# Leveraging Proximal Optimization for Differentiating Optimal Control Solvers 

Oumayma Bounou<br>oumayma.bounou@inria.fr *,1,2

Jean Ponce<br>jean.ponce@ens.fr ${ }^{2,3}$

Justin Carpentier<br>justin.carpentier@inria.fr ${ }^{1,2}$


#### Abstract

Over the past few years, differentiable optimization has gained in maturity and attractivity within both machine learning and robotics communities. It consists in computing the derivatives of a given optimization problem which can then be used by learning algorithms, and enables to generically plug computational blocks reflecting the solving of generic mathematical programming problems into a learning pipeline. Until now, dedicated approaches have been proposed to compute the derivatives of various types of optimization problems (LPs, QPs, SOCPs, etc.). However, these approaches assume the problems are well-posed (e.g., satisfaction of the linearly independent constraint qualifications), limiting de facto their application to ill-posed problems. In this work, we focus on the differentiation of optimal control solvers widely used in robotics. We notably introduce a differentiable proximal formulation for solving equality-constrained LQR problems that is effective in solving ill-posed and rank-deficient problems accurately. Importantly, we show that this proximal formulation allows us to compute accurate gradients even in the case of ill-posed problems which do not satisfy the classical constraints qualification. Because any optimal control problem can be casted as an equalityconstrained LQR problem in the vicinity of the optimal solution, ours robust LQR derivatives computation can then be exploited to obtain the derivatives of general optimal control problems. We demonstrate the effectiveness of our approach in dynamics learning and system parameters identification experiments in linear optimal control problems.


## I. Introduction

A key part of programming robot movements is often an instance of trajectory optimization, where the task to solve is encoded as a cost term and physical constraints are encoded as path constraints. Optimal control offers a sound mathematical framework to generically solve trajectory optimization problems. Over the past decades, differential dynamic programming (DDP) [1] and the iterative linear quadratic regulator (iLQR) [2] have become widespread and tractable approaches to solving real and complex robotic problems, ranging from UAVs navigation [3] to the control of legged locomotion [4]. While DDP was originally designed for unconstrained problems, variants have been proposed to account for various levels of path constraints, ranging from simple bounds on the control inputs [5] to the handling of generic equality [6], [7] and inequality constraints [8], [9], [10].

With recent progress in the development of advanced computational frameworks in machine learning (e.g., PyTorch

[^0][11], TensorFlow [12], JAX [13]), differentiable optimization appears as a generic approach to compute the derivatives of computational layers involving mathematical programming problems. It can be used, for instance, to find the optimal design parameters of a robotic mechanism (a.k.a. co-design) given a task to solve [14], or to estimate the sensitivities of an optimal solution with respect to the parameters of the problem [15]. In general, differentiable optimization relies on differentiating the optimality conditions (e.g., the Karush-Kuhn-Tucker conditions) associated to a given problem. In the field of convex optimization, recent works have introduced efficient approaches to compute derivatives of standard quadratic programming (QP) instances [16] and more general convex programming problems [17], which are very common in robotics. When deriving the KKT conditions, these works often assume the well-posedness of the KKT conditions at the optimum (e.g., linearly independent constraints qualification), which cannot be ensured for the large majority of control problems in robotics and more generally when integrating an optimization stage as a layer within a computational graph. For instance, in the case of redundant constraints (e.g., multiple contacts between a robot segment and the environment), one needs to adequately regularize the problem in order to robustly converge to an optimal solution satisfying the constraints [18]. This is even more true in machine learning applications, where mathematical programs are used as differentiable layers. Indeed, the intermediate problems that need to be solved during the learning process might likely not satisfy the required optimality qualifications as they are never enforced during training.

The closest work to ours is the work of Amos et al. in [19], where the authors introduce a generic approach to differentiating constrained LQR problems, and more generically, DDP-based approaches. In their experimental validation, the authors notably report a high failure rate of their differentiable solver when using its computed sensitivities in standard machine learning problems. In this paper, we argue that these failures are partly due to the lack of regularity when evaluating the derivatives of the considered optimal control problems. To overcome these limitations, we leverage the proximal method of multipliers [20] to robustly evaluate the derivatives of optimal control problems. Particularly, we develop a generic solution to compute these derivatives in the equality-constrained linear quadratic regulator (LQR), which is, as highlighted in [19], the core block required to compute the derivatives of nonlinear optimal control problems.

Following the set of experimental validations introduced in [19], we notably demonstrate that our proximal approach is able to properly cope with ill-posed or difficult system identification problems, unlike existing solutions of the state of the art.

## II. Proximal solving of LQR problems

Several approaches have been proposed to efficiently solve LQR problems numerically. In this work, we focus on the dynamic programming formulation of LQR which we regularize with the proximal method of multipliers [20]. Let us first briefly recall this method, and explain how we include it in our approach.

## A. Background on the proximal method of multipliers

Throughout this paper, we address convex minimization problems of the form:

$$
\begin{equation*}
\min _{x} \ell(x), \quad \text { s.t. } c(x)=0 \tag{1}
\end{equation*}
$$

where $\ell: \mathbb{R}^{n} \rightarrow \mathbb{R}$ is a convex function and $c$ is a linear constraint in $x$. Typically, in the case of quadratic programming problems (QP problems), we have $\ell(x)=$ $x^{T} H x+g^{T} x$, where $H$ is a square symmetric and semidefinite matrix of appropriate dimensions, $g$ is a vector and $c(x)=A x-b$, with $A$ a potentially rank-deficient matrix. Let us define the Lagrangian for this problem as

$$
\begin{equation*}
\mathcal{L}(x, \lambda)=\ell(x)+\lambda^{T} c(x) \tag{2}
\end{equation*}
$$

where $\lambda$ is a vector of multipliers. Rockafellar proves in [20] that solving Eq. (1) as a min max problem with saddle points solutions $x, \lambda$ :

$$
\begin{equation*}
x, \lambda=\arg \min _{x} \max _{\lambda} \mathcal{L}(x, \lambda) \tag{3}
\end{equation*}
$$

is equivalent to iteratively solving the regularized min max problem

$$
\begin{align*}
x^{k+1}, \lambda^{k+1}=\arg \min _{x} \max _{\lambda} \mathcal{L}(x, \lambda) & +\frac{1}{2 \mu}\left\|x-x^{k}\right\|_{2}^{2} \\
& -\frac{1}{2 \mu}\left\|\lambda-\lambda^{k}\right\|_{2}^{2} \tag{4}
\end{align*}
$$

starting from initial vectors $x^{0}, \lambda^{0}$ until convergence, with $\mu$ a fixed positive scalar. In (4), the iterative min max problem solved is the same as (3) with additional L2-norm regularization on the primal and dual variables $x$ and $\lambda$, which vanishes over the iterations. Essentially, this proximal regularization [21] of the saddle points associated with problem (1) allows us to solve potentially simpler problems at the price of more iterations. It is central to our approach and will be notably used to solve constrained LQR problems accurately and compute their sensitivities robustly.

