A Bilevel Approach for Parameter Learning in Inverse Problems

Gernot Holler[†¶] Karl Kunisch[§‖] Richard C. Barnard[‡**]

July 27, 2018

Abstract

A learning approach for selecting regularization parameters in multi-penalty Tikhonov regularization is investigated. It leads to a bilevel optimization problem, where the lower level problem is a Tikhonov regularized problem parameterized in the regularization parameters. Conditions which ensure the existence of solutions to the bilevel optimization problem are derived, and these conditions are verified for two relevant examples. Difficulties arising from the possible lack of convexity of the lower level problems are discussed. Optimality conditions are given provided that a reasonable constraint qualification holds. Finally, results from numerical experiments used to test the developed theory are presented.

Key words. parameter learning, Tikhonov regularization, bilevel optimization, multi-penalty regularization

1 Introduction

Tikhonov regularization is a well-known method for regularization of ill-posed inverse problems, see e.g. [3, 14, 27, 23, 31, 32]. Given only a noisy measurement y_δ of some outcome $y_\dagger \in Y$, and assuming that the inverse problem consists in finding $u_\dagger \in U_{ad}$ such that

$$S(u_\dagger) = y_\dagger, \tag{1.1}$$

[†]Institute for Mathematics and Scientific Computing, University of Graz, Heinrichstr. 36, 8010 Graz, Austria, email: gernot.holler@uni-graz.at .

[§]Institute for Mathematics and Scientific Computing, University of Graz, Heinrichstr. 36, 8010 Graz, Austria, and Radon Institute, Austrian Academy of Sciences, Linz, Austria, email: karl.kunisch@uni-graz.at .

[‡]Oak Ridge National Laboratory, Oak Ridge, TN 37831, USA, email: barnardrc@ornl.gov .

[¶]The author gratefully acknowledges support from the International Research Training Group IGDK1754, funded by the DFG and FWF.

[‖]The author gratefully acknowledges partial support by the ERC advanced grant 668998 (OCLOC) under the EU's H2020 research program.

where S is a mapping from a subset U_{ad} of a Banach space U to a Banach space Y, in (generalized) Tikhonov regularization a solution to (1.1) is approximated by solving

$$\min_{u \in U_{ad}} \mathcal{J}_{\alpha,y_\delta}(u) \equiv \|S(u) - y_\delta\|^2 + \alpha \cdot \Psi(u), \qquad (\mathcal{P}_{\alpha,y_\delta})$$

for suitable choices of i) a discrepancy function, here chosen to measure a squared norm distance to the noisy measurement, ii) a (vector valued) penalty function $\Psi \colon U \to [0,\infty]^r$, and iii) a vector of regularization parameters $\alpha \in (0,\infty)^r$. Regularization parameters are used to balance the relative importance of the discrepancy and the penalty term.

Which discrepancy and penalty functions should be chosen depends heavily on the specific application. For choosing the regularization parameters, many general strategies have been proposed. They are usually divided into a posteriori, a priori, and heuristic parameter choice strategies. When discussing parameter choice strategies, it is frequently assumed that a noise level $\delta > 0$ such that $\|y_\dagger - y_\delta\|^2 \le \delta$ is known. A posteriori strategies determine regularization parameters based on knowledge of the noisy measurement and the noise level. A famous example for an a posteriori parameter choice strategy is Morozov's discrepancy principle, where regularization parameters are chosen such that for solutions to $(\mathcal{P}_{\alpha,y_\delta})$ the discrepancy is of the same order as the noise level, see e.g. [14, Section 4.3]. In a priori strategies, regularization parameters are chosen depending only on the noise level. A popular a priori strategy is to choose regularizations parameters as certain powers of the noise level, which, under reasonable assumptions, implies that solutions to $(\mathcal{P}_{\alpha,y_\delta})$ converge to so-called Ψ-minimum solutions as the noise level vanishes and α goes to zero. We refer to [14, Theorem 10.3] for more details. Heuristic strategies determine regularization parameters based solely on the noisy measurement. Prominent examples for heuristic strategies are the L-curve method [18, 19], and generalized cross validation [15]. Most of the existing parameter choice strategies address the case of a single regularization parameter, or, in cases where generalizations for determining multiple regularization parameters exist, become quite involved when one has to deal with a larger number of parameters.

The learning problem In this paper we consider a basic learning approach for selecting regularization parameters. The idea is to choose these parameters based on their performance on a training set. In the simplest case, the training set consists of a single vector $(y_\dagger, u_\dagger, y_\delta)$, where $(u_\dagger, y_\dagger)$ is a ground truth input-output pair, i.e.

$$S(u_\dagger) = y_\dagger.$$

and y_δ is the noisy measurement of $y_\dagger$ usually available in practice. Given such data, for every choice of α we can compute the distance between solutions u_α to the regularized problem $(\mathcal{P}_{\alpha,y_\delta})$ and the ground truth $u_\dagger$. This is used in the learning process where we aim at finding the regularization parameter α^* for which a solution u_{α^*} to $(\mathcal{P}_{\alpha^*,y_\delta})$ has minimal distance to the ground truth $u_\dagger$ over all parameters within a parameter interval $[\underline{\alpha}, \bar{\alpha}]$, where $0 < \underline{\alpha} < \bar{\alpha} < \infty$ are a-priori chosen r-dimensional bounds. This leads to the following problem:

$$\text{"}\min_{\alpha \in [\underline{\alpha},\bar{\alpha}]}\text{"} \; \|u_\dagger - u_\alpha\|^2 \quad \text{s.t.} \quad u_\alpha \in \operatorname*{arg\,min}_{u \in U_{ad}} \|S(u) - y_\delta\|^2 + \alpha \cdot \Psi(u). \qquad (1.2)$$

The quotation marks are used, since, if solutions to the Tikhonov regularized problems $(\mathcal{P}_{\alpha,y_\delta})$ are not unique, then it is not a-priori clear which solutions to choose. One possibility to overcome this difficulty is to look for α such that the minimal distance to the ground truth solution over all solutions to ($\mathcal{P}_{\alpha,y_\delta}$) is small. This is called the optimistic position and leads to the following problem:

$$\min_{\alpha\in[\underline{\alpha},\bar{\alpha}]}\ \min_{u_\alpha\in U_{ad}} \|u_\dagger - u_\alpha\|^2 \quad \text{s.t.} \quad u_\alpha \in \operatorname*{arg\,min}_{u\in U_{ad}} \|S(u)-y_\delta\|^2 + \alpha\cdot\Psi(u). \tag{1.3}$$

Another possibility is to look for α such that the maximal distance to the ground truth solution over all solutions to ($\mathcal{P}_{\alpha,y_\delta}$) is small. This is called the pessimistic position and amounts to the following problem:

$$\min_{\alpha\in[\underline{\alpha},\bar{\alpha}]}\ \max_{u_\alpha\in U_{ad}} \|u_\dagger - u_\alpha\|^2 \quad \text{s.t.} \quad u_\alpha \in \operatorname*{arg\,min}_{u\in U_{ad}} \|S(u)-y_\delta\|^2 + \alpha\cdot\Psi(u). \tag{1.4}$$

Here, we only consider the optimistic position (1.3). From now on we call (1.3) the learning problem, since regularization parameters should be learned by solving it. Conceptually, the learning problem is an optimization problem in two variables, which is constrained by requiring that one variable is a solution to a lower level (optimization) problem which is parametrized in the other variable. In the literature, problems of this type are called bilevel optimization problems, see e.g. [13]. If the training set consists of multiple data vectors $\{({y_\dagger}^k, {u_\dagger}^k, {y_\delta}^k)\}_{k=1}^N$, we determine regularization parameters as solutions to

$$\min_{\alpha\in[\underline{\alpha},\bar{\alpha}]}\ \min_{u_\alpha^k\in U_{ad},\,1\le k\le N} \sum_{k=1}^{N} \|u_\dagger - u_\alpha^k\|^2 \quad \text{s.t.} \quad u_\alpha^k \in \operatorname*{arg\,min}_{u\in U_{ad}} \|S(u)-{y_\delta}^k\|^2 + \alpha\cdot\Psi(u). \tag{1.5}$$

In this work, except for the numerical experiments, we only analyze the single data case (1.3). However, the generalization to the case of multiple data vectors is straightforward, for instance, by considering the product spaces $\prod_{k=1}^{N} U$ and $\prod_{k=1}^{N} Y$ instead of U and Y, respectively.

Related work In terms of learning theory, the discussed method can be viewed as a supervised learning approach. From the large amount of work on supervised learning, we discuss only selected papers which are closest to our approach. The present work is motivated by a parameter learning method for variational image denoising models that has been studied in [24]. There, both smooth quadratic and non smooth ℓ_1 and $\ell_{\frac{1}{2}}$ type penalty functionals are analyzed in a finite dimensional setting. For efficient numerical treatment of the non smooth problems, semi smooth Newton methods are proposed. In a similar approach presented in [11], a bilevel problem is solved to learn suitable discrepancy functions. In the statistical sense, this can be interpreted as learning the noise model. The basic idea is to assume that the discrepancy function can be written as a sum of discrepancy terms, parametrized by (possibly spatially dependent) weights. The weights are then chosen based on their performance on a training set.

Following these works, learning regularization parameters for infinite dimensional problems is studied in [12], with a particular emphasis on image restoration problems. The authors prove the existence of solutions for a general class of non smooth regularization operators. Moreover, for a particular class of non smooth penalty functionals (that encompasses total variation, total generalized variation, and infimal-convolution total variation), it is shown that optimal regularization parameters are strictly positive under suitable assumptions. Numerical results are only presented for a finite set of admissible regularization parameters. However, in a follow up work [10], using a regularized version of the non smooth penalty functionals, an optimality system is derived. This optimality system is then used to perform numerical experiments in which the quality of obtained image reconstructions using optimal regularization parameters is compared for different choices of penalty functions. Finally, in [6], instead of regularizing the non-smooth lower level problem, it is suggested to replace the lower level problem constraint by a differentiable update rule, which is given as the n-th step in an iterative procedure to determine an approximate solution to the lower level problem.

Prior to the above mentioned works, the problem of learning optimal spectral filters for finite dimensional inverse problems was studied in [8]. This approach contains learning regularization parameters for standard Tikhonov regularization as a special case. To evaluate the performance of spectral filters, the authors use a variety of performance measures. Morever, they give a statistical motivation for their approach, which can be seen as average empirical (Bayes) risk minimization. This work was extended in [9] to include generalized Tikhonov regularization.

The problem of choosing penalty functionals from a parametrized class of functions, based on training data, is studied in [17]. In this work, the authors compare the learning approach to the Bayesian approach, for which a learning strategy could be used to estimate the prior probability density function. Finally, we would like to mention that our method is loosely related to the problem of determining an optimal experimental design, which, in the context of ill-posed inverse problems, is studied e.g. in [16, 30].

Contributions We extend the approach presented and further developed in [24] and [12, 10, 9], respectively, to the class of inverse problems, where the aim is to determine coefficients or controls in partial differential equations (PDEs). One of the novelties of this paper is, that in addition to continuous linear operators S with non closed range, we consider the case where S is nonlinear and $U_{ad} \neq U$. Moreover, in our context, S is typically only given as an implicit function, which we incorporate by assuming the relation

$$y = S(u) \quad \text{if and only if} \quad e(y, u) = 0 \quad \text{for } (y, u) \in Y \times U_{ad}.$$

Here, the state equation $e\colon Y \times U_{ad} \to Z$ corresponds to the weak-form of a PDE and Z is another Banach space. We state our hypotheses directly in terms of the state equation to facilitate the use of the presented results in practice.

A first contribution concerns the existence of solutions to the learning problem (1.3). We prove that existence of solutions can be established under essentially the same conditions which are typically needed to ensure that the lower level problem is well-posed. To emphasize this, we first recall some standard results and assumptions needed to show that the lower level problems has a desired structure. Then, we use these structural properties to give a simple proof for the existence of solutions.

