



HAL
open science

QPLayer: efficient differentiation of convex quadratic optimization

Antoine Bambade, Fabian Schramm, Adrien Taylor, Justin Carpentier

► **To cite this version:**

Antoine Bambade, Fabian Schramm, Adrien Taylor, Justin Carpentier. QPLayer: efficient differentiation of convex quadratic optimization. 2023. hal-04133055

HAL Id: hal-04133055

<https://inria.hal.science/hal-04133055>

Preprint submitted on 19 Jun 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Public Domain Mark 4.0 International License

QPLayer: efficient differentiation of convex quadratic optimization

Antoine Bambade*
Inria & ENS & ENPC
Paris

Fabian Schramm
Inria & ENS
Paris

Adrien Taylor
Inria & ENS
Paris

Justin Carpentier
Inria & ENS
Paris

June 19, 2023

Abstract

Optimization layers within neural network architectures have become increasingly popular for their ability to solve a wide range of machine learning tasks and to model domain-specific knowledge. However, designing optimization layers requires careful consideration as the underlying optimization problems might be infeasible during training. Motivated by applications in learning, control, and robotics, this work focuses on convex quadratic programming (QP) layers. The specific structure of this type of optimization layer can be efficiently exploited for faster computations while still allowing rich modeling capabilities. We leverage primal-dual augmented Lagrangian techniques for computing derivatives of both feasible and infeasible QPs. Not requiring feasibility allows, as a byproduct, for more flexibility in the QP to be learned. The effectiveness of our approach is demonstrated in a few standard learning experiments, obtaining three to ten times faster computations than alternative state-of-the-art methods while being more accurate and numerically robust. Along with these contributions, we provide an open-source C++ software package called QPLayer for efficiently differentiating convex QPs and which can be interfaced with modern learning frameworks.

1 Introduction

Incorporating differentiable optimization problems as layers within neural networks has recently become practical and effective for solving certain machine learning tasks, see, for instance [Geng et al., 2020, Amos and Kolter, 2017, Lee et al., 2019, Le Lidec et al., 2021, Donti et al., 2017, de Avila Belbute-Peres et al., 2018, Amos et al., 2018, Bounou et al., 2021]. Such layers allow capturing useful domain-specific knowledge or priors. Unlike conventional neural networks, where the output of each layer is provided by a simple (explicit) function of its input, the input of an optimization layer is the parameter of an optimization problem, and its output is a solution to this problem.

In this work, we focus on convex Quadratic Programming (QP) layers, a specific type of optimization layer that offers a rich modeling power [Amos and Kolter, 2017, Section 3.2]. A convex QP parameterized by θ is defined as follows

$$\begin{aligned} x^*(\theta) \in \operatorname{argmin}_{x \in \mathbb{R}^n} \left\{ f(x; \theta) := \frac{1}{2} x^\top H(\theta) x + x^\top g(\theta) \right\} \\ \text{s.t. } C(\theta) x \leq u(\theta), \end{aligned} \tag{QP(\theta)}$$

where $H(\theta) \in \mathcal{S}_+^n(\mathbb{R})$ is a real symmetric positive semi-definite matrix of $\mathbb{R}^{n \times n}$, $g(\theta) \in \mathbb{R}^n$, $C(\theta) \in \mathbb{R}^{n_i \times n}$ and $u(\theta) \in \mathbb{R}^{n_i}$. n is the problem dimension, while n_i is the number of inequality constraints. We will abusively denote H , g , C , and u without explicit dependence on θ when this dependence is clear from the context or does not generate any ambiguity. In order to use $\text{QP}(\theta)$ as a learning tool that can be trained with standard optimization techniques, we need to be able to differentiate $x^*(\theta)$ w.r.t. θ , which is challenging for a few reasons. First, there is usually no practical way to compute a closed-form for $x^*(\theta)$, even when $\text{QP}(\theta)$ is well-defined. Second, even when such an $x^*(\theta)$ exists, there is no guarantee for it to be unique nor differentiable w.r.t. θ (see, e.g., the assumptions of the implicit

*Communications regarding the paper should be addressed to antoine.bambade@inria.fr.

function theorem [Dontchev and Rockafellar, 2009, Theorem 1B.1]). As a consequence, concurrent approaches are generally based on architectures enforcing satisfaction of some strong assumptions. In particular, to the best of our knowledge, previous approaches enforce primal feasibility of the layer during training, which generally requires additional learning variables and limits the modeling power of those layers. For instance, as in [Amos and Kolter, 2017], learning $\text{QP}(\theta)$ requires imposing its feasible set to be non-empty. For imposing this while learning C , the authors also learn $z_0 \in \mathbb{R}^n$ and $s_0 \in \mathbb{R}_+^{n_i}$ and u of the form $u = Cz_0 + s_0$, thereby preventing, among others, u from being fixed independently of the learning.

This work makes the following contributions:

- We propose a unified approach to tackle the differentiability of both feasible and infeasible QPs. The main idea consists in extending the definition of $x^*(\theta)$ to be either a solution to $\text{QP}(\theta)$ when it is feasible or a solution of the closest feasible QP (in the least-square sense) when it is not. By relying on the notion of conservative Jacobian by [Bolte et al., 2021, Bolte and Pauwels, 2020], we notably show that the KKT map G of this extended problem is path differentiable w.r.t. θ and x^* (Section 3.2). In this context, the Jacobian $\frac{\partial x^*(\theta)}{\partial \theta}$ is defined as the least-square solution of the linear system formed by applying the implicit function theorem to G Section 3.3. We show that this definition consistently covers the differentiability of feasible QPs as with the traditional implicit differentiation [Amos and Kolter, 2017] when it is valid and with the least-square estimate proposed by [Agrawal et al., 2019, Appendix B] otherwise.
- In Section 3.4 we provide an efficient way to compute the Jacobian $\frac{\partial x^*(\theta)}{\partial \theta}$ in both forward and backward automatic differentiation modes.
- In Section 4 we demonstrate how the approach enables dealing with possibly infeasible $\text{QP}(\theta)$, allowing to train a broader range of QP layers (e.g., learning QPs that are not generically feasible).

Based on these developments, we provide QPLayer, an open-source implementation with efficient forward and backward passes, released upon paper acceptance. It takes advantage, among others, of recent advances in solving of QP problems to output in the forward pass the closest feasible QP solution in ℓ_2 -sense as soon as the program is primal infeasible [Chiche and Gilbert, 2016]. Section 4 highlights for different learning tasks the numerical robustness, accuracy, and speed of our approach against other state-of-the-art methods.

2 Related work

Differentiation for optimization layers. Under certain regularity conditions, it is possible to implicitly differentiate the optimality conditions of convex optimization problems. [Gould et al., 2016, Gilbert, 2021, Fiacco and McCormick, 1968, Robinson, 1980] present general conditions and techniques under which it is possible to differentiate through constrained optimization problems using the implicit function theorem. [Bolte et al., 2021, Bolte and Pauwels, 2020] present extensions to nonsmooth (not necessarily differentiable) functions for machine learning and optimization applications. [Amos and Kolter, 2017] treats the specific case of QPs with a dedicated network architecture, OptNet, and a specialized batched interior-point solver, Qpth, which allows for efficient backpropagation. More recently, [Butler and Kwon, 2023] proposed an alternative approach based on the differentiation of ADMM steps, yet dealing only with equality and box inequality constraints¹. For more general convex optimization, [Amos et al., 2018] proposed CvxpyLayer that differentiates through "Disciplined Convex Programs" using an LSQR [Paige and Saunders, 1982] solver to speed up the differentiation procedure. [Sun et al., 2022] has recently proposed an ADMM-type method, called Alt-Diff, to alternatively solving a constrained convex optimization program and obtaining approximate Jacobians at the current approximate solutions. [Blondel et al., 2022] also proposed a generic solver, JaxOpt, based on an implicit automatic differentiation mechanism leveraging the Jax framework. Finally, let us mention the work [Sharma et al.,

¹The associated R solver is not yet publicly available.

2022], which provided a Julia library, DiffOpt.jl, for differentiable QP and conic optimization (or any model that can be reformulated into these standard forms).

Unrolling methods. Argmin and Argmax operations can be approximated by first-order methods, which can be unrolled [Domke, 2012, Monga et al., 2021]. These architectures typically introduce an optimization procedure such as gradient descent into the inference procedure [Belanger and McCallum, 2016, Belanger et al., 2017, Amos et al., 2017, Metz et al., 2019], which is usually truncated to a predefined number of iterations. Recently, [Scieur et al., 2022] highlighted the "curse of unrolling" by showing that for unconstrained quadratic optimization, there is a tradeoff between the convergence speed of the iterates and that of the Jacobian. Although unrolling methods are easy to implement, most of their applications are limited to unconstrained problems. Indeed, if constraints are added, the unrolling solutions have to be projected into the feasible region, significantly increasing the computational burdens.

Implicit models. Implicit models replace explicit expressions in neural networks with layers defined by implicit functions [Zhang et al., 2020]. As for optimization layers, the backward pass requires solving nonlinear Jacobian-based equations arising from the implicit function theorem. Recently, there have been a growing number of applications using them, such as neural ODE [Chen et al., 2018], deep equilibrium models [Bai et al., 2019], logical reasoning in deep neural network (using MAXSAT SDP relaxation) [Wang et al., 2019], implicit surface representation [Michalkiewicz et al., 2019], attention mechanisms [Geng et al., 2021a], graph neural networks [Gu et al., 2020]. However, this method is not suitable for optimization layers with complicated constraints. Recently, [Fung et al., 2022] proposes a matrix-free approach to decrease computational costs, and [Geng et al., 2021b] proposes a phantom gradient that relies on fixed-point unrolling and a Neumann series for faster computations of approximate update directions.

3 The extended conservative Jacobian for convex QPs

This section introduces the main contribution of this work: an extended conservative Jacobian for the solutions to QP(θ) allowing to simultaneously deal with feasible and infeasible QPs, as provided in Section 3.2 and Section 3.3. Section 3.4 proposes efficient algorithms for computing them in forward and backward modes. For exposition purposes, Appendix C.1 illustrates the concepts on a few simple examples.

3.1 Problem formulation

For differentiating QPs, we solve a hierarchic problem QP-H(θ) which is equivalent to QP(θ) when QP(θ) is primal feasible (i.e., there exists x s.t. $C(\theta)x \leq u(\theta)$)

$$\begin{aligned}
 s^*(\theta) = \arg \min_{s \in \mathbb{R}^{n_i}} \frac{1}{2} \|s\|_2^2 \\
 \text{s.t. } x^*(\theta), z^*(\theta) \in \arg \min_{x \in \mathbb{R}^n} \max_{z \in \mathbb{R}_+^{n_i}} L(x, z, s; \theta),
 \end{aligned}
 \tag{QP-H(\theta)}$$

with $L(x, z, s; \theta) := \frac{1}{2}x^\top H(\theta)x + x^\top g(\theta) + z^\top (C(\theta)x - u(\theta) - s)$ (namely the Lagrangian of QP(θ) augmented with a slack variable s). The following assumption is necessary and sufficient for guaranteeing QP-H(θ) to have a solution. In this situation, QP-H(θ) is therefore well-posed and s^* is referred to as the optimal shift. It provides a measure of the distance of QP(θ) to be primal infeasible in ℓ_2 -sense (hence $s^* = 0$ iff QP(θ) is feasible).

Assumption 1. $H(\theta)$ is symmetric positive definite in the direction of $g(\theta)$ or $g(\theta)$ is orthogonal to the recession cone of QP(θ), i.e., $g(\theta) \perp C^\infty(\theta) := \{y \in \mathbb{R}^n | C(\theta)[x + \tau y] \leq u(\theta) \text{ s.t. } C(\theta)x \leq u(\theta), \tau \geq 0\}$.

The existence of a solution $(x^*(\theta), z^*(\theta), s^*(\theta))$ is also equivalent to the dual of QP(θ) having a non-empty domain (i.e., being proper), see [Chiche and Gilbert, 2016, Assumption 2.6 and Proposition 2.5]). So, the approach proposed here allows differentiating through dual feasible convex QPs.

3.2 The closest feasible QP

In what follows, we deal with QP-H(θ) via a nonlinear map G :

$$G(x, z, t; \theta) := \begin{bmatrix} H(\theta)x + g(\theta) + C(\theta)^\top z \\ C(\theta)x - u(\theta) - t \\ [[t]_- + z]_+ - z \\ C(\theta)^\top [t]_+ \end{bmatrix}, \quad (\text{G})$$

where $[\cdot]_+$ and $[\cdot]_-$ respectively correspond to component-wise projections on the non-negative and non-positive orthants. The following lemma guarantees solutions to QP-H(θ) to be zeros of G (see proof in Appendix A).

Lemma 1. *Let $H(\theta) \in \mathcal{S}_+^n(\mathbb{R})$, $g(\theta) \in \mathbb{R}^n$, $C(\theta) \in \mathbb{R}^{n_i \times n}$ and $u(\theta) \in \mathbb{R}^{n_i}$ be satisfying Assumption 1. It holds that (x^*, z^*, s^*) solves QP-H(θ) iff there exists $t^* \in \mathbb{R}^{n_i}$ s.t. $G(x^*, z^*, t^*; \theta) = 0$ and $s^* = [t^*]_+$.*

3.3 The extended conservative Jacobian

For differentiating through G , we rely on the notion of extended conservative Jacobian (ECJ). As provided by the following lemma, the nonlinear map $G(x, z, t; \theta)$ is path differentiable (see [Bolte and Pauwels, 2020, Definition 3]) w.r.t. x, z, t and also w.r.t. θ under the assumption that $H(\theta), g(\theta), C(\theta)$ and $u(\theta)$ are differentiable w.r.t. θ . This lemma is proved in Appendix B.

Lemma 2. *G is path differentiable w.r.t. x^*, z^* and t^* . Furthermore, if $H(\theta), g(\theta), C(\theta)$ and $u(\theta)$ are differentiable w.r.t. θ , then G is path differentiable w.r.t. θ .*

Definition 1. *Let $H(\theta), g(\theta), C(\theta)$ and $u(\theta)$ be differentiable w.r.t. θ and satisfying Assumption 1. Let $v^* = (x^*, z^*, t^*) \in \mathbb{R}^n \times \mathbb{R}_+^{n_i} \times \mathbb{R}^{n_i}$ s.t. $G(x^*, z^*, t^*; \theta) = 0$. We refer to the ECJs of x^*, z^* and t^* , respectively denoted by $\frac{\partial x^*}{\partial \theta}, \frac{\partial z^*}{\partial \theta}$ and $\frac{\partial t^*}{\partial \theta}$, as solutions of the following problem*

$$\left(\frac{\partial x^*}{\partial \theta}, \frac{\partial z^*}{\partial \theta}, \frac{\partial t^*}{\partial \theta} \right) \in \arg \min_w \left\| \frac{\partial G(x^*, z^*, t^*; \theta)}{\partial v^*} w + \frac{\partial G(x^*, z^*, t^*; \theta)}{\partial \theta} \right\|_2^2. \quad (1)$$

Furthermore, we refer to an ECJ of $s^* = [t^*]_+$, denoted by $\frac{\partial s^*}{\partial \theta}$, any element satisfying $\Pi \frac{\partial t^*}{\partial \theta} \in \frac{\partial s^*}{\partial \theta}$, with $\Pi \in \partial([\cdot]_+)(t^*)$ a subgradient of the positive orthant evaluated in t^* .

As shown in the next section the ECJs match the definitions of standard Jacobians under standard assumptions guaranteeing differentiability (when the QP is feasible), as provided by [Amos and Kolter, 2017, Dontchev and Rockafellar, 2009]. When the QP is feasible but not differentiable, the ECJ corresponds to a least-square approximation specialized for QPs. A similar practical least-square estimate was proposed in [Agrawal et al., 2019, Appendix B] for differentiating primal solutions of second-order cones (SOCs)². [Blondel et al., 2022, Section 2.1] proposed a similar estimate.

3.4 Deriving an extended conservative Jacobian

This section derives an ECJ and incorporates it in a backpropagation algorithm. It also shows how to efficiently compute this ECJ under primal feasibility.

3.4.1 General case: dealing with both feasible and infeasible QPs

In the following, we provide forward and backward pass algorithms to compute ECJs for both feasible and infeasible QPs.

²More precisely, [Agrawal et al., 2019, Appendix B] relies on a series of assumptions allowing to simplify the computations. In particular, they assume that $\frac{\partial z^*}{\partial \theta} = 0$ and $\frac{\partial t^*}{\partial \theta} = 0$, where t^* is a slack variable.