## B. Proximal LQR solving

Problem definition. We address the constrained LQR problem of finding the optimal sequence of states and controls that minimize a quadratic cost function $l_{t}$ under given linear dynamics $f_{t}$, path constraints encapsulated in $c_{t}$ (e.g., initial
and terminal constraints), with a control horizon of length $N$. The problem can be formulated as:

$$
\begin{align*}
& \min _{X, U} \sum_{t=0}^{N-1} l_{t}\left(x_{t}, u_{t}\right)+l_{N}\left(x_{N}\right) \\
& \text { s.t }\left\{\begin{array}{l}
x_{t+1}=f_{t}\left(x_{t}, u_{t}\right) \\
x_{0}=x_{0}^{*} \\
c_{t}\left(x_{t}, u_{t}\right)=0
\end{array}\right. \tag{5}
\end{align*}
$$

with $X=\left[\begin{array}{lll}x_{0}^{T} & \ldots & x_{N}^{T}\end{array}\right]^{T}$ and $U=\left[\begin{array}{lll}u_{0}^{T} & \ldots & u_{N-1}^{T}\end{array}\right]^{T}$.
While many methods can be used to solve problem (5], we follow here the dynamic programming approach [1]. Starting from $t=N$, we solve recursively backward in time:

$$
\begin{align*}
V_{t}\left(x_{t}\right)= & \min _{u_{t}, x_{t+1}} l_{t}\left(x_{t}, u_{t}\right)+V_{t+1}\left(x_{t+1}\right)  \tag{6}\\
& \text { s.t. } x_{t+1}=f_{t}\left(x_{t}, u_{t}\right) \text { and } c_{t}\left(x_{t}, u_{t}\right)=0 .
\end{align*}
$$

We define the problem cost and dynamics functions and the path constraints:

$$
\left\{\begin{array}{l}
l_{t}\left(x_{t}, u_{t}\right)=\frac{1}{2}\left(x_{t}^{T} Q_{t} x_{t}+u_{t}^{T} R_{t} u_{t}\right)+q_{x, t}^{T} x_{t}+q_{u, t}^{T} u_{t}  \tag{7}\\
l_{N}\left(x_{N}\right)=\frac{1}{2}\left(x_{N}^{T} Q_{N} x_{N}\right)+q_{x, N}^{T} x_{N} \\
f_{t}\left(x_{t}, u_{t}\right)=A_{t} x_{t}+B_{t} u_{t}+d_{t} \\
C_{t} x_{t}+D_{t} u_{t}=e_{t} \\
x_{0}=x_{0}^{*}
\end{array}\right.
$$

Here, $Q_{t}$ is a symmetric positive semi-definite matrix, $R_{t}$ is a symmetric positive-definite matrix, $A_{t}, B_{t}, C_{t}$ and $D_{t}$ are general matrices, and $e_{t}$ and $d_{t}$ are vectors, for all $t . x_{0}^{*}$ is the initial condition.
Relaxing the dynamic constraints. Following the approach introduced in [22], we relax the dynamics constraints using multiple shooting and introducing an auxiliary variable $y_{t}$, such that problem (6) becomes:

$$
\begin{align*}
& \quad V_{t}\left(x_{t}\right)=\min _{u_{t}, y_{t}} l_{t}\left(x_{t}, u_{t}\right)+V_{t+1}\left(y_{t}\right)  \tag{8}\\
& \text { s.t. } y_{t}=f_{t}\left(x_{t}, u_{t}\right) \text { and } c_{t}\left(x_{t}, u_{t}\right)=0 .
\end{align*}
$$

This inexact method aims at stabilizing the optimization procedure and does not prevent the solver from converging to an optimal solution, as demonstrated by Schmidt et al. in [23].
In the following, we use the usual shorthands $V^{\prime}, V_{x}$ and $V_{x x}$ to denote resepctively $V_{t+1}$, the first and the second derivatives of $V_{t}$. At time $t$, we solve problem (8) by solving the following min max problem:

$$
\begin{equation*}
V_{t}\left(x_{t}\right)=\min _{y_{t}, u_{t}} \max _{\lambda_{t}, \nu_{t}} \mathcal{L}_{t}\left(x_{t}, u_{t}, y_{t}, \lambda_{t}, \nu_{t}\right) \tag{9}
\end{equation*}
$$

where $\mathcal{L}_{t}$ is the Lagrangian of the problem, defined as:

$$
\begin{align*}
\mathcal{L}_{t}\left(x_{t}, u_{t}, y_{t}, \lambda_{t}, \nu_{t}\right)= & l_{t}\left(x_{t}, u_{t}\right)+V^{\prime}\left(y_{t}\right) \\
& +\lambda_{t}^{T}\left(y_{t}-f_{t}\left(x_{t}, u_{t}\right)\right)  \tag{10}\\
& +\nu_{t}^{T} c_{t}\left(x_{t}, u_{t}\right)
\end{align*}
$$

Here, $y_{t}$ and $u_{t}$ are the so-called primal variables, while $\lambda_{t}$ and $\nu_{t}$ are the dual ones. In the LQR context, $V_{N}$ is quadratic i.e., $V_{N}\left(x_{N}\right)=\frac{1}{2} x_{N}^{T} Q_{f} x_{N}+q_{N}^{T} x_{N}$. Thus, for each $t, V_{t}$ is also quadratic. The LQR problem can then be solved using
the dynamic programming principle, by solving at each time $t$ a linear system where the KKT matrix is

$$
\left[\begin{array}{cccc}
R_{t} & 0 & B_{t}^{T} & D_{t}^{T}  \tag{11}\\
0 & V_{x x}^{\prime} & -I & 0 \\
B_{t} & -I & 0 & 0 \\
D_{t} & 0 & 0 & 0
\end{array}\right] .
$$

Such a matrix can have very poor conditioning, especially in the cases where $V_{x x}^{\prime}$ and/or $R_{t}$ are not strictly symmetric positive definite or where the matrix $D_{t}$ is rank-deficient. As noted in the introduction, such a scenario is likely to happen in practice, especially in a learning framework where the cost matrices $R_{t}$ and $Q_{t}$ and the path constraints matrix $D_{t}$ are learnable parameters. To cope with this, we propose to solve a proximal reformulation of this problem that we detail in the next paragraph.

