



Circumcentering Reflection Methods for Nonconvex Feasibility Problems

Neil D. Dizon¹ · Jeffrey A. Hogan¹ · Scott B. Lindstrom² 

Received: 22 February 2021 / Accepted: 28 December 2021
© The Author(s) 2022

Abstract

Recently, circumcentering reflection method (CRM) has been introduced for solving the feasibility problem of finding a point in the intersection of closed constraint sets. It is closely related with Douglas–Rachford method (DR). We prove local convergence of CRM in the same prototypical settings of most theoretical analysis of regular nonconvex DR, whose consideration is made natural by the geometry of the phase retrieval problem. For the purpose, we show that CRM is related to the method of subgradient projections. For many cases when DR is known to converge to a feasible point, we establish that CRM locally provides a better convergence rate. As a root finder, we show that CRM has local convergence whenever Newton–Raphson method does, has quadratic rate whenever Newton–Raphson method does, and exhibits superlinear convergence in many cases when Newton–Raphson method fails to converge at all. We also obtain explicit regions of convergence. As an interesting aside, we demonstrate local convergence of CRM to feasible points in cases when DR converges to fixed points that are not feasible. We demonstrate an extension in higher dimensions, and use it to obtain convergence rate guarantees for sphere and subspace feasibility problems. Armed with these guarantees, we experimentally discover that CRM is highly sensitive to compounding numerical error that may cause it to achieve worse rates than those guaranteed by theory. We then introduce a numerical modification that enables CRM to achieve the theoretically guaranteed rates. Any future works that study CRM for product space formulations of feasibility problems should take note of this sensitivity and account for it in numerical implementations.

Keywords Douglas–Rachford · Feasibility · Projection methods · Reflection methods · Iterative methods · Circumcentering

✉ Scott B. Lindstrom
scott.lindstrom@curtin.edu.au

Neil D. Dizon
neilkristofer.dizon@uon.edu.au

Jeffrey A. Hogan
jeff.hogan@newcastle.edu.au

¹ The University of Newcastle Australia, Callaghan, Australia

² Hong Kong Polytechnic University, Hung Hom, Hong Kong

1 Introduction

The Douglas–Rachford method (DR) is frequently used to solve *feasibility problems* of the form

$$\text{find } \mathbf{x} \in A \cap B, \quad (1)$$

where, here and throughout, A and B are closed subsets of a finite dimensional Hilbert space \mathcal{H} and $A \cap B \neq \emptyset$. For such problems the method consists of iterating the DR operator, which is an averaged composition of two over-relaxed projection operators defined as follows:

$$T_{A,B} := \frac{1}{2}R_B R_A + \frac{1}{2}\text{Id} \quad \text{with} \quad R_C := 2P_C - \text{Id}, \quad (2)$$

where, here and throughout, Id is the identity map and the projection map P_S is as defined below in Eq. 3. The DR operator owes its colloquial name to its indirect introduction in the context of nonlinear heat flow problems [28], though it was independently discovered by Fienup in the nonconvex setting of *phase retrieval* [30], and so it has been known under various other names [39].

Its broader versatility in the nonconvex context was highlighted in by Elser and Gravel [34], who applied it to solve various nonconvex combinatorial problems modeled as feasibility problems with stochastic constraints. The method has since been applied to a host of other discrete feasibility problems, including Sudoku puzzles [3, 5], matrix completion [4, 21], graph coloring [6, 7], and bit retrieval [29], among others. For a more comprehensive overview of its history, including the broader context of DR as a splitting method in solving optimization problems, see for example, [39]. For more on the use of DR for solving both nonconvex and convex feasibility problems, refer to [9].

The aforementioned seminal work of Elser and Gravel [34] piqued the interests of Borwein and Sims, who in 2011 made the first rigorous attempt at analysing the behaviour of DR in the nonconvex setting of hypersurfaces [20]. The spiraling convergence pattern they observed characterizes performance when DR is applied to many other nonconvex hypersurface feasibility problems, which we recall in Section 4.

In the convex setting, the idea of *circumcentering* with the reflections has been recently introduced [12, 13, 15, 17]. Other methods have also been designed, based on using the past performance to predict future iterates [38, 44]. What motivates the *present* work is the tendency of splitting methods to solve *nonconvex* problems [2, 19, 20, 31, 36, 40], and the appetizing prospect of accelerating convergence in the nonconvex setting.

For the more general monotone inclusion problem, Douglas–Rachford method is dual to the *Alternating Direction Method of Multipliers (ADMM)* [32, 39]. This motivated the recent introduction of a class of novel methods [38] that, for some feasibility problems, includes CRM. The author introduced this class with the motivation of building algorithms that are primal/dual implementable. The simplest of them, L_T , may also be used for a feasibility problem. We will include it for comparison in our computed examples; for some problems, it exhibits apparently quadratic convergence, which is certainly interesting. However, as L_T was derived through a somewhat different framework to what we develop here, we will not make a theoretical comparison to it.

1.0.1 Goal

The goal of the present work is to furnish local convergence analysis for CRM in related nonconvex settings to those considered for DR, namely, the case when one set is a hyperplane and the other is a hypersurface that can be represented locally by the graph of a function [2, 18–20, 24, 40]. The purpose and value of this investigation (and those others like it) is, ultimately, not to develop a superior root finder on \mathbb{R} , but rather to understand the performance of CRM more generally. Since the construction of a single step of CRM is always computed in a 2-dimensional subspace, the study of 2-dimensional problems offers us such insights, as 2-dimensional problems are often prototypical of the 2-dimensional slices of problems in higher dimensions. This is why they have been used so often in the literature. As a natural example, we furnish Examples 1 and 2, which show that our results on plane curves admit quadratic convergence rate guarantees for spheres and subspaces. When our numerical experiment in Example 2 at first achieves only a linear rate, we then *know*, because a quadratic rate is *guaranteed* by theory, that the observed linear rate is a consequence of very small numerical errors. Armed with this knowledge, we locate the source of the error, and compensate for it to recover the superior convergence rate guaranteed by the theory. The discovery of this sensitivity to small numerical error will be invaluable to further efforts to study CRM with the usual product space formulation of the feasibility problem, for reasons we explain.

1.1 Outline

The remainder of this paper is outlined as follows. In Section 2, we provide the preliminaries on Douglas–Rachford and CRM. In Section 3, we introduce a modified version of CRM that is generically proper, and show that it has a reasonable fixed point property (Proposition 1). In Section 4, we show that CRM, in the case when the sets are a hyperplane and the graph of a function on \mathbb{R}^n , is related to subgradient projection on a lower dimensional space (Theorem 2). In \mathbb{R}^2 , the subgradient projection method is just Newton–Raphson method. In Section 5, we establish local convergence of CRM to a feasible point in the hypersurface settings considered previously for DR, along with convergence rate guarantees that are quadratic in many cases (Theorem 3). Our analysis, in cases with quadratic convergence, exploits the connection with Newton–Raphson method. More interestingly, in many cases when Newton–Raphson method fails to converge at all, we use the generalized angle bisector theorem to show local superlinear convergence of CRM (Lemma 5) in \mathbb{R}^2 . In Section 6, we provide some examples to show one way in which our rate guarantees from Section 5 may extend to guarantees for certain problems in \mathbb{R}^n . Specifically, we show quadratic convergence rates for spheres and affine subspaces, the problem that has gained particular interest for being prototypical of phase retrieval. We also provide numerical evidence that the convergence rate for CRM when one set is a subspace can be particularly sensitive to compounding numerical error, and we explain how we overcame this sensitivity to achieve the theoretical rate. We conclude in Section 7.

2 Preliminaries

Splitting methods such as the Douglas–Rachford method (DR) are frequently used to solve feasibility problems as in Eq. 1. When the sets of interests are closed, we often employ

iterative algorithms that make use of the projector operator

$$\mathbb{P}_S \mathbf{x} := \left\{ \mathbf{z} \in S : \|\mathbf{x} - \mathbf{z}\| = \inf_{\mathbf{v} \in S} \|\mathbf{x} - \mathbf{v}\| \right\}.$$

In the nonconvex setting, \mathbb{P}_S is a set-valued map where image values may contain more than one point. For the cases we will consider in this paper, \mathbb{P}_S is always nonempty, and we simplify the exposition by working with a selector

$$P_S : \mathcal{H} \rightarrow S : \mathbf{x} \mapsto P_S \mathbf{x} \in \mathbb{P}_S \mathbf{x}. \tag{3}$$

The classical result for the feasibility problem using DR, when A and B are closed and convex, is a consequence of a more general result of Lions & Mercier [41]. We provide the feasibility-specific version.

Theorem 1 (Lions & Mercier [41]) *Let $A, B \subset \mathcal{H}$ be closed, convex, and nonempty, with $A \cap B \neq \emptyset$ and the sum of their normal cone operators $N_A + N_B$ also maximal monotone. For any $\mathbf{x}_0 \in \mathcal{H}$, the sequence given by $\mathbf{x}_{n+1} = T_{A,B} \mathbf{x}_n$ converges to some $\mathbf{v} \in \mathcal{H}$ as $n \rightarrow \infty$ such that $P_A \mathbf{v} \in A \cap B$.*

Note that the condition that $N_A + N_B$ be monotone may be relaxed for the convex feasibility problem [10, Fact 5.9]. The operator described in Eq. 2 is called the Douglas–Rachford operator.

Generically, the fixed points of this operator may not themselves be feasible; see [39, Figure 5b] for a pictorial example. This example also illustrates the useful fact that, even in the nonconvex context, fixed points satisfy (see [39, Proposition 3.1])

$$(\mathbf{x} \in \text{Fix } T_{A,B}) \implies P_A \mathbf{x} \cap A \cap B \neq \emptyset. \tag{4}$$

Various extensions of DR to accommodate N -set (where $N > 2$) feasibility problems have also been considered. The most utilized is the product space reformulation [34, 43]. With such a reformulation, convergence is guaranteed for N closed and convex sets (with nonempty intersection) by the two set result in Theorem 1. An augmented discussion of the above details are given in the recent survey article [39].

2.1 Circumcentering

Given three points $\mathbf{u}, \mathbf{v}, \mathbf{w} \in \mathcal{H}$, we denote the circumcenter by $C(\mathbf{u}, \mathbf{v}, \mathbf{w})$ and define it to be the point equidistant to $\mathbf{u}, \mathbf{v}, \mathbf{w}$, and lying on the affine subspace they generate.

It may be readily verified that when \mathbf{u}, \mathbf{v} and \mathbf{w} are not colinear, $C(\mathbf{u}, \mathbf{v}, \mathbf{w})$ is the intersection of the perpendicular bisectors of the sides of the triangle formed by \mathbf{u}, \mathbf{v} and \mathbf{w} . Figure 1 illustrates that $C(\mathbf{u}, \mathbf{v}, \mathbf{w})$ is not necessarily contained within the convex hull of the triangle formed by \mathbf{u}, \mathbf{v} , and \mathbf{w} .

When $\{\mathbf{u}, \mathbf{v}, \mathbf{w}\}$ has cardinality 1, the definition clearly implies that $C(\mathbf{u}, \mathbf{v}, \mathbf{w}) = \mathbf{u} = \mathbf{v} = \mathbf{w}$. When $\{\mathbf{u}, \mathbf{v}, \mathbf{w}\}$ has cardinality 2, $C(\mathbf{u}, \mathbf{v}, \mathbf{w})$ is the average of the two distinct points.

When ‘‘circumcentering’’ the reflections, we compute a new iterate by taking the circumcenter of $\mathbf{x}, R_A \mathbf{x}$, and $R_B R_A \mathbf{x}$. Following [16], we denote this operation by

$$\text{CRM}(\mathbf{x}) := C(\mathbf{x}, R_A \mathbf{x}, R_B R_A \mathbf{x}). \tag{5}$$

The case where $\mathbf{x}, R_A \mathbf{x}$, and $R_B R_A \mathbf{x}$ are distinct and colinear does not occur when the sets in question are intersecting affine subspaces as in [15, 17]. It may also be avoided when the product space method is used for convex feasibility problems as in [16]. In such a case, the




Fig. 1 The circumcenter of a triangle

operator CRM is said to be *proper*. Sufficient conditions for the operator CRM to be proper are given in [12, 13]. For a convex example where the operator CRM is not proper, let A be the unit ball in \mathbb{R}^2 , $B = \{(\lambda, 3/4) \mid \lambda \in \mathbb{R}\}$, and $\mathbf{x} = (0, 2)$. Then $R_A\mathbf{x} = (0, 0)$ and $R_B R_A\mathbf{x} = (0, 3/2)$, and so $\mathbf{x}, R_A\mathbf{x}, R_B R_A\mathbf{x}$ are distinct and colinear.

For explicit formulas for computing the circumcenter, and for a generalization of the circumcenter of a triangle to the circumcenter of sets containing finitely many points, see [13]. For extensions of linear convergence results to infinite dimensional spaces, refer to [14]. A primal/dual centering approach that does not make use of reflection substeps was introduced in [38].

3 A Generically Implementable Nonconvex Adaptation

For nonconvex feasibility problems, CRM generically fails to be proper, and so we must choose a reasonable definition for the mapping in the nonconvex setting. This is important to ensure numerical stability in the computation, because applying the circumcentering operator to colinear (or nearly colinear) substeps may result in errors that disrupt computation. For a nonconvex problem, we may not know when this colinearity or near-colinearity might occur. The definition, therefore, should rely on conditions that are easy to check from a computational standpoint.

Fortunately, a clear choice presents itself. When $\mathbf{x}, R_A\mathbf{x}$, and $R_B R_A\mathbf{x}$ are colinear, the possibilities are as follow.

- (i) $R_B\mathbf{x} = R_A\mathbf{x} = \mathbf{x}$, in which case $\mathbf{x} \in A \cap B \cap \text{Fix } T_{A,B}$ and $\text{CRM}(\mathbf{x}) = T_{A,B}\mathbf{x}$.
- (ii) $R_B R_A\mathbf{x} \neq R_A\mathbf{x} = \mathbf{x}$ or $R_B R_A\mathbf{x} = R_A\mathbf{x} \neq \mathbf{x}$, in which case the average of the two distinct points is just $\frac{1}{2}\mathbf{x} + \frac{1}{2}R_B R_A\mathbf{x} = T_{A,B}\mathbf{x}$, and so again $\text{CRM}(\mathbf{x}) = T_{A,B}\mathbf{x}$.

- (iii) $R_B R_A \mathbf{x} = \mathbf{x} \neq R_A \mathbf{x}$, in which case $\mathbf{x} \in \text{Fix } T_{A,B}$, and so $P_A \mathbf{x} \in A \cap B$. In this case, $\text{CRM}(\mathbf{x}) = \frac{1}{2} R_A \mathbf{x} + \frac{1}{2} \mathbf{x} = P_A \mathbf{x} \in A \cap B$.
- (iv) $R_B R_A \mathbf{x}$, $R_A \mathbf{x}$, and \mathbf{x} are distinct, in which case $\text{CRM}(\mathbf{x}) = \emptyset$ while $T_{A,B} \mathbf{x} \neq \emptyset$.

