

SUBSPACE RECYCLING-BASED REGULARIZATION METHODS*

RONNY RAMLAU[†], KIRK M. SOODHALTER[‡], AND VICTORIA HUTTERER[§]

Abstract. Subspace recycling techniques have been used quite successfully for the acceleration of iterative methods for solving large-scale linear systems. These methods often work by augmenting a solution subspace generated iteratively by a known algorithm with a fixed subspace of vectors which are “useful” for solving the problem. Often, this has the effect of inducing a projected version of the original linear system to which the known iterative method is then applied, and this projection can act as a deflation preconditioner, accelerating convergence. Most often, these methods have been applied for the solution of well-posed problems. However, they have also begun to be considered for the solution of ill-posed problems.

In this paper, we consider subspace augmentation-type iterative schemes applied to linear ill-posed problems in a continuous Hilbert space setting, based on a recently developed framework describing these methods. We show that under suitable assumptions, a recycling method satisfies the formal definition of a regularization, as long as the underlying scheme is itself a regularization. We then develop an augmented subspace version of the gradient descent method and demonstrate its effectiveness, both on an academic Gaussian blur model and on problems arising from the adaptive optics community for the resolution of large sky images by ground-based extremely large telescopes.

Key words. ill-posed problems, augmented methods, recycling, iterative methods, Landweber, gradient descent

AMS subject classifications. 68Q25, 68R10, 68U05

1. Introduction. This paper concerns the use of subspace recycling iterative methods in the context of solving linear ill-posed problems. There have already been a number of augmented or subspace recycling methods proposed specifically to treat ill-posed problems (e.g., [4, 3, 9, 5]), but these methods have yet to be formally analyzed as regularization schemes. More recently, an augmented scheme to accelerate the regularization of a problem in adaptive optics, producing a large improvement in performance, was proposed [32]. However, the augmentation was done after the regularized problem was set up rather than actually combining these techniques. It would thus be helpful to place all these methods into a common framework to streamline the development of new methods and to enable them to be analyzed using a common language and framework. In particular, it enables us in this paper to prove under what conditions augmented approaches described by the framework are regularizations.

We take advantage of a newly-proposed general framework [36] which can be used to describe the vast majority of these methods. This framework allows us to perform a general analysis of augmentation and recycling schemes as regularizations. From this, we develop sufficient conditions for a recycled iterative method to be a regularization, namely that a subspace recycling scheme is a regularization if it is built from an underlying regularization method. As a proof-of-concept, we then leverage the framework and our analysis to propose augmented steepest descent and Landweber methods, which we then apply to some test problems.

*DATE.

[†]Industrial Mathematics Institute, Kepler University Linz, and Johann Radon Institute for Computational and Applied Mathematics (RICAM), Linz, Austria.(ronny.ramlau@jku.at).

[‡] School of Mathematics, Trinity College Dublin, The University of Dublin, College Green, Dublin 2, Ireland.(ksoodha@maths.tcd.ie).

[§]Industrial Mathematics Institute, Kepler University Linz, Austria.(victoria.hutterer@indmath.uni-linz.ac.at).

The linear inverse problem we consider is of the form

$$(1.1) \quad Tx = y$$

whereby we approximately reconstruct x from noisy data y^δ where $y^\delta = y + n$, and n represents the perturbation of our data which satisfies $\|n\| \leq \delta$. Here, $T : \mathcal{X} \rightarrow \mathcal{Y}$ is a continuous mapping between separable Hilbert spaces \mathcal{X}, \mathcal{Y} . In many relevant applications, the operator T is ill-posed, i.e., (1.1) violates Hadamard's conditions: a solution either might not exist, it might be non-unique or it may not depend continuously on the data. In this case regularization techniques have to be employed for the stable solution of (1.1), see [12, 13] and Section 2 for a summary on these methods.

