



# Variational regularization in inverse problems and machine learning

Martin Burger

**Abstract.** This paper discusses basic results and recent developments on variational regularization methods, as developed for inverse problems. In a typical setup we review basic properties needed to obtain a convergent regularization scheme and further discuss the derivation of quantitative estimates respectively the needed ingredients such as Bregman distances for convex functionals.

In addition to the approach developed for inverse problems, we will also discuss variational regularization in machine learning and work out some connections to the classical regularization theory. In particular we will discuss a reinterpretation of machine learning problems in the framework of regularization theory and a reinterpretation of variational methods for inverse problems in the framework of risk minimization. Moreover, we establish some previously unknown connections between error estimates in Bregman distances and generalization errors.

## 1. Introduction

Regularization methods are an approach of fundamental importance in the solution of ill-posed problems. Their main paradigm is to approximate an ill-posed problem by a parametrized family of well-posed problems, with appropriate convergence properties as the regularization parameter and the so-called noise level tend to zero. The noise level is a measure for the size of deterministic and stochastic errors in the data, which are usually the main cause of concern due to the ill-posedness.

A detailed theory of regularization has been developed in the typical setting of inverse problems, obviously with more precise results in the case of linear forward models than for nonlinear ones (cf. [3, 20, 24, 54, 56] and references therein). Regularization is however not only relevant in inverse problems, similar methods are now routinely used in machine learning, mainly from a practical point of view, with theoretical results often hidden in the statistical theory of generalization (cf. e.g. [27, 33, 40]). The role and objective of regularization is less clear and less developed

---

*2020 Mathematics Subject Classification.* Primary 65J20; Secondary 47A52.

*Keywords.* Regularization theory, variational methods, inverse problems, machine learning.

in the machine learning domain. In this paper we will thus aim to give a unified overview and present some links between the formulations and questions in inverse problems and those in machine learning. We will concentrate on the prominent class of variational regularization methods, which we interpret in a rather broad way.

## 2. Regularization theory

In order to present the basic ideas of regularization methods in a rather unified way for inverse and machine learning problems, we will first adopt a high-level point of view. Regularization theory is based on the following ingredients.

- An *ideal problem* respectively an *ideal solution*  $u^*$ . We can assume that the ideal problem is given by a map  $\Phi : \mathcal{V}_D \rightarrow \mathcal{U}$ , where  $\mathcal{V}_D$  is the space of ideal data and  $\mathcal{U}$  is the space of admissible solutions. The typical analysis is confined to Banach or at least metric spaces.
- A space  $\mathcal{V} \supset \mathcal{V}_D$  of *possible data* and a measure of noise between the ideal data  $v^* = \Phi^{-1}(u^*) \in \mathcal{V}_D$  and noisy data  $v \in \mathcal{V}$ . In the case of an ill-posed problem, the operator  $\Phi$  is not continuous when considered from (a subset of)  $\mathcal{V}$  to  $\mathcal{U}$ ; it may be continuous on bounded subsets of  $\mathcal{V}_D$  however. The latter leads to the concept of conditional stability (cf. [57, 58]) and corresponding stability estimates.
- A family of continuous, possibly multivalued, maps  $\Phi_\alpha : \mathcal{V} \rightarrow \mathcal{P}(\mathcal{U})$ ,  $\alpha \in \mathcal{A}$ , such that for a sequence  $(v_n) \subset \mathcal{V}$  converging to  $v^* \in \mathcal{V}_D$ , there exists a parameter sequence  $\alpha_n$  such that there is  $u_n \in \Phi_{\alpha_n}(v_n)$  converging to  $u^*$  (in a suitable metrizable topology, possibly weak or weak-star on bounded sets in the Banach space case). Sometimes the notion of convergence is restricted to subsequences.

To make these notions more concise we will discuss them in the setting of inverse problems as well as machine learning subsequently.

### 2.1. Inverse problems

In the typical case of inverse problems, there is first a (continuous) forward operator  $F : \mathcal{U} \rightarrow \mathcal{V}$ , which is typically not invertible and if it is on a subset of  $\mathcal{V}$ , the inverse is discontinuous. The set of ideal data is a subset of  $F(\mathcal{U})$ , and there the multivalued operator

$$\Phi_0 : \mathcal{V}_D \rightarrow \mathcal{P}(\mathcal{U}), \quad v \mapsto F^{-1}(v)$$

can be defined. In order to obtain a unique (generalized) inverse, a further selection operator  $\Sigma : \mathcal{P}(\mathcal{U}) \rightarrow \mathcal{U}$  is defined to obtain  $\Phi := \Sigma \circ \Phi_0$ . Let us mention that there are standard examples of the selection operator such as the minimum norm solution, but often this issue is treated in a hidden or unprecise way. We refer to [3] for a detailed discussion of selection operators in inverse problems.

The standard notion of noise is the perturbation of the data, i.e.  $v - v^*$ , either as a deterministic or a stochastic quantity. The norm of  $v - v^*$  in the Banach space  $\mathcal{V}$  (or the expectation of some power of the norm) serves as a definition of the noise level.

The solution of the inverse problem can then be cast as the solution of the ill-posed operator equation

$$F(u) = v$$

or as the minimization of

$$D(u) = L(F(u), v), \quad (2.1)$$

where  $L$  is some distance measure between the predicted data  $F(u)$  and the measured data  $v$ . If statistical information about the noise is available or the forward model contains other stochastic elements,  $L$  is typically a negative log-likelihood functional.

As mentioned above, regularization methods are families of multivalued operators  $\Phi_\alpha : \mathcal{V} \rightarrow \mathcal{P}(\mathcal{U})$ ; in most cases the parameter domain  $\mathcal{A}$  is a subset of the positive real numbers. The well-posedness of  $\Phi_\alpha$  is characterized by some set-valued continuity, e.g. if  $u_n \rightarrow u$ , then  $\Phi_\alpha(u_n)$  contains a convergent subsequence and each limit  $v$  of a convergent subsequence satisfies  $v \in \Phi_\alpha(u)$ . In most cases the regularization operator satisfies a stronger stability estimate of the form

$$d_U(u_1, u_2) \leq C_\alpha d_V(v_1, v_2) \quad \forall u_1 \in \Phi_\alpha(v_1), u_2 \in \Phi_\alpha(v_2), \quad (2.2)$$

where  $d_U$  and  $d_V$  are appropriate distance measures (that may be degenerate in the sense that  $d_U(u_1, u_2)$  can vanish also if  $u_1 \neq u_2$ ).

Regularization methods are constructed along several different paradigms.

- Data smoothing or mollifier methods, which are of the form  $\Phi_\alpha = F^{-1} \circ M_\alpha$ , where  $M_\alpha : \mathcal{V} \rightarrow \mathcal{V}_D$  is a family of mollifying (smoothing) operators into an appropriate subspace of  $\mathcal{V}$  on which there exists a continuous inverse of  $F$ . In order to obtain suitable regularization methods, a quite detailed characterization of the forward operator is needed in order to be sure to construct a mollification to the right subspace. Consequently, such methods became popular for inverse problems with well-understood forward operators such as tomography (cf. [43, 44]).
- Direct approximation of the operator  $F$  by continuously invertible operators (cf. [24, 37, 41, 56] and references therein). The construction of approximations is usually done only in the case of linear forward operators based on modifying (small) singular vectors or by approximating the normal equation, i.e.  $F^*F$ . The latter is however related to the minimization of the least-squares function  $\|F(u) - v\|^2$  and can thus be viewed as a variational method. Another approach modifying the forward operator is discretization, the regularization parameter thus being related to the discretization fineness.

- Variational methods are based on a perturbation of the likelihood minimization, with  $\Phi_\alpha$  mapping  $v$  to the set of minimizers of

$$D_\alpha(u) = L(F(u), v) + \alpha J(u)$$

for some regularization functional  $J$  that introduces the needed compactness properties for the existence of minimizers and  $\alpha \in \mathbb{R}_+$  being the regularization parameter (cf. [3, 54]).

