.5$ —in this example, overdamping the system yield better results. We report the convergence of different methods in Fig. 12. All methods successfully reach the termination tolerance, as for the previous test, but now achieve a better reconstruction accuracy. The total number of iterations in this case are shown in Fig. 13. Note that, in this experiment, the decaying damping variants did not improve over the nonaccelerated method, but the constant

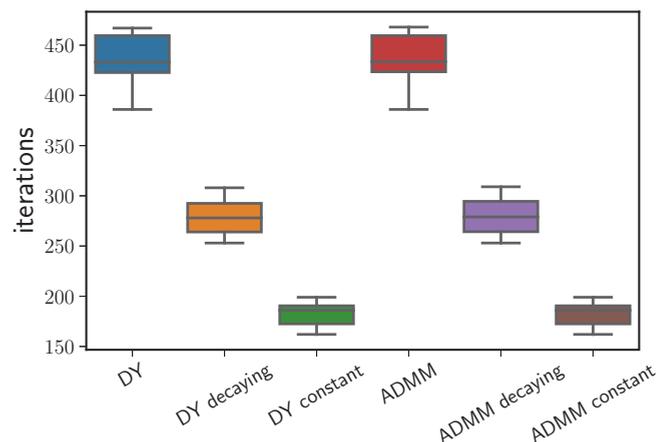


FIG. 11. Number of iterations needed to reach the termination tolerance for the problem in Fig. 10.

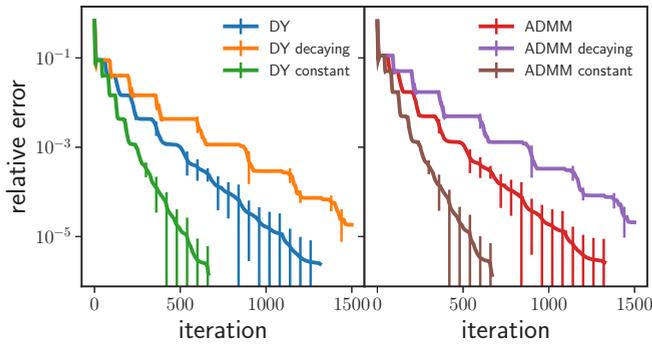


FIG. 12. Performance of different algorithms on problem (113) with annealing on α .

damping variants still provided a significant improvement. This has to do with the “hardness” of the problem, i.e., due to the choice of rank and sampling ratio the objective function should have high curvature and this is why the overdamped case is faster (recall Fig. 1).

We now provide a real world example where the goal is to reconstruct a partially observed image. We pick a gray scale image, and by a singular value decomposition we truncate its first 70 eigenvalues. The end result is the image shown in Fig. 14(a), represented by a 974×1194 matrix with entries between 0 and 1 and of low rank $r = 70$ — this image will not be used as input data but will serve as the ground truth M . Then, we sample some of its entries uniformly with sampling ratio $s = 0.3$, obtaining the image shown in Fig. 14(b)—this is the input data M_{obs} and has only 30% of the original entries. This problem is actually hard, e.g., the number of effective degrees of freedom is $r(n + m - r)/(nm) \approx 0.42$ (see Ref. [35] for details).

By solving the nonnegative matrix completion problem (113)—with $a = 0$ and $b = 1$ —our goal is to recover M from M_{obs} . As before, by running the previous algorithms a single time we recover the image shown in Fig. 14(c). The convergence rates of different methods are shown in Fig. 15—in this case we choose $\alpha = 1$, $\eta = 0.1$ for constant damping (24), $r = 3$ for decaying damping (25), step size $h = 1$, and

