

SURVEY OF SEQUENTIAL CONVEX PROGRAMMING AND GENERALIZED GAUSS-NEWTON METHODS *

FLORIAN MESSERER¹, KATRIN BAUMGÄRTNER¹ AND MORITZ DIEHL^{1,2}

Abstract. We provide an overview of a class of iterative convex approximation methods for nonlinear optimization problems with convex-over-nonlinear substructure. These problems are characterized by outer convexities on the one hand, and nonlinear, generally nonconvex, but differentiable functions on the other hand. All methods from this class use only first order derivatives of the nonlinear functions and sequentially solve convex optimization problems. All of them are different generalizations of the classical Gauss-Newton (GN) method. We focus on the smooth constrained case and on three methods to address it: Sequential Convex Programming (SCP), Sequential Convex Quadratic Programming (SCQP), and Sequential Quadratically Constrained Quadratic Programming (SQCQP). While the first two methods were previously known, the last is newly proposed and investigated in this paper. We show under mild assumptions that SCP, SCQP and SQCQP have exactly the same local linear convergence – or divergence – rate. We then discuss the special case in which the solution is fully determined by the active constraints, and show that for this case the KKT conditions are sufficient for local optimality and that SCP, SCQP and SQCQP even converge quadratically. In the context of parameter estimation with symmetric convex loss functions, the possible divergence of the methods can in fact be an advantage that helps them to avoid some undesirable local minima: generalizing existing results, we show that the presented methods converge to a local minimum if and only if this local minimum is stable against a mirroring operation applied to the measurement data of the estimation problem. All results are illustrated by numerical experiments on a tutorial example.

1. INTRODUCTION

Throughout this paper we consider nonlinear optimization problems of the form

$$\min_{w \in \mathbb{R}^n} \phi_0(F_0(w)) \quad \text{s.t.} \quad \begin{aligned} F_i(w) &\in \Omega_i, \quad i = 1, \dots, q, \\ g(w) &= 0, \end{aligned} \quad (1)$$

with nonlinear functions $\phi_0 : \mathbb{R}^{m_0} \rightarrow \mathbb{R}$, $v \mapsto \phi_0(v)$, $F_i : \mathbb{R}^n \rightarrow \mathbb{R}^{m_i}$ and $g : \mathbb{R}^n \rightarrow \mathbb{R}^p$. Critically, the function $\phi_0(v)$ and the sets Ω_i are assumed to be convex. The problem is thus characterized by “convex-over-nonlinear” substructures $\phi_0(F_0(w))$ and $F_i(w) \in \Omega_i$. We call ϕ_0 and the Ω_i “outer convexities” and F_i the “inner nonlinearities”. A special – but still quite general – case is when the sets Ω_i can be described by smooth

* This research was supported by the German Federal Ministry for Economic Affairs and Energy (BMWi) via DyConPV (0324166B), and by DFG via Research Unit FOR 2401.

¹ Department of Microsystems Engineering (IMTEK), University of Freiburg, 79110 Freiburg, Germany. {florian.messerer, katrin.baumgaertner, moritz.diehl}@imtek.de

² Department of Mathematics, University of Freiburg, 79104 Freiburg, Germany

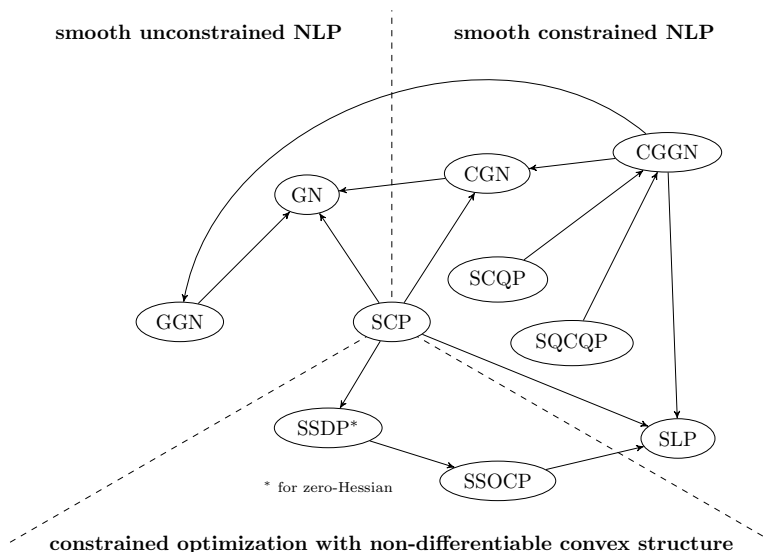


FIGURE 1. Overview of the methods discussed in this paper. Arrows indicate specializations, i.e., $M_1 \rightarrow M_2$ means method M_2 is a special case of method M_1 . SSDP is only a special case of SCP if the SSDP variant with a zero Hessian is used.

convex functions $\phi_i : \mathbb{R}^{m_i} \rightarrow \mathbb{R}$, i.e., $\Omega_i = \{v \in \mathbb{R}^{m_i} \mid \phi_i(v) \leq 0\}$ for $i = 1, \dots, q$. The inequality constraints could then be expressed as $\phi_i(F_i(w)) \leq 0$. We define $f_i(w) := \phi_i(F_i(w))$ as shorthand for this composition, for $i = 0, \dots, q$.

The problem structure (1) arises in various contexts and applications, e.g., nonlinear least-squares formulations for estimation and tracking [5, 24, 34], nonlinear matrix inequalities for reduced order controller design [12, 38], nonlinear model predictive control with ellipsoidal feasible sets [7, 18, 41, 44], robustified inequalities in nonlinear optimization [9, 29], tube-following optimal control problems [39], non-smooth composite minimization [2, 23], as well as training of deep neural networks [25–27, 33, 42].

In this paper, we will give an overview over several methods that aim to exploit this convex-over-nonlinear structure. They do so by sequentially solving a convex approximation to (1), in which the nonlinear functions F_i and g have been linearized. The methods differ in the way they handle the outer convexities. In the case of nonlinear least-squares, that is with $\phi_0(v) = \frac{1}{2}\|v\|_2^2$ and no constraints, all of these methods simplify to the classical Gauss-Newton method (GN). Thus, all of the presented approaches can be seen as generalizations of Gauss-Newton, though only one of these methods will be called the Generalized Gauss-Newton method (GGN) in the following. To avoid confusion it is important to be aware that so far there has been no generally used naming convention and the same name may refer to different methods. For example, GGN in [33] refers to a generalization of the Gauss-Newton method to general smooth outer convexities ϕ_0 , whereas GGN in [5] means the generalization to constrained nonlinear least-squares problems. In this paper we keep the name GGN for the method in [33] and call the method in [5] the Constrained Gauss-Newton method (CGN).

Outline

In this paper we distinguish between three classes of methods, each of them addressing a special case of the general problem (1): smooth unconstrained NLP, smooth constrained NLP and constrained problems with non-differentiable convex structure. Figure 1 gives an overview of the presented methods and the problem class which they address. The most central method is Sequential Convex Programming (SCP), which can address problems from all three classes. It obtains a convex approximation to (1) by linearizing the equality constraints $g(w)$ and the inner nonlinearities $F_i(w)$, but keeping the outer convexities $\phi_0(v)$ and Ω_i . In its most general

form, it can be stated as

$$w_{k+1} \in \arg \min_{w \in \mathbb{R}^n} \phi_0(F_0^{\text{lin}}(w; w_k)) \quad \text{s.t.} \quad \begin{aligned} F_i^{\text{lin}}(w; w_k) &\in \Omega_i, \quad i = 1, \dots, q, \\ g^{\text{lin}}(w; w_k) &= 0, \end{aligned} \quad (2)$$

Besides SCP, the first class, i.e., methods for smooth unconstrained NLP, contains the original Gauss-Newton method (GN) by Gauss [17] and the Generalized Gauss-Newton method (GGN) as introduced by Schraudolph [33]. Section 2 gives a detailed introduction into the algorithms from this class.

In Section 3, we focus on methods for smooth constrained NLP, which is the second problem class. The probably oldest method in this sector is Sequential Linear Programming (SLP), which Griffith and Stewart introduced in 1961 under the name “method of approximation programming” [19]. This sector also contains the classical Constrained Gauss-Newton method (CGN) proposed by Bock [5] for nonlinear least-squares problems with nonlinear constraints. We present several methods that can be seen as generalizations of the CGN method: we first present the smooth constrained variant of SCP, as well as Sequential Convex Quadratic Programming (SCQP) [40], and Sequential Quadratically Constrained Quadratic Programming (SQCQP), a novel method which can be seen as an intermediary between SCP and SCQP. These three methods share some intriguing properties and are the main focus of this paper. To complete the picture, we also introduce the Constrained Generalized Gauss-Newton (CGGN) method, a trivial combination of the CGN and GGN methods. We provide a local convergence analysis of three methods – SCP, SCQP, SQCQP – in Section 4. We show under mild assumptions that all of them share the same asymptotic linear contraction rate, which is exactly determined. This analysis trivially extends to the methods for smooth unconstrained NLP. Furthermore, we analyze a special case, in which the methods converge quadratically. In Section 5, we examine the interesting phenomenon of “mirroring” and desirable divergence, as first described by Bock et al. in the context of L_2 and L_1 estimation [3, 4]. We extend it to estimation problems with general smooth convex negative log-likelihood functions.

Algorithms of the third sector, i.e., constrained optimization problems with non-differentiable convex structures, are mentioned here for completeness, but are not discussed any further in the following. Besides SCP, this category comprises one variant of the sequential semidefinite programming (SSDP) algorithm proposed and analyzed by Fares et al [12]. While different, possibly indefinite Hessian approximations were discussed in the original SSDP paper, only the variant with a zero Hessian falls into the class discussed in this paper. In our notation, this SSDP variant corresponds to the SCP algorithm in Eq. (2) for the case in which there is no convex structure in the objective, $\phi_0(v) = v$, and there is only one structured convex constraint, i.e., $q = 1$, with the corresponding feasible set given by the positive-semidefinite cone, $\Omega_1 = \mathbb{S}_+$ of appropriate dimension. The SSDP algorithm was also used and investigated under the name SSP in [15], and a special SSDP variant exploiting convex-concave decompositions of nonconvex bilinear matrix inequalities was proposed in [36]. For an overview of nonlinear SDP algorithms like SSDP and its variants we refer to the overview paper by Yamashita and Yabe [43]. Another relevant subclass in the nondifferentiable sector is sequential second order cone programming (SSOCP) which can be regarded a special case of SSDP.

All of the methods discussed here formulate and solve convex subproblems in each iteration, and in particular use only positive semidefinite Hessian approximations. While this is desirable for both theoretical and practical reasons, there is a price to pay for convexity of the subproblems: it generally limits the convergence speed to only linear, even if arbitrary positive semidefinite – but bounded – Hessian approximations are used, as shown in [10] for the simple example problem of minimizing the nonconvex objective $f_0(w) = -w_1^2 - (w_2 + 1)^2$ over the Euclidean unit ball $\Omega_1 = \{w \in \mathbb{R}^2 \mid \|w\|_2 \leq 1\}$.

The discussed methods can be seen as Newton-type methods if classified according to [14] or [21]. For example, the methods solving QP subproblems (GN, GGN, CGN, CGGN, SCQP) are straightforward to recognize as variants of Sequential Quadratic Programming (SQP), using a specific positive semi-definite Hessian approximation, whereas SCP can be interpreted as a perturbed version of the Josephy-Newton method, cf. [21, 37]. As there are Newton-type methods which converge superlinearly or even quadratically, this opens the questions of what advantages are to be gained by paying the price of the restriction to a linear convergence rate. The first is the convexity of the subproblems. This makes the subproblems well behaved in the sense that, while

their solutions are not necessarily unique, they are at least not isolated from each other. Further, this means that convex solvers can be utilized, as well as all of the appropriate theory and results from the field of convex optimization [6]. Second, the subproblems are constructed mainly from first-order derivatives, such that almost no second-order derivatives need to be computed, which tend to be expensive. This is especially attractive when problems have to be solved in real time, for example in feedback control, or in the context of large scale optimization problems, where the cost of computing second-order derivatives might be prohibitively high, such as in the training of deep neural networks. Third, most, though not all, of the methods are multiplier-free in the sense that the construction of the subproblems depends only on the primal variables of the current iterate. Therefore the dual variables are not part of the memory.

Notation and preliminaries

We denote by $\frac{\partial f}{\partial w}(w) \in \mathbb{R}^{m \times n}$ the Jacobian of a function $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$, $w \mapsto f(w)$, and use the convention that the gradient symbol denotes its transpose, $\nabla f(w) = \frac{\partial f}{\partial w}(w)^\top$. Thus, for scalar f , the gradient is a column vector. The gradient operator ∇ always refers to the first argument of a function, unless the argument is explicitly written as subscript. For the Jacobian of $F_i(w)$ we use the shortcut $J_i(w)$. If $f(\cdot)$ is scalar valued, its Hessian is denoted by $\nabla^2 f(w)$. Linearizations are referred to as $f^{\text{lin}}(w; \bar{w}) := f(\bar{w}) + \nabla f(\bar{w})^\top (w - \bar{w})$. The notation $f_{\mathcal{S}}(w)$, for a set of indices $\mathcal{S} \subset \{0, \dots, q\}$, means the vertical concatenation of the $f_i(w)$, $i \in \mathcal{S}$. Similarly, for a vector $\mu \in \mathbb{R}^{q+1}$, $\mu_{\mathcal{S}}$ is the vector slice of the corresponding indices. Slightly less conventional, $\phi_{\mathcal{S}}(F_{\mathcal{S}}(w))$ denotes the concatenation of the corresponding $\phi_i(F_i(w))$. The cardinality of set \mathcal{S} is $|\mathcal{S}|$. For a more lightweight notation, the vertical concatenation $[x^\top, y^\top]^\top$ of two vectors $x \in \mathbb{R}^n$, $y \in \mathbb{R}^m$, is denoted by (x, y) . $\|\cdot\|$ denotes a general vector norm. More specifically, $\|\cdot\|_2$ is the Euclidean norm and $\|\cdot\|_\infty$ the maximum norm. For two vectors $x, y \in \mathbb{R}^n$ we use $x \succeq y$ and similar to denote elementwise inequality. For two symmetric matrices $A, B \in \mathbb{R}^{n \times n}$, the matrix inequality $A \succeq B$ means $A - B \succeq 0$, i.e., that $A - B$ is positive semidefinite, and accordingly for similar inequalities. The identity matrix in $\mathbb{R}^{n \times n}$ is denoted by I_n and the subscript n may be dropped if the dimension is clear from the context. The set of non-negative reals is \mathbb{R}_+ . Concepts from the field of numerical optimization, such as strict complementarity, the linear independence constraint qualification (LICQ) and the Karush Kuhn Tucker (KKT) conditions are defined as in textbooks such as [30] unless otherwise stated. For several important results of this paper we require the functions defining problem (1) to be smooth. Technically, requiring the $F_i(w)$ and $g(w)$ to be twice, the $\phi_i(v)$ three times continuously differentiable would be sufficient for all results. Since this difference is in practice often irrelevant, and for a simpler text flow, we decide to use “smooth” as requirement, but by this we actually mean the just stated weaker requirements.