- Iterative regularization methods use a well-defined iteration method such as a fixed-point iteration or some descent scheme for the likelihood minimization to define an approximation of the inverse of  $F$ , with the iteration number  $\alpha \in \mathbb{N}$  being the regularization parameter (cf. [14, 24, 34, 35, 49]). Since the majority of iterative methods, in particular in the nonlinear case, are iterative methods for variational problems, there is an intimate connection to variational regularization methods.
- Learned regularization methods are of increasing relevance recently (cf. [1, 3] and references therein), which are categorized into supervised and semi-supervised approaches. The supervised approach tries to learn the regularization operator  $\Phi_\alpha$  directly from a collection of pairs of training data  $(u_i, v_i)$ , e.g. by approximation with a deep neural network. Consistent data pairs are however difficult to obtain in many inverse problems, in particular with realistic input data  $u_i$  and realistic noise in  $v_i$ . The alternative semi-supervised approach mainly works on suitable solutions  $u_i$ , e.g. images for reconstruction tasks, and tries to learn a more conventional regularization approach, e.g. the regularization functional  $J$  in variational regularization methods. With certain restrictions such as convex networks those become accessible for theoretical arguments of regularization theory.

Besides providing a well-posed problem for fixed  $\alpha$ , which often requires some advanced analysis itself (e.g. existence of minimizers for variational problems), a major goal of regularization theory is to study the convergence of regularized solutions. While a qualitative convergence theory can be developed under generic conditions, it is well known that a quantitative theory will rely on additional assumptions on the ideal solution  $u^*$  due to the underlying ill-posedness. To understand the possibility to derive such estimates and the used assumptions from a generic point of view, let us consider a sequence of data  $v_n \rightarrow v^*$  and a parameter choice  $\alpha_n$ , assuming that  $\alpha_n$  is a nonnegative scalar sequence converging to zero (e.g. the regularization parameters in a variational regularization method or  $\alpha_n = \frac{1}{k_n}$  with  $k_n$  the maximal iteration number in an iterative regularization method). Now assume that the stability estimate (2.2) holds and that  $u^*$  satisfies a *range condition* for the regularization operator (cf. [3]).

**Definition 2.1.** An element  $u^* \in \mathcal{U}$  is said to satisfy a *range condition for the regularization operator*  $\Phi_\alpha$  if for all  $\alpha$  there exists  $v_\alpha^*$  such that  $u^* \in \Phi_\alpha(v_\alpha^*)$ .

Under a range condition we can write

$$u_n - u^* \in \Phi_{\alpha_n}(v_n) - \Phi_{\alpha_n}(v_{\alpha_n}^*)$$

and exploit the stability estimate (2.2) to obtain

$$d_U(u^*, u_n) \leq C_{\alpha_n} d_V(v_{\alpha_n}^*, v^*).$$

Thus, if we can control the range condition in the sense that we can construct an element  $v_{\alpha_n}^*$  out of  $v^*$  such that the distance can be estimated, we directly obtain an error estimate. This will be made more precise in the next section on variational regularization methods.

## 2.2. Learning and risk minimization

In the typical case of machine learning problems (cf. [33,45]) we are given (randomly sampled) input samples  $x_i \in \mathcal{X}$  and output samples  $y_i \in \mathcal{Y}$ ,  $i = 1, \dots, N$ , and want to infer a parametrized map  $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$  reasonably reproducing these training data and generalizing further to other data of the same kind. These properties are frequently obtained from risk minimization arguments. Given a loss  $\ell$  measuring deviations in the output space, the empirical risk is given by

$$\hat{R}(\theta) = \frac{1}{N} \sum_{i=1}^N \ell(f_\theta(x_i), y_i)$$

and approximate solutions are constructed as approximate minimizers of  $\hat{R}$ , e.g. via variational regularization methods minimizing

$$D_\alpha(\theta) = \hat{R}(\theta) + \alpha J(\theta)$$

or by iterative methods such as the gradient descent

$$\theta^{k+1} = \theta^k - \tau^k \hat{R}'(\theta^k)$$

or even more often by stochastic gradient descent, where the term implicit regularization is common (cf. [48]).

Generalization is usually measured by the behavior on the population risk, i.e.

$$R(\theta) = \mathbb{E}_{(x,y) \sim \mathbb{P}}(\ell(f_\theta(x), y));$$

in particular the generalization error defined by

$$G(\theta) = R(\theta) - \widehat{R}(\theta),$$

evaluated at a regularized solution. Note that the generalization error is actually a random variable depending on the samples  $(x_i, y_i)_{i=1}^N$ , hence it is relevant to consider its distribution among the random sampling.

The ideal model could be defined in two ways, depending on what variable is identified to be the relevant one. In any case the ideal solution is perceived as a minimizer of the population risk, however one could define  $u^*$  as the optimal parameter value or the optimal function. Thus we are led to the following cases.

- (i) The first case, corresponding to classical approaches in statistics such as regression, is to define  $\mathcal{U}$  as the set of possible parameters, genuinely a finite-dimensional space (with few generalizations to infinite-dimensional models recently, cf. [39, 47]). Thus, the ideal solution is given by

$$\theta^* \in \arg \min_{\theta \in \mathcal{U}} R(\theta).$$

- (ii) The second case rather corresponds to the perspective of modern learning theory; it extends the population risk to some function class  $\mathcal{F}$  and computes for  $f \in \mathcal{F}$

$$S(f) = \mathbb{E}_{(x,y) \sim \mathbb{P}}(\ell(f(x), y)).$$

The ideal solution is given by

$$f^* \in \arg \min_{f \in \mathcal{F}} S(f).$$

Another obvious question in this case is how to define the ideal and perturbed data. We follow a distributional viewpoint and define the ideal data  $v^*$  as the data distribution  $\mathbb{P}$ . Correspondingly, the perturbed data are given by the empirical distribution

$$\mathbb{P}^N = \frac{1}{N} \sum_{i=1}^N \delta_{(x_i, y_i)},$$

where  $\delta_z$  denotes the concentrated measure at  $z$ . Thus, the noise level becomes a distance between (probability) distributions, standard distances such as the total variation distance or Wasserstein metrics.

The regularization operator  $\Phi_\alpha$  maps from a space of probability distributions to (a set of) regularized solutions. Take the variational regularization of minimizing  $D_\alpha$  as an example. Then in case (i),  $\Phi_\alpha$  is given by

$$\Phi_\alpha : \mathbb{P}^N \mapsto \arg \min_{\theta} D_\alpha(\theta),$$

while in the second case (ii) we have

$$\Phi_\alpha : \mathbb{P}^N \mapsto \{f_\theta \mid \theta \in \arg \min_\theta D_\alpha(\theta)\}.$$

We finally mention that these models can obviously be generalized, in particular to the case of further data errors in the samples  $(x_i, y_i)$ . Then the samples can be considered to be drawn from a distribution  $\mathbb{P}'$  and the effective error is not just determined by sampling but also by the distance of  $\mathbb{P}$  and  $\mathbb{P}'$ .

Thus, we see that regularized learning problems can be reformulated in the language of regularization theory for inverse problems (see also [12, 53]). In turn we will see that many inverse problems can be reformulated as risk minimization problems, in particular if there is additional sampling of measurement points.

### 2.3. Risk minimization formulation of inverse problems

Many inverse problems are dealing with data being functions of a variable  $x$ , e.g. in integral equations of the first kind or tomography, where  $x$  is a set of distances and angles (cf. [46]). Denoting the unknown of the inverse problem by  $\theta$ , we thus obtain  $F(\theta)$  as a function of  $x$  and denote  $f(x; \theta) = F(\theta)(x)$ . Moreover, standard log-likelihood functionals in this setting are of the form

$$L(F(\theta), v) = \int_\Omega \ell(F(\theta)(x), v(x)) dx$$

for some function  $\ell$ . Thus, choosing  $\mathcal{P} = \mathcal{L}_\Omega \delta_{v(x)}$ , where  $\mathcal{L}_\Omega$  denotes the Lebesgue measure on  $\Omega$ , we obtain

$$L(F(\theta), v) = \mathbb{E}_{(x,y) \sim \mathcal{P}}(\ell(F(\theta)(x), y)) = \mathbb{E}_{(x,y) \sim \mathbb{P}}(\ell(f(x; \theta), y)).$$

The ideal problem is thus the minimization of the loss for appropriate data  $v^*$ .