A main contribution concerns the derivation of an optimality system. The main difficulty here is the non standard constraint, that feasible points must be solutions to the lower level problem. In cases where there are no control constraints, i.e. $U_{ad} = U$, and the lower level is convex and smooth, this difficulty can be overcome by utilizing that points are solutions to the lower level problem if and only if they satisfy a first order optimality condition, which in turn can be written as an equality constraint. However, in our context the lower level problem can not be expected to be convex, and thus the usual first order optimality conditions are no longer sufficient. For this reason, some additional work is required. Our approach can be summarized as follows: First, we provide a theorem which relates the learning problem (1.3) to a reformulation, in which the first order optimality condition of the lower level problem arises as a constraint. To be more precise, we show that under reasonable assumptions, a solution to the learning problem is at least a local solution to the reformulation. As a consequence, the problem of finding optimality conditions for the learning problem is reduced to finding optimality conditions for local solutions to the reformulation. To obtain optimality conditions for the reformulation, we use a Lagrange multiplier approach. For both, deriving optimality conditions and relating the learning problem to its reformulation, second order sufficient optimality conditions of the lower level problem are used. This is a natural requirement for such results.

We demonstrate the applicability of our assumptions by verifying them for two examples of PDE constrained optimization problems. In the first example, S is the linear solution operator to

$$-\gamma\Delta y + y = u \quad \text{in } \Omega, \quad \text{and} \quad y = 0 \quad \text{on } \partial\Omega, \tag{1.6}$$

for a given $\gamma > 0$. Choosing quadratic penalty functionals, the corresponding lower level problem is smooth and convex, and as we will see, the assumptions of our main results can be verified in a straightforward manner. In the second example, S is the non-linear solution operator to

$$-\nabla\cdot(u\,\nabla\, y) = f \quad \text{in } \Omega, \quad \text{and} \quad y = 0 \quad \text{on } \partial\Omega, \tag{1.7}$$

where $f \in L^2(\Omega)$ is a given function. For the reader interested in practical applications of parameter estimation in PDEs we refer to [4] and the references mentioned therein. While the requirements for existence of solutions can be verified relatively easily, ensuring the requirements for the results needed to derive optimality conditions is much harder for this constraint. For this purpose, we present a criterion involving the adjoint state of a critical point of the lower level problem, which, if satisfied, guarantees that all necessary requirements hold. This criterion can be interpreted as a smallness condition on the discrepancy.

Finally, we provide a numerical validation of our approach. Thereby we compute optimal regularization parameters for weighted H^1 regularization of the above examples. Additionally, for the first example, we create large training and validation sets to test how well learned regularization parameters are performing on structurally related data.

Structure of the work In Section 2 we provide a precise statement of the learning problem and introduce some notation. In Section 3 we recall some basic properties of the lower level problem, i.e. of the Tikhonov regularized problem. These properties are in turn used in Section 4 to show that the learning problem has a solution under standard assumptions. In Section 5 we discuss the derivation of optimality conditions for the learning problem. Examples for possible applications are given in Section 6. Results from numerical experiments are presented in Section 7. In Section 8, we round off our contributions with a brief discussion of possible future work.

2 Problem statement

In the following we present the general setting of the learning problem to be considered in this work.

$$(\mathcal{LP}) \qquad \begin{cases} \min\limits_{\alpha\in[\underline{\alpha},\bar{\alpha}],\,(y_\alpha,u_\alpha)\in Y\times U_{ad}} \|u_\alpha - u_\dagger\|_{\tilde{U}}^2 \qquad \text{subject to} \\[2ex] (y_\alpha,u_\alpha) \in \operatorname*{arg\,min}\limits_{\substack{u\in U_{ad}\\ y\in Y}} \left\{ \frac{1}{2m}\sum\limits_{j=1}^{m} \|y - y_{\delta j}\|_{\tilde{Y}}^2 + \alpha\cdot\Psi(u) \mid e(y,u) = 0 \right\}, \end{cases}$$

where $m, r \in \mathbb{N}$, and

- U_{ad} is a subset of a reflexive Banach space U,
- Y is a reflexive Banach space,
- $\tilde{U}$ is a Hilbert space such that U is continuously embedded in $\tilde{U}$,
- $\tilde{Y}$ is a Hilbert space such that Y is continuously embedded in $\tilde{Y}$,
- $u_\dagger \in \tilde{U}$ is the ground truth control, and $y_{\delta j} \in \tilde{Y}$, $1 \le j \le m$, are noisy measurements of the ground truth state,
- $e\colon Y \times U_{ad} \to Z$ represents equality constraints in a reflexive Banach space Z,
- $\Psi_i\colon U \to [0, \infty]$, $1 \le i \le r$, are penalty functionals, and

$$\Psi := (\Psi_1, \dots, \Psi_r)^T,$$

- $\underset{\sim}{\alpha}, \bar{\alpha} \in \mathbb{R}^r$ are bounds for the regularization parameters with

$$0 < \underset{\sim}{\alpha} \le \bar{\alpha} < \infty,$$

  where the inequalities should be understood element wise.

Instead of working with an explicit solution operator S, we consider a more general implicit formulation by requiring that for feasible $(y, u) \in Y \times U_{ad}$ it holds that

$$e(y, u) = 0. \tag{2.1}$$

If for each $u \in U_{ad}$ there exists a unique $y \in Y$ such that (2.1) holds, then a solution operator S can be defined by setting

$$y = S(u) \quad \text{if and only if} \quad e(y, u) = 0 \quad \text{for } (y, u) \in Y \times U_{ad}.$$

The so-called lower level problem

$$(\mathcal{P}_{\alpha, y_\delta}) \qquad \begin{cases} \min\limits_{(y,u) \in Y \times U} \mathcal{J}_{\alpha, y_\delta}(y, u) \equiv \frac{1}{2m} \sum\limits_{j=1}^{m} \|y - y_{\delta j}\|_{\tilde{Y}}^2 + \alpha \cdot \Psi(u) \quad \text{subject to} \\ u \in U_{ad} \quad \text{and} \quad e(y, u) = 0, \end{cases}$$

which depends on the parameter $\alpha \in [\underset{\sim}{\alpha}, \bar{\alpha}]$, is a multi-penalty Tikhonov regularized inverse problem. We let

$$F_{ad} := \{(y, u) \in Y \times U \mid u \in U_{ad} \text{ and } e(y, u) = 0\}$$

denote the set of feasible points of the lower level problem. To fix ideas, typical choices for the used spaces are

$$U = H^1(\Omega), \quad Y = H_0^1(\Omega), \quad \tilde{Y} = L^2(\Omega), \quad \tilde{U} = L^2(\Omega),$$

where Ω is a bounded Lipschitz domain. Concrete examples for the state equation are given in Section 6.

2.1 Basic assumptions

The following assumptions are frequently invoked throughout this work.

(H1) The feasible control set U_{ad} is convex and closed in U.

(H2) The feasible set of the lower level problem F_{ad} is non-empty.

(H3) For every sequence (y^n, u^n) in $Y \times U_{ad}$ and $(\bar{y}, \bar{u}) \in Y \times U_{ad}$ such that

$$e(y^n, u^n) = 0 \quad \text{for all } n \in \mathbb{N}, \quad \text{and} \quad (y^n, u^n) \rightharpoonup (\bar{y}, \bar{u}),$$

it follows that

$$e(\bar{y}, \bar{u}) = 0.$$

(H4) For every sequence (y^n, u^n) in F_{ad} it holds that if (u^n) is bounded in U, then (y^n) is bounded in Y.

(H5) The function

$$\sum_{i=1}^{r} \Psi_i$$

is coercive on U and proper on F_{ad}.

(H6) The penalty functionals Ψ_i, $1 \leq i \leq r$, are weakly lower semi continuous on U.

3 The lower level problem

When we discuss existence of solutions and optimality conditions for the learning problem in Section 4 and 5, respectively, we frequently make use of basic properties of the lower level problem. In this section these properties are derived. Throughout this section we always assume that $\alpha \in [\underline{\alpha}, \bar{\alpha}]$.

3.1 Existence of solutions

Proposition 3.1 (Existence of solutions) If (H1)–(H6) hold, then $(\mathcal{P}_{\alpha, y_\delta})$ has a solution.

Proof. By (H1) the set $Y \times U_{ad}$ is closed and convex, and thus weakly closed [5, Theorem 3.7 on p.60]. It is then a direct consequence of (H3) that F_{ad} is weakly sequentially closed. From (H4)–(H5) and the assumption that $\alpha > 0$, it follows that $\mathcal{J}_{\alpha, y_\delta}$ is coercive on F_{ad}. The mapping

$$(y, u) \mapsto \frac{1}{2m} \sum_{j=1}^{m} \|y - y_{\delta j}\|_{\tilde{Y}}^2$$

is weakly lower semi continuous as a convex continuous function [5, Corollary 3.9 on p.61]. In combination with (H6) this implies that $\mathcal{J}_{\alpha, y_\delta}$ is weakly lower semi continuous on $Y \times U$. Since it is well-known that a weakly lower semi continuous and coercive function attains a minimum on a non empty and weakly sequentially closed subset of a reflexive Banach space, the proof is complete. □

Remark 3.1. As an immediate consequence of Proposition 3.1, we obtain that the feasible set of the learning problem $(\mathcal{LP})$ is non empty.

3.2 Stability

One of the reasons for regularizing an inverse problem is lack of stability with respect to the data. It is thus expected that stability, at least in some sense, holds for the Tikhonov regularized problem ($\mathcal{P}_{\alpha,y_\delta}$). Indeed, as stated below in Corollary 3.1, stability can be guaranteed under reasonable assumptions. Before we begin working towards this result, we need to clarify what we mean by stability (in particular in the context of problems with possibly non unique solutions).

Definition 3.1 (Stability with respect to the data) We say that ($\mathcal{P}_{\alpha,y_\delta}$) is stable with respect to the data if and only if the following holds: ($\mathcal{P}_{\alpha,y_\delta}$) has a solution for every $(\alpha, y_\delta) \in [\underline{\alpha}, \bar{\alpha}] \times \tilde{Y}^m$, and for every sequence $({y_\delta}^n)$ in $\tilde{Y}^m$ such that

$$
{y_\delta}^n \to y_\delta,
$$

it follows that every sequence (y^n, u^n) of corresponding solutions to ($\mathcal{P}_{\alpha,{y_\delta}^n}$) has an accumulation point, and every such accumulation point is a solution to ($\mathcal{P}_{\alpha,y_\delta}$).

Remark 3.2. If ($\mathcal{P}_{\alpha,y_\delta}$) has a unique solution, then it is straightforward to verify that stability with respect to the data is equivalent to requiring that every sequence (y^n, u^n) as in Definition 3.1 is converging to the unique solution of ($\mathcal{P}_{\alpha,y_\delta}$).

Recall that in the learning problem we minimize the distance to the ground truth control over the set of all feasible regularization parameters and corresponding solutions to the lower level problem. It is useful to know whether the lower level problem is stable with respect to the regularization parameters.

Definition 3.2 (Stability with respect to the regularization parameters) We say that ($\mathcal{P}_{\alpha,y_\delta}$) is stable with respect to the regularization parameters if and only if the following holds: ($\mathcal{P}_{\alpha,y_\delta}$) has a solution for every $(\alpha, y_\delta) \in [\underline{\alpha}, \bar{\alpha}] \times \tilde{Y}^m$, and for every sequence (α^n) in $[\underline{\alpha}, \bar{\alpha}]$ such that

$$
\alpha^n \to \alpha,
$$

it follows, that every sequence (y^n, u^n) of corresponding solutions to ($\mathcal{P}_{\alpha^n,y_\delta}$) has an accumulation point, and every such accumulation point is a solution to ($\mathcal{P}_{\alpha,y_\delta}$).

As a first step towards showing stability, we prove the following lemma, which states that under standard assumptions at least weak stability can be guaranteed with respect to both the data and the regularization parameters.

Lemma 3.1 (Weak stability) Assume that (H1)–(H6) hold, and let (α^n, y_δ^n) be a sequence in $[\underline{\alpha}, \bar{\alpha}] \times \tilde{Y}^m$ such that

$$
(\alpha^n, y_\delta^n) \to (\alpha, y_\delta).
$$

Then every sequence (y^n, u^n) of solutions to ($\mathcal{P}_{\alpha^n,{y_\delta}^n}$) has a subsequence (y^{n_k}, u^{n_k}) converging weakly to a solution $(\bar{y}, \bar{u})$ of ($\mathcal{P}_{\alpha,y_\delta}$), and

$$
\lim_{k\to\infty} \Psi(u^{n_k}) = \Psi(\bar{u}).
$$

Proof. A proof is given in the Appendix. □

Strong convergence as in Definition 3.1 and 3.2, and thus stability, can be achieved if the following additional assumptions are satisfied.

