
Differential Inclusions for Modeling Nonsmooth ADMM Variants: A Continuous Limit Theory

Huizhuo Yuan¹ Yuren Zhou² Chris Junchi Li³ Qingyun Sun⁴

Abstract

Recently, there has been a great deal of research attention on understanding the convergence behavior of first-order methods. One line of this research focuses on analyzing the convergence behavior of first-order methods using tools from continuous dynamical systems such as ordinary differential equations and differential inclusions. These research results shed lights on better understanding first-order methods from a non-optimization point of view. The alternating direction method of multipliers (ADMM) is a widely used first-order method for solving optimization problems arising from machine learning and statistics, and it is important to investigate its behavior using these new techniques from dynamical systems. Existing works along this line have been mainly focusing on problems with smooth objective functions, which exclude many important applications that are traditionally solved by ADMM variants. In this paper, we analyze some well-known and widely used ADMM variants for nonsmooth optimization problems using tools of differential inclusions. In particular, we analyze the convergence behavior of linearized ADMM, gradient-based ADMM, generalized ADMM and accelerated generalized ADMM for nonsmooth problems and show their connections with dynamical systems. We anticipate that these results will provide new insights on understanding ADMM for solving nonsmooth problems.

1. Introduction

Recently, there has been tremendous interests in using continuous-time dynamical system tools to analyze first-

*Equal contribution ¹Peking University ²Duke University
³Tencent AI Lab ⁴Stanford University. Correspondence to: Chris Junchi Li <junchi.li.duke@gmail.com>.

order optimization algorithms such as Nesterov's accelerated gradient method (AGM) (Nesterov, 1983) and its variants. In the seminal work Su et al. (2016), the authors designed a differential equation for modeling AGM, and analyzed the connection between the solution of the differential equation and the continuous limit of the iterates of AGM. Their work provided new insights on understanding the convergence behavior of AGM. Later investigations along this line mainly focused on analyzing AGM and its variants such as FISTA and heavy ball method using the tools of ordinary differential equations, differential inclusions, and more generally, continuous dynamical systems (see, e.g., Shi et al. (2018); Wibisono et al. (2016); Wilson et al. (2016); Krichene et al. (2015); An et al. (2018); Zhou et al. (2017)). Very recently, França et al. (2018a;b) made a significant step towards understanding the alternating direction method of multipliers (ADMM) using the tools from continuous dynamical systems. ADMM is now a widely used algorithm for solving problems with separable structures, which include a lot of important applications arising from image processing, signal processing, machine learning, statistics, etc. It has a close connection with some classical operator-splitting methods in numerical PDEs such as Douglas-Rachford (Douglas & Rachford, 1956) and Peaceman-Rachford (Peaceman & Rachford, 1955) operator-splitting methods that dated back to the 1950s. These operator-splitting methods were later studied in Gabay & Mercier (1975); Glowinski & Marroco (1975); Gabay (1983); Fortin & Glowinski (1983); Glowinski & Le Tallec (1989); Eckstein & Bertsekas (1992). But the renaissance of ADMM was due to several works in 2007-2008 that introduced this algorithm to solving signal processing and image processing problems (Combettes & Pesquet, 2007; Goldstein & Osher, 2009; Wang et al., 2008; Mardani et al., 2018). Since then, ADMM was successfully used for solving important applications in many areas in science and engineering. The popularity and importance of ADMM has been partly demonstrated by the recognition of the highly influential survey paper Boyd et al. (2011). As a result, the works of França et.al. (França et al., 2018a;b) are very timely and important as they provided new tools for further understanding the convergence behavior of this influential algorithm.

However, one major drawback of França et al. (2018a;b) is that they assume that the objective function is smooth, which in fact rules out most of the applications solved by ADMM and its variants. More specifically, França et al. (2018a;b) consider the following problem:

$$\underset{x \in \mathbb{R}^d}{\text{minimize}} \quad f(x) + g(Ax), \quad (1)$$

where $A \in \mathbb{R}^{m \times d}$, $f : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$, $g : \mathbb{R}^m \rightarrow \mathbb{R} \cup \{+\infty\}$, and both f and g are convex and differentiable. We need to point out that under such assumptions, there exist many other efficient algorithms for solving (1) and ADMM may not be a good choice.

In this paper, we allow f and g to be nonsmooth functions in (1). To apply ADMM, one standard technique is to rewrite (1) as

$$\begin{aligned} & \underset{x \in \mathbb{R}^d, z \in \mathbb{R}^m}{\text{minimize}} && f(x) + g(z) \\ & \text{subject to} && Ax - z = 0. \end{aligned} \quad (2)$$

Note that the classical setting of linear constraint $Ax + Bz = c$ can be reformulated as $z = Ax$ by a simple linear transformation operation.

One typical iteration of ADMM for solving (2) is

$$x_{k+1} := \underset{x}{\text{argmin}} \left[f(x) + \frac{\rho}{2} \|Ax - z_k + u_k\|_2^2 \right], \quad (3a)$$

$$z_{k+1} := \underset{z}{\text{argmin}} \left[g(z) + \frac{\rho}{2} \|Ax_{k+1} - z + u_k\|_2^2 \right], \quad (3b)$$

$$u_{k+1} := u_k + Ax_{k+1} - z_{k+1}, \quad (3c)$$

where u is the (scaled) Lagrange multiplier and $\rho > 0$ is a penalty parameter in the augmented Lagrangian function:

$$\mathcal{L}_\rho(x, z, u) = f(x) + g(z) + \rho u^\top (Ax - z) + \frac{\rho}{2} \|Ax - z\|_2^2. \quad (4)$$

Throughout this paper, we assume that both f and g are proper closed convex functions.

1.1. Our Contributions

We extend the analysis of França et al. (2018a;b) to nonsmooth cases using a completely different tool: differential inclusion, which is motivated by the analysis of FISTA by Vassilis et al. (2018). More specifically, we analyze the convergence rate of continuous limit of two widely used ADMM variants and two generalizations for nonsmooth problems: linearized ADMM, gradient-based ADMM, generalized ADMM and accelerated generalized ADMM. We anticipate that these results will provide new insights on understanding ADMM for solving nonsmooth problems.