## C. Proximal regularization

We find a solution to problem (9) by iteratively solving for all $t$ the equivalent regularized problem

$$
\begin{align*}
\mathcal{L}_{\rho \mu}\left(x_{t}, u_{t}, y_{t}, \lambda_{t}\right) & =l_{t}\left(x_{t}, u_{t}\right)+V_{t+1}\left(y_{t}\right) \\
& +\lambda_{t}^{T}\left(y_{t}-f\left(x_{t}, u_{t}\right)\right)+\nu_{t} c_{t}\left(x_{t}, u_{t}\right) \\
& +\frac{\rho}{2}\left\|u_{t}-u_{t}^{-}\right\|_{2}^{2}+\frac{\rho}{2}\left\|y_{t}-y_{t}^{-}\right\|_{2}^{2} \\
& -\frac{\mu}{2}\left\|\lambda_{t}-\lambda_{t}^{-}\right\|_{2}^{2}-\frac{\mu}{2}\left\|\nu_{t}-\nu_{t}^{-}\right\|_{2}^{2} \tag{12}
\end{align*}
$$

starting from initial guesses $u_{t}^{0}, y_{t}^{0}, \lambda_{t}^{0}$ and $\nu_{t}^{0}$ as shown in [20] and detailed in Sec. II-A The superscript "-" is used for variables from the previous proximal iteration. $\mathcal{L}_{\rho \mu}$ is a regularized version of $\mathcal{L}_{t}$ with additional L2 penalties on primal $(u, y)$ and dual $(\lambda, \nu)$ variables, and $\rho$ and $\mu$ are positive regularization coefficients.
The new KKT conditions at time $t$ at the optimum are given by:

$$
\left[\begin{array}{cccc}
R_{t}+\rho I & 0 & B_{t}^{T} & D_{t}^{T}  \tag{13}\\
0 & V_{x x}^{\prime}+\rho I & -I & 0 \\
B_{t} & -I & -\mu I & 0 \\
D_{t} & 0 & 0 & -\mu I
\end{array}\right]\left[\begin{array}{c}
d u_{t} \\
d y_{t} \\
d \lambda_{t} \\
d \lambda_{t}
\end{array}\right]=-\left[\begin{array}{c}
R_{t} u_{t}^{-}+B_{t}^{T} \lambda_{t}^{-}+D_{t}^{T} \nu_{t}^{-}+q_{u, t} \\
V_{x x}^{\prime} y_{t}^{t}+V_{x}^{\prime}-\lambda_{t}^{-} \\
A_{t} x_{t}+B_{t} u_{t}^{-}+d_{t}-y_{t}^{-} \\
C_{t} x_{t}+D_{t} u_{t}^{-}-e_{t}
\end{array}\right],
$$

where $d v=v-v^{-}$, for $v$ in $\left\{u_{t}, y_{t}, \lambda_{t}, \nu_{t}\right\}$. When solving Eq. (13) backward in time, $x_{t}$ is unknown, but $d u_{t}, d y_{t}, d \lambda_{t}, d \nu_{t}$ can be expressed as affine functions of $x_{t}$ :

$$
\left\{\begin{align*}
d u_{t} & =\Gamma_{t} x_{t}+\gamma_{t}  \tag{14}\\
d y_{t} & =M_{t} x_{t}+m_{t}, \\
d \lambda_{t} & =\Xi_{t} x_{t}+\xi_{t}, \\
d \nu_{t} & =\Omega_{t} x_{t}+\omega_{t}
\end{align*}\right.
$$

and the coefficients $\Gamma_{t}, \gamma_{t}, M_{t}, m_{t}, \Xi_{t}$ and $\xi_{t}$ can be obtained by solving the linear system:

$$
\mathrm{K}\left[\begin{array}{cc}
\Gamma_{t} & \gamma_{t}  \tag{15}\\
M_{t} & m_{t} \\
\Xi_{t} & \xi_{t} \\
\Omega_{t} & \omega_{t}
\end{array}\right]=-\left[\begin{array}{cc}
0 & R_{t} u_{t}^{-}+B_{t}^{T} \lambda_{t}^{-}+D_{t}^{T} \nu_{t}^{-}+q_{u, t} \\
0 & V_{x x}^{\prime} y_{t}^{-}+V_{x}^{\prime}-\lambda_{t}^{-} \\
A_{t} & B u_{t}^{-}+d_{t}-y_{t}^{-} \\
-C_{t} & D_{t} u_{t}^{-}-e_{t}
\end{array}\right] .
$$

In the backward pass of the algorithm, we compute the coefficients from Eq. 15) and update the value function
first and second derivatives $V_{x x}$ and $V_{x}$ for each $t$. Then, we perform forward passes in time to update the variables $u_{t}, y_{t}, \lambda_{t}, \nu_{t}$, with

$$
\left\{\begin{array}{c}
u_{t}^{+}=u_{t}^{-}+d u_{t},  \tag{16}\\
y_{t}^{+}=y_{t}^{-}+d y_{t}, \\
\lambda_{t}^{+}=\lambda_{t}^{-}+d \lambda_{t}, \\
\nu_{t}^{+}=\nu_{t}^{-}+d \nu_{t} .
\end{array} .\right.
$$

We thus solve the LQR problem by iteratively repeating this backward / forward procedure until convergence. Our stopping criterion is the infinite norm of the gradient of $\mathcal{L}_{\rho, \mu}$ (i.e., the infinite norm of the right hand side of Eq. (13) over all times $t$ ):

$$
\left\|\left[\begin{array}{c}
\nabla_{y, u} \mathcal{L}(y, u)  \tag{17}\\
\nabla_{\lambda, \nu} \mathcal{L}(\lambda, \nu)
\end{array}\right]\right\|_{\infty} \leq \epsilon,
$$

where $\epsilon$ is a fixed tolerance parameter.