Altogether, in cases (i) and (ii) $\text{CRM}(\mathbf{x})$ and $T_{A,B} \mathbf{x}$ coincide, and in case (iii) it does not matter whether one updates with $\text{CRM}(\mathbf{x})$ or $T_{A,B} \mathbf{x}$, since $P_A \mathbf{x}$ solves the feasibility problem. We choose to update with $T_{A,B} \mathbf{x}$ in case (iv), which is consistent but is also a reasonable choice, given what is known about the “searching” behaviour of DR for many nonconvex problems (specifically, it sometimes walks in a straight line before finding the local basin; see [19]). Finally, for the sake of simplicity, we will also choose to “update” with $T_{A,B} \mathbf{x}$ in case (iii); in this way, our definition differs in the convex setting from that in [17], but it does not differ in a consequential way, since case (iii) only occurs when we have already solved the problem. The complete definition of our generically proper circumcentering reflection operator is

$$C_T : \mathcal{H} \rightarrow \mathcal{H} : \mathbf{x} \mapsto \begin{cases} T_{A,B} \mathbf{x} & \text{if } \mathbf{x}, R_A \mathbf{x}, \text{ and } R_B R_A \mathbf{x} \text{ are colinear;} \\ \text{CRM}(\mathbf{x}) & \text{otherwise.} \end{cases} \tag{6}$$

From a computational standpoint, this definition is easy to employ. One need only specify some small numerical tolerance $\epsilon > 0$ and verify noncolinearity by checking that

$$\frac{|(\mathbf{x} - R_B R_A \mathbf{x}, R_A \mathbf{x} - R_B R_A \mathbf{x})|}{\|\mathbf{x} - R_B R_A \mathbf{x}\| \|R_A \mathbf{x} - R_B R_A \mathbf{x}\|} < 1 - \epsilon,$$

before using the closed form for the circumcenter from [13, Theorems 8.4,8.5].

Since C_T specifies to CRM when $\text{CRM}(\mathbf{x})$ is proper—except in the uninteresting case of (iii) when the feasibility problem is essentially solved—we immediately have the following fixed point result.

Proposition 1 (Fixed points of C_T) *If $\mathbf{x} \in \text{Fix } C_T$ then $\mathbb{P}_A \mathbf{x} \cap A \cap B \neq \emptyset$.*

Proof Let $\mathbf{x} \in \text{Fix } C_T$. Then we have from Eq. 6 that either $\mathbf{x} = T_{A,B} \mathbf{x}$ or $\mathbf{x} = \text{CRM}(\mathbf{x})$. If $\mathbf{x} \in T_{A,B} \mathbf{x}$, then the result follows from Eq. 4. If $\mathbf{x} \in \text{CRM}(\mathbf{x})$, then \mathbf{x} is equidistant from \mathbf{x} , $R_A \mathbf{x}$, and $R_B R_A \mathbf{x}$, and so $\mathbf{x} = R_A \mathbf{x} = R_B R_A \mathbf{x}$. Thus $\mathbf{x} \in A \cap B$, and so $P_A \mathbf{x} = \mathbf{x} \in A \cap B$. □

Definition 1 (Generically proper CRM) Let $\mathbf{x}_0 \in \mathcal{H}$. Define $(\mathbf{x}_n)_{n \in \mathbb{N}}$ by

$$\mathbf{x}_{n+1} := C_T \mathbf{x}_n, \tag{7}$$

where C_T is as in Eq. 6.

For nonconvex problems, DR often has colinear substeps at the beginning of the search; see, for example, the example of Section 5.4 item 3 or the example of the ellipse and line from [19]. However, for many problems, this co-linear case never occurs (e.g. Examples 1, 2, and 3 from this paper), or at least never occurs when the algorithm starts sufficiently near to a solution (e.g. the ellipse and line example in [19]). For this reason, the local analysis of convergence for C_T often reduces to the analysis of CRM. For our part, we do not view “Generically proper CRM” to be a distinct algorithm from CRM, casewise definition notwithstanding. We simply consider it to be the natural way to adapt CRM for a nonconvex feasibility problem. Unless one knows a local theoretical guarantee that the co-linear case will not occur (as we use in our convergence analysis in Section 5), one simply *must*,

for computational safety, include a co-linearity check and have a plan in place for handling the co-linear case.

4 Hypersurface Feasibility Problems

For DR, Borwein and Sims considered in detail the case of a unit sphere A in \mathbb{R}^n and a line B [20]. Based on experimentation with the dynamical geometry software *Cinderella* [1], they hypothesized global convergence of the sequence for starting points not on the *singular set* – the line perpendicular to B and passing through the center of A . Aragón Artacho and Borwein later provided a conditional proof based on the piecewise study of regions [2], and Benoist showed convergence definitively by constructing a Lyapunov function [18].

With the proof of Benoist [18], the case of a 2-sphere and a line was mostly resolved, though Borwein and Sims’ conjecture of chaos on the singular manifold [20] was later disproven in the seemingly different context of [11], where Bauschke, Dao, and Lindstrom proved it to be aperiodic but fully describable in terms of generalized Beatty sequences. Two generalizations of the 2-sphere were considered in [19]. In this setting, the singular set has nonzero measure, and the dynamical system is characterized by basins of varying periodicities.

In contradistinction, as long as x_0 is not in the singular set of measure zero, CRM exhibits global convergence for any configuration of an ellipse and line, and spiralling is entirely absent. This is actually implicitly proven in [16, see Remark 1], because the CRM sequence, after the first update, coincides with the CRM sequence for the convex problem of the ball and a hyperplane. Figure 2 is a representative example of what the behaviour looks like, and it suggests the following result.

Lemma 1 *Let $x \in \mathcal{H}$. The following hold.*

1. *If A is a hyperplane and x , $R_A x$, $R_B R_A x$ are not colinear, then $C_T x \in A$;*




Fig. 2 Several iterates of CRM for a sphere and line

2. If B is a hyperplane and \mathbf{x} , $R_A\mathbf{x}$, $R_B R_A\mathbf{x}$ are not colinear, then $C_T\mathbf{x} \in B$;
3. If B is a hyperplane and $\mathbf{x} = R_A\mathbf{x}$, then $C_T\mathbf{x} \in B$;
4. If B is an affine subspace and $\mathbf{x} \in B$ and \mathbf{x} , $R_A\mathbf{x}$, $R_B R_A\mathbf{x}$ are not colinear, then $C_T\mathbf{x} \in B$.

Proof (i): Since \mathbf{x} , $R_A\mathbf{x}$, and $R_B R_A\mathbf{x}$ are not colinear, $C_T\mathbf{x} = \text{CRM}(\mathbf{x})$, and so $C_T\mathbf{x}$ is equidistant from $R_A\mathbf{x}$ and $R_B R_A\mathbf{x}$. Since $R_A\mathbf{x}$ and $R_B R_A\mathbf{x}$ are distinct and A is a hyperplane, we first claim that A is the set of points equidistant from $R_A\mathbf{x}$ and $R_B R_A\mathbf{x}$. To see why, we assume without loss of generality that $\mathbf{0} \in A$, to simplify our notation. We can decompose \mathbf{x} into components that are in A and its orthogonal complement A^\perp , that is, $\mathbf{x} = \mathbf{z}_A + \mathbf{z}_{A^\perp}$. This implies that $P_{A^\perp}\mathbf{x} = \mathbf{z}_{A^\perp}$, and so $R_A\mathbf{x} = \mathbf{z}_A - \mathbf{z}_{A^\perp}$. For any point $\mathbf{y} \in A$, the Pythagorean theorem yields

$$\|\mathbf{y} - \mathbf{x}\|^2 = \|\mathbf{y} - \mathbf{z}_A\|^2 + \|\mathbf{z}_{A^\perp}\|^2 = \|\mathbf{y} - R_A\mathbf{x}\|^2.$$

Since \mathbf{y} is arbitrary, then any point \mathbf{y} is equidistant from \mathbf{x} and $R_A\mathbf{x}$. Suppose now that $\mathbf{y} \in \mathcal{H}$ is equidistant from \mathbf{x} and $R_A\mathbf{x}$. We may similarly write $\mathbf{y} = \mathbf{y}_A + \mathbf{y}_{A^\perp}$ and obtain that

$$\begin{aligned} \|\mathbf{y}_A - \mathbf{z}_A\|^2 + \|\mathbf{y}_{A^\perp} - \mathbf{z}_{A^\perp}\|^2 &= \|\mathbf{y}_A + \mathbf{y}_{A^\perp} - \mathbf{z}_A - \mathbf{z}_{A^\perp}\|^2 = \|\mathbf{y} - \mathbf{x}\|^2 \stackrel{(\star)}{=} \|\mathbf{y} - R_A\mathbf{x}\|^2 \\ &= \|\mathbf{y}_A + \mathbf{y}_{A^\perp} - \mathbf{z}_A + \mathbf{z}_{A^\perp}\|^2 = \|\mathbf{y}_A - \mathbf{z}_A\|^2 + \|\mathbf{y}_{A^\perp} + \mathbf{z}_{A^\perp}\|^2, \end{aligned}$$

and so $\|\mathbf{y}_{A^\perp} - \mathbf{z}_{A^\perp}\|^2 = \|\mathbf{y}_{A^\perp} + \mathbf{z}_{A^\perp}\|^2$. (8)

Here (\star) uses the equidistance assumption. Notice that Eq. 8 forces $\langle \mathbf{z}_{A^\perp}, \mathbf{y}_{A^\perp} \rangle = 0$, and since A^\perp is one-dimensional, then $\mathbf{z}_{A^\perp} = \mathbf{0}$ or $\mathbf{y}_{A^\perp} = \mathbf{0}$. If $\mathbf{z}_{A^\perp} = \mathbf{0}$, then $\mathbf{x} = R_A\mathbf{x}$, which is a contradiction. Therefore $\mathbf{y}_{A^\perp} = \mathbf{0}$, and so $\mathbf{y} \in A$. Altogether, the set of equidistant points is exactly A , and so $C_T\mathbf{x} \in A$.

(ii): Since \mathbf{x} , $R_A\mathbf{x}$, and $R_B R_A\mathbf{x}$ are not colinear, $C_T\mathbf{x} = \text{CRM}(\mathbf{x})$, and so $C_T\mathbf{x}$ is equidistant from $R_A\mathbf{x}$ and $R_B R_A\mathbf{x}$. Since $R_A\mathbf{x}$ and $R_B R_A\mathbf{x}$ are distinct and B is a hyperplane, we first claim that B is the set of points equidistant from $R_A\mathbf{x}$ and $R_B R_A\mathbf{x}$. To see why, we assume without loss of generality that $\mathbf{0} \in B$, to simplify our notation. We can decompose $R_A\mathbf{x}$ into components that are in B and in its orthogonal complement B^\perp , that is, $R_A\mathbf{x} = \mathbf{z}_B + \mathbf{z}_{B^\perp}$. This implies that $P_B R_A\mathbf{x} = \mathbf{z}_B$, and so $R_B R_A\mathbf{x} = \mathbf{z}_B - \mathbf{z}_{B^\perp}$ and $P_{B^\perp} R_B R_A\mathbf{x} = \mathbf{z}_{B^\perp}$. For any point $\mathbf{y} \in B$, the Pythagorean theorem yields

$$\|\mathbf{y} - R_A\mathbf{x}\|^2 = \|\mathbf{y} - \mathbf{z}_B\|^2 + \|\mathbf{z}_{B^\perp}\|^2 = \|\mathbf{y} - R_B R_A\mathbf{x}\|^2.$$

Since \mathbf{y} is arbitrary, then any point $\mathbf{y} \in B$ is equidistant from $R_A\mathbf{x}$ and $R_B R_A\mathbf{x}$. Suppose now that $\mathbf{y} \in \mathcal{H}$ is equidistant from $R_A\mathbf{x}$ and $R_B R_A\mathbf{x}$. We may similarly write $\mathbf{y} = \mathbf{y}_B + \mathbf{y}_{B^\perp}$ and obtain that

$$\begin{aligned} \|\mathbf{y}_B - \mathbf{z}_B\|^2 + \|\mathbf{y}_{B^\perp} - \mathbf{z}_{B^\perp}\|^2 &= \|\mathbf{y}_B + \mathbf{y}_{B^\perp} - \mathbf{z}_B - \mathbf{z}_{B^\perp}\|^2 = \|\mathbf{y} - R_A\mathbf{x}\|^2 \stackrel{(\star)}{=} \|\mathbf{y} - R_B R_A\mathbf{x}\|^2 \\ &= \|\mathbf{y}_B + \mathbf{y}_{B^\perp} - \mathbf{z}_B + \mathbf{z}_{B^\perp}\|^2 = \|\mathbf{y}_B - \mathbf{z}_B\|^2 + \|\mathbf{y}_{B^\perp} + \mathbf{z}_{B^\perp}\|^2, \end{aligned}$$

and so $\|\mathbf{y}_{B^\perp} - \mathbf{z}_{B^\perp}\|^2 = \|\mathbf{y}_{B^\perp} + \mathbf{z}_{B^\perp}\|^2$. (9)

Here (\star) uses the equidistance assumption. Notice that Eq. 9 forces $\langle \mathbf{z}_{B^\perp}, \mathbf{y}_{B^\perp} \rangle = 0$, and since B^\perp is one-dimensional then $\mathbf{z}_{B^\perp} = \mathbf{0}$ or $\mathbf{y}_{B^\perp} = \mathbf{0}$. If $\mathbf{z}_{B^\perp} = \mathbf{0}$, then $R_A\mathbf{x} = R_B R_A\mathbf{x}$, which is a contradiction. Therefore $\mathbf{y}_{B^\perp} = \mathbf{0}$, and so $\mathbf{y} \in B$. Altogether, the set of equidistant points is B , and so $C_T\mathbf{x} \in B$.

(iii): Since $\mathbf{x} = R_A\mathbf{x}$, we have $C_T\mathbf{x} = T_{A,B}\mathbf{x}$ which expands to

$$C_T\mathbf{x} = (1/2)\mathbf{x} + (1/2)R_B R_A\mathbf{x} = (1/2)\mathbf{x} + (1/2)R_B\mathbf{x} = P_B\mathbf{x},$$

where the second equality uses the condition that $\mathbf{x} = R_A\mathbf{x}$. Thus, $C_T\mathbf{x} \in B$.