The quality of the reconstruction for a specific problem depends heavily on available additional information on the solution, which is often directly included into the chosen reconstruction approach. E.g., source conditions of the solution are used to estimate the reconstruction error [13, 27]. Alternatively, information on the smoothness of the solution can be incorporated into Tikhonov regularization by choosing a penalty term that enforces the reconstruction to be smooth. For iterative regularization methods, the iteration can be started with a rough guess for the solution, leading generally to shorter reconstruction times and better results. Those approaches are of particular interest for inverse problems that have to be solved repeatedly in time and where the previous reconstruction can be considered as a good guess for the solution of the current problem. This would be the case, e.g., for *Adaptive Optics* applications in Astronomy, where the image quality loss due to turbulences in the atmosphere is corrected by a suitable shaped deformable mirror. The shape of the deformable mirrors has to be adjusted within milliseconds. The optimal shape of the mirror is derived by solving the atmospheric tomography problem repeatedly (about every 1-2ms). Although the atmosphere changes rapidly it can be regarded *frozen* on a sub-millisecond timescale, meaning that the previous reconstruction of the atmospheric tomography problem can be regarded as a good approximation to the new solution [40, 31].

In this paper we want to consider the case that we know a finite-dimensional subspace $\mathcal{U} \subset \mathcal{X}$ that contains a significant part of the solution x^\dagger which we seek. As we will see, this part can be easily and stably be reconstructed. What then remains is to reconstruct the part that belongs to the orthogonal complement \mathcal{U}^\perp of \mathcal{U} , possibly induces an acceleration in the convergence rate of the resulting iterative method, cf. Section 5.

For finite dimensional problems, this approach is known as a *subspace recycling method* and is currently a hot topic in numerical linear algebra. The goal of this paper is therefore to extend the theory to ill-posed problems in an infinite dimensional setting. Specifically, we will show that the concept of subspace recycling can be used in connection with most of the standard regularization methods.

The paper is organized as follows: In Section 2 we give a short introduction to Inverse Problems and regularization methods, whereas Section 3 discusses the state of research of subspace recycling methods. In Sections 4 and 5 we show that standard linear regularization method can be combined with subspace recycling and prove that the resulting method is still a regularization. Finally, Section 6 considers recycled gradient descent methods, and Section 7 presents some numerical results.

2. Regularization of linear ill-posed problems. We begin with a brief review of some basic regularization theory, which we mainly base on [12, 13, 27]. As mentioned above, a solution to the problem $Tx = y$ might not exist; if it exists it might not be unique and/or might not depend continuously on the data. The first two

problems can be circumvented by using the generalized inverse T^\dagger , whereas stability is enforced by using *regularization methods*. To be more specific, an operator

$$(2.1) \quad R_\alpha : Y \rightarrow X$$

is a regularization for a linear operator $T : X \rightarrow Y$ iff for any right hand side y^δ with $\|y - y^\delta\| \leq \delta$ there exists a parameter choice rule

$$(2.2) \quad \alpha : \mathbb{R} \times Y \rightarrow \mathbb{R}^+$$

such that

$$(2.3) \quad \lim_{\delta \rightarrow 0} R_{\alpha(\delta)} y^\delta = T^\dagger y =: x^\dagger$$

holds. Here, the parameter is chosen such that $\lim_{\delta \rightarrow 0} \alpha(\delta) = 0$. Over the last decades, many regularization methods have been investigated, most prominently *Tikhonov regularization* and iterative methods like *Landweber iteration* [13, 24] and the *conjugate gradient* method. For Tikhonov regularization, the approximation x_α^δ to the generalized solution x^\dagger is computed as the minimizer of the Tikhonov functional

$$(2.4) \quad J_\alpha(x) = \|y^\delta - Tx\|_Y^2 + \alpha\Omega(x),$$

where $\Omega(x)$ is a suitable penalty term, e.g., $\Omega(x) = \|x\|^2$. For more general penalties we refer, e.g., to [12, 35]. For Tikhonov to be rendered a regularization method, one appropriate method of choosing the regularization parameter is the discrepancy principle, where α is chosen such that

$$(2.5) \quad \|y^\delta - Tx_\alpha^\delta\| = \tau\delta, \quad \tau > 1.$$

Iterative methods are popular as their computational realization is often straightforward. Landweber iteration computes as

$$(2.6) \quad x_{k+1}^\delta = x_k^\delta + \beta T^*(y^\delta - Tx_k^\delta), \quad 0 < \beta < 2/\|T\|^2.$$