Illustrative Example

We now introduce an example that will accompany us throughout this article. In this example we use the pseudo Huber loss function

$$\varphi_\delta : \mathbb{R}^N \rightarrow \mathbb{R}, \quad \varphi_\delta(v) := \sum_{i=1}^N \sqrt{\delta^2 + v_i^2} - \delta, \quad (3)$$

parameterized by $\delta \in \mathbb{R}_+$. For v close to 0, it approximates a least-squares loss function, whereas for v far from 0, it has linear behavior similar to the L_1 -norm. The size of the quadratic region is controlled by the Huber parameter δ . The larger δ , the larger the quadratic region. In the limit case of $\delta \rightarrow 0$, the pseudo Huber loss approaches the L_1 -norm. A visualization of this behavior is given in Figure 2 for a single component, $N = 1$. The example problems are implemented via the Python interface of CasADi [1], using the solver Gurobi [20] for second order cone programs (SOCP), and qpOASES for quadratic programs (QP) [13].

Example 1.1. Assume we have modeled the explicit time dependency of some output $\psi \in \mathbb{R}$ as

$$\psi(t) = \frac{3}{4}t + \sin(t). \quad (4)$$




FIGURE 2. The pseudo Huber loss $\varphi_\delta(v)$ as defined in (3) for two values of δ ($\delta_1 = 10^{-2}$, $\delta_2 = 10^{-1}$) compared to the absolute value $|v|$, which corresponds to the L_1 norm.

We have noisy measurements η_i of this output, obtained at times x_i , but they are associated with some unknown time delay w , i.e., $x_i = t_i - w$. Our aim is to identify this true time delay w from N input-output pairs (x_i, η_i) , $i = 1, \dots, N$, such that we obtain an estimate of the true time $t_i = x_i + w$ at which the output $\psi(t_i)$ occurred. We model the measurements as $\eta_i = \psi(x_i + w) + \nu_i$, where ν_i is unknown noise. The η_i are collected in $\eta \in \mathbb{R}^N$ and the model predictions in $M(w)$, $M : \mathbb{R} \rightarrow \mathbb{R}^N$, with $M_i(w) := \psi(x_i + w)$. If we choose the pseudo Huber loss (3) as penalty of the model-measurement mismatch, we obtain

$$\min_{w \in \mathbb{R}} \varphi_\delta(\eta - M(w)) \quad (5)$$

as our identification problem. This has convex-over-nonlinear structure, with outer convexity $\phi_0(v) = \varphi_\delta(v)$ and inner nonlinearity $F_0(w) = \eta - M(w)$. For the purpose of a clean demonstration of the concepts presented in this paper, we assume $N = 3$ with the data given as $x = (-0.5, 0, 0.5)$ and $\eta = (0, 0, 1)$. Unless otherwise stated, the Huber parameter is chosen as $\delta = 0.1$.

2. METHODS FOR SMOOTH UNCONSTRAINED NLP

Let us first regard only the unconstrained case,

$$\min_{w \in \mathbb{R}^n} \underbrace{\phi_0(F_0(w))}_{=: f_0(w)}, \quad (6)$$

with $\phi_0(v)$ and $F_0(w)$ smooth, and introduce two convexity exploiting methods, the unconstrained version of SCP and the Generalized Gauss-Newton method (GGN). Both of them are generalizations of the classical Gauss-Newton method to a general smooth outer convexity $\phi_0(v)$. At the end of the section we state a theorem that exactly characterizes their linear local convergence rate.

2.1. Sequential Convex Programming

Sequential Convex Programming (SCP) obtains a convex approximation of (6) by linearizing the inner nonlinearity $F_0(w)$ at the current iterate w_k , but keeping the full outer convexity $\phi_0(v)$. It thus iterates as

$$w_{k+1} \in \arg \min_{w \in \mathbb{R}^n} \underbrace{\phi_0(F_0^{\text{lin}}(w; w_k))}_{=: f_0^{\text{SCP}}(w; w_k)}, \quad (7)$$

solving a convex – but generally nonlinear – subproblem at every iteration. In case the minimizer of (7) is not uniquely defined, one has to take extra measures, such as picking the minimum norm solution, in order to

obtain a well defined algorithm. If we apply SCP to the nonlinear least-squares problem, i.e., if $\phi_0(v) = \frac{1}{2}\|v\|_2^2$, we recover the classical Gauss-Newton method, as in this case $f_0^{\text{SCP}}(w; w_k) = \frac{1}{2}\|F_0^{\text{lin}}(w; w_k)\|_2^2$.

2.2. Generalized Gauss-Newton

Generalized Gauss-Newton (GGN) iterates by solving a convex quadratic approximation of (6) at every iteration [33]. As in the classical Gauss-Newton method, this approximation is obtained by splitting the Hessian of the objective $f_0(w)$ into two terms, one of them positive semidefinite, and the other generally indefinite. The Hessian approximation is then obtained by neglecting the indefinite term. For the objective function $f_0(w)$ this split is given by

$$\nabla^2 f_0(w) = \underbrace{J_0(w)^\top \nabla^2 \phi_0(F(w)) J_0(w)}_{=: B_{\text{GGN}}(w)} + \underbrace{\sum_{i=1}^{m_0} \nabla^2 F_{0,i}(w) \frac{\partial \phi_0}{\partial v_i}(F_0(w))}_{=: E_{\text{GGN}}(w)}, \quad (8)$$

where $\nabla^2 F_{0,i}(w)$ denotes the Hessian of the i -th component of $F_0(w)$. The first term $B_{\text{GGN}}(w)$ is the Generalized Gauss-Newton Hessian approximation and contains the curvature of the outer convexity $\phi_0(v)$. Due to convexity of ϕ_0 , we have $B_{\text{GGN}}(w) \succeq 0$ for all w . Note that apart from $\nabla^2 \phi_0(v)$ one only needs first-order derivatives to evaluate $B_{\text{GGN}}(w)$. The ‘‘error matrix’’ $E_{\text{GGN}}(w)$ comprises the components we neglect by choosing $B_{\text{GGN}}(w)$ as our Hessian approximation. In particular, it contains the curvature of the inner nonlinearity $F_0(w)$. This error is small if the neglected curvature $\nabla^2 F_{0,i}(w)$ or the derivatives $\frac{\partial \phi_0}{\partial v_i}(F_0(w))$ are small. The second condition in particular includes the case that the residuals $F_0(w)$ are close to a minimizer of $\phi_0(v)$ (if it exists). Intuitively, the smaller the approximation error, and therefore the closer $B_{\text{GGN}}(w)$ is to the true Hessian, the better is the convergence behavior we would expect. We will exactly quantify this intuition in Theorem 2.2.

The GGN iterations are defined via the unconstrained Quadratic Program (QP)

$$w_{k+1} \in \arg \min_{w \in \mathbb{R}^n} \underbrace{\frac{1}{2}(w - w_k)^\top B_{\text{GGN}}(w_k)(w - w_k) + \nabla f_0(w_k)^\top (w - w_k) + f_0(w_k)}_{=: f_0^{\text{GGN}}(w; w_k)}. \quad (9)$$

Since $B_{\text{GGN}}(w) \succeq 0$ for all w , this QP is always convex. As for SCP, in case the solution to (9) is not unique, we need a procedure for picking exactly one solution, such that w_{k+1} is uniquely defined. A definition equivalent to (9) is obtained by noting that (9) in effect amounts to solving the linear system

$$w_{k+1} \in \{w \in \mathbb{R}^n \mid B_{\text{GGN}}(w_k)(w - w_k) = -\nabla f_0(w_k)\} \quad (10)$$

at every iteration. This is – in general – considerably faster than solving a full nonlinear optimization problem at every iteration, as is necessary in SCP. If the SCP subproblem is solved with a Newton-type algorithm, this in fact means several iterations with the structure of (10) per iteration of SCP. For the special case of nonlinear least-squares, i.e., $\phi_0(v) = \frac{1}{2}\|v\|_2^2$, we have $\nabla^2 \phi_0(v) = I$. We therefore obtain $B_{\text{GGN}}(w) = J(w)^\top J(w)$. This is the well-known Gauss-Newton Hessian approximation.

Example 2.1. Recall the guiding Example 1.1. We compute the SCP and GGN approximations to the objective function, i.e., $f_0^{\text{SCP}}(w; \bar{w})$ resp. $f_0^{\text{GGN}}(w; \bar{w})$ as defined in (7) and (9). We do so for two distinct linearization points, $\bar{w}_1 = 0.1$ and $\bar{w}_2 = 0.3$. The results are illustrated in Figure 3. It can be seen that SCP manages a closer approximation to the objective function $f_0(w)$, since it keeps the information about the specific shape of the outer convexity $\phi_0(v)$. GGN on the other hand always approximates the objective as a quadratic function.

2.3. Local Convergence Analysis

We state here already a theorem on the linear local convergence rate of unconstrained SCP and GGN. This is actually a special case of Theorem 4.5 which will be proven later for the smooth constrained case. We therefore refrain from giving the proof here.

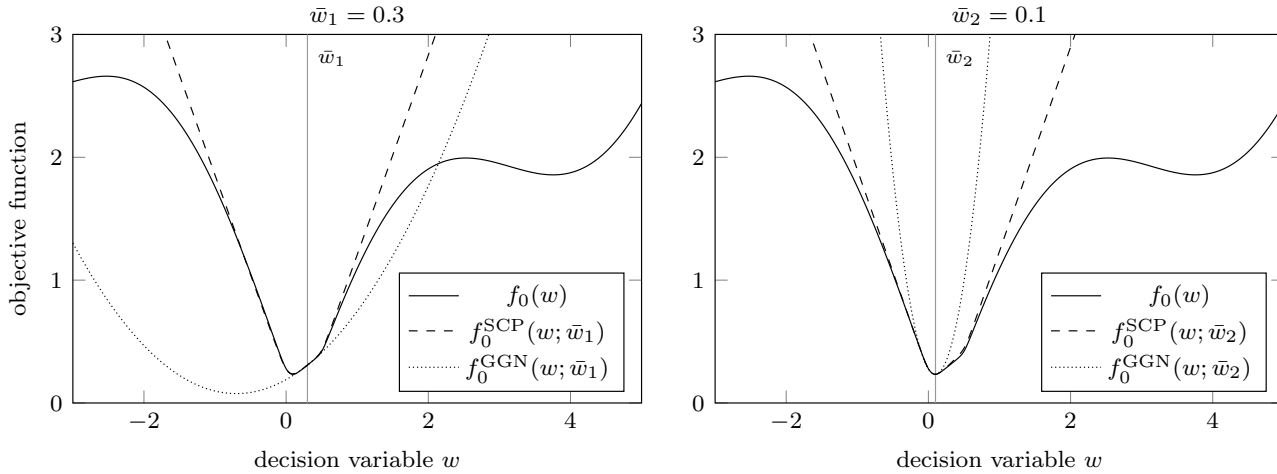


FIGURE 3. Illustration of the SCP and GGN approximations to the nonlinear objective $f_0(w)$ for two values of \bar{w} . GGN approximates $f_0(w)$ only quadratically, whereas SCP is able to match the characteristic shape of the outer convexity $\phi_0(v)$. We can also see that for \bar{w}_2 , which is close to a local minimum, this difference does not seem to be too important, since locally both methods provide a good approximation of the true objective.

Theorem 2.2 (Linear local convergence of SCP and GGN [11]). *Regard a local minimizer w^* of f_0 that satisfies $\nabla f_0(w^*) = 0$ and $B_{\text{GGN}}(w^*) \succ 0$. Then w^* is a fixed point for both the SCP and GGN iterations, the iterates of both methods are well-defined in a neighborhood of w^* , and the local linear contraction – or divergence – rates of SCP and GGN are equal to each other and given by the smallest $\alpha \geq 0$ that satisfies the linear matrix inequalities (LMI)*

$$-\alpha B_{\text{GGN}}(w^*) \preceq E_{\text{GGN}}(w^*) \preceq \alpha B_{\text{GGN}}(w^*). \quad (11)$$

As a consequence, a sufficient condition for Q -linear local convergence with contraction rate $\alpha < 1$ is given by the LMI

$$\frac{1}{1+\alpha} \nabla^2 f_0(w^*) \preceq B_{\text{GGN}}(w^*) \preceq \frac{1}{1-\alpha} \nabla^2 f_0(w^*). \quad (12)$$

Also, a necessary condition for local convergence is given by $B_{\text{GGN}}(w^*) \succeq \frac{1}{2} \nabla^2 f_0(w^*)$. If $\nabla^2 f_0(w^*) \succ 0$, a sufficient condition for local convergence is given by $B_{\text{GGN}}(w^*) \succ \frac{1}{2} \nabla^2 f_0(w^*)$.