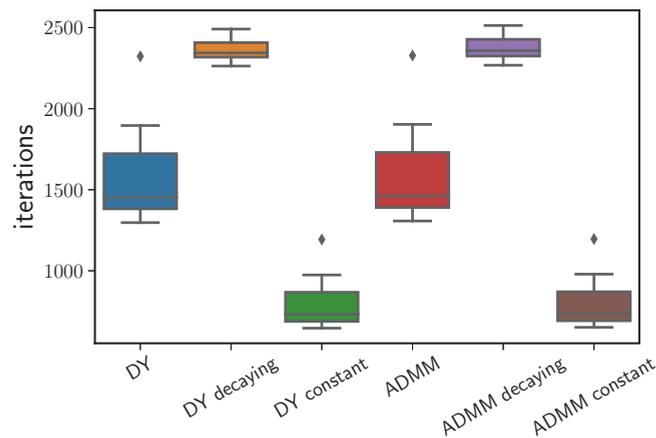


FIG. 13. Number of iterations needed to reach the termination tolerance for the problem in Fig. 12.

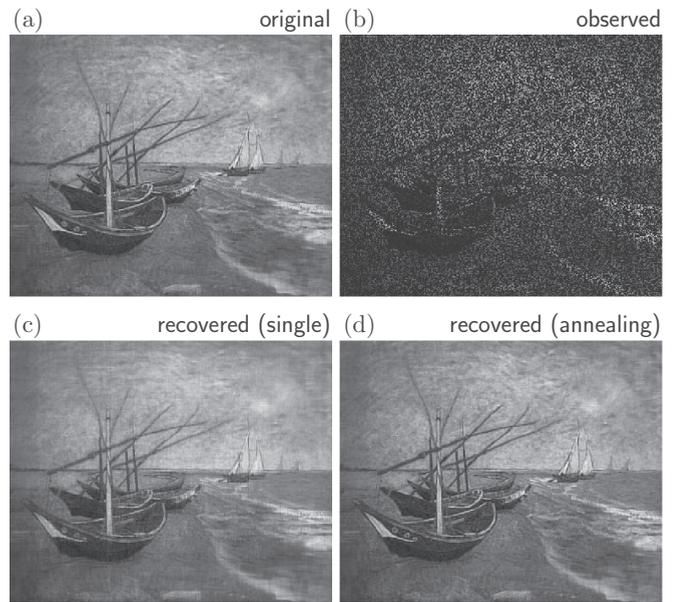


FIG. 14. (a) Image of 974×1194 pixels with rank $r = 70$. (b) Observed data with sampling ratio $s = 0.3$. (c) Reconstructed image by running the previous algorithms a single time. The relative error is 8×10^{-2} — see Fig. 15 for convergence rates. (d) Reconstructed image with annealing. The relative error is 1.6×10^{-4} —see Fig. 16 for convergence rates. The correct rank is also recovered by all methods.

tolerance $\epsilon = 10^{-6}$ in (116). Note how the accelerated variants achieve faster convergence. In this example, all methods were able to recover the correct rank.

To obtain more accurate results, we again consider annealing α according to Eq. (119)—we choose $\delta = 0.25$ and $\bar{\alpha} = 10^{-4}$, which was the smallest feasible value we found through several trials. An example of the recovered image—with the accelerated variants since the base methods were unable to achieve such a low error—is shown in Fig. 14(d). By a closer inspection one can verify that Fig. 14(d) is actually better than Fig. 14(c), e.g., has a better resolution. The convergence rates of each method under this setting are shown in Fig. 16. In this example, both accelerated variants—i.e.,

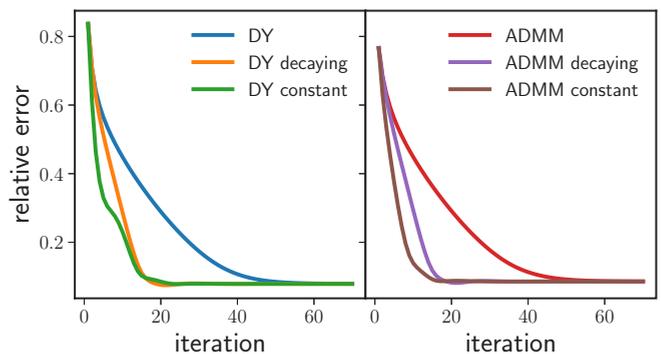


FIG. 15. Convergence rates of different methods when reconstructing the observed image in Fig. 14(b) by solving the optimization problem (113). An example of the recovered image is in Fig. 14(c).