Forward pass: Let $H(\theta)$, $g(\theta)$, $C(\theta)$ and $u(\theta)$ be differentiable w.r.t. θ and satisfying Assumption 1, and x^*, z^*, t^* s.t. $G(x^*, z^*, t^*; \theta) = 0$. We show in Appendix B.2.1.1 that we can efficiently derive ECJs of x^* , z^* and t^* by solving the following QP using an augmented Lagrangian-based algorithm [Rockafellar, 1976]

$$\frac{\partial x^*}{\partial \theta}, \frac{\partial z^*}{\partial \theta}, \frac{\partial t^*}{\partial \theta} \in \arg \min_{\frac{\partial x}{\partial \theta}, \frac{\partial z}{\partial \theta}, \frac{\partial t}{\partial \theta}} \left\| \begin{bmatrix} H & C^\top & 0 \\ C & 0 & -I \\ 0 & \Pi_1 - I & \Pi_1 \Pi_2 \\ 0 & 0 & C^\top (I - \Pi_2) \end{bmatrix} \begin{bmatrix} \frac{\partial x}{\partial \theta} \\ \frac{\partial z}{\partial \theta} \\ \frac{\partial t}{\partial \theta} \end{bmatrix} + \begin{bmatrix} \frac{\partial H}{\partial \theta} x^* + \frac{\partial g}{\partial \theta} + \frac{\partial C}{\partial \theta}^\top z^* \\ \frac{\partial C}{\partial \theta} x^* - \frac{\partial u}{\partial \theta} \\ 0 \\ \frac{\partial C}{\partial \theta}^\top [t^*]_+ \end{bmatrix} \right\|_2^2, \quad (2)$$

where Π_1 and Π_2 are binary diagonal matrices respectively corresponding to the subdifferentials $\partial([\cdot]_+)([t^*]_- + z^*)$ and $\partial([\cdot]_-)(t^*)$, with the following specific choices in zeros

$$(\Pi_1)_i = 1 \text{ when } [t_i^*]_- + z_i^* = 0, \quad (\Pi_2)_i = 1 \text{ when } t_i^* = 0. \quad (3)$$

Furthermore, an ECJ of s^* can be obtained via $\Pi_2 \frac{\partial t^*}{\partial \theta} \in \frac{\partial s^*}{\partial \theta}$.

Under nonsingularity assumptions and infeasibility of QP(θ), the next lemma ensures that the ECJs consistently correspond to conservative Jacobians [Bolte and Pauwels, 2020]. Its proof is available in Appendix B.2.1.

Lemma 3. *Let $H(\theta)$, $g(\theta)$, $C(\theta)$ and $u(\theta)$ be differentiable w.r.t. θ and satisfying Assumption 1. If $s^* > 0$ and x^* and z^* satisfy strict complementarity and C is nonsingular, then the ECJs of x^* , z^* , t^* , and s^* correspond to conservative Jacobians.*

As s^* is a direct output of an augmented Lagrangian-based algorithm [Chiche and Gilbert, 2016], in what follows, we work with s^* instead of t^* .

Backward pass: Let $h : \mathbb{R}^n \times (\mathbb{R}^{n_i})^2 \rightarrow \mathbb{R}$ be a differentiable function, and let $H(\theta)$, $g(\theta)$, $C(\theta)$ and $u(\theta)$ be differentiable w.r.t. θ and satisfying Assumption 1. Then, denoting $\mathcal{L}(\theta) := h(x^*(\theta), z^*(\theta), s^*(\theta))$ and following the methodology provided in [Amos and Kolter, 2017, Section 3], we show in Appendix B.2.1.2 that under assumptions of Lemma 3 a conservative Jacobian $\frac{\partial \mathcal{L}}{\partial \theta}$ can be obtained from the usual chain rule:

$$\frac{\partial \mathcal{L}}{\partial \theta} = (b_1^*)^\top \frac{\partial H}{\partial \theta} x^* + (b_1^*)^\top \frac{\partial g}{\partial \theta} + (b_2^*)^\top \frac{\partial C}{\partial \theta} x^* + (z^*)^\top \frac{\partial C}{\partial \theta} b_1^* + (s^*)^\top \frac{\partial C}{\partial \theta} b_4^* - (b_2^*)^\top \frac{\partial u}{\partial \theta}, \quad (4)$$

where b_1^* , b_2^* , b_3^* and b_4^* are solutions of the linear system

$$\begin{bmatrix} H & C^\top & 0 & 0 \\ C & 0 & (I - \Pi_1) & 0 \\ 0 & -I & -\Pi_1 \Pi_2 & (1 - \Pi_2)C \end{bmatrix} \begin{bmatrix} b_1^* \\ b_2^* \\ b_3^* \\ b_4^* \end{bmatrix} = - \begin{bmatrix} \frac{\partial \mathcal{L}}{\partial x^*} \\ \frac{\partial \mathcal{L}}{\partial z^*} \\ \frac{\partial \mathcal{L}}{\partial s^*} \end{bmatrix}. \quad (5)$$

If not, the methodology proposed in [Amos and Kolter, 2017, Section 3] cannot be used. Hence, ECJs of x^* , z^* and s^* are derived using forward mode and $\frac{\partial \mathcal{L}}{\partial \theta}$ is recovered from

$$\frac{\partial \mathcal{L}}{\partial \theta} = \frac{\partial \mathcal{L}}{\partial x^*}^\top \frac{\partial x^*}{\partial \theta} + \frac{\partial \mathcal{L}}{\partial z^*}^\top \frac{\partial z^*}{\partial \theta} + \frac{\partial \mathcal{L}}{\partial s^*}^\top \frac{\partial s^*}{\partial \theta}. \quad (6)$$

In practice, we provide a feasibility tolerance that can be set by the user and above which (5) is considered not accurate enough or infeasible. In such cases, the forward mode is used internally.

3.4.2 Exploiting primal feasibility of the QP

In this section, we exploit feasibility of the QP for simplifying the computations. First, for the forward pass, the QP needs only be feasible for the value of θ under consideration. For the backward pass, we exploit the standard assumption (see [Amos and Kolter, 2017]) of the QP being constructively feasible for all values of θ (which is of course restrictive, but which can be exploited for efficiency).

Forward pass: When $QP(\theta)$ is feasible and $H(\theta)$, $g(\theta)$, $C(\theta)$ and $u(\theta)$ are differentiable w.r.t. θ and satisfy Assumption 1, we show in Appendix B.2.2.1 that ECJs can be obtained as a solution to the simpler:

$$\frac{\partial x^*}{\partial \theta}, \frac{\partial z^*}{\partial \theta} \in \arg \min_{\frac{\partial x}{\partial \theta}, \frac{\partial z}{\partial \theta}} \left\| \begin{bmatrix} H & C^\top \\ \frac{\Pi_1}{\sqrt{1+\Pi_1}} C & \Pi_1 - I \end{bmatrix} \begin{bmatrix} \frac{\partial x}{\partial \theta} \\ \frac{\partial z}{\partial \theta} \end{bmatrix} + \begin{bmatrix} \frac{\partial H}{\partial \theta} x^* + \frac{\partial g}{\partial \theta} + \frac{\partial C}{\partial \theta}^\top z^* \\ \frac{\Pi_1}{\sqrt{1+\Pi_1}} \left(\frac{\partial C}{\partial \theta} x^* - \frac{\partial u}{\partial \theta} \right) \end{bmatrix} \right\|^2, \quad (7)$$

$$\frac{\partial t^*}{\partial \theta} = (I + \Pi_1)^{-1} \left(C \frac{\partial x^*}{\partial \theta} + \frac{\partial C}{\partial \theta} x^* - \frac{\partial u}{\partial \theta} \right), \quad (8)$$

where Π_1 is a binary diagonal matrices representing the subdifferential $\partial[\cdot]_+(Cx^* - u + z^*)$ with the following specific choice:

$$(\Pi_1)_i = 1 \text{ when } C_i x^* - u_i + z_i^* = 0. \quad (9)$$

The following lemma (see proof in Appendix B.2.1.2) guarantees that, under standard assumptions, solutions to (7) correspond to standard Jacobians (see, e.g., [Amos and Kolter, 2017]).

Lemma 4. *If $QP(\theta)$ is feasible and $H(\theta)$, $g(\theta)$, $C(\theta)$ and $u(\theta)$ are differentiable w.r.t. θ and satisfy Assumption 1, and if the KKT matrix of active constraints is nonsingular and x^* , z^* satisfy strict complementarity, then the ECJs matches the standard Jacobian, i.e., $\frac{\partial x^*(\theta)}{\partial \theta} = \nabla x^*(\theta)$ and $\frac{\partial z^*(\theta)}{\partial \theta} = \nabla z^*(\theta)$.*

Backward pass: If $QP(\theta)$ is by construction primal feasible for any θ , then for any θ , $s^*(\theta) = 0$. We can exploit this result for considering simpler losses not depending anymore of $s^*(\theta)$. More precisely, let $h : \mathbb{R}^n \times (\mathbb{R}^{n_i}) \rightarrow \mathbb{R}$ be a differentiable function, and let $H(\theta)$, $g(\theta)$, $C(\theta)$ and $u(\theta)$ be differentiable w.r.t. θ and satisfying Assumption 1. Then, denoting $\mathcal{L}(\theta) := h(x^*(\theta), z^*(\theta))$, we show in Appendix B.2.2.2 that when assumptions of Lemma 4 hold the backward pass can be evaluated by solving the following linear system

$$\begin{bmatrix} H & C_J^\top \\ C_J & 0 \end{bmatrix} \begin{bmatrix} b_x^* \\ b_{z_J}^* \end{bmatrix} = - \begin{bmatrix} \frac{\partial \mathcal{L}}{\partial x^*} \\ \frac{\partial \mathcal{L}}{\partial z_J^*} \end{bmatrix}, \quad b_{z_{J^c}}^* = \frac{\partial \mathcal{L}}{\partial z_{J^c}^*}, \quad (10)$$

where J is the set of constraints for which $(\Pi_1)_i = 1$ and J^c the one for which $(\Pi_1)_i = 0$. $\frac{\partial \mathcal{L}}{\partial \theta}$ is then retrieved from the chain rule

$$\frac{\partial \mathcal{L}}{\partial \theta} = (b_x^*)^\top \frac{\partial H}{\partial \theta} x^* + (b_x^*)^\top \frac{\partial g}{\partial \theta} + (\Pi_1 b_z^*)^\top \frac{\partial C}{\partial \theta} x^* + (z^*)^\top \frac{\partial C}{\partial \theta} b_x^* - (\Pi_1 b_z^*)^\top \frac{\partial u}{\partial \theta}. \quad (11)$$

Note that the assumptions of Lemma 4 are not necessarily met. Therefore, if (10) is found to be infeasible we provide an option to use forward mode in the backward pass as in the generic case (see Section 3.4.1). Such infeasibility can be detected easily using iterative refinement [Parikh and Boyd, 2014, Section 4.1.2]) as it converges to the least-square solution of (10) in the infeasible case (see [Güler, 1991, Theorem 2.3]).

3.4.3 Future work and potential improvements

Before moving to the experiments, let us mention a few potential directions for future work and improvements in our approach. There remain a few gaps in the theoretical foundations of our methodology, which we believe should be handled in the future. Indeed, we have not proved that we could apply the chain rule to ECJs in the general case in the spirit of CJs (see [Bolte et al., 2021]). Also, it should be confirmed that ECJ indeed reduces to CJ under weaker assumptions than those of Lemma 3. While those problems are present in most frameworks [Agrawal et al., 2019, Section B], [Blondel et al., 2022, Section 2.1], using the least-square estimate provides good practical results when non-differentiability occurs.

4 Experimental results

Our forward and backward mode differentiation of convex QP layers have been implemented in C++. We refer to it as QP layer in what follows. It will be made available at <https://github.com/Simple-Robotics/proxsuite>. Our code leverages the primal-dual augmented Lagrangian solver ProxQP [Bambade et al., 2022], also written in C++ as its internal QP solver. This section illustrates through different classic learning tasks that QP layer is faster than other state-of-the-art approaches. QP layer further allows relaxing primal feasibility constraints, thereby enabling the training of simplified layers. We also show in Appendix C.3 that QP layer can be used for solving tasks with high learning rates.

Benchmark setup. In the first set of experiments (see Section 4.1.1), QP layer is compared to OptNet, CvxpyLayer, JaxOpt, and Alt-Diff. For all the other experiments, QP layer is benchmarked only against approaches available within the PyTorch framework (i.e., OptNet and CvxpyLayer³). The experiments were conducted using all threads of an Intel(R) Core(TM) i7-4790 CPU @ 3.60GHz. The benchmark API will be released upon acceptance of this work. It was inspired by <https://github.com/locuslab/optnet> and has been extended to cover all our experiments. It should be noted that all the benchmarks were obtained on CPU whereas we plan to release a GPU extension of our techniques in future work.

4.1 Timings

In this section, we report our numerical results and compare them against state-of-the-art frameworks on a set of standard experiments. We outline detailed timings for differentiating solutions on a set of different QPs. First, Table 1 shows the results for a few randomly generated QPs. Second, Table 2 reports the average time spent in the differentiation procedure on four different learning tasks. Additional experiments are provided in Appendix C.2.

4.1.1 Random QPs

In this first set of experiments, we generated random QPs with 100 variables, 50 equality, and 50 inequality constraints. We solve and backpropagate through all those QPs for different forward pass accuracies. We then average the results over 5 trials and report the timings in the spirit of [Sun et al., 2022].

For a batch size of 100, Table 1 shows that QP layer is almost 4 times faster than OptNet (the second fastest approach). We can see similar performance for a batch size of 1 (see Table 4 in the appendix). We observe that the speed gain is mostly due to the forward pass speed-up, enabled by the use of ProxQP and thread parallelization. It is also confirmed by Table 5, which reproduces in the appendix the serial forward timing benchmark proposed in [Amos and Kolter, 2017, Section 4.1]. It exhibits from 5 to 9 times faster computation times.

4.1.2 Learning tasks

For this second set of experiments, we report the numerical results obtained on 4 traditional learning tasks (namely MNIST classification, signal denoising, Sudoku solving and cart-pole experiment). For all experiments, we report the average (over all epochs) time spent in the forward and backward passes. Table 2 reports that QP layer is 3 to 10 times faster than the second fastest approach (i.e., 3 times faster on the classification task, 4 times faster on the Sudoku, about 10 times faster for the denoising and 7 times faster for the cart-pole experiments). In all cases, the test

³Alt-Diff exhibited too slow performances for a fair and reasonable comparison. Note that [Sun et al., 2022] have not yet proposed an Alt-Diff layer deriving all QP Jacobians. Therefore, we have included in our benchmark an open-source implementation of Alt-Diff based on their work.

loss incurred using QP layer is either similar (for the denoising and cart-pole tasks) or far better than its competitor layers (about 2 times smaller and 3 orders of magnitudes better for the classification and Sudoku experiments).

More precisely, the first three experiments reproduce the ones originally described in [Amos and Kolter, 2017, Sections 4.2 to 4.4]. A complete description of the cart-pole swing-up task is detailed in Appendix D.2. These tasks involve learning convex feasible QPs using Adam optimizer [Kingma and Ba, 2014]. The first three experiments are run with the default batch sizes and the number of epochs fixed by the original authors (i.e., batch size equals 64 for classification, 150 for denoising and Sudoku tasks; 30 epochs for classification, and 20 for denoising and Sudoku tasks). We run the cart-pole example with batch size 1 for 800 epochs. For the first three experiments, we use the same QP layer models as [Amos and Kolter, 2017], except that we have changed the backends for executing the forward and backward passes (using either Qpth, QP layer⁴ or CvxpyLayer). Finally, we have used the following learning rates for running the benchmarks: 10^{-3} for classification, 10^{-5} for denoising, 5×10^{-2} for Sudoku, and 10^{-1} for cart-pole tasks.

Forward tolerance ϵ	10^{-1}	10^{-2}	10^{-3}
Forward (ms)	72.43 (± 7.28)	71.89 (± 3.0)	72.12 (± 4.21)
Backward (ms)	18.14 (± 1.01)	18.03 (± 0.17)	18.12 (± 0.37)
QP layer total (ms)	90.57	89.92	90.24
Forward (ms)	340.62 (± 7.06)	348.68 (± 3.51)	349.44 (± 2.44)
Backward (ms)	6.39 (± 0.16)	6.61 (± 0.41)	6.66 (± 0.38)
OptNet total (ms)	347.01	355.29	356.10
Forward (ms)	123.35 (± 13.70)	196.75 (± 38.13)	281.54 (± 64.18)
Backward (ms)	546.91 (± 55.58)	622.45 (± 66.63)	723.34 (± 66.26)
JaxOpt total (ms)	670.16	819.20	1004.88
Forward (ms)	$1.16(\pm 0.038) \times 10^3$	$1.19(\pm 0.015) \times 10^3$	$1.24(\pm 0.017) \times 10^3$
Backward (ms)	187.52 (± 8.59)	187.74 (± 11.54)	197.18 (± 7.99)
CvxpyLayer total (ms)	1.35×10^3	1.38×10^3	1.43×10^3
Forward (ms)	Time limit	Time limit	Time limit
Backward (ms)	Time limit	Time limit	Time limit
Alt-Diff total (ms)	Time limit	Time limit	Time limit

Table 1: Timings for deriving all Jacobians of random QPs with different forward pass accuracies and batch size 100.