For iterative methods the iteration index plays the role of the regularization parameter. According to the discrepancy principle, the iteration is terminated at step k when for the first time $\|y^\delta - Tx_k^\delta\| \leq \tau\delta$ with $\tau > 1$ fixed. The Landweber iteration with the discrepancy principle is a regularization method [13]. It is well known that the convergence of regularization methods for $\delta \rightarrow 0$ can be arbitrarily slow. Convergence rates can be obtained by using *source conditions*. Popular are Hölder type source conditions

$$(2.7) \quad x^\dagger = (T^*T)^\nu w$$

for some $\mu \geq 0$ and $w \in X$. A regularization method is called order optimal iff the reconstruction error can be estimated as

$$(2.8) \quad \|x^\dagger - R_{\alpha(\delta)} y^\delta\| = \mathcal{O}(\delta^{\nu/(\nu+1)})$$

provided a source condition (2.7) is fulfilled. Both Tikhonov and Landweber equipped with suitable source conditions are order optimal [13].

In this paper we will show that subspace recycling-based methods also form a regularization.

3. Augmented iterative methods. Augmented iterative methods, as we discuss them here, have been proposed in the literature beginning in the mid- to late-nineties; see [36] for a survey of this history. They can be considered as a specialized version of the class of iterative methods called *residual constraint* or *residual projection* methods. We discuss these methods briefly to put them into context, but for more details, the reader should see the survey paper [36], and references therein.

3.1. Iterations determined by residual constraint. Many iterative methods which are commonly used to treat both well- and ill-posed linear problems can be formulated according to a correction/constraint setup. Our discussion here in [Section 3](#) focuses on the well-posed problem case for simplicity. Let $(\cdot, \cdot)_{\mathcal{X}}$ and $(\cdot, \cdot)_{\mathcal{Y}}$, respectively, be the inner products for \mathcal{X} and \mathcal{Y} and $(Tx, y)_{\mathcal{X}} = (x, T^*y)_{\mathcal{Y}}$.

Consider the case wherein we are solving (1.1) with no initial approximation; i.e., $x_0 = 0$. The strategies we outline here take the abstract form of approximating the solution x^\dagger by $\hat{x} \in \mathcal{S}$ such that the residual satisfies $y - T\hat{x} \perp \tilde{\mathcal{S}}$, where $\mathcal{S} \subset \mathcal{X}$ and $\tilde{\mathcal{S}} \subset \mathcal{Y}$ are such that $\dim \mathcal{S} = \dim \tilde{\mathcal{S}} < \infty$. They are called the *correction* and *constraint* spaces, respectively. A common example of a case for which $\dim \mathcal{S} = 1$ is steepest descent (i.e., gradient descent with step size chosen to minimize the T^*T -norm of the error); see, e.g., [34, Section 5.3.1].