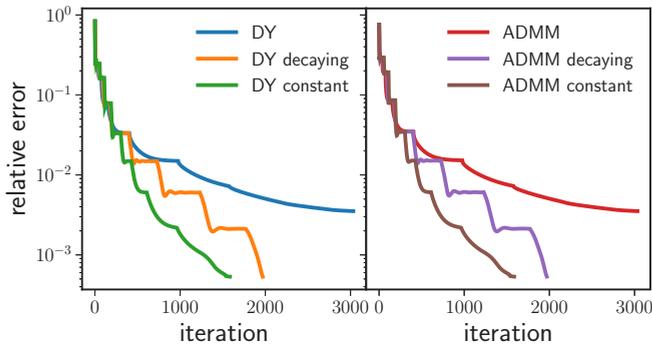


FIG. 16. Convergence rates for the same problem as in Fig. 15 but with annealing on α . An example of the recovered image with the accelerated methods is shown in Fig. 14(d)—note that the nonaccelerated methods were unable to achieve such a small error.

constant and decaying damping—were able to significantly improve over the base methods. Since this is a harder problem compared to the case of Fig. 12, the objective function should have less curvature thus underdamping the system yield better results.

X. DISCUSSION AND OUTLOOK

The main goal of this paper was to provide a unified perspective on standard optimization methods from a physics standpoint; these results are summarized in Fig. 4. We showed that there is indeed a unifying principle behind several well-known, but also new, optimization methods. The most general situation can be described by an underdamped Langevin equation, which accounts for accelerated methods in a stochastic setting. The randomness comes from a stochastic approximation to gradients or proximal operators—this has similarities to a Gaussian perturbation from a heat bath in standard Langevin dynamics. When the minibatch is sufficiently large, i.e., the “temperature” $T \sim S^{-1} \rightarrow 0$, we recover the deterministic setting that is described by a dissipative form of Newton’s equation. In addition, in the high friction limit $\eta \rightarrow \infty$, the dynamics becomes nonaccelerated and thus described by an overdamped Langevin or by a first-order gradient flow, in the stochastic and deterministic regimes, respectively. All the optimization methods we discussed arise as discretizations of such systems. Thus, the fact that many seemingly unrelated optimization algorithms consist of actual *simulations* of the same dissipative physical system is quite surprising given that most of these methods were originally introduced from a completely independent approach, and without any connection to physics.

Using ODE splitting ideas and implicit discretizations, we introduced several accelerated and stochastic generalizations of the most important proximal algorithms in the literature. Moreover, important gradient-based methods, such as gradient descent, heavy ball, and Nesterov, also fit our framework. In these cases there is no splitting of the objective function and they correspond to explicit discretizations (see the Appendix). Interestingly, the heavy ball method turns out to be a structure-preserving (i.e., conformal symplectic) integrator, while Nesterov introduces a spurious dissipation; see Ref. [21]

for details. Let us also point out that recently we generalized symplectic integrators to general dissipative Hamiltonian systems [33], which may offer a systematic approach to construct new gradient-based optimization methods based on simulations of physical systems.