## D. Derivatives

We follow the approach of [19] to compute derivatives of LQR outputs (i.e., the trajectories $\mathbf{x}$ and $\mathbf{u}$ ) with respect to the problem parameters (i.e., the dynamics and cost matrices $A_{t}, B_{t}, Q_{t}, R_{t}$, the path constraints $C_{t}, D_{t}, d_{t}, e_{t}$, and the initial condition $x_{0}^{*}$ ) to be able to plug the LQR solver as a differentiable layer in a model. Because our formulation of the LQR problem in Eq. (7) is generic and accounts for additional path constraints, our differentiation formulation can be used to differentiate through more general optimal control problems with potentially nonlinear equality and/or inequality constraints, by differentiating through the solutions at the optimum.
We first reformulate the resolution of the general LQR problem defined in Eq. (7) as the minimization of a large QP problem

$$
\begin{equation*}
\min _{X} \frac{1}{2} X^{T} H X+q^{T} X, \quad \text { s.t } \hat{A} X=b \tag{18}
\end{equation*}
$$

with the following augmented matrices:

$$
\begin{gathered}
H=\operatorname{diag}\left(Q_{t}, \ldots, R_{t}, \ldots\right), \\
q=\left[\begin{array}{lcccccccc}
q_{x, 0}^{T} & \ldots & q_{x, N}^{T} & q_{u, 0}^{T} & \ldots & q_{u, N-1}^{T}
\end{array}\right]^{T}, \\
\hat{A}=\left[\begin{array}{ccccccccc}
-I & 0 & 0 & \ldots & 0 & 0 & 0 & \ldots & 0 \\
A_{0} & -I & 0 & \ldots & 0 & B_{0} & 0 & \ldots & 0 \\
0 & A_{1} & -I & \ldots & 0 & 0 & B_{1} & \ldots & 0 \\
\ldots & \ldots & \ldots & \ldots & \ldots & \ldots & \ldots & \ldots & \ldots \\
\ldots & \ldots & \ldots & A_{N-1} & -I & 0 & \ldots & 0 & B_{N-1} \\
C_{0} & \ldots & \ldots & 0 & 0 & D_{0} & \ldots & 0 & 0 \\
\ldots & \ldots & \ldots & \ldots & \ldots & \ldots & \ldots & \ldots & \ldots \\
0 & \ldots & \ldots & C_{N-1} & 0 & D_{N-1} & \ldots & 0 & 0 \\
0 & \ldots & \ldots & 0 & C_{N} & 0 & \ldots & 0 & 0
\end{array}\right], \\
X=\left[\begin{array}{lllllll}
x_{0}^{T} & \ldots & x_{N}^{T} & u_{0}^{T} & \ldots & u_{N-1}^{T}
\end{array}\right]^{T},
\end{gathered}
$$

and

$$
b=-\left[\begin{array}{lllllll}
x_{0}^{*, T} & d_{1}^{T} & \ldots & d_{N}^{T} & e_{0}^{T} & \ldots & e_{N}^{T}
\end{array}\right] .
$$

The KKT system associated to this problem at the optimum is:

$$
\left[\begin{array}{cc}
H & \hat{A}^{T}  \tag{19}\\
\hat{A} & 0
\end{array}\right]\left[\begin{array}{c}
X^{*} \\
\Lambda^{*}
\end{array}\right]=\left[\begin{array}{c}
-q \\
b
\end{array}\right]
$$

where $\Lambda^{*}$ is a vector of multipliers associated to the constraints $\hat{A} X=b$ and containing the stacked multipliers $\lambda_{t}$, $\nu_{t}$ from Eq. (10). for all $t$. We denote as $\hat{K}$ the KKT matrix from Eq. 19p associated with the large QP problem. In the following, we drop the " $*$ " superscript for clarity. In practice, the derivative of interest in learning frameworks is $\partial r / \partial p$, where $p$ is in $\mathcal{P}_{t}=\left\{A_{t}, B_{t}, Q_{t}, q_{x, t}, q_{u, t}, C_{t}, D_{t}, x_{0}^{*}, d_{t}, e_{t}\right\}$, and $r$ is a scalar function of the parameters $\mathcal{P}_{t}$. To obtain $\partial r / \partial p$ directly, we use the same trick as in [19] and [24], and obtain expressions for the partial derivatives

$$
\begin{array}{rlrl}
\partial r / \partial A_{t} & =\lambda_{t+1} G_{x, t}^{T}+x_{t} G_{\lambda, t+1}^{T} \\
\partial r / \partial B_{t} & =\lambda_{t+1} G_{u, t}^{T}+G_{\lambda, t+1} u_{t}^{T} \\
\partial r / \partial Q_{t} & =x_{t} G_{x, t}^{T} & & \partial r / \partial R_{t}=G_{u, t} u_{t}^{T}  \tag{20}\\
\partial r / \partial C_{t} & =-x_{t} G_{\nu, t}^{T} & & \partial r / \partial e_{t}=G_{\nu, t} \\
\partial r / \partial D_{t} & =-u_{t} G_{\nu, t}^{T} & & \partial r / \partial d_{t}=G_{\lambda, t} \\
\partial r / \partial q_{x, t} & =G_{x, t} & & \partial r / \partial q_{u, t}=G_{u, t}
\end{array}
$$