Example 2.3. We return to our example problem defined in (5). Since $w \in \mathbb{R}$, the LMI in (11) simplify to scalar inequalities. We can thus explicitly compute the smallest α satisfying (11) as

$$\check{\alpha}(w) := \frac{|\nabla^2 f_0(w) - B_{\text{GGN}}(w)|}{|B_{\text{GGN}}(w)|}. \quad (13)$$

Only for a local minimizer w^* the interpretation of $\check{\alpha}(w^*)$ as linear local convergence rate is valid, but it is still interesting to visualize $\check{\alpha}(w)$ for any $w \in \mathbb{R}$. In Figure 4, the objective function $f_0(w)$ as well as $\check{\alpha}(w)$ are illustrated for the example problem. For the local minimum at $w_{\text{good}}^* \approx 0.1$ – which is actually the global minimum – we compute the theoretical contraction rate $\check{\alpha}(w_{\text{good}}^*) \approx 0.02$. Furthermore, we hope to awaken the reader's curiosity by pointing to an interesting observation: there is also a second, worse, local minimum at $w_{\text{bad}}^* \approx 3.7$, for which holds $\check{\alpha}(w_{\text{bad}}^*) \approx 3200 \gg 1$. This means that SCP and GGN would strongly diverge from this undesirable local minimum. This is actually not a coincidence and later in this paper we dedicate a full section to this behavior.




FIGURE 4. Visualization of the objective function and $\check{\alpha}(w)$. Note that $\check{\alpha}(w)$ attains its meaning as local contraction rate only at local minima.

3. METHODS FOR SMOOTH CONSTRAINED NLP

We will now return to the constrained case and consider methods that can be applied to NLP of the form

$$\min_{w \in \mathbb{R}^n} \phi_0(F_0(w)) \quad \text{s.t.} \quad \begin{aligned} \phi_i(F_i(w)) &\leq 0, & i = 1, \dots, q, \\ g(w) &= 0, \end{aligned} \quad (14)$$

composed of only smooth functions, and with $\phi_i(F_i(w)) =: f_i(w)$, $i = 0, \dots, q$. We define its Lagrangian as

$$\mathcal{L}(w, \lambda, \mu) = \phi_0(F_0(w)) + \sum_{i=1}^q \mu_i \phi_i(F_i(w)) + \sum_{i=1}^p \lambda_i g_i(w), \quad (15)$$

with Lagrange multipliers – or dual variables – $\mu \in \mathbb{R}^q$ and $\lambda \in \mathbb{R}^p$. In this section, we discuss four methods that exploit the convex substructure of (14), namely the smooth constrained version of SCP, Constrained Generalized Gauss-Newton (CGGN), Sequential Convex Quadratic Programming (SCQP) and Sequential Quadratically Constrained Quadratic Programming (SQCQP).

3.1. Sequential Convex Programming

In a straight-forward generalization of the unconstrained case, SCP approximates (14) by linearizing the inner nonlinearities $F_i(w)$, while keeping the outer convexities $\phi_i(v)$. In every iteration, SCP thus solves a nonlinear but convex optimization problem of the following form:

$$w_{k+1} \in \arg \min_{w \in \mathbb{R}^n} \phi_0(F_0^{\text{lin}}(w; w_k)) \quad \text{s.t.} \quad \begin{aligned} \phi_i(F_i^{\text{lin}}(w; w_k)) &\leq 0, & i = 1, \dots, q, \\ g^{\text{lin}}(w; w_k) &= 0, \end{aligned} \quad (16)$$

In accordance with the definition of $f_0^{\text{SCP}}(w; w_k)$, we introduce $f_i^{\text{SCP}}(w; w_k) := \phi_i(F_i^{\text{lin}}(w; w_k))$ for $i = 0, \dots, q$ as shorthand for the approximation functions. We point out that – though every solution to (16) is also associated with Lagrange multipliers $\mu_{k+1} \in \mathbb{R}^q$ and $\lambda_{k+1} \in \mathbb{R}^p$ – the iterations only depend on the primal variable w_k . In other words, SCP is a multiplier-free method. For constrained nonlinear least-squares with no (or ignored) convex substructure in the constraints, i.e., $\phi_0(v) = \frac{1}{2} \|v\|_2^2$ and $\phi_i(v) = v$ for $i = 1, \dots, q$, SCP

recovers the constrained Gauss-Newton (CGN) method [5], in which the objective function is approximated as $\frac{1}{2}\|F_0^{\text{lin}}(w; \bar{w})\|_2^2$ and the constraints are completely linearized. Another special case occurs if additionally there is no convex substructure in the objective, $\phi_0(v) = v$. Then also the objective function is completely linearized and SCP simplifies to Sequential Linear Programming (SLP) [19].

3.2. Constrained Generalized Gauss-Newton

As for the unconstrained case, SCP has to solve a generally nonlinear program at every iteration, which can be quite expensive. We therefore turn our attention to methods that approximate (14) by cheaper subproblems, similar to GGN. Unlike SCP, the generalization of GGN to constrained NLP can go into several directions, such that none of them can be considered to be the most straightforward. Maybe the simplest idea is to use the GGN Hessian approximation $B_{\text{GGN}}(w; \bar{w})$ – as defined in (8) – for a quadratic approximation of the objective function and to linearize the constraints in order to obtain a constrained QP. This yields a method that we call the Constrained Generalized Gauss-Newton method (CGGN), with the subproblem defined as

$$\begin{aligned} w_{k+1} \in \arg \min_{w \in \mathbb{R}^n} & \quad \frac{1}{2}(w - w_k)^\top B_{\text{GGN}}(w_k)(w - w_k) + \nabla f_0(w_k)^\top (w - w_k) + f_0(w_k) \\ \text{s.t.} & \quad f_i^{\text{lin}}(w; w_k) \leq 0, \quad i = 1, \dots, q, \\ & \quad g^{\text{lin}}(w; w_k) = 0. \end{aligned} \tag{17}$$

To the authors' knowledge, CGGN has never been explicitly discussed before. Given the straightforwardness of the approach, this seems surprising. It might be due to the fact that GGN originates from the neural network community, which usually deals with unconstrained optimization problems [26, 33]. Note that CGGN only uses information about the curvature of the objective function in its approximation, as $B_{\text{GGN}}(\bar{w})$ only depends on $f_0(w)$. Possible convex curvature information of the inequality constraints is thrown away. This method is therefore mainly interesting for problems without exploitable convex substructure in the constraints, i.e., when $\phi_i(v) = v$ for $i \geq 1$. In the case of nonlinear least-squares, $\phi_0(v) = \frac{1}{2}\|v\|_2^2$, the CGGN method simplifies to CGN [5]. CGGN can therefore be seen both as a generalization of GGN to constrained NLP and as generalization of CGN to general smooth convexities in the objective. When there is no convexity in the objective function, $\phi_0(v) = v$, the Hessian approximation is $B_{\text{GGN}}(\bar{w}) = 0$, and CGGN recovers SLP [19].

3.3. Sequential Convex Quadratic Programming

We will now try to improve upon CGGN by exploiting our knowledge of the convex-over-nonlinear structure in the constraints, while still only solving a QP at every iteration. This leads us to Sequential Convex Quadratic Programming (SCQP), which was introduced by Verschueren et al in [40]. Consider the Hessian of the Lagrangian (15) of the original problem, which is given by

$$\begin{aligned} \nabla^2 \mathcal{L}(w, \lambda, \mu) &= \overbrace{J_0(w)^\top \nabla^2 \phi_0(F_0(w)) J_0(w) + \sum_{i=1}^q \mu_i J_i(w)^\top \nabla^2 \phi_i(F_i(w)) J_i(w)}^{=: B_{\text{SCQP}}(w, \mu)} \\ &+ \underbrace{\sum_{j=1}^{m_0} \nabla^2 F_{0,j}(w) \frac{\partial \phi_0}{\partial v_j}(F_0(w)) + \sum_{i=1}^q \mu_i \sum_{j=1}^{m_i} \nabla^2 F_{i,j}(w) \frac{\partial \phi_i}{\partial v_j}(F_i(w)) + \sum_{i=1}^p \lambda_i \nabla^2 g_i(w)}_{=: E_{\text{SCQP}}(w, \mu, \lambda)}, \end{aligned} \tag{18}$$

where $F_{i,j}$ denotes the j -th entry of F_i . In the spirit of Gauss-Newton, we split the true Hessian into the SCQP Hessian approximation $B_{\text{SCQP}}(w, \mu)$ – containing the curvature of all outer convexities $\phi_i(v)$ – and the SCQP

Hessian approximation error $E_{\text{SCQP}}(w, \mu, \lambda)$, which contains the curvature of the inner nonlinearities $F_i(w)$ and the equality constraints $g(w)$. We introduce the shorthands

$$B_i(w) := J_i(w)^\top \nabla^2 \phi_i(F_i(w)) J_i(w) \quad \text{for } i = 0, \dots, q, \quad (19)$$

such that $B_{\text{SCQP}}(w, \mu) = B_0(w) + \sum_{i=1}^q \mu_i B_i(w)$, and see that $B_0(w)$ is exactly the GGN Hessian as defined in (8), i.e., $B_0(w) = B_{\text{GGN}}(w)$. An immediate consequence is the following proposition.

Proposition 3.1. *Assume (w^*, λ^*, μ^*) is a KKT point of (14). Then it holds that $B_{\text{SCQP}}(w^*, \mu^*) \succeq B_{\text{GGN}}(w^*)$.*

Proof.

$$B_{\text{SCQP}}(w^*, \mu^*) = B_0(w^*) + \sum_{i=1}^q \mu_i^* B_i(w^*) = B_{\text{GGN}}(w^*) + \sum_{i=1}^q \underbrace{\mu_i^*}_{\geq 0} \underbrace{B_i(w^*)}_{\succeq 0} \succeq B_{\text{GGN}}(w^*) \quad (20)$$

□

We can now use $B_{\text{SCQP}}(w, \mu)$ as Hessian in a QP approximation to (14), and define the SCQP subproblem as

$$\begin{aligned} w_{k+1} \in \arg \min_{w \in \mathbb{R}^n} & \quad \frac{1}{2} (w - w_k)^\top B_{\text{SCQP}}(w_k, \mu_k) (w - w_k) + \nabla f_0(w_k)^\top (w - w_k) + f_0(w_k) \\ \text{s.t.} & \quad f_i^{\text{lin}}(w; w_k) \leq 0, \quad i = 1, \dots, q, \\ & \quad g^{\text{lin}}(w; w_k) = 0, \end{aligned} \quad (21)$$

with $f_0^{\text{SCQP}}(w; \bar{w})$ the shorthand for the objective function. Here, we point to a subtle, but important, difference that should be kept in mind: all convex approximation schemes introduced so far, and also all that will follow, only depend on the current primal iterate w_k as linearization point. To construct $B_{\text{SCQP}}(w_k, \mu_k)$ on the other hand, we also need the dual iterate μ_k . The full SCQP iteration scheme thus needs to extract μ_{k+1} from the multipliers satisfying the KKT conditions of (21). An immediate consequence is that all components of μ_{k+1} are nonnegative. It follows that, similar to Proposition 3.1, we have $B_{\text{SCQP}}(w_{k+1}, \mu_{k+1}) \succeq 0$ for all $k \geq 0$. Therefore we only need to pick an initialization $\mu_0 \geq 0$ to ensure convexity of (21) at all iterations. If $\phi_i(v) = v$ for all $i \geq 1$, we recover CGGN, since then $\nabla^2 \phi_i(v) = 0$ for all $i \geq 1$.

3.4. Sequential Quadratically Constrained Quadratic Programming

A different approach to exploit the known curvature information would be to approximate both the objective function and the inequality constraints by convex quadratic functions. This leads to the Sequential Quadratically Constrained Quadratic Programming method (SQCQP). To our knowledge this method – using generalized Gauss-Newton Hessian approximations to ensure convexity of the subproblems – was not yet presented in the literature. An exact-Hessian variant has been proposed and discussed in [16]. SQCQP solves a Quadratically Constrained Quadratic Program (QCQP) at every iteration:

$$\begin{aligned} w_{k+1} \in \arg \min_{w \in \mathbb{R}^n} & \quad \frac{1}{2} (w - w_k)^\top B_0(w_k) (w - w_k) + \nabla f_0(w_k)^\top (w - w_k) + f_0(w_k) \\ \text{s.t.} & \quad \frac{1}{2} (w - w_k)^\top B_i(w_k) (w - w_k) + \nabla f_i(w_k)^\top (w - w_k) + f_i(w_k) \leq 0, \quad i = 1, \dots, q, \\ & \quad g^{\text{lin}}(w; w_k) = 0, \end{aligned} \quad (22)$$

with the Hessian approximations $B_i(w)$ as defined in (19). We abbreviate these quadratic approximations by $f_i^{\text{quad}}(w; \bar{w}) := \frac{1}{2} (w - w_k)^\top B_i(w_k) (w - w_k) + \nabla f_i(w_k)^\top (w - w_k) + f_i(w_k)$. Compared to SCQP, this method should lead to a slightly closer approximation of the original NLP, as the constraints are approximated quadratically instead of linearized. On the other hand, each iteration is in general more expensive, as a QCQP has to be solved instead of a QP. Compared to SCP, the QCQP approximation is generally worse, while each iteration of

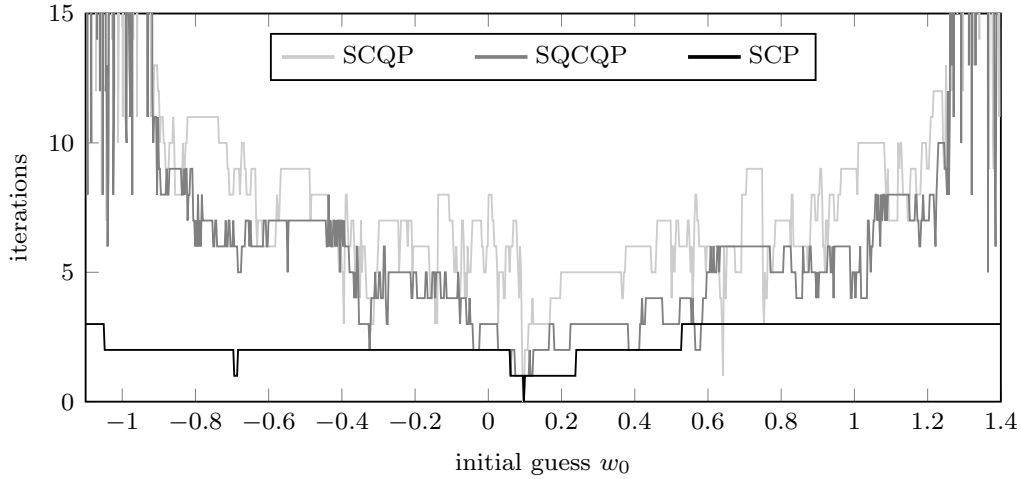


FIGURE 5. Number of iterations until convergence to $w^* \approx 0.1$ depending on the initial guess w_0 . SCP converged to w^* in 100.0% of the cases in 100 iterations or less, SQCQP in 95.7%, and SCQP in 90.3%.