The results of this paper explain why some different optimization methods may actually behave quite similarly in practice; the reason is because they are numerical integrators to the same physical system and have the same order of accuracy. Therefore, at this level, it is not possible to distinguish between different methods. However, the connections we established allow one to perform *backward error analysis*, which would result into modified ODEs that capture how a particular discretization perturbs the original system—we carried out such an analysis in the case of heavy ball and Nesterov [21]. It would thus be interesting to consider a similar analysis to some of the proximal methods previously discussed. This would bring refined insights and a better understanding of these methods. Another interesting question concerns the choice of damping. We considered constant and decaying damping since these cases are well-understood, e.g., we know some convergence rates and exact solutions in the quadratic case, as described in Sec. III. However, other choices are possible and establishing the “optimal” damping strategy for a given problem class is an interesting and nontrivial problem.

It is worth mentioning that the connections between stochastic optimization and Langevin dynamics offers an interesting opportunity for bringing techniques from perturbation theory and nonequilibrium statistical mechanics into machine learning. For instance, any Fokker-Planck equation is equivalent to a path integral through the Martin-Siggia-Rose formalism [37]; this is a good starting point to apply field theory techniques to such problems. However, the optimization methods we considered may be applied to numerical problems in physics. For instance, there is a close connection between statistical mechanics and combinatorial optimization problems, such as finding ground states of spin glasses [38,39]. The proximal methods considered in this paper can be adapted to tackle combinatorial problems such as max-cut and graph partitioning, which are somehow equivalent. Notably, in the context of machine learning and signal processing, proximal methods have been applied to large-scale problems with great success and proved to be quite efficient. It seems interesting to consider whether such an approach can lead to scalable and cheaper methods for disordered systems compared to standard Monte Carlo approaches.

Finally, it is thus clear that standard optimization methods are nothing but simulations of classical dissipative systems. A natural question concerns extending this perspective to the quantum level. Can we construct quantum optimization algorithms by simulating dissipative quantum systems? Open quantum systems is still a controversial topic, but they may offer an alternative to quantum annealing, whose discretizations may yield general quantum optimization algorithms. Thus, in the same way that a classical computer is employed to implement standard—i.e., classical—optimization algorithms, or equivalently to simulate a classical dissipative systems, perhaps a quantum computer may be necessary to simulate a dissipative quantum system [40].

ACKNOWLEDGMENTS

G.F. thanks Michael I. Jordan for helpful discussions and Patrick Johnstone for comments on an earlier draft of this paper. We also thank the anonymous referees for insightful comments that helped to improve the paper. This work was supported by Grants ARO MURI No. W911NF-17-1-0304, NSF No. 2031985, and NSF No. 1934931.

APPENDIX: STOCHASTIC GRADIENT DESCENT, HEAVY BALL, AND NESTEROV

Since stochastic versions of gradient descent, heavy ball, and Nesterov's method are widely used in machine learning, here we show how these methods arise from the previous physical systems.

Let us start with gradient descent. An explicit Euler-Maruyama discretization of the Langevin Eq. (103) yields

$$\begin{aligned} x_{k+1} &= x_k - h\nabla\varphi(x_k) + \frac{h}{\sqrt{S}}C(x_k)\epsilon_k \\ &= x_k - h\widetilde{\nabla}\varphi(x_k), \end{aligned} \quad (\text{A1})$$

where in the last passage we used Eqs. (102) and (95). This is nothing but the well-known *stochastic gradient descent* (SGD) method, widely used in training neural networks. When $S \rightarrow \infty$, i.e., the stochastic gradient becomes the true gradient, $\widetilde{\nabla}\varphi \rightarrow \nabla\varphi$, we recover the deterministic gradient descent.

Let us now consider Nesterov's method. Writing the underdamped Langevin Eq. (104) as

$$\ddot{x} = -\nabla\varphi(x) - \eta(t)\dot{x} + \sqrt{\frac{h}{S}}C(x)\dot{W}, \quad (\text{A2})$$

where this should be understood in the Itô sense, we discretize it with the help of Eqs. (21), (102), and (95) to obtain

$$\begin{aligned} \frac{x_{k+1} - \hat{x}_k}{h^2} &= -\nabla\varphi(\hat{x}_k) + \frac{1}{\sqrt{S}}C(\hat{x}_k)\epsilon_k \\ &= -\widetilde{\nabla}\varphi(\hat{x}_k). \end{aligned} \quad (\text{A3})$$