1.2. Notations

Let $F, f, g : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ be function mappings from \mathbb{R}^d to the extended real space $\mathbb{R} \cup \{+\infty\}$. Let $x_k \in$

\mathbb{R}^d , $\{z_k, u_k, \hat{z}_k, \hat{u}_k\} \in \mathbb{R}^m$ be solution sequences of the ADMM iterates, while the capital letters $X(t), Z(t), U(t)$ denote the solutions of the continuous dynamical systems. Let $\langle \cdot, \cdot \rangle$ be the inner product, and let $\|\cdot\|_2$ and $\|\cdot\|_1$ be the L_2 and L_1 norms, respectively. Let τ_L and τ_G be parameters controlling the quadratic penalties in the linearized and gradient-based cases, respectively. We use k to denote the discrete timestep and t denotes the continuous time. Other notations will be explained at their first entrances.

2. Continuous Limit of Linearized ADMM and Gradient-Based ADMM

We need to point out that the ADMM given in (3) is rarely used in practice, because for most applications, the x -subproblem does *not* have closed-form solutions and an iterative solver is still needed to solve it. Note that although the z -subproblem in (3) corresponds to the proximal mapping of function g , the x -subproblem does *not* correspond to the proximal mapping of f because of the presence of matrix A . Moreover, it is possible that in some applications, f does *not* have an easy proximal mapping. Two most commonly used nonsmooth ADMM variants in practice are linearized ADMM and gradient-based ADMM, and they are suitable for the following two cases, which cover most applications of ADMM:

- Case (i): *Linearized ADMM* is suitable for the case where f is nonsmooth with easy proximal mappings; one representative application in this case is the Lasso problem where $f(x) = \|x\|_1$ and $g(z) = \frac{1}{2}\|z - b\|_2^2$ (Tibshirani, 1996).¹
- Case (ii): *Gradient-based ADMM* is suitable for the case where g is nonsmooth with easy proximal mapping, f is differentiable but does not have an easy proximal mapping; one representative application in this case is the sparse logistic regression problem where g is the ℓ_1 norm and f is the logistic loss function (Liu et al., 2009). Note that $A = I$ in this particular application.

We now provide more details about the applicability of linearized ADMM and gradient-based ADMM.

In Case (i), where both f and g are nonsmooth with easy proximal mappings, the z -subproblem (3b) corresponds to the proximal mapping of g and is thus easy to solve; however the presence of matrix A brings difficulty to solving the x -subproblem (3a). Linearized ADMM addresses this issue by adding a suitably chosen proximal

¹Both f and g can be nonsmooth in this case, so we could also write $f = \frac{1}{2}\|\cdot - b\|_2^2$ and $g = \|\cdot\|_1$. We represent f and g the other way around to follow the notations in Yang & Zhang (2011) as an illustration.

term $\frac{1}{2}\|x - x_k\|_{\tau_L I - \rho A^\top A}$ to the objective function of (3a),² which results in the following subproblem whose solution corresponds to the proximal mapping of f :

$$\underset{x \in \mathbb{R}^d}{\text{minimize}} \quad f(x) + \frac{\tau_L}{2} \left\| x - \left(x_k - \frac{\rho}{\tau_L} A^\top (Ax_k - z_k + u_k) \right) \right\|_2^2,$$

where $1/\tau_L$ can be viewed as the step size of the gradient step of the quadratic penalty. Notice that by the above equation we are making a first-order Taylor approximation to the second term of (3a) to avoid trouble caused by matrix A . Combining this subproblem with (3b) and (3c) yields the linearized ADMM. Here we consider a slightly more general version of linearized ADMM by adding a relaxation term to the intermediate residual $Ax_{k+1} - z_k$, which is summarized in (5):

$$x_{k+1} = \underset{x}{\text{argmin}} \left\{ f(x) + \frac{\tau_L}{2} \left\| x - \left(x_k - \frac{\rho}{\tau_L} A^\top (Ax_k - z_k + u_k) \right) \right\|_2^2 \right\}, \quad (5a)$$

$$z_{k+1} = \underset{z}{\text{argmin}} \left\{ g(z) + \frac{\rho}{2} \|\alpha Ax_{k+1} + (1 - \alpha)z_k - z + u_k\|_2^2 \right\}, \quad (5b)$$

$$u_{k+1} = u_k + (\alpha Ax_{k+1} + (1 - \alpha)z_k - z_{k+1}), \quad (5c)$$

where $\alpha \in (0, 2)$ is a relaxation parameter, and when $\alpha = 1$, it reduces to the classical linearized ADMM. As a widely used nonsmooth ADMM variant, linearized ADMM has been studied by many researchers, (see, e.g., [Chen & Teboulle \(1994\)](#); [Eckstein \(1994\)](#); [He et al. \(2002\)](#); [Zhang et al. \(2010\)](#); [Yang & Zhang \(2011\)](#); [Lin et al. \(2011\)](#); [Ma \(2016\)](#); [Xu \(2015\)](#); [Yang & Yuan \(2013\)](#); [Ouyang et al. \(2015\)](#)). The difference between (3b)-(3c) and (5b)-(5c) is that Ax_{k+1} is replaced by $\alpha Ax_{k+1} + (1 - \alpha)z_k$. This is called *relaxation*, which has been suggested in many papers (see, e.g., [Eckstein & Bertsekas \(1992\)](#)) to provide more flexibility and potentially improve the convergence speed of the algorithm.

We now use the total variation minimization problem ([Rudin et al., 1992](#)) as an example to show how (5) works for a particular problem. The total variation minimization problem can be casted as the following form after variable splitting,

$$\begin{aligned} & \underset{x, z}{\text{minimize}} \quad \frac{1}{2} \|x - b\|_2^2 + \lambda \|z\|_1 \\ & \text{subject to} \quad z = Dx, \end{aligned} \quad (6)$$

which is in the form of (1) with $A = D$, $f(\cdot) = \frac{1}{2} \|\cdot - b\|_2^2$ and $g(\cdot) = \lambda \|\cdot\|_1$. When linearized ADMM (5) is applied to solve (6), the two subproblems (5a) and (5b) respectively

correspond to the proximal mappings of $\|\cdot\|_1$ and $\frac{1}{2} \|\cdot - b\|_2^2$, which are both very easy to compute.