In a practical setting we have a finite sampling of data with additional noise, which we consider to be additive for simplicity in the following. This means the practical data are a finite number  $N$  of samples  $y_i = F(\theta)(x_i) + n_i$ , where  $n_i$  are the noise samples drawn from some distribution. The practical distribution of samples and data is of the form

$$\mathbb{P}^N = \frac{1}{N} \sum_{i=1}^N \delta_{x_i} \otimes \delta_{F(\theta^*)(x_i) + n_i},$$

where the  $x_i$  are drawn from a prior distribution (usually a deterministic one) and the  $n_i$  are drawn from the noise distribution.

**Example 2.2.** As a simple example consider the inversion of the Radon transform on a domain  $\Omega \subset \mathbb{R}^2$ . Then in the standard parametrization we can choose  $x \in [0, \pi) \times$

$[0, L]$  as the angle and distance to origin of the lines to be integrated on. Correspondingly  $F(\theta)(x)$  is the line integral of the density function  $\theta$  on the line parametrized by  $x$ . Now let  $x$  be drawn from the uniform distribution on  $[0, \pi) \times [0, L]$  and let each  $n$  be drawn from a Gaussian distribution  $G_\sigma$  with zero mean and finite variance. Then the population risk becomes

$$\begin{aligned} R(\theta) &= \frac{1}{2L\pi} \int_{[0,\pi) \times [0,L]} \int_{\mathbb{R}} |F(\theta)(x) - F(\theta^*)(x) - n|^2 dG_\sigma(n) dx \\ &= \frac{1}{2L\pi} \int_{[0,\pi) \times [0,L]} |F(\theta)(x) - F(\theta^*)(x)|^2 dx + \int_{\mathbb{R}} n^2 dG_\sigma(n). \end{aligned}$$

Hence, after affine transform with terms independent of  $\theta$ , the population risk equals the squared  $L^2$ -distance of the Radon transforms of  $\theta$  and  $\theta^*$ , which is the usual data discrepancy  $L$ . The empirical risk on the other hand is of the form

$$\hat{R}(\theta) = \frac{1}{2N} \sum_{i=1}^N |F(\theta)(x_i) - y_i|^2,$$

which is the standard functional minimized in practice.

For a more general noise model one may construct the conditional distribution for  $y$  based on using the appropriate push-forward of the noise distribution based on applying the noise to  $F(\theta^*)(x)$  and an appropriately chosen loss function. Moreover, errors in the forward model could be included in the stochastic model, which will imply that even in the ideal model the conditional distribution of  $y$  given  $x$  is not concentrated.

### 3. Variational regularization

In the following we present some key steps in the analysis of iterative regularization methods, for the sake of a simpler presentation restricting ourselves to a linear forward model and a quadratic data fidelity in a Hilbert space, i.e.

$$D_\alpha(u) = \frac{1}{2} \|Fu - v\|^2 + \alpha J(u), \tag{3.1}$$

where  $J : \mathcal{U} \rightarrow \mathbb{R} \cup \{+\infty\}$  is assumed to be convex and proper. Moreover, we assume  $\mathcal{V}$  to be a Hilbert space and  $\mathcal{U}$  a Banach space being the dual of some Banach space  $\mathcal{W}$ , with the additional property that the weak-star topology on  $\mathcal{U}$  is metrizable on bounded sets. The operator  $F : \mathcal{U} \rightarrow \mathcal{V}$  is assumed to be bounded and the adjoint of a bounded linear operator  $E : \mathcal{V} \rightarrow \mathcal{W}$ . With abuse of notation we shall write



$F^* = E$ . Finally, we need some additional property of the regularization functional; we assume that it is the convex conjugate of some other functional  $H : \mathcal{W} \rightarrow \mathbb{R}$ , i.e.

$$J(u) = \sup_{w \in \mathcal{W}} \langle u, w \rangle - H(w).$$

Let us mention that convex conjugates are weak-star lower semicontinuous, which is obviously an important property of the functional and can be inferred by similar arguments as the weak lower semicontinuity results in [23]. Finally, a coercivity property is needed to apply weak-star compactness arguments (based on the Banach–Alaoglu theorem); we assume that the sublevel sets

$$M_C = \{u \in \mathcal{U} \mid J(u) \leq C\}$$

are bounded in  $\mathcal{U}$  for  $C > 0$ . The final property we need is that  $J$  is bounded below; we can assume directly that  $J$  is nonnegative.

There are various important examples in literature motivating the above model and assumptions. A popular and reasonably easy to compute approach is classical Tikhonov–Phillips regularization with  $\mathcal{U}$  being a Hilbert space and

$$J(u) = \frac{1}{2} \|u\|^2.$$

Possibly the most prominent example with a variety of applications is total variation regularization (cf. [16, 19]), i.e.  $\mathcal{U} = BV(\Omega)$  and

$$J(u) = \sup_{g \in C_0^\infty(\Omega)^d, \|g\|_\infty \leq 1} \int_{\Omega} u \nabla \cdot g \, d\mathcal{L}_{\Omega},$$

where  $\Omega \subset \mathbb{R}^d$  is the domain on which the function to be reconstructed is defined. There are various variants of total variation, including higher order versions, which received considerable attention. Another class of important regularization methods are sparsity-enforcing priors (cf. [50]), in the simplest setup  $\mathcal{U} = \ell^1$  and

$$J(u) = \sum |u_i|.$$

An interesting case in deconvolution problems as well as mean-field approaches to learning with neural networks is the continuum variant, the total variation norm of Radon measures (cf. [7, 22]). Here we have  $\mathcal{U} = \mathcal{M}(\Omega)$  and

$$J(u) = \sup_{w \in C_0(\Omega)} \int_{\Omega} w \, du.$$

### 3.1. Basic properties of variational regularization methods

A key result, often found for special cases in literature (cf. e.g. [16, 55]), is the existence of a minimizer and some stability, which verifies the well-posedness of the regularization operator  $\Phi_\alpha(v) := \arg \min_u D_\alpha(u)$ .

**Theorem 3.1.** *Under the above assumptions on  $\mathcal{U}$ ,  $\mathcal{V}$ ,  $F$ , and  $J$  there exists a minimizer of  $D_\alpha(u)$  for every  $v \in \mathcal{V}$  and every  $\alpha > 0$ . Moreover, if  $\alpha > 0$ ,  $v_n \rightarrow v$ , and  $u_n \in \Phi_\alpha(v_n)$ , then there exists a weak-star convergent subsequence  $v_{n_k}$  and the limit  $u$  of every weak-star convergent subsequence satisfies  $u \in \Phi_\alpha(v)$ .*

In general, no uniqueness can be shown under the above conditions, which is anyway not to be expected for the rather degenerate examples above. However, a weaker type of uniqueness can be inferred from the convexity and optimality condition

$$F^*(Fu - v) + \alpha p = 0, \quad p \in \partial J(u),$$

where  $\partial J(u)$  denotes the subdifferential

$$\partial J(u) = \{w \in \mathcal{U}^* \mid J(u) + \langle w, \tilde{u} - u \rangle \leq J(\tilde{u}) \ \forall \tilde{u} \in \mathcal{U}\}.$$

From the assumptions on  $F$  we see that  $F^*$  effectively maps to the predual space  $\mathcal{W}$ , thus the subgradients in the optimality condition effectively satisfy  $p \in \mathcal{W}$ , which is a weak regularity condition. A key concept needed in the following is the Bregman distance or generalized Bregman distance (cf. [11, 38]).