(iv): Let L be the set of points that are equidistant from $R_A\mathbf{x}$ and $R_B R_A\mathbf{x}$. To show that $C_T(\mathbf{x}) \in B$, it suffices to show that

$$L \cap \text{aff}(\mathbf{x}, R_A\mathbf{x}, R_B R_A\mathbf{x}) \subset B.$$

Because $\mathbf{x}, R_A\mathbf{x}, R_B R_A\mathbf{x}$ are not colinear, it must be true that $R_A\mathbf{x} \neq R_B R_A\mathbf{x}$, and so the inclusion

$$L \cap \text{aff}(\mathbf{x}, R_A\mathbf{x}, R_B R_A\mathbf{x}) \subsetneq \text{aff}(\mathbf{x}, R_A\mathbf{x}, R_B R_A\mathbf{x})$$

is not an equality. Combining with the fact that the affine subspace on the right side is 2-dimensional, we have

$$\dim(L \cap \text{aff}(\mathbf{x}, R_A\mathbf{x}, R_B R_A\mathbf{x})) = 1.$$

Notice that in our proof of (ii), the argument that shows that B is contained in the subset of equidistant points (i.e. $B \subset L$) depended only upon B being an affine subspace (only the reverse inclusion \supset relied upon B being a hyperplane) and so $B \subset L$ holds now by the same argument. More specifically, since $\mathbf{x} \in B$ and $P_B R_A\mathbf{x} \in B$, we have that

$$B^* := \text{aff}\{\mathbf{x}, P_B R_A\mathbf{x}\} \subset B \subset L.$$

Of course, we also have

$$B^* \subset \text{aff}\{\mathbf{x}, R_A\mathbf{x}, R_B R_A\mathbf{x}\}.$$

Moreover, since $\mathbf{x}, R_A\mathbf{x}$, and $R_B R_A\mathbf{x}$ are not colinear B^* is of dimension 1. Altogether, we have that B^* is a one-dimensional affine subspace contained in $L \cap \text{aff}(\mathbf{x}, R_A\mathbf{x}, R_B R_A\mathbf{x})$. As the two affine subspaces are of dimension 1, and one is contained in the other, they must be equal. This concludes the result. □

Of the above proposition, the following should be noted:

1. Items (i) and (ii) of the above proposition are not necessarily true if the set that is assumed to be a hyperplane is, instead, only assumed to be an affine subspace of arbitrary dimension. In case (ii), for instance, if one assumes B is only an affine subspace instead of a hyperplane, then one can only obtain $\mathbf{y}_{B^\perp} \in \{\mathbf{z}_{B^\perp}\}^\perp$. As an example, the perpendicular bisector of a reflection across a line in \mathbb{R}^3 is a hyperplane containing the line; the circumcenter will be in the hyperplane, but not necessarily the line.
2. Item (iv) is already known in the case when A is convex, due to [16, Lemma 3]. That result uses in its proof the convexity of A ; if it did not, then we would simply cite the extension as vacuous.

Throughout the remaining part of this section, we write $\mathbf{x} = (x, x') \in \mathbb{R}^\eta \times \mathbb{R} = \mathbb{R}^{\eta+1}$ where x is the component of \mathbf{x} in \mathbb{R}^η and x' is the component of \mathbf{x} in \mathbb{R} . Furthermore, we set $A, B \subset \mathbb{R}^{\eta+1}$ as

$$A = \text{gra } f = \{(y, f(y)) \mid y \in \text{dom } f \subset \mathbb{R}^\eta\} \quad \text{and} \quad B = \mathbb{R}^\eta \times \{0\}.$$

Here, $\text{gra } f$ denotes the *graph* of a function $f : \mathbb{R}^\eta \rightarrow \mathbb{R}$. Furthermore, we use $\partial^0 f$ to denote the symmetric subdifferential of f , which is given by $\partial^0 f := \partial f \cup (-\partial(-f))$, where $\partial f(x) := \{x^* \in X \mid (x^*, 1) \in N_{\text{epi } f}(x, f(x))\}$ is the *limiting subdifferential* of f at x [24, Section 2.3]. The following lemma establishes a relationship between CRM and subgradient projections in \mathbb{R}^η .

Theorem 2 Let $\mathbf{x} = (x, x') \in \mathbb{R}^\eta \times \mathbb{R}$, $B = \{(t, 0) : t \in \mathbb{R}^\eta\}$ and $A = \text{gra } f$ where $f : \mathbb{R}^\eta \rightarrow \mathbb{R}$ is proper and has a closed graph. Suppose further that f is Lipschitz continuous locally at y where $(y, f(y)) =: P_A \mathbf{x}$. Then the following hold.

(i) If $\mathbf{x}, R_A \mathbf{x}, R_B R_A \mathbf{x}$ are not colinear, then

$$C_T \mathbf{x} = \left(y - \frac{f(y)}{\|y^*\|^2} y^*, 0 \right)$$

where $y^* \in \partial^0(f(y))$ satisfies $x = y + (f(y) - x')y^*$

(ii) Otherwise, $C_T \mathbf{x} = (y, x' - f(y))$.

Proof First note that the existence of y^* that satisfies

$$x = y + (f(y) - x')y^* \tag{10}$$

is assured by [24, Lemma 3.4].

(i) The assumption of non-colinearity forces $C_T \mathbf{x} = \text{CRM}(\mathbf{x})$. We must have $y^* \neq 0$, since $(y^* = 0) \implies (x = y)$, which forces $\mathbf{x}, R_A \mathbf{x}, R_B R_A \mathbf{x}$ to be colinear, a contradiction. Similarly note also that $(f(y) - x' = 0) \implies (x = y)$, which forces the same contradiction; therefore $f(y) - x' \neq 0$. Any point $\mathbf{w} = (w, w')$ in the perpendicular bisector of $(\mathbf{x}, R_A \mathbf{x})$ satisfies

$$\begin{aligned} 0 &= \langle \mathbf{x} - P_A \mathbf{x}, (w, w') - P_A \mathbf{x} \rangle \\ &= \langle (x - y, x' - f(y)), (w - y, w' - f(y)) \rangle \\ &= \langle ((f(y) - x')y^*, x' - f(y)), (w - y, w' - f(y)) \rangle \quad (\text{using Eq. 10}) \\ &= (f(y) - x') \langle (y^*, -1), (w - y, w' - f(y)) \rangle \\ &\stackrel{(a)}{=} (f(y) - x') \langle (w - y, y^*) - w' + f(y) \rangle \end{aligned}$$

$$\iff (w, w') \stackrel{(b)}{=} (y + \lambda, f(y) + \langle \lambda, y^* \rangle) \text{ for some } \lambda \in \mathbb{R}^\eta.$$

To see why (b) holds, remember that $f(y) - x' \neq 0$, and so the equality (a) is equivalent to $\langle (w - y, y^*) - w' + f(y) \rangle = 0$, which is equivalent to $w' = f(y) + \langle w - y, y^* \rangle$. Defining $\lambda := w - y$, (b) and (a) are equivalent.

Altogether, we have shown that the perpendicular bisector of the segment $(\mathbf{x}, R_A \mathbf{x})$ is the tangent plane $(y + \lambda, f(y) + \langle \lambda, y^* \rangle)_{\lambda \in \mathbb{R}^\eta} =: H$. Combining with Lemma 1 and the definition of CRM, we have that $\text{CRM}(\mathbf{x}) \in H \cap B \cap \text{aff}(\mathbf{x}, R_A \mathbf{x}, R_B R_A \mathbf{x})$. We also have that

$$\text{aff}(\mathbf{x}, R_A \mathbf{x}, R_B R_A \mathbf{x}) = \{(y + \alpha y^*, \beta) \mid (\alpha, \beta) \in \mathbb{R}^2\}. \tag{11}$$

To see why Eq. 11 holds, let $\mathbf{v} \in \text{aff}(\mathbf{x}, R_A \mathbf{x}, R_B R_A \mathbf{x})$. Then

$$\mathbf{v} = \alpha_x \mathbf{x} + \alpha_{R_A \mathbf{x}} R_A \mathbf{x} + \alpha_{R_B R_A \mathbf{x}} R_B R_A \mathbf{x} \text{ with } \alpha_x + \alpha_{R_A \mathbf{x}} + \alpha_{R_B R_A \mathbf{x}} = 1 \tag{12a}$$

$$= \alpha_x(x, x') + \alpha_{R_A \mathbf{x}}(2y - x, 2f(y) - x') + \alpha_{R_B R_A \mathbf{x}}(2y - x, -2f(y) + x') \tag{12b}$$

$$= (y + \alpha^* y^*, \beta^*) \tag{12c}$$

$$\text{where } \alpha^* := (2\alpha_x - 1)(f(y) - x') \tag{12d}$$

$$\text{and } \beta^* := \alpha_x x' + (2f(y) - x')(\alpha_{R_A \mathbf{x}} - \alpha_{R_B R_A \mathbf{x}}). \tag{12e}$$

Here Eq. 12a simply uses the definition of the affine hull; Eq. 12b uses the definition of B , whereby $R_B R_A \mathbf{x}$ is obtained by taking $R_A \mathbf{x}$ and simply reversing the sign on its $(\eta + 1)$ th coordinate; Eq. 12c uses the fact that $x = y + (f(y) - x')y^*$ and the equality $\alpha_x + \alpha_{R_A \mathbf{x}} + \alpha_{R_B R_A \mathbf{x}} = 1$. Altogether, we have shown $\mathbf{v} = (y + \alpha^* y^*, \beta^*) \in \{(y + \alpha y^*, \beta) \mid (\alpha, \beta) \in \mathbb{R}^2\}$, showing the inclusion \subset in Eq. 11. Moreover, as both affine subspaces in Eq. 11 are of the

same dimension and one is contained in the other, they must be equal, and so the equality holds in Eq. 11.

Altogether, we have

$$\begin{aligned} \text{aff}(\mathbf{x}, R_A \mathbf{x}, R_B R_A \mathbf{x}) \cap H \cap B &= \left\{ \underbrace{(y + \lambda, f(y) + \langle \lambda, y^* \rangle)}_{\text{constraint } H} \mid \underbrace{\lambda = \alpha y^*}_{\text{constraint aff}(\dots)}, \underbrace{f(y) + \langle \lambda, y^* \rangle = 0}_{\text{constraint } B} \right\} \\ &= \left\{ \left(y - \frac{f(y)}{\|y^*\|^2} y^*, 0 \right) \right\}. \end{aligned}$$

This shows the desired result.

(ii): In the colinear case, $C_T \mathbf{x} = T_{A,B} \mathbf{x}$. Computing, one has $R_A \mathbf{x} = (2y - x, 2f(y) - x')$, and $R_B R_A \mathbf{x} = (2y - x, -2f(y) + x')$, and finally $(1/2)R_B R_A \mathbf{x} + (1/2)\mathbf{x} = (y, x' - f(y))$, which shows the result. \square

Remark 1 (A subgradient projections characterization) Notice that $y \mapsto y - \frac{f(y)}{\|y^*\|^2} y^*$ is a step of subgradient projections applied to the function f at y . For a differentiable function $f : \mathbb{R} \rightarrow \mathbb{R}$, it reduces to a step of Newton–Raphson method: $y \mapsto y - f(y)/f'(y)$ where $y \in \mathbb{R}$. Thus, Theorem 2 shows that, for problems when $C_T \mathbf{x} = \text{CRM}(\mathbf{x})$ and $\mathbf{x} \in \mathbb{R}^2$, the method reduces to a step of alternating projections $P_B P_A \mathbf{x}$, followed by a step of subgradient descent for the function f . Figure 3 illustrates this in the case when A is the graph of the function $x \mapsto x/\sqrt{|x|}$ ($x \in \mathbb{R}$) and B is the horizontal axis. Solving the feasibility problem amounts to finding the root of the function $x \mapsto x/\sqrt{|x|}$ ($x \in \mathbb{R}$), a problem that Newton’s method on \mathbb{R} fails to solve.

Interestingly, [38, Proposition 3.6] has described a relationship between CRM and the method of subgradient projections applied for a different function: one that describes the dynamical system admitted by repeated application of the DR operator. This function, called a Lyapunov function, is defined on a lifted space. The author used the theory of Lyapunov functions to motivate their introduction of different algorithm, L_T , which we will include for comparison in our computed examples.

The use of DR to find roots of functions on \mathbb{R} , as well as for $x/\sqrt{|x|}$ ($x \in \mathbb{R}$) in particular, was first considered in [40]. Dao and Tam adapted Benoist’s Lyapunov function approach in order to demonstrate local convergence for non-tangentially intersecting cases thereof [24].

5 A Line and a Plane Curve

Most local convergence results in the nonconvex setting of hypersurfaces have focused on hyperplanes and graphs of functions [2, 18–20, 24, 40] with \mathbb{R}^2 as a setting of particular focus. One motivation is that the *phase retrieval problem* may be thought of as a feasibility problem on n -tuples of \mathbb{R}^2 . Another reason is that DR has variously been observed to spiral in lower dimensional subspaces, a phenomenon theorized to often occur in a lower dimensional affine subspace of \mathbb{R}^n . For example, Aragón Artacho has created an image that shows this behaviour for a line and sphere in \mathbb{R}^3 [22, slide 33]. For CRM, \mathbb{R}^2 is especially a natural context for investigation because the circumcenter construction for \mathbf{x} is performed in the affine hull of \mathbf{x} , $R_A \mathbf{x}$, and $R_B R_A \mathbf{x}$, which in the non-colinear case is always a copy of \mathbb{R}^2 . Thus, sets in \mathbb{R}^2 generically represent slices of sets in the affine hull of reflection substeps for higher dimensional problems. Similar planar results motivated, in part, the introduction of the first primal/dual centering method in [38].




Fig. 3 CRM applied to find a root of $x \mapsto x/\sqrt{|x|}$ where $(x \in \mathbb{R})$

Under mild assumptions, we will show local convergence in \mathbb{R}^2 for the algorithm generated by iteratively applying the map C_T with a curve A and line B . Our approach differs from those in previous works both in terms of the methods used and the results obtained. In particular, the use of the generalized angle bisector theorem suggests a path forward for more complicated problems.

5.0.1 Comparing Approaches

The first theoretical result on the convergence of DR in the nonconvex setting is that of Borwein and Sims [20] who used Perron theorem [20, Theorem 6.1] [35, Corollary 4.7.2] on the stability of almost linear difference equations to show local convergence in the setting where A is a unit sphere in \mathbb{R}^n and B is a line. This approach has since been adapted [19] to show local convergence for plane curves more generally. The strategy relies upon the fact that $T_{A,B}\mathbf{x}$ may be described as a continuous function of the 3-tuple $(\mathbf{x}, R_A\mathbf{x}, R_B R_A\mathbf{x})$. This continuity does not extend to the case of the operator C_T , and so the same approach does not immediately extend to this new context.

The approach of Benoist, Dao, and Tam relies on a Lyapunov function [18, 24, 38]. It is less clear how to adapt such an approach when the spiral itself is obviated by circumcentering the method.