4.2 Learning capabilities

Differentiable optimization for neural network layers has shown great representational power for learning problems that are fundamentally rooted in optimization. The Sudoku problem is one such problem, which can naturally be cast as a mixed integer linear program (MILP). For the Sudoku, OptNet recently showed better robustness and prediction accuracy results than traditional neural networks [Amos and Kolter, 2017, Section 4.4]. This section shows that QP layer generalizes even better by exploiting the fact that it allows learning LPs (and not only QPs). Further, the ability of QP layer to deal with possibly primal infeasible problems during the learning process appears to be key.

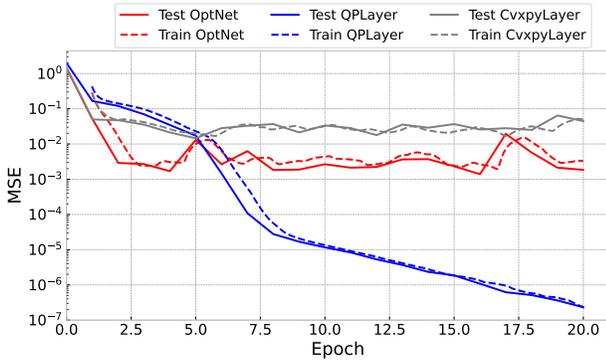
4.2.1 Learning linear programs

The Sudoku problem is detailed in [Amos and Kolter, 2017, Section 4.4]. We reproduce those experiments with OptNet and CvxpyLayer, while letting QP layer learn linear programs (LPs) instead of strictly convex QPs (the layer model is detailed in Figure 9 of the appendix). OptNet, CvxpyLayer, and QP layer were trained using Adam with a batch of size 150 and a learning rate of 0.05 to minimize an MSE loss on the dataset created by [Amos and Kolter, 2017]. The dataset contains 9000 training puzzles and 1000 held-out puzzles for testing. First, Figure 1a

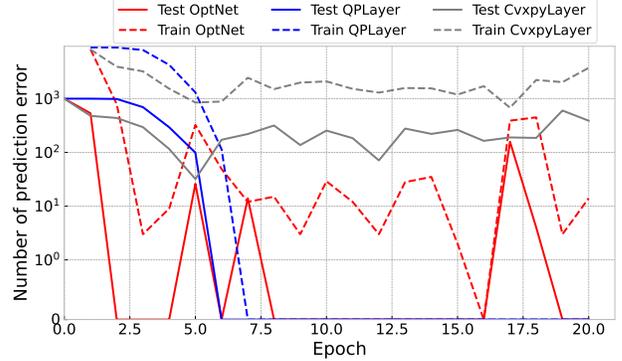
⁴Yet two differences should be noted: QP layer learns LP for the Sudoku experiment. We have not imposed zero Hessian for the CvxpyLayer even if it could learn it, as it would display worse results. Furthermore, for the denoising experiment, QP layer learns the lower and upper bounds simultaneously.

Learning Tasks	QP layer	OptNet	CvxpyLayer
cart-pole			
Forward (ms)	55.20 (± 6.93)	615.84 (± 16.15)	629.70 (± 30.42)
Backward (ms)	39.27 (± 2.17)	61.26 (± 2.84)	101.88 (± 1.11)
Final test loss	0.02556	0.02604	0.02566
Sudoku			
Forward (ms)	87.39 (± 16.69)	454.48 (± 22.22)	772.77 (± 33.89)
Backward (ms)	20.80 (± 1.70)	8.07 (± 0.42)	192.99 (± 9.81)
Final test loss	2.31×10^{-7}	0.0017	0.389
denoising			
Forward (ms)	247.89 (± 17.03)	2827.48 (± 618.78)	Error
Backward (ms)	52.99 (± 2.75)	40.57 (± 2.98)	Error
Final test loss	3281.84	3529.2029	Error
classification			
Forward (ms)	26.77 (± 3.63)	102.62 (± 18.23)	Error
Backward (ms)	11.12 (± 3.01)	16.58 (± 12.05)	Error
Final test loss	0.1363	0.3264	Error

Table 2: Timings and final loss of 4 learning tasks. CvxpyLayer errors arise due to failures in filling the disciplined parametrized programming (DPP) form of the quadratic cost.



(a) Test and training MSE losses of QPLayer, CvxpyLayer and OptNet layers.



(b) Test and training prediction errors of QPLayer, CvxpyLayer, and OptNet over 1000 and 9000 puzzles.

Figure 1: Sudoku training and test plots using QPLayer, OptNet, and CvxpyLayer. QPLayer can learn LPs (which appear more appropriate), whereas OptNet is restricted to strictly convex QPs.

shows that QPLayer minimizes the training and test loss without ending up over-fitting to the training data, contrary to OptNet and CvxpyLayer which appear to saturate. Second, Figure 1b shows that QPLayer achieves significantly more accurate and robust training and test error predictions than OptNet and CvxpyLayer.

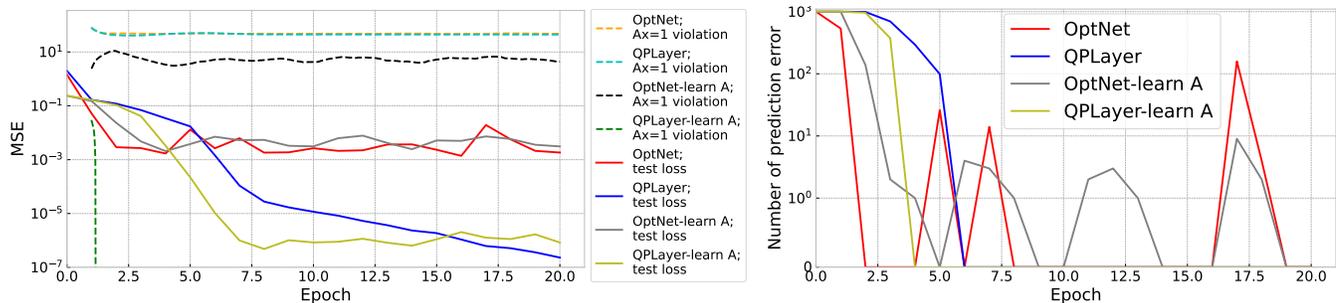
4.2.2 Handling primal infeasibility

As outlined in Section 4.2.1, forcing primal feasibility while learning is a common algorithmic strategy. For the Sudoku problems, those techniques enforcing primal feasibility typically involve neglecting a linear equality constraint $Ax = 1$ (we learn A) which corresponds to Sudoku rules. As shown in Figure 2b, this means that the learning procedures do not respect the Sudoku rule constraint (see green and orange dashed lines—labeled "QPLayer; $Ax = 1$ violation" and "OptNet; $Ax = 1$ violation"). In the end, neglecting this constraint ultimately leads to learning a constraint matrix that is inconsistent with the Sudoku rules. Relaxing the primal feasibility imposed by differentiation procedures of previous solvers thereby appears to be key.

By incorporating a potential optimal shift s^* in its formulation (as exposed in Section 3.1), QPLayer allows dealing with infeasible problems. Numerical performances are reported in Figure 2a and Figure 2b. It is apparent that the dark green curve labeled "QPLayer-learn A ; $Ax = 1$ violation" converges after the end of the first epoch towards a model satisfying Sudoku rules. It also converges slightly faster towards a regime without any prediction errors. The steeper slope observed in the graph suggests that it might be worthwhile to train a layer that more accurately adheres to the Sudoku rules, as this could potentially lead to faster puzzle-solving and more interpretable outcomes. For comparison, the gray curves correspond to results obtained by training a strictly convex QP layer with relaxed primal feasibility constraint (labeled "OptNet-learn A ; $Ax = 1$ violation") using OptNet. As expected, it fails to satisfy the primal feasibility constraint (i.e., the black dashed curve value is around 10) and displays a worse prediction error.

4.3 Additional experiments

Appendix C.1 provides different simple experiments with parametric QPs to illustrate ECJs concept. Appendix C.2 contains additional timing experiments. Appendix C.3 illustrates through several experiments that QPLayer is numerically more robust and can thereby be trained with large learning rates.



(a) Test MSE loss of QP layer, OptNet, QP layer-learn A, and OptNet-learn A specialized for learning A. It includes Sudoku $Ax = 1$ violation.

(b) Test prediction errors over 1000 puzzles of OptNet, QP layer, QP layer-learn A and OptNet-learn A specialized for learning A.

Figure 2: Sudoku training and test plots using QP layer and OptNet layers. QP layer can learn LPs, whereas OptNet is restricted to strictly convex QPs, which limits its representational power. Contrary to OptNet, QP layer can be specialized to learn models satisfying specific linear constraints.

5 Conclusion

In this work, we introduced an approach for differentiating both feasible and infeasible convex quadratic programs in a unified fashion. This approach is particularly relevant for learning with optimization layers through differentiable optimization.

In particular, by leveraging augmented Lagrangian techniques for solving QP layers that are potentially infeasible, we propose an extended conservative Jacobian formulation for differentiating convex QPs, covering both feasible and infeasible problems. For feasible problems, and when the solution is differentiable, this reduces to standard Jacobians. We further provide an open-source C++ framework, referred to as “QP layer”, which efficiently implements the approach. Through extensive benchmarks and experiments, we have demonstrated how QP layer is simultaneously faster and numerically more robust than alternative state-of-the-art optimization layers.

As for future plans, we will extend QP layer to operate on GPUs and to deal with a broader range of optimization layers that include second-order cones.

Acknowledgments

This work was supported in part by the French government under the management of Agence Nationale de la Recherche through the project INEXACT (ANR-22-CE33-0007-01) and as part of the "Investissements d’avenir" program, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute), by the European Union through the AGIMUS project (GA no.101070165) and the Louis Vuitton ENS Chair on Artificial Intelligence. Views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

Bibliography

- A. Agrawal, B. Amos, S. T. Barratt, S. P. Boyd, S. Diamond, and J. Z. Kolter. Differentiable convex optimization layers. In *NeurIPS*, pages 9558–9570, 2019. URL <https://dblp.org/rec/conf/nips/AgrawalABBDK19>.
- B. Amos and J. Z. Kolter. Optnet: Differentiable optimization as a layer in neural networks. In *ICML*, volume 70 of *Proceedings of Machine Learning Research*, pages 136–145. PMLR, 2017. URL <https://dblp.org/rec/conf/icml/AmosK17>.

- B. Amos, L. Xu, and J. Z. Kolter. Input convex neural networks. In *ICML*, volume 70 of *Proceedings of Machine Learning Research*, pages 146–155. PMLR, 2017. URL <https://dblp.org/rec/conf/icml/AmosXK17>.
- B. Amos, I. D. J. Rodriguez, J. Sacks, B. Boots, and J. Z. Kolter. Differentiable MPC for end-to-end planning and control. In *NeurIPS*, pages 8299–8310, 2018. URL <https://dblp.org/rec/conf/nips/AmosRSBK18>.
- C. Anderson. Learning to control an inverted pendulum using neural networks. *IEEE Control Systems Magazine*, 9(3):31–37, 1989. doi: 10.1109/37.24809. URL <https://ieeexplore.ieee.org/abstract/document/24809>.
- S. Bai, J. Z. Kolter, and V. Koltun. Deep equilibrium models. In *NeurIPS*, pages 688–699, 2019. URL <https://dblp.org/rec/conf/nips/BaiKK19>.
- A. Bambade, S. El-Kazdadi, A. Taylor, and J. Carpentier. PROX-QP: Yet another Quadratic Programming Solver for Robotics and beyond. In *RSS 2022 - Robotics: Science and Systems*, New York, United States, June 2022. URL <https://hal.inria.fr/hal-03683733>.
- D. Belanger and A. McCallum. Structured prediction energy networks. In *ICML*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 983–992. JMLR.org, 2016. URL <https://dblp.org/rec/conf/icml/BelangerM16>.
- D. Belanger, B. Yang, and A. McCallum. End-to-end learning for structured prediction energy networks. In *ICML*, volume 70 of *Proceedings of Machine Learning Research*, pages 429–439. PMLR, 2017. URL <https://dblp.org/rec/conf/icml/BelangerYM17>.
- M. Blondel, Q. Berthet, M. Cuturi, R. Frostig, S. Hoyer, F. Llinares-López, F. Pedregosa, and J. Vert. Efficient and modular implicit differentiation. In *NeurIPS*, 2022. URL <https://dblp.org/rec/conf/nips/BlondelBCFHLPV22>.
- J. Bolte and E. Pauwels. Conservative set valued fields, automatic differentiation, stochastic gradient method and deep learning. *Mathematical Programming*, 188(19-51), Apr. 2020. doi: 10.1007/s10107-020-01501-5. URL <https://hal.science/hal-02521848>.
- J. Bolte, T. Le, E. Pauwels, and A. Silveti-Falls. Nonsmooth implicit differentiation for machine-learning and optimization. In *NeurIPS*, pages 13537–13549, 2021. URL <https://dblp.org/rec/conf/nips/BolteLPS21>.
- O. Bounou, J. Ponce, and J. Carpentier. Online learning and control of complex dynamical systems from sensory input. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=x6tV8QhHjs1>.
- A. Butler and R. H. Kwon. Efficient differentiable quadratic programming layers: an ADMM approach. *Comput. Optim. Appl.*, 84(2):449–476, 2023. URL <https://dblp.org/rec/journals/coap/ButlerK23>.
- T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. Duvenaud. Neural ordinary differential equations. In *NeurIPS*, pages 6572–6583, 2018. URL <https://dblp.org/rec/conf/nips/ChenRBD18>.
- A. Chiche and J. C. Gilbert. How the augmented Lagrangian algorithm can deal with an infeasible convex quadratic optimization problem. *Journal of Convex Analysis*, 23(2), 2016. URL <https://hal.inria.fr/hal-01057577>.
- F. de Avila Belbute-Peres, K. A. Smith, K. R. Allen, J. Tenenbaum, and J. Z. Kolter. End-to-end differentiable physics for learning and control. In *NeurIPS*, pages 7178–7189, 2018. URL <https://dblp.org/rec/conf/nips/Belbute-PeresSA18>.
- A. De Marchi. On a primal-dual Newton proximal method for convex quadratic programs. *Computational Optimization and Applications*, 81(2):369–395, 2022. ISSN 0926-6003, 1573-2894. doi: 10.1007/s10589-021-00342-y. URL <https://link.springer.com/10.1007/s10589-021-00342-y>.
- J. Domke. Generic methods for optimization-based modeling. In *AISTATS*, volume 22 of *JMLR Proceedings*, pages 318–326. JMLR.org, 2012. URL <https://dblp.org/rec/journals/jmlr/Domke12>.

- A. L. Dontchev and R. T. Rockafellar. Implicit functions and solution mappings. 2009.
- P. L. Donti, J. Z. Kolter, and B. Amos. Task-based end-to-end model learning in stochastic optimization. In *NIPS*, pages 5484–5494, 2017. URL <https://dblp.org/rec/conf/nips/DontiKA17>.
- A. V. Fiacco and G. P. McCormick. *Nonlinear Programming: Sequential Unconstrained Minimization Techniques*. John Wiley & Sons, New York, NY, USA, 1968. Reprinted by SIAM Publications in 1990.
- S. W. Fung, H. Heaton, Q. Li, D. McKenzie, S. J. Osher, and W. Yin. JFB: jacobian-free backpropagation for implicit networks. In *AAAI*, pages 6648–6656. AAAI Press, 2022. URL <https://dblp.org/rec/conf/aaai/FungHLM0Y22>.
- Z. Geng, D. Johnson, and R. Fedkiw. Coercing machine learning to output physically accurate results. *J. Comput. Phys.*, 406:109099, 2020. URL <https://dblp.org/rec/journals/jcphy/GengJF20>.
- Z. Geng, M. Guo, H. Chen, X. Li, K. Wei, and Z. Lin. Is attention better than matrix decomposition? In *ICLR*. OpenReview.net, 2021a. URL <https://dblp.org/rec/conf/iclr/GengGCLWL21>.
- Z. Geng, X. Zhang, S. Bai, Y. Wang, and Z. Lin. On training implicit models. In *NeurIPS*, pages 24247–24260, 2021b. URL <https://dblp.org/rec/conf/nips/GengZBWL21>.
- J. C. Gilbert. Fragments d’Optimisation Différentiable - Théories et Algorithmes. Lecture, Mar. 2021. URL <https://hal.inria.fr/hal-03347060>.
- S. Gould, B. Fernando, A. Cherian, P. Anderson, R. S. Cruz, and E. Guo. On differentiating parameterized argmin and argmax problems with application to bi-level optimization. *CoRR*, abs/1607.05447, 2016. URL <https://dblp.org/rec/journals/corr/GouldFCACG16>.
- F. Gu, H. Chang, W. Zhu, S. Sojoudi, and L. E. Ghaoui. Implicit graph neural networks. In *NeurIPS*, 2020. URL <https://dblp.org/rec/journals/corr/abs-2009-06211>.
- O. Güler. On the convergence of the proximal point algorithm for convex minimization. *SIAM Journal on Control and Optimization*, 29(2):403–419, 1991. doi: 10.1137/0329022. URL <https://doi.org/10.1137/0329022>.
- D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014.
- Q. Le Lidec, I. Kalevatykh, I. Laptev, C. Schmid, and J. Carpentier. Differentiable simulation for physical system identification. *IEEE Robotics and Automation Letters*, 6(2):3413–3420, 2021.
- K. Lee, S. Maji, A. Ravichandran, and S. Soatto. Meta-learning with differentiable convex optimization. *CoRR*, abs/1904.03758, 2019. URL <http://arxiv.org/abs/1904.03758>.
- Q. L. Lidec, L. Montaut, C. Schmid, I. Laptev, and J. Carpentier. Leveraging randomized smoothing for optimal control of nonsmooth dynamical systems. *CoRR*, abs/2203.03986, 2022. URL <https://dblp.org/rec/journals/corr/abs-2203-03986>.
- L. Metz, N. Maheswaranathan, J. Nixon, C. D. Freeman, and J. Sohl-Dickstein. Understanding and correcting pathologies in the training of learned optimizers. In *ICML*, volume 97 of *Proceedings of Machine Learning Research*, pages 4556–4565. PMLR, 2019. URL <https://dblp.org/rec/conf/icml/MetzMNFS19>.
- M. Michalkiewicz, J. K. Pontes, D. Jack, M. Baktashmotlagh, and A. P. Eriksson. Implicit surface representations as layers in neural networks. In *ICCV*, pages 4742–4751. IEEE, 2019. URL <https://dblp.org/rec/conf/iccv/MichalkiewiczPJ19>.
- V. Monga, Y. Li, and Y. C. Eldar. Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing. *IEEE Signal Process. Mag.*, 38(2):18–44, 2021. URL <https://dblp.org/rec/journals/spm/MongaLE21>.