(H7) For every sequence (u^n) in U and $u \in U$ it holds that, if

$$u^n \rightharpoonup u \quad \text{and} \quad \Psi(u^n) \to \Psi(u),$$

then it follows that $u^n \to u$.

(H8) For each $u \in U_{ad}$ there exists a unique $y(u) \in Y$ such that

$$e(y(u), u) = 0,$$

and the mapping

$$u \mapsto y(u)$$

is continuous from U_{ad} to Y.

Remark 3.3. Condition (H7) is known to hold, for instance, if

$$\| \cdot \|_U = \sum_{i=1}^{r} \Psi_i$$

and U is a uniformly convex Banach space [5, Proposition 3.32. on p.78]. Note that every Hilbert space is a uniformly convex Banach space.

The following corollary, which, under reasonable assumptions, guarantees stability for the lower level problem, summarizes the considerations in this subsection.

Corollary 3.1 (Stability) If (H1)–(H8) hold, then $(\mathcal{P}_{\alpha,y_\delta})$ is stable with respect to both the data and the regularization parameters.

Proof. In combination with (H7) and (H8) this is a direct consequence of Lemma 3.1. □

4 Existence of solutions of the learning problem

Using results from the previous section, we can now apply standard arguments to prove that the learning problem has a solution.

Theorem 4.1 If (H1)–(H6) hold, then $(\mathcal{LP})$ has a solution.

Proof. We begin by showing that the feasible set of $(\mathcal{LP})$, which is given by

$$\mathcal{F} := \{(\alpha, y, u) \in [\underline{\alpha}, \bar{\alpha}] \times F_{ad} \mid (y, u) \text{ solves } (\mathcal{P}_{\alpha,y_\delta})\},$$

is non empty and weakly sequentially compact. The non emptiness of $\mathcal{F}$ follows from Proposition 3.1. In order to prove that $\mathcal{F}$ is weakly sequentially compact, we argue as follows: As a consequence of the Bolzano-Weierstraß theorem, every sequence (α^n, y^n, u^n) in $\mathcal{F}$ has a subsequence $(\alpha^{n_k}, y^{n_k}, u^{n_k})$ such that for some $\alpha^* \in [\underline{\alpha}, \bar{\alpha}]$

$$\alpha^{n_k} \to \alpha^*. \tag{4.1}$$

Utilizing that $(\mathcal{P}_{\alpha,y_\delta})$ is weakly stable with respect to the regularization parameters (Lemma 3.1) we can assume, possibly after taking another subsequence, that in addition to (4.1)

$$(y^{n_k}, u^{n_k}) \rightharpoonup (y^*, u^*) \tag{4.2}$$

for some $(y^*, u^*) \in F_{ad}$ which solves $(\mathcal{P}_{\alpha^*,y_\delta})$. Since $(\alpha^*, y^*, u^*) \in \mathcal{F}$, this proves that $\mathcal{F}$ is weakly sequentially compact. In view of the fact that a weakly lower semi continuous function attains a minimum on a non empty and weakly sequentially compact set (see e.g. [22, Theorem 2.3 on p.8]), it remains to show that the mapping

$$(\alpha, y, u) \mapsto \|u - u_\dagger\|_{\tilde{U}}^2$$

is weakly lower semi continuous on $\mathcal{F}$. This follows from [5, Corollary 3.9 on p.61], and thus the proof complete. □

5 Optimality conditions

The aim of this section is to derive optimality conditions for the learning problem. In this context, as we will see, a second order sufficient optimality condition for the lower level problem plays the role of a constraint qualification. We begin by stating our basic hypotheses, which are assumed to hold throughout this section: There exists an open neighborhood V of U_{ad} with the following properties:

(B1) The state equation e is defined and twice continuously F-differentiable on $Y \times V$.

(B2) The penalty function Ψ is twice continuously F-differentiable on V.

(B3) For each $u \in V$ there is a unique $y \in Y$ such that

$$e(y, u) = 0. \tag{5.1}$$

Moreover, $e_y(y, u) \in \mathcal{L}(Y, Z)$ is bijective for all $(y, u) \in Y \times V$ satisfying (5.1).

5.1 The lower level problem

Let us briefly recall some standard results on optimality conditions for the lower level problem. First order necessary optimality conditions are usually expressed using the notion of a Karush–Kuhn–Tucker (KKT) point.

Definition 5.1 (KKT point) We say that $(y, u, \lambda) \in Y \times U \times Z'$ is a KKT point of $(\mathcal{P}_{\alpha,y_\delta})$, if the following statements hold:

$$\langle \alpha \cdot \Psi_u(u) + \lambda e_u(y, u), v - u\rangle_{U',U} \geq 0, \quad \text{for all } v \in U_{ad}, \tag{optimality}$$

$$\frac{1}{m}\sum_{j=1}^{m}(y - y_{\delta j}) + \lambda e_y(y, u) = 0, \tag{adjoint equation}$$

$$u \in U_{ad}, \quad e(y, u) = 0. \tag{state equation}$$

It is well known that if a solution (y^*, u^*) to $(\mathcal{P}_{\alpha,y_\delta})$ satisfies the constraint qualification in the sense that $e_y(y^*, u^*)$ is bijective, then there exists a unique $\lambda^* \in Z'$ such that (y^*, u^*, λ^*) is a KKT point of $(\mathcal{P}_{\alpha,y_\delta})$, see e.g. [25]. Henceforth we refer to the system in Definition 5.1 as first order necessary optimality condition. In order to write second order sufficient optimality conditions in a compact form, we now introduce the Lagrange function.

Definition 5.2 (Lagrange function) We define the Lagrange function $\mathcal{L}\colon [\underline{\alpha}, \bar{\alpha}] \times Y \times U_{ad} \times Z' \to \mathbb{R}$ of the lower level problem by

$$\mathcal{L}(\alpha, y, u, \lambda) := \mathcal{J}_{\alpha,y_\delta}(y, u) + \lambda e(y, u)$$

for $(\alpha, y, u, \lambda) \in [\underline{\alpha}, \bar{\alpha}] \times Y \times U_{ad} \times Z'$.

Second order sufficient optimality conditions have many important practical implications. The condition presented in the following definition is one of the most often used ones in practice. For a discussion of the role of second order sufficient optimality conditions in PDE constrained optimization, we refer to [7] and the references given therein.

Definition 5.3 (Second order sufficient optimality condition) We say that (y, u) satisfies the second order sufficient optimality condition of $(\mathcal{P}_{\alpha,y_\delta})$ if $e_y(y, u)$ is bijective and there exists $\lambda \in Z'$ and $\eta > 0$ such that (y, u, λ) is a KKT point and

$$D^2_{(y,u)}\mathcal{L}(\alpha, y, u, \lambda)[(\delta_y, \delta_u), (\delta_y, \delta_u)] \geq \eta \|(\delta_y, \delta_u)\|^2_{Y\times U}$$

for all $(\delta_y, \delta_u) \in \ker De(y, u) \cap (Y \times (U_{ad} - U_{ad}))$.

If (y^*, u^*) satisfies the second order sufficient optimality condition of $(\mathcal{P}_{\alpha,y_\delta})$, then (y^*, u^*) is a local solution to $(\mathcal{P}_{\alpha,y_\delta})$, see e.g. [28].

5.2 KKT reformulation and optimality conditions

A major obstacle when deriving optimality conditions for bilevel optimization problems, consists in the fact that the constraint that feasible points must be solutions to a lower level problem can not always be expressed in standard form, i.e. via convex, equality and cone constraints. A standard technique to partially overcome this issue, consists in introducing the following auxiliary problem, which is usually referred to as the KKT reformulation of a bilevel problem:

$$(\mathcal{LP}_{\mathrm{KKT}}) \qquad \min_{\alpha\in[\underline{\alpha},\bar{\alpha}],\,(y,u,\lambda)\in Y\times U_{ad}\times Z'} \|u - u_\dagger\|^2_{\tilde{U}} \quad \text{s.t.} \quad (y, u, \lambda) \text{ is a KKT point of } (\mathcal{P}_{\alpha,y_\delta}).$$

Note, that in order to be feasible for the KKT reformulation, points only need to satisfy the first order necessary optimality condition, instead of the generally stronger requirement that they must be (global) solutions to the lower level problem.

Remark 5.1. In $(\mathcal{LP}_{\mathrm{KKT}})$, the Lagrange multiplier λ is included as an additional optimization variable. In our case there exists a unique feasible state $y(u)$ and a unique solution to the adjoint equation $\lambda(u)$ for every $u \in U_{ad}$. Thus, it is possible to equivalently rewrite $(\mathcal{LP}_{\mathrm{KKT}})$ as a problem in the controls u only. This will be used in the theorem below.

The main reason for introducing the KKT reformulation, is that in general its constraints are easier to handle than the constraints of the original problem. For example, if there are no control constraints in the lower level problem, then the constraints of the KKT reformulated problem consist only of equality and convex constraints. Before discussing how to obtain optimality conditions for $(\mathcal{LP}_{\mathrm{KKT}})$, we need to investigate the relation between the KKT reformulation and the original problem. Clearly, if the lower level problem is convex for every parameter $\alpha \in [\underline{\alpha}, \bar{\alpha}]$, then both are equivalent. In general, this is not the case since KKT points of the lower level problem need not be solutions to the lower level problem. However, since we are only interested in $(\mathcal{LP}_{\mathrm{KKT}})$ to obtain optimality conditions for the learning problem, for our purposes it is sufficient to know under which conditions a solution to the learning problem is guaranteed to be a local solution to $(\mathcal{LP}_{\mathrm{KKT}})$. This question is addressed in the following theorem. Thereby, it is important to keep in mind that feasible points of the learning problem are defined as global solutions to the lower level problem.

Theorem 5.1 Let (α^*, y^*, u^*) be a solution to $(\mathcal{LP})$ and assume that the following statements hold:

(C1) $(\mathcal{P}_{\alpha^*,y_\delta})$ is stable with respect to the regularization parameters.

(C2) (y^*,u^*) satisfies the second order sufficient optimality condition of $(\mathcal{P}_{\alpha^*,y_\delta})$.

(C3) (y^*,u^*) is the unique solution to $(\mathcal{P}_{\alpha^*,y_\delta})$.

Then there exists a unique $\lambda^* \in Z'$ such that $(\alpha^*, y^*, u^*, \lambda^*)$ is a local solution to $(\mathcal{LP}_{\text{KKT}})$.

Proof. We define

$$F(\alpha,u) \coloneqq \mathcal{J}_{\alpha,y_\delta}(y(u),u) \tag{5.3}$$

as the reduced cost functional of the lower level problem, and for every $u \in V$ we let $\mathcal{N}(u)$ denote the polar cone to the set U_{ad} at the point u, i.e.

$$\mathcal{N}(u) = \begin{cases} \{u' \in U' \mid \langle u', v-u\rangle \leq 0 \text{ for all } v \in U_{ad}\} & \text{if} \quad u \in U_{ad} \\ \emptyset & \text{if} \quad u \notin U_{ad}. \end{cases}$$

Note that for every $u \in V$ there exists a unique state $y(u)$ and Lagrange multiplier $\lambda(u)$ satisfying the associated state and adjoint equation, respectively. It is straightforward to show that $(y(u), u, \lambda(u))$ is a KKT point of $(\mathcal{P}_{\alpha,y_\delta})$ if and only if u satisfies the variational inequality

$$0 \in D_u F(\alpha,u) + \mathcal{N}(u). \tag{5.4}$$

A reduced KKT formulation of $(\mathcal{LP})$ can thus be written as follows.

$$\min_{(\alpha,u)\in[\underline{\alpha},\bar{\alpha}]\times V} \|u-u_\dagger\|^2_{\tilde{U}} \quad \text{subject to} \quad 0 \in D_u F(\alpha,u) + \mathcal{N}(u). \tag{5.5}$$

We now divide the proof into three steps.