PROPOSITION 3.1. *The $(j+1)$ st step of steepest descent can be formulated as*

$$(3.1) \quad \text{select } t_{j+1} \in \text{span}\{T^*r_j\} \quad \text{such that } y - T(x_j + t_{j+1}) \perp \text{span}\{TT^*r_j\},$$

where $r_j = y - Tx_j$; i.e., $\mathcal{S} = \text{span}\{T^*r_j\}$, and $\tilde{\mathcal{S}} = \text{span}\{TT^*r_j\}$.¹

Proof. Suppose we have constructed the approximation x_j at the previous iteration. Recall that gradient descent methods compute $x_{j+1} = x_j + \alpha T^*r_j$, where $r_j = y - Tx_j$. The choice of α (called the step size) determines the method's behavior. Steepest descent computes

$$\alpha_{j+1} = \underset{\alpha \in \mathbb{R}}{\operatorname{argmin}} \|y - T(x_j + \alpha T^*r_j)\|_{\mathcal{Y}}.$$

Differentiating the right-hand side, setting it equal to zero, and solving for α_{j+1} yields the minimizer $\alpha_{j+1} = \frac{(r_j, TT^*r_j)_{\mathcal{Y}}}{(TT^*r_j, TT^*r_j)_{\mathcal{Y}}}$. One must simply show that (3.1) produces the same choice of α_{j+1} . However, this constraint condition is equivalent to solving the equation

$$(r_j - \alpha_j TT^*r_j, TT^*r_j)_{\mathcal{Y}} = 0,$$

which when solved for α_j produces the same T^*T -norm-minimizing choice of α_j . \square

Steepest descent is an example of a method for which $\dim \mathcal{S} = 1$. There are many methods for which the size of the correction and constraint spaces grows at each iteration. It is a well-known result that one can characterize the method of conjugate gradients in this way, which we present without proof; see [33, 34] for proofs.

PROPOSITION 3.2. *The $(j+1)$ st iterate produced by the method of conjugate gradients applied to the normal equations associated to (1.1) can be formulated as*

$$(3.2) \quad \text{select } x_{j+1} \in \mathcal{K}_{j+1}(T^*T, T^*y) \quad \text{such that } y - Tx_{j+1} \perp T\mathcal{K}_{j+1}(T^*T, T^*y),$$

¹One could also, with some work, use this result to obtain a residual constraint formulation for a more general Landweber method, but it is not clear that this brings an analytic advantage.

where $\mathcal{K}_{j+1}(T^*T, T^*y) = \text{span} \{y, T^*Ty, (T^*T)^2y, \dots, (T^*T)^jy\}$.

3.2. Augmentation via constraint over a pair of spaces. Many flavors of augmented iterative methods have been proposed to treat both well- and ill-posed problems. These methods treat the situation in which we wish to approximate the solution using an iterative method, but we also have a fixed subspace \mathcal{U} which we also want to use to build the approximation. Using the correction/constraint formulation in [Subsection 3.1](#), we can describe many of these methods as selecting approximations from a correction space $\mathcal{S}_j = \mathcal{U} + \mathcal{V}_j \subset \mathcal{X}$ with a residual constraint space $\tilde{\mathcal{S}}_j = \tilde{\mathcal{U}} + \tilde{\mathcal{V}}_j \subset \mathcal{Y}$, where \mathcal{U} and $\tilde{\mathcal{U}}$ are spaces of fixed dimension k and \mathcal{V}_j and $\tilde{\mathcal{V}}_j$ are generated iteratively, with dimension j at step j . It has been shown that this correction/constraint over the sum of subspaces reduces mathematically to applying the underlying iterative method to a projected problem, obtaining part of the approximation from \mathcal{V}_j and then obtaining the corrections from \mathcal{U} afterward. This will be explained in more detail in [Section 5](#), but we give a brief outline here.

A projector is uniquely defined by its range and its null space. A standard result from linear algebra shows that this is equivalent to defining the ranges of the projector and its complement. In this setting, consider the projectors $P \in \mathcal{L}(\mathcal{X})$ and $Q \in \mathcal{L}(\mathcal{Y})$ having the relationship that