- C. C. Paige and M. A. Saunders. LSQR: an algorithm for sparse linear equations and sparse least squares. *ACM Trans. Math. Softw.*, 8(1):43–71, 1982. URL <https://dblp.org/rec/journals/toms/PaigeS82>.
- N. Parikh and S. Boyd. Proximal algorithms. *Foundations and Trends® in Optimization*, 1(3):127–239, 2014. ISSN 2167-3888. doi: 10.1561/2400000003. URL <http://dx.doi.org/10.1561/2400000003>.
- S. M. Robinson. Strongly regular generalized equations. *Math. Oper. Res.*, 5(1):43–62, 1980. URL <https://dblp.org/rec/journals/mor/Robinson80>.
- R. T. Rockafellar. Augmented Lagrangians and Applications of the Proximal Point Algorithm in Convex Programming. *Mathematics of Operations Research*, 1(2):97–116, 1976. URL <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.298.6206&rep=rep1&type=pdf>.
- D. Scieur, G. Gidel, Q. Bertrand, and F. Pedregosa. The curse of unrolling: Rate of differentiating through optimization. In *NeurIPS*, 2022. URL <https://dblp.org/rec/conf/nips/ScieurGBP22>.
- A. Sharma, M. Besançon, J. D. Garcia, and B. Legat. Flexible differentiable optimization via model transformations. *CoRR*, abs/2206.06135, 2022. URL <https://dblp.org/rec/journals/corr/abs-2206-06135>.
- H. J. T. Suh, T. Pang, and R. Tedrake. Bundled gradients through contact via randomized smoothing. *CoRR*, abs/2109.05143, 2021. URL <https://dblp.org/rec/journals/corr/abs-2109-05143>.
- D. Sun and L. Qi. On ncp-functions. *Comput. Optim. Appl.*, 13(1-3):201–220, 1999. URL <https://dblp.org/rec/journals/coap/SunQ99>.
- H. Sun, Y. Shi, J. Wang, H. D. Tuan, H. V. Poor, and D. Tao. Alternating differentiation for optimization layers. *CoRR*, abs/2210.01802, 2022. URL <https://dblp.org/rec/journals/corr/abs-2210-01802>.
- P. Wang, P. L. Donti, B. Wilder, and J. Z. Kolter. Satnet: Bridging deep learning and logical reasoning using a differentiable satisfiability solver. In *ICML*, volume 97 of *Proceedings of Machine Learning Research*, pages 6545–6554. PMLR, 2019. URL <https://dblp.org/rec/conf/icml/WangDWK19>.
- Q. Zhang, Y. Gu, M. Michalkiewicz, M. Baktashmotlagh, and A. P. Eriksson. Implicitly defined layers in neural networks. *CoRR*, abs/2003.01822, 2020. URL <https://dblp.org/rec/journals/corr/abs-2003-01822>.

Organization of the appendix

Section	Content
Appendix A	Lemma 1: solutions to QP-H(θ).
Appendix B ECJs and automatic differentiation.	Lemma 2: path differentiability of G (Appendix B.1).
	Forward AD (general case, Appendix B.2.1.1).
	Backward AD (general case, Appendix B.2.1.2).
	Forward AD (feasible QPs, Appendix B.2.2.1).
	Backward AD (structurally feasible QPs, Appendix B.2.2.2).
Appendix C.1 Pedagogical examples.	Lemma 3: when do ECJs reduce to CJs? (Appendix B.3).
	Lemma 4: when do ECJs reduce to Jacobians? (Appendix B.4).
	Strictly convex QP (parameterized constraints) (Appendix C.1.1).
	Strictly convex QP (parameterized objective) (Appendix C.1.2).
Appendix C.2	Parameterized linear program (Appendix C.1.3).
Appendix C.3	Additional benchmarks (timing).
Appendix D	Training with large learning rates experiments.
Appendix E	Layer architecture for the Sudoku problem (Appendix D.1).
Experimental setups.	Description of the cart-pole problem (Appendix D.2).

Table 3: Organization of the appendix. QP stands for “quadratic programming”, AD stands for “automatic differentiation”, and (E)CJ stands for (extended) conservative Jacobian.

A Proof of Lemma 1

For proving Lemma 1, we first show that solutions to QP-H(θ) are zeros of the map \mathcal{G} :

$$\mathcal{G}(x, z; \theta) := \begin{bmatrix} H(\theta)x + g(\theta) + C(\theta)^\top z \\ [[C(\theta)x - u(\theta)]_- + z]_+ - z \\ C(\theta)^\top [C(\theta)x - u(\theta)]_+ \end{bmatrix}. \quad (12)$$

Then, a suitable change of variable shows that finding a zero of \mathcal{G} is equivalent to finding a zero of map G .

Lemma 1. *Let $H(\theta) \in \mathcal{S}_+^n(\mathbb{R})$, $g(\theta) \in \mathbb{R}^n$, $C(\theta) \in \mathbb{R}^{n_i \times n}$ and $u(\theta) \in \mathbb{R}^{n_i}$ be satisfying Assumption 1. It holds that (x^*, z^*, s^*) solves QP-H(θ) iff there exists $t^* \in \mathbb{R}^{n_i}$ s.t. $G(x^*, z^*, t^*; \theta) = 0$ and $s^* = [t^*]_+$.*

Proof. We first show that $(x^*, z^*, [Cx^* - u]_+)$ solves (QP-H(θ)) if and only if $\mathcal{G}(x^*, z^*; \theta) = 0$.

The optimal shift s^* (that corresponds to the closest feasible QP) is equal to $[Cx^* - u]_+$ and is characterized by the ℓ_2 optimality condition [Chiche and Gilbert, 2016, Lemma 2.13]:

$$C^\top [Cx^* - u]_+ = 0.$$

Furthermore, for a feasible problem, the KKT conditions using nonlinear complementarity formulation [Sun and Qi, 1999] reads [De Marchi, 2022, Section 2.1]:

$$\begin{aligned} Hx^* + g + C^\top z^* &= 0, \\ [Cx^* - u + z^*]_+ - z^* &= 0. \end{aligned} \tag{13}$$

For showing equivalence, it is thereby sufficient to show that the second line of (13) corresponds to:

$$[[Cx^* - u]_- + z^*]_+ - z^* = 0. \tag{14}$$

This equivalence is straightforward when $(QP(\theta))$ is feasible, it therefore follows that we only need to handle the infeasible case. When $(QP(\theta))$ is primal infeasible, then $t^* = Cx^* - u$ has a set of components $I \subset [1, n_i]$ strictly positive, hence $s_I^* = [t_I^*]_+ = t_I^* > 0$. For these components, a solution x^* of the closest feasible QP lies on the border $C_I x^* = u_I + t_I^*$. The complementarity condition [De Marchi, 2022, Section 2.1] for these components reads:

$$\underbrace{[Cx^* - u_I - t_I^* + z_I^*]_+}_{=0} - z_I^* = 0,$$

and we thus have $[z_I^*]_+ = z_I^*$. For the other set of components, which we denote by I^c , we have that $s_{I^c}^* = 0$, and hence x^* follows the complementary conditions as in the feasible case

$$[Cx^* - u_{I^c} + z_{I^c}^*]_+ - z_{I^c}^* = 0.$$

Therefore, it follows that (14) captures the two cases (that is, both feasible and infeasible QPs), which concludes the first part of the proof.

Finally, introducing the slack variable $t^* = C(\theta)x^* - u(\theta)$, we have

$$\mathcal{G}(x^*, z^*; \theta) = 0 \tag{15}$$

$$\Leftrightarrow \begin{bmatrix} H(\theta)x^* + g(\theta) + C(\theta)^\top z^* \\ [[C(\theta)x^* - u(\theta)]_- + z^*]_+ - z^* \\ C(\theta)^\top [C(\theta)x^* - u(\theta)]_+ \end{bmatrix} = 0 \tag{16}$$

$$\Leftrightarrow \begin{bmatrix} H(\theta)x^* + g(\theta) + C(\theta)^\top z^* \\ C(\theta)x^* - u(\theta) - t^* \\ [[t^*]_- + z^*]_+ - z^* \\ C(\theta)^\top [t^*]_+ \end{bmatrix} = 0 \tag{17}$$

$$G(x^*, z^*, t^*; \theta) = 0. \tag{18}$$

Hence $G(x^*, z^*, t^*) = 0$ iff $(x^*, z^*, [t^*]_+)$ solves QP-H(θ), which concludes.

□

B ECJs and automatic differentiation

This section provides proofs of the different results used in Section 3. In particular, we define ECJs and provide algorithms for computing them (both in forward and backward AD modes).

B.1 Proof of Lemma 2

Lemma 2. *G is path differentiable w.r.t. x^* , z^* and t^* . Furthermore, if $H(\theta)$, $g(\theta)$, $C(\theta)$ and $u(\theta)$ are differentiable w.r.t. θ , then G is path differentiable w.r.t. θ .*

Proof. We start with the first claim of the lemma. The non-negative projector $[\cdot]_+$ is (component-wise) convex, and hence path differentiable [Bolte and Pauwels, 2020, Proposition 2(i)]. Thus, it remains to show that the third component of G is path differentiable for reaching the desired conclusion.

To do so, we show that the third component is Lipschitz continuous and real semialgebraic [Bolte and Pauwels, 2020, Proposition 2(iv)]. Without loss of generality, we restrict ourselves to the case with 2 components (one for the dual variables, and one for the slack variables) using the following function $h : \mathbb{R}^2 \rightarrow \mathbb{R}$, s.t. $h(z, s) := [[s]_- + z]_+ - z$. Then, the following Lipschitzness argument applies component-wise.

Let $(s_1, z_1) \in \mathbb{R}^2$ and $(s_2, z_2) \in \mathbb{R}^2$:

$$\begin{aligned}
|h(z_1, s_1) - h(z_2, s_2)| &\leq |[[s_1]_- + z_1]_+ - [[s_2]_- + z_2]_+| + |z_1 - z_2| \\
&\leq |[s_1]_- + z_1 - [s_2]_- - z_2| + |z_1 - z_2| \text{ by monotonicity of } [\cdot]_+ \\
&\leq |[s_1]_- - [s_2]_-| + 2|z_1 - z_2| \\
&= |s_1 - [s_1]_+ - (s_2 - [s_2]_+)| + 2|z_1 - z_2| \\
&\leq |s_1 - s_2| + |[s_1]_+ - [s_2]_+| + 2|z_1 - z_2| \\
&\leq 2|s_1 - s_2| + 2|z_1 - z_2| \text{ by monotonicity of } [\cdot]_+ \\
&\leq 2\sqrt{2} \left\| \begin{bmatrix} s_1 \\ z_1 \end{bmatrix} - \begin{bmatrix} s_2 \\ z_2 \end{bmatrix} \right\|_2.
\end{aligned}$$

Therefore, h is Lipschitz continuous. For showing that the third component G describes a semi-algebraic set, we explicitly formulate the graph of h as a finite union of base semi-algebraic sets, as follows:

$$\begin{aligned}
\text{gph}(h) &= \{(z, s, y) \in \mathbb{R}^3 \mid y = [s]_- \text{ and } [s]_- + z > 0\} \\
&\quad \cup \{(z, s, y) \in \mathbb{R}^3 \mid y = -z \text{ and } [s]_- + z = 0\} \\
&\quad \cup \{(z, s, y) \in \mathbb{R}^3 \mid y = -z \text{ and } [s]_- + z < 0\} \\
&= \{(z, s, y) \in \mathbb{R}^3 \mid y = s \text{ and } s + z > 0 \text{ and } s < 0\} \\
&\quad \cup \{(z, s, y) \in \mathbb{R}^3 \mid y = s \text{ and } s = 0\} \\
&\quad \cup \{(z, s, y) \in \mathbb{R}^3 \mid y = 0 \text{ and } z > 0 \text{ and } s > 0\} \\
&\quad \cup \{(z, s, y) \in \mathbb{R}^3 \mid y = -z \text{ and } s + z = 0 \text{ and } s < 0\} \\
&\quad \cup \{(z, s, y) \in \mathbb{R}^3 \mid y = -z \text{ and } z = 0 \text{ and } s = 0\} \\
&\quad \cup \{(z, s, y) \in \mathbb{R}^3 \mid y = -z \text{ and } z = 0 \text{ and } s > 0\} \\
&\quad \cup \{(z, s, y) \in \mathbb{R}^3 \mid y = -z \text{ and } s + z < 0 \text{ and } s < 0\} \\
&\quad \cup \{(z, s, y) \in \mathbb{R}^3 \mid y = -z \text{ and } z < 0 \text{ and } s = 0\} \\
&\quad \cup \{(z, s, y) \in \mathbb{R}^3 \mid y = -z \text{ and } z < 0 \text{ and } s > 0\}.
\end{aligned}$$

Hence, $\text{gph}(h)$ is real and semi-algebraic as it is a finite union of sets defined by polynomial equalities and inequalities. Hence h is Lipschitz continuous and real semi-algebraic, thereby reaching the target conclusion for the first part of Lemma 2.

As for the second part of Lemma 2. G is linear, and hence differentiable, w.r.t. $H(\theta)$, $g(\theta)$, $C(\theta)$ and $u(\theta)$. Furthermore, by assumption $H(\theta)$, $g(\theta)$, $C(\theta)$ and $u(\theta)$ are differentiable w.r.t. θ . As differentiability implies path-differentiability [Bolte and Pauwels, 2020, Remark 3b], we arrive at the desired claim by the conservativity of the chain rule for path-differentiable functions [Bolte and Pauwels, 2020, Proposition 2]. \square

B.2 Forward and backward AD for computing ECJs

This section provides technical details for the computation of ECJs in forward and backward modes for both primal feasible and infeasible problems. We further include the proofs of Lemma 3 and Lemma 4.

B.2.1 General case

B.2.1.1 Forward pass. Let $H(\theta)$, $g(\theta)$, $C(\theta)$ and $u(\theta)$ be differentiable w.r.t. θ and satisfying Assumption 1, and x^*, z^*, t^* s.t. $G(x^*, z^*, t^*; \theta) = 0$.

As G is path-differentiable w.r.t. $v^* := (x^*, z^*, t^*)$ (see Lemma 2), we have

$$\begin{bmatrix} H & C^\top & 0 \\ C & 0 & -I \\ 0 & \Pi_1 - I & \Pi_1 \Pi_2 \\ 0 & 0 & C^\top \Pi_3 \end{bmatrix} \in \frac{\partial G(x^*, z^*, t^*; \theta)}{\partial v^*},$$

for some $\Pi_1 \in \partial[\cdot]_+([t^*]_- + z^*)$, $\Pi_2 \in \partial[\cdot]_-(t^*)$ and $\Pi_3 \in \partial[\cdot]_+(t^*)$.