SQCQP is possibly cheaper. We might say that SQCQP is in between SCP and SCQP. As for SCQP, for the case that $\phi_i(v) = v$ for $i = 1, \dots, q$, we get $\nabla^2 \phi_i(v) = 0$ for all $i \geq 1$ and therefore recover the CGGN method.

Example 3.2. We now illustrate the global convergence behavior of SCP, SCQP and SQCQP. To this end, we revisit Example 1.1. We reformulate the parameter estimation problem (5) as

$$\min_{w, s} \sum_{i=1}^N s_i \quad \text{s.t.} \quad \begin{aligned} 0 &\leq s_i, & i = 1, \dots, N, \\ \sqrt{\delta^2 + F_i(w)^2} - \delta &\leq s_i, & i = 1, \dots, N, \end{aligned} \quad (23)$$

where we introduced slack variables $s \in \mathbb{R}^N$, and subsumed the model-measurement residual in $F_i(w) = \eta_i - M_i(w)$. Note that the positiveness constraint on s is not strictly necessary as it is implicitly enforced by the first constraint. For the SCQP and SQCQP approximations to (23) though, this is not the case and including the constraint significantly improves their global convergence behavior. Note that the SCP subproblem is a second order cone program (SOCP) and in consequence solved as such. We initialize all three methods at 1000 different values of w_0 , linearly spaced between -1.1 and 1.5. The slack variables are initialized as zero, and, in the case of SCQP, all multipliers by one. The resulting number of iterations is shown in Figure 5. As heuristically predicted, SCP shows the best global convergence behavior, both in terms of the number of iterations and the fraction of initializations for which it converges. Furthermore, it has the steadiest behavior, in the sense that the number of iterations does not jump around wildly for small changes in w_0 . SCQP shows the worst behavior for these three measures, whereas SQCQP is in between the other two methods. Nonetheless we point out that the take-away message of this example is not that SCP should always be the method of choice. This would ignore the fact that the iterations of SCQP are significantly cheaper than those of SCP. A globalized SCQP method might outperform full step SCP on the considered properties.

4. LOCAL CONVERGENCE ANALYSIS

We will now investigate the local convergence behavior of the introduced methods. The most general are SCP, SCQP and SQCQP, so we will focus our attention on them. All other methods, namely GGN and CGGN, and of course also GN and CGN, are special cases in the absence of convex structure in the constraints. Therefore

the results trivially extend to them as well. For unconstrained SCP and GGN, the results on linear convergence have been obtained already in [11]. Tran-Dinh et al. proved linear convergence of general SCP methods under mild assumptions, but without a tight characterization of the rate [35]. For SCQP a tight characterization has been obtained in [40], and for constrained SCP in [28], while SQCQP has not yet been investigated before.

We start by establishing stationarity of the methods at a solution to (14). Collecting both primal and dual variables in $z_k = (w_k, \mu_k, \lambda_k)$, the subproblems of each method, i.e., (16), (21), resp. (22), define iteration maps

$$z_{k+1} = z_{\text{SCP}}^{\text{sol}}(z_k), \quad z_{k+1} = z_{\text{SCQP}}^{\text{sol}}(z_k) \quad \text{resp.} \quad z_{k+1} = z_{\text{SQCQP}}^{\text{sol}}(z_k). \quad (24)$$

We assume uniqueness of the next iterate has been assured by taking the solution closest to z_k . More specifically, if a subproblem has a set of solutions \mathcal{Z} , we define $z_{k+1} := \arg \min_{z \in \mathcal{Z}} \|z - z_k\|$. Note that in order to obtain a simpler notation, the iteration maps have been defined as taking the full z_k as input. This hides the fact that $z_{\text{SCP}}^{\text{sol}}(z_k)$ and $z_{\text{SQCQP}}^{\text{sol}}(z_k)$ only depend on w_k , since they are multiplier-free, and that $z_{\text{SCQP}}^{\text{sol}}(z_k)$ additionally depends on μ_k , but none of the methods depends on λ_k .

We denote the set of KKT points of (14) as \mathcal{Z}^* . The active set $\mathcal{A}(z) \subseteq \{1, \dots, q\}$ contains the indices of all inequality constraints which are active at z . For a specific $z^* \in \mathcal{Z}^*$, we use the shorthand $\mathcal{A}^* := \mathcal{A}(z^*)$, with corresponding multipliers $\mu_{\mathcal{A}^*}$. Accordingly, we have the inactive set $\mathcal{I}(z)$ with $\mathcal{I}^* := \mathcal{I}(z^*)$ and $\mu_{\mathcal{I}^*}$.

Lemma 4.1. *Let z^* be a feasible point of (14) at which LICQ holds. Then z^* is a fixed point of all three iteration maps in (24) if and only if $z^* \in \mathcal{Z}^*$. Furthermore, strict complementarity of the SCP, SCQP and SQCQP subproblems (16), (21) and (22) holds at $z^* \in \mathcal{Z}^*$ if and only if it also holds for (14) at z^* . More specifically, the active set \mathcal{A}^* and the corresponding multipliers are identical.*

Proof. To increase readability we only state the proof for SCP. The proofs for SCQP and SQCQP are completely analogous. Assume \bar{z} is a fixed point of $z_{\text{SCP}}^{\text{sol}}(z_k)$, i.e., $z_{\text{SCP}}^{\text{sol}}(\bar{z}) = \bar{z}$. This means that \bar{z} solves the SCP subproblem (16) and, due to LICQ, is also a KKT point. Then, by substituting $z_{k+1} = z_k = \bar{z}$ into the KKT conditions of (16) – which we know to hold in this case – they collapse to the KKT conditions of (14) – which therefore also hold. It follows that $\bar{z} \in \mathcal{Z}^*$. Vice versa, assume that $z^* \in \mathcal{Z}^*$. Writing down the KKT conditions of (16), we can see that they hold at $z_{k+1} = z_k = z^*$. Therefore, due to convexity, z^* solves (16) at z^* . Denote by \mathcal{Z} the set of all solutions to (16) at z^* . We know that $z^* \in \mathcal{Z}$. From $z_{k+1} = \arg \min_{z \in \mathcal{Z}} \|z - z^*\| = z^*$ it follows that z^* is stationary w.r.t. the SCP iteration defined in (24). For the statement on strict complementarity we point again to the fact that the KKT conditions of (14) and (16) are identical at $z_{k+1} = z_k = z^* \in \mathcal{Z}^*$. \square

In order to prepare the statement and proof of the main theorem of this paper, Theorem 4.5, we start by pointing out that for a KKT point z^* the inactive constraints have no influence on its position (apart from necessitating $\mu_{\mathcal{I}^*} = 0$). The active inequality constraints on the other hand can be treated exactly like equality constraints. In fact, if the active set $\mathcal{A}(w^*)$ of a local minimizer w^* is known, finding w^* simplifies to the solution of an equality constrained problem. Furthermore, if the active set is stable, i.e., if it does not change close to w^* , then the inactive inequality constraints will also not influence the local convergence behavior, and for the purpose of local convergence analysis we can completely disregard them. Having this already in mind, we need to set up some notation and intermediate results. We consider a local minimizer w^* of (14) with active set \mathcal{A}^* and define the corresponding extended equality constraints as

$$G(w) := \begin{bmatrix} f_{\mathcal{A}^*}(w) \\ g(w) \end{bmatrix}, \quad G_{\text{SCP}}(w; \bar{w}) := \begin{bmatrix} f_{\mathcal{A}^*}^{\text{SCP}}(w; \bar{w}) \\ g^{\text{lin}}(w; \bar{w}) \end{bmatrix}, \quad G_{\text{SCQP}}(w; \bar{w}) := \begin{bmatrix} f_{\mathcal{A}^*}^{\text{lin}}(w; \bar{w}) \\ g^{\text{lin}}(w; \bar{w}) \end{bmatrix}, \quad G_{\text{SQCQP}}(w; \bar{w}) := \begin{bmatrix} f_{\mathcal{A}^*}^{\text{quad}}(w; \bar{w}) \\ g^{\text{lin}}(w; \bar{w}) \end{bmatrix}, \quad (25)$$

with output dimension $r := |\mathcal{A}^*| + p$. The associated Lagrange multipliers are collected in $\gamma = (\mu_{\mathcal{A}^*}, \lambda)$. We note that if their Jacobians with respect to the first argument are evaluated at the given linearization point \bar{w} , we obtain

$$\Gamma(\bar{w}) := \left. \frac{\partial G(w)}{\partial w} \right|_{w=\bar{w}} = \left. \frac{\partial G_{\text{SCP}}(w; \bar{w})}{\partial w} \right|_{w=\bar{w}} = \left. \frac{\partial G_{\text{SCQP}}(w; \bar{w})}{\partial w} \right|_{w=\bar{w}} = \left. \frac{\partial G_{\text{SQCQP}}(w; \bar{w})}{\partial w} \right|_{w=\bar{w}} = \begin{bmatrix} \nabla f_{\mathcal{A}^*}(\bar{w})^\top \\ \nabla g(\bar{w})^\top \end{bmatrix}. \quad (26)$$

This is not surprising, given that at the linearization point all three approximations match the original NLP up to first order. We also regard the Hessians of the Lagrangians of the subproblems (16), (21) resp. (22) at \bar{z} , which are

$$B_{\text{SCP}}(w, \mu; \bar{w}) := J_0(\bar{w})^\top \nabla^2 \phi_0(F_0^{\text{lin}}(w; \bar{w})) J_0(\bar{w}) + \sum_{i=1}^q \mu_i J_i(\bar{w})^\top \nabla^2 \phi_i(F_i^{\text{lin}}(w; \bar{w})) J_i(\bar{w}), \quad (27)$$

$$B_{\text{SCQP}}(\bar{w}, \bar{\mu}) := J_0(\bar{w})^\top \nabla^2 \phi_0(F_0(\bar{w})) J_0(\bar{w}) + \sum_{i=1}^q \bar{\mu}_i J_i(\bar{w})^\top \nabla^2 \phi_i(F_i(\bar{w})) J_i(\bar{w}), \quad (28)$$

$$B_{\text{SQCQP}}(\mu; \bar{w}) := J_0(\bar{w})^\top \nabla^2 \phi_0(F_0(\bar{w})) J_0(\bar{w}) + \sum_{i=1}^q \mu_i J_i(\bar{w})^\top \nabla^2 \phi_i(F_i(\bar{w})) J_i(\bar{w}). \quad (29)$$

As the expressions look confusingly similar, we explicitly point out the differences. First, recall that w and μ are variables of the subproblems, whereas \bar{w} and $\bar{\mu}$ are parameters defining the point of approximation. SCP and SQCQP are multiplier-free methods, so their Hessians do not depend on $\bar{\mu}$. B_{SCP} depends on w via the functions $F_i^{\text{lin}}(w; \bar{w})$. As these are the arguments of the $\nabla_i^2 \phi(\cdot)$, SCP keeps some curvature information of ϕ_i even when moving away from \bar{w} . The SCQP subproblem is a QP. Therefore its Hessian is constant with respect to the decision variables, and only depends on the linearization point $(\bar{w}, \bar{\mu})$. The expressions for SCQP and SQCQP are almost identical, with the slight difference that $\bar{\mu}$ is substituted by μ . All three methods linearize the equality constraints $g(w)$ without keeping any curvature information, so none of the Hessians depends on the corresponding multiplier λ . An important property is that for $z = (w, \lambda, \mu) = (\bar{w}, \bar{\lambda}, \bar{\mu}) = \bar{z}$, all three Hessians are identical, i.e.,

$$B_{\text{SCP}}(\bar{w}, \bar{\mu}; \bar{w}) = B_{\text{SCQP}}(\bar{w}, \bar{\mu}) = B_{\text{SQCQP}}(\bar{\mu}; \bar{w}) =: B(\bar{z}), \quad (30)$$

which follows directly from $F_i^{\text{lin}}(\bar{w}; \bar{w}) = F_i(\bar{w})$. Finally, we define $B_* := B(z^*)$, $\Gamma_* := \Gamma(w^*) \in \mathbb{R}^{r \times n}$ and Z as a basis of the null space of Γ_* . The reduced Hessian approximation is denoted by $\tilde{B}_* := Z^\top B_* Z$. One way to obtain Z would be via the QR factorization

$$\Gamma_*^\top = Q\bar{R} = \begin{bmatrix} Y & Z \end{bmatrix} \begin{bmatrix} R \\ 0 \end{bmatrix}, \quad (31)$$

where $Q \in \mathbb{R}^{n \times n}$ is orthogonal, $Q^\top Q = I$, with block components $Y \in \mathbb{R}^{n \times r}$ and $Z \in \mathbb{R}^{n \times (n-r)}$. The non-zero block of $\bar{R} \in \mathbb{R}^{n \times r}$ is the upper triangular $R \in \mathbb{R}^{r \times r}$. Note that $\Gamma_* Z = 0$ and $\Gamma_* Y = R^\top$, which is invertible if Γ_* has full rank (LICQ holds). We can now state the following lemma on stability of the active set. This has in fact been proven by Robinson for a more general class of iterative methods for solving NLPs [32]. We thus refrain from giving a proof here. The reader interested in a proof for our specific case and notation is referred to [28].