Our approach is to use Theorem 2 and employ trigonometry to show results about the intersections of the tangents taken for the curve A with the line B . Without loss of generality, we let B be the horizontal axis. We then use the observation in Remark 1 that, under mild assumptions, C_T locally behaves like a step of alternating projections for A and B followed by a step of Newton–Raphson method employed to find a root of the function whose graph is the curve A . Figure 4 is helpful in understanding this observation. Here, when B is the horizontal axis and A is the graph of a function $f : \mathbb{R} \rightarrow \mathbb{R}$, we may associate the points $\mathbf{x}_n, \mathbf{x}_{n+1}$, and $(y_n, f(y_n))$ in $\mathbb{R}^2 = \{\mathbf{t} = (t, t') \mid t, t' \in \mathbb{R}\}$ with their horizontal components x_n, x_{n+1} , and y_n . We then have the relationship

$$x_{n+1} = g_f(y_n) := y_n - \frac{f(y_n)}{f'(y_n)}, \quad \text{where } (y_n, 0) = P_B P_A(\mathbf{x}_n). \tag{13}$$

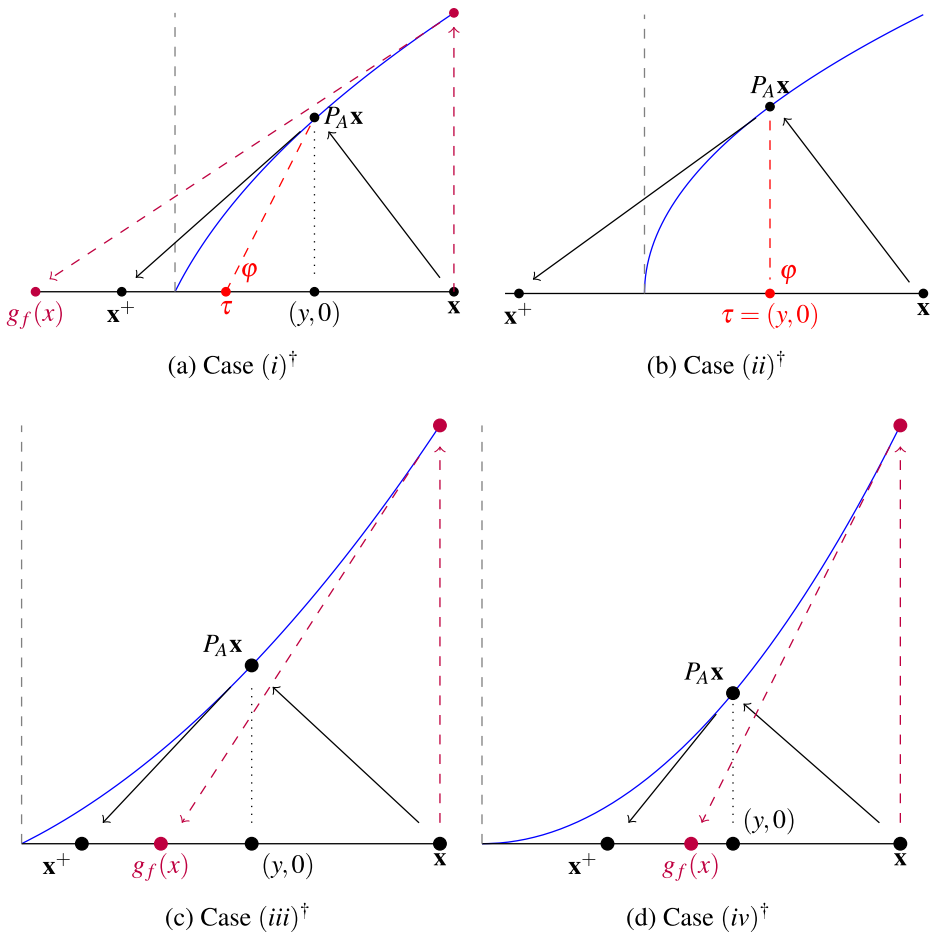


Fig. 4 Cases for proof of local convergence

Here, g_f is the Newton–Raphson operator that is commonly used to search for a root of the function f , and for which convergence results are well known. For the sake of cleanliness in Fig. 4, we abuse notation slightly by assigning the label $g_f(x_n)$ to the point that is actually $(g_f(x_n), 0)$. Note that when f is strictly differentiable at $y \in \mathbb{R}$, $\partial^0 f(y) = f'(y)$ [42, Corollary 1.82], an identity we will implicitly use in this section whenever appealing to Theorem 2.

5.1 Auxiliary Results on Newton–Raphson Method

Because our results will exploit the connection with Newton–Raphson method, we will need two preliminary lemmas that tailor the classical theory to our purpose.

Lemma 2 (When Newton–Raphson rate is quadratic) *Let $\theta : \mathbb{R} \rightarrow \mathbb{R}$ and $g(t) = t - \theta(t)/\theta'(t)$ for $t \in [-\varepsilon, \varepsilon]$ and $\theta(0) = 0$. Suppose there exist $h_{LEFT}, h_{RIGHT} \in C^2[-\varepsilon, \varepsilon]$*

such that

$$(\forall t \in [0, \varepsilon]) \quad h_{RIGHT}(t) = \theta(t) \quad \text{and} \quad (\forall t \in [-\varepsilon, 0]) \quad h_{LEFT}(t) = \theta(t),$$

with $h'_{LEFT}(0), h'_{RIGHT}(0) \in \mathbb{R} \setminus \{0\}$. Then there exists $\delta > 0$ such that for $p_0 \in [-\delta, \delta]$, the sequence defined by $p_n = g(p_{n-1})$, when $n \geq 1$, converges at least quadratically to 0. Moreover, for t sufficiently near 0,

$$|g(t)| < \frac{M}{2}|t|^2,$$

for some $M \in \mathbb{R}$.

Proof By the classical result on Newton–Raphson (see [23, Theorem 2.9]) , there exist $\delta_{LEFT}, \delta_{RIGHT}, M_{LEFT}, M_{RIGHT}$ such that

$$t \in [-\delta_{LEFT}, \delta_{LEFT}] \implies |g_{LEFT}(t)| < \frac{M_{LEFT}}{2}|t|^2$$

$$\text{and } t \in [-\delta_{RIGHT}, \delta_{RIGHT}] \implies |g_{RIGHT}(t)| < \frac{M_{RIGHT}}{2}|t|^2$$

$$\text{where } g_{RIGHT} : t \rightarrow t - h_{RIGHT}(t)/h'_{RIGHT}(t)$$

$$\text{and } g_{LEFT} : t \rightarrow t - h_{LEFT}(t)/h'_{LEFT}(t).$$

Letting $M_{MAX} := \max\{M_{LEFT}, M_{RIGHT}\}$ and $\delta_{MIN} := \min\{\delta_{LEFT}, \delta_{RIGHT}\}$ and using the fact that

$$g(t) = \begin{cases} g_{LEFT}(t) & \text{if } t \leq 0 \\ g_{RIGHT}(t) & \text{if } t \geq 0, \end{cases}$$

we have that

$$t \in [-\delta_{MIN}, \delta_{MIN}] \implies |g(t)| < \frac{M_{MAX}}{2}|t|^2. \tag{14}$$

□

Lemma 3 (When Newton–Raphson rate is linear) *Let $h \in C^m[-\varepsilon, \varepsilon]$ and $p = 0$ be a root of h of multiplicity m . Let $g : t \mapsto t - h(t)/h'(t)$. Then there exists $k < 1$ and $\delta_{LIN} > 0$ such that*

$$t \in [-\delta_{LIN}, \delta_{LIN}] \implies |g(t) - p| < k|t - p|.$$

Proof Using the fact that

$$g'(t) = \frac{h(t)h''(t)}{h'(t)^2},$$

we have that

$$\begin{aligned} g'(p) &= \lim_{t \rightarrow 0} \frac{h(p+t)h''(p+t)}{h'(p+t)^2} \\ &= \lim_{t \rightarrow 0} \frac{\left(\sum_{j=0}^m \frac{h^j(p)}{j!} t^j + O(t^{m+1})\right) \left(\sum_{j=2}^m \frac{h^j(p)}{(j-2)!} t^{j-2} + O(t^{m-1})\right)}{\left(\sum_{j=1}^m \frac{h^j(p)}{(j-1)!} t^{j-1} + O(t^m)\right)^2} \\ &\stackrel{(**)}{=} \lim_{t \rightarrow 0} \frac{\frac{h^m(p)^2}{m!(m-2)!} t^{2m-2} + O(t^{2m-1})}{\frac{h^m(p)^2}{(m-1)!^2} t^{2m-2} + O(t^{2m-1})} \\ &= \frac{(m-1)!^2}{m!(m-2)!} = \frac{m-1}{m} < 1. \end{aligned}$$

Notice that $(\star\star)$ uses the fact that $h^j(p) = 0$ for $j < m$ (since p is a root of h of multiplicity m). Letting $k = \frac{1}{2}(\frac{m-1}{m} + 1)$, we have that, for t sufficiently near to p , $g'(t) < k < 1$. In particular, we may choose δ_{LIN} so that

$$t \in [-\delta_{\text{LIN}}, \delta_{\text{LIN}}] \implies g'(t) < k.$$

Applying [23, Theorem 2.8], we have that

$$t \in [-\delta_{\text{LIN}}, \delta_{\text{LIN}}] \implies |g(t) - p| < k|t - p|,$$

which shows the result. □

5.2 Convergence of CRM: Basic Conditions

Throughout the rest of this section, A is the graph of a function $f : \mathbb{R} \rightarrow \mathbb{R}$, B is the horizontal axis $\mathbb{R} \times \{0\} \subset \mathbb{R}^2$, $\mathbf{x} = (x, x') \in \mathbb{R}^2$, $(y, f(y)) := P_A \mathbf{x}$, and $\mathcal{B}_0(r) := [-r, r]$. As the projection onto B is computable, one can always start with the first iterate $\mathbf{x}_0 \in B$. In particular, we will show that for $\mathbf{x}_0 \in B$ started sufficiently near to the solution, the iterated scheme $\mathbf{x}_n := (C_T)^n \mathbf{x}_0$ reduces to $\mathbf{x}_n := (\text{CRM})^n \mathbf{x}_0$ and converges to the solution $\mathbf{0}$. In particular, the conditions we will impose on f will locally prevent colinearity of \mathbf{x} , $R_A \mathbf{x}$ and $R_B R_A \mathbf{x}$ (as will be shown in Lemma 4) so that we always have Theorem 2(i), which simplifies our analysis greatly. We will consider the case when the following hold.

- (I) the curve A is the graph of a continuous function f and where $\mathbf{0} \in A \cap B$ is an isolated feasible point;
- (II) there exists $\epsilon_1 > 0$ such that 0 is the only root of f on $\mathcal{B}_0(\epsilon_1)$;
- (III) f is continuous on $\mathcal{B}_0(\epsilon_1)$ and differentiable on $\mathcal{B}_0(\epsilon_1) \setminus \{0\}$;
- (IV) there exists $\epsilon_2 > 0$ such that $\text{zer } f' \cap \mathcal{B}_0(\epsilon_2) \subset \{0\}$, that is, f' has no other roots in $\mathcal{B}_0(\epsilon_2)$ except possibly for 0 .

The isolated root condition we have imposed on f excludes such pathological cases as $f : t \mapsto \sin(1/t)$ (where $t \in \mathbb{R} \setminus \{0\}$), while the analogous condition we have imposed on f' further excludes such pathological cases as $f : t \mapsto t \sin(1/t)$ (where $t \in \mathbb{R} \setminus \{0\}$). Finally, we will assume a similar condition about f'' namely

- (V) there exists $\epsilon_3 > 0$ such that f' is continuous and differentiable on $\mathcal{B}_0(\epsilon_3) \setminus \{0\}$ with $\text{zer } f'' \cap \mathcal{B}_0(\epsilon_3) \subset \{0\}$.

In other words, f'' has no other roots in $\mathcal{B}_0(\epsilon_3)$ except, possibly, for 0 . Since the sign of f does not change on $]0, \epsilon_1]$, we may by symmetry work in the first quadrant by assuming (without loss of generality) further that

- (VI) f is positive on $]0, \epsilon_1]$.

When f satisfies the conditions (I)–(VI), we write $\epsilon_f := \min\{\epsilon_1, \epsilon_2/2, \epsilon_3\}$. The reason for the choice of $\epsilon_2/2$ will be apparent in Lemma 4. Since f is positive on $]0, \epsilon_f]$ and f' does not have a root on $]0, \epsilon_f]$, then we further have that f is increasing on $]0, \epsilon_f]$, or equivalently, f' is positive on $]0, \epsilon_f]$.

Remark 2 (When A is a line segment locally) Our conditions have excluded the possibility that there exists $\delta > 0$ such that f'' is zero on $[0, \delta]$. Of course, in this case, the graph of A is a line segment locally, and so we have convergence in a single step, by Theorem 2. Our use of the graph of a single-valued function f to represent A also precludes the possibility

that A is locally a vertical line segment, another case when local convergence is immediate. Thus, we do not lose any (new) insights by disregarding such cases.

5.3 Convergence of CRM

In what follows, $\mathbf{x}_+ := C_T \mathbf{x}$. For the supposed conditions in the following Lemma 4, if one remembers that we are only writing $x \geq 0$ because we may do so by symmetry and without loss of generality, then the conditions we have assumed, in essence, are implied by \mathbf{x} lying in B and within a certain local ball about the solution, a standard assumption used for studying local convergence.

Lemma 4 (Newton–Raphson as a convergence rate bound) *Suppose f satisfies the basic conditions (I)–(VI) and that $\mathbf{x} = (x, 0)$ with $0 < x \leq \epsilon_f$. Then, whenever $\mathbf{0} \neq (y, f(y)) = P_A \mathbf{x}$, the following hold.*

- (i) $y \in]0, x[$
- (ii) $C_T \mathbf{x} = (g_f(y), 0)$
- (iii) If $f''(t) > 0$ for $t \in]0, \epsilon_3]$, then $0 \leq g_f(y) < g_f(x)$;
- (iv) If $f''(t) < 0$ for $t \in]0, \epsilon_3]$, then $g_f(x) < g_f(y) \leq 0$;
- (v) $\frac{\|\mathbf{x}_+\|}{\|\mathbf{x}\|} = \frac{|g_f(y)|}{|x|} \leq \frac{|g_f(x)|}{|x|}$.

Proof (i) Since $x \leq \epsilon_2/2$ and $(0, 0) \in A$, one has $\|\mathbf{x} - P_A(\mathbf{x})\| \leq \|(x, 0) - (0, 0)\| = x \leq \epsilon_2/2$. Since $(x, f(x)) \in A$, one also has $\|\mathbf{x} - P_A(\mathbf{x})\| \leq \|(x, 0) - (x, f(x))\| = f(x)$. It then follows that

$$\|\mathbf{x} - P_A(\mathbf{x})\| \leq \min\{x, f(x)\} \leq \epsilon_2/2. \tag{15}$$

Suppose, for a contradiction, that $P_A(\mathbf{x}) = (t, f(t))$ for $t \notin]0, x]$. If $t > \epsilon_2$, then $\|\mathbf{x} - P_A(\mathbf{x})\| = \|(x - t, -f(t))\| \geq \|t - x\| > \epsilon_f$ which contradicts Eq. 15. Thus $t \leq \epsilon_2$. Since f is monotone increasing on $[0, \epsilon_2]$, any point $(t, f(t)) \in A$ that satisfies $\epsilon_2 > t > x$ must also satisfy $f(t) > f(x)$, which forces $\|\mathbf{x} - P_A(\mathbf{x})\| = \|(x, 0) - (t, f(t))\| > f(t) > f(x)$ contradicting Eq. 15. Lastly, if $t < 0$, one has $\|\mathbf{x} - P_A(\mathbf{x})\| > x$, which is again a contradiction of Eq. 15. Thus, we have established that $y \in]0, x]$. Consequently, we also have that f is Lipschitz continuous on an open ball about y . Since $(y, f(y)) = P_A((x, 0))$, we may apply [24, Lemma 3.4] to obtain Eq. 10, which simplifies to

$$x = y + f(y)f'(y).$$

If $f(y) = 0$, then $y = 0$, which is a case we may ignore. Otherwise, we have $f(y)f'(y) > 0$, and so $x > y$. Thus $y \in]0, x[$.

(ii): Having shown(i), we have that

$$R_A(\mathbf{x}) = (2y - x, 2f(y)), \quad R_B R_A(\mathbf{x}) = (2y - x, -2f(y)), \quad \text{and } \mathbf{x} = (x, 0)$$

are not colinear. Thereupon applying Theorem 2, we obtain (ii).

