Nr. 89
29. Jul. 2022

universität innsbruck
Institut für Mathematik

Preprint-Series: Department of Mathematics - Applied Mathematics

Convergence analysis of critical point regularization with non-convex regularizers

D. Obmann, M. Haltmeier

AppliedMathematics

Technikerstraße 13 - 6020 Innsbruck - Austria
Tel.: +43 512 507 53803 Fax: +43 512 507 53898
https://applied-math.uibk.ac.at

Convergence analysis of critical point regularization with non-convex regularizers

Daniel Obmann and Markus Haltmeier

Department of Mathematics, University of Innsbruck
Technikerstrasse 13, 6020 Innsbruck, Austria
{daniel.obmann, markus.haltmeier}@uibk.ac.at

August 1, 2022

Abstract

In recent years, several methods using regularizers defined by neural networks as penalty terms in variational methods have been developed. One of the key assumptions in the stability and convergence analysis of these methods is the ability of finding global minimizers. However, such an assumption is often not feasible when the regularizer is a black box or non-convex, making the search for global minimizers of the involved Tikhonov functional a challenging task. Instead, standard minimization schemes are applied which typically only guarantee that a critical point is found. To address this issue, in this paper we study stability and convergence properties of critical points of Tikhonov functionals with a possible non-convex regularizer. To this end, we introduce the concept of relative sub-differentiability and study its basic properties. Based on this concept, we develop a convergence analysis assuming relative sub-differentiability of the regularizer. For the case where the noise level tends to zero, we derive a limiting problem representing first-order optimality conditions of a related restricted optimization problem. Finally, we provide numerical simulations that support our theoretical findings and the need for the sort of analysis that we provide in this paper.

Keywords: Inverse problems, regularization, critical points, stability guarantees, learned regularizer, non-convex regularizer, neural networks, variational methods

1 Introduction

In various scientific fields and applications, such as medical imaging or remote sensing, it is often not possible to obtain the desired quantity of interest directly. Assuming a linear measurement model, recovering the quantity of interest requires solving an inverse problem of the form

$$y^\delta = \mathbf{K}x + \eta^\delta \,, \tag{1.1}$$

where $\mathbf{K}\colon \mathbb{X} \to \mathbb{Y}$ is a linear operator between Hilbert spaces modeling the forward problem, η^δ is the data perturbation, $y^\delta \in \mathbb{Y}$ is the noisy data and $x \in \mathbb{X}$ is the sought for signal. In many cases these problems are ill-posed, meaning that no continuous right inverse of the operator $\mathbf{K}$ exists. To overcome such issues several established approaches for the stable approximation of solutions of inverse problems exist.

1.1 Regularization with non-convex penalties

Particularly popular regularization techniques are variational methods [9, 18]. These methods recover regularized solutions x_α^δ as global minimizers of the Tikhonov functional

$$\mathcal{T}_{\alpha,y^\delta}(x) := \frac{1}{2}\|\mathbf{K}x - y^\delta\|^2 + \alpha\mathcal{R}(x)\,. \tag{1.2}$$

Here, $\mathcal{R}$ is a regularizer which encodes prior information about the desired solution and $\frac{1}{2}\|\mathbf{K}x - y^\delta\|^2$ plays the role of a data-discrepancy measure. Classically, regularizers have been hand-crafted, including L^2-penalties, sparse regularization techniques or total variation [1,9,11]. While such hand-crafted regularizers are often convex and hence global minima can be computed by classical convex optimization, hand-crafted priors typically lack adaptability to available data.

In more recent years, there has been a shift to learned and potentially non-convex priors [3, 12–14, 16]. It has been observed that these methods often outperform classical methods. Moreover, a full convergence analysis has been be provided [12, 16]. However, such an analysis assumes minimizers of the Tikhonov functional to be known or at least be given within a certain accuracy. For non-convex regularizers such an assumption is unrealistic and global minimizers are challenging to compute. Instead, when trying to find a regularized solution one often employs minimization algorithms such as gradient descent or variations thereof which converge to critical points (such as local minimizers close to the initial guess) rather than to global minimizers of the Tikhonov functional. While one could constrain the learned regularizers to only include convex functionals [2, 14] this might result in suboptimal reconstructions when the underlying signal class is inherently non-convex. For such classes non-convexity of the regularizer can be a highly desirable property and as such a convergence analysis for this case is needed. Importantly, such an analysis should not rely on the strict assumption that the regularized solutions are global minimizers of the underlying Tikhonov functional.

1.2 Proposed critical point regularization

In this paper we present a convergence analysis of critical points of the Tikhonov functional $\mathcal{T}_{\alpha,y^\delta}$ for the stable solution of inverse problems of the form (1.1). We refer to any such method which recovers a critical point as regularized solutions as critical point regularization. In fact, we show stability and convergence for a relaxed notion of critical points. More precisely we study stability and convergence of ϕ-critical points, namely elements satisfying $0 \in \partial_\phi \mathcal{T}_{\alpha,y^\delta}(x_\alpha^\delta)$. Here ∂_ϕ is the ϕ-relative sub-differential, a novel concept that we introduce and study in this paper. Whenever the classical norm-discrepancy is used to measure similarity, as the noise level tends to zero, we show that regularized elements converge to elements $x^+ \in \mathbb{X}$ with

$$-\partial_\phi \mathcal{R}(x^+) \cap \ker(\mathbf{K})^\perp \neq \emptyset\,, \tag{1.3}$$

resembling first order conditions of the constraint optimization problem $\arg\min\{\mathcal{R}(x) \mid \mathbf{K}x = y\}$ defining $\mathcal{R}$-minimizing solutions.

Note that actually we give our analysis for more general data discrepancy measures $\mathcal{S}(x, y^\delta)$ for which $\frac{1}{2}\|\mathbf{K}x - y^\delta\|^2$ is only a special case. Further, we mention that in [8] an analysis of stability for the case of local minima has been done. Opposed to our work the authors of [8] restrict themselves to the finite dimensional setting and do not provide convergence results for the case that the noise-level tends to zero. We are not aware of any other study which includes stability and convergence of critical points and to the best of our knowledge the present analysis is the first to attempt this.

1.3 Main contributions

In this paper we introduce the concept of relative sub-differentiability as a generalization of sub-differentiability of convex functions to the non-convex case. We develop theory for relative sub-differentiability and show that corresponding ϕ-critical points can be found by employing a generalized gradient descent method. From the viewpoint of regularization theory we give existence, stability and convergence results for ϕ-critical points and derive the limiting problem for critical point regularization. As opposed to the convex case where the solutions one obtains are $\mathcal{R}$-minimizing solutions we get as a limiting problem a related first order optimality condition. As a special case of our analysis we derive stability and convergence results for critical points of differentiable Tikhonov functionals. For example, in this case, we get that $-\mathcal{R}'(x_+)$ is in the normal cone of the set of all solutions.

Finally, we provide numerical simulations which support our theoretical findings, in particular the stability, convergence and the limiting problem. Moreover, the results of our numerical simulations show that even in simple cases of non-convex regularizers the assumption of obtaining global minima or even local minima is infeasible thus further emphasizing the need for the analysis we provide in this paper. Besides, the numerical results show that the solutions we obtain cannot be expected to be $\mathcal{R}$-minimizing solutions and may even be local maxima of the regularizer whenever the initialization is chosen inappropriately and the algorithm does not guarantee that local minima are obtained.

1.4 Overview

The rest of the paper is organized as follows. In Section 2 we motivate and introduce the concept of relative sub-differentiability and corresponding ϕ-critical points. Moreover, we study basic properties of relative sub-differentiability and show that ϕ-critical points can be achieved by employing a generalized gradient descent method. Section 3 builds on this concept of relative sub-differentiability and gives a convergence analysis for critical point regularization. Moreover, we take a closer look at the differentiable case and identify the limiting problem in this case. In Section 4 we provide numerical experiments which support our theoretical findings such as stability and convergence. Finally, we conclude the paper by giving a brief summary and outlook in Section 5.

2 Relative sub-differentiability

In this section and in the rest of the paper, unless stated otherwise, we assume that $\mathbb{X}$ is a Banach space, denote by $\mathbb{X}^*$ its dual and by $\langle \cdot, \cdot \rangle$ the dual pairing of $\mathbb{X}$ and $\mathbb{X}^*$, i.e. for $\varphi \in \mathbb{X}^*$ and $x \in \mathbb{X}$ we have $\langle \varphi, x \rangle = \varphi(x)$. Moreover, we denote by $\mathcal{R}'$ the derivative of any differentiable function $\mathcal{R}\colon \mathbb{X} \to \mathbb{R}$ and for any similarity measure $\mathcal{S}\colon \mathbb{X} \times \mathbb{Y} \to \mathbb{R}$ we denote by $\mathcal{S}'$ the derivative with respect to its first argument.