Lemma 4.2. *Assume LICQ and strict complementarity hold at $z^* \in \mathcal{Z}^*$, and that $\tilde{B}_* \succ 0$. Then, for \bar{z} sufficiently close to z^* , the active sets of problems (16), (21) and (22) are stable and correspond to \mathcal{A}^* , the active set of the original NLP (14) at z^* .*

Lemma 4.2 implies that we can disregard the inactive inequality constraints if we are close enough to z^* . This leads us to the definition of $y := (w, \mu_{\mathcal{A}^*}, \lambda) = (w, \gamma)$, where we have dropped the inactive multipliers $\mu_{\mathcal{I}^*}$. Note that $\mu_{\mathcal{I}^*} = 0$. If we furthermore treat the active inequality constraints as equality constraints, the KKT

conditions of the resulting NLP and the corresponding subproblems define the following residual maps:

$$R(y) := \begin{bmatrix} \nabla f_0(w) + \nabla G(w)\gamma \\ G(w) \end{bmatrix}, \quad (32)$$

$$R_{\text{SCP}}(y; \bar{y}) := \begin{bmatrix} \nabla f_0^{\text{SCP}}(w; \bar{w}) + \nabla G_{\text{SCP}}(w; \bar{w})\gamma \\ G_{\text{SCP}}(w; \bar{w}) \end{bmatrix}, \quad (33)$$

$$R_{\text{SCQP}}(y; \bar{y}) := \begin{bmatrix} \nabla f_0^{\text{SCQP}}(w; \bar{w}) + \nabla G_{\text{SCQP}}(w; \bar{w})\gamma \\ G_{\text{SCQP}}(w; \bar{w}) \end{bmatrix}, \quad (34)$$

$$R_{\text{SQCQP}}(y; \bar{y}) := \begin{bmatrix} \nabla f_0^{\text{quad}}(w; \bar{w}) + \nabla G_{\text{SQCQP}}(w; \bar{w})\gamma \\ G_{\text{SQCQP}}(w; \bar{w}) \end{bmatrix}. \quad (35)$$

Since SCQP approximates (14) by a QP, its residual map is affine in y and can be written as

$$R_{\text{SCQP}}(y; \bar{y}) = \begin{bmatrix} B_{\text{SCQP}}(\bar{w}, \bar{\mu}) & \nabla G(\bar{w}) \\ \nabla G(\bar{w})^\top & 0 \end{bmatrix} \begin{bmatrix} w - \bar{w} \\ \gamma \end{bmatrix} + \begin{bmatrix} \nabla f_0(\bar{w}) \\ G(\bar{w}) \end{bmatrix}. \quad (36)$$

For SCP and SQCQP this map is generally nonlinear. All KKT points y^* of the reduced problem are characterized by $R(y^*) = 0$ and each method iterates by solving the root-finding problem $R_M(y_{k+1}; y_k) = 0$, for $M = \text{SCP}, \text{SCQP}, \text{SQCQP}$, until a stationary point is found. Before taking a closer look at these iterations, we state a lemma on the residual maps, which is central to our later theorem on the local convergence rates.

Lemma 4.3. *Consider the just defined residual maps. The residual maps of SCP, SCQP and SQCQP are identical up to second order,*

$$R_{\text{SCP}}(y; \bar{y}) = R_{\text{SCQP}}(y; \bar{y}) + \mathcal{O}(\|y - \bar{y}\|^2) \quad (37)$$

$$R_{\text{SQCQP}}(y; \bar{y}) = R_{\text{SCQP}}(y; \bar{y}) + \mathcal{O}(\|y - \bar{y}\|^2), \quad (38)$$

but match the residual map of the original problem (14) only up to first order,

$$R_{\text{SCQP}}(y; \bar{y}) = R(y) + \mathcal{O}(\|y - \bar{y}\|). \quad (39)$$

Proof. The proof will go by showing identity of the residuals and their partial derivatives at $y = \bar{y}$. From (26) we already know that $\nabla G_{\text{SCP}}(\bar{w}; \bar{w}) = \nabla G_{\text{SCQP}}(\bar{w}; \bar{w}) = \nabla G_{\text{SQCQP}}(\bar{w}; \bar{w}) = \Gamma(\bar{w})^\top$. From $g^{\text{lin}}(\bar{w}; \bar{w}) = g(\bar{w})$ and $f_i^{\text{SCP}}(\bar{w}; \bar{w}) = f_i^{\text{quad}}(\bar{w}; \bar{w}) = f_i^{\text{lin}}(\bar{w}; \bar{w}) = f(\bar{w})$, we have that $G_{\text{SCP}}(\bar{w}; \bar{w}) = G_{\text{SCQP}}(\bar{w}; \bar{w}) = G_{\text{SQCQP}}(\bar{w}; \bar{w}) = G(\bar{w})$. It follows that $R_{\text{SCP}}(\bar{y}; \bar{y}) = R_{\text{SCQP}}(\bar{y}; \bar{y}) = R_{\text{SQCQP}}(\bar{y}; \bar{y}) = R(\bar{y})$. The partial derivative w.r.t. w of the first row of each residual map is exactly the Hessian of the Lagrangian of the corresponding subproblem. They are given in (27) to (29) and from (30) we know them to be identical at $y = \bar{y}$. In consequence,

$$\left. \frac{\partial R_{\text{SCP}}(y; \bar{y})}{\partial y} \right|_{y=\bar{y}} = \left. \frac{\partial R_{\text{SCQP}}(y; \bar{y})}{\partial y} \right|_{y=\bar{y}} = \left. \frac{\partial R_{\text{SQCQP}}(y; \bar{y})}{\partial y} \right|_{y=\bar{y}} = \begin{bmatrix} B(\bar{y}) & \Gamma(\bar{w})^\top \\ \Gamma(\bar{w}) & 0 \end{bmatrix}. \quad (40)$$

□

Before moving to the next theorem, we consider the special case in which a local minimizer w^* is fully determined by the active constraints.

Definition 4.4 ([22]). We say a feasible point w of NLP (14) is fully determined by the active constraints if LICQ holds at w and $|\mathcal{A}(w)| = n - p$.

We introduce this special case for two reasons. First, if a local minimizer w^* is fully determined by the active constraints, the constraint Jacobian $\Gamma_* \in \mathbb{R}^{(|\mathcal{A}^*|+p) \times n}$ is a square matrix, since for the fully determined case $|\mathcal{A}^*| + p = n$. Due to LICQ, Γ_* is invertible and its null space contains only the zero vector. Therefore the null space basis Z would be empty, $Z \in \mathbb{R}^{0 \times 0}$, and in consequence the reduced Hessian approximation $\tilde{B} := Z^\top B Z$, which plays a central role in the following theorem, would be ill-defined, $\tilde{B}_* \in \mathbb{R}^{0 \times 0}$. Second, this case distinction is not only technical, but interesting in its own right. In a later theorem, we will actually obtain quadratic convergence for this special case. For now, we only focus on the case that w^* is *not* fully determined by the active constraints.

4.1. Linear Convergence Rate

We are now ready to state the main theorem of this paper. Similarly to the reduced Hessian approximation \tilde{B}_* , we define $\tilde{\Lambda}_* := Z^\top \nabla^2 \mathcal{L}(w^*, \mu^*, \lambda^*) Z$ for the true Hessian of the original problem (14), and $E_* := E_{\text{SCQP}}(w^*, \mu^*, \lambda^*)$ as well as $\tilde{E}_* := Z^\top E_* Z$ for the corresponding Hessian approximation error. Note that $\tilde{\Lambda}_* = \tilde{B}_* + \tilde{E}_*$.

Theorem 4.5. *Assume $z^* = (w^*, \mu^*, \lambda^*)$ is local minimizer of (14), at which LICQ and strict complementarity hold. Then z^* is a fixed point of SCP, SCQP and SQCQP. If w^* is not fully determined by the active constraints and $\tilde{B}_* \succ 0$ holds, then all three methods have the same asymptotic local linear contraction – or divergence – rate. This asymptotic contraction rate is given by the smallest α that fulfills the condition*

$$-\alpha \tilde{B}_* \preceq \tilde{E}_* \preceq \alpha \tilde{B}_*. \quad (41)$$

In consequence, if

$$\frac{1}{1+\alpha} \tilde{\Lambda}_* \preceq \tilde{B}_* \preceq \frac{1}{1-\alpha} \tilde{\Lambda}_* \quad (42)$$

holds for some $\alpha < 1$, the methods converge Q -linearly with contraction rate α and a necessary condition for local convergence is given by $\tilde{B}_* \succeq \frac{1}{2} \tilde{\Lambda}_*$. Finally, if $\tilde{\Lambda}_* \succ 0$, then $\tilde{B}_* \succ \frac{1}{2} \tilde{\Lambda}_*$ is sufficient for local linear convergence.

Proof. Stationarity follows from Lemma 4.1. From Lemma 4.2 we know that the active set \mathcal{A}^* of all three methods is stable close to z^* . It follows that the $\mu_{\mathcal{I}}$ will be zero, $\mu_{\mathcal{I}^*} = \mu_{\mathcal{I}^*}^* = 0$, for z sufficiently close to z^* . We can therefore disregard the inactive constraints and characterize the convergence behavior from the residuals maps $R_M(y; \bar{y})$ as defined in (33) to (35) for $M = \text{SCP}, \text{SCQP}, \text{SQCQP}$. As a consequence of Lemma 4.3 their partial derivatives are identical at $y = \bar{y} = y^*$ and given by

$$K_* := \frac{\partial R_M(y; \bar{y})}{\partial y} \Big|_{\substack{y=y^* \\ \bar{y}=y^*}} = \begin{bmatrix} B_* & \Gamma_*^\top \\ \Gamma_* & 0 \end{bmatrix}, \quad L_* := \frac{\partial R_M(y; \bar{y})}{\partial \bar{y}} \Big|_{\substack{y=y^* \\ \bar{y}=y^*}} = \begin{bmatrix} E_* & 0 \\ 0 & 0 \end{bmatrix} \quad (43)$$

for $M = \text{SCP}, \text{SCQP}, \text{SQCQP}$. We have $R_M(y^*; y^*) = 0$, and $R_M(\cdot; \cdot)$ is continuously differentiable with respect to both arguments, due to the assumptions made on the composing functions. Furthermore, K_* is invertible due to the KKT Lemma [30, Lem. 16.1], which holds because of LICQ and $\tilde{B}_* \succ 0$. We thus know from the implicit function theorem that there exists a neighborhood of y^* in which the continuously differentiable functions $y_M^{\text{sol}}(\bar{y})$ are uniquely defined by $R_M(y_M^{\text{sol}}(\bar{y}); \bar{y}) = 0$, with $M = \text{SCP}, \text{SCQP}, \text{SQCQP}$. For \bar{y} sufficiently close to y^* , their Jacobians are given by

$$\frac{dy_M^{\text{sol}}}{d\bar{y}}(\bar{y}) = - \left(\frac{\partial R_M}{\partial y}(y_M^{\text{sol}}(\bar{y}); \bar{y}) \right)^{-1} \frac{\partial R_M}{\partial \bar{y}}(y_M^{\text{sol}}(\bar{y}); \bar{y}). \quad (44)$$

At the fixed point y^* , with $y_M^{\text{sol}}(y^*) = y^*$, this evaluates to

$$\left. \frac{dy_M^{\text{sol}}(\bar{y})}{d\bar{y}} \right|_{\bar{y}=y^*} = - \underbrace{K_*^{-1} L_*}_{=: A_*} = - \begin{bmatrix} B_* & \Gamma_*^\top \\ \Gamma_* & 0 \end{bmatrix}^{-1} \begin{bmatrix} E_* & 0 \\ 0 & 0 \end{bmatrix} \quad (45)$$

for $M = \text{SCP}, \text{SCQP}, \text{SQCQP}$. Now consider the iteration $y_{k+1} = y_M^{\text{sol}}(y_k)$, with y_0 sufficiently close to y^* . Taylor expansion around y^* yields

$$y_{k+1} - y^* = -A_*(y_k - y^*) + \mathcal{O}(\|y_k - y^*\|^2). \quad (46)$$

As is a standard result of linear stability analysis of nonlinear systems, convergence of y_k to y^* is determined by the spectral radius $\rho(A_*)$ [31, Chap. 22]. If $0 < \rho(A_*) < 1$, the sequence converges linearly with asymptotic contraction rate $\rho(A_*)$. When $\rho(A_*) = 0$, it converges faster than linear. If $\rho(A_*) = 1$, we cannot decide about convergence by considering only $\rho(A_*)$, and if $\rho(A_*) > 1$, then y^* is an unstable stationary point.