$$(3.3) \quad TP = QT.$$

Let $\text{Range}(P) = \mathcal{U}$ and $\text{Null}(P) = (T^*\tilde{\mathcal{U}})^\perp$, and $\text{Range}(Q) = T\mathcal{U}$ and $\text{Null}(Q) = (\tilde{\mathcal{U}})^\perp$. One can easily show that defining the null spaces is equivalent to defining $\text{Range}(I_{\mathcal{X}} - P) = (T^*\tilde{\mathcal{U}})^\perp$ and $\text{Range}(I_{\mathcal{Y}} - Q) = \tilde{\mathcal{U}}^\perp$. For example, let $v \in \mathcal{X}$ be chosen arbitrarily. It can be decomposed as $v = v_{\mathcal{R}} + v_{\mathcal{N}}$ with $v_{\mathcal{R}} \in \text{Range}(P)$ and $v_{\mathcal{N}} \in \text{Null}(P)$. As expected, $Pv = Pv_{\mathcal{R}} + Pv_{\mathcal{N}} = v_{\mathcal{R}}$. We can then observe that $(I_{\mathcal{X}} - P)v = v_{\mathcal{R}} + v_{\mathcal{N}} - Pv_{\mathcal{R}} - Pv_{\mathcal{N}} = v_{\mathcal{N}}$. As $v \in \mathcal{X}$ was chosen arbitrarily, we see that $\text{Range}(I_{\mathcal{X}} - P) = \text{Null}(P)$. The same can be shown for Q .

Approximating the solution of [\(1.1\)](#) by $x_j \in \mathcal{S}_j$ subject to the constraint $r_j \perp \tilde{\mathcal{S}}_j$ has been shown to be equivalent to approximating the solution t to the projected problem

$$(3.4) \quad (I - Q)Tt = (I - Q)y$$

by $t_j \in \mathcal{V}_j$ with residual constraint space $\tilde{\mathcal{V}}_j$ and setting the approximation $x_j \in \mathcal{U} + \mathcal{V}_j$ to be

$$(3.5) \quad x_j = Px^\dagger + (I - P)t_j;$$

see, [\[7\]](#) and [\[15, 16, 18\]](#).

For well-posed problems, the space \mathcal{U} often contains vectors from previous iteration cycles applied to the current or previously-solved related problem; see, e.g., [\[2, 23, 29, 38\]](#). These vectors are often (but not always) approximate eigenvectors (of a square operator) or left singular vectors, guided by Krylov subspace method convergence theory. For ill-posed problems, the strategy is generally to include some known features of the image/signal being reconstructed (see, e.g., [\[3, 4, 9\]](#)), but some current work by [\[5\]](#) still seeks to augment a Lanczos bidiagonalization-based solver for ill-posed, tall, skinny least-squares problems using approximate left singular vectors

to accelerate semi-convergence and reduce the influence of smaller singular values to achieve a regularizing effect. See [Subsection 5.1](#) for further discussion.

There is an alternative approach (originally proposed in the context of domain decomposition in, e.g., [\[14\]](#)) which begins with the projectors P and Q (rather than treating them as a consequence of the correction/constraint formulation) and produces the same algorithms but provides more flexibility to use methods like Landweber which do not have a clean residual constraint condition. This will be derived in detail in [Section 5](#), but it is built on the simple idea that the projector P can be used to get the decomposition $x^\dagger = Px^\dagger + (I - P)x^\dagger$. We will show that Px^\dagger can be cheaply computed and we can use an iterative method to approximate $(I - P)x^\dagger$.

4. Linear Projection Methods. We consider here the case that $\tilde{\mathcal{U}} = T\mathcal{U}$. For a finite dimensional subspace $\mathcal{U} \subset \mathcal{X}$ with $\dim \mathcal{U} = k > 0$, we denote the sibling subspace $\mathcal{C} = T\mathcal{U} \subset \mathcal{Y}$ and assume it also has dimension k . We can represent these subspaces in terms of bases, which we denote with a matrix-like notation. Let

$$(4.1) \quad \mathbf{U} = [u_1 \quad u_2 \quad \cdots \quad u_k] \in \mathcal{X}^k \quad \text{and} \quad \mathbf{C} = [c_1 \quad c_2 \quad \cdots \quad c_k] \in \mathcal{Y}^k$$

be k -tuples of elements from \mathcal{U} and \mathcal{C} , respectively, which are bases for these spaces.