Before giving the crucial definition of relative sub-differentiability we recall the importance of classical sub-differentiabilty in the context of convex functions. Recall that $r \in \mathbb{X}^*$ is called subgradient of some functional $\mathcal{F}\colon \mathbb{X} \to \mathbb{R}$ at $x \in \mathbb{X}$ if $\mathcal{F}(x) + \langle r, u - x \rangle \leq \mathcal{F}(u)$ for all $u \in \mathbb{X}$ and that $\mathcal{F}$ is sub-differentiable whenever the set of subgradients is non-empty for all $x \in \mathbb{X}$. Minimizers x of $\mathcal{F}$ are characterized by the optimality condition $0 \in \partial_0 \mathcal{F}(x)$ where $\partial_0 \mathcal{F}(x)$ denotes the set of all subgradients at point x. However, sub-differentiability implies convexity. We will therefore develop a relaxed concept of sub-differentiability relative to some functional $\phi\colon \mathbb{X} \to [0, \infty)$ by replacing the right hand side in the definition of subgradients by $\mathcal{F}(u) + \phi(u)$.

2.1 Definition and basic propertties

The following concept generalizing sub-differentiablity is also applicable to non-convex functions.

Definiton 2.1 (Relative sub-differentiability)**.** Let $\mathcal{F}\colon \mathbb{X} \to \mathbb{R}$ and $\phi\colon \mathbb{X} \to [0, \infty)$.

(a) $r \in \mathbb{X}^*$ is called ϕ-relative subgradient of $\mathcal{F}$ at $x \in \mathbb{X}$ if

$$u \in \mathbb{X}: \quad \mathcal{F}(x) + \langle r, u - x \rangle \leq \mathcal{F}(u) + \phi(u). \tag{2.1}$$

(b) The set set of all ϕ-relative subgradients at x is denoted by $\partial_\phi \mathcal{F}(x)$ and called ϕ-relative sub-differential of $\mathcal{F}$ at x.

(c) The functional $\mathcal{F}$ is called ϕ-relative sub-differentiable if $\partial_\phi \mathcal{F}(x) \neq \emptyset$ for all $x \in \mathbb{X}$.

Some remarks about Definition 2.1 are in order.

Remark 2.2.

- We call any such function ϕ a bound. It is clear that such a bound cannot be unique, since whenever $\mathcal{F}$ is a relatively sub-differentiable function with bound ϕ then it is also relatively sub-differentiable with bound $\phi + c$ for any $c \in [0, \infty)$.

- Choosing $\phi = 0$ we see that any convex and sub-differentiable function $\mathcal{F}$ is relatively sub-differentiable, i.e. the class of all relative sub-differentiable functions includes the set of convex sub-differentiable functions.
- The relative subgradients depend on the function ϕ. This shows that whenever we choose a larger ϕ then we generally also increase the set of possible relative subgradients, i.e. if $\phi_1 \le \phi_2$ then $\partial_{\phi_1}\mathcal{F} \subseteq \partial_{\phi_2}\mathcal{F}$.
- Similar to the concept of subgradients for convex functions, the concept of relative subgradients is a global property since the defining inequality has to hold for any point $u \in \mathbb{X}$.

Another approach of generalizing convexity and subgradients (and as a consequence critical points) is given in [10, 20] where convexity with respect to a set of functions W is defined. In such a setting $w \in W$ is a subgradient of $\mathcal{F}$ at x whenever $\mathcal{F}(u) \ge \mathcal{F}(x) + w(u) - w(x)$ for any $u \in \mathbb{X}$. As a consequence any critical point, i.e. a point where 0 is a subgradient, will also be a global minimizer and hence such a generalization cannot be used for our purposes. In [6] another concept of generalized gradients is discussed. In this setting the definition of the gradient depends only on neighborhoods around the point of interest. As a consequence we cannot expect the critical points to have any global properties which are necessary for the analysis in Section 3 hence making this generalization unfit for our analysis. However, it should be noted, that whenever convenient one might substitute any differentiability assumption on the involved functionals with Clarke's generalized gradient concept in any of the following discussions.

In what follows we will assume that $\mathcal{F}\colon \mathbb{X} \to \mathbb{R}$ is ϕ-relatively sub-differentiable for some fixed ϕ. Based on this definition we generalize the concept of critical points as follows.

Definiton 2.3 (ϕ-critical points)**.** We call $x \in \mathbb{X}$ a ϕ-critical point of $\mathcal{F}$ if $0 \in \partial_\phi\mathcal{F}(x)$. Moreover, we denote by $\operatorname{crit}_\phi \mathcal{F}$ the set of all ϕ-critical points of $\mathcal{F}$.

It should be noted that the definition of ϕ-critical points depends on ϕ and in practical applications one might not have access to ϕ. In such cases evaluating or finding relative subgradients might be infeasible. Nevertheless, the concept of ϕ-critical points is general enough to include an important class of points as the following remark illustrates.

Remark 2.4 (Critical points of differentiable functions)**.** Let us assume that $\mathcal{F}\colon \mathbb{X} \to \mathbb{R}$ is a differentiable function which satisfies the inequality $\mathcal{F}(x) + \langle \mathcal{F}'(x), u - x\rangle \le \mathcal{F}(u) + \phi(u)$ for any $x, u \in \mathbb{X}$ and some $\phi\colon \mathbb{X} \to [0,\infty)$. Then we have $\mathcal{F}'(x) \in \partial_\phi\mathcal{F}(x)$. This shows that in this special case we have access to at least one element of $\partial_\phi\mathcal{F}$. In particular, any critical point of $\mathcal{F}$, i.e. a point $x \in \mathbb{X}$ with $\mathcal{F}'(x) = 0$, will always yield a ϕ-critical point of $\mathcal{F}$ in the sense of Definition 2.3 and hence Definition 2.3 is a generalization of the classical concept of critical points for differentiable functions satisfying above inequality.
This shows that for a class of functions we have access to at least one element of the relative subgradient of $\mathcal{F}$. More importantly, for this class of functions we can make assertions about the points $x \in \mathbb{X}$ where $\mathcal{F}'(x) = 0$ holds, i.e. points which are reachable by use of a (minimization) algorithm which guarantees to find a critical point.

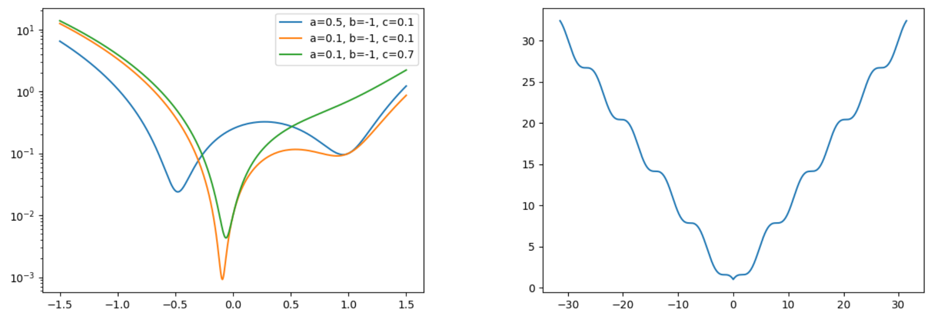

Figure 2.1: **Left:** Example of relatively sub-differentiable functions where the classical gradient is contained in the relative sub-gradient. **Right:** A function for which classical derivative cannot be in the relative sub-gradient.

Before we move on, we briefly give a prototypical example of a non-convex function for which a bound ϕ can be chosen, such that $\mathcal{F}'(x) \in \partial_\phi \mathcal{F}(x)$.

Remark 2.5 (Examples of relative sub-differentiability)**.** We start by giving a simple example of a function which is non-convex, but relatively sub-differentiable. To this end, let $a, b \in \mathbb{R}$ be given and define $\mathcal{F}(t) = (t+a)^2(t+b)^2$. It is readily seen, that $\mathcal{F}(t) + \mathcal{F}'(t)(s-t)$ is a polynomial of degree 4 with negative leading coefficient. Hence, this function is bounded from above and the relative sub-differentiability immediately follows by for example choosing $\phi(s) = \sup_t \mathcal{F}(t) + \mathcal{F}'(t)(s-t)$. Clearly then, the function $g(t) = \mathcal{F}(t) + ct^2$ is also sub-differentiable for $c > 0$. The function g is plotted in Figure 2.1 on the left side for different parameters a, b, c on a semi-logarithmic scale to emphasize the non-convexity.

Now let us consider the function $\mathcal{F}(t) = \cos(t) + |t|$, see Figure 2.1 on the right. Then, due to the coercivity and the existence of critical points "at infinity", the derivative of $\mathcal{F}$ cannot be in the relative sub-gradient of $\mathcal{F}$ for any ϕ. This example illustrates what types of functions are not included in the concept of relatively sub-differentiable functions for which the derivative is supposed to lie in the relative subgradient. In particular, the concept of relative sub-differentiability excludes coercive functionals which have critical points "at infinity".

Before discussing how one might obtain ϕ-critical points of relatively sub-differentiable functions, we list some useful properties of which we make constant use during the rest of the paper.