In Case (ii), gradient-based ADMM is suitable for the case where f does *not* have an easy proximal mapping. In this case, the linearized ADMM (5) is not a good choice, because the x -subproblem (5a) is still not easy to solve. As a result, the gradient-based ADMM is proposed to address this issue. A typical iteration of gradient-based ADMM is as follows:

$$x_{k+1} = x_k - \frac{1}{\tau_G} \left(\nabla f(x_k) + \rho A^\top (Ax_k - z_k + u_k) \right), \quad (7a)$$

$$z_{k+1} = \underset{z}{\text{argmin}} \left\{ g(z) + \frac{\rho}{2} \|\alpha Ax_{k+1} + (1 - \alpha)z_k - z + u_k\|_2^2 \right\}, \quad (7b)$$

$$u_{k+1} = u_k + (\alpha Ax_{k+1} + (1 - \alpha)z_k - z_{k+1}), \quad (7c)$$

where we again used a relaxation term to make the algorithm more general, and $1/\tau_G$ is the step size of the negative gradient step taken on the augmented Lagrangian function $\mathcal{L}_\rho(x_k, z_k, u_k)$. Note that since we assume that f is differentiable in this case, the objective function of the x -subproblem in (3a) becomes a differentiable function. As a result, gradient-based ADMM suggests that a gradient step is taken instead of minimizing the augmented Lagrangian function directly, which results in the new x -subproblem in (7a). This gradient-based ADMM has been studied in the literature extensively, (see, e.g., [Condat \(2013\)](#); [Vu \(2013\)](#); [Davis & Yin \(2017\)](#); [Lin et al. \(2017\)](#)).

We now use sparse logistic regression ([Koh et al., 2007](#)) as an example to show how (7) works for a particular problem. The sparse logistic regression problem can be casted as

$$\underset{x}{\text{minimize}} \quad \frac{1}{N} \sum_{i=1}^N \log(1 + \exp(-b_i(a_i^\top x + v))) + \lambda \|x\|_1, \quad (8)$$

which is in the form of (1) with f being the logistic loss function, g being the ℓ_1 norm and $A = I$. Note that the logistic loss function f does not adopt an easy proximal mapping, but it is differentiable and thus a gradient step can be taken for the x -subproblem. When the gradient-based ADMM (7) is applied to solve (8), the two subproblems (7a) and (7b) are both easy to be implemented.

2.1. Main Results

In this subsection, we present the main results on the convergence of the continuous limit of the iterates of linearized ADMM (5) and gradient-based ADMM (7). We focus on the continuous approximation results when $t = \rho^{-1}k$ and $\rho \rightarrow \infty$, where t denotes the time and k is the iteration counter. The following definitions and assumption are needed for our results.

²The norm $\|x\|_{\tau_L I - \rho A^\top A}$ is defined as $x^\top (\tau_L I - \rho A^\top A)x$.

Definition 1. (Nesterov, 2004). A vector v is called a subgradient of the function f at a point x_0 satisfying $f(x_0) < \infty$ if for any y satisfying $f(y) < \infty$, we have

$$f(y) \geq f(x_0) + v^\top (y - x_0).$$

The set of all subgradients of f at x_0 , $\partial f(x_0)$, is called the subdifferential of the function f at the point x_0 .

Definition 2. A function f defined on \mathbb{R}^d and taking values in $\mathbb{R} \cup \{+\infty\}$ is called closed, if its epigraph

$$\{(x, t) \in \mathbb{R}^d \times \mathbb{R} \mid f(x) \leq t\}$$

is a closed set in $\mathbb{R}^d \times \mathbb{R}$.

Definition 3. A function f defined on \mathbb{R}^d and taking values in $\mathbb{R} \cup \{+\infty\}$ is called convex, if for every $x, y \in \mathbb{R}^d$ and $\theta \in (0, 1)$,

$$f(\theta x + (1 - \theta)y) \leq \theta f(x) + (1 - \theta)f(y).$$

Assumption 4. We assume that f and g are closed convex functions taking values in $\mathbb{R} \cup \{+\infty\}$ and properly defined over \mathbb{R}^d (in the sense that $f(x) < +\infty$ for at least one x), and $F(x) := f(x) + g(Ax)$ has a global minimum. We assume further that matrix A has full column rank and singular values $\sigma_1 \geq \dots \geq \sigma_d > 0$.

Assumption 4 implies that f, g, F are proper, lower semi-continuous and subdifferentiable. (Nesterov, 2004)

For linearized ADMM, we have the following approximation theorem.

Theorem 5. Let Assumption 4 hold, and the relaxation parameter $\alpha \in (0, 2)$. Rescaling the time by setting $t = \rho^{-1}k$, the continuous-time limit of the iterates $\{x_k\}$ of linearized ADMM (5) as $\rho \rightarrow \infty$ and $\tau_L/\rho \rightarrow c \in (0, \infty)$ is given by the differential inclusion

$$0 \in \partial F(X(t)) + \left(cI + \frac{1-\alpha}{\alpha} A^\top A \right) \dot{X}(t), \quad (9)$$

with initial value $X(0) = x_0$.

Similarly, we have the following Theorem 6 for the continuous limit of iterates of gradient-based ADMM (7).

Theorem 6. Let Assumption 4 hold, and the relaxation parameter $\alpha \in (0, 2)$. In addition, we assume that f is smooth. Rescaling the time by setting $t = \rho^{-1}k$, the continuous-time limit of the iterates $\{x_k\}$ of gradient-based ADMM (7) as $\rho \rightarrow \infty$ and $\tau_G/\rho \rightarrow c \in (0, \infty)$ is given by the same differential inclusion (9) in Theorem 5 with initial value $X(0) = x_0$.

Remark 7. Here we assume mildly that the solution of the differential inclusion (9) exists and is unique. The existence and uniqueness of a solution are usually dealt with classical

monotone theory (Vassilis et al., 2018). However in some cases, the uniqueness of a solution does not necessarily hold. Theorem 1 in Aubin (1984) provides an additional assumption that can ensure the uniqueness of a solution within the context of this work.³

The next theorem shows the convergence rate of the continuous-time limit of the iterates $\{x_k\}$ generated by linearized ADMM (5) and gradient-based ADMM (7) in the convex case. We use x^* to denote an arbitrary minimizer of F . We recall that σ_1 is the largest singular value of A and σ_d is the smallest singular value of A , and assume that $cI + \frac{1-\alpha}{\alpha} A^\top A$ is a well defined positive definite matrix. For simplicity of notations, we define κ_1 and κ_d the largest and smallest singular value of $cI + \frac{1-\alpha}{\alpha} A^\top A$, respectively.