**Definition 3.2.** Let  $J : \mathcal{U} \rightarrow \mathbb{R} \cup \{+\infty\}$  be a convex proper functional and let  $u, \tilde{u} \in \mathcal{U}$  with  $p \in \partial J(u)$ . Then the (generalized) Bregman distance  $d_J^p(\tilde{u}, u)$  is defined by

$$d_J^p(\tilde{u}, u) = J(\tilde{u}) - J(u) - \langle p, \tilde{u} - u \rangle.$$

If  $\tilde{p} \in \partial J(\tilde{u})$ , the symmetric Bregman distance  $d_J^{\tilde{p}, p}(\tilde{u}, u)$  is defined by

$$d_J^{\tilde{p}, p}(\tilde{u}, u) = \langle \tilde{p} - p, \tilde{u} - u \rangle.$$

Now assume that there are two minimizers  $u_1$  and  $u_2$  of the variational regularization problem, then the difference in optimality conditions yields

$$F^*F(u_1 - u_2) + \alpha(p_1 - p_2) = 0$$

and from a duality product with  $u_1 - u_2$  we infer

$$\|F(u_1 - u_2)\|^2 + \alpha d_J^{p_1, p_2}(u_1, u_2) = 0.$$

Hence, by the nonnegativity of both terms we obtain uniqueness of the output value, i.e.,  $Fu_1 = Fu_2$  as well as a vanishing symmetric Bregman distance between  $u_1$  and  $u_2$ .

Finally, we can turn our attention to convergence properties of the regularization method. For this sake we use an exposition based on  $\Gamma$ -convergence (cf. [6]).

**Lemma 3.3.** *Let  $v_n \rightarrow v^* = Fu^*$  in  $\mathcal{V}$  and  $\alpha_n \rightarrow 0$ . Then the sequence of functionals  $D_{\alpha_n}$  defined by*

$$D_{\alpha_n}(u) = \frac{1}{2} \|Fu - v_n\|^2 + \alpha_n J(u)$$

$\Gamma$ -converges to

$$D_0(u) = \frac{1}{2} \|Fu - v^*\|^2$$

with respect to the weak-star topology in  $\mathcal{U}$ .

This kind of convergence is not strong enough to infer convergence of minimizers, in particular since there is no equicoercivity property. To achieve this, we need to rescale the functional, i.e. use  $\Gamma$ -convergence by development to the next order.

**Lemma 3.4.** *Let  $v_n \rightarrow v^* = Fu^*$  in  $\mathcal{V}$  and  $\alpha_n \rightarrow 0$  such that*

$$\frac{\|v_n - v^*\|^2}{\alpha_n} \rightarrow 0.$$

Then the sequence of functionals  $E_{\alpha_n}$  defined by

$$E_{\alpha_n}(u) = \frac{1}{2\alpha_n} \|Fu - v_n\|^2 + J(u)$$

$\Gamma$ -converges to

$$E_0(u) = \begin{cases} J(u) & \text{if } Fu = v^*, \\ +\infty & \text{else,} \end{cases}$$

with respect to the weak-star topology in  $\mathcal{U}$ .

Let us mention that we obtain divergence, i.e.  $E_{\alpha_n}$  converges to the functional identically equal to  $+\infty$ , if the condition on the parameter choice is violated, i.e.  $\liminf \frac{\|v_n - v^*\|^2}{\alpha_n} > 0$ . Since  $E_{\alpha_n} \geq J$  and  $J$  is coercive, we immediately conclude the equicoercivity of the sequence  $E_{\alpha_n}$ .

**Corollary 3.5.** *Let  $v_n \rightarrow v^* = Fu^*$  in  $\mathcal{V}$  and  $\alpha_n \rightarrow 0$  such that*

$$\frac{\|v_n - v^*\|^2}{\alpha_n} \rightarrow 0.$$

Moreover, let  $u_n$  be a sequence of minimizers of  $D_{\alpha_n}$  (or equivalently  $E_{\alpha_n}$ ), then there exists a subsequence converging with respect to the weak-star topology in  $\mathcal{U}$  and the limit  $u^{**}$  of each weakly convergent subsequence is a minimizer of  $E_0$ . Moreover,  $J(u_n) \rightarrow J(u^{**})$ .

Corollary 3.5 confirms that indeed the regularization operator defined by

$$\Phi_\alpha(v) = \arg \min_u D_\alpha(u)$$

yields a convergent regularization. Let us mention some further direct consequences.

- If the  $J$ -minimizing solution is unique, i.e.  $u^{**}$  is the unique minimizer of  $E_0$ , then the whole sequence  $u_n$  converges weakly-star to  $u^{**}$ . Moreover, if there is  $p^{**} \in \partial J(u) \cap \mathcal{W}$ , then due to the convergence of  $J$  and the weak-star convergence we conclude

$$d_J^{p^{**}}(u_n, u^{**}) \rightarrow 0.$$

- If  $u^*$  satisfies  $Fu^* = v^*$ , but is not a  $J$ -minimizing solution (a minimizer of  $E_0$ ), it cannot be reconstructed by the regularization method, i.e. it is not the limit of minimizers of the variational regularization for positive  $\alpha$ . This is related to the question whether the regularization functional introduces the right type of prior knowledge. If we are interested in reconstructing a solution like  $u^*$  that is not  $J$ -minimizing, then  $J$  is not a suitable choice.
- If  $J$  is the norm in  $\mathcal{U}$  as in many frequent examples and  $\mathcal{U}$  satisfies a Radon-Riesz property, the previous result indeed implies a strong convergence of subsequences.

The above analysis was based on a deterministic approach, but in a similar way a stochastic theory can be developed, e.g. for a sequence of random variables  $v_n$  with variance  $\mathbb{E}(\|v_n - v^*\|^2)$  converging to zero.

### 3.2. Quantitative estimates

As mentioned above, it is important to derive quantitative estimates between solutions of the regularized problem and ideal solutions, which we present here based on using range conditions as sketched above. In the following we denote by  $u_\alpha$  a regularized solution, i.e. a minimizer of  $D_\alpha$ . Due to convexity  $u_\alpha \in \Phi_\alpha(v)$  is characterized as the solution of the optimality condition

$$F^*(Fu_\alpha - v) + \alpha p_\alpha = 0, \quad p_\alpha \in \partial J(u_\alpha).$$

Taking two such solutions one can establish a stability estimate for the Bregman distance (cf. [3]).

**Theorem 3.6.** *Let  $u_\alpha \in \Phi_\alpha(v)$  and  $\tilde{u}_\alpha \in \Phi_\alpha(\tilde{v})$ . Then the estimate*

$$\frac{1}{2} \|Fu_\alpha - F\tilde{u}_\alpha\|^2 + \alpha d_J^{p_\alpha, \tilde{p}_\alpha}(u_\alpha, \tilde{u}_\alpha) \leq \frac{1}{2} \|v - \tilde{v}\|^2$$

*holds, where  $p_\alpha$  respectively  $\tilde{p}_\alpha$  are the subgradients appearing in the optimality condition for  $u_\alpha$  respectively  $\tilde{u}_\alpha$ .*

Now we turn to the range condition, effectively reformulating a result from [15].

**Lemma 3.7.** *An element  $u^* \in \mathcal{U}$  with  $v^* = Fu^*$  satisfies the range condition for the variational regularization operator  $\Phi_\alpha$  if and only if it satisfies the source condition*

$$\exists z^* \in \mathcal{V} : F^*z^* \in \partial J(u^*).$$

The key part of the proof is the explicit construction  $v_\alpha^* = v^* + \alpha z^*$ , which allows to obtain an estimate of the right-hand side in the error estimate, due to

$$\|v - v_\alpha^*\| \leq \|v - v^*\| + \|v^* - v_\alpha^*\| = \|v - v^*\| + \alpha \|z^*\|.$$

This leads to the error estimates as derived in [15].