as functions of a vector $G$ that verifies

$$
\begin{equation*}
\hat{K} G=Z \tag{21}
\end{equation*}
$$

where $Z$ is the vector resulting from stacking all the vectors $\partial r / \partial v_{t}$ with $v_{t}$ in $\left\{x_{t}, u_{t}, \lambda_{t}, \nu_{t}\right\}$, for all $t$, which can be obtained with backpropagation. We can also write $G$ as the stacking of vectors $G_{v, t}$ for all $t$, with $v_{t}$ in $\left\{x_{t}, u_{t}, \lambda_{t}, \nu_{t}\right\}$, where each vector $G_{v, t}$ is the same size as $v_{t}$. In [19], Amos et al. notice that $G$ verifies the same equation the vector $\left[\begin{array}{ll}X^{T} & \Lambda^{T}\end{array}\right]^{T}$ verifies in Eq. 19 . Thus, $G$ is solution to a new LQR problem, and can be obtained efficiently without explicitly inverting the large KKT matrix $\hat{K}$ in Eq. (21). Instead, we solve a similar LQR problem to the one $\mathbf{x}$ and $\mathbf{u}$ are solutions to. The KKT matrices of this new problem, which we refer to as the LQR derivatives problem, are the same as the original one, and only the values of the right hand side term in the large QP formulation are modified, which correspond to the parameters $q_{x, t}, q_{u, t}, d_{t}$ and $e_{t}$ in the LQR formulation. Their values for the LQR derivatives problem are:

$$
\left\{\begin{array}{lc}
q_{x, t}=\partial r / \partial x_{t} & q_{u, t}=\partial r / \partial u_{t}  \tag{22}\\
d_{t}=\partial r / \partial \lambda_{t} & e_{t}=\partial r / \partial \nu_{t}
\end{array}\right.
$$

Unlike [19], we solve the new LQR derivatives problem using the regularized solver formulation we introduced in sections II-B and II-C We solve the LQR problem the vectors $G_{x, t}, G_{u, t}$ are solutions to, and the multipliers $G_{\lambda, t}$ and $G_{\nu, t}$ associated to the constraints. Starting from initial vectors $G_{u, t}^{0}, G_{y, t}^{0}, G_{\lambda, t}^{0}, G_{\nu, t}^{0}$, we solve iteratively and until convergence the following linear system for all $t$, backward in time starting from $t=N$ :
$\mathrm{K}_{t}\left[\begin{array}{c}d G_{u, t} \\ d G_{y, t} \\ d G_{\lambda, t} \\ d G_{\nu, t}\end{array}\right]=-\left[\begin{array}{c}R_{t} G_{u, t}^{-}+B_{t}^{T} G_{\lambda, t}^{-}+D_{t}^{T} G_{\nu, t}^{-}+\partial r / \partial u_{t} \\ V_{x x}^{\prime} G_{y, t}^{-}+V_{x}^{\prime}-G_{\lambda, t}^{-} \\ A_{t} G_{x, t}+B_{t} G_{u, t}^{-}+\partial r / \partial \lambda_{t}-G_{y, t}^{-} \\ C_{t} G_{x, t}+D_{t} G_{u, t}^{-}-\partial r / \partial \nu_{t}\end{array}\right]$,
where

$$
K_{t}=\left[\begin{array}{cccc}
R_{t}+\rho I & 0 & B_{t}^{T} & D_{t}^{T}  \tag{24}\\
0 & V_{x x}^{\prime}+\rho I & -I & 0 \\
B_{t} & -I & -\mu I & 0 \\
D_{t} & 0 & 0 & -\mu I
\end{array}\right] .
$$

This allows us to find the optimal variable $G$ accurately, even in ill-posed problems (e.g., rank-deficient matrices). This is extremely important in practice since the gradients of interest, $\partial r / \partial p$ with $p$ in $\mathcal{P}$, are functions of the values of $G$ (equation (20). Thus, any inaccurate solution $G$ would lead to backpropagating wrong gradients, resulting in unstable training procedures.

## E. Leveraging $L Q R$ derivatives for nonlinear optimal control problems

The iLQR algorithm introduced in [2] solves problem (5) with general nonlinear dynamics and cost functions $f$ and $l$. It linearizes the dynamics around nominal state and control trajectories, makes quadratic approximations of the cost around these trajectories, then iteratively solves the LQR problem obtained with this approximation. Derivatives of iLQR solutions with respect to the problem parameters can be obtained from the LQR approximation of the iLQR problem at optimality.
Let us first introduce the notations for the iLQR problem. We use the usual $Q_{t}$ notation for the Lagrangian of problem (5) at time $t$ and introduce

$$
\begin{aligned}
Q_{\rho \mu}\left(x_{t}, u_{t}, y_{t}, \lambda_{t}\right) & =l_{t}\left(x_{t}, u_{t}\right)+V_{t+1}\left(y_{t}\right)+\lambda_{t}^{T}\left(f\left(x_{t}, u_{t}\right)-y_{t}\right) \\
& +\nu_{t}^{T} c\left(x_{t}, u_{t}\right)+\frac{\rho}{2}\left\|u_{t}-u_{t}^{-}\right\|_{2}^{2}+\frac{\rho}{2}\left\|y_{t}-y_{t}^{-}\right\|_{2}^{2}(25) \\
& -\frac{\mu}{2}\left\|\lambda_{t}-\lambda_{t}^{-}\right\|_{2}^{2}-\frac{\mu}{2}\left\|\nu_{t}-\nu_{t}^{-}\right\|_{2}^{2}
\end{aligned}
$$

with $\rho$ and $\mu$ positive regularization scalars. The goal is to find the optimal solutions $y_{t}, u_{t}, \lambda_{t}, \nu_{t}$ to problem

$$
\begin{equation*}
\arg \min _{y_{t}, u_{t}} \max _{\lambda_{t}, \nu_{t}} Q_{\rho \mu}\left(x_{t}, u_{t}, y_{t}, \lambda_{t}, \nu_{t}\right), \tag{26}
\end{equation*}
$$

using the dynamic programming recursion

$$
\begin{equation*}
V_{t}\left(x_{t}\right)=\min _{u_{t}, y_{t} \lambda_{t}, \nu_{t}} \max _{\rho \mu}\left(x_{t}, u_{t}, y_{t}, \lambda_{t}, \nu_{t}\right) \tag{27}
\end{equation*}
$$