Furthermore, as G is linear w.r.t. $H(\theta)$, $g(\theta)$, $C(\theta)$ and $u(\theta)$ and $H(\theta)$, $g(\theta)$, $C(\theta)$ and $u(\theta)$ are differentiable w.r.t. θ , the usual chain rule dictates that

$$\frac{\partial G(x^*, z^*, t^*; \theta)}{\partial \theta} = \begin{bmatrix} \frac{\partial H}{\partial \theta} x^* + \frac{\partial g}{\partial \theta} + \frac{\partial C}{\partial \theta}^\top z^* \\ \frac{\partial C}{\partial \theta} x^* - \frac{\partial u}{\partial \theta} \\ 0 \\ \frac{\partial C}{\partial \theta}^\top [t^*]_+ \end{bmatrix}.$$

Finally, as $\partial[\cdot]_+(0) = \partial[\cdot]_-(0) = [0, 1]$, we make the following arbitrary choices in zeros

$$\begin{aligned} \Pi_1 &= I \text{ when } [t^*]_- + z^* = 0, \\ \Pi_2 &= I \text{ when } t^* = 0, \\ \Pi_3 &= 0 \text{ when } t^* = 0, \end{aligned} \tag{19}$$

so that $\Pi_3 = I - \Pi_2$. ECJs are thus retrieved as solutions to:

$$\frac{\partial x^*}{\partial \theta}, \frac{\partial z^*}{\partial \theta}, \frac{\partial t^*}{\partial \theta} \in \arg \min_{\substack{\frac{\partial x}{\partial \theta}, \frac{\partial z}{\partial \theta}, \frac{\partial t}{\partial \theta}}} \left\| \begin{bmatrix} H & C^\top & 0 \\ C & 0 & -I \\ 0 & \Pi_1 - I & \Pi_1 \Pi_2 \\ 0 & 0 & C^\top (I - \Pi_2) \end{bmatrix} \begin{bmatrix} \frac{\partial x}{\partial \theta} \\ \frac{\partial z}{\partial \theta} \\ \frac{\partial t}{\partial \theta} \end{bmatrix} + \begin{bmatrix} \frac{\partial H}{\partial \theta} x^* + \frac{\partial g}{\partial \theta} + \frac{\partial C}{\partial \theta}^\top z^* \\ \frac{\partial C}{\partial \theta} x^* - \frac{\partial u}{\partial \theta} \\ 0 \\ \frac{\partial C}{\partial \theta}^\top [t^*]_+ \end{bmatrix} \right\|_2^2. \tag{20}$$

In practice, solving this problem can be done via an augmented Lagrangian-based solver for the problem:

$$\begin{aligned} \frac{\partial x^*}{\partial \theta}, \frac{\partial z^*}{\partial \theta}, \frac{\partial t^*}{\partial \theta} &\in \arg \min_{\substack{\frac{\partial x}{\partial \theta}, \frac{\partial z}{\partial \theta}, \frac{\partial t}{\partial \theta}}} 0 \\ \text{s.t.} &\begin{bmatrix} H & C^\top & 0 \\ C & 0 & -I \\ 0 & \Pi_1 - I & \Pi_1 \Pi_2 \\ 0 & 0 & C^\top (I - \Pi_2) \end{bmatrix} \begin{bmatrix} \frac{\partial x}{\partial \theta} \\ \frac{\partial z}{\partial \theta} \\ \frac{\partial t}{\partial \theta} \end{bmatrix} = - \begin{bmatrix} \frac{\partial H}{\partial \theta} x^* + \frac{\partial g}{\partial \theta} + \frac{\partial C}{\partial \theta}^\top z^* \\ \frac{\partial C}{\partial \theta} x^* - \frac{\partial u}{\partial \theta} \\ 0 \\ \frac{\partial C}{\partial \theta}^\top [t^*]_+ \end{bmatrix} \end{aligned} \tag{21}$$

when this problem is feasible. When it is not feasible, the augmented Lagrangian naturally converges to the solution of the more general (20), see [Chiche and Gilbert, 2016, Proposition 4.2]. In comparison to (20), a notable advantage of the formulation (21) is that it is naturally numerically more stable and sparse—by avoiding square matrix products from the objective.

B.2.1.2 Backward pass. The following lemma formally details the results from Section 3.4.1.

Lemma 5. *Let $h : \mathbb{R}^n \times (\mathbb{R}^{n_i})^2 \rightarrow \mathbb{R}$ be a differentiable function, and let $H(\theta)$, $g(\theta)$, $C(\theta)$ and $u(\theta)$ be differentiable w.r.t. θ and satisfying Assumption 1. Then, denoting $\mathcal{L}(\theta) := h(x^*(\theta), z^*(\theta), s^*(\theta))$ and under assumptions of Lemma 3, we have that $\frac{\partial \mathcal{L}}{\partial \theta}$ can be derived as follows*

$$\frac{\partial \mathcal{L}}{\partial \theta} = (b_1^*)^\top \frac{\partial H}{\partial \theta} x^* + (b_1^*)^\top \frac{\partial g}{\partial \theta} + (b_2^*)^\top \frac{\partial C}{\partial \theta} x^* + (z^*)^\top \frac{\partial C}{\partial \theta} b_1^* + (s^*)^\top \frac{\partial C}{\partial \theta} b_4^* - (b_2^*)^\top \frac{\partial u}{\partial \theta},$$

where b_1^* , b_2^* , b_3^* and b_4^* are the solutions of the linear system

$$\begin{bmatrix} H & C^\top & 0 & 0 \\ C & 0 & (I - \Pi_1) & 0 \\ 0 & -I & -\Pi_1 \Pi_2 & (1 - \Pi_2)C \end{bmatrix} \begin{bmatrix} b_1^* \\ b_2^* \\ b_3^* \\ b_4^* \end{bmatrix} = - \begin{bmatrix} \frac{\delta \mathcal{L}}{\delta \theta} \\ \frac{\delta \mathcal{L}}{\delta x^*} \\ \frac{\delta \mathcal{L}}{\delta z^*} \\ \frac{\delta \mathcal{L}}{\delta s^*} \end{bmatrix}.$$

Proof. Under the assumptions of Lemma 3, it holds that $\begin{bmatrix} \frac{\partial x^*}{\partial \theta} \\ \frac{\partial z^*}{\partial \theta} \\ \frac{\partial s^*}{\partial \theta} \end{bmatrix}$ is a CJ as $\begin{bmatrix} \frac{\partial x^*}{\partial \theta} \\ \frac{\partial z^*}{\partial \theta} \\ \frac{\partial t^*}{\partial \theta} \end{bmatrix}$ and $\Pi_1 \frac{\partial t^*}{\partial \theta} \in \frac{\partial s^*}{\partial \theta}$ are CJs ([Bolte and Pauwels, 2020, Proposition 2]). Furthermore, as \mathcal{L} is differentiable w.r.t. x^* , z^* and s^* , it is path-differentiable [Bolte and Pauwels, 2020, Remark 3b] w.r.t. x^* , z^* , s^* , so we can apply chain rule ([Bolte and Pauwels, 2020, Proposition 2]):

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \theta} &= \frac{\partial \mathcal{L}}{\partial x^*}^\top \frac{\partial x^*}{\partial \theta} + \frac{\partial \mathcal{L}}{\partial z^*}^\top \frac{\partial z^*}{\partial \theta} + \frac{\partial \mathcal{L}}{\partial s^*}^\top \frac{\partial s^*}{\partial \theta} \\ &= - \left(- \begin{bmatrix} \frac{\delta \mathcal{L}}{\delta x^*} \\ \frac{\delta \mathcal{L}}{\delta z^*} \\ \frac{\delta \mathcal{L}}{\delta s^*} \end{bmatrix} \right)^\top \begin{bmatrix} \frac{\partial x^*}{\partial \theta} \\ \frac{\partial z^*}{\partial \theta} \\ \frac{\partial t^*}{\partial \theta} \end{bmatrix} \\ &= - \left(\begin{bmatrix} H & C^\top & 0 \\ C & 0 & -I \\ 0 & \Pi_1 - I & \Pi_1 \Pi_2 \\ 0 & 0 & C^\top (I - \Pi_2) \end{bmatrix} \begin{bmatrix} b_1^* \\ b_2^* \\ b_3^* \\ b_4^* \end{bmatrix} \right)^\top \begin{bmatrix} \frac{\partial x^*}{\partial \theta} \\ \frac{\partial z^*}{\partial \theta} \\ \frac{\partial t^*}{\partial \theta} \end{bmatrix} \\ &= - \begin{bmatrix} b_1^* \\ b_2^* \\ b_3^* \\ b_4^* \end{bmatrix}^\top \left(- \begin{bmatrix} \frac{\partial H}{\partial \theta} x^* + \frac{\partial g}{\partial \theta} + \frac{\partial C}{\partial \theta}^\top z^* \\ \frac{\partial C}{\partial \theta} x^* - \frac{\partial u}{\partial \theta} \\ 0 \\ \frac{\partial C}{\partial \theta}^\top [t^*]_+ \end{bmatrix} \right) \\ &= (b_1^*)^\top \left(\frac{\partial H}{\partial \theta} x^* + \frac{\partial g}{\partial \theta} + \frac{\partial C}{\partial \theta}^\top z^* \right) \\ &\quad + (b_2^*)^\top \left(\frac{\partial C}{\partial \theta} x^* - \frac{\partial u}{\partial \theta} \right) \\ &\quad + (b_4^*)^\top \left(\frac{\partial C}{\partial \theta}^\top [t^*]_+ \right) \\ &= (b_1^*)^\top \frac{\partial H}{\partial \theta} x^* + \left(\frac{\partial C}{\partial \theta} b_1^* \right)^\top z^* + (b_1^*)^\top \frac{\partial g}{\partial \theta} \\ &\quad + (b_2^*)^\top \left(\frac{\partial C}{\partial \theta} x^* - \frac{\partial u}{\partial \theta} \right) \\ &\quad + \left(\frac{\partial C}{\partial \theta} b_4^* \right)^\top [t^*]_+, \end{aligned}$$

with b_1^*, b_2^*, b_4^* solution of the nonsingular system (see Lemma 3)

$$\begin{bmatrix} H & C^\top & 0 & 0 \\ C & 0 & (I - \Pi_1) & 0 \\ 0 & -I & -\Pi_1 \Pi_2 & (1 - \Pi_2)C \end{bmatrix} \begin{bmatrix} b_1^* \\ b_2^* \\ b_3^* \\ b_4^* \end{bmatrix} = - \begin{bmatrix} \frac{\delta \mathcal{L}}{\delta x^*} \\ \frac{\delta \mathcal{L}}{\delta z^*} \\ \frac{\delta \mathcal{L}}{\delta s^*} \end{bmatrix},$$

thereby reaching the desired statement with $s^* = [t^*]_+$. \square

B.2.2 Simplification of when QP is feasible

B.2.2.1 Simplification of the forward pass When QP(θ) is feasible and $H(\theta)$, $g(\theta)$, $C(\theta)$ and $u(\theta)$ are differentiable w.r.t. θ and satisfy Assumption 1, then $[t^*]_+ = 0$. Further, our choices of subgradients at zero (see (19)) imply that $\Pi_2 = I$ and hence the following simplifications

$$\begin{bmatrix} H & C^\top & 0 \\ C & 0 & -I \\ 0 & \Pi_1 - I & \Pi_1 \\ 0 & 0 & 0 \end{bmatrix} \in \frac{\partial G(x^*, z^*, t^*; \theta)}{\partial v^*},$$

$$\begin{bmatrix} \frac{\partial H}{\partial \theta} x^* + \frac{\partial g}{\partial \theta} + \frac{\partial C}{\partial \theta}^\top z^* \\ \frac{\partial C}{\partial \theta} x^* - \frac{\partial u}{\partial \theta} \\ 0 \\ 0 \end{bmatrix} = \frac{\partial G(x^*, z^*, t^*; \theta)}{\partial \theta}.$$

Moreover, the optimality conditions of

$$\frac{\partial x^*}{\partial \theta}, \frac{\partial z^*}{\partial \theta}, \frac{\partial t^*}{\partial \theta} \in \arg \min_{\frac{\partial x}{\partial \theta}, \frac{\partial z}{\partial \theta}, \frac{\partial t}{\partial \theta}} \left\| \begin{bmatrix} H & C^\top & 0 \\ C & 0 & -I \\ 0 & \Pi_1 - I & \Pi_1 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \frac{\partial x}{\partial \theta} \\ \frac{\partial z}{\partial \theta} \\ \frac{\partial t}{\partial \theta} \end{bmatrix} + \begin{bmatrix} \frac{\partial H}{\partial \theta} x^* + \frac{\partial g}{\partial \theta} + \frac{\partial C}{\partial \theta}^\top z^* \\ \frac{\partial C}{\partial \theta} x^* - \frac{\partial u}{\partial \theta} \\ 0 \\ 0 \end{bmatrix} \right\|_2^2,$$

write down

$$\begin{aligned} (H^2 + C^\top C) \frac{\partial x^*}{\partial \theta} + HC^\top \frac{\partial z^*}{\partial \theta} - C^\top \frac{\partial t^*}{\partial \theta} + [H(\frac{\partial H}{\partial \theta} x^* + \frac{\partial g}{\partial \theta} + \frac{\partial C}{\partial \theta}^\top z^*) + C^\top(\frac{\partial C}{\partial \theta} x^* - \frac{\partial u}{\partial \theta})] &= 0, \\ CH \frac{\partial x^*}{\partial \theta} + (C^\top C + I - \Pi_1) \frac{\partial z^*}{\partial \theta} + C(\frac{\partial H}{\partial \theta} x^* + \frac{\partial g}{\partial \theta} + \frac{\partial C}{\partial \theta}^\top z^*) &= 0, \\ -C \frac{\partial x^*}{\partial \theta} + (I + \Pi_1) \frac{\partial t^*}{\partial \theta} - (\frac{\partial C}{\partial \theta} x^* - \frac{\partial u}{\partial \theta}) &= 0. \end{aligned} \tag{22}$$

Third equation of (22) leads to

$$\frac{\partial t^*}{\partial \theta} = \frac{1}{1 + \Pi_1} (C \frac{\partial x^*}{\partial \theta} + (\frac{\partial C}{\partial \theta} x^* - \frac{\partial u}{\partial \theta})),$$

Hence, optimality conditions without variable $\frac{\partial t^*}{\partial \theta}$ reduce to

$$\begin{aligned} (H^2 + C^\top \frac{\Pi_1}{1 + \Pi_1} C) \frac{\partial x^*}{\partial \theta} + HC^\top \frac{\partial z^*}{\partial \theta} + C^\top \frac{\Pi_1}{1 + \Pi_1} (\frac{\partial C}{\partial \theta} x^* - \frac{\partial u}{\partial \theta}) + H(\frac{\partial H}{\partial \theta} x^* + \frac{\partial g}{\partial \theta} + \frac{\partial C}{\partial \theta}^\top z^*) &= 0, \\ CH \frac{\partial x^*}{\partial \theta} + (C^\top C + I - \Pi_1) \frac{\partial z^*}{\partial \theta} + C(\frac{\partial H}{\partial \theta} x^* + \frac{\partial g}{\partial \theta} + \frac{\partial C}{\partial \theta}^\top z^*) &= 0. \end{aligned} \tag{23}$$

Furthermore, the following problem

$$\min_{\frac{\partial x}{\partial \theta}, \frac{\partial z}{\partial \theta}} \left\| \begin{bmatrix} H & C^\top \\ \frac{\Pi_1}{\sqrt{1+\Pi_1}} C & \Pi_1 - I \end{bmatrix} \begin{bmatrix} \frac{\partial x}{\partial \theta} \\ \frac{\partial z}{\partial \theta} \end{bmatrix} + \begin{bmatrix} \frac{\partial H}{\partial \theta} x^* + \frac{\partial g}{\partial \theta} + \frac{\partial C}{\partial \theta}^\top z^* \\ \frac{\Pi_1}{\sqrt{1+\Pi_1}} \left(\frac{\partial C}{\partial \theta} x^* - \frac{\partial u}{\partial \theta} \right) \end{bmatrix} \right\|_2^2,$$

have the same KKT conditions as (23), thereby allowing to simplify the problem as follows:

$$\begin{aligned} & \min_{\frac{\partial x}{\partial \theta}, \frac{\partial z}{\partial \theta}, \frac{\partial t}{\partial \theta}} \left\| \frac{\partial G(x^*, z^*, s^*; \theta)}{\partial \hat{v}^*} \begin{bmatrix} \frac{\partial x}{\partial \theta} \\ \frac{\partial z}{\partial \theta} \\ \frac{\partial t}{\partial \theta} \end{bmatrix} + \frac{\partial G(x^*, z^*, s^*; \theta)}{\partial \theta} \right\|_2^2 \\ &= \min_{\frac{\partial x}{\partial \theta}, \frac{\partial z}{\partial \theta}} \left\| \begin{bmatrix} H & C^\top \\ \frac{\Pi_1}{\sqrt{1+\Pi_1}} C & \Pi_1 - I \end{bmatrix} \begin{bmatrix} \frac{\partial x}{\partial \theta} \\ \frac{\partial z}{\partial \theta} \end{bmatrix} + \begin{bmatrix} \frac{\partial H}{\partial \theta} x^* + \frac{\partial g}{\partial \theta} + \frac{\partial C}{\partial \theta}^\top z^* \\ \frac{\Pi_1}{\sqrt{1+\Pi_1}} \left(\frac{\partial C}{\partial \theta} x^* - \frac{\partial u}{\partial \theta} \right) \end{bmatrix} \right\|_2^2, \end{aligned} \quad (24)$$

Hence

$$\frac{\partial x^*}{\partial \theta}, \frac{\partial z^*}{\partial \theta}, \frac{\partial t^*}{\partial \theta} \in \arg \min_{\frac{\partial x}{\partial \theta}, \frac{\partial z}{\partial \theta}, \frac{\partial t}{\partial \theta}} \left\| \frac{\partial G(x^*, z^*, s^*; \theta)}{\partial \hat{v}^*} \begin{bmatrix} \frac{\partial x}{\partial \theta} \\ \frac{\partial z}{\partial \theta} \\ \frac{\partial t}{\partial \theta} \end{bmatrix} + \frac{\partial G(x^*, z^*, s^*; \theta)}{\partial \theta} \right\|_2^2$$

is equivalent to

$$\begin{aligned} \frac{\partial x^*}{\partial \theta}, \frac{\partial z^*}{\partial \theta} &\in \arg \min_{\frac{\partial x}{\partial \theta}, \frac{\partial z}{\partial \theta}} \left\| \begin{bmatrix} H & C^\top \\ \frac{\Pi_1}{\sqrt{1+\Pi_1}} C & \Pi_1 - I \end{bmatrix} \begin{bmatrix} \frac{\partial x}{\partial \theta} \\ \frac{\partial z}{\partial \theta} \end{bmatrix} + \begin{bmatrix} \frac{\partial H}{\partial \theta} x^* + \frac{\partial g}{\partial \theta} + \frac{\partial C}{\partial \theta}^\top z^* \\ \frac{\Pi_1}{\sqrt{1+\Pi_1}} \left(\frac{\partial C}{\partial \theta} x^* - \frac{\partial u}{\partial \theta} \right) \end{bmatrix} \right\|_2^2, \\ \frac{\delta t^*}{\delta \theta} &= (I + \Pi_1)^{-1} \left(C \frac{\delta x^*}{\delta \theta} + \frac{\partial C}{\partial \theta} x^* - \frac{\partial u}{\partial \theta} \right). \end{aligned}$$

B.2.2.2 Simplification of the backward pass. This sections details the results from Section 3.4.2.