Step 1: We claim: If (α^*, u^*) is a local solution to (5.5), then $(\alpha^*, y^*, u^*, \lambda^*)$ is a local solution to $(\mathcal{LP}_{\text{KKT}})$. In fact, if (α^*, u^*) is a solution to (5.5) restricted to $J(\alpha^*) \times V(\alpha^*)$, where $J(\alpha^*)$ and $V(u^*)$ are neighbourhoods of α^* and u^*, respectively, then $(\alpha^*, y^*, u^*, \lambda^*)$ is solution to $(\mathcal{LP}_{\text{KKT}})$ restricted to $J(\alpha^*) \times Y \times V(u^*) \times Z'$.

Step 2: We claim: There exist neighborhoods $J(\alpha^*)$ of α^* and $V(u^*)$ of u^*, and a Lipschitz continuous function $\Phi\colon J(\alpha^*) \to V(u^*)$ such that for all $(\alpha,u) \in J(\alpha^*) \times V(u^*)$ it holds that

$$0 \in D_u F(\alpha,u) + \mathcal{N}(u) \quad \text{if and only if} \quad u = \phi(\alpha). \tag{5.6}$$

This follows immediately from a generalized implicit function theorem ([29, Corollary 2.2]) if we can verify the necessary requirements for this result. Thereby, the only requirement, which does not follow immediately from our assumptions, is the strong regularity of (5.4) at (α^*, u^*). However, to verify strong regularity it suffices to show that for every $u' \in U'$ there exists a unique $u \in U$ such that

$$u' \in D_{uu} F(\alpha^*, u^*)(u-u^*) + \mathcal{N}(u), \tag{5.7}$$

and that the mapping $u' \to u$ is Lipschitz continuous. To prove this, we consider the problem

$$\min_{u\in U_{ad}} -u'(u) + \frac{1}{2} D_{uu} F(\alpha^*, u^*)[u-u^*, u-u^*]. \tag{5.8}$$

The second order sufficient optimality condition (C2) implies that (5.8) has a unique solution. Moreover, u solves (5.8) if and only if (5.7) holds. It remains to show Lipschitz continuity of the mapping. For this purpose let $u_1', u_2' \in U'$ be arbitrary, and let $u_1, u_2 \in U$, respectively, be the corresponding solutions to (5.7). We have

$$D_{uu}F(\alpha^*, u^*)(u_1 - u_2) - u_1' + u_2' \in \mathcal{N}(u_2) - \mathcal{N}(u_1),$$

from which it follows that

$$D_{uu}F(\alpha^*, u^*)[u_1 - u_2, u_1 - u_2] \leq \langle u_1' - u_2', u_1 - u_2 \rangle.$$

Again, we apply the second order sufficient optimality condition to arrive at

$$\|u_1 - u_2\|_U \leq \frac{1}{\eta} \|u_1' - u_2'\|_{U'}.$$

This shows that the mapping $u' \to u$ is Lipschitz continuous, which concludes the second step.

Step 3: This is the final step of the proof. We let $J(\alpha^*)$, $V(u^*)$, and $\phi\colon J(\alpha^*) \to V(u^*)$ be as in the second step and claim that their exists an open neighbourhood $I(\alpha^*) \subseteq J(\alpha^*)$ of α^* such that $\phi(\alpha)$ is a global solution to the reduced lower level problem for every $\alpha \in I(\alpha^*)$. This we prove by contradiction. If our claim was false, then there would be a sequence (α^n) in $J(\alpha^*)$ such that

$$\alpha^n \to \alpha^*,$$

and an associated sequence (u^n) of solutions to $(\mathcal{P}_{\alpha^n, y_\delta})$ such that (u^n) does not intersect $V(u^*)$. Using the stability assumption (C1) and the uniqueness assumption (C3), it follows that (u^n) must converge to u^*. However, since the sequence (u^n) was chosen such that $u^n \notin V(u^*)$ for all $n \in \mathbb{N}$, this leads to a contradiction. This proves that (α^*, u^*) is a solution to the reduced bilevel problem (5.5) restricted to $I(\alpha^*) \times V(u^*)$, and consequently, using the first step, it follows that $(\alpha^*, y^*, u^*, \lambda^*)$ is a local solution to $(\mathcal{LP}_{\text{KKT}})$. This completes the proof.

□

As an important consequence of the above theorem, we get the following: Any solution to $(\mathcal{LP})$, for which the assumptions of Theorem 5.1 hold, satisfies the optimality conditions to be a local solution to $(\mathcal{LP}_{\text{KKT}})$. The derivation of optimality conditions for $(\mathcal{LP}_{\text{KKT}})$ which are convenient for numerical realization, are still impeded by the presence of control constraints in the lower level problem, which in turn lead to set valued constraints in $(\mathcal{LP}_{\text{KKT}})$. We therefore first consider the case without constraints on u. In Section 6, we show how one can still obtain optimality conditions for a particular example where control constraints are essential to ensure that the lower level problem is well posed. This will be done using a smoothed point wise projection on the set of feasible controls. To conclude this section, we finally present optimality conditions for $(\mathcal{LP}_{\text{KKT}})$ in the absence of control constraints, which in turn gives us optimality conditions for $(\mathcal{LP})$ under the conditions of Theorem 5.1.

Lemma 5.1 Assume that $U_{ad} = U$. Let $(\alpha^*, y^*, u^*, \lambda^*)$ be a local solution to $(\mathcal{LP}_{\text{KKT}})$ with (y^*, u^*) satisfying the second order sufficient optimality condition of $(\mathcal{P}_{\alpha^*, y_\delta})$. Then there exists a unique triple $(p^*, q^*, z^*) \in Y \times U \times Z'$ such that

$$\langle \Psi_u(u^*) q^*, \alpha - \alpha^* \rangle_2 \geq 0, \quad \forall \alpha \in [\underline{\alpha}, \bar{\alpha}], \tag{5.9a}$$

$$p^* + \lambda^* e_{yy}(y^*, u^*) p^* + \lambda^* e_{yu}(y^*, u^*) q^* + z^* e_y(y^*, u^*) = 0, \tag{5.9b}$$

$$u^* - u_\dagger + \lambda^* e_{uy}(y^*, u^*) p^* + \alpha^* \cdot \Psi_{uu}(u^*) q^* + \lambda^* e_{uu}(y^*, u^*) q^* + z^* e_u(y^*, u^*) = 0, \tag{5.9c}$$

$$e_y(y^*, u^*) p^* + e_u(y^*, u^*) q^* = 0. \tag{5.9d}$$

Proof. We define

$$F(\alpha, y, u) := \mathcal{J}_{\alpha,y_\delta}(y, u).$$

In view of [25], it suffices to verify the regularity assumption consisting in the bijectivity of the mapping

$$(\delta_y, \delta_u, \delta_\lambda) \to \begin{pmatrix} F_{yy} + \lambda^* e_{yy} & F_{yu} + \lambda^* e_{yu} & e_y^* \\ F_{yu} + \lambda^* e_{yu} & F_{uu} + \lambda^* e_{yu} & e_u^* \\ e_y & e_u & 0 \end{pmatrix} \begin{pmatrix} \delta_y \\ \delta_u \\ \delta_\lambda \end{pmatrix}$$

from $Y \times U \times Z'$ to $Y' \times U' \times Z$, where we write $F_{yy} = F_{yy}(\alpha^*, y^*, u^*)$, $e_{yy} = e_{yy}(\alpha^*, y^*, u^*)$ et cetera. This follows from observing that for every $(y', u', z) \in Y' \times U' \times Z$, the quadratic problem

$$\begin{aligned} &\min_{(\delta_y,\delta_u)\in Y\times U} \quad D^2_{(y,u)}\mathcal{L}(\alpha^*, y^*, u^*, \lambda^*)[(\delta_y, \delta_u), (\delta_y, \delta_u)] - y'(\delta_y) - u'(\delta_u) \\ &\text{subject to} \quad De(y, u)(\delta_y, \delta_u) = z \end{aligned}$$

has a unique solution $(\delta_y^*, \delta_u^*, \delta_\lambda^*)$, which is characterized by

$$\begin{pmatrix} F_{yy} + \lambda e_{yy} & F_{yu} + \lambda^* e_{yu} & e_y^* \\ F_{yu} + \lambda^* e_{yu} & F_{uu} + \lambda^* e_{yu} & e_u^* \\ e_y & e_u & 0 \end{pmatrix} \begin{pmatrix} \delta_y^* \\ \delta_u^* \\ \delta_\lambda^* \end{pmatrix} = \begin{pmatrix} y' \\ u' \\ z \end{pmatrix}.$$

□

Let us briefly discuss the assumptions (C1)-(C3) needed in Theorem 5.1. As we have already seen in Corollary 3.1, if (H1)–(H8) are satisfied, then the stability required in (C1) holds. Condition (C2), which requires that the second order sufficient optimality condition holds, is quite standard, and also needed to ensure the existence of an optimality system for ($\mathcal{LP}_{\text{KKT}}$). Condition (C3) might seem very restrictive. First, let us point out that, unfortunately, as indicated by a counterexample in [20, Example 4.2.1], without the third condition the conclusion of Theorem 5.1 no longer remains true. However, notice that in (C3) we only require that the lower level problem has a unique solution for the optimal regularization parameter α^*. While in general unique solvability of the lower level problem for every regularization parameter can not be expected, there are ways to ensure unique solvability a-posteriori for a specific regularization parameter. To be more precise, given a critical point of the corresponding lower level problem, one can validate a criterion involving its adjoint state, which, if satisfied, guarantees that the given critical point is the unique global solution to the lower level problem. We provide such a criterion for a particular example in Section 6.2. This criterion, which can also be seen as a smallness condition on the discrepancy, will also be sufficient to ensure that the second order sufficient optimality condition (C2) is met. A similar criterion for a different class of optimal control problems is provided and thoroughly investigated in [2].

6 Examples

We now apply the developed theory for learning parameters for the regularization of two PDE constrained problems. In the first example, the lower level problem is quadratic and, as we will see, both existence of solutions and optimality conditions can be established quite easily using the results from the previous sections. In the second example, the state equation is bilinear. In this case, the lower level problem is not convex. For this example, we provide a criterion, which guarantees that the assumptions needed in Theorem 5.1 are satisfied, which in turn allows us to use the KKT reformulation to derive an optimality system.

6.1 Linear state equation

We consider ($\mathcal{LP}$) with $Y = H_0^1(\Omega)$, $U = H^1(\Omega)$, $\tilde{Y} = \tilde{U} = L^2(\Omega)$ and $Z = H^{-1}(\Omega)$. Here, $\Omega \subset \mathbb{R}^d$ is a bounded Lipschitz domain, where $d \in \mathbb{N}$. Furthermore, given $\gamma > 0$ we let the state equation $e\colon H_0^1(\Omega) \times H^1(\Omega) \to H^{-1}(\Omega)$ be defined as

$$e(y,u) = -\gamma\Delta y + y - u \quad \text{for } (y,u) \in H_0^1(\Omega) \times H^1(\Omega),$$

and assume that weighted H^1-regularization is used, i.e.

$$\Psi_i = \frac{1}{2}\|K_i \cdot\|_{L^2(\Omega)}^2, \quad \text{where} \quad K_i = \mathrm{I} \text{ and } \quad K_i = \partial_{x_{i-1}}, \quad \text{for } 2 \le i \le d+1.$$

In this example, the lower level problem is uniformly convex. Thus, it suffices to verify (H1)-(H8) and (B1)-(B3) to guarantee the existence of solutions and to obtain optimality conditions for this problem. First, we note that (H1) holds by definition and (B1) follows easily using that e is a linear and continuous operator. To verify (H3), one uses that linear continuous operators are weak-weak continuous. Conditions (H2), (H4), (H8), and (B3) can be seen as straightforward consequence of the Lax-Milgram lemma. Since the operators $K_i\colon H^1(\Omega) \to L^2(\Omega)$, $1 \le i \le d+1$, are linear and continuous, and

$$\|u\|_{H^1(\Omega)}^2 = \sum_{i=1}^{d+1} \|K_i u\|_{L^2(\Omega)}^2,$$

one can verify (H5), (H6), and (B2) by standard arguments. (H7) is of course satisfied since $U = H^1(\Omega)$ is a Hilbert space (see Remark 3.3).