In our formalism, the approximation of the function $Q_{t}$ is:

$$
\begin{aligned}
\delta Q_{t}= & \frac{1}{2}\left[\begin{array}{l}
\delta x_{t} \\
\delta y_{t} \\
\delta u_{t} \\
\delta \lambda_{t} \\
\delta \nu_{t}
\end{array}\right]^{T}\left[\begin{array}{lllll}
Q_{x x} & Q_{x y} & Q_{x u} & Q_{x \lambda} & Q_{x \nu} \\
Q_{y x} & Q_{y y} & Q_{y u} & Q_{y \lambda} & Q_{y \nu} \\
Q_{u x} & Q_{u y} & Q_{u u} & Q_{u \lambda} & Q_{u \nu} \\
Q_{\lambda x} & Q_{\lambda y} & Q_{\lambda u} & Q_{\lambda \lambda} & Q_{\lambda_{\nu}} \\
Q_{\nu_{x}} & Q_{\nu y} & Q_{\nu u} & Q_{\nu \lambda} & Q_{\nu \nu}
\end{array}\right]\left[\begin{array}{l}
\delta x_{t} \\
\delta y_{t} \\
\delta u_{t} \\
\delta \lambda_{t} \\
\delta \nu_{t}
\end{array}\right] \\
& Q_{x}^{T} \delta x_{t}+Q_{y}^{T} \delta y_{t}+Q_{u}^{T} \delta u_{t}+Q_{\lambda}^{T} \delta \lambda_{t}+Q_{\nu} \delta \nu_{t} .
\end{aligned}
$$

Let us assume the dynamics and cost function $f$ and $l$ are parameterized by some parameter $\theta . Q$ and its set of first and second order derivatives are evidently also parameterized by $\theta$, namely the ones at optimality. In a learning framework, the derivative of interest, i.e., $\partial r / \partial \theta$, with $r$ some scalar loss function over the parameters can be obtained through the chain rule

$$
\begin{equation*}
\frac{\partial r}{\partial \theta}=\sum_{P \in \mathcal{D}_{Q}^{*}} \frac{\partial r}{\partial P} \frac{\partial P}{\partial \theta} \tag{29}
\end{equation*}
$$

where $\mathcal{D}_{Q}^{*}=\left\{Q_{x}^{*}, Q_{x x}^{*}, \ldots, Q_{\lambda \lambda}^{*}\right\}$. The derivatives $\{\partial r / \partial P\}$ can be obtained with the method described in paragraph II-D.
and the derivatives $\partial P / \partial \theta$ are obtained with backpropagation.

## III. Experiments

## A. LQR problems with terminal constraints

We solve an LQR problem with a terminal constraint ( (5) and (7)) following the approach described in section II-C We compare the solutions of both our solver and CVXPY [25], [26], which is based on OSQP [27]. We run both solvers on three sets of parameter sizes $\{n, d, N\}$, with $n$ the system dimension, $d$ the control dimension and $N$ the time horizon. For each set of parameter sizes, we run 100 experiments with randomly generated time-invariant LQR problems. The dynamics matrix $A$ is forced to have all singular values lower than 1 , and $B$ is a matrix with random coefficients sampled uniformly in $[0,1]$. The cost matrices $Q$ and $R$ are set to respectively $10^{-2} I_{n}$ and $10^{-1} I_{d}$. All experiments are run on a single CPU. We report the average results in Table I A solver is considered successful when both primal and dual constraints are satisfied. Feasibility denotes the infinite norm on primal constraints (dynamics constraints $\left.\left\|x_{t+1}^{*}-A x_{t}^{*}-B u_{t}^{*}\right\|_{\infty}\right)$. Distance to goal is the average infinite norm $\left\|x_{N}-x_{N}^{*}\right\|_{\infty}$, with $x_{N}^{*}$ the target terminal constraint. In the rank-deficient cases (first and second rows), our solver converges to a solution for almost every experiment, while CVXPY only converges to a solution for up to $32 \%$ of them on average. On examples where both solvers converge, ours converges to a more accurate solution. We also report the average results on well-posed cases (third row). On such cases, both solvers converge to equally good solutions, but CVXPY is much faster, since our implementation is in an interepreted language (Python) that could easily be moved to a compiled one (e.g., $\mathrm{C}++$ ) for much better efficiency.

TABLE I: Solvers performance. Comparison on LQR problems with terminal constraints.

| Parameters | solver | success | feasibility | distance to goal |
| :--- | :---: | :---: | :---: | :---: |
| $N=20$, | cvxpy | 32 | $2.10^{-8}$ | $2.10^{-8}$ |
| $n=10, d=2$ | ours | $\mathbf{9 8}$ | $\mathbf{4 . 1 0}^{-9}$ | $\mathbf{4 . 1 0}^{-9}$ |
| $N=20$, | cvxpy | 2 | $6.10^{-6}$ | $6.10^{-6}$ |
| $n=15, d=3$ | ours | $\mathbf{9 3}$ | $\mathbf{9 . 1 0}^{-9}$ | $\mathbf{9 . 1 0}^{-9}$ |
| $N=20$, | cvxpy | $\mathbf{1 0 0}$ | $9.10^{-12}$ | $9.10^{-12}$ |
| $n=15, d=5$ | ours | $\mathbf{1 0 0}$ | $\mathbf{2 . 1 0}^{-14}$ | $\mathbf{3 . 1 0}^{-15}$ |

## B. System identification

Identifying dynamics and cost matrices $\left(A, B, Q, Q_{f}, R\right)$. We reproduce the experimental setting of [19]. Given optimal trajectories in states and controls of systems with linear dynamics and quadratic cost, the goal is to identify these dynamics and cost parameters. Formally, we solve

$$
\begin{aligned}
\min _{X, U} & \frac{1}{2} \sum_{t=0}^{N-1}\left(x_{t}^{T} Q x_{t}+u_{t}^{T} R u_{t}\right)+\frac{1}{2} x_{N}^{T} Q_{f} x_{N} \\
\text { s.t } & x_{t+1}=A x_{t}+B u_{t} \text { for } t \text { in } 0, \ldots, N-1, \\
& x_{0}=x_{0}^{*}
\end{aligned}
$$