Remark 4.1. To get a sense of these objects, observe that if $\mathcal{X} = \mathcal{Y} = \mathbb{R}^n$ then we would have $\mathbf{U}, \mathbf{C} \in \mathbb{R}^{n \times k}$; so this is just the generalization of the notion of a tall, skinny matrix.

We can consider \mathbf{U} and \mathbf{C} to be linear mappings induced by a basis expansion operation, with $\mathbf{U} \in \mathcal{L}(\mathbb{R}^k, \mathcal{X})$ and $\mathbf{C} \in \mathcal{L}(\mathbb{R}^k, \mathcal{Y})$, defined by

$$(4.2) \quad \mathbf{U}\mathbf{z} = \sum_{i=1}^k z_i u_i \quad \text{and} \quad \mathbf{C}\mathbf{z} = \sum_{i=1}^k z_i c_i \quad \text{for} \quad \mathbf{z} = [z_1 \quad z_2 \quad \cdots \quad z_k]^T \in \mathbb{R}^k.$$

Let

$$(4.3) \quad (x, \mathbf{U})_{\mathcal{X}} = \begin{bmatrix} (x, u_1)_{\mathcal{X}} \\ \vdots \\ (x, u_k)_{\mathcal{X}} \end{bmatrix} \in \mathbb{R}^k, \quad \text{and} \quad (y, \mathbf{C})_{\mathcal{Y}} = \begin{bmatrix} (y, c_1)_{\mathcal{Y}} \\ \vdots \\ (y, c_k)_{\mathcal{Y}} \end{bmatrix} \in \mathbb{R}^k.$$

Furthermore, for $\mathbf{M}, \mathbf{L} \in \mathcal{X}^k$, we define the bilinear mapping

$$(4.4) \quad (\cdot, \cdot)_{\mathcal{X}} : \mathcal{X}^k \times \mathcal{X}^k \rightarrow \mathbb{R}^{k \times k} \quad \text{with} \quad (\mathbf{M}, \mathbf{L}) \mapsto ((m_i, \ell_j)_{\mathcal{X}})_{i,j=1}^k \in \mathbb{R}^{k \times k}.$$

We similarly define the corresponding bilinear mapping for two elements from $\mathbf{F}, \mathbf{G} \in \mathcal{Y}^k$

$$(\cdot, \cdot)_{\mathcal{Y}} : \mathcal{Y}^k \times \mathcal{Y}^k \rightarrow \mathbb{R}^{k \times k} \quad \text{with} \quad (\mathbf{F}, \mathbf{G}) \mapsto ((f_i, g_j)_{\mathcal{Y}})_{i,j=1}^k \in \mathbb{R}^{k \times k}.$$

In particular, we have the following relation:

LEMMA 4.2. *For the operator \mathbf{U} (and similarly \mathbf{C}) holds*

$$(4.5) \quad \|\mathbf{U}z\|_{\mathcal{X}}^2 \leq \|(\mathbf{U}, \mathbf{U})\|_F \|z\|_{\mathbb{R}^k}^2$$

Proof. With $z = (z_1, \dots, z_k) \in \mathbb{R}^k$, $\mathbf{U}z = \sum_{i=1}^k z_i u_i \in \mathcal{X}$ follows

$$\begin{aligned} \|\mathbf{U}z\|_{\mathcal{X}}^2 &= \sum_{i=1}^k \sum_{j=1}^k z_i z_j \langle u_i, u_j \rangle_{\mathcal{X}} \\ &= \langle z, (\mathbf{U}, \mathbf{U})z \rangle_{\mathbb{R}^k} \leq \|(\mathbf{U}, \mathbf{U})z\|_{\mathbb{R}^k} \|z\|_{\mathbb{R}^k} \\ &\leq \|(\mathbf{U}, \mathbf{U})\|_F \|z\|_{\mathbb{R}^k}^2. \end{aligned}$$