Recalling the definition Eq. (20), and redefining the step size as usual, $h^2 \mapsto h$, we obtain

$$x_{k+1} = \hat{x}_k - h\widetilde{\nabla}\varphi(\hat{x}_k), \quad (\text{A4a})$$

$$\hat{x}_{k+1} = x_{k+1} + \gamma_{k+1}(x_{k+1} - x_k). \quad (\text{A4b})$$

This is a stochastic version of Nesterov's method—the original method is recovered in the deterministic case where $S \rightarrow \infty$.

Let us now consider Polyak's heavy ball method. From a physics point of view this case is more interesting since the deterministic method was recently shown by us [33] to be a conformal symplectic integrator. Consider the underdamped Langevin Eq. (104). We write the second equation as

$$\frac{d}{dt}(e^{\theta(t)}p) = -e^{\theta(t)}\nabla\varphi(x) + e^{\theta(t)}\sqrt{\frac{h}{S}}C(x)\dot{W}, \quad (\text{A5})$$

where $\dot{\theta}(t) \equiv \eta(t)$. Integrating this from t_k to $t_{k+1} = t_k + h$, and keeping terms up to $\mathcal{O}(h^2)$, we get

$$\begin{aligned} p_{k+1} &= e^{-h\eta_k}p_k - h\nabla\varphi(x_k) + \frac{h}{\sqrt{S}}C(x_k)\epsilon_k \\ &= e^{-h\eta_k}p_k - h\widetilde{\nabla}\varphi(x_k). \end{aligned} \quad (\text{A6})$$

The first equation in Eq. (104) gives

$$x_{k+1} = x_k + hp_{k+1} + \mathcal{O}(h^2). \quad (\text{A7})$$

Now we define the following variables:

$$\mu_k \equiv e^{-h\eta_k}, \quad v_k \equiv hp_k, \quad h^2 \mapsto h. \quad (\text{A8})$$

This allows us to write Eqs. (A6) and (A7) as

$$v_{k+1} = \mu_k v_k - h\widetilde{\nabla}\varphi(x_k), \quad (\text{A9a})$$

$$x_{k+1} = x_k + v_{k+1}. \quad (\text{A9b})$$

This is precisely a stochastic version of the heavy ball method, often referred to as momentum method, or SGD with momentum, in deep learning literature. Usually this method is used with a constant $\eta_k = \eta$, in which case $\mu \equiv e^{-h\eta} \in (0, 1]$ is known as the “momentum factor.” When $S \rightarrow \infty$, i.e., the stochastic gradient becomes the true gradient, the above recovers the deterministic heavy ball method. Therefore, once again, the scheme of Fig. 4 also apply to gradient-based methods. Compared to the proximal methods previously discussed, the difference in these cases is that the gradient is always computed explicitly as opposed to implicitly.

[1] A. Cauchy, C. R. Hebd. Seances Acad. Sci. **25**, 536 (1847).
 [2] J. Douglas and H. H. Rachford, *Trans. Am. Math. Soc.* **82**, 421 (1956).
 [3] R. Glowinski and A. Marroco, *ESAIM: Math. Model. Numer. Anal.* **9**, 41 (1975).
 [4] D. Gabay and B. Mercier, *Comput. Math. Appl.* **2**, 17 (1976).
 [5] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, *Found. Trends Mach. Learn.* **3**, 1 (2011).
 [6] P. L. Lions and B. Mercier, *SIAM J. Numer. Anal.* **16**, 964 (1979).
 [7] P. Tseng, *SIAM J. Control Optim.* **38**, 431 (2000).