(iii): By the mean value theorem, there exists $c \in]y, x[$ with

$$(x - y)f'(c) = f(x) - f(y). \tag{16}$$

Combining with the fact that $f''(t) > 0$ for $t \in [y, x]$, we have that f' is monotone increasing on $[y, x]$, and so $f'(y) < f'(c) < f'(x)$. Combining with Eq. 16,

$$\begin{aligned} (x - y)f'(x) &> f(x) - f(y), \\ \text{and so } x - y - \frac{f(x)}{f'(x)} &> -\frac{f(y)}{f'(x)} > -\frac{f(y)}{f'(y)}, \\ \text{which shows } x - \frac{f(x)}{f'(x)} &> y - \frac{f(y)}{f'(y)}. \end{aligned} \tag{17}$$

Applying the mean value theorem again, there exists $c' \in]0, y[$ such that

$$(y - 0)f'(c') = f(y) - f(0) \tag{18}$$

Since f' is monotone increasing on $[0, y]$, we have $f'(c') < f'(y)$. Combining this fact with Eq. 18, we have $yf'(y) \geq f(y)$, and so

$$0 \leq y - f(y)/f'(y). \tag{19}$$

Together, Eq. 17 and Eq. 19 show (iii). Combining with (ii), we have

$$\mathbf{x}_+ = (x_+, 0), \quad \text{and} \quad 0 \leq x_+ = g_f(y) \leq g_f(x),$$

which together yield (v).

(iv): The proof is similar to (iii), with the only change being that f' is monotone *decreasing* instead of *increasing*, which reverses the directions of the analogous inequalities. The proof of (v) then follows in the same way. □

As we have already noted, because f'' has no roots on $]0, \epsilon_f]$, f' is monotone on $]0, \epsilon_f]$. Consequently, if f' is bounded on $]0, \epsilon_f]$, then we have $\lim_{t \downarrow 0} f'(t) \in [0, \infty[$ exists by the monotone convergence theorem, and if f' is unbounded, then $\lim_{t \downarrow 0} f'(t) = \infty$. Taken together with Lemma 4, the conditions we have assumed about f'' leave us with the following four natural cases to consider – those illustrated in Fig. 4.

(i)[†] $f''(t) < 0$ for $t \in [0, \epsilon_f]$, and f is right differentiable at 0 with

$$\lim_{t \downarrow 0} f'(t) \in]0, \infty[.$$

(ii)[†] $f''(t) < 0$ for $t \in [0, \epsilon_f]$, and

$$\lim_{t \downarrow 0} f'(t) = \infty.$$

(iii)[†] $f''(t) > 0$ for $t \in [0, \epsilon_f]$, and $f'(0) \neq 0$.

(iv)[†] $f''(t) > 0$ for $t \in [0, \epsilon_f]$, and $f'(0) = 0$ and there exists $h \in C^m[-\epsilon_f, \epsilon_f]$ such that $h(t) = f(t)$ for $t > 0$ and 0 is a root of h of multiplicity m (for example, any function analytic at 0 whose Taylor series has 0 as a root of multiplicity greater than 1 will meet this criterion).

Lemma 4 essentially shows that the convergence rate of Newton–Raphson serves as an upper bound on the convergence rate of CRM in many cases. Now we use the generalized angle bisector theorem to obtain superlinear results in cases when Newton–Raphson cycles. Here, $\angle(\mathbf{a}, \mathbf{b}, \mathbf{c})$ denotes, as usual, the angle at which the segments \mathbf{ab} and \mathbf{bc} meet.

Lemma 5 (Cases (i)[†] and (ii)[†] superlinear) *Let f satisfy the basic conditions (I)–(VI), $f''(t) < 0$ for $t \in]0, \epsilon_3]$, and $\mathbf{x} = (x, 0)$. Then for any $K > 0$ there exists $\epsilon_{K_f} > 0$ such*

that

$$(x \in]0, \varepsilon_{K_f}]) \implies \frac{\|\mathbf{x}_+\|}{\|\mathbf{x}\|} \leq \frac{1}{K}. \tag{20}$$

Proof We illustrate our construction in Figs. 4a and 4b. We can and do assume that $0 < x < \varepsilon_f$ so that the characterizations in Lemma 4 always hold. For the sake of simplicity, when $\lim_{t \downarrow 0} f'(t) < \infty$, we write $f'(0) := \lim_{t \downarrow 0} f'(t)$. We also denote

$$(\varphi, \tau) := \begin{cases} (\arctan(f'(0)), (y - f(y)/f'(0), 0)) & \text{if } \lim_{t \downarrow 0} f'(t) < \infty \\ (\pi/2, (y, 0)) & \text{otherwise,} \end{cases}$$

where φ is an angle measure and τ is a point. By considering these two cases, we will first show that

$$\exists \delta_f > 0 \text{ such that } (x \in]0, \delta_f]) \implies \cot(\angle(\mathbf{x}, \mathbf{x}_+, P_A \mathbf{x})) \leq \cot((1/2)\varphi). \tag{21}$$

Case: $\lim_{t \downarrow 0} f'(t) = \infty$.

Combining this assumption with the continuity of f' , and the fact that f'' is negative on $[0, \varepsilon]$,

$$(\exists \delta_f \in [0, \varepsilon]) \text{ such that } (t \leq \delta_f) \implies f'(t) \geq 1. \tag{22}$$

Letting $x \in]0, \delta_f]$, we have from Lemma 4 that $y \in [0, \delta_f]$, and so $f'(y) \geq 1$. Consequently, we have that

$$\begin{aligned} \arctan(f'(y)) &= \angle(\mathbf{x}, \mathbf{x}_+, P_A \mathbf{x}) \geq \pi/4 \geq (1/2)(\pi/2) = (1/2)\varphi, \\ \text{and so } \cot(\angle(\mathbf{x}, \mathbf{x}_+, P_A \mathbf{x})) &\leq \cot((1/2)\varphi), \end{aligned}$$

which shows Eq. 21.

Case: $\lim_{t \downarrow 0} f'(t) < \infty$.

Since $f'(0) < \infty$ and $\varphi := \arctan(f'(0))$, there exists $\delta_f \in [0, \varepsilon]$ so that

$$(x \leq \delta_f) \implies \arctan(f'(t)) \geq (1/2) \arctan(f'(0)) = (1/2)\varphi,$$

whereupon $\arctan(f'(y)) \geq (1/2)\varphi$. Consequently,

$$\begin{aligned} \arctan(f'(y)) &= \angle(\mathbf{x}, \mathbf{x}_+, P_A \mathbf{x}) \geq (1/2)\varphi, \\ \text{and so } \cot(\angle(\mathbf{x}, \mathbf{x}_+, P_A \mathbf{x})) &\leq \cot((1/2)\varphi), \end{aligned}$$

which again shows Eq. 21. This concludes our need for separate cases.

Next, notice that

$$\|\mathbf{x}_+\| \leq \|\mathbf{x}_+ - \tau\| \tag{23a}$$

$$\text{and } \|\mathbf{x}\| \geq \|\mathbf{x} - \tau\|. \tag{23b}$$

To see why, recall that $\mathbf{x}_+ = (g_f(x), 0)$ with $g_f(x) \leq 0$ (Lemma 4(iv)). By definition, either $\tau = (\hat{\tau}, 0)$ satisfies $\hat{\tau} = y \in]0, x[$ or $\hat{\tau} = y - f(y)/f'(0)$. In the case when $\hat{\tau} = y \in]0, x[$, we have

$$\begin{aligned} \|\mathbf{x}_+\| &= |g_f(x)| < |y| + |g_f(x)| = |y - g_f(x)| = \|\mathbf{x}_+ - \tau\|, & \text{which shows Eq. 23a} \\ \text{and } \|\mathbf{x}\| &= |x| > |x - y| = \|\mathbf{x} - \tau\|, & \text{which shows Eq. 23b.} \end{aligned}$$

In the other case when $\hat{\tau} = y - f(y)/f'(0)$, we know that $\infty > f'(0) \geq f'(y) > 0$ (this is just from the definition of τ). We have from the mean value theorem that some $c \in [0, y]$

satisfies $f(y) - f(0) = f'(c)(y - 0)$, which is just $y = f(y)/f'(c)$. By monotonicity, $f'(0) \geq f'(c)$. Altogether,

$$\hat{\tau} = y - \frac{f(y)}{f'(0)} \geq y - \frac{f(y)}{f'(c)} = y - y = 0.$$

Combined with the fact that $\mathbf{x}_+ = (g_f(x), 0)$ with $g_f(x) \leq 0$, we again obtain Eq. 23a. Noticing further that

$$\hat{\tau} = y - \frac{f(y)}{f'(0)} \leq y < x,$$

we have that $\hat{\tau} \in [0, x]$, which shows Eq. 23b.

Now we have the following.

$$\frac{\|\mathbf{x}_+\|}{\|\mathbf{x}\|} \leq \frac{\|\mathbf{x}_+ - \tau\|}{\|\mathbf{x} - \tau\|} \tag{24a}$$

$$= \frac{\|\mathbf{x}_+ - P_A\mathbf{x}\|}{\|\mathbf{x} - P_A\mathbf{x}\|} \cdot \frac{\sin(\angle(\mathbf{x}_+, P_A\mathbf{x}, \tau))}{\sin(\angle(\mathbf{x}, P_A\mathbf{x}, \tau))} \tag{24b}$$

$$= \cot(\angle(\mathbf{x}, \mathbf{x}_+, P_A\mathbf{x})) \cdot \frac{\sin(\angle(\mathbf{x}_+, P_A\mathbf{x}, \tau))}{\sin(\angle(\mathbf{x}, P_A\mathbf{x}, \tau))} \tag{24c}$$

$$\leq \cot((1/2)\varphi) \cdot \frac{\sin(\angle(\mathbf{x}_+, P_A\mathbf{x}, \tau))}{\sin(\angle(\mathbf{x}, P_A\mathbf{x}, \tau))}. \tag{24d}$$

Here Eq. 24a follows from Eq. 23. We have the identity Eq. 24b from the generalized angle bisector theorem, and Eq. 24c is mere trigonometry. Finally, Eq. 24d is from Eq. 21. Now notice that

$$\angle(\mathbf{x}_+, P_A\mathbf{x}, \tau) = \pi - \angle(\mathbf{x}_+, \tau, P_A\mathbf{x}) - \angle(P_A\mathbf{x}, \mathbf{x}_+, \tau) \tag{25a}$$

$$= \pi - (\pi - \angle(P_A\mathbf{x}, \tau, (y, 0))) - \angle(P_A\mathbf{x}, \mathbf{x}_+, \tau) \tag{25b}$$

$$= \angle(P_A\mathbf{x}, \tau, (y, 0)) - \angle(P_A\mathbf{x}, \mathbf{x}_+, \tau) \tag{25c}$$

$$= \arctan(f'(0)) - \arctan(f'(y)) \tag{25d}$$

$$= \downarrow 0 \text{ as } x \downarrow 0. \tag{25d}$$

Here Eq. 25a is simply the fact that the sum of the angles of the triangle $(\mathbf{x}_+, P_A\mathbf{x}, \tau)$ is π , Eq. 25b uses the fact that $\angle(P_A\mathbf{x}, \tau, (y, 0))$ is complementary to $\angle(\mathbf{x}_+, \tau, P_A\mathbf{x})$, and Eq. 25c simply uses the fact that $\tan(\angle(P_A\mathbf{x}, \tau, (y, 0))) = \|P_A\mathbf{x} - (y, 0)\|/\|(y, 0) - \tau\| = f'(0)$ while $\tan(\angle(P_A\mathbf{x}, \mathbf{x}_+, \tau)) = \tan(\angle(P_A\mathbf{x}, \mathbf{x}_+, (y, 0))) = \|P_A\mathbf{x} - (y, 0)\|/\|(y, 0) - \mathbf{x}_+\| = f'(y)$. Finally, Eq. 25d uses continuity and the fact that $0 \leq y < x$. Additionally, we have that

$$\angle(\mathbf{x}, P_A\mathbf{x}, \tau) = \pi/2 - \angle(\mathbf{x}_+, P_A\mathbf{x}, \tau) \tag{26a}$$

$$\uparrow \pi/2 \text{ as } x \downarrow 0, \tag{26b}$$

where Eq. 26a uses the fact that the sum of the two angles $\angle(\mathbf{x}, P_A\mathbf{x}, \tau)$ and $\angle(\mathbf{x}_+, P_A\mathbf{x}, \tau)$ is $\angle(\mathbf{x}_+, P_A\mathbf{x}, \mathbf{x}) = \pi/2$, and Eq. 26b uses Eq. 25. Combining Eqs. 25 and 26 and the continuity of the sine function, we have that

$$\frac{\sin(\angle(\mathbf{x}_+, P_A\mathbf{x}, \tau))}{\sin(\angle(\mathbf{x}, P_A\mathbf{x}, \tau))} \downarrow 0 \text{ as } x \downarrow 0. \tag{27}$$

As a consequence of Eq. 27, for any $K > 0$, we may choose $\varepsilon_{K_f} > 0$ small enough to ensure that

$$(x \leq \varepsilon_{K_f}) \implies \frac{\sin(\angle(\mathbf{x}_+, P_A\mathbf{x}, \tau))}{\sin(\angle(\mathbf{x}, P_A\mathbf{x}, \tau))} \leq \frac{1}{\cot((1/2)\varphi) K},$$

which combines with Eq. 24d to ensure that

$$(x \leq \varepsilon_{K_f}) \implies \frac{\|\mathbf{x}_+\|}{\|\mathbf{x}\|} \leq \frac{1}{K},$$

This shows the desired result. □

Lemma 5 essentially shows local superlinear convergence rate for a setting in which Newton–Raphson method is not guaranteed to converge at all, and for settings where it has provably failed. One such example is $t \mapsto t/\sqrt{|t|}$ (where $t \in \mathbb{R}$). The behaviour of CRM for this problem is shown in Fig. 3.

Note that the basic conditions (I)–(V) in Section 5.2 allowed us to prove Lemmas 4 and 5 under the assumption of condition (VI), requiring that the graph of f sits (at least locally) in the first quadrant. Of course, there is nothing special about this quadrant. We now formalize our more general results in Theorem 3. Owing to the sheer number of cases that can be reduced to our framework, Theorem 3 is wide but not all-encompassing.

Theorem 3 *Let $\theta : \mathbb{R} \rightarrow \mathbb{R}$ be defined such that the functions $f = |\theta|$ and $f = |\theta \circ (-Id)|$ satisfy the basic conditions (I)–(V). Then the following hold.*

- (I) *Let θ be such that the functions $f = |\theta|$ and $f = |\theta \circ (-Id)|$ satisfy (i)[†]. Let $K' > 1$ and*

$$\varepsilon = \min \left\{ \varepsilon_{|\theta|}, \varepsilon_{|\theta \circ (-Id)|}, \varepsilon_{K'_{|\theta|}}, \varepsilon_{K'_{|\theta \circ (-Id)|}} \right\}, \tag{28}$$

where $\varepsilon_{|\theta|}, \varepsilon_{|\theta \circ (-Id)|}$ are as defined in the basic conditions for the functions $|\theta|$ and $|\theta \circ (-Id)|$ respectively, and $\varepsilon_{K'_{|\theta|}}, \varepsilon_{K'_{|\theta \circ (-Id)|}}$ are as defined in Lemma 5. Let $\mathbf{x}_0 = (x_0, 0)$ with $x_0 \in [-\varepsilon, \varepsilon]$. Then $\mathbf{x}_n \rightarrow \mathbf{0}$ as $n \rightarrow \infty$ with convergence rate that is never worse than $1/K'$ and ultimately either finite or quadratic.