**Corollary 3.8.** *Let  $u_\alpha \in \Phi_\alpha(v)$  and let  $v^* = Fu^*$ , with  $u^*$  satisfying the source condition  $p^* = F^*z^* \in \partial J(u^*)$ . Then the estimate*

$$\frac{1}{2} \|Fu_\alpha - Fu^*\|^2 + \alpha d_J^{p_\alpha, p^*}(u_\alpha, u^*) \leq \|v - v^*\|^2 + \alpha^2 \|z^*\|^2.$$

In the error estimate we see again the condition on the choice of  $\alpha$  needed for the convergence of regularization methods. While the estimate on the output error  $\|Fu_\alpha - Fu^*\|$  is uniform in  $\alpha$ ; the effective estimate for the Bregman distance is of the form

$$d_J^{p_\alpha, p^*}(u_\alpha, u^*) \leq \frac{\|v - v^*\|^2}{\alpha} + \alpha \|z^*\|^2,$$

which is small again only if  $\alpha$  and the quotient  $\frac{\|v - v^*\|^2}{\alpha}$  are small.

One also observes a bias-variance decomposition inherent in the estimate, even more clearly when we assume an underlying stochastic noise model, i.e.,  $v$  is a random variable. Without systematic errors in the measurements, we have  $\mathbb{E}(v) = v^*$  and hence

$$\mathbb{E}(d_J^{p_\alpha, p^*}(u_\alpha, u^*)) \leq \frac{\mathbb{E}(\|v - v^*\|^2)}{\alpha} + \alpha \|z^*\|^2.$$

The measure on the left-hand side is the natural generalization of the mean-squared error to the case of convex variational regularization, and the right-hand side is composed of the data variance and the bias term  $\|z^*\|^2$ , scaled by the regularization parameter.

Let us mention that the above estimates in Bregman distances lead to estimates in norms if  $J$  satisfies strong convexity conditions (cf. [54]). In the case of not strictly convex functionals, the Bregman distance can vanish even if  $u_\alpha \neq u^*$ , e.g. in total variation regularization if they differ by a change of contrast  $u_\alpha = h(u)$  with a monotone function  $h$ , but rather measures a deviation of the discontinuity sets (cf. [3, 16]). In such cases the multivaluedness of the subdifferential can even be an advantage that needs to be exploited, since we do not have just a single estimate, but actually an estimate for each  $p^*$  satisfying a source condition. Estimates for other quantities can then be derived from the Bregman distance estimates by optimizing over the possible  $p^*$  and the associated source elements  $z^*$  (respectively, their norm appearing in the error estimates). An example are estimates for the total variation regularization for piecewise constant functions; it has been shown already in [15] how the total variation of  $u_\alpha$  away from the discontinuity set of  $u^*$  can be estimated by choosing appropriate subgradients.

Again the above type of conditions and estimates are the canonical ones, but can be developed much farther (cf. e.g. [2, 25, 26, 28, 30–32, 51, 52, 57]). The first issue is the question of having better estimates under stronger conditions, and a typical example is an improved source condition  $p^* = F^*F\eta^* \in \partial J(u^*)$  for some  $\eta^* \in \mathcal{U}$ . In this case the element  $\eta^*$  can be used to construct an approximate solution  $u_\alpha^* = u^* - \alpha\eta^*$  instead of approximate data for a range condition. This was carried out in [51] (see also [29]) to obtain the estimate

$$d_J^{p^*}(u_\alpha, u^*) \leq d_J^{p^*}(u^* - \alpha\eta^*, u^*) + \frac{\|v - v^*\|^2}{2\alpha}.$$

The exact characterization of  $d_J^{p^*}(u^* - \alpha\eta^*, u^*)$  depends on the properties of the functional and may be on  $u^*$  itself. For  $J$  being Fréchet-differentiable with Lipschitz-continuous (or Hölder-continuous) derivative, it is always quadratic in  $\alpha$ ; hence the estimate is of higher order in  $\alpha$ . For the nonsmooth functionals like total variation or the  $\ell^1$ -norm the situation is different; at a first glance it cannot be expected that  $d_J^{p^*}(u^* - \alpha\eta^*, u^*)$  is of higher order in  $\alpha$ . However, in such situations we can even have  $d_J^{p^*}(u^* - \alpha\eta^*, u^*) = 0$  for  $\alpha$  small, e.g. in  $\ell^1$  regularization if the support of  $\eta^*$  is contained in the support of  $u^*$ .

The opposite question of weaker estimates arises if  $u^*$  does not satisfy the source condition  $p^* = F^*z^*$ . In this case approximate source conditions are used, which measure the deviation from the source condition. A frequently used concept is the so-called *distance function*

$$D_\rho(p^*) = \inf \{ \|F^*z - p^*\| \mid z \in \mathcal{V}, \|z\| \leq \rho \},$$

which is useful in particular under strong convexity assumptions and allows to build a theory in a similar way by optimizing the value  $\rho$  that finally appears in the error esti-

mate. For functionals not being strictly convex and in particular the one-homogeneous cases like total variation, a reformulation in terms of a dual problem is more suitable as seen in [13]. There the measure

$$e_{\alpha, \nu}(p^*) = \inf_{z \in \mathcal{V}} \nu J^* \left( \frac{F^* z - p^*}{\nu} \right) + \frac{\alpha}{2} \|z\|^2$$

was used to derive estimates. One observes some duality to the concept of distance functions, noticing that for  $J$  being a norm in Banach space we just have

$$e_{\alpha, \nu}(p^*) = \alpha \inf \{ \|z\| \mid z \in \mathcal{V}, \|F^* z - p^*\|_* \leq \nu \},$$

where  $\|\cdot\|_*$  is the dual norm to  $J$ . It was also shown that approximate source conditions are inherently related to the case of large noise, which is particularly relevant for stochastic models like white noise having non-finite variance (cf. [5, 13, 36]).

While the literature was focused on asymptotic results for a long time, the specific shape of solutions at a fixed positive  $\alpha$  became a more attractive topic in the last two decades. In order to understand this issue, a better understanding of the range condition for the regularization method is needed, which means the source condition  $p^* = F^* z^*$  in the case of variational regularization. Since  $F$  is modeled as a smoothing operator in inverse problems,  $F^*$  is smoothing as well, which implies that the source condition is an abstract smoothness condition. However, the smoothness is rather indirect, since it concerns the subgradient  $p^*$  and not directly  $u^*$ . Various results on the structure of minimizers, from sparsity properties for  $J = \ell^1$  or its counterpart in the space of measures to total variation and staircasing phenomena can be found in literature (cf. [18, 19]).

Another issue that found strong recent interest is debiasing, since in the case of large noise the bias caused by the regularization term (and the large value of  $\alpha$  that is needed to achieve stability) spoils the possible quality of regularized solutions. The influence of bias can also be seen from the term depending on  $\|z^*\|$  in the error estimates, and in practice it is often observed that the reconstruction of the subgradient is better than the one of the primal solution due to bias. First debiasing methods (also called refitting) appeared in  $\ell^1$  regularization, where in a first step the variational regularization is used and in a second step a simple least-squares problem is used on the support obtained from the first step, sometimes also with a sign constraint as obtained from the subgradient in the first step (cf. [21, 42]). This approach can be translated to a more general two-step approach for debiasing as worked out in [8], which computes

$$\Phi_\alpha(v) = \arg \min \{ d_J^{P_\alpha}(u, u_\alpha) \mid u_\alpha \in \Phi_\alpha^0(v) \},$$

with  $\Phi_\alpha^0$  being the regularization operator from the variational regularization method.