We now derive an explicit expression of A_* , to be able to characterize its spectral radius $\rho(A_*)$. Note that to compute A_* – due to the special structure of L_* – we only need the first block column of

$$K_*^{-1} = \begin{bmatrix} D & C^\top \\ C & F \end{bmatrix}, \quad (47)$$

with $D \in \mathbb{R}^{n \times n}$, $C \in \mathbb{R}^{r \times n}$, $F \in \mathbb{R}^{r \times r}$. Noting that $K_* K_*^{-1} = I$, we can compute the first block column of K_*^{-1} by solving the linear system

$$K_* \begin{bmatrix} D \\ C \end{bmatrix} = \begin{bmatrix} B_* & \Gamma_*^\top \\ \Gamma_* & 0 \end{bmatrix} \begin{bmatrix} D \\ C \end{bmatrix} = \begin{bmatrix} I \\ 0 \end{bmatrix} \quad (48)$$

via the null space method [30, Chap. 16.2]. We split D into two components,

$$D = Y D_Y + Z D_Z \quad (49)$$

with $D_Y \in \mathbb{R}^r$, $D_Z \in \mathbb{R}^{n-r}$, $Z \in \mathbb{R}^{n \times (n-r)}$ an orthonormal basis of the null space of Γ_* and $Y \in \mathbb{R}^{n \times r}$ such that $[Y \ Z]^\top [Y \ Z] = I$. Y and Z might be obtained via QR factorization of Γ_*^\top , see (31). From the second block row of (48) we have

$$\Gamma_* Y D_Y + \Gamma_* Z D_Z = 0 \quad \Leftrightarrow \quad D_Y = 0 \quad (50)$$

since $\Gamma_* Z = 0$ and $\Gamma_* Y \in \mathbb{R}^{m \times m}$ has full rank due to LICQ. Left-multiplying the first block row of (48) by Z^\top we get

$$Z^\top B_* D + Z^\top \Gamma_*^\top C = Z^\top. \quad (51)$$

Substituting $D = Z D_Z$ and noting that $Z^\top \Gamma_*^\top = 0$ and that $Z^\top B_* Z \succ 0$ is invertible by assumption, $D_Z = (Z^\top B_* Z)^{-1} Z^\top$ follows, and therefore

$$D = Z D_Z = Z (Z^\top B_* Z)^{-1} Z^\top. \quad (52)$$

Similarly, left-multiplying the first block row of (48) by Y^\top and substituting the above result for D yields

$$C = (\Gamma_* Y)^{-\top} Y^\top (I - B_* Z (Z^\top B_* Z)^{-1} Z^\top). \quad (53)$$

Finally, we can explicitly write A_* as

$$A_* := K_*^{-1} L_* = \begin{bmatrix} D E_* & 0 \\ C E_* & 0 \end{bmatrix}. \quad (54)$$

The next step is to find the spectral radius $\rho(A_*)$. Due to the zero blocks on the right-hand side, the non-zero eigenvalues of A_* coincide with those of its upper left block,

$$\hat{A}_* := DE_* = Z(Z^\top B_* Z)^{-1} Z^\top E_*. \quad (55)$$

Now \hat{A}_* is similar to

$$\hat{A}'_* := \begin{bmatrix} Y^\top \\ Z^\top \end{bmatrix} \hat{A}_* \begin{bmatrix} Y & Z \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ \tilde{B}_*^{-1} Z^\top E_* Y & \tilde{B}_*^{-1} \tilde{E}_* \end{bmatrix}, \quad (56)$$

where we used $Y^\top Z = 0$, $Z^\top Z = I$, $\tilde{B}_* = Z^\top B_* Z$ and $\tilde{E}_* = Z^\top E_* Z$. It follows that \hat{A}_* and \hat{A}'_* have the same eigenvalues. The non-zero eigenvalues of \hat{A}'_* are given by the eigenvalues of its lower right block $\tilde{A}'_* := \tilde{B}_*^{-1} \tilde{E}_*$. Due to $\tilde{B}_* \succ 0$, the symmetric positive definite square root $\tilde{B}_*^{\frac{1}{2}}$ exists and is uniquely defined. Therefore

$$\tilde{A}'_* := \tilde{B}_*^{\frac{1}{2}} \tilde{A}_* \tilde{B}_*^{-\frac{1}{2}} = \tilde{B}_*^{-\frac{1}{2}} \tilde{E}_* \tilde{B}_*^{-\frac{1}{2}} \quad (57)$$

is similar to \tilde{A}'_* and they share the same eigenvalues. From the rightmost expression in (57) we can see that \tilde{A}'_* is symmetric, so its eigenvalues are real. Its spectral radius $\rho(\tilde{A}'_*) = \rho(\tilde{A}_*) = \rho(\hat{A}'_*) = \rho(\hat{A}_*) = \rho(A_*)$ is then given by the smallest α satisfying

$$-\alpha I \preceq \tilde{B}_*^{-\frac{1}{2}} \tilde{E}_* \tilde{B}_*^{-\frac{1}{2}} \preceq \alpha I. \quad (58)$$

These two LMI are equivalent to non-negativity of $\xi^\top (\alpha I \pm \tilde{B}_*^{-\frac{1}{2}} \tilde{E}_* \tilde{B}_*^{-\frac{1}{2}}) \xi$ for all $\xi \in \mathbb{R}^{n-r}$. Substituting $\zeta = \tilde{B}_*^{-\frac{1}{2}} \xi$, non-negativity of $\zeta^\top (\alpha \tilde{B}_* \pm \tilde{E}_*) \zeta$ for all $\zeta \in \mathbb{R}^{n-r}$ follows, which is equivalent to $-\alpha \tilde{B}_* \preceq \tilde{E}_* \preceq \alpha \tilde{B}_*$. Recalling that $\tilde{E}_* = \tilde{\Lambda}_* - \tilde{B}_*$, (42) is a direct consequence. Substituting $\alpha = 1$ into the left-hand side LMI of (42) we can see that $\tilde{B}_* \succeq \frac{1}{2} \tilde{\Lambda}_*$ is a necessary condition for convergence: if we had $\tilde{B}_* \prec \frac{1}{2} \tilde{\Lambda}_*$, the LMI could only hold for $\alpha > 1$. Finally, if $\tilde{\Lambda}_* \succ 0$, we know that the right-hand side LMI of (42) has to hold for some $\alpha < 1$, and the same goes for the left-hand side LMI if $\tilde{B}_* \succ \frac{1}{2} \tilde{\Lambda}_*$ holds. Therefore $\tilde{B}_* \succ \frac{1}{2} \tilde{\Lambda}_*$ is sufficient for linear convergence. \square

Example 4.6. To illustrate the local convergence behavior, we revisit the problem formulation introduced in Example 3.2. We initialize the three methods at $w_0 = 0$, $s_0 = 0$, and, in the case of SCQP, $\mu_0 = 1$. For the obtained iteration sequences we compute the empirical contraction rate as

$$\kappa_k = \frac{\|\omega_{k+1} - \omega_k\|_2}{\|\omega_k - \omega_{k-1}\|_2}, \quad (59)$$

where $\omega_k = (w_k, s_k)$. The theoretical asymptotic rate $\check{\alpha}(w^*)$ is computed as defined by the LMI in (41). The results are shown in Figure 6. Note how the empirical rates approach the theoretically predicted rate in the final iterations.

4.2. Analysis of the Fully Determined Case

In this subsection, we consider the case in which the solution is fully determined by the active constraints, as defined in Definition 4.4. As we will show, this is an interesting special case in which the KKT conditions are sufficient for local optimality and SCP, SCQP, and SQCQP converge quadratically. We motivate the analysis with the following example.

Example 4.7. We continue with the just introduced slack reformulation (23). This time we want to investigate the convergence behavior when varying the Huber parameter δ . Recalling that for $\delta \rightarrow 0$ the pseudo Huber penalty approaches the L_1 -norm, we also consider a variation of our example in which the residuals are penalized

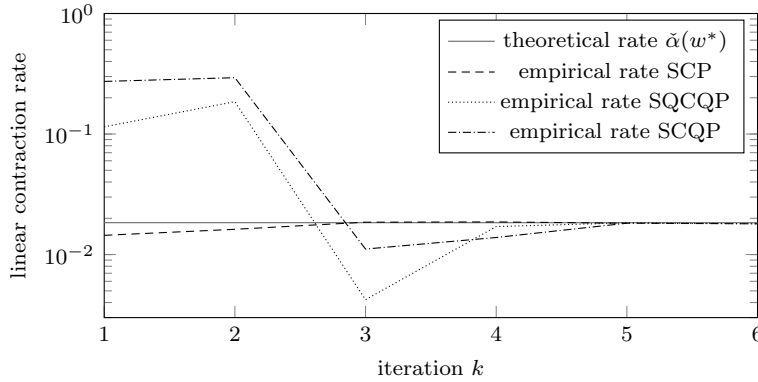


FIGURE 6. Convergence to the local minimum at $w^* \approx 0.1$. The observed contraction rate of the three methods is compared to the theoretically predicted rate $\tilde{\alpha}(w^*)$. During the final iterations, all three methods approach this rate. Also note how SCP has the fastest rate during the initial iterations, whereas SCQP is the slowest.

by the L_1 -norm. This leads us to the problem $\min_{w \in R} \|\eta - M(w)\|_1$, for which the smooth epigraph reformulation is

$$\min_{w, s} \sum_{i=1}^N s_i \quad \text{s.t.} \quad \begin{aligned} F_i(w) &\leq s_i, & i = 1, \dots, N, \\ -F_i(w) &\leq s_i, & i = 1, \dots, N. \end{aligned} \quad (60)$$

We use SCP to solve the above problem, as well as the problem given in (23) for values of $\delta \in [10^{-6}, 10^2]$. Note that applying SCP to (60) actually simplifies to SLP [19]. For each δ we compute the theoretical contraction rate. The results are visualized in Figure 7 on the left side. For $\delta \gg 1$ the contraction rate flatlines at $\tilde{\alpha} \approx 0.04$. This happens when δ is so large that all residuals are penalized quadratically, i.e., in a least-squares fashion. For $\delta \rightarrow 0$ something much more interesting happens: it seems that $\tilde{\alpha} \rightarrow 0$ as δ approaches 0, i.e., in the limit we would obtain convergence faster than linear. As in Example 4.6 we also compute the empirical contraction rate κ_k of SCP for a few values of δ and the L_1 -norm, initializing at $w_0 = 2$. The resulting contraction rates are compared to the theoretically predicted rate $\tilde{\alpha}$ in Figure 7 on the right side. What happens here is that as $\delta \rightarrow 0$, we approach the SLP iterations defined by the L_1 -norm. For a linear program, the solution – if it exists and is unique – always lies in a vertex of the feasible set. This corresponds to the fully determined case and therefore the convergence rate is quadratic, as we will see in the following.

Intuitively, the quadratic convergence can be explained as follows: If the active set \mathcal{A}^* is known, i.e., we know $G(w^*)$, and w^* is fully determined by the active constraints, i.e., $G(w^*)$ has n components and Γ_* is of full rank, then solving (14) simplifies to solving the feasibility problem $G(w) = 0$, which is independent of the objective function. Solving this nonlinear root-finding problem with the classical Newton's method would yield locally quadratic convergence [8]. Now SCQP linearizes all constraints, and thus in effect solves this special case with Newton's method, and should therefore converge quadratically. SCP and SQCQP keep some curvature in the constraints, so this reasoning is not fully applicable, but we will see that nonetheless they converge quadratically as well. A second consequence is the following proposition, which we might call the first order sufficient conditions (FOSC).

Proposition 4.8. *Suppose $z^* = (w^*, \lambda^*, \mu^*)$ is a KKT point of (14). If w^* is fully determined by the active constraints and LICQ and strict complementarity hold at z^* , then w^* is a strict local minimizer of (14), i.e., the KKT conditions are necessary and sufficient conditions for optimality.*

Proof. LICQ and strict complementarity hold at z^* , so the critical cone $\mathcal{C}(w^*, \mu^*)$ is given by the null space of the Jacobian of the extended equality constraints $G(w^*)$, i.e., $\mathcal{C}(w^*, \mu^*) = \text{null}(\Gamma_*)$. As w^* is fully determined

by the active constraints, Γ_* is an $n \times n$ matrix and due to LICQ has full rank. Thus, the null space of Γ_* contains only the zero vector, $\mathcal{C}(w^*, \mu^*) = \{0\}$. For any non-zero feasible direction $d \in \mathcal{F}(w^*)$, with tangent cone

$$\mathcal{F}(w^*) = \{d \in \mathbb{R}^n \mid \nabla g(w^*)^\top d = 0 \text{ and } \nabla f_i(w^*)^\top d \leq 0 \text{ for } i \in \mathcal{A}(w^*)\}, \quad (61)$$

there exists at least one index $j \in \mathcal{A}^*$ such that $d^\top \nabla f_j(w^*) < 0$. From $\nabla \mathcal{L}(w^*, \lambda^*, \mu^*) = 0$, we can conclude that

$$d^\top \nabla f_0(w^*) = -\underbrace{\mu_j^* d^\top \nabla f_j(w^*)}_{<0} - \sum_{i \in \mathcal{A}^* \setminus \{j\}} \underbrace{\mu_i^* d^\top \nabla f_i(w^*)}_{\leq 0} - \sum_{i=1}^d \underbrace{\lambda_i^* d^\top \nabla g_i(w^*)}_{=0} > 0. \quad (62)$$

Thus, any non-zero feasible direction d is an ascent direction, which implies that w^* is a strict local minimizer of (14). \square

Theorem 4.9. *Assume $z^* = (w^*, \mu^*, \lambda^*) = (w^*, \vartheta^*)$ is a KKT point of (14), at which LICQ and strict complementarity hold. If w^* is fully determined by the active constraints, then w^* is a local minimizer and SCP, SCQP and SQCQP converge Q-quadratically in the primal variable w , and R-quadratically in the dual variable $\vartheta = (\mu, \lambda)$, i.e., there are constants $c_1, c_2 \in \mathbb{R}_+$ such that*

$$\|w_{k+1} - w^*\| \leq c_1 \|w_k - w^*\|^2 \quad (63)$$

$$\text{while } \|\vartheta_{k+1} - \vartheta^*\| \leq c_2 \|w_k - w^*\|. \quad (64)$$