- (II) *Let θ be such that the functions $f = |\theta|$ and $f = |\theta \circ (-Id)|$ satisfy (ii)[†]. Let $K' > 1$ and ϵ be as in Eq. 28. Let $\mathbf{x}_0 = (x_0, 0)$ with $x_0 \in [-\varepsilon, \varepsilon]$. Then $\mathbf{x}_n \rightarrow \mathbf{0}$ as $n \rightarrow \infty$ with convergence rate that is never worse than $1/K'$ and ultimately either finite or superlinear.*
- (III) *If $f = |\theta|$ satisfies (iii)[†], and $\mathbf{x}_0 = (x_0, 0)$ with $x_0 \in [0, \min\{\varepsilon_{|\theta|}, \delta_{RIGHT}\}]$ where δ_{RIGHT} is as in Lemma 2, then $\mathbf{x}_n \rightarrow \mathbf{0}$ as $n \rightarrow \infty$ with convergence rate that is ultimately either finite or quadratic.*
- (IV) *Let $f = |\theta|$ satisfy (iv)[†] and $\mathbf{x}_0 = (x_0, 0)$ with $x_0 \in [0, \min\{\varepsilon_{|\theta|}, \delta_{LIN}\}]$ where δ_{LIN} is as defined in Lemma 3 for the function h specified in (iv)[†] with $f = |\theta|$. Then $\mathbf{x}_n \rightarrow \mathbf{0}$ as $n \rightarrow \infty$ with convergence rate that is ultimately either finite or linear.*

Proof (II): Suppose the convergence is not finite. Combining our assumptions with Lemma 5,

$$\frac{\|\mathbf{x}_{n+1}\|}{\|\mathbf{x}_n\|} \leq \frac{1}{K'},$$

which shows that $x_n \rightarrow 0$ with convergence rate that is no worse than $1/K'$. From Lemma 5 we also have that, for any $K > 0$, there exist constants $\varepsilon_{K_{|\theta|}}$ and $\varepsilon_{K_{|\theta \circ (-Id)|}}$ such that

$$(x_n \in [-\varepsilon_{K_{|\theta|}}, \varepsilon_{K_{|\theta|}}] \cap [-\varepsilon_{K_{|\theta \circ (-Id)|}}, \varepsilon_{K_{|\theta \circ (-Id)|}}]) \implies \frac{\|\mathbf{x}_{n+1}\|}{\|\mathbf{x}_n\|} \leq \frac{1}{K}. \tag{29}$$

Since $x_n \rightarrow 0$, we have that there exists N_K sufficiently large that

$$(n \geq N_K) \implies x_n \in [-\varepsilon_{K_{|\theta|}}, \varepsilon_{K_{|\theta|}}] \cap [-\varepsilon_{K_{|\theta \circ (-Id)|}}, \varepsilon_{K_{|\theta \circ (-Id)|}}] \tag{30}$$

Combining Eqs. 29 and 30, we have that

$$(n \geq N_K) \implies \frac{\|\mathbf{x}_{n+1}\|}{\|\mathbf{x}_n\|} \leq \frac{1}{K},$$

which shows the claimed superlinear convergence.

(I): By the same arguments we just used for (II), we have that $x_n \rightarrow 0$. Suppose the convergence is not finite. For n sufficiently large, $x_n \in [-\delta_{\text{MIN}}, \delta_{\text{MIN}}]$ where δ_{MIN} is as defined in Lemma 2, and so we may combine with Lemma 4(v) to obtain

$$\frac{\|\mathbf{x}_{n+1}\|}{\|\mathbf{x}_n\|^2} \leq \frac{|g_\theta(x_n)|}{|x_n|^2} < \frac{M_{\text{MAX}}}{2},$$

where M_{MAX} is as given in Lemma 2, which shows the eventual quadratic convergence rate.

(III): Suppose the convergence is not finite. Because $x_n \in [0, \min\{\epsilon_{|\theta|}, \delta_{\text{RIGHT}}\}]$, applying Lemma 2 together with Lemma 4(v) yields

$$\frac{\|\mathbf{x}_{n+1}\|}{\|\mathbf{x}_n\|^2} = \frac{|x_{n+1}|}{|x_n|^2} \leq \frac{|g_\theta(x_n)|}{|x_n|^2} < \frac{M_{\text{RIGHT}}}{2},$$

which shows the quadratic convergence.

(IV): Suppose the convergence is not finite. Using the fact that $f = |\theta|$ satisfies $(iv)^\dagger$, we apply Lemma 3 for the corresponding function h and combine Lemma 4(v) with the equality of f and h on $[0, \min\{\epsilon_{|\theta|}, \delta_{\text{LIN}}\}]$ to obtain

$$\frac{\|\mathbf{x}_{n+1}\|}{\|\mathbf{x}_n\|} = \frac{|x_{n+1}|}{|x_n|} \leq \frac{|g_\theta(x_n)|}{|x_n|} = \frac{|gh(x_n)|}{|x_n|} < k,$$

where $k = \frac{1}{2}(\frac{m-1}{m} + 1) < 1$. This shows the linear convergence. □

5.4 Plane Curve Examples

Local convergence of CRM to a feasible point for a line B together with many plane curves A may be handled by piecewise appeal to Theorem 3. We will mention a few examples that highlight the importance of these results.

1. We have convergence whenever A is an algebraic curve and B is a line. This includes the classical problems of ellipses and p-spheres [19]. We also have the following results about the rate of convergence.
 - (a) When A and B meet with multiplicity greater than one, as when B is the horizontal axis and A is the graph of $t \mapsto t^2$ (where $t \in \mathbb{R}$), we have linear convergence to the feasible point. This is in contrast with the setting of DR , where convergence is usually observed to be to a fixed point rather than a feasible point. Local convergence results about DR with plane curves have typically excluded such cases.
 - (b) When A and B meet with multiplicity one, as when B is the horizontal axis and A is the graph of $t \mapsto t^2 - 1$ (where $t \in \mathbb{R}$), we have superlinear convergence. When, additionally, the curves are not perpendicular at their point of meeting, the convergence is quadratic.
2. We have local superlinear convergence for the case where A is the graph of $t \mapsto t/\sqrt{|t|}$ (where $t \in \mathbb{R} \setminus \{0\}$) and B is the horizontal axis, a case where the Newton–Raphson method cycles. This particular example is shown in Fig. 3. Such cases highlight the importance of the projection onto the curve for preventing instability.

3. As an example where the generically proper characterization of C_T is useful, let $B = \mathbb{R} \times \{0\}$ and $A = \text{graf}$ with

$$f : t \mapsto \begin{cases} -t & \text{if } t \leq 1; \\ -1 & \text{otherwise.} \end{cases}$$

Then, for any x_n such that $P_A \mathbf{x}_n \in \{(t, -1) \mid t \geq 1\}$ has the same first coordinate as \mathbf{x}_n , one has a colinear case, whereupon $\mathbf{x}_{n+1} = T_{A,B} \mathbf{x}_n = \mathbf{x}_n + (0, 1)$. Consequently, for n sufficiently large, $P_A \mathbf{x}_n \in \{(t, -t) \mid t \in \mathbb{R}\}$, whereupon $\mathbf{x}_{n+1} = \text{CRM}(\mathbf{x}_n) = (0, 0) \in A \cap B$.

Succinctly put, for all cases when Douglas–Rachford exhibits local convergence to a *feasible* point, CRM provides a better convergence rate. Finally, as a root finder, CRM has local convergence in all cases when Newton–Raphson does, exhibits quadratic convergence in all cases when Newton–Raphson does, and exhibits superlinear convergence in many cases when Newton–Raphson fails to converge at all. Figure 5 shows the performance of the three methods, as measured by distance to solution, when B is the horizontal axis and A is the unit sphere centered at $(0, -1/2)$; the starting point used is $(0.9999, 0)$. These results reveal a local robustness of CRM that may be prototypical of convergence for more general feasibility problems for more complicated problems of interest, such as phase retrieval. The method L_T from [38] is also shown; as its performance appears similar to CRM, it is worth remembering that the cost of computing an update of L_T is double the cost of computing an update of CRM, because it requires the computation of 4 projections instead of 2.

6 Rate Guarantees and Numerical Discoveries in \mathbb{R}^n

While we were able to furnish Theorem 2 in \mathbb{R}^n , the fast rate results of Section 5 are (initially) limited to \mathbb{R}^2 . This restriction was somewhat unavoidable, because the framework we built was designed to be compatible with tying rate results explicitly to those of Newton–Raphson method. Interestingly, it is still possible to lift some of these results back into \mathbb{R}^n for certain problems. Of course, given that superlinear and quadratic rate guarantees are *very*




Fig. 5 Finding a point in the intersection of the sphere and line

strong, and also that they have been experimentally observed not to hold for many problems, one would expect that meaningful structural assumptions would be needed in order to guarantee a quadratic convergence result.

Owing to complexity, a “most general extension” in \mathbb{R}^η is not practicable. Indeed, even in \mathbb{R}^2 , we reduced the statement of Theorem 3 to just four cases (I)–(IV)) for simplicity, but one could afford far greater generality by stating it in 16 cases (simply mix and match the properties of $|\theta|$ and $|\theta \circ (-\text{Id})|$ to obtain many more guarantees).

What we will do in this section, then, is demonstrate one approach to adapting the theory from Section 5 for a selection of feasibility problems that have received significant attention in the literature: (1) spheres and hyperplanes, and (2) spheres and subspaces. Accordingly, we build one particular extension of our theory in \mathbb{R}^η (Theorem 4) that is designed with the specific purpose of working for the examples we care about, and use this specific extension to provide rates for these problems.

In the process, we stumble upon a much more significant discovery. Our numerical experiment for Example 2 reveals an extreme sensitivity to numerical error that may cost CRM the superior convergence rates that are guaranteed by the theory. We correct for the numerical error, and recover the superior rate guaranteed by the theory. As the sensitivity is especially relevant to problem architectures that involve subspaces, all future works that use numerical experiment to study CRM with the product space formulation of feasibility problems (e.g. [8, 16, 25, 26]) should take note of this sensitivity and account for it appropriately.

The particular extension theorem we introduce will be easier to understand if we first present one of the examples that motivates it. Spheres and subspaces are of interest because they are prototypical of phase retrieval. They were studied for the Douglas–Rachford method in [2, 18, 20], and the Lyapunov function discovered in [18] has been catalytic in other nonconvex investigations [24, 33]. Interestingly, global convergence for CRM for spheres and hyperplanes is implicitly shown in [16, see Remark 1]. We can furnish the first quadratic rate guarantees for these problems. For each of our computed examples, we will include the method L_T from [38]; for comparison, it is useful to remember that computing a step of L_T has double the cost (4 projections instead of 2) of computing a step of CRM.

Example 1 (Sphere and hyperplane) By symmetry, the (consistent) feasibility problem for a unit sphere A and hyperplane not containing zero, B , in $\mathbb{R}^{\eta+1}$ may always be rotated so that $B := \mathbb{R}^\eta \times \{b\}$ where $b \in]0, 1]$. For \mathbf{x}_0 started in the axis of symmetry $\{0\}^\eta \times \mathbb{R}$, the Douglas–Rachford method is known to fail to converge, producing colinear reflections for all n (see [11]).

For \mathbf{x}_0 started in B and also outside of this chaotic set, it suffices by symmetry to consider $\mathbf{x}_0 = (0, \dots, 0, t, b)$ for some $t \in]0, \infty[$. Suppose $\mathbf{x}_n = (0, \dots, 0, t, b)$ for some $t \in]0, \infty[$. Because the projection of \mathbf{x}_n onto A is given by $\mathbf{x}_n / \|\mathbf{x}_n\|$, its reflection across the sphere is

$$\begin{aligned}
 R_A \mathbf{x}_n &= 2\mathbf{x}_n / \|\mathbf{x}_n\| - \mathbf{x}_n = (0, \dots, 0, (2/\|\mathbf{x}_n\| - 1)t, (2/\|\mathbf{x}_n\| - 1)b) \\
 \text{and so } R_B R_A \mathbf{x}_n &= (0, \dots, 0, (2/\|\mathbf{x}_n\| - 1)t, 2b - (2/\|\mathbf{x}_n\| - 1)b) \\
 \text{and so } \text{CRM}(\mathbf{x}_n) &\in \{0\}^{\eta-1} \times \mathbb{R} \times \{b\} =: D \quad (\text{by Lemma 1(iv)}).
 \end{aligned}$$

Consequently, for all n , $\mathbf{x}_n \in \{0\}^{\eta-1} \times \mathbb{R} \times \{b\}$, while $P_A \mathbf{x}_n \in A \cap (\{0\}^{\eta-1} \times \mathbb{R} \times \mathbb{R})$. In other words, the computation \mathbf{x}_n , including all of the associated projections and reflections, takes place entirely within a subspace of dimension 2, and this subspace is the same for all n . The projections $P_A \mathbf{x}_n$ are projections onto the graph $\{(x, \Theta(x)) \in \mathbb{R}^{\eta+1} \mid x \in \mathbb{R}^\eta\}$ of the function $\Theta : \mathbb{R}^\eta \rightarrow \mathbb{R} : x \mapsto \sqrt{1 - \|x\|_{\mathbb{R}^\eta}^2}$. More specifically, though, they are projections

onto the slice of this graph that is given by $\{(x, \Theta(x)) \in \mathbb{R}^\eta \times \mathbb{R} \mid x \in \{0\} \times \dots \times \{0\} \times \mathbb{R} \times \{b\}\}$. This slice is the embedded lower-dimensional graph $\{\mathbf{v} + \theta(\mathbf{v}) \cdot (0, \dots, 0, 1) \in \mathbb{R}^{\eta+1} \mid \mathbf{v} \in D\}$ of the function $\theta : D \rightarrow \mathbb{R} : \mathbf{v} \rightarrow \sqrt{1 - |\mathbf{v}_\eta|^2} - b$. Moreover for all n , the point $\mathbf{s} := (0, \dots, 0, \sqrt{1 - b^2}, b)$ belongs to $A \cap B \cap \text{aff}\{\mathbf{x}_n, R_A \mathbf{x}_n, R_B R_A \mathbf{x}_n\}$. With these conditions, Theorem 4, which we will introduce momentarily, guarantees that \mathbf{x}_n converges to \mathbf{s} with a quadratic rate (where the set $\{\mathbf{s}_n\}_n$ in the theorem is exactly equal to the singleton $\{\mathbf{s}\}$).

Computed example: In Fig. 6 we show performance for a 10-sphere centered at $(0, \dots, 0, -1/2)$. In this example, we include the method of subgradient projections for comparison, where we treat the first 9 variables as the domain space and the function for which we apply subgradient projections is $f : \mathbf{x} \mapsto \sqrt{1 - (\mathbf{x}_1^2 + \dots + \mathbf{x}_9^2)} - 1/2$. Of course, this algorithm is only defined on $\mathbb{B} \setminus \{\mathbf{0}\}$ where \mathbb{B} is the unit ball in \mathbb{R}^9 . Moreover, for any dampened version of subgradient projections, if we start near enough to $\mathbf{0}$, our first update will be outside of this domain. Consequently, in order for subgradient projections to converge for this problem, we must start very close to the solution. For our experiments, we actually used CRM to obtain a starting point very near to the feasible set, and started all algorithms thereat.