6.2 Bilinear state equation

In this example, we discuss the estimation of the diffusion coefficient in a second order elliptic PDE using weighted H^k regularization. That is, we consider ($\mathcal{LP}$) with $Y = H_0^1(\Omega)$, $U = H^k(\Omega)$, $\tilde{Y} = \tilde{U} = L^2(\Omega)$ and $Z = H^{-1}(\Omega)$. Here, we assume that $\Omega \subset \mathbb{R}^d$ is a bounded Lipschitz domain, where $d \in \mathbb{N}$, and let the state equation $e\colon H_0^1(\Omega) \times U_{ad} \to H^{-1}(\Omega)$ be given by

$$e(y,u) = -\nabla \cdot (u \nabla y) - f \quad \text{for } (y,u) \in H_0^1(\Omega) \times U_{ad},$$

for some $f \in L^2(\Omega)$. We define the penalty function Ψ by setting

$$\Psi_\beta(u) = \frac{1}{2}\|\partial_\beta u\|_2^2 \quad \text{for } \beta \in \mathbb{N}_0^d \text{ with } |\beta|_1 \le k.$$

The set of feasible controls is chosen to be

$$U_{ad} \coloneqq \{u \in H^k(\Omega) \cap L^\infty(\Omega) \mid a \le u \le b \quad \text{a.e. in } \Omega\}$$

for given $0 < a < b < \infty$. While in the following we prove the existence of solutions for arbitrary dimension d and order of regularization k, to obtain optimality conditions we assume that $d = k = 2$. Similar arguments can be used to obtain optimality conditions whenever $k > d/2$. In the following discussion, we systematically identify the space of Lagrange multipliers $Z' = (H^{-1}(\Omega))'$ with $H_0^1(\Omega)$.

Existence of solutions Let us first establish the existence of solutions. In view of Theorem 4.1, it suffices to verify (H1)–(H6). Showing (H1) is straightforward. Conditions (H5) and (H6) can be verified similarly as for the previous example. Note that since for all $u \in U_{ad}$ and $y \in H_0^1(\Omega)$ we have

$$\int_\Omega u \nabla y \nabla y \geq a \| \nabla y\|_2^2,$$

where a is the pointwise lower bound on the diffusion coefficient, one obtains (H2) and (H4) as a direct consequence of the Lax-Milgram lemma in combination with the Poincaré inequality. To prove (H3), we argue as follows: Let $(y^n, u^n) \in F_{ad}$ and $(\bar{y}, \bar{u}) \in Y \times U_{ad}$ such that $(y^n, u^n) \rightharpoonup (\bar{y}, \bar{u})$. Since $H^k(\Omega)$ is compactly embedded in $L^2(\Omega)$, it follows that

$$u^n \to \bar{u} \quad \text{in } L^2(\Omega).$$

Now for arbitrary $v \in C_c^\infty(\Omega)$ we conclude that

$$u_n \nabla v \to u \nabla v \quad \text{in } (L^2(\Omega))^d,$$

using Hölder's inequality. This implies that

$$\int_\Omega u^n \nabla y^n \cdot \nabla v \to \int_\Omega \bar{u} \nabla \bar{y} \cdot \nabla v.$$

Thus, we have proven that

$$\int_\Omega \bar{u} \nabla \bar{y} \cdot \nabla v - fv = 0 \quad \text{for all } v \in C_c^\infty(\Omega).$$

Since $C_c^\infty(\Omega)$ is dense in $H_0^1(\Omega)$, a simple continuity argument shows that

$$\int_\Omega \bar{u} \nabla \bar{y} \cdot \nabla v - fv = 0 \quad \text{for all } v \in H_0^1(\Omega),$$

where we use that $\bar{u} \in L^\infty(\Omega)$. This is precisely what we needed to show.

Relation to the KKT reformulation We aim at applying the results from Section 5. For this purpose, we first need to show that (B1)–(B3) hold. To guarantee (B1) and (B3) we require that $H^k(\Omega)$ can be continuously embedded into $L^\infty(\Omega)$, which is the case if and only if $k > d/2$, see e.g. [1, Theorem 5.4, Example 5.25, and 5.26]). Clearly, (B2) is always satisfied.

For simplicity, from now on we let $d = 2$ and $k = 2$. In order to relate the problem to its KKT reformulation, we have to verify the conditions (C1)-(C3) in Theorem 5.1. Since we already know that (H1)–(H6) hold, the stability assumption (C1) follows from Corollary 3.1, if (H7) and (H8) hold. Since $H^2(\Omega)$ is a Hilbert space, (H7) follows directly. (H8) is a consequence of the implicit function theorem in combination with the Lax-Milgram lemma. To verify the remaining conditions, we make use of results on higher regularity. In particular, note that under our assumptions, for each $u \in U_{ad}$ both the associated state y and the solution to the adjoint equation p are in $H^2(\Omega)$, assuming that Ω has sufficient regularity, see [26, Theorem 10.1 on p. 188]. Using standard Sobolev embeddings, we may also assume that

$$\|v\|_{L^\infty(\Omega)} \leq C_1 \|v\|_{H^2(\Omega)} \quad \text{and} \quad \|v\|_{W^{1,4}(\Omega)} \leq C_2 \|v\|_{H^2(\Omega)}. \tag{6.1}$$

for constants $C_1, C_2 > 0$. The following estimate will be useful later on.

Lemma 6.1 For all $(y,u,p) \in H_0^1(\Omega) \times H^2(\Omega) \times H_0^2(\Omega)$ it holds that

$$| \int_\Omega u \nabla y \nabla p \,| \leq C \|u\|_{H^2(\Omega)} \|y\|_{L^2(\Omega)} \|p\|_{H^2(\Omega)}, \tag{6.2}$$

where $C := C_1 + C_2^2$ for constants C_1, C_2 as in (6.1).

Proof. Using partial integration, Hölders inequality, and the embeddings in (6.1), we estimate

$$\begin{aligned} | \int_\Omega u \nabla p \nabla y \,| = | \int_\Omega \nabla \cdot (u \nabla p)\, y \,| = | \int_\Omega y \nabla u \nabla p + \int_\Omega y\, u\, \Delta p \,| \leq \\ \|y\|_{L^2(\Omega)} \left(\| \nabla p \nabla u\|_{L^2(\Omega)} + \|u\|_{L^\infty(\Omega)} \|\Delta p\|_{L^2(\Omega)} \right) \\ \leq (C_2^2 + C_1) \|y\|_{L^2(\Omega)} \|p\|_{H^2(\Omega)} \|u\|_{H^2(\Omega)}. \end{aligned}$$

This is what we needed to show. □

In the following proposition we provide a criterion involving the adjoint state of a critical point of the lower level problem, which, if satisfied, guarantees that the second order sufficient optimality condition (C2) holds.

Proposition 6.1 Let $(\bar{y}, \bar{u}, \bar{p}) \in H_0^1(\Omega) \times H^2(\Omega) \times H_0^1(\Omega)$ be a KKT point of $(\mathcal{P}_{\alpha, y_\delta})$ and

$$\alpha_{\min} := \min\{\alpha_i \mid 1 \leq i \leq r\}.$$

If it holds that

$$\|\bar{p}\|_{H^2(\Omega)} < \frac{2\sqrt{\alpha_{min}}}{C},$$

where C is as in Lemma 6.1, then $(\bar{y}, \bar{u})$ satisfies the second order sufficient optimality condition of $(\mathcal{P}_{\alpha, y_\delta})$.

Proof. For arbitrary $\delta_u \in U_{ad} - U_{ad}$ we let $\delta_y \in Y$ be such that $(\delta_y, \delta_u) \in \ker De(\bar{y}, \bar{u})$, i.e.

$$\int_\Omega u \nabla \delta_y \nabla v + \int_\Omega \delta_u \nabla y \nabla v = 0 \quad \text{for all } v \in H_0^1(\Omega).$$

The claim now follows from the estimate

$$\begin{aligned} D^2_{(y,u)} \mathcal{L}(\alpha, \bar{y}, \bar{u}, \bar{p})[(\delta_y, \delta_u), (\delta_y, \delta_u)] = \|\delta_y\|^2_{L^2(\Omega)} + 2\alpha \cdot\, \Psi(\delta_u) + \int_\Omega \delta_u \nabla \delta_y \nabla \bar{p} \geq \\ \|\delta_y\|^2_{L^2(\Omega)} + \alpha_{\min} \|\delta_u\|^2_{H^2(\Omega)} - C \|\bar{p}\|_{H^2(\Omega)} \|\delta_u\|_{H^2(\Omega)} \|\delta_y\|_{L^2(\Omega)} \geq \\ \left(1 - \frac{C}{2\sqrt{\alpha_{\min}}} \|\bar{p}\|_{H^2(\Omega)} \right) \left(\|\delta_y\|^2_{L^2(\Omega)} + \alpha_{\min} \|\delta_u\|^2_{H^2(\Omega)} \right), \end{aligned}$$

where we use Lemma 6.1 and the Cauchy-type inequality

$$ab \leq \frac{1}{2\varepsilon^2}(a^2 + \varepsilon^4 b^2)$$

for $a = \|\delta_y\|_{L^2(\Omega)}$, $b = \|\delta_y\|_{H^2(\Omega)}$, and $\varepsilon = (\alpha_{\min})^{1/4}$.

□

The same condition on the adjoint state is also sufficient to ensure (C3), i.e. that the lower level problem has a unique solution for the optimal parameter.

Proposition 6.2 Let $(\bar{y}, \bar{u}, \bar{p}) \in H_0^1(\Omega) \times H^2(\Omega) \times H_0^1(\Omega)$ be a KKT point of $(\mathcal{P}_{\alpha,y_\delta})$ and

$$\alpha_{\min} := \min\{\alpha_i \mid 1 \leq i \leq r\}.$$

If it holds that

$$\|\bar{p}\|_{H^2(\Omega)} < \frac{2\sqrt{\alpha_{min}}}{C},$$

where C is as in Lemma 6.1, then $(\bar{y}, \bar{u})$ is the unique global solution to the lower level problem $(\mathcal{P}_{\alpha,y_\delta})$.

Proof. For every $(y, u) \in F_{ad}$ we have

$$0 = e(y,u) - e(\bar{y}, \bar{u}) = e_y(\bar{y}, \bar{u})[y - \bar{y}] + e_u(\bar{y}, \bar{u})[u - \bar{u}] + e_{yu}(\bar{y}, \bar{u})[y - \bar{y}, u - \bar{u}],$$

where we use that e is an affine bilinear function. This implies that

$$-\bar{p}(e_y(\bar{y}, \bar{u})[y - \bar{y}] + e_u(\bar{y}, \bar{u})[u - \bar{u}]) = \bar{p}e_{yu}(\bar{y}, \bar{u})[y - \bar{y}, u - \bar{u}]. \tag{6.3}$$

Using that $(\bar{y}, \bar{u}, \bar{p})$ is a KKT point and a quadratic expansion of $\mathcal{J}_{\alpha,y_\delta}$ around $(\bar{y}, \bar{u})$, we can estimate

$$\begin{aligned}
&\mathcal{J}_{\alpha,y_\delta}(y,u) - \mathcal{J}_{\alpha,y_\delta}(\bar{y}, \bar{u}) \geq \\
&\quad D_y\mathcal{J}_{\alpha,y_\delta}(\bar{y}, \bar{u})[y - \bar{y}] + D_u\mathcal{J}_{\alpha,y_\delta}(\bar{y}, \bar{u})[u - \bar{u}] + \|y - \bar{y}\|^2_{L^2(\Omega)} + \alpha_{\min}\|u - \bar{u}\|^2_{H^2(\Omega)} \geq \\
&\qquad -\bar{p}(e_y(\bar{y}, \bar{u})[y - \bar{y}] + e_u(\bar{y}, \bar{u})[u - \bar{u}]) + \|y - \bar{y}\|^2_{L^2(\Omega)} + \alpha_{\min}\|u - \bar{u}\|^2_{H^2(\Omega)}.
\end{aligned}$$

By (6.3), Lemma 6.1, and similar arguments as in the proof of Proposition 6.1, we arrive at

$$\mathcal{J}_{\alpha,y_\delta}(y,u) - \mathcal{J}_{\alpha,y_\delta}(\bar{y}, \bar{u}) \geq \left(1 - \frac{C}{2\sqrt{\alpha_{\min}}}\|\bar{p}\|_{H^2(\Omega)}\right)\left(\|y - \bar{y}\|^2_{L^2(\Omega)} + \alpha_{\min}\|u - \bar{u}\|^2_{H^2(\Omega)}\right),$$

from which the claim easily follows. □

Optimality conditions In this example, control constraints are essential to ensure that the lower level problem is well-posed. However, the case of control constraints is not covered in Lemma 5.1. To circumvent this issue, we consider a relaxed version of the problem, in which the state equation e is replaced by

$$\tilde{e}(y,u) = -\nabla \cdot (\phi_\varepsilon(u) \nabla y) - f \quad \text{for} \quad (y,u) \in H_0^1(\Omega) \times H^k(\Omega),$$

where $\phi_\varepsilon \colon \mathbb{R} \to \mathbb{R}$ is a smoothed point wise projection on $[a, b]$. The precise definition of ϕ_ε is given in the Appendix. One can show that the learning problem with this relaxed state equation fulfills the assumptions of Theorem 4.1, and thus has a solution. Optimality conditions can the be obtained via Theorem 5.1 and Lemma 5.1.