Lemma 2.6 (Basic properties of relative subgradients)**.** *Let $\mathcal{F}, \mathcal{F}_i \colon \mathbb{X} \to \mathbb{R}$ and $\phi, \phi_i \colon \mathbb{X} \to [0, \infty)$ be bounds of $\mathcal{F}, \mathcal{F}_i$ for $i = 1, \ldots, n$ and $w > 0$. Moreover, set $c := \inf \mathcal{F} + \phi$. Then the following hold*

(1) $\sum \partial_{\phi_i} \mathcal{F}_i \subseteq \partial_{\sum \phi_i} \sum \mathcal{F}_i$

(2) $w \partial_\phi \mathcal{F} = \partial_{w\phi}(w\mathcal{F})$

(3) If $\mathcal{F}$ is convex then $\partial_0 \mathcal{F}(x) \subseteq \partial_\phi \mathcal{F}(x)$ for any $x \in \mathbb{X}$

(4) $\partial_\phi \mathcal{F}(x)$ is convex and (weak) closed*

(5) If $x_* \in \arg\min \mathcal{F}(x)$ then $0 \in \partial_\phi \mathcal{F}(x_*)$

(6) $0 \in \partial_\phi \mathcal{F}(x) \iff \mathcal{F}(x) \leq c$

(7) If $\mathcal{F}$ *is Lipschitz and* ϕ *bounded on bounded subsets, then* $\partial_\phi \mathcal{F}$ *is bounded. In particular, in this case the set* $\partial_\phi \mathcal{F}$ *is weak*-compact.*

(8) Let $p_k = g_k + z_k$ *where* $g_k \in \partial_\phi \mathcal{F}(x_k)$ *and* $\|z_k\| \leq \varepsilon_k$ *with* $\varepsilon_k \to 0$*. Assume that* x_k *converge weakly to* x_+ *and* g_k *converge to* g *and that* $\mathcal{F}$ *is weakly lower semi-continuous. Then* $g \in \partial_\phi \mathcal{F}(x_+)$*. If, instead,* x_k *converge strongly to* x_+*,* g_k *converge weakly to* g *and* $\mathcal{F}$ *is lower semi-continuous then we also have* $g \in \partial_\phi \mathcal{F}(x_+)$*.*

Proof. (1) Let $p_i \in \partial \mathcal{F}_i(x)$ and define $p = \sum_i p_i$. Then we have

$$\sum_i \mathcal{F}_i(x) + \langle p, u - x \rangle = \sum_i (\mathcal{F}_i(x) + \langle p_i, u - x \rangle) \leq \sum_i (\mathcal{F}_i(u) + \phi_i(u))$$

and hence the claim follows.

(2) Assume that $p \in \partial_\phi \mathcal{F}(x)$. Then we have $w\mathcal{F}(x) + w\langle p, u - x \rangle \leq w\left(\mathcal{F}(u) + \phi(u)\right)$ by non-negativity of w and hence $w\partial_\phi \mathcal{F} \subseteq \partial_{w\phi}(w\mathcal{F})$. Now let $p \in \partial_{w\phi}(w\mathcal{F})$ then we define $q = \frac{p}{w}$ and it follows

$$w\left(\mathcal{F}(x) + \langle q, u - x \rangle\right) = w\mathcal{F}(x) + \langle p, u - x \rangle \leq w\left(\mathcal{F}(u) + \phi(u)\right)$$

which shows that $p \in w\partial_\phi \mathcal{F}(x)$.

(3) This is an immediate consequence of $\mathcal{F}(u) \leq \mathcal{F}(u) + \phi(u)$ by non-negativity of ϕ.

(4) Let $p_1, p_2 \in \partial_\phi \mathcal{F}(x)$ and $\lambda \in (0, 1)$. Then we have

$$\begin{aligned} \mathcal{F}(x) + \langle \lambda p_1 + (1-\lambda)p_2, u - x \rangle &= \lambda(\mathcal{F}(x) + \langle p_1, u - x \rangle) + (1-\lambda)(\mathcal{F}(x) + \langle p_2, u - x \rangle) \\ &\leq \lambda(\mathcal{F}(u) + \phi(u)) + (1-\lambda)(\mathcal{F}(u) + \phi(u)), \end{aligned}$$

which proves the convexity of $\partial_\phi \mathcal{F}$. Now let us assume that $p_k \in \partial \mathcal{F}(x)$ with p_k (weak*) converges to p. By (weak*) convergence we have $\langle p_k, u - x \rangle \to \langle p, u - x \rangle$ and hence p is also an relatively sub-differentiable subgradient.

(5) This is also a consequence of the non-negativity of ϕ and the assumption that x_+ is a global minimizer.

(6) Let $0 \in \partial_\phi \mathcal{F}(x)$. Then by definition we have $\mathcal{F}(x) \leq \mathcal{F}(u) + \phi(u)$ for any $u \in \mathbb{X}$ and hence also $\mathcal{F}(x) \leq c$. On the other hand, if $\mathcal{F}(x) \leq c$ then we have $\mathcal{F}(x) \leq \mathcal{F}(u) + \phi(u)$ for any $u \in \mathbb{X}$ and hence $0 \in \partial_\phi \mathcal{F}(x)$.

(7) Let $p \in \partial_\phi \mathcal{F}(x)$ and set $u = x + v$ with $\|v\| = 1$. Using the defining inequality we find

$$\langle p, v \rangle \leq \mathcal{F}(x + v) - \mathcal{F}(x) + \phi(x + v) \leq L + \phi(x + v)$$

and thus by taking the supremum over v we find that $\|p\|$ is bounded. Using Banach-Alaouglu we see that $\partial_\phi \mathcal{F}$ must be weak*-compact.

(8) By assumption x_k is bounded. Thus, we have

$$\begin{aligned}\mathcal{F}(x_+) + \langle g, u - x_+\rangle &\leq \liminf_k \mathcal{F}(x_k) + \langle p_k, u - x_k\rangle \\ &\leq \liminf_k \mathcal{F}(u) + \phi(u) + \langle z_k, u - x_k\rangle \\ &\leq \mathcal{F}(u) + \phi(u) + \liminf_k \varepsilon_k \|u - x_k\| \\ &= \mathcal{F}(u) + \phi(u),\end{aligned}$$

which proves the claim. □

Lemma 2.6 gives us a characterization of ϕ-critical points as points x for which $\mathcal{F} + \phi$ is an upper bound of $\mathcal{F}(x)$. This characterization in particular implies that for any differentiable and relatively sub-differentiable function $\mathcal{F}$ we have that the points x with $\mathcal{F}'(x) = 0$ must have bounded value independent of x. Comparing this to the convex case we have that x is a critical point of the function $\mathcal{F}$ if and only if x is a global minimizer. In some sense, the definition of ϕ-critical points allows for some error to be made and guarantees that ϕ-critical points cannot have arbitrarily large $\mathcal{F}$-value. Moreover, whenever $\mathcal{F}$ is coercive then all ϕ-critical points must be inside some ball $B_r(0)$ for some $r > 0$.

2.2 Computation of ϕ-critical points

We next answer the question of how to obtain ϕ-critical points at least for the case where $\mathbb{X}$ is a Hilbert space. Clearly, if $\mathcal{F}$ is differentiable then one could consider classical gradient descent methods. Since we are also interested in non-differentiable functions, gradient descent in its classical form may not be applicable. Below we show that a generalized gradient method using relative subgradients instead of gradients will yield ϕ-critical points in the sense of Definition 2.3. This shows that Algorithm 1 is a natural extension of subgradient descent [5, 19].

Algorithm 1 Relative subgradient descent

Require: Starting point $x_0 \in \mathbb{X}$, stepsizes $\eta_n > 0$
  $n \leftarrow 0$
  while $0 \notin \partial_\phi \mathcal{F}(x_n)$ **do**
    Choose $g_n^* \in \partial_\phi \mathcal{F}(x_n)$ and $g_n \in \mathbb{X}$ such that $\langle g_n^*, g_n\rangle > 0$
    $x_{n+1} = x_n - \eta_n g_n$
    $n \leftarrow n + 1$
  end while

The following results shows that Algorithm 1 converges to a ϕ-critical point of the function $\mathcal{F}$. The given proof closely follows the one given in [5] but does not assume a finite dimensional setting and considers relatively sub-differentiable functionals instead of sub-differentiable function.