Theorem 8. Let Assumption 4 hold. Assume that c and α are chosen such that the matrix $(cI + \frac{1-\alpha}{\alpha} A^\top A)$ is positive definite, with largest and smallest eigenvalues being κ_1^2, κ_d^2 ($\kappa_1, \kappa_d > 0$). Let $X(t)$ be any shock solution⁴ of the differential inclusion (9) with convex objective function F and initial value $X(0) = x_0$. Then $X(t)$ has bounded trajectory and the function value gap has $\mathcal{O}(t^{-1})$ convergence rate almost everywhere, i.e., for a minimizer x^* and a.e. $t > 0$ it holds that $\|X(t) - x^*\|_2 \leq \frac{\kappa_1}{\kappa_d} \|x_0 - x^*\|_2$, and

$$F(X(t)) - F(x^*) \leq \frac{\kappa_1^2 \|x_0 - x^*\|_2^2}{2t}.$$

Moreover, we have

$$\int_0^{+\infty} [F(X(t)) - F(x^*)] \leq \frac{\kappa_1^2}{2} \|x_0 - x^*\|_2^2,$$

and

$$\int_0^{+\infty} t \|\dot{X}(t)\|_2^2 \leq \frac{\kappa_1^2}{2\kappa_d^2} \|x_0 - x^*\|_2^2.$$

The convergence rate of $\mathcal{O}(t^{-1})$ in the continuous case is the same as the $\mathcal{O}(k^{-1})$ rate in the discrete case. We notice that the rate depends heavily on α : when $0 < \alpha \leq 1$, $\kappa_1 = \sqrt{c + \frac{1-\alpha}{\alpha} \sigma_1^2}$ and $\kappa_d = \sqrt{c + \frac{1-\alpha}{\alpha} \sigma_d^2}$, and when $1 < \alpha < 2$, $\kappa_1 = \sqrt{c + \frac{1-\alpha}{\alpha} \sigma_d^2}$ and $\kappa_d = \sqrt{c + \frac{1-\alpha}{\alpha} \sigma_1^2}$, where σ_1, σ_d are singular value of matrix A .

3. Continuous Limit of Generalized ADMM

In this section, we study the continuous limit of the generalized ADMM (G-ADMM) proposed by Eckstein & Bertsekas

³ For conditions that guarantee the existence and uniqueness of the solution to differential inclusion, we refer the readers to Adly et al. (2006); Attouch et al. (2002); Paoli (2000).

⁴ The concept of shock solution is given in Attouch et al. (2002); Paoli (2000); Vassilis et al. (2018). Since the goal of this paper is to illustrate the main idea, we include its definition and existence result in Appendix A for completeness.

(1992). We point out that this has been studied by França et al. (2018a;b) for the cases where f and g are both smooth. We now extend the analysis to problems with nonsmooth f and g . G-ADMM allows more flexibility of ADMM by introducing a new relaxation parameter $\alpha \in (0, 2)$, and it updates the iterates as

$$x_{k+1} = \operatorname{argmin}_x \left\{ f(x) + \frac{\rho}{2} \|Ax - z_k + u_k\|_2^2 \right\}, \quad (10a)$$

$$z_{k+1} = \operatorname{argmin}_z \left\{ g(z) + \frac{\rho}{2} \left\| \alpha Ax_{k+1} + (1 - \alpha)z_k - z + u_k \right\|_2^2 \right\}, \quad (10b)$$

$$u_{k+1} = u_k + (\alpha Ax_{k+1} + (1 - \alpha)z_k - z_{k+1}). \quad (10c)$$

3.1. Differential Inclusion for G-ADMM

Following França et al. (2018a), we rescale the time by a factor of ρ^{-1} , i.e. let $t = \rho^{-1}k$, and obtain the continuous-time approximation. One can see from (10a) that larger parameter $\rho > 0$ results in smaller-pace updates of x_k . We study the limit of updates in the regime $\rho \rightarrow \infty$. Our main result on the differential inclusion approximation of G-ADMM is as follows:

Theorem 9. *Let Assumption 4 hold, and the relaxation parameter $\alpha \in (0, 2)$. Rescale the time by setting $t = \rho^{-1}k$. The continuous limit of iterates of $\{x_k\}$ in Algorithm (14) as $\rho \rightarrow \infty$ is given by the following differential inclusion:*

$$\frac{1}{\alpha} (A^\top A) \dot{X}(t) + \partial F(X(t)) \ni 0, \quad (11)$$

with $X(0) = x_0$.

We move forward and analyze the convergence property of differential inclusion (11). Recall that σ_1 and σ_d are the largest and smallest singular values of A , respectively.

Theorem 10. *When the Assumption 4 holds, the shock solution $X(t)$ of differential inclusion (11) has bounded trajectory and $\mathcal{O}(t^{-1})$ convergence rate almost everywhere, i.e., for a.e. $t \geq 0$, $\|X(t) - x^*\|_2 \leq \frac{\sigma_1}{\sigma_d} \|x_0 - x^*\|_2$, and*

$$F(X(t)) - F(x^*) \leq \frac{\sigma_1^2 \|x_0 - x^*\|_2^2}{2\alpha t}.$$

Moreover,

$$\int_0^{+\infty} [F(X(t)) - F(x^*)] dt \leq \frac{\sigma_1^2}{2\alpha} \|x_0 - x^*\|_2^2,$$

$$\int_0^{+\infty} t \|\dot{X}(t)\|_2^2 dt \leq \frac{\sigma_1^2}{2\sigma_d^2} \|x_0 - x^*\|_2^2.$$

Here, the key idea of the convergence analysis of differential inclusion (11) is to use a sequence of approximating differential equations (ADE) that approaches the differential inclusion.