Another approach effectively leading to debiasing, but also with other advantages, are iterative regularization methods such as the Bregman iteration (cf. [49]). In the case of a quadratic functional, it can be formulated as an augmented Lagrangian method for computing the  $J$ -minimizing solution of  $Fu = v$ , i.e.

$$\begin{aligned} u^{k+1} &\in \arg \min_u \frac{1}{2} \|Fu - v^k\|^2 + \alpha J(u), \\ v^{k+1} &= v^k + v - Fu^{k+1}, \end{aligned}$$

with  $v^0 = v$ . To have a suitable generalization also for other loss functionals this can be reformulated as

$$\begin{aligned} u^{k+1} &\in \arg \min_u \frac{1}{2} \|Fu - v\|^2 + \alpha d_f^k(u, u^k), \\ p^{k+1} &= p^k + \frac{1}{\alpha} F^*(v - Fu^{k+1}) \in \partial J(u^{k+1}). \end{aligned}$$

The regularization parameter in this case is not  $\alpha$ , which is to be chosen rather larger in order to achieve good results, but the number of iterations carried out. Due to the variational structure in each iteration step, variational methods can be employed to prove well-definedness of the regularization operator, convergence, and error estimates. We refer to [3, 17, 49] for a detailed discussion of such iterative approaches and their analysis. Let us finally mention that in this respect there is another relation to machine learning, since Bregman iterations for  $\ell^1$  regularizations have been developed further recently for the training of sparse deep neural networks and their architecture design (cf. [9, 10]).

## 4. Variational regularization and generalization

In this final part we discuss some possible relations between the setup in machine learning and the above results on variational regularization theory. In particular we highlight some connections between the typical error measures used in the two fields, namely generalization errors on the one hand and Bregman distances on the other.

### 4.1. Error decomposition and generalization error

Let us return to the setup of machine learning with the minimization of the empirical risk with a convex loss  $\ell$ , taking the viewpoint that the ideal solution is the function  $f^*$ . While we have seen that naturally Bregman distances are estimated in the theory of variational regularization, the generalization error

$$G = \mathbb{E}_{(x,y) \sim \mathbb{P}} (\ell(f(x; \theta), y)) - \mathbb{E}_{(x,y) \sim \mathbb{P}^N} (\ell(f(x; \theta), y))$$

is the commonly used quantity in machine learning.



In order to understand the connections to Bregman distances, consider an ideal solution  $f^* \in \mathcal{F}$  minimizing the population risk, i.e.

$$f^* \in \arg \min_{f \in \mathcal{F}} \mathbb{E}_{(x,y) \sim \mathbb{P}} (\ell(f(x), y)) = \arg \min_{f \in \mathcal{F}} R^*(f).$$

Since the population risk is convex with respect to  $f$ , we conclude that  $0 \in \partial R^*(f^*)$ , which implies that

$$d_{R^*}^0(f(\cdot, \theta), f^*) = \mathbb{E}_{(x,y) \sim \mathbb{P}} (\ell(f(x; \theta), y)) - \mathbb{E}_{(x,y) \sim \mathbb{P}} (\ell(f^*(x), y)).$$

The latter can be decomposed in a similar spirit to the error decomposition in [4]:

$$\begin{aligned} d_{R^*}^0(f(\cdot, \theta), f^*) &= \mathbb{E}_{(x,y) \sim \mathbb{P}} (\ell(f(x; \theta), y)) - \mathbb{E}_{(x,y) \sim \mathbb{P}^N} (\ell(f(x; \theta), y)) \\ &\quad + \mathbb{E}_{(x,y) \sim \mathbb{P}^N} (\ell(f(x; \theta), y) - \ell(f^*(x), y)) \\ &\quad + \mathbb{E}_{(x,y) \sim \mathbb{P}^N} (\ell(f^*(x), y)) - \mathbb{E}_{(x,y) \sim \mathbb{P}} (\ell(f^*(x), y)). \end{aligned}$$

We see that the Bregman distance is decomposed into three parts: in addition to the generalization error in the first line, we have an approximation error in the second line (or rather a term that can be controlled with an approximation error in standard spaces) and a sampling error in the last line. The approximation error can be estimated beforehand or is often even negligible, since overparametrized models such as deep neural networks can usually be trained to have  $\mathbb{E}_{(x,y) \sim \mathbb{P}^N} (\ell(f(x; \theta), y)) \approx 0$  and the second part is nonpositive. Moreover, the last term vanishes on expectation over the sampling if  $\mathbb{P}^N$  is obtained from i.i.d. samples. Thus, in order to control the expected Bregman distance, the most important term is indeed the expected generalization error.

## 4.2. Estimates with operator errors and generalization

Errors due to sampling are effectively related to operator errors in inverse problems, which we see also from Example 2.2, where effectively the operator  $F$  is replaced by an operator  $\tilde{F}$  being the concatenation of  $F$  with a random sampling operator. Moreover, we assume again a source condition of the form  $p^* = F^* z^* \in \partial J(u^*)$ .

The generalization error in this notation is given by (noticing that we might need to use different norms for the two terms)

$$G(u) = \|Fu - v\|^2 - \|\tilde{F}u - \tilde{v}\|^2.$$

Hence, let us start again with the optimality condition of a regularized solution

$$u_\alpha \in \arg \min_u \frac{1}{2} \|\tilde{F}u - \tilde{v}\|^2 + \alpha J(u),$$

which is given by

$$\tilde{F}^*(\tilde{F}u_\alpha - \tilde{v}) + \alpha p_\alpha = 0, \quad p_\alpha \in \partial J(u_\alpha).$$

Rewriting to

$$F^*F(u_\alpha - u^*) + \alpha(p_\alpha - p^*) = F^*(Fu_\alpha - v) - \tilde{F}^*(\tilde{F}u_\alpha - \tilde{v}) - \alpha F^*z^*,$$

we are in a position to derive the kind of estimate we are after. A duality product with  $u_\alpha - u^*$  and several applications of Young's inequality imply

$$\frac{1}{4}\|F(u_\alpha - u^*)\|^2 + \alpha d_J^{p_\alpha, p^*}(u_\alpha, u^*) \leq \alpha^2\|z^*\|^2 + \|\tilde{F}u^* - \tilde{v}\|^2 + \frac{1}{2}G(u_\alpha).$$

In the case of consistent data, such as obtained from sampling  $F$ , we further have  $\tilde{v} = \tilde{F}u^*$ ; i.e., we obtain in particular

$$d_J^{p_\alpha, p^*}(u_\alpha, u^*) \leq \alpha\|z^*\|^2 + \frac{1}{2\alpha}G(u_\alpha).$$

Thus, the error in the Bregman distance is controlled by the systematic error and the generalization error.

### 4.3. Regularized risk minimization problems

The above arguments can be extended to convex risk minimization problems of the form

$$D_\alpha(\theta) = \mathbb{E}_{(x,y) \sim \mathbb{P}^N}(\ell(f(x; \theta), y) + \alpha J(\theta)).$$

For simplicity we assume that the model  $f$  is linear, i.e.  $f(x; \theta) = (F\theta)(x)$  with a linear operator  $F$  mapping to an appropriate function space  $\mathcal{F}$ , and  $\ell$  is the squared Euclidean norm. Consequently, we will consider  $F$  as a bounded linear operator from some parameter space  $\Theta$  to  $L_{\mathbb{P}}^2(\Omega)^m$  for some domain  $\Omega \subset \mathbb{R}^d$ . The ideal solution  $\theta^*$  is a minimizer of the population risk

$$R(\theta) = \mathbb{E}_{(x,y) \sim \mathbb{P}}(\|(F\theta)(x) - y\|^2).$$

With this setup, the regularization operator is given by

$$\Phi_\alpha(\mathbb{P}^N) = \arg \min_{\theta} \mathbb{E}_{(x,y) \sim \mathbb{P}^N} \left( \frac{1}{2} \|(F\theta)(x) - y\|^2 + \alpha J(\theta) \right). \quad (4.1)$$

Moreover, the source condition becomes

$$p^* = F^*z^* \in \partial J(\theta^*) \quad \text{with } z^* \in L_{\mathbb{P}}^2(\Omega)^m. \quad (4.2)$$

Similar to the reasoning in the previous section we can use the optimality condition

$$\mathbb{E}_{(x,y)\sim\mathbb{P}^N} \left( (F\theta_\alpha)(x) - y, F\theta' \right) = +\alpha p_\alpha = 0, \quad p_\alpha \in \partial J(\theta_\alpha)$$

for all  $\theta' \in \Theta$  to derive the following result.