We now introduce our particular extension theorem, whose many conditions can all be verified for the sphere and subspace feasibility problems in Examples 1 and 2. As we have already explained, many other generalizations in higher dimensions are possible, but Example 1 may be read as a template for motivating the specific conditions we chose to work with in this particular extension.

Theorem 4 *Let A be a subset, and B a proper subspace, of \mathbb{R}^η . Let $\mathbf{x}_0 \in B$ and $\mathbf{x}_n := \text{CRM}^n(\mathbf{x}_0)$. Suppose further that the following hold:*

- (i) *For all n , $\mathbf{x}_n \in B$;*
- (ii) *$\mathbf{x}_n \rightarrow \mathbf{y} \in B \cap A$ with $\xi > 0$ satisfying $\mathbf{x}_n \in \mathbb{B}_\xi(\mathbf{y})$ for all n ;*




Fig. 6 Finding a point in the intersection of the 10-sphere and hyperplane from Example 1

- (iii) There exists $\mathbf{u} \in B^\perp \cap \mathbb{B}_1(\mathbf{0})$ such that, for all n , $P_A(\mathbf{x}_n) = P_{\text{gra}\Theta}(\mathbf{x}_n)$ where $\Theta : B \rightarrow \mathbb{R}$ and $\text{gra}\Theta := \{\mathbf{x} + \Theta(\mathbf{x}) \cdot \mathbf{u} \mid \mathbf{x} \in B\}$;
- (iv) For all n , the sets

$$\{r(\mathbf{x}_n - \mathbf{x}_{n+1}) \mid r \in \mathbb{R}\} \cap A \cap L \cap \mathbb{B}_\xi(\mathbf{x})$$

are singletons, which we name \mathbf{s}_{n+1} respectively;

- (v) For the functions

$$\theta_n : r \mapsto \Theta \left(r \frac{(\mathbf{x}_n - \mathbf{s}_n)}{\|\mathbf{x}_n - \mathbf{s}_n\|} + \mathbf{s}_n \right),$$

(which, as we have defined them, possess a root at 0), it holds that $|\theta_n|$ and $|\theta_n \circ \text{Id}|$ satisfy the basic conditions (I)–(V).

- (vi) The functions $f_n = |\theta_n|$ satisfy (iii)[†] with corresponding $\epsilon_{f_n} > 0$ and $\delta_{\text{RIGHT}}^{f_n} > 0$ (where $\delta_{\text{RIGHT}}^{f_n}$ is as in Lemma 2), as well as $\min\{\epsilon_{f_n}, \delta_{\text{RIGHT}}^{f_n}\} \geq \|\mathbf{x}_n - \mathbf{s}_n\| > 0$ for all $n \geq N$ for some $N > 0$.

Then for all $n > N$, it holds that \mathbf{x}_n satisfies

$$\frac{\|\mathbf{x}_{n+1} - \mathbf{s}_n\|}{\|\mathbf{x}_n - \mathbf{s}_n\|^2} \leq \frac{M_{\text{RIGHT}}}{2}.$$

Proof Suppose convergence is not finite. Let $n > N$. By a suitable translation, we can and do without loss of generality let $\mathbf{s}_n = \mathbf{0}$, which simplifies the notation and assures that $\|\mathbf{x}_n\| \leq \min\{\epsilon_{f_n}, \delta_{\text{RIGHT}}^{f_n}\}$. Define

$$\text{gra}\theta_n := \{\mathbf{x} + \Theta(\mathbf{x}) \cdot \mathbf{u} \mid \mathbf{x} \in \text{aff}\{\mathbf{x}_n, \mathbf{x}_{n+1}\}\}.$$

Suppose for a contradiction that $P_A(\mathbf{x}_n) \notin \text{gra}\theta_n$. Then, combining with the fact that $P_A(\mathbf{x}_n) \in \text{gra}\Theta$, there exists $\mathbf{y} \in B \setminus \text{aff}\{\mathbf{x}_n, \mathbf{x}_{n+1}\}$ with $P_A(\mathbf{x}_n) = \mathbf{y} + \Theta(\mathbf{y}) \cdot \mathbf{u}$. Consequently,

$$\begin{aligned} R_A(\mathbf{x}_n) &= 2\mathbf{y} + 2\Theta(\mathbf{y}) \cdot \mathbf{u} - \mathbf{x} \\ \text{and so } R_B R_A(\mathbf{x}_n) &= 2\mathbf{y} - 2\Theta(\mathbf{y}) \cdot \mathbf{u} - \mathbf{x} \\ \text{whereupon } \text{aff}\{\mathbf{x}_n, R_A \mathbf{x}_n, R_B R_A \mathbf{x}_n\} \cap B &= \text{aff}\{2\mathbf{y} - \mathbf{x}_n, \mathbf{x}_n\} \\ &= \text{aff}\{\mathbf{y}, \mathbf{x}_n\} \not\supseteq \mathbf{x}_{n+1}, \end{aligned} \tag{31}$$

where the final $\not\supseteq$ must hold because we assumed $\mathbf{y} \in B \setminus \text{aff}\{\mathbf{x}_n, \mathbf{x}_{n+1}\}$. However, in view of the definition of CRM, it must hold that $\mathbf{x}_{n+1} \in \text{aff}\{\mathbf{x}_n, R_A \mathbf{x}_n, R_B R_A \mathbf{x}_n\}$, while in view of Lemma 1(iv), it must also hold that $\mathbf{x}_{n+1} \in B$. Consequently,

$$\mathbf{x}_{n+1} \in \text{aff}\{\mathbf{x}_n, R_A \mathbf{x}_n, R_B R_A \mathbf{x}_n\} \cap B,$$

and this contradicts Eq. 31. Thus we have shown that $P_A(\mathbf{x}_n) \in \text{gra}\theta_n$. This fact combines with the fact that $P_A(\mathbf{x}_n) = P_{\text{gra}\Theta}(\mathbf{x}_n)$ and the fact that $\text{gra}\theta_n \subset \text{gra}\Theta$, to guarantee that

$$P_A(\mathbf{x}_n) = P_{\text{gra}\theta_n}(\mathbf{x}_n).$$

Therefore the entire CRM construction at step $n + 1$ lives entirely in the 2-dimensional subspace $\text{aff}\{\mathbf{x}_n, \mathbf{x}_{n+1}\} \oplus \text{span}\{\mathbf{u}\}$. As we have $\mathbf{0} = \mathbf{s}_n \in \text{aff}\{\mathbf{x}_n, \mathbf{x}_{n+1}\}$, we may use the change of coordinates:

$$B \oplus \text{span}\{\mathbf{u}\} \rightarrow \mathbb{R} \times \mathbb{R} : (\mathbf{x}, \mathbf{x}) \mapsto (\langle \mathbf{x}, \mathbf{x}_n \rangle / \|\mathbf{x}_n\|, \langle \mathbf{x}, \mathbf{u} \rangle),$$

whereupon $\mathbf{x}_n = (x_n, 0)$ with $x_n = \|\mathbf{x}_n\| \leq \min\{\epsilon_{f_n}, \delta_{\text{RIGHT}}^{f_n}\} \leq \epsilon_{f_n}$. Thus we have all the necessary conditions to apply Lemma 4(v), whereupon $\|\mathbf{x}_{n+1}\| \leq g_{\theta_n}(\|\mathbf{x}_n\|)$. Combining

with the fact that $\|\mathbf{x}_n\| \leq \min\{\epsilon_{f_n}, \delta_{\text{RIGHT}}^{f_n}\}$, Lemma 2 yields

$$\frac{\|\mathbf{x}_{n+1}\|}{\|\mathbf{x}_n\|^2} \leq \frac{g_{\theta_n}(\|\mathbf{x}_n\|)}{\|\mathbf{x}_n\|^2} < \frac{M_{\text{RIGHT}}}{2}.$$

□

The more general subspace case in Example 2 will be a straightforward adaptation of the hyperplane case in Example 1. However, the phenomenon we observed in our numerical experiments is quite distinct from the hyperplane case. It demonstrates an extreme sensitivity of convergence rate for methods like CRM to small compounding numerical error that may cause sequences to depart from subspaces in which they are *theoretically* guaranteed to remain. This small, numerically introduced, departure from the subspace results in a completely different convergence rate.

Example 2 (A sphere and an affine subspace) In $\mathbb{R}^{\eta+m}$, let A be the unit sphere and B be a subspace not containing zero. More specifically, let B have dimension η . Then by symmetry, the (consistent) feasibility problem may always be rotated so that $B := \mathbb{R}^\eta \times \{\mathbf{b}\}$ where $\mathbf{b} \in]0, 1]^m$ and $\|\mathbf{b}\|_{\mathbb{R}^m} \leq 1$. We will consider the subset of cases where $\|\mathbf{b}\|_{\mathbb{R}^m} < 1$. For \mathbf{x}_0 started in the axis of symmetry $\{0\}^\eta \times \mathbb{R}^m$, the Douglas–Rachford method is known to fail to converge, producing colinear reflections for all n (see [11]).

For \mathbf{x}_0 started in B and also outside of the axis of symmetry, it suffices by symmetry to consider $\mathbf{x}_0 = (0, \dots, 0, t, b_1, \dots, b_m) =: (0, \dots, 0, t, \mathbf{b})$ for some $t \in]0, \infty[$. Suppose $\mathbf{x}_n = (0, \dots, 0, t, \mathbf{b})$ for some $t \in]0, \infty[$. Because the projection of \mathbf{x}_n onto A is given by $\mathbf{x}_n / \|\mathbf{x}_n\|$, its reflection across the sphere is

$$\begin{aligned} R_A \mathbf{x}_n &= 2\mathbf{x}_n / \|\mathbf{x}_n\| - \mathbf{x}_n = (0, \dots, 0, (2/\|\mathbf{x}_n\| - 1)t, (2/\|\mathbf{x}_n\| - 1)\mathbf{b}) \\ \text{and so } R_B R_A \mathbf{x}_n &= (0, \dots, 0, (2/\|\mathbf{x}_n\| - 1)t, 2\mathbf{b} - (2/\|\mathbf{x}_n\| - 1)\mathbf{b}) \\ \text{and so } \text{CRM}(\mathbf{x}_n) &\in \{0\}^{\eta-1} \times \mathbb{R} \times \{\mathbf{b}\} =: D. \quad (\text{by Lemma 1(iv)}) \end{aligned}$$

Consequently, for all n , $\mathbf{x}_n \in \{0\}^{\eta-1} \times \mathbb{R} \times \{\mathbf{b}\}$, while $P_A \mathbf{x}_n \in A \cap (\{0\}^{\eta-1} \times \mathbb{R} \times \text{span}\{\mathbf{b}\})$. In other words, the computation \mathbf{x}_n , including all of the associated projections and reflections, takes place entirely within a subspace of dimension 2, and this subspace is the same for all n . The projections $P_A \mathbf{x}_n$ are projections onto the graph $\{(x, \Theta(x)) \in \mathbb{R}^{\eta+m} \mid x \in \mathbb{R}^{\eta+m-1}\}$ of the function $\Theta : \mathbb{R}^{\eta+m-1} \rightarrow \mathbb{R} : x \mapsto \sqrt{1 - \|x\|}$. More specifically, though, they are projections onto the slice of this graph that is given by $\{(x, \Theta(x)) \in \mathbb{R}^{\eta+m-1} \times \mathbb{R} \mid x \in \{0\}^{\eta-1} \times \mathbb{R} \times \{\mathbf{b}\}\}$. This slice is the embedded lower-dimensional graph $\{\mathbf{v} + \theta(\mathbf{v}) \cdot (0, \dots, 0, 1) \in \mathbb{R}^{\eta+m} \mid \mathbf{v} \in D\}$ of the function $\theta : D \rightarrow \mathbb{R} : \mathbf{v} \rightarrow \sqrt{1 - \|\mathbf{v}_\eta\|^2 - \|\mathbf{b}\|_{\mathbb{R}^m}}$. Moreover $\mathbf{s}_n = (0, \dots, 0, \sqrt{1 - \|\mathbf{b}\|_{\mathbb{R}^m}^2}, \mathbf{b}) =: \mathbf{s}$ for all n . Applying Theorem 4, we have that \mathbf{x}_n converges to \mathbf{s} with a quadratic rate.

Computed example: For our subspace B , we used a line. As there is no longer a single obvious way to define subgradient projections for this problem, we omit it. In our first numerical experiment, we begin with a $x_0 \notin B$, and we observe apparently only linear convergence in Fig. 7a. This is interesting, because it suggests that the condition $x_0 \in B$ cannot be relaxed without losing the quadratic rate. In our second experiment in Fig. 7b, we begin with $x_0 \in B$ and see early apparent quadratic convergence degrade into only linear convergence. However, our theoretical results guarantee a quadratic rate. Investigating further, we discover that compounding numerical error has ultimately caused \mathbf{x}_n to be outside of B , and this small difference has caused the loss of the superior convergence rate. For our final

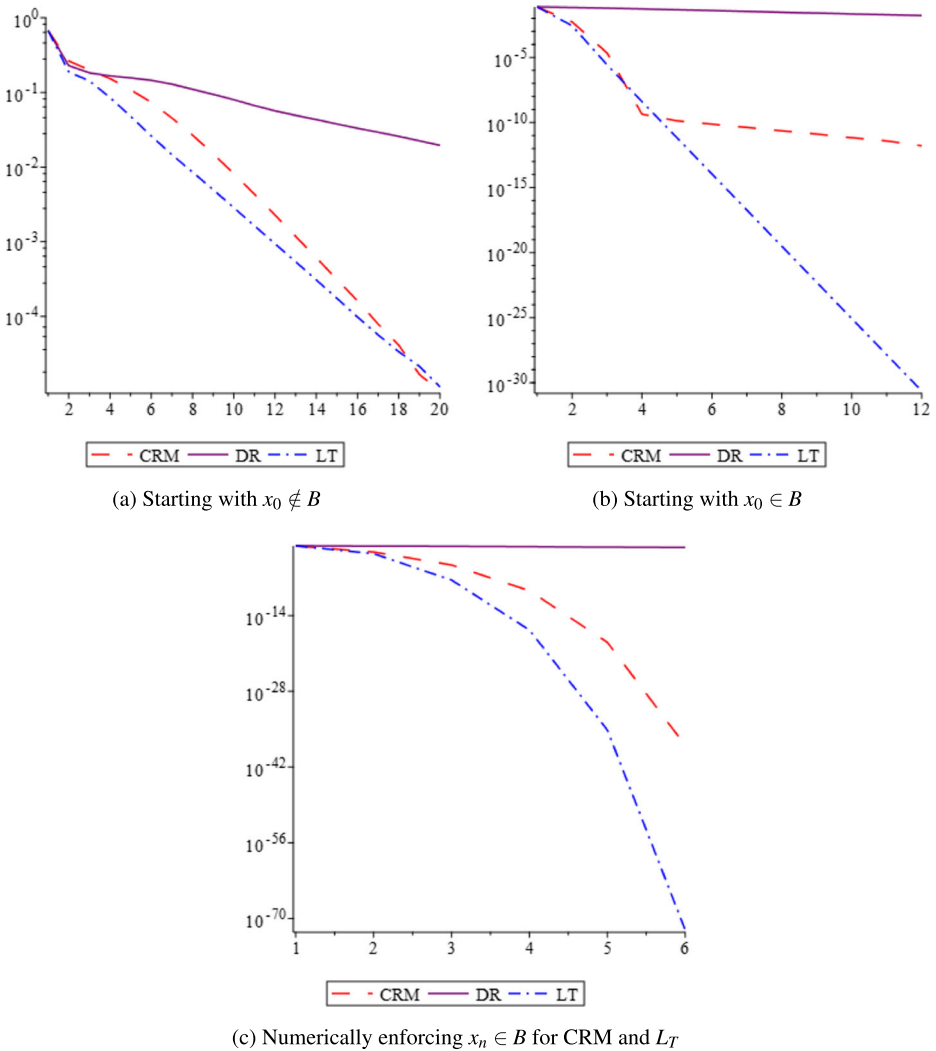


Fig. 7 Computational results from Example 2

experiment in Fig. 7c, we correct this compounding numerical error with an extra projection onto B ; in other words, we are iteratively applying $CRM \circ P_B$. Having corrected for the numerical error in this way, we see the quadratic convergence that the theory guarantees. This is quite interesting, because it illustrates the sensitivity of the convergence rate to very small numerical error. Purely for the sake of curiosity, in Fig. 7c we also replace L_T with $L_T \circ P_B$; in this case, the modification does not simply correct for numerical error, but actually defines a new algorithm. The apparent quadratic convergence of this algorithm is interesting and may merit future investigation.