Theorem 2.7 (Convergence of Algorithm 1). *Assume that $\mathbb{X}$ is a Hilbert space and that $\mathcal{F}$ is relatively sub-differentiable with bound ϕ. Moreover, choose $g_n = \lambda_n g_n^*$ in Algorithm 1 with $\lambda_n > 0$ such that $\|g_n\| \leq C$ for all $n \in \mathbb{N}$. Then for any point $u \in \mathbb{X}$ and any step*

$N \in \mathbb{N}$ we have

$$\min_{i=1,\dots,N} \mathcal{F}(x_i) \le \mathcal{F}(u) + \phi(u) + \frac{\|x_0 - u\|^2 + C^2 \sum_{i=1}^N \eta_i^2}{2\sum_{i=1}^N \eta_i}.$$

Proof. Let $u \in \mathbb{X}$. After rescaling of g_n^* according to assumption we may assume that $g_n = g_n^*$. Then by definition of x_{n+1} we have

$$\begin{aligned}\|x_{n+1} - u\|^2 &= \|x_n - u\|^2 - 2\eta_n \langle g_n, x_n - u\rangle + \eta_n^2 \|g_n\|^2 \\ &\le \|x_n - u\|^2 - 2\eta_n(\mathcal{F}(x_n) - \mathcal{F}(u) - \phi(u)) + \eta_n^2\|g_n\|^2.\end{aligned}$$

Applying this inequality recursively and using the fact that $\|x_{n+1} - u\|^2 \ge 0$ we find

$$2\sum_{i=1}^n \eta_i(\mathcal{F}(x_i) - \mathcal{F}(u) - \phi(u)) \le \|x_0 - u\|^2 + \sum_{i=1}^n \eta_i^2 \|g_i\|^2,$$

which together with the inequalities $\|g_i\| \le C$ and $\sum_{i=1}^n \eta_i \mathcal{F}(x_i) \ge \min_{i=1,\dots,n} \mathcal{F}(x_i) \sum_{i=1}^n \eta_i$ shows the desired result. □

Theorem 2.7 shows that under the assumption that the sequence of step-sizes $(\eta_n)_{n\in\mathbb{N}}$ is square-summable but not summable, then in the limit we have $\lim_{n\to\infty} \min_{i=1,\dots,n} \mathcal{F}(x_i) \le \mathcal{F}(u) + \phi(u)$ for any $u \in \mathbb{X}$. Note the analysis and the proofs heavily rely on the usage of the functional ϕ, but we note that at no point during Algorithm 1 do we need explicit knowledge of the functional ϕ but only access to elements of $\partial_\phi \mathcal{F}$. In particular, in the case of Remark 2.4 when using the gradient of $\mathcal{F}$ as the update direction the generated sequence will yield a ϕ-critical point.

Finally, assume that we have a functional of the form $\mathcal{F}(x) = \mathcal{S}(x) + \alpha\mathcal{R}(x)$ where each term is relatively sub-differentiable with bounds $\phi_\mathcal{S}$ and $\phi_\mathcal{R}$. Then Lemma 2.6 shows that for $s \in \partial_{\phi_\mathcal{S}}\mathcal{S}$ and $r \in \partial_{\phi_\mathcal{R}}\mathcal{R}$ we have $s + \alpha r \in \partial_{\phi_\mathcal{S} + \alpha\phi_\mathcal{R}}(\mathcal{S} + \alpha\mathcal{R})$. This implies that Algorithm 1 can be applied in the case where we are looking for a ϕ-critical point of the sum of two relatively sub-differentiable functionals and only have access to elements of $\partial_{\phi_\mathcal{S}}\mathcal{S}$ and $\partial_{\phi_\mathcal{R}}\mathcal{R}$.

3 Regularizing properties of ϕ-critical points

In this section we present a convergence analysis for ϕ-critical points of Tikhonov-type functionals $\mathcal{T}_{\alpha,y^\delta}$ extending the existing analysis for global minima [18]. At this point we want to emphasize again that the assumption of being able to obtain global minima of $\mathcal{T}_{\alpha,y^\delta}$ can be extremely restrictive when $\mathcal{R}$ is non-convex and the main goal of our analysis is to discard this assumption. Instead we focus only on ϕ-critical points of $\mathcal{T}_{\alpha,y^\delta}$ which may include local minimizers, saddle points or even local maxima.

Recall that we are interested in Tikhonov-type functionals $\mathcal{T}_{\alpha,y^\delta} \colon \mathbb{X} \to [0,\infty)$ of the form

$$\mathcal{T}_{\alpha,y^\delta}(x) = \mathcal{S}(x, y^\delta) + \alpha\mathcal{R}(x)\,, \tag{3.1}$$

for given $\mathcal{S}\colon \mathbb{X}\times\mathbb{Y}\to[0,\infty)$ and $\mathcal{R}\colon \mathbb{X}\to[0,\infty)$. Here, $\mathcal{S}$ is a similarity measure between x and y^δ and a standard situation we are interested in is $\mathcal{S}(x,y^\delta)=\frac{1}{2}\|\mathbf{K}(x)-y^\delta\|^2$ where $\mathbf{K}\colon\mathbb{X}\to\mathbb{Y}$ is the forward operator of the inverse problem of interest. Instead of working with global minima of the functional (3.1) we consider regularized solutions x_α^δ as $\alpha\phi$-critical points of $\mathcal{T}_{\alpha,y^\delta}$, meaning

$$0\in\partial_{\alpha\phi}\left(\mathcal{S}(\cdot,y^\delta)+\alpha\mathcal{R}(\cdot)\right)(x_\alpha^\delta) \tag{3.2}$$

We will analyze stability and convergence of such critical points.

For the analysis we make the following assumptions.

Condition 3.1 (Critical point regularization).

(C1) $\mathbb{X}$ is a reflexive Banach spaces and $\mathbb{Y}$ is a metric space with metric $\mathcal{D}$

(C2) $\mathcal{R}$ is weakly sequentially lower semi-continuous

(C3) $\mathcal{R}$ is relatively sub-differentiable with bound ϕ

(C4) $\mathcal{S}$ is weakly sequentially lower semi-continuous, convex in its first argument and continuous in its second argument

(C5) $\exists C>0\ \exists p\geq 1\ \forall z\in\mathbb{X}\ \forall y,y^\delta\in\mathbb{Y}\colon \mathcal{S}(z,y)\leq C\left(\mathcal{S}(z,y^\delta)+\mathcal{D}(y,y^\delta)^p\right)$

(C6) $\forall\alpha>0$ and $\forall y^\delta\in\mathbb{Y}$ the functional $\mathcal{T}_{\alpha,y^\delta}$ is coercive

Most of the assumptions in Condition 3.1 are classical assumptions (or generalizations thereof), e.g. [12,16,18], made for the analysis of variational methods. The major difference in the analysis provided here is that $\mathcal{R}$ is relatively sub-differentiable, which we have motivated in Section 2, and the assumption that in general the regularized solutions are not global minima but only ϕ-critical points.

One of the simplest and commonly used example of a similarity measure which satisfies Assumptions (C4) and (C5) is given by $\mathcal{S}(x,y^\delta)=\|\mathbf{K}x-y^\delta\|^p$ whenever $\mathbf{K}\colon\mathbb{X}\to\mathbb{Y}$ is the linear forward operator of the underlying inverse problem and $\mathbb{Y}$ is a Banach space. In general, any similarity measure of the form $\|\mathbf{L}(\mathbf{K}x-y^\delta)\|^p$ satisfies these assumptions, if $\mathbf{L}$ is a linear and bounded operator, e.g. a reweighting of the residual $(\mathbf{K}x-y^\delta)$.

We now turn our focus to the stability and convergence analysis of the considered method, i.e. $x_\alpha^\delta\in\operatorname{crit}_{\alpha\phi}\mathcal{T}_{\alpha,y^\delta}$. We start with existence and stability results.

3.1 Existence and stability

Theorem 3.2 (Existence). *Under Assumption 3.1 the problem is well-posed, i.e. for every $\alpha>0$ and $y^\delta\in\mathbb{Y}$ the set $\operatorname{crit}_{\alpha\phi}\mathcal{T}_{\alpha,y^\delta}$ is non-empty.*

Proof. This is an immediate consequence of the existence of minimizers of $\mathcal{T}_{\alpha,y^\delta}$ which follows from the coercivity and the continuity assumptions on the functional $\mathcal{T}_{\alpha,y^\delta}$. A more detailed proof can be found in [18]. □

Clearly, ϕ-critical points may exist under weaker assumptions than a coercivity assumption. However, the coercivity is an important property in the following analysis which guarantees the existence of a weakly convergent subsequence whenever the sequence is bounded. As such we have also derived existence of ϕ-critical points using the coercivity. Extending the current analysis to the case of non-coercive functionals $\mathcal{T}_{\alpha,y^\delta}$ is subject to future work.

Another advantage of using ϕ-critical points opposed to global minima, besides being numerically and hence practically more tractable for non-convex functionals, is that we have a simple way of talking about "inexact" critical points, i.e. points where the gradient is small but not necessarily 0. As it turns out, the following analysis can be performed under the even weaker assumption that the stabilized solutions are "inexact" critical points instead of exact critical points.