3.2. Differential Inclusion for Accelerated G-ADMM

Goldstein et al. (2014) proposed an accelerated ADMM by incorporating Nesterov's extrapolation technique. (Nesterov, 1983) This method is generalized by França et al. (2018a;b) which jointly consider relaxation and acceleration. The accelerated G-ADMM considered in França et al. (2018a;b) uses $\gamma_{k+1} = \frac{k}{k+r}$ as the momentum coefficient (r being a positive constant), and is described as follows:

$$x_{k+1} = \operatorname{argmin}_x \left\{ f(x) + \frac{\rho}{2} \|Ax - \hat{z}_k + \hat{u}_k\|_2^2 \right\}, \quad (12a)$$

$$z_{k+1} = \operatorname{argmin}_z \left\{ g(z) + \frac{\rho}{2} \left\| \alpha Ax_{k+1} + (1 - \alpha)\hat{z}_k - z + \hat{u}_k \right\|_2^2 \right\}, \quad (12b)$$

$$u_{k+1} = \hat{u}_k + (\alpha Ax_{k+1} + (1 - \alpha)\hat{z}_k - z_{k+1}), \quad (12c)$$

$$\hat{u}_{k+1} = u_{k+1} + \gamma_{k+1}(u_{k+1} - u_k), \quad (12d)$$

$$\hat{z}_{k+1} = z_{k+1} + \gamma_{k+1}(z_{k+1} - z_k), \quad (12e)$$

Here we extend the results in França et al. (2018a;b) to the case where f and g are nonsmooth functions. We present our first main result in the following theorem:

Theorem 11. *Let Assumption 4 hold, and let the relaxation parameter $\alpha \in (0, 2)$. Rescale the time by setting $t = \rho^{-1/2}k$, the continuous-time approximation of Algorithm (12) iteration as $\rho \rightarrow \infty$ is given by the differential inclusion*

$$\frac{1}{\alpha} (A^\top A) \left(\ddot{X}(t) + \frac{r}{t} \dot{X}(t) \right) + \partial F(X(t)) \ni 0, \quad (13)$$

with $X(t_0) = x_0$ and $\dot{X}(t_0) = 0$. Here t_0 is an arbitrary positive starting time.

Unlike all previous time rescaling scheme where $k = \rho t$, for accelerated G-ADMM the time rescaling is $k = \rho^{1/2}t$, which is in accordance with the idea of *acceleration*. r/t is the damping ratio of differential inclusion (13). We refer to the case $r \geq 3$ as high friction case, $0 < r < 3$ as low friction case and provide two separate convergence theorems under these two cases. Moreover, we specifically point out that the convergence rate can be *sharper* when $r > 3$.

In the following, we define $\Delta_0^2 = \max\{t_0^2(F(x_0) - F(x^*)), \|x_0 - x^*\|_2^2\}$. We also remark that all the factors C used later are functions of the constant r , the relaxation parameter α , and singular values σ_1 and σ_d .

Firstly, we show that $X(t)$ has almost surely bounded trajectory for $t \geq t_0$, and the convergence rate is $\mathcal{O}(t^{-2})$ in terms of both $F(X(t)) - F(x^*)$ and $\|\dot{X}(t)\|_2^2$.

Theorem 12 (High Friction). *When $r \geq 3$, the shock solution $X(t)$ of Differential Inclusion (13) has bounded trajectory and $\mathcal{O}(t^{-2})$ convergence rate almost everywhere, i.e.*

there exists positive factors C_1, C_2, C_3 depending on r, α such that, for a.e. $t \geq t_0$, $\|X(t) - x^*\|_2 \leq C_1 \Delta_0$, and

$$F(X(t)) - F(x^*) \leq \frac{C_2 \Delta_0^2}{t^2}, \quad \|\dot{X}(t)\|_2 \leq \frac{C_3 \Delta_0}{t}.$$

When $r > 3$, there exist positive factors C_4, C_5 depending on r, α such that

$$\int_{t_0}^{\infty} t(F(X(t)) - F(x^*)) dt \leq C_4 \Delta_0^2, \\ \int_{t_0}^{\infty} t \|\dot{X}(t)\|_2^2 dt \leq C_5 \Delta_0^2.$$

Using the bounded integration result in Theorem 12, we could show that, when $r > 3$, the convergence rate is actually $o(t^{-2})$:

Remark 13. *The convergence rate in Theorem 12 is not sharp when $r > 3$, in the sense that $F(X(t)) - F(x^*) = o(t^{-2})$ and $\|\dot{X}(t)\|_2 = o(t^{-1})$ as $t \rightarrow \infty$, i.e. for all $r > 3$,*

$$\lim_{t \rightarrow \infty} t^2(F(X(t)) - F(x^*)) = 0, \quad \lim_{t \rightarrow \infty} t \|\dot{X}(t)\|_2 = 0.$$

The constant r has the so-called *magic number* 3, as discussed in details by Su et al. (2016), in the sense that an $\mathcal{O}(t^{-2})$ convergence rate could be guaranteed for AGM in high friction case $r \geq 3$, but not for the low friction case $0 < r < 3$. The follow-up works Attouch et al. (2017), Attouch et al. (2018), Vassilis et al. (2018), França et al. (2018a) derive the $\mathcal{O}(t^{-2r/3})$ convergence rate for AGM under low friction case and extend $\mathcal{O}(t^{-2})$ convergence rate from AGM to accelerated G-ADMM in high friction case, respectively.

For nonsmooth and low friction case, we show a $\mathcal{O}(t^{-2r/3})$ convergence rate for $F(X(t)) - F(x^*)$, and for $\|\dot{X}(t)\|_2^2$ if the trajectory is almost surely bounded.

Theorem 14 (Low Friction). *When $0 < r < 3$, the shock solution $X(t)$ of Differential Inclusion (13) has $\mathcal{O}(t^{-2r/3})$ convergence rate almost everywhere, i.e. there exists positive factor C_6 depending on r, α such that, for a.e. $t \geq t_0$,*

$$F(X(t)) - F(x^*) \leq \frac{C_6 t_0^{-2(3-r)/3} \Delta_0^2}{t^{2r/3}}.$$

If in addition the trajectory $\{X(t)\}_{t \geq t_0}$ is bounded almost everywhere for $t \geq t_0$, then there also exists some positive factor C_7 depending on r, α such that for a.e. $t \geq t_0$,

$$\|\dot{X}(t)\|_2 \leq \frac{C_7 t_0^{-(3-r)/3} \Delta_0}{t^{r/3}}.$$

Theorems 12 and 14 provide convergence results for accelerated G-ADMM in continuous-time scheme with the second-order differential inclusion and accompanying tools, which sheds light on the discrete-time algorithm (Su et al., 2016).