**Theorem 4.1.** *Let  $\theta_\alpha \in \Phi_\alpha(\mathbb{P}^N)$  be defined by (4.1) and let the source condition (4.2) be satisfied. Then for appropriate  $p_\alpha \in \partial J(u_\alpha)$  the estimate*

$$\begin{aligned} & \frac{1}{4} \mathbb{E}_{(x,y)\sim\mathbb{P}} \left( \| (F\theta_\alpha)(x) - (F\theta^*)(x) \|^2 \right) + \alpha d_{J^{\alpha,p^*}}^{\alpha,p^*}(\theta_\alpha, \theta^*) \\ & \leq \frac{1}{2} G(\theta_\alpha) + \alpha^2 \|z^*\|^2 + \mathbb{E}_{(x,y)\sim\mathbb{P}^N} \left( \| (F\theta^*)(x) - y \|^2 \right), \end{aligned}$$

with the generalization error

$$G(\theta_\alpha) = \mathbb{E}_{(x,y)\sim\mathbb{P}} \left( \frac{1}{2} \| (F\theta_\alpha)(x) - y \|^2 \right) - \mathbb{E}_{(x,y)\sim\mathbb{P}^N} \left( \frac{1}{2} \| (F\theta_\alpha)(x) - y \|^2 \right).$$

**Funding.** Research supported in part by ERC via Grant EU FP7 ERC Consolidator Grant 615216 LifeInverse, by the German Ministry of Science and Technology (BMBF) under grant 05M2020 - DELETO, and by the EU under grant 2020 NoMADS - DLV-777826.

## References

- [1] S. Arridge, P. Maass, O. Öktem, and C.-B. Schönlieb, Solving inverse problems using data-driven models. *Acta Numer.* **28** (2019), 1–174 Zbl [1429.65116](#) MR [3963505](#)
- [2] M. Benning and M. Burger, Error estimates for general fidelities. *Electron. Trans. Numer. Anal.* **38** (2011), 44–68 Zbl [1285.47015](#) MR [2871859](#)
- [3] M. Benning and M. Burger, Modern regularization methods for inverse problems. *Acta Numer.* **27** (2018), 1–111 Zbl [1431.65080](#) MR [3826506](#)
- [4] J. Berner, P. Grohs, G. Kutyniok, and P. Petersen, The modern mathematics of deep learning. 2021, arXiv:[2105.04026](#)
- [5] N. Bissantz, T. Hohage, A. Munk, and F. Ruymgaart, Convergence rates of general regularization methods for statistical inverse problems and applications. *SIAM J. Numer. Anal.* **45** (2007), no. 6, 2610–2636 Zbl [1234.62062](#) MR [2361904](#)
- [6] A. Braides,  *$\Gamma$ -Convergence for Beginners*. Oxford Lecture Ser. Math. Appl. 22, Oxford University Press, Oxford, 2002 Zbl [1198.49001](#) MR [1968440](#)
- [7] K. Bredies and H. K. Pikkarainen, Inverse problems in spaces of measures. *ESAIM Control Optim. Calc. Var.* **19** (2013), no. 1, 190–218 Zbl [1266.65083](#) MR [3023066](#)

- [8] E.-M. Brinkmann, M. Burger, J. Rasch, and C. Sutour, Bias reduction in variational regularization. *J. Math. Imaging Vision* **59** (2017), no. 3, 534–566 Zbl [1385.49002](#) MR [3712429](#)
- [9] L. Bungert, T. Roith, D. Tenbrinck, and M. Burger, A Bregman learning framework for sparse neural networks. 2021, arXiv:[2105.04319](#)
- [10] L. Bungert, T. Roith, D. Tenbrinck, and M. Burger, Neural architecture search via Bregman iterations. 2021, arXiv:[2106.02479](#)
- [11] M. Burger, Bregman distances in inverse problems and partial differential equations. In *Advances in Mathematical Modeling, Optimization and Optimal Control*, pp. 3–33, Springer Optim. Appl. 109, Springer, Cham, 2016 Zbl [1346.90007](#) MR [3526561](#)
- [12] M. Burger and H. W. Engl, Training neural networks with noisy data as an ill-posed problem. *Adv. Comput. Math.* **13** (2000), no. 4, 335–354 Zbl [1126.41301](#) MR [1826332](#)
- [13] M. Burger, T. Helin, and H. Kekkonen, Large noise in variational regularization. *Trans. Math. Appl.* **2** (2018), no. 1, tny002 Zbl [1404.49025](#) MR [3899949](#)
- [14] M. Burger, B. Kaltenbacher, and A. Neubauer, Iterative solution methods. In *Handbook of Mathematical Methods in Imaging. Vol. 1, 2, 3*, pp. 431–470, Springer, New York, 2nd edn., 2015 Zbl [1331.65080](#) MR [3560073](#)
- [15] M. Burger and S. Osher, Convergence rates of convex variational regularization. *Inverse Problems* **20** (2004), no. 5, 1411–1421 Zbl [1068.65085](#) MR [2109126](#)
- [16] M. Burger and S. Osher, A guide to the TV zoo. In *Level Set and PDE Based Reconstruction Methods in Imaging*, pp. 1–70, Lecture Notes in Math. 2090, Springer, Cham, 2013 Zbl [1342.94014](#) MR [3203352](#)
- [17] M. Burger, E. Resmerita, and L. He, Error estimation for Bregman iterations and inverse scale space methods in image restoration. *Computing* **81** (2007), no. 2-3, 109–135 Zbl [1147.68790](#) MR [2354192](#)
- [18] V. Caselles, A. Chambolle, and M. Novaga, The discontinuity set of solutions of the TV denoising problem and some extensions. *Multiscale Model. Simul.* **6** (2007), no. 3, 879–894 Zbl [1145.49024](#) MR [2368971](#)
- [19] A. Chambolle, V. Caselles, D. Cremers, M. Novaga, and T. Pock, An introduction to total variation for image analysis. In *Theoretical Foundations and Numerical Methods for Sparse Recovery*, pp. 263–340, Radon Ser. Comput. Appl. Math. 9, Walter de Gruyter, Berlin, 2010 Zbl [1209.94004](#) MR [2731599](#)
- [20] J. Cheng and B. Hofmann, Regularization methods for ill-posed problems. In *Handbook of Mathematical Methods in Imaging. Vol. 1, 2, 3*, pp. 91–123, Springer, New York, 2nd edn., 2015 Zbl [1331.65081](#) MR [3560066](#)
- [21] C.-A. Deledalle, N. Papadakis, J. Salmon, and S. Vaiter, CLEAR: Covariant LEAsT-square Refitting with applications to image restoration. *SIAM J. Imaging Sci.* **10** (2017), no. 1, 243–284 Zbl [1365.49034](#) MR [3615463](#)
- [22] Q. Denoyelle, V. Duval, and G. Peyré, Support recovery for sparse super-resolution of positive measures. *J. Fourier Anal. Appl.* **23** (2017), no. 5, 1153–1194 Zbl [1417.65223](#) MR [3704760](#)