Theorem 3.3 (Stability). *Let $y^\delta \in \mathbb{Y}, \alpha > 0$ and $y_k \to y^\delta$ and assume that $x_k \in \mathbb{X}$ is such that $z_k \in \partial_{\alpha\phi}\left(\mathcal{S}(\cdot, y_k) + \alpha\mathcal{R}(\cdot)\right)(x_k)$ with $\|z_k\| \to 0$ and $\langle z_k, x_k\rangle \leq 0$. Then the sequence $(x_k)_k$ has a weakly convergent subsequence and the limit x_+ of every weakly convergent subsequence is an $\alpha\phi$-critical point of $\mathcal{T}_{\alpha,y^\delta}$.*

Proof. To show the existence of a weakly convergent subsequence, using the reflexivity of $\mathbb{X}$, it is enough to show that $(x_k)_k$ is a bounded sequence. By coercivity of $\mathcal{T}_{\alpha,y^\delta}$ it is enough to show that $(\mathcal{T}_{\alpha,y^\delta}(x_k))_k$ is bounded. We have for any $u \in \mathbb{X}$

$$\mathcal{S}(x_k, y_k) + \alpha\mathcal{R}(x_k) + \langle z_k, u - x_k\rangle \leq \mathcal{S}(u, y_k) + \alpha\mathcal{R}(u) + \alpha\phi(u)$$

and using $\langle z_k, x_k\rangle \leq 0$ it follows

$$\mathcal{S}(x_k, y_k) + \alpha\mathcal{R}(x_k) \leq \mathcal{S}(u, y_k) + \alpha\mathcal{R}(u) + \alpha\phi(u) + \|z_k\|\|u\|.$$

By assumption on $\mathcal{S}$ we have $\mathcal{S}(x_k, y^\delta) \leq C(\mathcal{S}(x_k, y_k) + \mathcal{D}(y_k, y^\delta)^p)$ which yields

$$\begin{aligned}\mathcal{S}(x_k, y^\delta) + \alpha\mathcal{R}(x_k) &\leq C\left(\mathcal{S}(x_k, y_k) + \alpha\mathcal{R}(x_k) + \mathcal{D}(y_k, y^\delta)^p\right)\\ &\leq C\left(\mathcal{S}(u, y_k) + \alpha\mathcal{R}(u) + \alpha\phi(u) + \|z_k\|\|u\| + \mathcal{D}(y_k, y^\delta)^p\right)\\ &\leq \tilde{C}\left(\mathcal{S}(u, y^\delta) + \alpha\mathcal{R}(u) + \alpha\phi(u) + \|z_k\|\|u\| + \mathcal{D}(y_k, y^\delta)^p\right)\end{aligned}$$

for any $u \in \mathbb{X}$. By assumption we have $\|z_k\| \to 0$ and $\mathcal{D}(y_k, y^\delta) \to 0$ so the right hand side is bounded for k large enough. This shows that there exists some weakly convergent subsequence.

Let now $(x_k)_k$ denote such a subsequence and denote by x_+ its limit. Using the weak lower semi-continuity of the involved functionals it follows for any $u \in \mathbb{X}$

$$\begin{aligned}\mathcal{S}(x_+, y^\delta) + \alpha\mathcal{R}(x_+) &\leq \liminf_k \mathcal{S}(x_k, y_k) + \alpha\mathcal{R}(x_k) + \langle z_k, u - x_k\rangle\\ &\leq \liminf_k \mathcal{S}(u, y_k) + \alpha\mathcal{R}(u) + \alpha\phi(u) + \|z_k\|\|u\|\\ &= \mathcal{S}(u, y^\delta) + \alpha\mathcal{R}(u) + \alpha\phi(u)\end{aligned}$$

where the last equality follows from continuity of $\mathcal{S}$ in its second argument. This shows that $0 \in \partial_{\alpha\phi}\left(\mathcal{S}(\cdot, y^\delta) + \alpha\mathcal{R}(\cdot)\right)(x_+)$. □

Clearly, whenever $z_k = 0$, i.e. x_k is an $\alpha\phi$-critical point, then the assumptions on z_k in Theorem 3.3 are satisfied. It follows that $\alpha\phi$-critical points are stable in the above sense. However, Theorem 3.3 also shows that we do not need access to exact $\alpha\phi$-critical points but rather points which are in some sense close to an $\alpha\phi$-critical point.

Remark 3.4 (Inexact critical points obtained by use of minimization schemes)**.** Consider once again the case of Remark 2.4 and assume that the $\alpha\phi$-critical points are obtained by using gradient descent or any other algorithm which finds zeros of the gradient. Then we have $z_k = \mathcal{S}'(x_k, y_k) + \alpha\mathcal{R}'(x_k)$ and whenever $\|z_k\| \to 0$ and $\langle z_k, x_k\rangle \leq 0$ we have that the considered points have a weakly convergent subsequence. For practical applications this means, that we have an easily verifiable condition which can be used as a kind of stopping criterion when searching for critical points. As a consequence, we do not have to guarantee that the regularized solutions are critical points but rather are "close" to being a critical point.

3.2 Convergence

The next goal is to show the convergence of the regularized solutions to a solution of the original problem in the case that the noise-level δ tends to 0. Here, we call $z \in \mathbb{X}$ an $\mathcal{S}$-solution of $y \in \mathbb{Y}$ if $\mathcal{S}(z, y) = 0$. Like in the case of Theorem 3.3, the proof can be done under the weaker assumption of only having access to "inexact" ϕ-critical points (see Remark 3.4).

Theorem 3.5 (Convergence)**.** *Let $y \in \mathbb{Y}$ and assume it has an $\mathcal{S}$-solution. Further, let $y_k \in \mathbb{Y}$ with $\mathcal{D}(y_k, y) \leq \delta_k$ with $\delta_k \to 0$. Choose $\alpha = \alpha(\delta)$ such that for $\alpha_k = \alpha(\delta_k)$ we have $\lim_k \alpha_k = \lim_k \delta_k^p/\alpha_k = 0$. Assume that the regularized solutions $x_k \in \mathbb{X}$ are such that $z_k \in \partial_{\alpha_k\phi}(\mathcal{S}(\cdot, y_k) + \alpha_k\mathcal{R}(\cdot))(x_k)$ with $\|z_k\|/\alpha_k \to 0$ and $\langle z_k, x_k\rangle \leq 0$.*
Then the sequence $(x_k)_k$ has a weakly convergent subsequence and the limit x_+ of any such sequence is an $\mathcal{S}$-solution of y. Moreover, we have $\mathcal{R}(x_+) \leq \mathcal{R}(u) + \phi(u)$ for any $\mathcal{S}$-solution u. Finally, whenever the $\mathcal{S}$-solution is unique then $(x_k)_k$ converges weakly to this solution.

Proof. Similar to the stability proof we show that $(x_k)_k$ is bounded by using the coercivity of the functionals $\mathcal{T}_{\alpha,y^\delta}$. Following the above proof we find for any $u \in \mathbb{X}$

$$\mathcal{S}(x_k, y_k) + \alpha_k\mathcal{R}(x_k) \leq \mathcal{S}(u, y_k) + \alpha_k\mathcal{R}(u) + \alpha_k\phi(u) + \|z_k\|\|u\|$$

and by choosing u such that $\mathcal{S}(u, y) = 0$ we find $\mathcal{S}(u, y_k) \leq C\delta_k^p$ which implies

$$\mathcal{S}(x_k, y_k) + \alpha_k\mathcal{R}(x_k) \leq C\delta_k^p + \alpha_k(\mathcal{R}(u) + \phi(u)) + \|z_k\|\|u\|.$$

Since both $\mathcal{S}$ and $\mathcal{R}$ are non-negative it then follows

$$\lim_k \mathcal{S}(x_k, y_k) = 0$$
$$\limsup_k \mathcal{R}(x_k) \leq \mathcal{R}(u) + \phi(u),$$

where we have used the assumptions $\lim_k \delta_k^p/\alpha_k = \lim_k \|z_k\|/\alpha_k = 0$. This shows that $(\mathcal{R}(x_k))_k$ is a bounded sequence and using once again $\mathcal{S}(x_k, y) \leq C\left(\mathcal{S}(x_k, y_k) + \delta_k^p\right)$ we find that for

$\alpha^+ = \max\{\alpha_k \colon k \in \mathbb{N}\}$ the sequence $(\mathcal{S}(x_k, y) + \alpha^+\mathcal{R}(x_k))_k$ is bounded. Using the coercivity of $\mathcal{S}(\cdot, y) + \alpha^+\mathcal{R}(\cdot)$ we get that the sequence $(x_k)_k$ is bounded and hence has a weakly convergent subsequence.

Finally, using the weak lower-semicontinuity of $\mathcal{S}$ and $\mathcal{R}$ we have that for any such weakly convergent subsequence with limit x_+

$$\mathcal{S}(x_+, y) \leq \liminf_k \mathcal{S}(x_k, y_k) = 0$$
$$\mathcal{R}(x_+) \leq \liminf_k \mathcal{R}(x_k) \leq \mathcal{R}(u) + \phi(u)$$

for any $u \in \mathbb{X}$ with $\mathcal{S}(u, y) = 0$.

Whenever the solution is unique, then every subsequence of $(x_k)_k$ has a subsequence converging to this solution. This shows that $(x_k)_k$ converges weakly to the unique solution. □

At this point, we want to emphasize once again, that the assumptions on the choice of points x_k in Theorem 3.5 are weaker than the assumption that x_k is an $\alpha_k\phi$-critical point and that in particular the analysis also holds for these points.

Since for this section we only assume that $\mathcal{R}$ is relatively sub-differentiable without explicit knowledge of the bound ϕ Theorem 3.5 gives a somewhat intangible condition on the type of solutions we obtain in the limit $\delta \to 0$. A more tangible condition, and more importantly one independent of ϕ, is given by the next theorem, where we assume a separability condition on the gradients $z_k \in \partial_{\alpha_k\phi}\left(\mathcal{S}(\cdot, y_k) + \alpha_k\mathcal{R}(\cdot)\right)(x_k)$. This separability assumption can be satisfied in many cases, e.g. when $\mathcal{S}$ and $\mathcal{R}$ are (relatively sub-)differentiable and the ϕ-critical points arise due to some algorithm such as gradient descent. Using these algorithms we are often in the situation that $z_k = s_k + \alpha_k r_k$ where s_k is a (sub-)gradient of $\mathcal{S}$ and r_k is an (relatively sub-differentiable sub-)gradient of $\mathcal{R}$. Assuming that the gradients $(r_k)_k$ of $\mathcal{R}$ have a cluster point, we get the additional following property of these cluster points.