4. Numerical Experiments

According to the convergence theorems in previous sections, both the linearized ADMM and the gradient ADMM share the same differential inclusion (9). In the following two examples, we show that for $\rho = 10$ the trajectory of the linearized or gradient-based ADMM algorithm is very close to the trajectory of the corresponding differential inclusion.

To plot the trajectory of the differential inclusion numerically, we use entropic approximation Teboulle (1992) of the nonsmooth objective to compute the trajectory of the differential inclusion, we remark that entropic approximation is parallel in theory to Moreau-Yosida approximation, which use quadratic approximation. Entropic approximation would provide a smooth approximation of the sub-gradient for $\|z\|_1$ as $\tanh \beta z$, where we choose $\beta = 10^6$ in the following experiments, so that the numerical approximation error is of the order 10^{-6} .

4.1. Total Variation Minimization with Linearized ADMM

Consider the total variation minimization problem (6), see Rudin et al. (1992). Using variable splitting, we can write the problem as

$$\begin{aligned} \underset{x, z}{\text{minimize}} \quad & \frac{1}{2} \|x - b\|_2^2 + \lambda \|z\|_1 \\ \text{subject to} \quad & z = Dx, \end{aligned}$$

where D is the finite difference matrix, $x, z \in \mathbb{R}^n$. This problem fits to the general framework of ADMM with $A = D$, $f(x) = \frac{1}{2} \|x - b\|_2^2$ and $g(z) = \lambda \|z\|_1$, since f is quadratic with easy proximal and g is nonsmooth with easy proximal, namely, its proximal operator is the soft thresholding operator $S_\lambda(z) = \text{sign}(z)(|z| - \lambda)_+$, we could use linearized ADMM to solve this problem.

We generated a problem instance, where the true signal x^{true} is a piece-wise constant signal, the observation is $b = x^{\text{true}} + n$, $n \sim \mathcal{N}(0, I)$. We solve the total variation minimization with $\lambda = 0.01$ to recover the true signal x^{true} from b , following the MATLAB examples of the paper Boyd et al. (2011).

In the following, we run both the linearized ADMM algorithm with $\rho = 10$ and the differential inclusion. We demonstrate the trajectory of linearized ADMM algorithm and the differential inclusion for different parameter configuration α and c . First of all, we fix $c = 10$, and vary α from 2^{-3} to 2; then, we fix $\alpha = 1.6$, and vary c from 1 to 32. Figures 1 are the trajectory of linearized ADMM with $\rho = 10$ and the corresponding differential inclusion, with the x axis being the iteration count, the y axis being the relative error $\frac{F(x_k) - F(x^*)}{F(x_0) - F(x^*)}$, where $F(x^*)$ is the function value at optimal solution, $F(x_0)$ is the function value at initializa-




Figure 1. On total variation minimization problem, the plots are the trajectory of linearized ADMM with $\rho = 10$ and the corresponding differential inclusion, the first plot is for different α from 2^{-3} to 2 when $c = 10$, second plot is for different c from 1 to 32 when $\alpha = 1.6$.




Figure 2. On sparse logistic regression, the plots are gradient ADMM and the differential inclusion when $\rho = 10$, first plot is for different α from 2^{-3} to 2 when $c = 10$, second plot is for different c from 1 to 32 when $\alpha = 1.6$.




Figure 3. On Lasso problem, the plots are the trajectory of G-ADMM with $\rho = 50$ and the corresponding differential inclusion. We can see that the differential inclusion matches the trajectory of the linearized ADMM algorithm very closely for all parameter settings.

4.2. Sparse Logistic Regression with Gradient ADMM

Consider the sparse logistic regression problem (8) (see Koh et al. (2007) for more details). Using variable splitting, we can write the problem as

$$\begin{aligned} & \underset{x, z}{\text{minimize}} && \frac{1}{N} \sum_{i=1}^N \log(1 + \exp(-b_i(a_i^\top x + v))) + \lambda \|z\|_1 \\ & \text{subject to} && z = x, \end{aligned}$$

with variable $x \in \mathbb{R}^{d-1}$, $v \in \mathbb{R}$.

This problem fits to the general framework with variable $\bar{x} = (x, v)$, $A = I$, $f(\bar{x}) = \log(1 + \exp(-b_i(a_i^\top x + v)))$ and $g(\bar{x}) = \lambda \|\bar{x}_{1:n}\|_1$, since f is differentiable but does not

have an easy proximal mapping and g is nonsmooth with easy proximal, we could use gradient ADMM to solve this problem.

We generated a problem instance following the MATLAB examples of the paper Boyd et al. (2011). More specifically, we chose a true weight vector x^{true} sampled from Bernoulli-Gaussian distribution with mean 0, variance 1 and sparsity level 0.1, along with the true intercept v^{true} sampled from standard normal. Each feature vector a_i was generated from Bernoulli-Gaussian distribution at sparsity level 0.2. The labels were then generated using $b_i = \text{sign}(a_i^\top x^{\text{true}} + v^{\text{true}} + \nu_i)$, where $\nu_i \sim \mathcal{N}(0, 0.1)$. The regularization parameter is set to $\lambda = 0.1 \lambda_{\max}$ according to Koh et al. (2007), where $\lambda_{\max} = \|A^\top \tilde{b}\|_\infty$ is the critical value above which the solution of the problem is $x = 0$, where \tilde{b} is defined in page 93 of Boyd et al. (2011).

We run both the gradient ADMM algorithm with $\rho = 10$ and the differential inclusion. We demonstrate the trajectory of gradient ADMM algorithm and the differential inclusion for different parameter configuration α and c . First of all, we fix $c = 10$, and vary α from 2^{-3} to 2; then, we fix $\alpha = 1.6$, and vary c from 1 to 32. Figures 2 are the trajectory of gradient ADMM with $\rho = 10$ and the corresponding differential inclusion, with the x axis being the iteration count, the y axis being the relative error $\frac{F(x_k) - F(x^*)}{F(x_0) - F(x^*)}$, where $F(x^*)$ is the function value at optimal solution, $F(x_0)$ is the function value at initialization. We can see that the differential inclusion matches the trajectory of the gradient

ADMM algorithm very closely for all parameter settings.