7 Numerical experiments

In this section we present two numerical experiments regarding learning regularization parameters in weighted H^1-regularization.

7.1 Linear state equation

In the first experiment, we consider the problem from Section 6.1, where the inverse problem for which we learn regularization parameters consists in estimating the forcing function in a second order elliptic PDE. Here, we let $\Omega = (-1,1) \times (-1,1)$ and $\gamma = 0.1$.

Single vector of data To create a single vector of data, we set $c_1 = 0.4$, $c_2 = 0.4$, $h = 2.5$, $d = 0.3$, $\kappa = 2$, and let $y_\dagger$ be the corresponding state for the ground truth control

$$u_\dagger(x_1, x_2) = \begin{cases} h & \text{if } |x_1 - c_1| < d \quad \text{and} \quad |x_2 - c_2| < d, \\ 2.5(\sin^2(\kappa\pi x_1) + x_2{}^2) & \text{else.} \end{cases}$$

The ground truth control is shown in Figure 1a. We discretize the problem on a 128×128 mesh using the standard five-point stencil for the Laplace operator. Noisy data measurements $y_{\delta j}$ are generated by point wise setting

$$y_{\delta j} = y_\dagger + \varepsilon \xi_j,$$

for $1 \leq j \leq m$, where ξ_j follows a normal distribution with mean 0 and standard deviation 1, and $\varepsilon := \epsilon \max |y_\dagger|$ with ϵ being the relative noise level. We consider the following regularization operators

$$K_1 := \mathrm{I}, \quad K_2 := \partial_{x_1}, \quad K_3 := \partial_{x_2}.$$

In Figure 2a, we plot the values of the bilevel cost functional, i.e. the squared distance between the recovered and ground truth control, in dependence of the regularization parameter when using the single operator K_1 for different noise levels. Figure 2b shows a contour plot of the bilevel cost functional in dependence of the regularization parameters using the operators K_2 and K_3 for 1% noise. Note that in both figures the bilevel cost functional attains a distinct minimum. This confirms the feasibility of the formulation of finding regularization parameters as a learning problem.

Training and validation sets To create multiple vectors of data for training and validation, we proceed similarly as for the single dataset. The difference consists in randomly choosing $\kappa \in [0,3]$, $h \in [0,4]$, $d \in [0,0.8]$ and $c_1, c_2 \in [0,1]$ following a uniform distribution.

Used methods and the solution algorithm To solve the learning problem we use a globalized quasi-Newton method. Since modifying the approximate Hessian to be positive definite would result in quite poor performance, we use a different strategy: We perform a regular BFGS update, unless we detect that a descent condition in the BFGS update direction is violated. In that case, instead, we perform a gradient descent update and reset the approximate Hessian (compare [33, Algorithm 11.5 on p.60]). In both cases, we perform an Armijo backtracking line search along the search directions. For a warm start, we always begin the iteration with 5 initial gradient descent steps. We terminated the algorithm, if the norm of the gradient fell below a certain threshold. In addition, for finer discretizations, we also terminated the algorithm if the Armijo backtracking line search was unsuccessful (which also indicates that we are close to a solution).

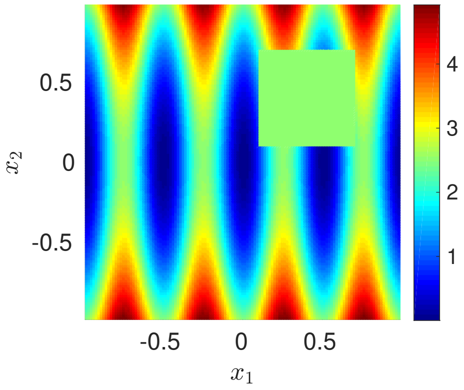

(a) Ground truth control $u_\dagger$.

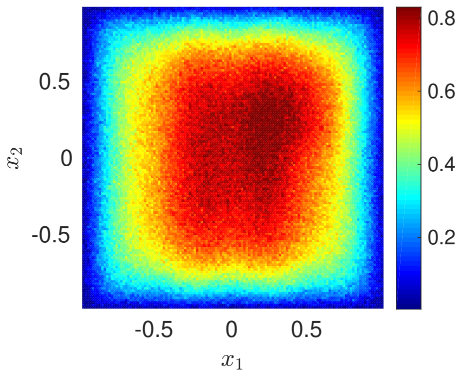

(b) Ground truth state with 3% noise added.

Figure 1 Ground truth data used for the linear state equation.

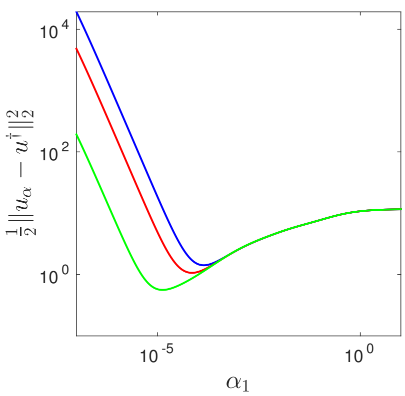

(a) Using only $K_1 = \mathrm{I}$ with 10% noise (blue), 5% noise (red), and 1% noise (green).

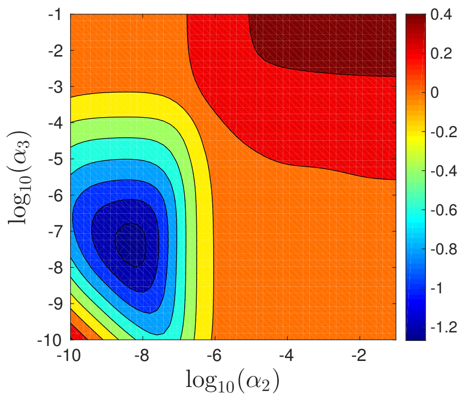

(b) Contour plot (in a logarithmic scale), using $K_2 = \partial_{x_2}$ and $K_3 = \partial_{x_3}$, with 1% noise.

Figure 2 Values of the bilevel cost functional in dependence of the regularization parameters using one or two penalty functionals. In both graphs we use a single noisy data measurement.

Algorithm 1: Iterative method for parameter learning with a linear state equation.

Data: Let α^0 be given.

Define $H^0 := \mathrm{I}$. Compute u^0 solving the lower level problem for α^0 and the corresponding Lagrange multiplier q_0 from the optimality system in Lemma 5.1. For $1 \leq i \leq 3$, set $g_i^0 := \langle K_i q^0, K_i u^0 \rangle$, whenever the operator K_i should be used, and set $d^0 := -g^0$, $k := 0$.

while $\|g^k\|_2^2 > \mathit{tolerance}$ **do**

  if $\langle g^k, d^k \rangle < -\min\{c_1, c_2\|d^k\|_2^2\}\|d^k\|_2^2$ *and* $k \geq 5$ **then**

    Perform Armijo backtracking line search along d^k, set $k = k + 1$ and update α^k;

  else

    $H^k = \mathrm{I}$;

    Perform Armijo backtracking line search along $-g^k$, set $k = k + 1$ and update α^k;

  Compute u^k solving the lower level problem for α^k and the corresponding Lagrange multiplier q^k. Set $g_i^k := \langle K_i q^k, K_i u^k \rangle$, whenever the operator K_i should be used. Update the approximate Hessian H^k and compute the BFGS update direction d^k.

Results for a single vector of data We tested the algorithm in MATLAB for various choices of operators K_i, noise levels, and numbers of available noisy measurements. To be able to compare results for the different settings, we used a fixed seed for random number generation for each dataset. We noticed the following behavior:

- In all tested cases $K_1 = \mathrm{I}$ is the best operator to use, if only one operator should be used. Using only $K_2 = \partial_{x_1}$ or $K_3 = \partial_{x_2}$ results in quite poor performances (see Table 1), which is not surprising since in these cases the lower level problem might be ill-posed.

- Adding another regularization operator K_i to any choice of one or two regularization operators improves tracking of the exact control (see Table 1).

- Using $K_2 = \partial_{x_1}$ and $K_3 = \partial_{x_2}$ is the best choice amongst the two operator cases. The performance using these two operators is only slightly inferior to the performance using all three operators (Figure 3 and Table 1).

- When using multiple noisy data measurements $y_{\delta j}$ with the same statistical structure, the ability to track $u_\dagger$ is significantly improved, as we would expect (compare Table 1 and Figure 3).

- When we only use unilateral regularization associated to K_2 or K_3, the optimal u^* seems to have jumps in the direction which is not penalized (see Figure 3).

Results for training and validation sets Here, we always use all regularization operators K_1, K_2, and K_3. We create $N_T = 100$ training and $N_E = 500$ validation data vectors. The training set is devided into $N_B \in \{1, 2, 5, 10, 100\}$ training batches. Each training batch then consists of $N = N_T/N_B$ training data vectors. For $1 \leq i \leq N_B$ an optimal regularization parameter $\alpha^{*,i}$ is computed for the i-th batch by solving the associated

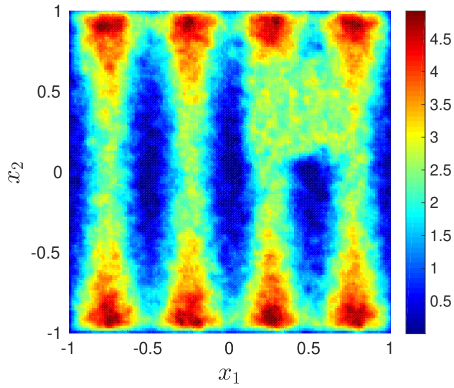

(a) Using $K_1 = \mathrm{I}$ $m = 5$, with optimal $\alpha = (5.73 \times 10^{-6})$ and $\|u^* - u_\dagger\|_2^2 = 0.8336$.

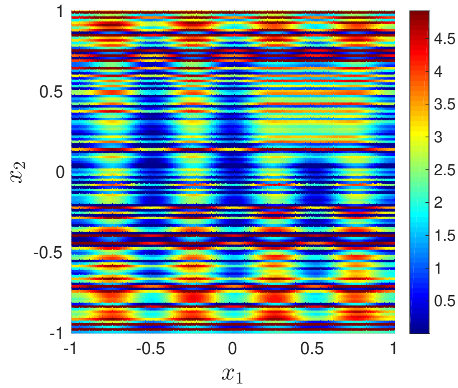

(b) Using $K_2 = \partial_{x_1}$ $m = 5$, with optimal $\alpha = (3.98 \times 10^{-8})$ and $\|u^* - u_\dagger\|_2^2 = 8.7794$.

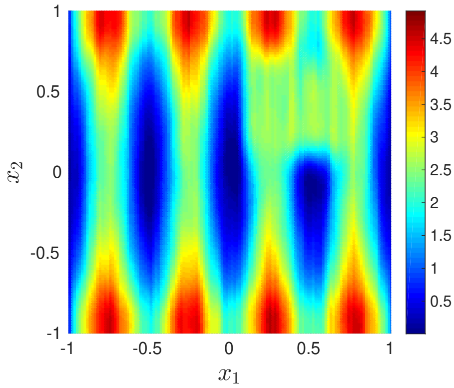

(c) Using $K_1 = \mathrm{I}$ $K_3 = \partial_{x_2}$, $m = 5$, with optimal $\alpha = (2.25 \times 10^{-6}, 5.63 \times 10^{-8})$ and $\|u^* - u_\dagger\|_2^2 = 0.13642$.