Theorem 3.6 (Normality property of the solution)**.** *Let the same assumptions as in Theorem 3.5 hold and denote by $(x_k)_k$ a weakly convergent subsequence with limit x_+. Let $z_k = s_k + \alpha_k r_k$ where $s_k \in \partial_0\mathcal{S}(\cdot, y_k)(x_k)$ and $r_k \in \partial_\phi\mathcal{R}(x_k)$.*
Then any cluster point r of the sequence $(r_k)_k$ satisfies $-r \in -\partial_\phi\mathcal{R}(x_+) \cap N_{L(y)}(x_+)$, where $N_{L(y)}(x_+)$ is the normal cone of the convex set of all $\mathcal{S}$-solutions of y at x_+.

Proof. Let r be a cluster point of the sequence $(r_k)_k$. Then by weak lower semi-continuity of $\mathcal{R}$ and by assumption on r_k we have

$$\mathcal{R}(x_+) + \langle r, u - x_+\rangle \leq \liminf_k \mathcal{R}(x_k) + \langle r_k, u - x_k\rangle \leq \mathcal{R}(u) + \phi(u),$$

which shows that $r \in \partial_\phi\mathcal{R}(x_+)$.

Now assume that $u \in \mathbb{X}$ is such that $\mathcal{S}(u, y) = 0$. Then we have $\mathcal{S}(u, y_k) \leq C\delta_k^p$ and it follows

$$\begin{aligned}
\langle -r, u - x_+\rangle &= \lim_k \langle \frac{s_k - z_k}{\alpha_k}, u - x_k\rangle \\
&\leq \lim_k \frac{1}{\alpha_k} \left(\mathcal{S}(u, y_k) - \mathcal{S}(x_k, y_k) + \|z_k\|\|u\| + \langle z_k, x_k\rangle\right) \\
&\leq \lim_k \frac{1}{\alpha_k} \left(\mathcal{S}(u, y_k) + \|z_k\|\|u\|\right) \\
&\leq \lim_k C\frac{\delta_k^p}{\alpha_k} + \|u\|\frac{\|z_k\|}{\alpha_k} \\
&= 0,
\end{aligned}$$

where we have used the convexity of $\mathcal{S}$ in its first argument and the assumption on the limits of the sequences $(\alpha_k)_k$ and $(\|z_k\|/\alpha_k)_k$. ☐

Theorem 3.6 shows that the solution x_+ obtained by critical point regularization satisfies some form of first order optimality conditions, see e.g. [17].

Also note that in the case where $\mathcal{R}$ is convex and we choose $\phi = 0$, both properties in Theorem 3.5 and 3.6 reduce to the common property that x_+ is an $\mathcal{R}$-minimizing solution, i.e. $\mathcal{R}(x_+) \leq \mathcal{R}(u)$ for any $u \in \mathbb{X}$ with $\mathcal{S}(u, y) = 0$.

Remark 3.7 (Convex regularizers)**.** Clearly, any sub-differentiable convex function is relatively sub-differentiable with the choice $\phi = 0$. Nevertheless, one could also choose $\phi = \varepsilon > 0$. With this choice we see that the results in Theorem 3.5 roughly state that the solutions x_+ we approximate by using critical point regularization are $\mathcal{R}$-minimizing solutions up to an error ε whenever the regularized solutions are minimizers up to an error of ε.
This result, as opposed to classical variational regularization theory e.g. [18], has the advantage that at no point do we require exact global minimizers of the functionals $\mathcal{T}_{\alpha,y}$ but only approximate minimizers, which may be more easily reachable in practical applications. Consider for example the case where we employ an iterative algorithm which has convergence guarantees of the form $\mathcal{F}(x_n) - \mathcal{F}(x_*) \leq C/n^r$ for the n-th iterate and x_* being a minimizer of $\mathcal{F}$. Applying this algorithm to $\mathcal{F} = \mathcal{T}_{\alpha,y}$ and requiring that $C/n^r \leq \alpha\varepsilon$ in order to get that x_n is an $\alpha\phi$-critical point, we see that the above theory shows that one might stop the iterative algorithm after a finite amount of steps, i.e. we do not necessarily need to run the algorithm until it converges and we still get a stable and convergent regularization method.
At first it might seem that a disadvantage of this is that we do not achieve an $\mathcal{R}$-minimizing solutions in the limit. However, this can also be circumvented by considering a variable ε. To be more precise, following the proof of Theorem 3.5 with $\varepsilon = \varepsilon(\delta)$ and the condition $\varepsilon(\delta) \to 0$ as $\delta \to 0$, it is easy to see that to obtain a sequence $(x_k)_k$ weakly converging to an $\mathcal{R}$-minimizing solution it is enough to run the iterative algorithm for a number of iterations steps n_k such that $C/n_k^r \leq \alpha_k\varepsilon_k$.

We next discuss another special case of our analysis which pertains to functionals such as the one in Remark 2.5.

3.3 Differentiable regularizers and classical critical points

In this subsection we consider the important special case where the ϕ-critical points are given by classical critical points, i.e. by points x for which $\mathcal{T}'_{\alpha,y}(x) = 0$ and we give stability and convergence results for this case. To this end, we assume that the bound ϕ can be constructed in such a way that $\mathcal{R}(x) + \langle \mathcal{R}'(x), u - x\rangle \leq \mathcal{R}(u) + \phi(u)$, see e.g. Remark 2.4. Then, if $\mathcal{S}$ is differentiable in its first argument, by convexity of $\mathcal{S}$, we have $\mathcal{S}'(x, y^\delta) + \alpha\mathcal{R}'(x) \in \partial_{\alpha\phi}(\mathcal{S}(\cdot, y^\delta) + \alpha\mathcal{R}(\cdot))(x)$. This shows, that whenever we employ some algorithm which finds a classical critical point, we also obtain an $\alpha\phi$-critical point in the sense of Definition 2.1 which satisfies the separability assumption necessary for Theorem 3.6. In particular, these points are amenable to the analysis above.

Nevertheless, the analysis relies on an abstract concept of ϕ-critical points and even in the case where the involved functionals are differentiable we cannot guarantee that the limits will again be ϕ-critical points without any additional assumptions. In order to guarantee this we need the assumption that $\mathcal{S}'$ and $\mathcal{R}'$ are weakly (sequentially) continuous. Combining the above theorems we then get the following result.

Proposition 3.8 (Existence, stability and convergence for classical critical points)**.** Assume that $\mathcal{S}$ and $\mathcal{R}$ are differentiable with weakly continuous derivatives and let Condition 3.1 hold. Moreover, let $y, y^\delta \in \mathbb{Y}$ and $\alpha > 0$ and assume that y has an $\mathcal{S}$-solution. Then the following hold

1. **Existence:** $\mathcal{T}_{\alpha,y^\delta}$ has at least one ϕ-critical point.
2. **Stability:** If $(y_k)_k \subseteq \mathbb{Y}$ is a sequence converging to y^δ and x_k is such that $z_k = \mathcal{S}'(x_k, y_k) + \alpha\mathcal{R}'(x_k) \to 0$ as $k \to \infty$ and $\langle z_k, x_k\rangle \leq 0$. Then $(x_k)_k$ has a weakly convergent subsequence and any weak clusterpoint x_+ of $(x_k)_k$ is a critical point of $\mathcal{T}_{\alpha,y^\delta}$.
3. **Convergence:** Let $(y_k)_k \subseteq \mathbb{Y}$ be a sequence with $\mathcal{D}(y_k, y) \leq \delta_k$ and $\alpha = \alpha(\delta)$ be such that for $\alpha_k = \alpha(\delta_k)$ we have $\lim_k \alpha_k = \lim_k \delta_k^p/\alpha_k = 0$. Then, if we choose $x_k \in \mathbb{X}$ such that $z_k = \mathcal{S}'(x_k, y_k) + \alpha_k\mathcal{R}'(x_k)$ satisfies $\lim_k \|z_k\|/\alpha_k = 0$ and $\langle z_k, x_k\rangle \leq 0$ the sequence $(x_k)_k$ has at least one weak clusterpoint and any such clusterpoint x_+ is an $\mathcal{S}$-solution of y with the following additional properties
   (a) $\mathcal{R}(x_+) \leq \inf_{\mathcal{S}(u,y)=0} \mathcal{R}(u) + \phi(u)$
   (b) $\langle -\mathcal{R}'(x_+), z - x_+\rangle \leq 0$ for any $z \in \mathbb{X}$ with $\mathcal{S}(z, y) = 0$, i.e. $-\mathcal{R}'(x_+) \in N_{L(y)}(x_+)$.

   Finally, whenever the $\mathcal{S}$-solution is unique then $(x_k)_k$ converges weakly to this solution.