4.3. Lasso with G-ADMM

Additionally, we show that the trajectory of G-ADMM algorithm with different α is close to the trajectory of the differential inclusion in Lasso example Tibshirani (1996).

The Lasso problem can be casted as

$$\begin{aligned} & \underset{x,z}{\text{minimize}} && \frac{1}{2}\|Ax - b\|_2^2 + \lambda\|z\|_1 \\ & \text{subject to} && z = x, \end{aligned}$$

which fits to the general framework with $f(\cdot) = \frac{1}{2}\|A \cdot - b\|_2^2$ and $g(\cdot) = \lambda\|\cdot\|_1$. When G-ADMM (14) is applied to solve this problem, the two subproblems (10a) and (10b) respectively correspond to the proximal mappings of $\frac{1}{2}\|A \cdot - b\|_2^2$ and $\|\cdot\|_1$, which are both very easy to compute.

We generated a problem instance following the MATLAB examples of the paper Boyd et al. (2011). Specifically, we sample true sparse signal x_0 from Bernoulli-Gaussian distribution with mean 0, variance 1 and sparsity level 0.05, A is sampled from Gaussian random matrix of size 100 by 400 with columns norm normalized to one, $b = Ax_0 + v$, where $v \sim \mathcal{N}(0, 0.001)$. The regularization parameter is set to $\lambda = 0.1\lambda_{\max}$ according to Koh et al. (2007), where $\lambda_{\max} = \|A^\top b\|_\infty$ is the critical value above which the solution of the problem is $x = 0$.

We run both the G-ADMM algorithm with $\rho = 50$ and the differential inclusion. We vary α from 1 to 1.8 as the range that people typically use for G-ADMM(Boyd et al., 2011). Figure 3 is the trajectory of G-ADMM with $\rho = 50$ and the corresponding differential inclusion, with the x axis being the iteration count, the y axis being the relative error $\frac{F(x_k) - F(x^*)}{F(x_0) - F(x^*)}$, where $F(x^*)$ is the function value at optimal solution, $F(x_0)$ is the function value at initialization. We can see that the differential inclusion matches the trajectory of the linearized ADMM algorithm very closely for all parameter settings.

5. Conclusions

In this paper, we analyzed the convergence behavior of the continuous limits of some widely used nonsmooth ADMM variants: linearized ADMM, gradient-based ADMM, as well as G-ADMM and its Nesterov's acceleration. Such continuous limits are characterized by the tool of differential inclusion and promote understandings of these ADMM variants from the angles of dynamical systems. Our novel continuous-time convergence theorems characterize these ADMM variants, which is further supported by experimental results. The differential inclusion for linearized and gradient ADMM variants suggests that we could choose the algo-

rithmic parameters via a principled approach that uses the condition number of the matrix, which serves as a practical guidance learned from theoretical insights.

References

- Adly, S., Attouch, H., & Cabot, A. (2006). Finite time stabilization of nonlinear oscillators subject to dry friction. In *Nonsmooth mechanics and analysis* (pp. 289–304). Springer.
- Afriat, S. (1971). Theory of maxima and the method of lagrange. *SIAM Journal on Applied Mathematics*, 20(3), 343–357.
- An, W., Wang, H., Sun, Q., Xu, J., Dai, Q., & Zhang, L. (2018). A pid controller approach for stochastic optimization of deep networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 8522–8531).
- Attouch, H., Cabot, A., & Redont, P. (2002). The dynamics of elastic shocks via epigraphical regularization of a differential inclusion barrier and penalty approximations. *Adv. Math. Sci. Appl.*, 12, 273–306.
- Attouch, H., Chbani, Z., Peypouquet, J., & Redont, P. (2018). Fast convergence of inertial dynamics and algorithms with asymptotic vanishing viscosity. *Mathematical Programming*, 168(1-2), 123–175.
- Attouch, H., Chbani, Z., & Riahi, H. (2017). Rate of convergence of the nesterov accelerated gradient method in the subcritical case $\alpha \leq 3$. *arXiv preprint arXiv:1706.05671*.
- Aubin, J.-P. (1984). Cellina. *Differential Inclusions*, 264.
- Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J., et al. (2011). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine learning*, 3(1), 1–122.
- Chen, G. & Teboulle, M. (1994). A proximal-based decomposition method for convex minimization problems. *Mathematical Programming*, 64, 81–101.
- Combettes, P. L. & Pesquet, J.-C. (2007). A Douglas-Rachford splitting approach to nonsmooth convex variational signal recovery. *IEEE Journal of Selected Topics in Signal Processing*, 1(4), 564–574.
- Condat, L. (2013). A primal-dual splitting method for convex optimization involving Lipschitzian, proximable and linear composite terms. *J. Optimization Theory and Applications*, 158(2), 460–479.
- Davis, D. & Yin, W. (2017). A three-operator splitting scheme and its optimization applications. *Set-Valued and Variational Analysis*, 25(4), 829–858.