Proof. Proposition 4.8 implies that w^* is a local minimizer of (14). For the convergence analysis, we follow the proof of Theorem 4.5: We disregard the inactive multipliers $\mu_{\mathcal{I}^*} = 0$ and focus the analysis on $y = (w, \mu_{\mathcal{A}^*}, \lambda) = (w, \gamma)$. Consider again the iterations defined by the solution operators $y_{k+1} = y_M^{\text{sol}}(y_k)$, for $M = \text{SCP}, \text{SCQP}, \text{SQCQP}$, and their Jacobian at y^* , as given in (45) as $-A^* = -K_*^{-1}L_*$. For the fully determined case, Γ_* is a square matrix and due to LICQ invertible. We can therefore obtain the first block column of K_*^{-1} as $D = 0$ and $C = \Gamma_*^{-\top}$, cf. (47), and subsequently

$$A_* = \begin{bmatrix} 0 & 0 \\ \Gamma_*^{-\top} E_* & 0 \end{bmatrix}. \quad (65)$$

The only non-zero block of A_* is in its lower left, from which $\rho(A_*) = 0$ and therefore at least superlinear convergence follows. Now consider again the Taylor expanded iteration map (46),

$$y_{k+1} - y^* = -A_*(y_k - y^*) + \mathcal{O}(\|y_k - y^*\|^2) = -\begin{bmatrix} 0 & 0 \\ \Gamma_*^{-\top} E_* & 0 \end{bmatrix} (y_k - y^*) + \mathcal{O}(\|y_k - y^*\|^2). \quad (66)$$

We now make two observations: (a) both SCP and SQCQP are multiplier-free methods, i.e., their solution operators depend only on w_k , but not on γ_k ; (b) while SCQP is in general not multiplier-free, for the fully determined case we can explicitly obtain the primal part of its solution operator as $w_{k+1} = w_k - \Gamma(w_k)^{-1}G(w_k)$ from the second block row of (36), which does not depend on γ_k . Combining these observations with the first block row of (66), we can write $w_{k+1} - w^* = \mathcal{O}(\|w_k - w^*\|^2)$, from which Q-quadratic convergence of the primal variable w follows. The second block row is given as

$$\gamma_{k+1} - \gamma^* = -\Gamma_*^{-\top} E_*(w_k - w^*) + \mathcal{O}(\|y_k - y^*\|^2). \quad (67)$$

For y_k sufficiently close to y^* , we can find a constant $c_2 \in \mathbb{R}_+$ such that $\|\gamma_{k+1} - \gamma^*\| \leq c_2 \|w_k - w^*\|$. Since the sequence $\|w_k - w^*\|$ converges Q-quadratically to zero, R-quadratic convergence of the sequence γ_k to γ^* follows. \square




FIGURE 7. Left: Theoretical linear contraction rate for varying the pseudo Huber loss parameter δ . Right: Empirical convergence rate of SCP for several δ and the limit case of the L_1 -norm (which corresponds to $\delta = 0$). The dotted lines indicate the theoretically predicted rates. For the L_1 -norm this rate is 0, and therefore not visualizable on a log-scale.

5. MIRRORING AND DESIRABLE DIVERGENCE

We now turn our attention to an interesting phenomenon that arises in the context of parameter estimation. As we have seen in the previous section, the discussed methods do not always converge locally. At first glance this might seem like a weakness. In this section we argue why – in the context of maximum likelihood estimation – this is actually a strength. This idea has first been presented by Bock in the context of L_2 estimation [3] and was later extended to L_1 estimation [4]. In [11] we generalized to unconstrained NLP with general convex symmetric penalty functions, and now generalize this further to problems with equality constraints.

Imagine we have a model function $m(x; \bar{w})$ that maps an input $x \in \mathbb{R}^{n_x}$ to an output $m(x; \bar{w}) \in \mathbb{R}$, depending on a parameter $\bar{w} \in \mathbb{R}^n$. Of this output we only have noisy measurements $\eta_i = m(x_i; \bar{w}) + \nu_i$, where $\nu_i \in \mathbb{R}$ is some form of noise. Our wish is now to identify this parameter from input-output pairs (x_i, η_i) , $i = 1, \dots, N$. In the framework of maximum-likelihood estimation this leads to an objective function in the form of $f(w) = \sum_{i=1}^N \varphi(\eta_i - m(x_i; w))$ with a cost function $\varphi: \mathbb{R} \rightarrow \mathbb{R}$ determined by the assumed probability distribution of ν_i . More generally we may phrase this as

$$\min_w \phi_0(\eta - M(w)) \quad \text{s.t.} \quad g(w) = 0. \quad (68)$$

Here $M(w)$ is the vertical concatenation of the $m(x_i; w)$ and correspondingly for η . The constraint $g(w)$ encodes further prior knowledge on the parameter, that might be derived from physical insight into the problem. As should be straightforward to see, (68) includes scenarios more general than the one described, e.g., it also includes scenarios with multi-output models or models with time dependency. Central to our results here is the observation that most commonly used noise distributions are radially symmetric. That is, if a measurement is perturbed by some noise realization ν , its negative version $-\nu$ should be equally likely. In consequence, the cost function defined by the corresponding distribution is symmetric, $\phi_0(y) = \phi_0(-y)$. With these properties in mind we define the “mirror problem”.

Definition 5.1 (Mirror problem). Consider a point $w' \in \mathbb{R}^n$. The corresponding residual error is $e(w') := \eta - M(w')$, which we can rephrase as $\eta = M(w') + e(w')$. We obtain “mirrored” measurements $\check{\eta}$ by using this error to perturb the model into the opposite direction, $\check{\eta}(w') := M(w') - e(w') = 2M(w') - \eta$. These mirrored measurements then define the mirror problem at w' as

$$\min_w \phi_0(\check{\eta}(w') - M(w)) \quad \text{s.t.} \quad g(w) = 0. \quad (69)$$

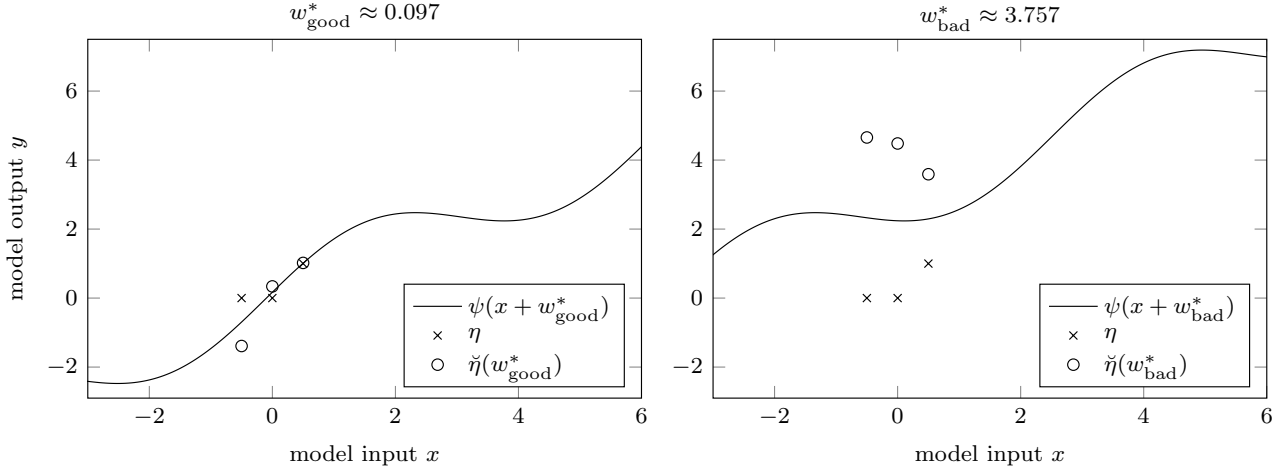


FIGURE 8. Illustration of the mirror problem for two local minima, w_{good}^* and w_{bad}^* . The mirrored measurements $\check{\eta}$ are obtained by mirroring the original measurements η vertically at the model function.

Example 5.2. We return to our example as defined in (5). For the two local minima, w_{good}^* and w_{bad}^* , we compute the mirrored measurements $\check{\eta}(w_{\text{good}}^*)$ resp. $\check{\eta}(w_{\text{bad}}^*)$ as explained in Definition 5.1. This is visualized in Figure 8 alongside the measurement model for the respective minimizer. In this context the model is $m(x_i; w) := \psi(x_i + w)$ as defined in (4). Note that varying the model parameter w corresponds to shifting ψ horizontally.

Lemma 5.3. Assume $z^* = (w^*, \lambda^*)$ is a KKT point of (68). Then $\check{z} = (w^*, -\lambda^*)$ is a KKT point of its mirror problem (69) at w^* and vice versa.

Proof. We start by defining the Lagrangian of (68) as

$$\mathcal{L}(z) = \mathcal{L}(w, \lambda) = \phi_0((\eta - M(w))) + \lambda^\top g(w), \quad (70)$$

and correspondingly for its mirror problem (69). Now consider the KKT conditions of the mirror problem:

$$J(w)^\top \nabla \phi_0(\check{\eta}(w^*) - M(w)) + \nabla g(w) \lambda = 0, \quad (71a)$$

$$g(w) = 0. \quad (71b)$$

In these, first substitute $\check{z} = (w^*, -\lambda^*)$ for z . Then note that due to symmetry $\nabla \phi_0(v) = -\nabla \phi_0(-v)$, and therefore $\nabla \phi_0(\check{\eta}(w^*) - M(w)) = -\nabla \phi_0(M(w) - \check{\eta}(w^*))$. Finally substitute $\check{\eta}(w^*) = 2M(w^*) - \eta$, to obtain

$$-(J(w^*)^\top \nabla \phi_0(\eta - M(w^*)) + \nabla g(w^*) \lambda^*) = 0, \quad (72a)$$

$$g(w^*) = 0. \quad (72b)$$

These are exactly the KKT conditions of (68) at z^* , which hold by assumption. Therefore (71) holds at \check{z} and \check{z} is a KKT point of the mirror problem. For the reverse direction note that the mirror problem of the mirror problem is the original problem. \square

Lemma 5.4. Assume z^* is a KKT point of (68), and \check{z} the corresponding KKT point of the mirror problem at w^* as obtained in Lemma 5.3. Define $B_* := B(z^*)$ and $E_* := E(z^*)$. Then the Hessians of the Lagrangians

of the original resp. the mirror problem are given by

$$\Lambda_* := \nabla^2 \mathcal{L}(z^*) = B_* + E_* \quad \text{and} \quad \check{\Lambda}_* := \nabla^2 \check{\mathcal{L}}(\check{z}) = B_* - E_*. \quad (73)$$

Proof. The Hessians of the Lagrangians of (68) and (69) are

$$\nabla^2 \mathcal{L}(z) = \overbrace{J_0(w)^\top \nabla^2 \phi_0(\eta - M(w)) J_0(w)}{=:B(z)} + \overbrace{\sum_{i=1}^N \nabla^2 M_i(w) \frac{\partial \phi_0}{\partial v_i} (\eta - M(w)) + \sum_{i=1}^p \lambda_i \nabla^2 g_i(w)}{=:E(z)} \quad (74)$$

$$\text{and} \quad \nabla^2 \check{\mathcal{L}}(z) = J_0(w)^\top \nabla^2 \phi_0(\check{\eta} - M(w)) J_0(w) + \sum_{i=1}^N \nabla^2 M_i(w) \frac{\partial \phi_0}{\partial v_i} (\check{\eta} - M(w)) + \sum_{i=1}^p \lambda_i \nabla^2 g_i(w), \quad (75)$$

where we dropped the argument of $\check{\eta}(w^*)$ for readability. The expression for $\Lambda_* = \nabla^2 \mathcal{L}(z^*)$ follows directly from (74). For $\check{\Lambda}_* = \nabla^2 \check{\mathcal{L}}(\check{z})$ we recall that $\check{z} = (w^*, -\lambda^*)$, $\nabla \phi_0(-v) = -\nabla \phi_0(v)$, $\nabla^2 \phi_0(v) = \nabla^2 \phi_0(-v)$, and $\check{\eta} = 2M(w^*) - \eta$. Substituting all of this in (75) we see that

$$\check{\Lambda}_* = J_0(w^*)^\top \nabla^2 \phi_0(\eta - M(w^*)) J_0(w^*) - \sum_{i=1}^N \nabla^2 M_i(w^*) \frac{\partial \phi_0}{\partial v_i} (\eta - M(w^*)) - \sum_{i=1}^p \lambda_i^* \nabla^2 g_i(w^*) = B_* - E_*. \quad (76)$$

□

We are now ready to state the main theorem of this section. But before we do so, we would like the reader to note that the estimation problem (68) has no inequality constraints. This leads to the fact that in this case both SCQP and SQCQP simplify to CGGN. As in the previous section, a tilde is used to denote the reduced Hessians, e.g., $\tilde{B}_* = Z^\top B_* Z$.