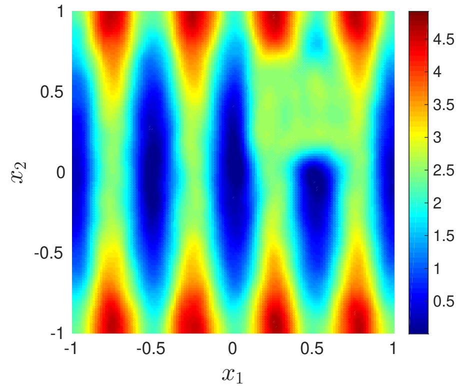

(d) Using $K_2 = \partial_{x_1}$ $K_3 = \partial_{x_2}$, $m = 5$, with optimal $\alpha = (2.11 \times 10^{-9}, 1.47 \times 10^{-8})$ and $\|u^* - u_\dagger\|_2^2 = 0.068485$.

Figure 3 Optimal u^* for the linear state equation, various choices of K_i, and $m = 5$; 1% additive noise was used.

Used Operators	(Locally) Optimal α^*	$\|u^* - u_\dagger\|_2^2$
K_1	(1.33×10^{-5})	1.1235
K_2	(9.08×10^{-8})	38.5143
K_3	(4.02×10^{-7})	37.2593
K_1, K_2	$(2.66 \times 10^{-6}, 2.10 \times 10^{-8})$	0.73592
K_1, K_3	$(5.50 \times 10^{-6}, 1.42 \times 10^{-7})$	0.20574
K_2, K_3	$(5.17 \times 10^{-9}, 4.24 \times 10^{-8})$	0.10707
K_1, K_2, K_3	$(-5.60 \times 10^{-7}, 5.30 \times 10^{-9}, 4.19 \times 10^{-8})$	0.10657

Table 1 Locally optimal α^* for different sets of operators K_i with 1% noise and $m = 1$.

number of batches	training error	validation error
1	2.08×10^{-1}	2.02×10^{-1}
2	2.08×10^{-1}	2.02×10^{-1}
5	2.05×10^{-1}	2.05×10^{-1}
10	2.00×10^{-1}	2.10×10^{-1}
100	1.64×10^{-1}	3.50×10^{-1}

Table 2 Average training and validation error using $N_T = 100$ training and $N_E = 500$ validation data points. The training set is divided into $N_B \in \{1, 2, 5, 10, 100\}$ batches. Here we have used a 64×64 mesh.

(multiple data) learning problem (1.5). Subsequently the optimal regularization parameters are tested on the validation set. Thus for each validation vector $({y_\dagger}^V, {u_\dagger}^V, {y_\delta}^V)$ and each parameter $\alpha^{*,i}$, we compute a solution $u^V_{\alpha^{*,i}}$ to the corresponding lower level problem. We then compute the validation error given by $\|u^V_{\alpha^{*,i}} - {u_\dagger}^V\|^2_{L^2(\Omega)}$. The average validation error is obtained by averaging the validation error over all validation vectors and parameters $\alpha^{*,i}$. In Table 2 the average validation error is compared to the average training error. It is somewhat surprising that the average error over the training set does not significantly differ from the average validation error, except when regularization parameters are determined separately for each data vector ($N = 1$), which is the case for $N_B = 100$. However, this suggests that using learned regularization parameters in structurally related problems can lead to good results.

7.2 Bilinear state equation

In the second numerical experiment the inverse problem is to estimate the diffusion coefficient in a second order elliptic partial differential equation.

Problem setting We consider the learning problem with $\Omega = (-1, 1) \times (-1, 1)$, $Y = H^1_0(\Omega)$, $\tilde{Y} = \tilde{U} = L^2(\Omega)$, and let the state equation $e\colon H^1_0(\Omega) \times U \to H^{-1}(\Omega)$ be given by

$$e(y, u) = \nabla \cdot (\phi_\varepsilon(u)\, \nabla\, y) - f \quad \text{for } (y, u) \in H^1_0(\Omega) \times U.$$

In principle, following the considerations in Section 6.2, we should choose $U = H^2(\Omega)$. However, to avoid the resulting extra numerical difficulties, we work with $U = H^1(\Omega)$. This choice still guarantees the existence of optimal regularization parameters, but the derivation of the optimality system is formal. The function $\phi_\varepsilon\colon \mathbb{R} \to \mathbb{R}$ is a smoothed point wise projection on the interval $[a, b]$. Its precise definition is given in the Appendix. In the state equation, we choose $f \in L^2(\Omega)$ such that for the ground truth control $u_\dagger$ given by

$$u_\dagger(x_1, x_2) = \begin{cases} 1 + {x_2}^2 & \text{if } \sqrt{{x_1}^2 + {x_2}^2} \le \frac{1}{2}, \\ 0.1 + {x_1}^2 & \text{else,} \end{cases}$$

the corresponding state $y_\dagger$ is given by

$$y_\dagger(x_1, x_2) = ({x_1}^4 - {x_1}^2)({x_2}^2 - 1).$$

The ground truth control is shown in Figure 4a. We discretize the problem on a 64×64 mesh using Lagrange P_1 finite elements. Noisy data measurements $y_{\delta j}$ are generated as for the previous example and the same regularization operators are considered.

Used Operators	(Locally) Optimal α^*	$\|u^* - u_\dagger\|_2^2$
K_1	(1.15×10^{-4})	0.73853
K_1, K_2	$(3.52 \times 10^{-5}, 3.33 \times 10^{-6})$	0.23291
K_1, K_3	$(3.94 \times 10^{-5}, 9.22 \times 10^{-7})$	0.48599
K_2, K_3	$(2.49 \times 10^{-7}, 1.52 \times 10^{-7})$	0.16802
K_1, K_2, K_3	$(9.62 \times 10^{-8}, 4.45 \times 10^{-7}, 1.09 \times 10^{-7})$	0.16603

Table 3 Locally optimal α^* for different sets of K_i with 3% noise and $m = 5$.

Used methods and the solution algorithm We use nearly the same globalized quasi-Newton method as for the linear state equation. The only difference is that here a solution to the lower level problem is computed using a sequential programming method (SP method for short) from [21].

Results We performed similar experiments as for the linear state equation and noticed the same behaviour with respect to using different numbers and types of regularization operators. Additionally, we made the following observations:

- $K_1 = \mathrm{I}$ was the only choice of a single operator which lead to meaningful results (see Figure 4b).
- Using the operator $K_2 = \partial_{x_1}$ seems to be more significant for the quality of the reconstructions than using $K_3 = \partial_{x_2}$ (see Table 3). This is also indicated by observing that the obtained optimal regularization parameter for K_2 is usually larger than the optimal regularization parameter for K_3 when using both operators.
- We expect difficulties reconstructing $u_\dagger$ at stationary points of $y_\dagger$ (see [14, p.24]). A simple computation shows that (x_1, x_2) is a stationary point of $y_\dagger$ if and only if one of the following statements is true:

  a) $x_1 = 0$ (line segment along the x_2-axis)
  b) $|x_1| = 1$ and $|x_2| = 1$ (edges of the domain)
  c) $|x_1| = \sqrt{1/2}$ and $x_2 = 0$

  Here, we have continuously extended the gradient of $y_\dagger$ to the boundary of the domain. Difficulties reconstructing $u_\dagger$ near the edges of the domain can be seen in Figure 4b. Since in this case there is no additional smoothing in any of the directions, the values of the reconstructed u^* near the edges tend to zero. Difficulties reconstructing $u_\dagger$ near the x_2-axis can be seen in Figure 4b and 4c. Note that smoothing in the x_1-direction, however, largely prevents the issues near the x_2-axis, as we can see in Figure 4d.

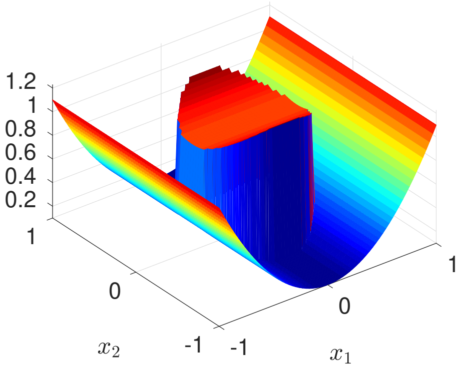

(a) Ground truth control $u_\dagger$.

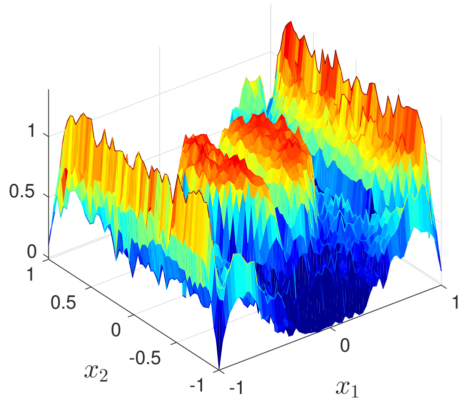

(b) Using $K_1 = \mathrm{I}$ $m = 1$, with optimal $\alpha = (1.13 \times 10^{-4})$ and $\|u^* - u_\dagger\|_2^2 = 0.6862$.

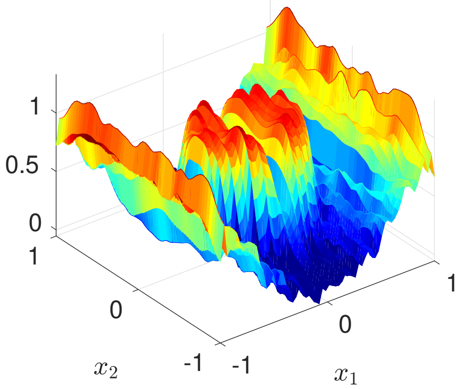

(c) Using $K_1 = \mathrm{I}$ $K_3 = \partial_{x_2}$, $m = 1$, with optimal $\alpha = (2.05 \times 10^{-5}, 7.14 \times 10^{-7})$ and $\|u^* - u_\dagger\|_2^2 = 0.40535$.

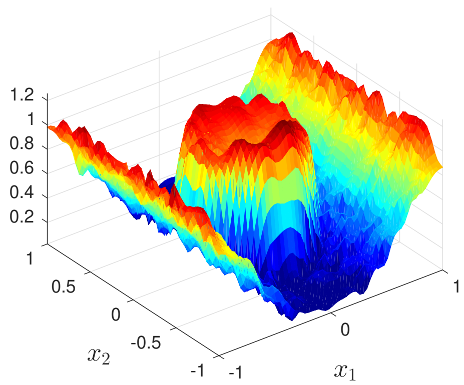

(d) Using $K_2 = \partial_{x_1}$ $K_3 = \partial_{x_2}$, $m = 1$, with optimal $\alpha = (1.96 \times 10^{-7}, 7.26 \times 10^{-8})$ and $\|u^* - u_\dagger\|_2^2 = 0.13812$.

Figure 4 Ground truth and reconstructed controls for the bilinear state equation, various choices of operators K_i and $m = 1$; 1% additive noise was used.

8 Outlook

Possible directions for further research include the use of nonlocal operators as penalty functionals as well as non smooth penalty functionals for nonlinear PDE constraints. It can also be of interest to consider regularization terms of the form $\beta \sum_{i=1}^{r} \alpha_i \Psi_i(u)$, where the α_i are determined by learning techniques based on a training set, and, given new data, β is chosen by classical parameter choice strategies to adapt the amount of regularization to possibly changing noise levels and different scaling. We also point out that the ability to compute optimal regularization parameters provides the opportunity to evaluate how well classical parameter choice strategies are performing.

A Proof of Lemma 3.1

Proof. The proof is divided into three steps.

Step 1: We first aim at showing that the sequence (y^n, u^n) has a weakly convergent subsequence in F_{ad}. Since F_{ad} is a weakly closed subset of a reflexive Banach space, for this purpose it is sufficient to show that (y^n, u^n) is bounded. Utilizing (H4), in turn, the boundedness of (y^n, u^n) follows if we can prove that (u^n) is bounded.