- Douglas, J. & Rachford, H. H. (1956). On the numerical solution of the heat conduction problem in 2 and 3 space variables. *Transactions of the American Mathematical Society*, 82, 421–439.
- Eckstein, J. (1994). Some saddle-function splitting methods for convex programming. *Optimization Methods and Software*, 4(1), 75–83.
- Eckstein, J. & Bertsekas, D. P. (1992). On the Douglas–Rachford splitting method and the proximal point algorithm for maximal monotone operators. *Mathematical Programming*, 55(1-3), 293–318.
- Fortin, M. & Glowinski, R. (1983). *Augmented Lagrangian methods: applications to the numerical solution of boundary-value problems*. North-Holland Pub. Co.
- França, G., Robinson, D. P., & Vidal, R. (2018a). ADMM and accelerated ADMM as continuous dynamical systems. In *Proceedings of the 35th International Conference on Machine Learning* (pp. 1559–1567).
- França, G., Robinson, D. P., & Vidal, R. (2018b). Relax, and accelerate: A continuous perspective on ADMM. *arXiv preprint arXiv:1808.04048*.
- Gabay, D. (1983). Applications of the method of multipliers to variational inequalities. In M. Fortin & R. Glowinski (Eds.), *Augmented Lagrangian Methods: Applications to the Solution of Boundary Value Problems*. Amsterdam: North-Holland.
- Gabay, D. & Mercier, B. (1975). *A dual algorithm for the solution of non linear variational problems via finite element approximation*. Institut de recherche d’informatique et d’automatique.
- Glowinski, R. & Le Tallec, P. (1989). *Augmented Lagrangian and Operator-Splitting Methods in Nonlinear Mechanics*. Philadelphia, Pennsylvania: SIAM.
- Glowinski, R. & Marroco, A. (1975). Sur l’approximation, par éléments finis d’ordre un, et la résolution, par pénalisation-dualité d’une classe de problèmes de Dirichlet non linéaires. *Revue française d’automatique, informatique, recherche opérationnelle. Analyse numérique*, 9(R2), 41–76.
- Goldstein, T., O’Donoghue, B., Setzer, S., & Baraniuk, R. (2014). Fast alternating direction optimization methods. *SIAM Journal on Imaging Sciences*, 7(3), 1588–1623.
- Goldstein, T. & Osher, S. (2009). The split Bregman method for L1-regularized problems. *SIAM J. Imaging Sci.*, 2, 323–343.
- He, B. S., Liao, L., Han, D., & Yang, H. (2002). A new inexact alternating direction method for monotone variational inequalities. *Mathematical Programming*, 92, 103–118.
- Koh, K., Kim, S.-J., & Boyd, S. (2007). An interior-point method for large-scale L1-regularized logistic regression. *Journal of Machine Learning Research*, 8(Jul), 1519–1555.
- Krichene, W., Bayen, A., & Bartlett, P. L. (2015). Accelerated mirror descent in continuous and discrete time. In *NIPS*.
- Lin, T., Ma, S., & Zhang, S. (2017). An extragradient-based alternating direction method for convex minimization. *Foundations of Computational Mathematics*, 17(1), 35–59.
- Lin, Z., Liu, R., & Su, Z. (2011). Linearized alternating direction method with adaptive penalty for low-rank representation. In *Advances in neural information processing systems* (pp. 612–620).
- Liu, J., Chen, J., & Ye, J. (2009). Large-scale sparse logistic regression. In *SIGKDD*.
- Ma, S. (2016). Alternating proximal gradient method for convex minimization. *Journal of Scientific Computing*, 68(2), 546–572.
- Mardani, M., Sun, Q., Donoho, D., Pappas, V., Monajemi, H., Vasanawala, S., & Pauly, J. (2018). Neural proximal gradient descent for compressive imaging. In *Advances in Neural Information Processing Systems* (pp. 9596–9606).
- Nesterov, Y. (1983). A method for unconstrained convex minimization problem with the rate of convergence $O(1/k^2)$. In *Soviet Mathematics Doklady*, volume 27 (pp. 372–376).
- Nesterov, Y. (2004). *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media.
- Ouyang, Y., Chen, Y., Lan, G., & Pong, E. (2015). An accelerated linearized alternating direction method of multipliers. *SIAM Journal on Imaging Sciences*, 8(1), 644–681.
- Paoli, L. A. (2000). An existence result for vibrations with unilateral constraints: case of a nonsmooth set of constraints. *Math. Models Methods Appl. Sci.*, 10, 815–831.
- Peaceman, D. H. & Rachford, H. H. (1955). The numerical solution of parabolic elliptic differential equations. *SIAM Journal on Applied Mathematics*, 3, 28–41.
- Rudin, L. I., Osher, S., & Fatemi, E. (1992). Nonlinear total variation based noise removal algorithms. *Physica D: nonlinear phenomena*, 60(1-4), 259–268.

- Shi, B., Du, S. S., Jordan, M. I., & Su, W. J. (2018). Understanding the acceleration phenomenon via high-resolution differential equations. *arXiv preprint arXiv:1810.08907*.
- Su, W., Boyd, S., & Candes, E. J. (2016). A differential equation for modeling Nesterov's accelerated gradient method: theory and insights. *Journal of Machine Learning Research*, 17(153), 1–43.
- Takayama, A. (1985). *Mathematical economics*. Cambridge University Press.
- Teboulle, M. (1992). Entropic proximal mappings with applications to nonlinear programming. *Mathematics of Operations Research*, 17(3), 670–690.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, (pp. 267–288).
- Vassilis, A., Jean-François, A., & Charles, D. (2018). The differential inclusion modeling FISTA algorithm and optimality of convergence rate in the case $b \leq 3$. *SIAM Journal on Optimization*, 28(1), 551–574.
- Vu, B. C. (2013). A splitting algorithm for dual monotone inclusions involving cocoercive operators. *Advances in Computational Mathematics*, 38(3), 667–681.
- Wang, Y., Yang, J., Yin, W., & Zhang, Y. (2008). A new alternating minimization algorithm for total variation image reconstruction. *SIAM Journal on Imaging Sciences*, 1(3), 248–272.
- Wibisono, A., Wilson, A. C., & Jordan, M. I. (2016). A variational perspective on accelerated methods in optimization. *Proceedings of the National Academy of Sciences*, (pp. 201614734).
- Wilson, A. C., Recht, B., & Jordan, M. I. (2016). A Lyapunov analysis of momentum methods in optimization. *arXiv preprint arXiv:1611.02635*.
- Xu, Y. (2015). Alternating proximal gradient method for sparse nonnegative Tucker decomposition. *Mathematical Programming Computation*, 7(1), 39–70.
- Yang, J. & Yuan, X. (2013). Linearized augmented Lagrangian and alternating direction methods for nuclear norm minimization. *Mathematics of Computation*, 82(281), 301–329.
- Yang, J. & Zhang, Y. (2011). Alternating direction algorithms for ℓ_1 -problems in compressive sensing. *SIAM journal on scientific computing*, 33(1), 250–278.
- Zhang, X., Burger, M., Bresson, X., & Osher, S. (2010). Bregmanized nonlocal regularization for deconvolution and sparse reconstruction. *SIAM Journal on Imaging Science*, 3, 253–276.
- Zhou, Z., Mertikopoulos, P., Bambos, N., Boyd, S., & Glynn, P. (2017). Mirror descent in non-convex stochastic programming. *arXiv preprint arXiv:1706.05681*.