[8] D. Davis and W. Yin, *Set-Valued Variation. Anal.* **25**, 829 (2017).
 [9] E. K. Ryu and S. Boyd, *Appl. Comput. Math.* **15**, 3 (2016).
 [10] N. Parikh and S. Boyd, *Found. Trends Optim.* **1**, 127 (2014).
 [11] F. E. Browder, *Bull. Amer. Math. Soc.* **69**, 862 (1963).
 [12] F. E. Browder, *Duke Math J.* **30**, 557 (1963).
 [13] F. E. Browder, *Proc. Nat. Acad. Sci. USA* **50**, 31 (1963).
 [14] G. J. Minty, *Duke Math. J.* **29**, 341 (1962).
 [15] H. H. Bauschke and P. L. Combettes, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces* (Springer International Publishing, Berlin, 2017).

- [16] E. Zeidler, *Nonlinear Functional Analysis and Its Applications, II/B: Nonlinear Monotone Operators* (Springer-Verlag, Berlin, 1990).
- [17] B. T. Polyak, *USSR Comp. Math. Math. Phys.* **4**, 1 (1964).
- [18] Y. Nesterov, *Sov. Math. Doklady* **27**, 372 (1983).
- [19] W. Su, S. Boyd, and E. J. Candès, *J. Mach. Learn. Res.* **17**, 1 (2016).
- [20] A. Wibisono, A. C. Wilson, and M. I. Jordan, *Proc. Nat. Acad. Sci. USA* **113**, E7351 (2016).
- [21] G. França, J. Sulam, D. P. Robinson, and R. Vidal, *J. Stat. Mech.* (2020) 124008.
- [22] P. Caldirola, *Nuovo Cim.* **18**, 393 (1941).
- [23] E. Kanai, *Prog. Theor. Phys.* **3**, 440 (1948).
- [24] A. O. Caldeira and A. J. Leggett, *Phys. Rev. Lett.* **46**, 211 (1981).
- [25] R. I. McLachlan and G. R. W. Quispel, *Acta Numerica* **11**, 341 (2002).
- [26] R. L. Speth, W. H. Green, S. MacNamara, and G. Strang, *SIAM J. Numer. Anal.* **51**, 3084 (2013).
- [27] G. França, D. P. Robinson, and R. Vidal, *Proceedings of the International Conference on Machine Learning* (ACM, New York, NY, 2018).
- [28] G. França, D. P. Robinson, and R. Vidal, [arXiv:1808.04048](https://arxiv.org/abs/1808.04048).
- [29] E. Hairer, C. Lubich, and G. Wanner, *Geometric Numerical Integration* (Springer, Berlin, 2006).
- [30] T. Goldstein, B. O'Donoghue, S. Setzer, and R. Baraniuk, *SIAM J. Imag. Sci.* **7**, 1588 (2014).
- [31] P. Patrinos, L. Stella, and A. Bemporad, in *Proceedings of the 53rd IEEE Conference on Decision and Control* (IEEE, 2014), pp. 4234–4239.
- [32] A. Beck and M. Teboulle, *SIAM J. Imag. Sci.* **2**, 183 (2009).
- [33] G. França, M. I. Jordan, and R. Vidal, *J. Stat. Mech.* (2021) 043402.
- [34] H. Robbins and S. Monro, *Ann. Math. Stat.* **22**, 400 (1951).
- [35] J. Cai, E. Candès, and Z. Shen, *SIAM J. Optim.* **20**, 1956 (2010).
- [36] S. Ma, D. Goldfarb, and L. Chen, *Math. Prog.* **128**, 321 (2011).
- [37] J. Zinn-Justin, *Quantum Field Theory and Critical Phenomena* (Clarendon Press, Oxford, UK, 2002).
- [38] M. Mezard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987).
- [39] H. Nishimori, *Statistical Physics of Spin Glasses and Information Processing: An Introduction* (Clarendon Press, Oxford, UK, 2001).
- [40] R. P. Feynman, *Int. J. Theor. Phys.* **21**, 467 (1982).