Lemma 6. *Let $h : \mathbb{R}^n \times (\mathbb{R}^{n_i}) \rightarrow \mathbb{R}$ be a differentiable function, and let $H(\theta)$, $g(\theta)$, $C(\theta)$ and $u(\theta)$ be differentiable w.r.t. θ and satisfying Assumption 1. Then, denoting $\mathcal{L}(\theta) := h(x^*(\theta), z^*(\theta))$, we have under assumptions of Lemma 4 that $\frac{\partial \mathcal{L}}{\partial \theta}$ can be derived as follows*

$$\frac{\partial \mathcal{L}}{\partial \theta} = (b_x^*)^\top \frac{\partial H}{\partial \theta} x^* + (b_x^*)^\top \frac{\partial g}{\partial \theta} + (\Pi_1 b_z^*)^\top \frac{\partial C}{\partial \theta} x^* + (z^*)^\top \frac{\partial C}{\partial \theta} b_x^* - (\Pi_1 b_z^*)^\top \frac{\partial u}{\partial \theta},$$

with b_x^* , b_z^* , the solution of the following linear system

$$\begin{bmatrix} H & C^\top \Pi_1 \\ C & -(I - \Pi_1) \end{bmatrix} \begin{bmatrix} b_x \\ b_z \end{bmatrix} = - \begin{bmatrix} \frac{\delta \mathcal{L}}{\delta x^*} \\ \frac{\delta \mathcal{L}}{\delta z^*} \end{bmatrix},$$

Furthermore, this latter linear system can be solved using iterative refinement.

Proof. Under the assumptions of Lemma 4, it holds that $\begin{bmatrix} \frac{\partial x^*}{\partial \theta} \\ \frac{\partial z^*}{\partial \theta} \end{bmatrix}$ is a Jacobian. Furthermore, as \mathcal{L} is differentiable, the chain rule implies that:

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial \theta} &= \begin{bmatrix} \frac{\partial \mathcal{L}}{\partial x^*} \\ \frac{\partial \mathcal{L}}{\partial z^*} \end{bmatrix}^\top \begin{bmatrix} \frac{\partial x^*}{\partial \theta} \\ \frac{\partial z^*}{\partial \theta} \end{bmatrix} \\
&= - \left(- \begin{bmatrix} \frac{\partial \mathcal{L}}{\partial x^*} \\ \frac{\partial \mathcal{L}}{\partial z^*} \end{bmatrix} \right)^\top \begin{bmatrix} \frac{\partial x^*}{\partial \theta} \\ \frac{\partial z^*}{\partial \theta} \end{bmatrix} \\
&= - \left(\begin{bmatrix} H & C^\top \\ \Pi_1 C & \Pi_1 - I \end{bmatrix} \begin{bmatrix} b_x \\ b_z \end{bmatrix} \right)^\top \begin{bmatrix} \frac{\partial x^*}{\partial \theta} \\ \frac{\partial z^*}{\partial \theta} \end{bmatrix} \text{ as the matrix is nonsingular (see Lemma 4)} \\
&= - \begin{bmatrix} b_x \\ b_z \end{bmatrix}^\top \left(- \begin{bmatrix} \frac{\partial H}{\partial \theta} x^* + \frac{\partial g}{\partial \theta} + \frac{\partial C}{\partial \theta}^\top z^* \\ \Pi_1 \left(\frac{\partial C}{\partial \theta} x^* - \frac{\partial u}{\partial \theta} \right) \end{bmatrix} \right) \\
&= (b_x)^\top \left(\frac{\partial H}{\partial \theta} x^* + \frac{\partial g}{\partial \theta} + \frac{\partial C}{\partial \theta}^\top z^* \right) \\
&\quad + (b_z)^\top (\Pi_1 (dC x^* - du)) \\
&= [(b_x)^\top \frac{\partial H}{\partial \theta} x^* + (\frac{\partial C}{\partial \theta} b_x)^\top z^* + (b_x)^\top \frac{\partial g}{\partial \theta}] \\
&\quad + (b_z)^\top (\Pi_1 (\frac{\partial C}{\partial \theta} x^* - \frac{\partial u}{\partial \theta})),
\end{aligned}$$

where (b_x, b_z) is a solution to

$$\begin{bmatrix} H & C^\top \Pi_1 \\ C & -(I - \Pi_1) \end{bmatrix} \begin{bmatrix} b_x \\ b_z \end{bmatrix} = - \begin{bmatrix} \frac{\delta \mathcal{L}}{\delta x^*} \\ \frac{\delta \mathcal{L}}{\delta z^*} \end{bmatrix}.$$

As detailed in the proof of Lemma 4 (see details in Appendix B.4), one can equivalently solve

$$\begin{bmatrix} H & C_J^\top \\ C_J & 0 \end{bmatrix} \begin{bmatrix} b_x \\ b_{z_J} \end{bmatrix} = - \begin{bmatrix} \frac{\delta \mathcal{L}}{\delta x^*} \\ \frac{\delta \mathcal{L}}{\delta z_J^*} \end{bmatrix},$$

$$b_{z_J^c} = \frac{\delta \mathcal{L}}{\delta z_{J^c}^*},$$

with J^c the index set for which the solution is strictly feasible (i.e., $i \in [1, n_i], (\Pi_1)_i = 0$), and J the set of active constraints (i.e., for which $(\Pi_1)_i = 1$). Such linear systems can be solved e.g., via iterative refinement (as the matrix involved is symmetric positive semi-definite [Parikh and Boyd, 2014, Section 4.1.2]). \square

B.3 Proof of Lemma 3

This section provides a proof for Lemma 3, ensuring that ECJs are CJs under some regularity assumptions.

Lemma 3. *Let $H(\theta)$, $g(\theta)$, $C(\theta)$ and $u(\theta)$ be differentiable w.r.t. θ and satisfying Assumption 1. If $s^* > 0$ and x^* and z^* satisfy strict complementarity and C is nonsingular, then the ECJs of x^* , z^* , t^* , and s^* correspond to conservative Jacobians.*

Proof. If $s^* > 0$, then $\Pi_2 = 0$. Furthermore, if x^* and z^* satisfy strict complementarity then $\Pi_1 = I$. Hence, ECJs can be computed as least-square solutions to

$$\underbrace{\begin{bmatrix} H & C^\top & 0 \\ C & 0 & -I \\ 0 & 0 & C^\top \end{bmatrix}}_{:=\Delta} \begin{bmatrix} \frac{\partial x^*}{\partial \theta} \\ \frac{\partial z^*}{\partial \theta} \\ \frac{\partial t^*}{\partial \theta} \end{bmatrix} = - \begin{bmatrix} \frac{\partial H}{\partial \theta} x^* + \frac{\partial g}{\partial \theta} + \frac{\partial C}{\partial \theta}^\top z^* \\ \frac{\partial C}{\partial \theta} x^* - \frac{\partial u}{\partial \theta} \\ \frac{\partial C}{\partial \theta}^\top [t^*]_+ \end{bmatrix}.$$

If C is nonsingular, then $\frac{\partial t^*}{\partial \theta}$ is uniquely defined through equation

$$C^\top \frac{\partial t^*}{\partial \theta} = \frac{\partial C}{\partial \theta}^\top [t^*]_+.$$

Furthermore, nonsingularity of C implies nonsingularity of the resulting system

$$\begin{bmatrix} H & C^\top \\ C & 0 \end{bmatrix} \begin{bmatrix} \frac{\partial x^*}{\partial \theta} \\ \frac{\partial z^*}{\partial \theta} \end{bmatrix} = - \begin{bmatrix} \frac{\partial H}{\partial \theta} x^* + \frac{\partial g}{\partial \theta} + \frac{\partial C}{\partial \theta}^\top z^* \\ \frac{\partial C}{\partial \theta} x^* - \frac{\partial u}{\partial \theta} - \frac{\partial t^*}{\partial \theta} \end{bmatrix}.$$

Thus Δ is nonsingular and the path-differentiable implicit function theorem [Bolte and Pauwels, 2020, Corollary 1]

can be applied to $G(x^*, z^*, s^*; \theta)$. Hence $\begin{bmatrix} \frac{\partial x^*}{\partial \theta} \\ \frac{\partial z^*}{\partial \theta} \\ \frac{\partial t^*}{\partial \theta} \end{bmatrix}$ corresponds to a CJ. \square

B.4 Proof of Lemma 4

This section details the proof of Lemma 4, ensuring that, under some regularity assumptions, ECJs reduce to standard Jacobians.

Lemma 4. *If $QP(\theta)$ is feasible and $H(\theta)$, $g(\theta)$, $C(\theta)$ and $u(\theta)$ are differentiable w.r.t. θ and satisfy Assumption 1, and if the KKT matrix of active constraints is nonsingular and x^* , z^* satisfy strict complementarity, then the ECJs matches the standard Jacobian, i.e., $\frac{\partial x^*(\theta)}{\partial \theta} = \nabla x^*(\theta)$ and $\frac{\partial z^*(\theta)}{\partial \theta} = \nabla z^*(\theta)$.*

Proof. If $(QP(\theta))$ is feasible, then the ECJs of x^* and z^* w.r.t. θ are provided by

$$\frac{\partial x^*}{\partial \theta}, \frac{\partial z^*}{\partial \theta} \in \arg \min_{\frac{\partial x}{\partial \theta}, \frac{\partial z}{\partial \theta}} \left\| \underbrace{\begin{bmatrix} H & C^\top \\ \frac{\Pi_1}{\sqrt{1+\Pi_1}} C & \Pi_1 - I \end{bmatrix}}_{:=\Delta} \begin{bmatrix} \frac{\partial x}{\partial \theta} \\ \frac{\partial z}{\partial \theta} \end{bmatrix} + \begin{bmatrix} \frac{\partial H}{\partial \theta} x^* + \frac{\partial g}{\partial \theta} + \frac{\partial C}{\partial \theta}^\top z^* \\ \frac{\Pi_1}{\sqrt{1+\Pi_1}} \left(\frac{\partial C}{\partial \theta} x^* - \frac{\partial u}{\partial \theta} \right) \end{bmatrix} \right\|_2^2, \quad (25)$$

where Π_1 corresponds to a binary diagonal matrix of the complementarity conditions $Cx^* - u + z^* \geq 0$. Denoting by J^c the index set for which the solution is strictly feasible (i.e., $i \in [1, n_i]$, $(\Pi_1)_i = 0$), and by J the index set of active constraints (i.e., for which $(\Pi_1)_i = 1$) then Δ can be reformulated as follows (by strict complementary)

$$\Delta = \begin{bmatrix} H & C_J^\top & C_{J^c}^\top \\ \frac{1}{\sqrt{2}} C_J & 0 & 0 \\ 0 & 0 & -I \end{bmatrix},$$

with I being the identity matrix of appropriate dimension. Furthermore, the right-hand side of the linear system within the ℓ_2 norm becomes

$$\begin{bmatrix} \frac{\partial H}{\partial \theta} x^* + \frac{\partial g}{\partial \theta} + \frac{\partial C}{\partial \theta}^\top z^* \\ \frac{1}{\sqrt{2}} \left(\frac{\partial C_J}{\partial \theta} x^* - \frac{\partial u_{J^c}}{\partial \theta} \right) \\ 0 \end{bmatrix}.$$

$\begin{bmatrix} H & C_J^\top \\ C_J & 0 \end{bmatrix}$ corresponds to the KKT matrix of active constraints, and is nonsingular by assumption. As it implies nonsingularity of $\begin{bmatrix} H & C_J^\top \\ \frac{1}{\sqrt{2}}C_J & 0 \end{bmatrix}$, it follows that :

$$\frac{\partial z_{J^c}^*}{\partial \theta} = 0,$$

and the solution to (25) is uniquely determined as the solution of the following linear system (as in [Amos and Kolter, 2017, Appendix A], after multiplying second row block by $\sqrt{2}$):

$$\begin{bmatrix} H & C_J^\top \\ C_J & 0 \end{bmatrix} \begin{bmatrix} \frac{\partial x^*}{\partial \theta} \\ \frac{\partial z^*}{\partial \theta} \end{bmatrix} = - \begin{bmatrix} \frac{\partial H}{\partial \theta} x^* + \frac{\partial g}{\partial \theta} + \frac{\partial C}{\partial \theta}^\top z^* \\ \frac{\partial C_J}{\partial \theta} x^* - \frac{\partial u_J}{\partial \theta} \end{bmatrix},$$

Hence, we arrive at the desired conclusion that ECJ coincides with the usual Jacobian in this case. □

C Additional experimental results

Appendix C.1 provides a few simple experiments with parametric QPs to illustrate the concept of ECJs. Appendix C.2 contains additional benchmarks. Appendix C.3 illustrates through several experiments that QPLayer can be trained with large learning rates.

C.1 Pedagogical examples of parametric QPs

A few numerical examples illustrate the concept of ECJ in different simple scenarios. The first example corresponds to a strictly convex parametric QP which can be either feasible or infeasible. In this example, a linear constraint depends on a parameter θ . Depending on the value of this parameter, the QP can either be feasible or infeasible.

The second example is a strictly convex QP with a parameterized linear cost. This problem is always feasible.

The last example is a parametric LP with possibly multiple solutions. For appropriate values of the parameters, the LP is feasible but not differentiable.

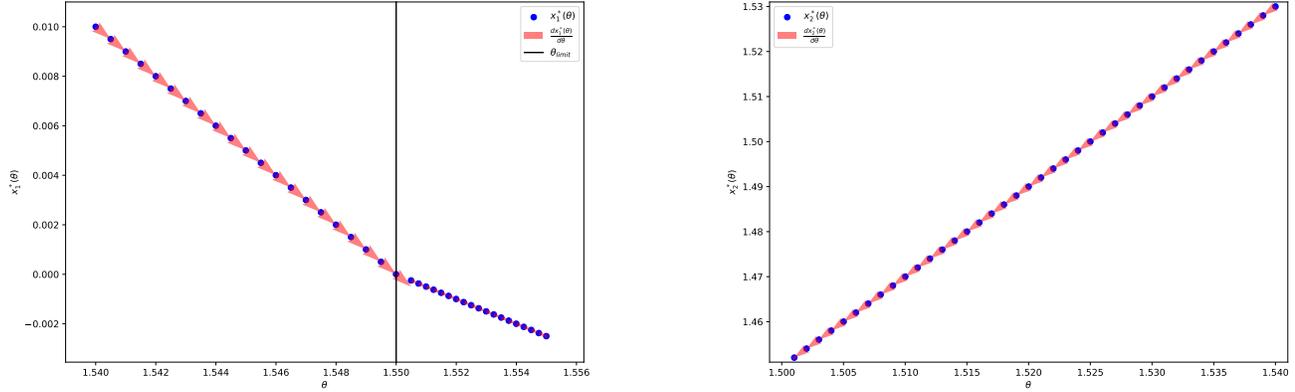
C.1.1 Strictly convex QP (parameterized constraints)

Consider the following strictly convex QP parameterized by a scalar value θ

$$\begin{aligned} x^*(\theta) = \arg \min_{x_1, x_2 \in \mathbb{R}^2} & \frac{1}{2}(x_1^2 + x_2^2) \\ \text{s.t. } & \theta \leq x_1 + x_2 \leq 1.55, \\ & 1.5 \leq 2x_1 + x_2 \leq 1.55 \end{aligned} \tag{26}$$

Notice that for $\theta > \theta_{\text{limit}} := 1.55$, the QP becomes primal infeasible. We use gradient descent to minimize two scalar losses $\mathcal{L}_1(\theta) = x_1^*(\theta)$ and $\mathcal{L}_2(\theta) = x_2^*(\theta)$, starting from a predefined value θ_0 . More precisely we have launched gradient descent for 40 steps with a learning rate 5×10^{-4} starting from $\theta_0 = 1.54$. Figure 3a illustrates the results by showing the iterates of gradient descent for minimizing $x_1^*(\theta)$ (as well as the search direction—minus the ECJs). By doing so θ increases and eventually becomes larger than θ_{limit} .