For our next and final example, we choose an application that has *some* similarities to the context we have worked in (one set is a subspace and the other an implicitly defined surface),


but for which the rate is unknown. Naively, we are hoping to be pleasantly surprised by an improved rate, without necessarily expecting to be.

Example 3 (A boundary value ordinary differential equation) The task of numerically finding a discretized solution for a boundary value ordinary differential equation may be reformulated as a 2-set feasibility problem, where one of the sets is an implicitly defined surface and the other is a subspace [36, 37]. In [36], the Douglas–Rachford method and method of alternating projections are applied to such problems. We test CRM and L_T for the specific boundary value problem from [36, Example 6.6] (with a grid consisting of 20 mesh points for 21 segments; problem dimension 20^2). The details for how to compute projections are described in [36].

Computed example: We use the starting point that corresponds to a constant function that returns the value 0.5 on the open interval between the boundary points (see [36] for full details). A traditional solver—based on Newton’s method—finds a false solution from this starting point, but each of the methods DR, CRM, and L_T solved the problem in our experiment. Having learned the lesson about compounding numerical error from Example 2, we iterate $\text{CRM} \circ P_B$, where B is the *agreement subspace* and the inclusion of P_B serves to prevent compounding numerical error from pushing the sequence of updates out of B . We also replace L_T with $L_T \circ P_B$.

Partial solutions obtained by DR, CRM, and L_T (paused at 2,000 iterations), in addition to the true solution to the discretised problem and the false solution obtained by the traditional solver, are shown in Fig. 8a.

In spite of the aforementioned similarities with our framework, we observe only linear convergence in Fig. 8b. Here we record the norm distance between the shadow sequence of partial solutions for each method ($P_B x_n$) and the true solution to the discretised problem. For a more detailed explanation of why error is reported this way, see [36]. The



(a) Partial solutions corresponding to 2,000th iterate (b) Distance from discretized solution at n th iterate

Fig. 8 Computational results from Example 3

progress of DR towards the solution is best measured by tracking the “peaks” of the visible “tombstones” in its performance profile; for more information about this, see [36] or [38].

The fact that we observe only (apparently) linear convergence for CRM is unsurprising; it indicates that many conditions must all be met in order for CRM to have a superlinear convergence rate for a more general problem. When interpreting this graph, it is valuable to remember that one update of L_T is twice as costly (requires 4 projections to be computed instead of 2) to compute as one update of CRM. It is also worth noting that the case where $x_n, R_A x_n, R_B R_A x_n$ are distinct and colinear did not occur in our experiment.

7 Conclusion

The results in Sections 4 and 5 are natural analogs of those already in the literature for DR [2, 19, 20, 24, 40]. The framework, which relies on Theorem 2, is novel, and it illuminates a connection between CRM and subgradient projections in \mathbb{R}^n . In Section 6, the roadmap for adaptation in \mathbb{R}^n , and the convergence rate guarantees for spheres and hyperplanes, while interesting in their own regards, are both overshadowed in importance by the numerical revelations of Example 2. Now that this numerical sensitivity has been documented, all future investigations of CRM on feasibility problems built with the traditional agreement subspace architecture *must account for such numerical deviations*. If they do not do so, it is entirely possible that experiments will fail to reveal convergence rates that are theoretically achievable. The apparent quadratic convergence of the modified version of algorithm L_T also clearly merits further investigation.

Even as recently as while this paper was in peer review, [8] has used the epigraphical subgradient-descent characterization from Theorem 2 to furnish rates for convex feasibility problems that are related to those here. For the product space formulation of a feasibility problem, one of the constraint sets is a subspace B . Under some local assumptions, (Lemma 1), subsequences $(\mathbf{x}_j)_j$ that have $\mathbf{x}_0 \in B$ satisfy $(\mathbf{x}_j)_j \subset B$, whereupon the 2-dimensional affine subspace $\text{aff}(\mathbf{x}_j, R_A \mathbf{x}_j, R_B R_A \mathbf{x}_j)$ always contains a line $B \cap \text{aff}(\mathbf{x}_j, R_A \mathbf{x}_j, R_B R_A \mathbf{x}_j)$. The rate guarantee from [8], when the angle between sets “vanishes” is related to Theorem 3(IV), while the superlinear guarantees when the angle does not vanish are related to Theorem 3(I)–(III).

This is just one example of how these results are also of broader interest, because plane curve problems shed light on how such methods are thought to behave more generally. They motivated the introduction of 2 stage DR–CRM search algorithm in [25, 26]. Now that the groundwork has been laid for the prototypical settings, another natural next step of investigation is to conduct a detailed study on CRM for the phase retrieval problem specifically. Methods for accelerating convergence for Newton–Raphson are well known and are found in any numerical calculus textbook. The connections discussed in Section 5.2 indicate that another natural possible step is to attempt such methods with CRM.

Acknowledgements SBL was supported in part by Hong Kong Research Grants Council PolyU153085/16p and by an AustMS Lift-Off Fellowship; his collaboration in this project was also made possible in part by funding from CARMA Priority Research Centre at University of Newcastle. JAH and NDD are supported by Australian Research Council Grant DP160101537.

Funding Open Access funding enabled and organized by CAUL and its Member Institutions.

Data Availability The code used to generate the numerical results is available at [27].

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Cinderella (software): Available at <https://cinderella.de/tiki-index.php> (2016)
2. Artacho, F.J.A., Borwein, J.M.: Global convergence of a non-convex douglas–Rachford iteration. *J. Glob. Optim.* **57**(3), 753–769 (2013)
3. Artacho, F.J.A., Borwein, J.M., Tam, M.K.: Recent results on Douglas–Rachford methods. *Serdica Math. J.* **39**, 313–330 (2013)
4. Artacho, F.J.A., Borwein, J.M., Tam, M.K.: Douglas–Rachford feasibility methods for matrix completion problems. *ANZIAM J.* **55**(4), 299–326 (2014)
5. Artacho, F.J.A., Borwein, J.M., Tam, M.K.: Recent results on Douglas–Rachford methods for combinatorial optimization problems. *J. Optim Theory Appl.* **163**(1), 1–30 (2014)
6. Artacho, F.J.A., Campoy, R.: Solving graph coloring problems with the douglas–Rachford algorithm. *Set-Valued and Variational Analysis* **26**(2), 277–304 (2018)
7. Artacho, F.J.A., Campoy, R., Elser, V.: An enhanced formulation for solving graph coloring problems with the douglas–Rachford algorithm. *J. Glob. Optim.* **77**(2), 383–403 (2020)
8. Arefidamghani, R., Behling, R., Bello-Cruz, Y., Iusem, A.N., Santos, L.R.: The circumcentered-reflection method achieves better rates than alternating projections. *Comput. Optim. Appl.* **79**(2), 383–403 (2021)
9. Artacho, F.J.A., Campoy, R., Tam, M.K.: The douglas–Rachford algorithm for convex and nonconvex feasibility problems. *Math. Methods Oper. Res.* **91**(2), 201–240 (2020)
10. Bauschke, H.H., Combettes, P.L., Luke, D.R.: Phase retrieval, error reduction algorithm, and Fienup variants: a view from convex optimization. *J. Opt. Soc. Amer. A* **19**(7), 1334–1345 (2002)
11. Bauschke, H.H., Dao, M.N., Lindstrom, S.B.: The douglas–Rachford algorithm for a hyperplane and a doubleton. *J. Glob. Optim.* **74**(1), 79–93 (2019)
12. Bauschke, H.H., Ouyang, H., Wang, X.: On circumcenter mappings induced by nonexpansive operators, *Pure and Applied Functional Analysis*, in press (2018)
13. Bauschke, H.H., Ouyang, H., Wang, X.: On circumcenters of finite sets in Hilbert spaces. *Linear and Nonlinear Anal.* **4**, 271–295 (2018)
14. Bauschke, H.H., Ouyang, H., Wang, X.: On the linear convergence of circumcentered isometry methods. *Numer. Algo.* **87**(1), 263–297 (2021)
15. Behling, R., Bello-Cruz, J.Y., Santos, L.R.: On the linear convergence of the circumcentered-reflection method. *Oper. Res. Lett.* **46**(2), 159–162 (2018)
16. Behling, R., Bello-Cruz, J.Y., Santos, L.R.: On the circumcentered-reflection method for the convex feasibility problem. *Numer. Algo.* **86**(4), 1475–1494 (2021)
17. Behling, R., Bello-Cruz, J.Y., Santos, L.R.: Circumcentering the douglas–Rachford method. *Numer. Algo.* **78**, 759–776 (2018)
18. Benoist, J.: The douglas–Rachford algorithm for the case of the sphere and the line. *J. Glob. Optim.* **63**, 363–380 (2015)
19. Borwein, J.M., Lindstrom, S.B., Brailey, S., Skerritt, M.P., Schneider, A.: Dynamics of the Douglas–Rachford method for ellipses and p-spheres. *Set-Valued Anal.* **26**(2), 385–403 (2018)
20. Borwein, J.M., Brailey, S., The Douglas–Rachford algorithm in the absence of convexity. In: Heinz H. Bauschke, Regina S. Burachik, Patrick L. Combettes, Veit Elser, D. Russell Luke, and Henry Wolkowicz, editors, *Fixed Point Algorithms for Inverse Problems in Science and Engineering*, volume 49 of *Springer Optimization and Its Applications*, pages 93–109. Springer Optimization and Its Applications (2011)
21. Borwein, J.M., Tam, M.K.: Reflection Methods for Inverse Problems with Applications to Protein Conformation Determination. In: *Springer Volume on the CIMPA School Generalized Nash Equilibrium Problems, Bilevel Programming and MPEC*, New Delhi, India (2012)

22. Borwein, J.M., Tam, M.K.: Douglas–Rachford iterations in the absence of convexity: Jonathan Borwein’s lecture at 2015 spring school on variational analysis in Czech Republic. Available at [https://carma.newcastle.edu.au/\[DR\]methods/paseky/\[III\]-nonconvex{DRS}mall.pdf](https://carma.newcastle.edu.au/[DR]methods/paseky/[III]-nonconvex{DRS}mall.pdf) (2015)
23. Burden, R.L., Douglas Faires, J.: Numerical analysis. cengage learning (2016)
24. Dao, M.N., Tam, M.K.: A Lyapunov-type approach to convergence of the Douglas–Rachford algorithm. *J. Glob Optim.* **73**(1), 83–112 (2019)
25. Dizon, N., Hogan, J., Lindstrom, S.B.: Centering projection methods for wavelet feasibility problems. ISAAC 2019: The 30th International Symposium on Algorithms and Computation proceeding volume Current Trends in Analysis, its Applications and Computation, in press (2020)
26. Dizon, N., Hogan, J., Lindstrom, S.B.: Circumcentered reflections method for wavelet feasibility problems. *ANZIAM J.* **22** (2022)
27. Dizon, N., Hogan, J., Lindstrom, S.B.: Code for the article “Circumcentering reflection methods for nonconvex feasibility problems”. Available at <https://github.com/lindstromscott/Circumcentering-Reflection-Methods-for-Nonconvex-Feasibility-Problems> (2021)
28. Douglas Jr, J., Rachford Jr, H.H.: On the numerical solution of heat conduction problems in two and three space variables. *Trans. Amer. Math. Soc.* **82**, 421–439 (1956)
29. Elser, V.: The complexity of bit retrieval. *IEEE Trans. Inf. Theory* **64**(1), 412–428 (2018)
30. Fienup, J.R.: Phase retrieval algorithms: a comparison. *Appl. optics* **21**(15), 2758–2769 (1982)
31. Franklin, D.J.: Projection Algorithms for Non-separable Wavelets and Clifford Fourier Analysis. PhD thesis, University of Newcastle. Available at <http://hdl.handle.net/1959.13/1395028> (2018)
32. Gabay, D.: Applications of the Method of Multipliers to Variational Inequalities. In: *Studies in Mathematics and Its Applications*, Volume 15, Chapter IX, Pages 299–331. Elsevier (1983)
33. Giladi, O., Rüffer, B.S.: A Lyapunov function construction for a non-convex douglas–Rachford iteration. *J. Optim. Theory Appl.* **180**(3), 729–750 (2019)
34. Gravel, S., Elser, V.: Divide and concur: A general approach to constraint satisfaction. *Phy. Rev. E* **78**(3), 036706 (2008)
35. Lakshmikantham, V., Trigiante, D.: *Theory of difference equations - numerical methods and applications* marcel dekker (2002)
36. Lamichhane, B.P., Lindstrom, S.B., Brailey, S.: Application of projection algorithms to differential equations: boundary value problems. *The ANZIAM Journal* **61**(1), 23–46 (2019)
37. Lindstrom, S.B.: Proximal Point algorithms, dynamical systems, and Associated Operators: Modern Perspectives from Experimental Mathematics. PhD thesis, University of Newcastle (2019)
38. Lindstrom, S.B.: Computable centering methods for spiraling algorithms and their duals, with motivations from the theory of Lyapunov functions. arXiv:2001.10784 (2020)
39. Lindstrom, S.B., Brailey, S.: Survey: sixty years of douglas–Rachford. *J. Aust. Math. Soc.* **110**(3), 333–370 (2021)
40. Lindstrom, S.B., Brailey, S., Skerritt, M.P.: Computing intersections of implicitly specified plane curves. *Nonlinear and Conv. Anal.* **18**(3), 347–359 (2017)
41. Lions, P.L., Mercier, B.: Splitting algorithms for the sum of two nonlinear operators. *SIAM J. Numer. Anal.* **16**(6), 964–979 (1979)
42. Boris, S.M.: *Variational analysis and generalized differentiation i: Basic theory*, volume 330 Springer Science & Business Media (2006)
43. Pierra, G.: Decomposition through formalization in a product space. *Math. Program.* **28**(1), 96–115 (1984)
44. Poon, C., Liang, J.: Trajectory of alternating direction method of multipliers and adaptive acceleration. *NeurIPS 2019* (2019)