Proof. This follows immediately by applying Theorems 3.2, 3.3, 3.5 and 3.6. □

For the differentiable case this identifies the limiting problem we solve by regularizing the inverse problems with critical points, i.e. in the limit we find solutions which satisfies a first order optimality condition of the constrained optimization problem

$$\inf_u \mathcal{R}(u) \quad \text{such that} \quad \mathcal{S}(u, y) = 0.$$

We now briefly discuss the special case where $\mathcal{S}$ is given as the norm-discrepancy, e.g. in the case where $\mathbb{Y}$ is a Hilbert-space.

Lemma 3.9 (Solution for norm discrepancy)**.** *Let the same assumptions as in Proposition 3.8 hold and assume that $\mathcal{S}(x, y^\delta) = \frac{1}{p}\|\mathbf{K}x - y^\delta\|_{\mathbb{Y}}^p$ for some $p > 1$ where $\mathbf{K}\colon \mathbb{X} \to \mathbb{Y}$ is a linear and bounded forward operator between Banach spaces and $\|\cdot\|_{\mathbb{Y}}^p$ is differentiable. Furthermore, denote by x_+ a solution according to Proposition 3.8.*
Then we have $-\mathcal{R}'(x_+) \in \ker(\mathbf{K})^\perp = \{p \in \mathbb{X}^* : \forall x_0 \in \ker(\mathbf{K})\colon \langle p, x_0\rangle = 0\}$.

Proof. Any solution z can be written as $z = x_+ + x_0$ where $x_0 \in \ker(\mathbf{K})$. By using x_0 and $-x_0$, Proposition 3.8 shows that $\langle -\mathcal{R}'(x_+), x_0\rangle = 0$ for any $x_0 \in \ker(\mathbf{K})$ and hence the claim follows. □

4 Numerical simulations

The goal of this section is not to show that non-convex regularizers can improve the reconstructions, but rather to test the theory derived in the previous sections and to show what may happen when non-convex regularizers are chosen.
To this end, we consider the discretzied version of two toy-problems in 1D. We consider an inapinting problem where around 50% of the signal entries were randomly removed. In this case the kernel of the forward operator $\mathbf{K}$ is simple to compute and by using a separable prior we can easily study the properties of the solution we obtain in the limit. This makes the first toy problem ideal for testing whether the properties (as described in the theory section) of the limiting solution hold true or not.
Further, we consider recovering a signal from its cumulative sum. Since this matrix is invertible there is a unique solution and following Theorem 3.5 we should observe convergence to this solution in the limit $\delta \to 0$. This toy problem is therefore well suited to study if the given ϕ-critical points actually converge to the unique solution.

For both problems we consider as the signal to recover the discretization of the function $f(t) = \exp(-t^2) \cdot \cos(t) \cdot (t - 0.5)^2 + \sin(t^2)$ on $t \in [-1, 1]$ using $N = 512$ equidistant sample points. We denote this signal by x_{true} and the true underlying data by $y_{\text{true}} = \mathbf{K}x_{\text{true}}$ where $\mathbf{K}$ is the forward operator of the corresponding problem.

For each problem we consider the similarity measure given by $\mathcal{S}(x, y) = \frac{1}{2}\|\mathbf{K}x - y\|^2$ and we construct a regularizer by $\mathcal{R}(x) = \sum_{i=1}^N \psi_{\rho,\beta}(x_i)$. Here, we define $\psi_{\rho,\beta}(t) = (t-\rho)^2 \cdot (t+\frac{\rho}{2})^2 + \frac{\beta}{2}t^2$ for $\rho, \beta > 0$. The function $\psi_{\rho,\beta}$ is constructed in such a way that it is non-convex but relatively sub-differentiable, see Remark 2.5. Figure 4.3 shows the function $\psi_{\rho,\beta}(t)$ with parameters $\rho = 2$ and $\beta = 10^{-1}$ for $t \in [-3, 3]$ where the y-axis is plotted on a logarithmic scale in order to emphasize the non-convexity. We can see that this function has a global minimum at around $t = -\frac{\rho}{2}$, a local minimum close to $t = \rho$ and another critical point in the interval $[0, 1]$. The parameters $\rho = 2$ and $\beta = 10^{-1}$ are used for all the following simulations.

As a separable sum of relatively sub-differentiable and non-convex terms the regularizer $\mathcal{R}$ as defined above is relatively sub-differentiable and non-convex. By definition of $\mathcal{R}$ it is further

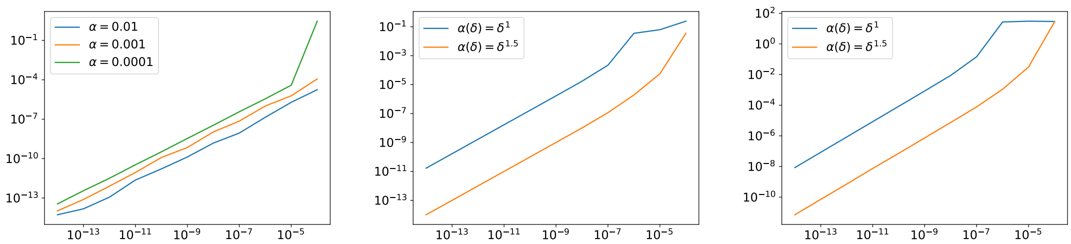

Figure 4.1: **Stability and convergence for the cumulative sum problem.** Each value is plotted in dependence on δ. **Left:** $||x_\alpha^\delta - x_\alpha||$ for different but fixed values of α. **Middle:** $||\mathbf{K}x_{\alpha(\delta),\delta} - y||$ for different $\alpha(\delta)$. **Right:** $||x_{\alpha(\delta),\delta} - x_+||$ for different $\alpha(\delta)$.

coercive and hence the functional $\mathcal{T}_{\alpha,y^\delta}$ is coercive. This shows that Condition 3.1 is satisfied and we consider the stability and convergence of the ϕ-critical points according to Theorem 3.3 and 3.5.

To simulate noisy data we consider the data $y_k = y_{\texttt{true}} + \delta_k \cdot n$ where $n = \frac{\xi}{||\xi||}$, ξ is a normally distributed random variable and $\delta_k = 10^{-k}$ for $k \in \{4, \dots, 14\}$.

Since a bound ϕ for $\mathcal{R}$ can be chosen such that $\mathcal{R}'(x) \in \partial_\phi \mathcal{R}(x)$ (see Remark 2.5), we can simply search for a classical critical point of $\mathcal{T}_{\alpha,y^\delta}$ in order to obatin ϕ-critical points. To achieve this, we apply Newton's method, e.g. [7], and we find an initial guess for Newton's method by applying Nesterov accelerated gradient descent [15] to the starting point $x_0 = 0$. In the following we denote by x_α^δ a critical point of $\mathcal{T}_{\alpha,y^\delta}$ and by x_α a critical point of $\mathcal{T}_{\alpha,y_{\texttt{true}}}$. Here, x_α is considered as the limit point for the stability considerations for which we consider the choices $\alpha \in \{10^{-2}, 10^{-3}, 10^{-4}\}$. In order to test for convergence we chosen $\alpha = \alpha(\delta) = \delta^q$ for $q \in \{1, \frac{3}{2}\}$. For the convergence simulations we consider as the limit point the signal $x_{\texttt{true}}$ for the cumulative sum problem, as in this case the solution is unique, and we construct an approximate solution for the inpainting problem by finding a critical point of the function $\mathcal{T}_{\alpha(\delta),\delta}$ for $\delta = 10^{-16}$ and we denote this solution by x_+. Implementation details and code are publicly available[1].

4.1 Results

Figure 4.1 depicts the value $||x_\alpha^\delta - x_\alpha||$ for different values of $\alpha > 0$ (left), $||\mathbf{K}x_{\alpha(\delta),\delta} - y_{\texttt{true}}||$ (middle) and $||x_{\alpha(\delta),\delta} - x_+||$ (right). Each of these values is plotted against δ on a log-log scale. The plot in the left shows that for any chosen α we can observe the convergence of the sequence x_α^δ to the critical point x_α as the noise-level tends to 0. The plots in the middle and right show the convergence behaviour for different choices $\alpha(\delta) = \delta^q$ as specified above. All of the sequences can be observed to converge, i.e. $||\mathbf{K}x_{\alpha(\delta),\delta} - y_{\texttt{true}}|| \to 0$ and $||x_{\alpha(\delta),\delta} - x_+|| \to 0$ as the noise δ tends to 0.

Figure 4.2 shows the same behaviour for the stability and convergence plots for the inpainting problem as Figure 4.1 in the limit $\delta \to 0$. In particular convergence to a solution x_+ of the problem $\mathbf{K}x = y_{\texttt{true}}$ can be observed.