We observe $M$ trajectories $\left[x_{0}^{*, i}, \ldots x_{N}^{*, i}\right] \quad$ and $\left[u_{0}^{*, i}, \ldots, u_{N-1}^{*, i}\right](i$ in $1, \ldots, M)$ that are optimal solutions to problem (30), with different random initial conditions $x_{0}^{*}$. The states $x_{t}$ are vectors of size $n$ (between 2 and 10), the controls $u_{t}$ are vectors of size $d$ (between 3 and 10), and the control horizon is $N$ (between 5 and 20). Our goal is to identify the dynamics matrices $(A, B)$ and the cost matrices $\left(Q, Q_{f}, R\right)$. In other words, we want to estimate $\theta$, where $\theta$ can be any of $A, B, Q, Q_{f}, R$, or a combination of two or more of these matrices. We solve:

$$
\begin{equation*}
\min _{\theta} \sum_{i} \sum_{t}\left\|x_{t}^{*, i}-x_{t}^{i}(\theta)\right\|_{2}^{2}+\left\|u_{t}^{*, i}-u_{t}^{i}(\theta)\right\|_{2}^{2} \tag{30}
\end{equation*}
$$

where $x_{t}^{i}$ and $u_{t}^{i}$ are solutions to the LQR problem parameterized by $\theta$ with initial condition $x_{0}^{*, i}$. Using the same notations as in [19], we define the optimal trajectory vector $i, \tau^{*, i}$ as

$$
\tau^{*, i}=\left[\begin{array}{llllll}
x_{0}^{*, i T} & \ldots & x_{T}^{*, i T} & u_{0}^{*, i T} & \ldots & u_{T-1}^{*, i} \tag{31}
\end{array}\right]^{T} .
$$

Here, $\tau^{*, i}$ is a vector of size $p=(N+1) n+N d$. The problem is now reduced to solving

$$
\begin{equation*}
\min _{\theta} \sum_{i=1}^{M} \sum_{k=0}^{p}\left\|\tau_{k}^{*, i}-\tau_{k}^{i}(\theta)\right\|_{2}^{2}=\min _{\theta} \sum_{k=0}^{N p}\left\|r_{k}(\theta)\right\|_{2}^{2} \tag{32}
\end{equation*}
$$

where $r_{k}(\theta)=\Gamma_{k}^{*}-\Gamma_{k}, \tau_{k}^{*, i}$ (respectively $\left.\tau_{k}^{i}(\theta)\right)$ is the $k$-th component of vector $\tau^{*, i}$ (respectively $\tau^{i}$ ), and $\Gamma_{k}^{*}$ (respectively $\Gamma_{k}$ ) contains stacked vectors $\tau^{*, i}$ (respectively $\tau^{i}$ ). Problem (32) is a non-linear least-squares problem that can be solved using methods such as Gauss-Newton [28] or LevenbergMarquardt [29]. In fact, when a step in a solution of an optimal control problem boils down to least squares, stochastic gradient descent should not be used for this step (as expected from the optimization literature), which partly explains what happens in [19], where the identification experiments fail in half the trials when using gradient descent to optimize Eq. 30. We reproduced their experiments and show results

TABLE II: Success rate of identification methods. We compare Levenberg-Marquardt and RMSProp optimization success rates (in \%).

|  | $\mathrm{N}=5, \mathrm{n}=3, \mathrm{~d}=3$ | $\mathrm{~N}=20, \mathrm{n}=3, \mathrm{~d}=3$ | $\mathrm{~N}=10, \mathrm{n}=5, \mathrm{~d}=2$ |
| :---: | :---: | :---: | :---: |
| RMSProp | 21 | 20 | 21 |
| LM | $\mathbf{9 1}$ | $\mathbf{8 5}$ | $\mathbf{9 0}$ |

averaged on 100 experiments in Table II An experiment is considered successful if the identification error $\left\|\theta-\theta^{*}\right\|_{\infty}$ reaches the threshold $5.10^{-6}$ in less than 50 iterations of the Levenberg-Marquardt algorithm or 2000 epochs of RMSProp. Table II shows that the identification succeeds $85 \%$ of the time when using a least-squares method, while it only succeeds $20 \%$ of the time when using RMSProp, which demonstrate the ineffectiveness of stochastic gradient methods in leastssquares problems. It should be noted however that the implementation of Levenberg-Marquardt-like methods require computing the full Jacobian matrix of the residual function $r$ (instead of just Jacobian-vector products as in SGD methods), which is a function that scales linearly with $N, n$ and $d$.

In the following experiments, the identification problem was solved using the Levenberg-Marquardt method for both our solver and the solver from [19].


Fig. 1: Identification error on identifying the matrices $A$ and $B$ with $Q=10^{-4} I$. Pairs of same color curves are identification experiments on the same problem parameters solved using different solvers: diff-mpc in dashed-lines, and ours in solid lines.


Fig. 2: Optimization loss on identifying the matrices $A$ and $B$ with $Q=10^{-4} I$. Pairs of same color curves are identification experiments on the same problem parameters solved using different solvers: diff-mpc in dashed-lines, and ours in solid lines.

Figures 1 and 2 show the system identification and optimization errors as functions of the number of iterations using both the solver from [19], which we refer to as diffmpc , and our solver on four trials ( 100 experiments were run, but we only show four randomly selected ones here). In this experiment, the parameters to identify were the dynamics and control matrices $A$ and $B$, the control cost matrix $R$ was set to the identity matrix, and the states cost matrix $Q$ was set to $10^{-4}$. In this experiment, both solvers converge, but our regularized solver converges in twice less iterations on all the trials.
Figures 3 and 4 show the system identification and optimization errors as functions of the number of iterations using both solvers again, on experiments where the parameter to identify was the matrix $Q$. The approach of [19] fails on all trials, demonstrating the importance of regularizing the LQR solver


Fig. 3: Identification error. Identification of $Q$.Pairs of same color curves are identification experiments on the same problem parameters solved using different solvers: diff-mpc in dashed-lines, and ours in solid lines.


Fig. 4: Optimization loss. Identification of $Q$. Pairs of same color curves are identification experiments on the same problem parameters solved using different solvers: diff-mpc in dashed-lines, and ours in solid lines.
to avoid backpropagation of wrong gradients.

## IV. Conclusion

We have introduced a regularized differentiable equalityconstrained LQR solver with a generic formulation that handles path constraints. When used in learning frameworks, our solver is robust to ill-posed problems and achieves better performance on system identification experiments than the authors in [19]. Since our formulation of the LQR problem is generic, it can be applied to differentiate accurately through more general optimal control problems. Future work should include experiments on such problems, and ultimately on real robotic systems to demonstrate the effectiveness and robustness of this approach on real-world robotic tasks.