Theorem 5.5 (Divergence from undesirable local minima). *Assume z^* is a local minimizer of (68), at which LICQ, strict complementarity, $\tilde{\Lambda}_* \succeq 0$ and $\tilde{B}_* \succ 0$ hold. Then z^* is a KKT point of (68) and \check{z} as defined in Lemma 5.3 is a KKT point of its mirror problem. Both are therefore stationary points of SCP and CGGN applied to the respective problem. But if these methods applied to (68) diverge from z^* , then \check{z} is not a local minimizer of the mirror problem.*

Proof. z^* is a KKT point by definition, and that \check{z} is a KKT point follows from Lemma 5.3. Stationarity of the methods at z^* resp. \check{z} follows from Lemma 4.1. From Theorem 4.5 we know that local divergence means that $-\alpha \tilde{B}_* \preceq \tilde{E}_* \preceq \alpha \tilde{B}_*$ only holds for some $\alpha > 1$. Therefore there exists some $\alpha' \in (1, \alpha)$ such that one of the two LMI is violated. The left-hand side can be reformulated as $\alpha \tilde{B}_* + \tilde{E}_* \succeq 0$, and due to $\tilde{B}_* + \tilde{E}_* = \tilde{\Lambda}_* \succeq 0$ and $\tilde{B}_* \succ 0$ we see that this can never be violated for $\alpha > 1$. It follows that the right-hand side must be violated. This means that there exists some $p \in \mathbb{R}^{n_Z} \setminus \{0\}$ – with n_Z the dimension of the null space – such that

$$0 > p^\top (\alpha' \tilde{B}_* - \tilde{E}_*) p = p^\top ((\alpha' - 1) \tilde{B}_* + \tilde{B}_* - \tilde{E}_*) p = \underbrace{(\alpha' - 1) p^\top \tilde{B}_* p}_{>0} + \underbrace{p^\top \tilde{B}_* p}_{>0} + p^\top \check{\Lambda}_* p > p^\top \check{\Lambda}_* p, \quad (77)$$

where we used $\check{\Lambda}_* = B_* - E_*$ from Lemma 5.4 and introduced its reduced version $\check{\tilde{\Lambda}} := Z^\top \check{\Lambda}_* Z$. Therefore $\check{\tilde{\Lambda}} \not\succeq 0$. Since $\check{\tilde{\Lambda}}$ is the reduced Hessian of the mirror problem evaluated at \check{z} , it follows that the second order necessary conditions do not hold and \check{z} cannot be a minimizer of the mirror problem. □

Example 5.6. Continuing directly on Example 5.2, we now investigate the mirror problem as defined in (69). From Example 2.3 we recall that for the asymptotic linear contraction rates we have $\check{\alpha}(w_{\text{good}}^*) \approx 0.02$ and

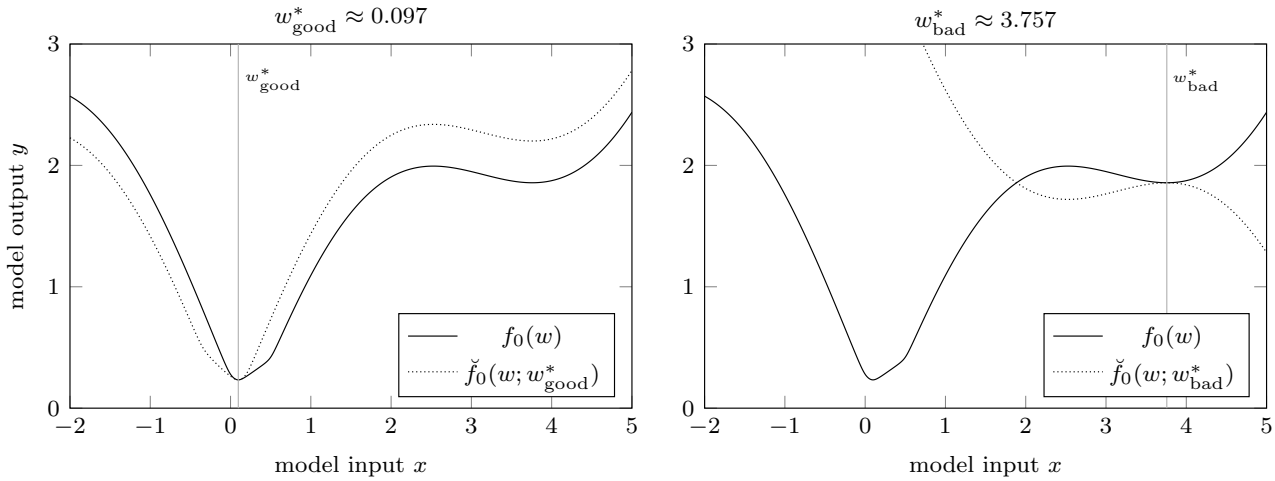


FIGURE 9. Comparison of the original objective function $f_0(w)$ to the objective functions $\check{f}_0(w; w^*)$ of the mirror problems for the two local minima w_{good}^* and w_{bad}^* . The bad local minimum at w_{bad}^* turns into a maximum of the corresponding mirror problem.

$\check{\alpha}(w_{\text{bad}}^*) \approx 3200 \gg 1$. In Figure 9, we visualize the objective functions of the mirror problems at both local minima and see how w_{good}^* remains a clear minimizer for the mirror problem, whereas the local minimum at w_{bad}^* is transformed into a local maximum of the mirror problem.

6. CONCLUSIONS

We provided an overview of methods that exploit convex-over-nonlinear substructure in optimization problems. These were divided into three sectors: smooth unconstrained NLP, smooth constrained NLP, and optimization problems with non-smooth convex substructures. For all three sectors we discussed methods for the most general problem formulation as well as special cases, of which the simplest is nonlinear least-squares. A special focus lay on methods for smooth constrained problems and for these we provided an analysis of local convergence. We proved that under mild assumptions SCP, SCQP and SQCQP have the same local linear contraction rate. For a fully determined solution, they even converge quadratically. For a simple numerical example we have seen that SCP has better global convergence behavior, whereas SCQP has cheaper iterations. SQCQP can be seen to provide a trade-off between these two cases. It would be interesting to further validate this on a broad set of example problems and with a rigorous comparison of the computational cost. As we have only considered the full step methods here, an analysis of globalized versions might also yield important results. Furthermore, we have shown that the generalized CGN and SCP methods locally converge to a local minimizer for equality constrained estimation problems with symmetric convex penalties if and only if this minimizer is stable under a mirroring operation. Thus, the methods are only attracted by those local minima that satisfy the desirable property of stability under mirroring.

The authors would like to thank Matilde Gargiani, Joris Gillis, James Martens, Dominikus Noll and Stephen Wright for inspiring discussions.

REFERENCES

- [1] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl. CasADi: a software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*, 2018.

- [2] P. Apkarian, D. Noll, and A. Rondepierre. Mixed H_2/H_∞ control via nonsmooth optimization. *SIAM Journal on Control and Optimization*, 47(3):1516–1546, 2008.
- [3] H. Bock. *Randwertproblemmethoden zur Parameteridentifizierung in Systemen nichtlinearer Differentialgleichungen*, volume 183 of *Bonner Mathematische Schriften*. Universität Bonn, Bonn, 1987.
- [4] H. Bock, E. Kostina, and J. Schlöder. Numerical methods for parameter estimation in nonlinear differential algebraic equations. *GAMM Mitteilungen*, 30/2:376–408, 2007.
- [5] H. G. Bock. Recent advances in parameter identification techniques for ODE. In *Numerical Treatment of Inverse Problems in Differential and Integral Equations*, pages 95–121. Birkhäuser, 1983.
- [6] S. Boyd and L. Vandenberghe. *Convex Optimization*. University Press, Cambridge, 2004.
- [7] H. Chen and F. Allgöwer. A quasi-infinite horizon nonlinear model predictive control scheme with guaranteed stability. *Journal of Process Control*, 8(5-6):475–485, 1998.
- [8] P. Deufhard. *Newton methods for nonlinear problems: affine invariance and adaptive algorithms*, volume 35. Springer, 2011.
- [9] M. Diehl, H. Bock, and E. Kostina. An approximation technique for robust nonlinear optimization. *Mathematical Programming*, 107:213–230, 2006.
- [10] M. Diehl, F. Jarre, and C. Vogelbusch. Loss of superlinear convergence for an SQP-type method with conic constraints. *SIAM Journal on Optimization*, 16(4):1201–1210, 2006.
- [11] M. Diehl and F. Messerer. Local convergence of generalized Gauss-Newton and sequential convex programming. In *Proceedings of the IEEE Conference on Decision and Control (CDC)*, 2019.
- [12] B. Fares, D. Noll, and P. Apkarian. Robust Control via Sequential Semidefinite Programming. *SIAM Journal on Control and Optimization*, 40(6):1791–1820, 2002.
- [13] H. J. Ferreau, C. Kirches, A. Potschka, H. G. Bock, and M. Diehl. qpOASES: a parametric active-set algorithm for quadratic programming. *Mathematical Programming Computation*, 6(4):327–363, 2014.
- [14] A. Fischer. Local behavior of an iterative framework for generalized equations with nonisolated solutions. *Mathematical Programming*, 94(1):91–124, 2002.
- [15] R. Freund, F. Jarre, and C. Vogelbusch. Nonlinear semidefinite programming: sensitivity, convergence, and an application in passive reduced-order modeling. *Math. Program. Ser. B.*, 109:581–611, 2007.
- [16] M. Fukushima, Z.-Q. Luo, and P. Tseng. A sequential quadratically constrained quadratic programming method for differentiable convex minimization. *SIAM Journal on Optimization*, 13(4):1098–1119, 2003.
- [17] C. F. Gauss. *Theoria motus corporum coelestium in sectionibus conicis solem ambientium*, volume 7. Perthes et Besser, 1809.
- [18] D. Geweth, A. Zanelli, G. Frison, U. Vollmer, and M. Diehl. Field oriented economic model predictive control for permanent magnet synchronous motors. In *Proceedings of the IFAC World Congress*, Berlin, Germany, July 2020.
- [19] R. E. Griffith and R. Stewart. A nonlinear programming technique for the optimization of continuous processing systems. *Management science*, 7(4):379–392, 1961.
- [20] Gurobi Optimization. <http://www.gurobi.com>.
- [21] A. F. Izmailov and M. V. Solodov. Newton-type methods: a broader view. *Journal of Optim*, 164(2):577–620, 2015.
- [22] T. Kim and S. J. Wright. An S 11 LP-active set approach for feasibility restoration in power systems. *Optimization and Engineering*, 17(2):385–419, 2016.
- [23] A. S. Lewis and S. J. Wright. A proximal method for composite minimization. *Mathematical Programming*, 158(1–2):501–546, 2016.
- [24] W. C. Li and L. T. Biegler. Multistep, Newton-type control strategies for constrained nonlinear processes. *Chem. Eng. Res. Des.*, 67:562–577, 1989.
- [25] J. Martens. Deep learning via hessian-free optimization. In *27th International Conference on Machine Learning*, pages 35–742, 2010.
- [26] J. Martens. *Second-Order Optimization for Neural Networks*. PhD thesis, University of Toronto, 2016.
- [27] J. Martens and I. Sutskever. Learning recurrent neural networks with hessian-free optimization. In *28th International Conference on Machine Learning*, 2011.
- [28] F. Messerer and M. Diehl. Determining the exact local convergence rate of sequential convex programming. In *Proceedings of the European Control Conference (ECC)*, 2020.
- [29] Z. Nagy and R. Braatz. Worst-case and distributional robustness analysis of finite-time control trajectories for nonlinear distributed parameter systems. *IEEE Transactions on Control Systems Technology*, 11:494–504, 2003.
- [30] J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering. Springer, 2 edition, 2006.
- [31] A. M. Ostrowski. *Solutions of Equations in Euclidean and Banach Spaces*. Academic Press, New York and London, 1973.
- [32] S. Robinson. Perturbed Kuhn-Tucker points and rates of convergence for a class of nonlinear programming algorithms. *Mathematical Programming*, 7:1–16, 1974.
- [33] N. N. Schraudolph. Fast curvature matrix-vector products for second-order gradient descent. *Neural Computation*, 14(7):1723–1738, 2002.
- [34] A. Sideris and J. Bodrow. An efficient sequential linear quadratic algorithm for solving unconstrained nonlinear optimal control problems. *IEEE Transactions on Automatic Control*, 50(12):2043–2047, 2005.

- [35] Q. Tran-Dinh and M. Diehl. Local convergence of sequential convex programming for nonconvex optimization. In M. Diehl, F. Glineur, E. Jarlebring, and W. Michiels, editors, *Recent advances in optimization and its application in engineering*, pages 93–103. Springer-Verlag, 2010.
- [36] Q. Tran-Dinh, S. Gumussoy, W. Michiels, and M. Diehl. Combining convex–concave decompositions and linearization approaches for solving BMIs, with application to static output feedback. *IEEE Trans. Automat. Control*, 57(6), 2012.
- [37] Q. Tran-Dinh, C. Savorgnan, and M. Diehl. Adjoint-based predictor-corrector sequential convex programming for parametric nonlinear optimization. *SIAM J. Optimization*, 22(4):1258–1284, 2012.
- [38] Q. Tran-Dinh, C. Savorgnan, and M. Diehl. Real-time sequential convex programming for optimal control applications. In H. G. Bock, H. X. Phu, R. Rannacher, and J. P. Schlöder, editors, *Modeling, Simulation and Optimization of Complex Processes*. Springer, 2012.
- [39] N. van Duijkeren. *Online Motion Control in Virtual Corridors – for Fast Robotic Systems*. PhD thesis, KU Leuven, 2019.
- [40] R. Verschueren, N. van Duijkeren, R. Quirynen, and M. Diehl. Exploiting convexity in direct optimal control: a sequential convex quadratic programming method. In *Proceedings of the IEEE Conference on Decision and Control (CDC)*, 2016.
- [41] R. Verschueren, M. Zanon, R. Quirynen, and M. Diehl. A sparsity preserving convexification procedure for indefinite quadratic programs arising in direct optimal control. *SIAM Journal of Optimization*, 27(3):2085–2109, 2017.
- [42] C.-C. Wang, K. L. Tan, and C.-J. Lin. Newton methods for CNNs. *ACM Trans. Intell. Syst. Technol.*, 11(2), 2020.
- [43] H. Yamashita and H. Yabe. A survey of numerical methods for nonlinear semidefinite programming. *Journal of the Operations Research Society of Japan*, 58(1):24–60, 2015.
- [44] A. Zanelli, J. Kullick, H. Eldeeb, G. Frison, C. Hackl, and M. Diehl. Continuous control set nonlinear model predictive control of reluctance synchronous machines. *IEEE Transactions on Control Systems Technology*, 2021.