[1] https://git.uibk.ac.at/c7021101/critical-point-regularization

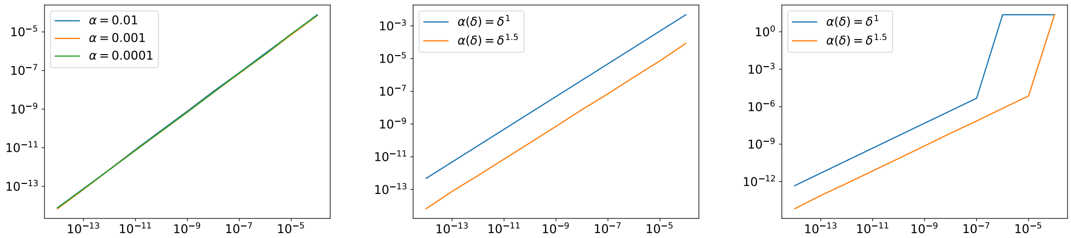

Figure 4.2: **Stability and convergence for the inpainting problem.** Each value is plotted in dependence on δ. **Left:** $\|x_\alpha^\delta - x_\alpha\|$ for different but fixed values of α. **Middle:** $\|\mathbf{K}x_{\alpha(\delta),\delta} - y\|$ for different $\alpha(\delta)$. **Right:** $\|x_{\alpha(\delta),\delta} - x_+\|$ for different $\alpha(\delta)$.

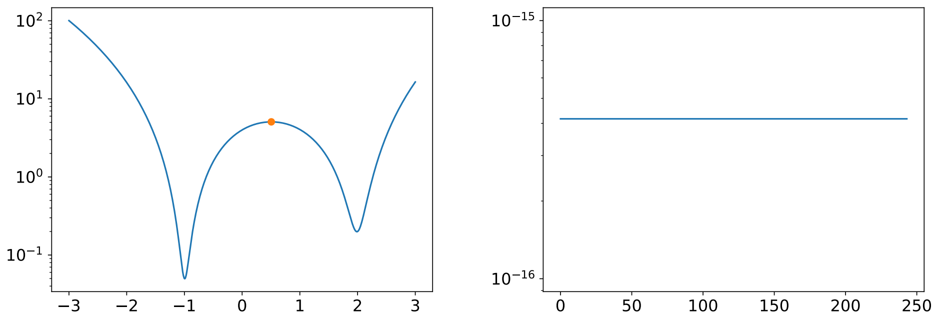

Figure 4.3: **Regularizer and properties of the inpainting solution. Left:** $\psi_{2,10^{-1}}(t)$ for $t \in [-3, 3]$ on a logarithmic scale to emphasize the local minimum. The dot is the value of $(x_+)_i$ where x_+ is the solution of the inpainting problem and i is an index chosen such that e_i is in the kernel of the inpainiting problem. **Right:** the values $|\langle \mathcal{R}'(x_+), e_i\rangle|$ where $(e_i)_{i\in I}$ is a basis of the kernel of the inpainting problem.

A closer look at the inpainting problem reveals that the limit point x_+ is, however, not an $\mathcal{R}$-minimizing solution. This can easily be checked due to the separability of the regularizer and the simple representation of the kernel of the inpainting problem. The orange dot in Figure 4.3 (left) is the $\psi_{\rho,\beta}$-value of $(x_+)_i$ where i is chosen as an index in the kernel of the inpainting matrix $\mathbf{K}$, i.e. such that $\mathbf{K}e_i = 0$ where e_i is the i-th standard basis vector. Due to the separability of the regularizer we clearly have that x_+ is not an $\mathcal{R}$-minimizing solution which arises due to the non-convexity of the regularizer.
Moreover, we have observed that if we initialize the values in the kernel close to -1 or 2 then the limit x_+ will have entries at these ϕ-critical points of $\psi_{\rho,\beta}$. This shows that in such cases the solution we obtain in the limit heavily depends on the initialization we choose and that, depending on this initialization, the recovered solution may not be an $\mathcal{R}$-minimizing solution and potentially even a local maximum or a saddle point.

Finally, Figure 4.3 (right) shows the values $|\langle \mathcal{R}'(x_+), e_i\rangle|$ where $(e_i)_i$ is a basis of the kernel of $\mathbf{K}$. Up to numerical accuracy we see that we have $\langle \mathcal{R}'(x_+), e_i\rangle = 0$ for each such index i which shows that $-\mathcal{R}'(x_+) \in \ker(\mathbf{K})^\perp$ as in Lemma 3.9.

5 Conclusion and outlook

We have introduced and studied the concept of critical point regularization, which, opposed to classical variational regularization, considers (ϕ-)critical points of Tikhonov-functionals as regularized solutions. The advantage of this approach is that it completely discards the strong and typically unrealistic assumption of being able to achieve global minimizers of these functionals. Our theory shows that under reasonable assumptions on the involved functionals the resulting method will nevertheless be a stable and convergent regularization method. Further, we have shown that the solutions in the limit $\delta \to 0$ satisfy some form of first order optimality conditions of the constrained optimization problem $\inf_x \mathcal{R}(x)$ subject to the constraint $\mathcal{S}(x, y) = 0$. Besides this, the theory presented here extends the theory of convex functionals by showing that at no point does one require global minimizers, but only points which are close to a global minimum in some sense. For practical applications this means that minimization algorithms do not need to be run until convergence but may be stopped early, if easily verifiable conditions are met. Additionally, under assumptions on the regularizer $\mathcal{R}$ this theory is directly applicable to regularized solutions which are classical critical points of the involved functionals. As such our theory gives stability and convergence results for critical points of potentially non-convex functionals.

Finally, we have provided numerical simulations which support our theoretical findings, i.e. the stability and convergence of critical point regularization. Depending on the algorithm used for obtaining critical points, these numerical examples show that one cannot expect to find global or even local minima which further supports the arguments for the need of a theory based on (ϕ-)critical points, which we have developed in this paper.

As the main concern of this paper was to introduce the concept of using (ϕ-)critical points as regularized solutions, we have not derived any stability- or convergence-rates and studying such rates is subject to future work. Besides this, deriving conditions under which learned regularizers, e.g. [4, 14, 16], give rise to relatively sub-differentiable functions is also subject of future work.

References

[1] R. Acar and C. R. Vogel. Analysis of bounded variation penalty methods for ill-posed problems. *Inverse problems*, 10(6):1217, 1994.

[2] B. Amos, L. Xu, and J. Z. Kolter. Input convex neural networks. In *International Conference on Machine Learning*, pages 146–155. PMLR, 2017.

[3] S. Antholzer and M. Haltmeier. Discretization of learned NETT regularization for solving inverse problems. *arXiv:2011.03627*, 2020.

[4] S. Antholzer, J. Schwab, J. Bauer-Marschallinger, P. Burgholzer, and M. Haltmeier. NETT regularization for compressed sensing photoacoustic tomography. In *Photons Plus Ultrasound: Imaging and Sensing 2019*, volume 10878, page 108783B, 2019.

[5] S. Boyd, L. Xiao, and A. Mutapcic. Subgradient methods. *lecture notes of EE392o, Stanford University, Autumn Quarter*, 2004:2004–2005, 2003.

[6] F. H. Clarke. Generalized gradients and applications. *Transactions of the American Mathematical Society*, 205:247–262, 1975.

[7] J. E. Dennis Jr and R. B. Schnabel. *Numerical methods for unconstrained optimization and nonlinear equations*. SIAM, 1996.

[8] S. Durand and M. Nikolova. Stability of the minimizers of least squares with a non-convex regularization. part i: Local behavior. *Applied Mathematics and Optimization*, 53(2):185–208, 2006.

[9] H. W. Engl, M. Hanke, and A. Neubauer. *Regularization of inverse problems*, volume 375 of *Mathematics and its Applications*. Kluwer Academic Publishers Group, Dordrecht, 1996.

[10] M. Grasmair. Generalized bregman distances and convergence rates for non-convex regularization methods. *Inverse Problems*, 26(11):115014, 2010.

[11] M. Grasmair, M. Haltmeier, and O. Scherzer. Sparse regularization with lq penalty term. *Inverse Problems*, 24(5):055020, 2008.

[12] H. Li, J. Schwab, S. Antholzer, and M. Haltmeier. NETT: Solving inverse problems with deep neural networks. *Inverse Probl.*, 36(6):065005, 2020.

[13] S. Lunz, O. Öktem, and C.-B. Schönlieb. Adversarial regularizers in inverse problems. In *Advances in Neural Information Processing Systems*, pages 8507–8516, 2018.

[14] S. Mukherjee, S. Dittmer, Z. Shumaylov, S. Lunz, O. Öktem, and C.-B. Schönlieb. Learned convex regularizers for inverse problems. *arXiv:2008.02839*, 2020.

[15] Y. E. Nesterov. A method for solving the convex programming problem with convergence rate o (1/kˆ 2). In *Dokl. akad. nauk Sssr*, volume 269, pages 543–547, 1983.

[16] D. Obmann, L. Nguyen, J. Schwab, and M. Haltmeier. Augmented nett regularization of inverse problems. *Journal of Physics Communications*, 2021.

[17] R. T. Rockafellar. *Convex analysis*. Princeton university press, 2015.

[18] O. Scherzer, M. Grasmair, H. Grossauer, M. Haltmeier, and F. Lenzen. *Variational methods in imaging*, volume 167 of *Applied Mathematical Sciences*. Springer, New York, 2009.

[19] N. Z. Shor. *Minimization methods for non-differentiable functions*, volume 3. Springer Science & Business Media, 2012.

[20] I. Singer. *Abstract convex analysis*, volume 25. John Wiley & Sons, 1997.