Figure 3b reports a similar experiment for when minimizing $x_2^*(\theta)$.



(a) 40 steps of gradient descent for minimizing $x_1^*(\theta)$ starting from $\theta_0 = 1.54$. When $\theta > 1.55$, (26) is differentiated though infeasible.

(b) 40 steps of gradient descent for minimizing $x_2^*(\theta)$ starting from $\theta_0 = 1.54$. The QPs remain feasible.

C.1.2 Strictly convex QP (parameterized objective)

Consider the following strictly convex QP parametrized by a scalar value θ

$$\begin{aligned}
 x^*(\theta) = \arg \min_{x_1, x_2 \in \mathbb{R}^2} & \frac{1}{2} \left\| \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + \begin{bmatrix} \theta \\ -2 \end{bmatrix} \right\|_2^2 \\
 \text{s.t.} & -300 \leq x_1 + x_2 \leq 400, \\
 & -200 \leq 2x_1 + x_2 \leq 500
 \end{aligned} \tag{27}$$

We use gradient descent to minimize the loss $\mathcal{L}_1(\theta) = x_1^*(\theta)$. More precisely, we run 40 iterations of gradient descent with learning rate 5×10^{-4} starting from $\theta_0 = 1.54$, as reported by Figure 4. As expected, we see that $x_1^*(\theta) = -\theta$, hence increasing θ decreases $x_1^*(\theta)$.

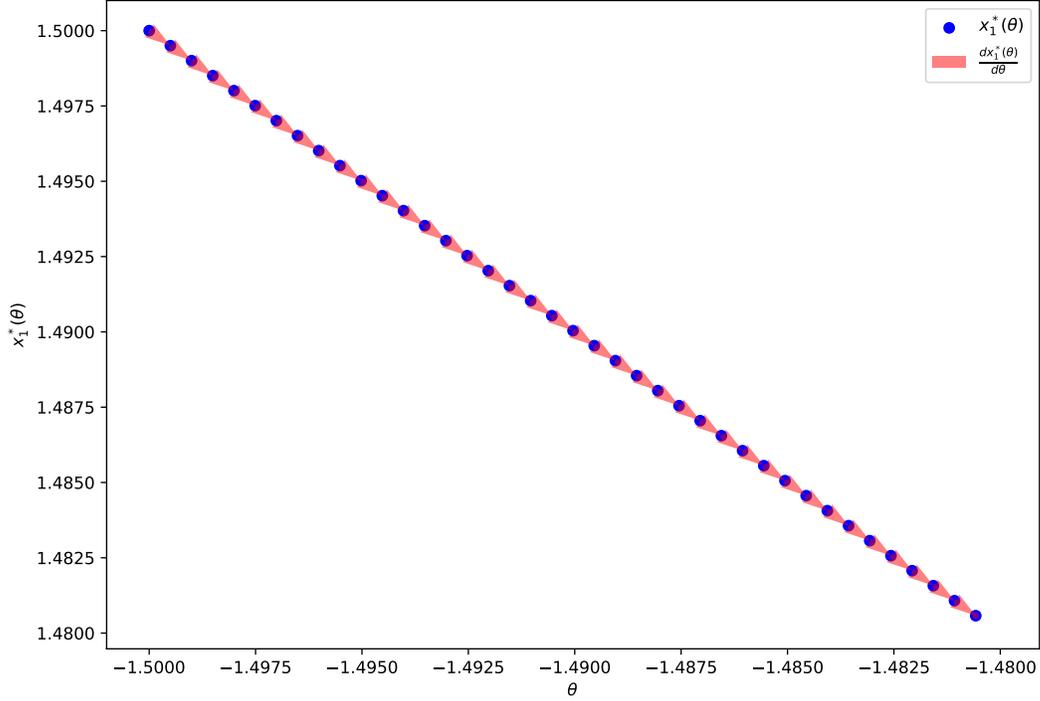


Figure 4: Gradient descent for minimizing $x_1^*(\theta)$ (solution to (27)) starting from $\theta_0 = -1.5$.

C.1.3 Parameterized linear program

Consider the following LP parameterized by a scalar parameter $\theta > 0$

$$\begin{aligned}
 x^*(\theta) \in \arg \min_{x_1, x_2 \in \mathbb{R}^2} x_1 + x_2 \\
 \text{s.t. } \theta &\leq x_1 + x_2, \\
 0 &\leq x_1 \leq 1, \\
 0 &\leq x_2 \leq 1.
 \end{aligned} \tag{28}$$

Note that this LP is always well-defined for any θ since the linear cost is orthogonal to the recession cone (which is empty), thereby satisfying the technical requirements from Assumption 1. We use gradient descent for minimizing two scalar losses $\mathcal{L}_1(\theta) = x_1^*(\theta)$ and $\mathcal{L}_2(\theta) = x_2^*(\theta)$.

C.1.3.1 Feasible case ($\theta \leq 2$). For any $\theta \in]0, 2]$, there are infinitely many solutions to (28) which are defined by the segment equation

$$\begin{aligned}
 x_1^* + x_2^* &= \theta, \\
 0 &\leq x_1^* \leq 1, \\
 0 &\leq x_2^* \leq 1.
 \end{aligned}$$

We can see in Figure 5a and Figure 5b that the forward pass chooses as solution $x_1^* = x_2^* = \frac{\theta}{2}$. Hence, only the constraint $\theta \leq x_1 + x_2$ is active. Following the formalism from Section 3.1 we have

$$C = \begin{bmatrix} -1 & -1 \\ 1 & 0 \\ -1 & 0 \\ 0 & 1 \\ 0 & -1 \end{bmatrix}, \quad u = \begin{bmatrix} -\theta \\ 1 \\ 0 \\ 1 \\ 0 \end{bmatrix}.$$

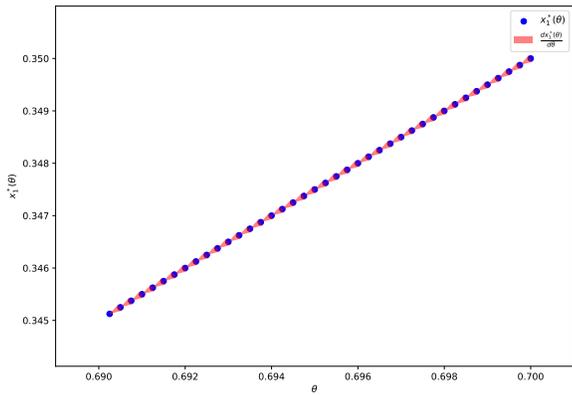
The ECJs of \mathcal{L}_1 and \mathcal{L}_2 w.r.t θ are the solutions to

$$\begin{aligned} \begin{bmatrix} (b_x^*)_1 \\ (b_x^*)_2 \\ b_z^* \end{bmatrix} &\in \arg \min_{b_x, b_z} \left\| \begin{bmatrix} 0 & 0 & -1 \\ 0 & 0 & -1 \\ -1 & -1 & 0 \end{bmatrix} \begin{bmatrix} (b_x)_1 \\ (b_x)_2 \\ b_z \end{bmatrix} + \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \right\|_2^2, \\ \begin{bmatrix} (d_x^*)_1 \\ (d_x^*)_2 \\ d_z^* \end{bmatrix} &\in \arg \min_{d_x, d_z} \left\| \begin{bmatrix} 0 & 0 & -1 \\ 0 & 0 & -1 \\ -1 & -1 & 0 \end{bmatrix} \begin{bmatrix} (d_x)_1 \\ (d_x)_2 \\ d_z \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} \right\|_2^2. \end{aligned}$$

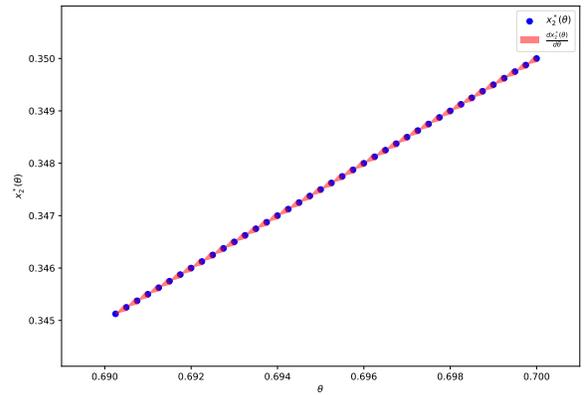
As the corresponding linear systems involved within the ℓ_2 norm are infeasible, the least square estimates do not correspond to solutions to the linear system. That is, the corresponding least-square solutions are respectively the solutions of the following projected linear systems:

$$\begin{aligned} \begin{bmatrix} 1 & 1 & 0 \\ 1 & 1 & 0 \\ 0 & 0 & 2 \end{bmatrix} \begin{bmatrix} (b_x^*)_1 \\ (b_x^*)_2 \\ b_z^* \end{bmatrix} &= \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}, \\ \begin{bmatrix} 1 & 1 & 0 \\ 1 & 1 & 0 \\ 0 & 0 & 2 \end{bmatrix} \begin{bmatrix} (d_x^*)_1 \\ (d_x^*)_2 \\ d_z^* \end{bmatrix} &= \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}, \end{aligned}$$

which leads to $\frac{\partial \mathcal{L}_1}{\partial \theta} = \frac{\partial \mathcal{L}_2}{\partial \theta} = b_z^* = d_z^* = \frac{1}{2}$. Figure 5a and Figure 5b show that those directions allow minimizing \mathcal{L}_1 and \mathcal{L}_2 through gradient descent.



(a) 40 iterations of gradient descent for minimizing $x_1^*(\theta)$ for the problem (28).



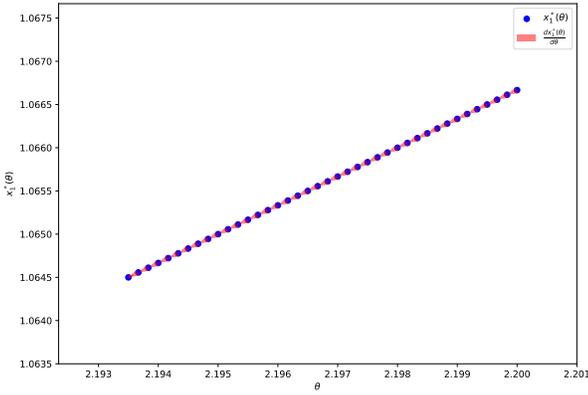
(b) 40 iterations of gradient descent for minimizing $x_2^*(\theta)$ for the problem (28).

C.1.3.2 Infeasible case ($\theta > 2$). For any $\theta > 2$, the LP is infeasible. The corresponding ECJs are the least-square solutions to

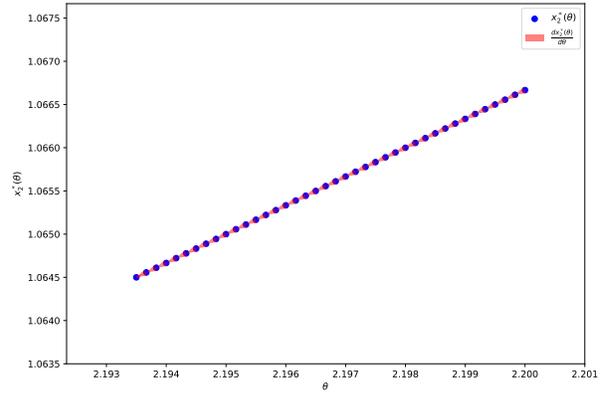
$$\arg \min_{b_1, b_2, b_3, b_4} \left\| \begin{bmatrix} 0 & C^\top & 0 & 0 \\ C & 0 & (I - \Pi_1) & 0 \\ 0 & -I & -\Pi_1 \Pi_2 & (1 - \Pi_2)C \end{bmatrix} \begin{bmatrix} b_1 \\ b_2 \\ b_3 \\ b_4 \end{bmatrix} + \begin{bmatrix} \frac{\delta \mathcal{L}_i}{\delta x} \\ 0 \\ 0 \end{bmatrix} \right\|_2^2 \text{ for } i \in \{1, 2\},$$

with $P_1 = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$ and $P_2 = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}$. The linear system within the ℓ_2 norm is feasible and

QPLayer outputs as ECJs $\frac{\partial \mathcal{L}_1}{\partial \theta} = \frac{\partial \mathcal{L}_2}{\partial \theta} = \frac{1}{3}$. Figure 6a and Figure 6b show that following such directions allows using gradient descent for minimizing \mathcal{L}_1 and \mathcal{L}_2 .



(a) 40 steps of gradient descent applied to minimize $x_1^*(\theta)$ starting from $\theta_0 = 2.2$, when considering the infeasible LP defined by (28).



(b) 40 steps of gradient descent applied to minimize $x_2^*(\theta)$ starting from $\theta_0 = 2.2$, when considering the infeasible LP defined by (28).

C.2 Additional benchmarks

This section reports two additional benchmarks. We outline detailed timings for differentiating solutions on a set of different QPs. First, Table 4 shows the results of our first random QP benchmark with a batch size of 1 (see Section 4.1.1). Second, Table 5 reports the average time spent for executing the forward pass on another set of randomly generated QPs proposed by [Amos and Kolter, 2017].

Table 4 shows that QPLayer is about 4 times faster than OptNet (the second fastest approach) for all target accuracies. Table 5 illustrates that QPLayer is from 5 to 9 times faster than OptNet (also the second fastest approach) for running the forward pass.

Forward tolerance ϵ	10^{-1}	10^{-2}	10^{-3}
Forward (ms)	1.30 (± 0.17)	1.42 (± 0.19)	1.55 (± 0.20)
Backward (ms)	0.77 (± 0.01)	0.70 (± 0.02)	0.71 (± 0.02)
QPLayer total (ms)	2.07	2.12	2.26
Forward (ms)	6.93 (± 0.82)	6.85 (± 0.05)	7.49 (± 0.04)
Backward (ms)	0.75 (± 12)	0.70 (± 0.01)	0.70 (± 0.01)
OptNet total (ms)	7.68	7.55	8.19
Forward (ms)	7.99 (± 0.93)	12.82 (± 2.67)	18.82 (± 4.31)
Backward (ms)	24.71 (± 2.53)	30.47 (± 2.71)	36.43 (± 3.93)
JaxOpt total (ms)	32.70	43.29	55.25
Forward (ms)	411.20 (± 3.90)	415.94 (± 2.63)	422.28 (± 2.65)
Backward (ms)	6.13 (± 0.02)	6.34 (± 0.33)	6.26 (± 0.06)
CvxpyLayer total (ms)	417.33	422.24	428.54
Forward (ms)	$6.00 (\pm 0.85) \times 10^3$	$38.66 (\pm 7.11) \times 10^3$	$114.20 (\pm 31.48) \times 10^3$
Backward (ms)	0.99 (± 0.21)	0.98 (± 0.13)	1.05 (± 0.14)
Alt-Diff total (ms)	6.00×10^3	36.66×10^3	114.20×10^3

Table 4: Averaged timings for the forward and backward passes when computing all Jacobians of randomly generated feasible QPs (with 100 primal variables, 50 equality constraints and 50 inequality constraints) considering different forward pass accuracies. The batch size is 1. QPLayer has the best total timings for all accuracies.