To show that (u^n) is bounded, we argue as follows: Since the sequence $({y_\delta}^n)$ is convergent, there exists $M > 0$ such that

$$\|{y_{\delta}}_j^n\|_{\tilde{Y}} \leq M \quad \text{for all } n \in \mathbb{N} \text{ and } 1 \leq j \leq m.$$

A simple computation now shows that for every $(y, u) \in F_{ad}$ and every $n \in \mathbb{N}$ we have

$$\underset{\sim}{\alpha} \cdot \Psi(u^n) \leq \mathcal{J}_{\alpha^n, {y_\delta}^n}(y^n, u^n) \leq \mathcal{J}_{\alpha^n, {y_\delta}^n}(y, u) \leq \frac{1}{2}(\|y\|_{\tilde{Y}} + M)^2 + \bar{\alpha} \cdot \Psi(u).$$

Using that Ψ is proper on F_{ad}, we can choose $(y, u) \in F_{ad}$ such that the right-hand side of this chain of inequalities is finite. Since the right-hand side is independent of n, and $\underset{\sim}{\alpha} > 0$, this shows that

$$\sum_{i=1}^{r} \Psi_i(u^n)$$

is bounded. Consequently, from (H5) it follows that (u^n) is bounded; and thus the first step is complete.

Step 2: Using the first step, we can assume that there exists a subsequence of (y^n, u^n), which, for simplicity, we again denote by (y^n, u^n), and $(\bar{y}, \bar{u}) \in F_{ad}$ such that

$$(y^n, u^n) \rightharpoonup (\bar{y}, \bar{u}).$$

Our goal in the second step is to show that $(\bar{y}, \bar{u})$ solves $(\mathcal{P}_{\alpha, y_\delta})$. For this purpose, since (y^n, u^n) solves $(\mathcal{P}_{\alpha^n, {y_\delta}^n})$, note that

$$\mathcal{J}_{\alpha^n, {y_\delta}^n}(y^n, u^n) \leq \mathcal{J}_{\alpha^n, {y_\delta}^n}(y, u) \tag{A.1}$$

for all $(y, u) \in F_{ad}$ and $n \in \mathbb{N}$. Using that

$$(\alpha^n, {y_\delta}^n, y^n, u^n) \mapsto \mathcal{J}_{\alpha^n, {y_\delta}^n}(y^n, u^n)$$

is weakly lower semi continuous on $[\underline{\alpha}, \bar{\alpha}] \times \tilde{Y}^m \times Y \times U$, and that for every $(y, u) \in F_{ad}$ the mapping

$$(\alpha^n, {y_\delta}^n) \mapsto \mathcal{J}_{\alpha^n, {y_\delta}^n}(y, u)$$

is continuous on $[\underline{\alpha}, \bar{\alpha}] \times \tilde{Y}^m$, taking the limit $n \to \infty$ in (A.1) we arrive at

$$\mathcal{J}_{\alpha, y_\delta}(\bar{y}, \bar{u}) \leq \liminf_{n\to\infty} \mathcal{J}_{\alpha^n, {y_\delta}^n}(y^n, u^n) \leq \lim_{n\to\infty} \mathcal{J}_{\alpha^n, {y_\delta}^n}(y, u) = \mathcal{J}_{\alpha, y_\delta}(y, u). \tag{A.2}$$

As a consequence of this estimate, we have

$$\lim_{n\to\infty} \mathcal{J}_{\alpha^n, {y_\delta}^n}(y^n, u^n) = \mathcal{J}_{\alpha, y_\delta}(\bar{y}, \bar{u}) = \min_{(y,u)\in F_{ad}} \mathcal{J}_{\alpha, y_\delta}(y, u) < \infty, \tag{A.3}$$

which shows that $(\bar{y}, \bar{u})$ solves $(\mathcal{P}_{\alpha, y_\delta})$. This finishes the second step.

Step 3: In order to complete the proof it remains to show that

$$\lim_{n\to\infty} \Psi(u^n) = \Psi(\bar{u}),$$

which is done now. First, observe that due to weak lower semi continuity of the involved functions

$$\|\bar{y} - y_{\delta j}\|_{\tilde{Y}}^2 \leq \liminf_{n\to\infty} \|y^n - y_{\delta j}\|_{\tilde{Y}}^2 \quad \text{for } 1 \leq j \leq m \tag{A.4}$$

and

$$\Psi_i(\bar{u}) \leq \liminf_{n\to\infty} \Psi_i(u^n) \quad \text{for } 1 \leq i \leq r. \tag{A.5}$$

We now argue as follows: If for some $1 \leq i \leq r$ it holds that

$$\Psi_i(\bar{u}) < \liminf_{n\to\infty} \Psi_i(u^n),$$

then in view of (A.4)–(A.5), and using $\mathcal{J}_{\alpha, y_\delta}(\bar{y}, \bar{u}) < \infty$, this implies

$$\mathcal{J}_{\alpha, y_\delta}(\bar{y}, \bar{u}) < \lim_{n\to\infty} \mathcal{J}_{\alpha^n, {y_\delta}^n}(y^n, u^n).$$

Since we have already shown that

$$\lim_{n\to\infty} \mathcal{J}_{\alpha^n, {y_\delta}^n}(y^n, u^n) = \mathcal{J}_{\alpha, y_\delta}(\bar{y}, \bar{u}),$$

this leads to a contradiction. Consequently, we must have

$$\Psi_i(\bar{u}) = \liminf_{n\to\infty} \Psi_i(u^n) \quad \text{for all } 1 \leq i \leq r. \tag{A.6}$$

Since (A.6) is also true for every subsequence of (u^n), this implies that

$$\Psi_i(\bar{u}) = \lim_{n\to\infty} \Psi_i(u^n) \quad \text{for all } 1 \leq i \leq r,$$

which is what was left to show.

□

B Smoothed projection

Here, we provide the precise definition of the function $\phi_\varepsilon \colon \mathbb{R} \to \mathbb{R}$ for given $\varepsilon > 0$, which can be seen as a smoothed projection on the interval $[a, b]$. We have determined ϕ_ε by choosing it as the polynomial with (piecewise) lowest degree matching the following conditions: $\phi_\varepsilon \in C^3(\mathbb{R})$, $a \leq \phi_\varepsilon \leq b$, and

- $\phi_\varepsilon(t) = t \quad \text{for } t \in [a+\varepsilon, b-\varepsilon]$,
- $\phi_\varepsilon(t) = a \quad \text{for } t \leq a$,
- $\phi_\varepsilon(t) = b \quad \text{for } b \leq t$.

The resulting function can be written as follows:

$$\phi_\varepsilon(t) = \begin{cases} t & \text{if} \quad a+\varepsilon \leq t \leq b-\varepsilon \\ a & \text{if} \quad t \leq a \\ b & \text{if} \quad b \leq t \\ f_\varepsilon(t-a) & \text{if} \quad a \leq t \leq a+\varepsilon \\ -f_\varepsilon(-t+b)+a+b & \text{if} \quad b-\varepsilon \leq t \leq b, \end{cases}$$

where

$$f_\varepsilon(x) = \frac{-10}{\varepsilon^6}x^7 + \frac{36}{\varepsilon^5}x^6 - \frac{45}{\varepsilon^4}x^5 + \frac{20}{\varepsilon^3}x^4 + a.$$

References

[1] R. A. Adams. *Sobolev Spaces.* Pure and applied mathematics. Academic Press, 1975.

[2] A. A. Ali, K. Deckelnick, and M. Hinze. Global minima for semilinear optimal control problems. *Computational Optimization and Applications*, 65(1):261–288, 2016.

[3] V. Y. Arsenin and A. N. Tikhonov. *Solutions of ill-posed problems*, volume 14. Winston Washington, DC, 1977.

[4] H. T. Banks and K. Kunisch. *Estimation techniques for distributed parameter systems.* Springer, 1989.

[5] H. Brezis. *Functional Analysis, Sobolev Spaces and Partial Differential Equations.* Universitext. Springer New York, 2010.

[6] T. Brox, P. Ochs, T. Pock, and R. Ranftl. Bilevel optimization with nonsmooth lower level problems. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pages 654–665. Springer, 2015.

[7] E. Casas and F. Tröltzsch. Second order optimality conditions and their role in PDE control. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 117(1):3–44, 2015.

[8] J. Chung, M. Chung, and D. P. O'Leary. Designing optimal spectral filters for inverse problems. *SIAM Journal on Scientific Computing*, 33(6):3132–3152, 2011.

[9] J. Chung and M. I. Español. Learning regularization parameters for general-form Tikhonov. *Inverse Problems*, 33(7):074004, 2017.

[10] J. C. De los Reyes, , C.-B. Schönlieb, and T. Valkonen. Bilevel parameter learning for higher-order total variation regularisation models. *Journal of Mathematical Imaging and Vision*, 57(1):1–25, 2017.

[11] J. C. De los Reyes and C.-B. Schönlieb. Image denoising: Learning the noise model via nonsmooth PDE-constrained optimization. *Inverse Problems & Imaging*, 7(4), 2013.

[12] J. C. De los Reyes, C.-B. Schönlieb, and T. Valkonen. The structure of optimal parameters for image restoration problems. *Journal of Mathematical Analysis and Applications*, 434(1):464 – 500, 2016.

[13] S. Dempe. *Foundations of bilevel programming*. Springer Science & Business Media, 2002.

[14] H. W. Engl, M. Hanke, and A. Neubauer. *Regularization of Inverse Problems*. Mathematics and Its Applications. Springer Netherlands, 1996.

[15] G. H. Golub, M. Heath, and G. Wahba. Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics*, 21(2):215–223, 1979.

[16] E. Haber, L. Horesh, and L. Tenorio. Numerical methods for experimental design of large-scale linear ill-posed inverse problems. *Inverse Problems*, 24(5):055012, 2008.

[17] E. Haber and L. Tenorio. Learning regularization functionals—a supervised training approach. *Inverse Problems*, 19(3):611, 2003.

[18] P. C. Hansen. Analysis of discrete ill-posed problems by means of the L-curve. *SIAM review*, 34(4):561–580, 1992.

[19] P. C. Hansen and D. P. O'Leary. The use of the L-curve in the regularization of discrete ill-posed problems. *SIAM Journal on Scientific Computing*, 14(6):1487–1503, 1993.

[20] G. Holler. A bilevel approach for parameter learning in inverse problems. Master's thesis, Karl-Franzens Universität Graz, Austria, 2017.

[21] K. Ito and K. Kunisch. A sequential method for a class of stable mathematical programming problems. *SIAM Journal on Optimization 26(2) 2016, 1262-1292*, 2016.

[22] J. Jahn. *Introduction to the theory of nonlinear optimization*. Springer Science & Business Media, 2007.

[23] B. Kaltenbacher, A. Neubauer, and O. Scherzer. *Iterative regularization methods for nonlinear ill-posed problems*, volume 6. Walter de Gruyter, 2008.

[24] K. Kunisch and T. Pock. A bilevel optimization approach for parameter learning in variational models. *SIAM Journal on Imaging Sciences*, 6(2):938–983, 2013.

[25] S. Kurcyusz and J. Zowe. Regularity and stability for the mathematical programming problem in Banach spaces. *Applied Mathematics and Optimization*, 5(1):49–62, 1979.

[26] O. A. Ladyzhenskaya and N. N. Uraltseva. *Linear and quasilinear elliptic equations*. Academic Press, New York, 1968.

[27] A. K. Louis. *Inverse und schlecht gestellte Probleme*. Springer-Verlag, 2013.

[28] H. Maurer and J. Zowe. First and second-order necessary and sufficient optimality conditions for infinite-dimensional programming problems. *Mathematical Programming*, 16(1):98–110, 1979.

[29] S. M. Robinson. Strongly regular generalized equations. *Mathematics of Operations Research*, 5(1):43–62, 1980.

[30] L. Ruthotto, J. Chung, and M. Chung. Optimal experimental design for constrained inverse problems. *arXiv preprint arXiv:1708.04740*, 2017.

[31] O. Scherzer, M. Grasmair, H. Grossauer, M. Haltmeier, and F. Lenzen. *Variational methods in imaging*. Springer, 2009.

[32] T. Schuster, B. Kaltenbacher, B. Hofmann, and K. S. Kazimierski. *Regularization methods in Banach spaces*, volume 10. Walter de Gruyter, 2012.

[33] M. Ulbrich and S. Ulbrich. *Nichtlineare Optimierung*. Springer-Verlag, 2012.