## ACKNOWLEDGMENT

This work was supported in part by the Inria/NYU collaboration, the Louis Vuitton/ENS chair on artificial intelligence, the HPC resources from GENCI-IDRIS (Grand 2020-AD011011263R1) and the French government under management of Agence Nationale de la Recherche as part of the "Investissements d'avenir" program, reference ANR19P3IA0001 (PRAIRIE 3IA Institute).

## References

[1] D. Q. Mayne, "Differential Dynamic Programming-A Unified Approach to the Optimization of Dynamic Systems," vol. 10 of Control and Dynamic Systems, pp. 179-254, Academic Press, 1973.
[2] W. Li and E. Todorov, "Iterative linear quadratic regulator design for nonlinear biological movement systems.," in ICINCO (1), pp. 222-229, Citeseer, 2004.
[3] J. Marti-Saumell, J. Solà, C. Mastalli, and A. Santamaria-Navarro, "Squash-box feasibility driven differential dynamic programming," in 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 7637-7644, IEEE.
[4] R. Budhiraja, J. Carpentier, C. Mastalli, and N. Mansard, "Differential dynamic programming for multi-phase rigid contact dynamics," in 2018 IEEE-RAS 18th International Conference on Humanoid Robots (Humanoids), pp. 1-9, IEEE, 2018.
[5] Y. Tassa, N. Mansard, and E. Todorov, "Control-limited differential dynamic programming," in 2014 IEEE International Conference on Robotics and Automation (ICRA), pp. 1168-1175, IEEE, 2014.
[6] M. Giftthaler and J. Buchli, "A projection approach to equality constrained iterative linear quadratic optimal control," in 2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids), pp. 61-66, IEEE, 2017.
[7] S. El Kazdadi, J. Carpentier, and J. Ponce, "Equality constrained differential dynamic programming," in 2021 IEEE International Conference on Robotics and Automation (ICRA), pp. 8053-8059, IEEE, 2021.
[8] G. Lantoine and R. P. Russell, "A hybrid differential dynamic programming algorithm for constrained optimal control problems. part 1: Theory," Journal of Optimization Theory and Applications, vol. 154, no. 2, pp. 382-417, 2012.
[9] T. A. Howell, B. E. Jackson, and Z. Manchester, "Altro: A fast solver for constrained trajectory optimization," in 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 7674-7679, IEEE, 2019.
[10] W. Jallet, A. Bambade, N. Mansard, and J. Carpentier, "Constrained Differential Dynamic Programming: A primal-dual augmented Lagrangian approach," in 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems, (Kyoto, Japan), Oct. 2022.
[11] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, highperformance deep learning library," 2019.
[12] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, "TensorFlow: Largescale machine learning on heterogeneous systems," 2015. Software available from tensorflow.org.
[13] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, "JAX: composable transformations of Python+NumPy programs," 2018.
[14] G. Fadini, T. Flayols, A. Del Prete, N. Mansard, and P. Souères, "Computational design of energy-efficient legged robots: Optimizing for size and actuators," in 2021 IEEE International Conference on Robotics and Automation (ICRA), pp. 9898-9904, IEEE, 2021.
[15] A. Oshin, M. Houghton, M. Acheson, I. Gregory, and E. Theodorou, "Parameterized Differential Dynamic Programming," in Proceedings of Robotics: Science and Systems, (New York City, NY, USA), June 2022.
[16] B. Amos and J. Z. Kolter, "Optnet: Differentiable optimization as a layer in neural networks," in International Conference on Machine Learning, pp. 136-145, PMLR, 2017.
[17] A. Agrawal, B. Amos, S. Barratt, S. Boyd, S. Diamond, and J. Z. Kolter, "Differentiable convex optimization layers," Advances in Neural Information Processing Systems, vol. 32, 2019.
[18] D. P. Bertsekas, Constrained optimization and Lagrange multiplier methods. Academic press, 2014.
[19] B. Amos, I. Jimenez, J. Sacks, B. Boots, and J. Z. Kolter, "Differentiable mpc for end-to-end planning and control," Advances in Neural Information Processing Systems, vol. 31, 2018.
[20] R. T. Rockafellar, "Augmented Lagrangians and applications of the proximal point algorithm in convex programming," Mathematics of Operations Research, vol. 1, no. 2, pp. 97-116, 1976.
[21] N. Parikh and S. Boyd, "Proximal algorithms," Found. Trends Optim., vol. 1, p. 127-239, Jan. 2014.
[22] W. Jallet, N. Mansard, and J. Carpentier, "Implicit Differential Dynamic Programming," in International Conference on Robotics and Automation (ICRA 2022), (Philadelphia, United States), IEEE Robotics and Automation Society, May 2022.
[23] M. Schmidt, N. Roux, and F. Bach, "Convergence rates of inexact proximal-gradient methods for convex optimization," Advances in neural information processing systems, vol. 24, 2011.
[24] Q. Le Lidec, I. Kalevatykh, I. Laptev, C. Schmid, and J. Carpentier, "Differentiable simulation for physical system identification," IEEE Robotics and Automation Letters, vol. 6, no. 2, pp. 3413-3420, 2021.
[25] S. Diamond and S. Boyd, "CVXPY: A Python-embedded modeling language for convex optimization," Journal of Machine Learning Research, vol. 17, no. 83, pp. 1-5, 2016.
[26] A. Agrawal, R. Verschueren, S. Diamond, and S. Boyd, "A rewriting system for convex optimization problems," Journal of Control and Decision, vol. 5, no. 1, pp. 42-60, 2018.
[27] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd, "OSQP: an operator splitting solver for quadratic programs," Mathematical Programming Computation, vol. 12, no. 4, pp. 637-672, 2020.
[28] C. A. Floudas and P. M. Pardalos, Encyclopedia of Optimization. Springer Science \& Business Media, 2008.
[29] J. J. Moré and D. C. Sorensen, "Computing a trust region step," Siam Journal on Scientific and Statistical Computing, vol. 4, pp. 553-572, 1983.


[^0]:    *Corresponding author
    ${ }^{1}$ Inria
    ${ }^{2}$ Département d'Informatique de l'École normale supérieure, (ENS-PSL, CNRS, Inria)
    ${ }^{3}$ Center for Data Science, New York University