QP and Batch sizes	QPLayer (ms)	OptNet (ms)	CvxpyLayer (ms)
Batch = 1, $n = 100$, $n_e = 0$, $n_i = 50$	0.88 (± 0.06)	8.01 (± 0.89)	334.06 (± 7.83)
Batch = 1, $n = 100$, $n_e = 50$, $n_i = 50$	1.09 (± 0.10)	8.35 (± 0.67)	406.18 (± 1.44)
Batch = 1, $n = 100$, $n_e = 0$, $n_i = 100$	1.22 (± 0.08)	7.62 (± 0.32)	422.59 (± 4.34)
Batch = 1, $n = 100$, $n_e = 50$, $n_i = 100$	1.75 (± 0.17)	13.95 (± 2.03)	522.80 (± 3.96)
Batch = 1, $n = 100$, $n_e = 100$, $n_i = 50$	1.26 (± 0.14)	16.77 (± 3.41)	521.18 (± 16.15)
Batch = 1, $n = 100$, $n_e = 100$, $n_i = 100$	1.45 (± 0.28)	18.33 (± 5.35)	621.25 (± 9.67)
Batch = 64, $n = 100$, $n_e = 0$, $n_i = 50$	30.16 (± 5.51)	173.95 (± 12.82)	771.49 (± 59.74)
Batch = 64, $n = 100$, $n_e = 50$, $n_i = 50$	48.60 (± 9.19)	183.84 (± 9.99)	818.75 (± 5.42)
Batch = 64, $n = 100$, $n_e = 0$, $n_i = 100$	44.68 (± 5.95)	176.88 (± 2.30)	820.85 (± 12.38)
Batch = 64, $n = 100$, $n_e = 50$, $n_i = 100$	55.77 (± 6.38)	243.15 (± 27.86)	1126.99 (± 34.33)
Batch = 64, $n = 100$, $n_e = 100$, $n_i = 50$	51.66 (± 8.51)	231.11 (± 13.75)	1177.89 (± 132.99)
Batch = 64, $n = 100$, $n_e = 100$, $n_i = 100$	62.52 (± 7.84)	276.18 (± 31.29)	1314.47 (± 69.60)
Batch = 128, $n = 100$, $n_e = 0$, $n_i = 50$	62.94 (± 9.52)	620.91 (± 3.52)	1202.61 (± 88.32)
Batch = 128, $n = 100$, $n_e = 50$, $n_i = 50$	78.01 (± 4.06)	667.55 (± 5.86)	1295.50 (± 24.63)
Batch = 128, $n = 100$, $n_e = 0$, $n_i = 100$	89.48 (± 7.09)	653.32 (± 5.54)	1267.83 (± 30.85)
Batch = 128, $n = 100$, $n_e = 50$, $n_i = 100$	111.01 (± 8.21)	774.08 (± 20.90)	1739.41 (± 70.53)
Batch = 128, $n = 100$, $n_e = 100$, $n_i = 50$	96.23 (± 9.47)	811.44 (± 25.29)	1896.52 (± 284.89)
Batch = 128, $n = 100$, $n_e = 100$, $n_i = 100$	119.13 (± 10.17)	934.26 (± 87.79)	2059.44 (± 147.77)

Table 5: Timing benchmarks of different backends used for solving a forward pass for different batch sizes. n_e stands for the number of equality constraints.

C.3 Training with large learning rates

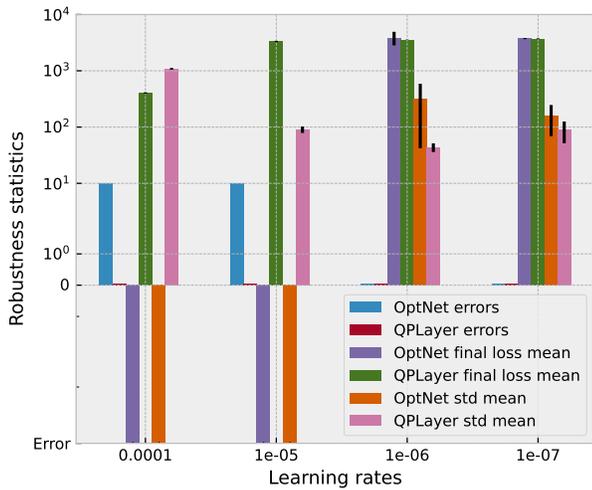
In this section, we assess the numerical robustness of QPLayer on traditional learning tasks by demonstrating that it can be trained with larger learning rates than other approaches, potentially resulting in improved attraction pools.

More precisely, we ran the MNIST classification and denoising tasks described in Section 4.1 with SGD and larger learning rates and reported the corresponding results. We measured the final test loss and error reached after 30 epochs and the standard deviation over the last 10 epochs of the test loss and test error. The results are averaged over 10 seeds and the experiment is performed for different learning rates.

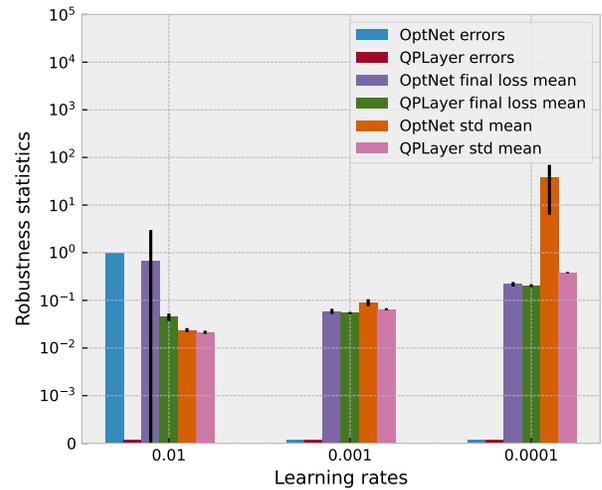
As described in Section 4.1.2, for those tasks, the QP layers need to learn all the model parameters (i.e., H , g , C , and u). We observe that it generates potentially very ill-conditioned problems when the forward or the backward

passes are not solved accurately enough. This phenomenon appears to be amplified with larger learning rates. In those situations, it appears that robust solution methods (e.g., allowing for temporary infeasible, or ill-conditioned problems) are critical.

Figure 7a and Figure 7b show that for too high learning rates (i.e., 10^{-4} or 10^{-5} for denoising task and 10^{-2} for the classification task) the OptNet layer generate errors, whereas it is never the case for QPLayer. Furthermore, for low learning rate levels (i.e., 10^{-6} or 10^{-7} for the denoising task and 10^{-3} and 10^{-4} for the classification task), the final loss reached is similar but with a less important noise amplitude level when using QPLayer (Figure 8 provides robustness statistics of the classification task using the prediction error rate of the two layers). Finally, QPLayer is capable of being trained with a larger learning rate (i.e., 10^{-4} for the denoising task and 10^{-2} for the classification task). Also, note that CvxpyLayer fails to be run in all these robustness experiments because the Hessian part of the quadratic model fails to fit the required DPP form [Amos et al., 2018, Section 4.1].



(a) Robustness statistics for denoising task.



(b) Robustness statistics for MNIST classification task.

Figure 7: Robustness statistics for denoising and MNIST classification tasks: number of errors (i.e., NaNs), averaged final loss reached after 30 epochs (with 95% confidence intervals), and averaged standard deviations over the last 10 epochs (with 95% confidence intervals). Results are averaged over 10 seeds. CvxpyLayer fails to be trained for all these tasks.

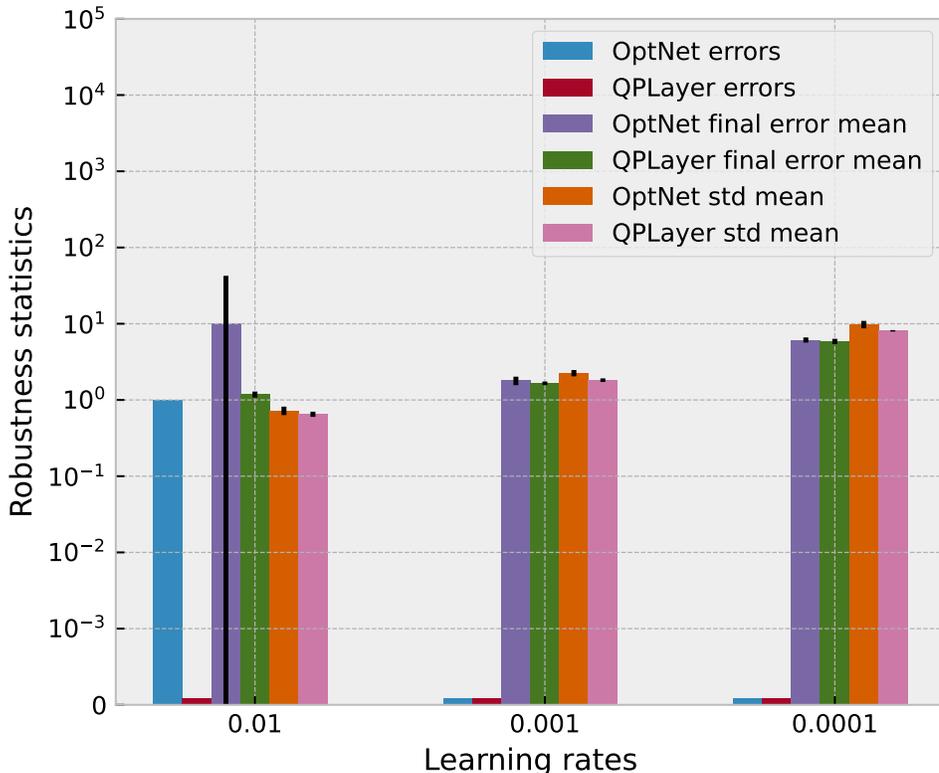


Figure 8: Robustness statistics for the MNIST classification: number of errors (i.e., NaNs), averaged final prediction error reached after 30 epochs (with 95% confidence intervals), and averaged standard deviations over the last 10 epochs (with 95% confidence intervals). Results are averaged over 10 seeds.

Remark 1 (Numerical differences with OptNet). *Our approach offers a few numerical advantages compared to [Amos and Kolter, 2017]. In particular, a numerical matrix factorization is at the center of most popular techniques for differentiating through QPs. This factorization procedure represents one of the main bottlenecks in the computational costs. In our approach, we need to factorize smaller and better-conditioned symmetric matrices.*

The formulation exploited by OptNet consists of a larger linear system to compute its Jacobians. More precisely, they factorize a matrix of the form:

$$K = \begin{bmatrix} H & 0 & C^\top \\ 0 & D(z^*) & D(t^*) \\ C & I & 0 \end{bmatrix},$$

where $t^ = Cx^* - u$ and $D(z^*)$ corresponds to a diagonal matrix whose diagonal entries correspond to z^* . For obtaining a symmetrized version that can be factorized with efficient methods, the second row block is scaled by $D(1/t^*)$ [Amos and Kolter, 2017, Section 3.1]. Yet, symmetrization comes at the price of being more sensitive to the localization of the solution w.r.t the constraints. Indeed, if x^* lies on the boundary, i.e., $C_I x^* - u_I = 0$ for some component index I , the conditioning of the matrix is degraded, as OptNet needs to divide by zeros (or small clamped numbers in practice).*

On our side, the formulation for feasible QPs relies on a smaller matrix, which is symmetric and better-conditioned (it does not require scaling rows by values that are potentially zeros, see (10)).

D Experimental setup

This section details the optimization architectures used for the Sudoku tasks described in Section 4.2. The cart-pole task mentioned in Section 4.1.2 is detailed in Appendix D.2.

D.1 Layer architecture for the Sudoku problem

The layer architecture used by OptNet Amos and Kolter [2017] for the Sudoku problem is described in Figure 9. In contrast, the QPLayer architecture (which does not require structural feasibility of the QPs) is described in Figure 10. The QPLayer architecture allows learning more structured constraints, such as the Sudoku constraint “ $Ax = 1$ ”, which cannot be done, as is, with OptNet (which requires the QPs to be structurally feasible).

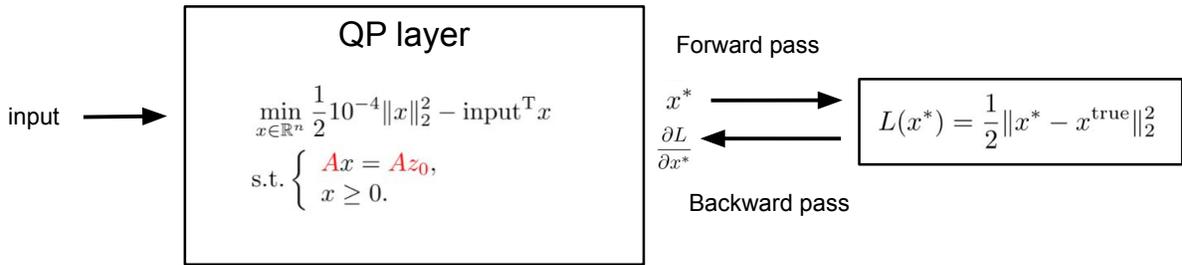


Figure 9: Strictly convex QP layer (as in OptNet Amos and Kolter [2017]). The constraint matrix and an extra variable z_0 are learned in order to be sure that the QP is always primal feasible (structural feasibility).

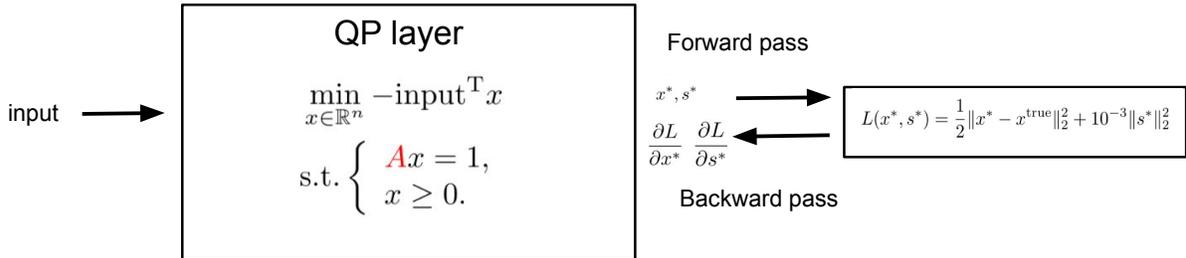


Figure 10: An LP layer (as allowed by QPLayer) allows for more flexibility in the problem to be learned (only the constraint matrix A is learned). The optimal shift s^* is a new output variable minimized in the loss, in order to learn at the end a feasible LP layer.

D.2 Description of the cart-pole problem

The cart-pole system [Anderson, 1989] is a classic control problem used for benchmarking control algorithms. The system we consider consists of an extended model with dry friction on the joints of the cart-pole, namely on the prismatic cart joint and the revolute joint of the pole. It makes the dynamics non-smooth. It is described by a set of differential equations relating the position, velocity, acceleration, angle, and angular velocity of the cart and pole plus the additional friction forces. The static friction forces on each joint can be obtained by solving a QP problem, see (29) and [Le Lidec et al., 2021].

Task: The initial position of the cart-pole system consists of the pole hanging down vertically. The objective of the task is to move the cart in such a way as to swing the pole up and keep it balanced in the upright position. To swing the pole up, the control inputs may involve moving the cart back and forth in a particular pattern that generates the necessary forces to overcome the friction and accelerate the pole in the desired direction.

The forward dynamics with friction

$$Ma = \tau + \lambda,$$

can be re-written in terms of velocity and impulses with timestep Δt as

$$v = v_f + M^{-1}\lambda\Delta t = v_f + M^{-1}\Lambda,$$

where M is the inertia matrix of the system, $a \in \mathbb{R}^{n_v}$ is the joint acceleration, $\tau \in \mathbb{R}^{n_v}$ the joint torque, v_f the free velocity of the system without friction and $\lambda \in \mathbb{R}^{n_v}$ the dry friction force on every joint. To obtain the friction impulse Λ corresponding to the friction coefficient η , the following quadratic problem can be solved:

$$\begin{aligned} \min_{\Lambda} \quad & \frac{1}{2}\Lambda^T M^{-1}\Lambda + v_f^T \Lambda \\ \text{s.t.} \quad & |\Lambda| \leq \eta. \end{aligned} \tag{29}$$

Its Lagrangian L can be written as follows

$$L(\Lambda, y) := \frac{1}{2}\Lambda^T M^{-1}\Lambda + v_f^T \Lambda + y^T (|\Lambda| - \eta),$$

which leads to the KKT system:

$$M^{-1}\Lambda + v_f + \text{diag}(\text{sign}(\Lambda))y = 0, \tag{30}$$

$$|\Lambda| \leq \eta, \tag{31}$$

$$y \geq 0, \tag{32}$$

$$y \odot [|\Lambda| - \eta] = 0, \tag{33}$$

where \odot stands for the standard Hadamard product. Considering the case where the friction force is within the friction cone for a specific joint j , i.e., $|\Lambda_j| < \eta_j$, the joint is then not moving. We see from (33) that $y_j = 0$ satisfies (32) and we get from (30)

$$M^{-1}\Lambda_j = -v_{f,j}.$$

Consequently, the friction impulse is acting in the opposite direction than the joint torque τ_j and with a magnitude that is canceling out the free velocity. If $|\Lambda_j| = \eta_j$, the joint will no longer be blocked by the friction forces and will thus start moving. We see from (30)

$$M^{-1}\Lambda_j + v_{f,j} = v_j = -\text{diag}(\text{sign}(\Lambda_j))y_j. \tag{34}$$

As $y \geq 0$, (34) shows that Λ is opposed to the velocity of the joint.

Optimal control algorithms such as Differential Dynamic Programming (DDP) can be used to compute the optimal trajectory that minimizes a cost function over a finite time horizon, subject to the dynamics of the cart-pole system and control input constraints.

These methods take advantage of the derivatives of the dynamics to efficiently control physical systems. In the presence of non-smooth dynamics, such a class of algorithms is likely to fail due, for instance, to the presence of discontinuities in the dynamics derivatives or because of the non-informative gradient [Lidec et al., 2022]. In the cart-pole benchmark, randomized smoothing, as proposed by [Lidec et al., 2022, Suh et al., 2021] is used to cope with the non-smooth dynamical system. For the optimal swing-up trajectory with 20 timesteps, 5 random samples with a uniform Gaussian noise are generated. The random noise is applied to the input controls, and the dynamics are calculated for each of them and afterward averaged in the forward pass, resulting in informative gradients in the backward pass. (29), has to be solved for every random sample in every timestep.