- [23] I. Ekeland and R. Témam, *Convex Analysis and Variational Problems*. Classics Appl. Math. 28, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1999 Zbl [0939.49002](#) MR [1727362](#)
- [24] H. W. Engl, M. Hanke, and A. Neubauer, *Regularization of Inverse Problems*. Math. Appl. 375, Kluwer Academic Publishers Group, Dordrecht, 1996 Zbl [0859.65054](#) MR [1408680](#)
- [25] J. Flemming, Variational smoothness assumptions in convergence rate theory—an overview. *J. Inverse Ill-Posed Probl.* **21** (2013), no. 3, 395–409 Zbl [1293.47058](#) MR [3069356](#)
- [26] J. Flemming and B. Hofmann, A new approach to source conditions in regularization with general residual term. *Numer. Funct. Anal. Optim.* **31** (2010), no. 1-3, 254–284 Zbl [1195.47040](#) MR [2677255](#)
- [27] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Adapt. Comput. Mach. Learn., MIT Press, Cambridge, MA, 2016 Zbl [1373.68009](#) MR [3617773](#)
- [28] M. Grasmair, Linear convergence rates for Tikhonov regularization with positively homogeneous functionals. *Inverse Problems* **27** (2011), no. 7, 075014 Zbl [1219.35359](#) MR [2817427](#)
- [29] M. Grasmair, Variational inequalities and higher order convergence rates for Tikhonov regularisation on Banach spaces. *J. Inverse Ill-Posed Probl.* **21** (2013), no. 3, 379–394 Zbl [1288.47016](#) MR [3069355](#)
- [30] B. Hofmann, B. Kaltenbacher, C. Pöschl, and O. Scherzer, A convergence rates result for Tikhonov regularization in Banach spaces with non-smooth operators. *Inverse Problems* **23** (2007), no. 3, 987–1010 Zbl [1131.65046](#) MR [2329928](#)
- [31] B. Hofmann and P. Mathé, Parameter choice in Banach space regularization under variational inequalities. *Inverse Problems* **28** (2012), no. 10, 104006 Zbl [1253.47041](#) MR [2987901](#)
- [32] T. Hohage and F. Weidling, Characterizations of variational source conditions, converse results, and maxisets of spectral regularization methods. *SIAM J. Numer. Anal.* **55** (2017), no. 2, 598–620 Zbl [1432.65070](#) MR [3623207](#)
- [33] G. James, D. Witten, T. Hastie, and R. Tibshirani, *An Introduction to Statistical Learning*. Springer Texts Statist. 103, Springer, New York, 2013 Zbl [1281.62147](#) MR [3100153](#)
- [34] B. Kaltenbacher, A. Neubauer, and O. Scherzer, *Iterative Regularization Methods for Non-linear Ill-Posed Problems*. Radon Ser. Comput. Appl. Math. 6, Walter de Gruyter, Berlin, 2008 Zbl [1145.65037](#) MR [2459012](#)
- [35] B. Kaltenbacher, F. Schöpfer, and T. Schuster, Iterative methods for nonlinear ill-posed problems in Banach spaces: convergence and applications to parameter identification problems. *Inverse Problems* **25** (2009), no. 6, 065003 Zbl [1176.65070](#) MR [2506848](#)
- [36] H. Kekkonen, M. Lassas, and S. Siltanen, Analysis of regularized inversion of data corrupted by white Gaussian noise. *Inverse Problems* **30** (2014), no. 4, 045009 Zbl [1287.35101](#) MR [3178119](#)

- [37] A. Kirsch, *An Introduction to the Mathematical Theory of Inverse Problems*. 2nd edn., Appl. Math. Sci. 120, Springer, New York, 2011 Zbl [1213.35004](#) MR [3025302](#)
- [38] K. C. Kiwiel, Proximal minimization methods with generalized Bregman functions. *SIAM J. Control Optim.* **35** (1997), no. 4, 1142–1168 Zbl [0890.65061](#) MR [1453294](#)
- [39] N. Kovachki, Z. Li, B. Liu, K. Azizzadenesheli, K. Bhattacharya, A. Stuart, and A. Anandkumar, Neural operator: Learning maps between function spaces. 2021, arXiv:[2108.08481](#)
- [40] J. Kukačka, V. Golkov, and D. Cremers, Regularization for deep learning: A taxonomy. 2017, arXiv:[1710.10686](#)
- [41] M. M. Lavrentiev, *Some Improperly Posed Problems of Mathematical Physics*. Springer Tracts Nat. Philos. 11, Springer, New York, 1967 Zbl [0149.41902](#)
- [42] J. Lederer, Trust, but verify: benefits and pitfalls of least-squares refitting in high dimensions. 2013, arXiv:[1306.0113](#)
- [43] A. K. Louis, Approximate inverse for linear and some nonlinear problems. *Inverse Problems* **12** (1996), no. 2, 175–190 Zbl [0851.65036](#) MR [1382237](#)
- [44] A. K. Louis and P. Maass, A mollifier method for linear operator equations of the first kind. *Inverse Problems* **6** (1990), no. 3, 427–440 Zbl [0713.65040](#) MR [1057035](#)
- [45] M. Mohri, A. Rostamizadeh, and A. Talwalkar, *Foundations of Machine Learning*. Adapt. Comput. Mach. Learn., MIT Press, Cambridge, MA, 2018 Zbl [1407.68007](#) MR [3931734](#)
- [46] F. Natterer, *The Mathematics of Computerized Tomography*. Classics Appl. Math. 32, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2001 Zbl [0973.92020](#) MR [1847845](#)
- [47] N. H. Nelsen and A. M. Stuart, The random feature model for input-output maps between Banach spaces. *SIAM J. Sci. Comput.* **43** (2021), no. 5, A3212–A3243 Zbl [07398767](#) MR [4313849](#)
- [48] B. Neyshabur, R. Tomioka, and N. Srebro, In search of the real inductive bias: On the role of implicit regularization in deep learning. In *International Conference on Learning Representations (ICLR)*, 2015
- [49] S. Osher, M. Burger, D. Goldfarb, J. Xu, and W. Yin, An iterative regularization method for total variation-based image restoration. *Multiscale Model. Simul.* **4** (2005), no. 2, 460–489 Zbl [1090.94003](#) MR [2162864](#)
- [50] R. Ramlau and G. Teschke, Sparse recovery in inverse problems. In *Theoretical Foundations and Numerical Methods for Sparse Recovery*, pp. 201–262, Radon Ser. Comput. Appl. Math. 9, Walter de Gruyter, Berlin, 2010 Zbl [1210.65107](#) MR [2731600](#)
- [51] E. Resmerita, Regularization of ill-posed problems in Banach spaces: convergence rates. *Inverse Problems* **21** (2005), no. 4, 1303–1314 Zbl [1082.65055](#) MR [2158110](#)
- [52] E. Resmerita and O. Scherzer, Error estimates for non-quadratic regularization and the relation to enhancement. *Inverse Problems* **22** (2006), no. 3, 801–814 Zbl [1103.65062](#) MR [2235638](#)

- [53] L. Rosasco, A. Caponnetto, E. Vito, F. Odone, and U. Giovannini, Learning, regularization and ill-posed inverse problems. In *Advances in Neural Information Processing Systems*, edited by L. Saul, Y. Weiss, and L. Bottou, pp. 1145–1152, 17, 2005
- [54] T. Schuster, B. Kaltenbacher, B. Hofmann, and K. S. Kazimierski, *Regularization Methods in Banach Spaces*. Radon Ser. Comput. Appl. Math. 10, Walter de Gruyter, Berlin, 2012 Zbl [1259.65087](#) MR [2963507](#)
- [55] T. I. Seidman and C. R. Vogel, Well-posedness and convergence of some regularisation methods for nonlinear ill posed problems. *Inverse Problems* **5** (1989), no. 2, 227–238 Zbl [0691.35090](#) MR [991919](#)
- [56] A. N. Tikhonov and V. Y. Arsenin, *Solutions of Ill-Posed Problems*. Scripta Series in Mathematics, John Wiley & Sons, New York, 1977 Zbl [0354.65028](#) MR [0455365](#)
- [57] F. Weidling, *Variational source conditions and conditional stability estimates for inverse problems in PDEs*. Ph.D. thesis, Göttingen University, 2019
- [58] F. Werner and B. Hofmann, Convergence analysis of (statistical) inverse problems under conditional stability estimates. *Inverse Problems* **36** (2020), no. 1, 015004 Zbl [07153871](#) MR [4047904](#)

**Martin Burger**

Department of Mathematics and Research Center for Mathematics of Data,  
Friedrich-Alexander-Universität Erlangen-Nürnberg, Cauerstr. 11, 91058 Erlangen, Germany;  
[martin.burger@fau.de](mailto:martin.burger@fau.de)