□

From Lemma 4.2 follows in particular

$$(4.6) \quad \|\mathbf{U}\|_{\mathbb{R}^k \rightarrow \mathcal{X}} \leq \sqrt{\|(\mathbf{U}, \mathbf{U})\|_F}.$$

We will further on also need

LEMMA 4.3. *For $y \in \mathcal{Y}$ and \mathbf{C} as above holds*

$$(4.7) \quad \|(y, \mathbf{C})\|_{\mathbb{R}^k} \leq \|T\| \|y\|_{\mathcal{Y}} \sqrt{\sum_{l=1}^k \|u_l\|^2}$$

Proof. With (4.3) yields

$$(4.8) \quad \|(y, \mathbf{C})\|_{\mathbb{R}^k}^2 = \sum_{i=1}^k |(y, c_i)_{\mathcal{Y}}|^2 \leq \sum_{i=1}^k \|y\|_{\mathcal{Y}}^2 \|c_i\|_{\mathcal{Y}}^2$$

$$(4.9) \quad = \|y\|_{\mathcal{Y}}^2 \sum_{i=1}^k \|T u_i\|^2 \leq \|T\|_{\mathcal{X} \rightarrow \mathcal{Y}}^2 \|y\|_{\mathcal{Y}}^2 \sum_{i=1}^k \|u_i\|_{\mathcal{X}}^2$$

□

5. Projected Regularization Method. Let $\mathcal{U} \subset \mathcal{X}$ be a k -dimensional subspace, and $\mathcal{C} = T\mathcal{U} \subset \mathcal{Y}$ be its image under the action of T . Here it is important that we assume that the restricted operator $T|_{\mathcal{U}}$ is invertible, so as to guarantee that $\dim \mathcal{U} = \dim \mathcal{C} = k$ holds. This allows us express clearly the structure of projectors onto these spaces.

PROPOSITION 5.1. *Assume that the vectors $u_i \in \mathbf{U}$ are linear independent, and that*

$$(5.1) \quad T|_{\mathcal{U}} : \mathcal{U} \rightarrow \mathcal{Y}$$

is invertible on its range $\mathcal{R}(T|_{\mathcal{U}})$. Then $(\mathbf{C}, \mathbf{C}) \in \mathbb{R}^{k \times k}$ is continuously invertible.

Proof. According to (4.1), $\mathbf{C} = [c_1 \ c_2 \ \dots \ c_k] \in \mathcal{Y}^k$ with $c_i = T u_i$. As T is invertible on $\mathcal{R}(T|_{\mathcal{U}})$, the vectors in \mathbf{C} are linear independent if the vectors in \mathbf{U} are linear independent. Moreover,

$$(5.2) \quad (\mathbf{C}, \mathbf{C}) = ((c_i, c_j))_{i,j=1,\dots,k}$$

is the Gramian matrix of the linear independent system contained in \mathbf{C} and therefore an invertible matrix. Additionally, as $(\mathbf{C}, \mathbf{C}) : \mathbb{R}^k \rightarrow \mathbb{R}^k$ maps between finite dimensional spaces, its inverse is also bounded. □

In general, the operator T is ill posed and might have a non-trivial null space, i.e., T is not continuously invertible. However, on a finite dimensional subspace invertibility

might be achieved easily.

Now let Q be the \mathcal{Y} -orthogonal projector onto \mathcal{C} and P be the $(\cdot, T^*T\cdot)_{\mathcal{X}}$ -orthogonal projector onto \mathcal{U} . The orthogonal projector Q induces a direct sum splitting $\mathcal{Y} = \mathcal{C} \oplus \mathcal{C}^{\perp(\cdot, \cdot)_{\mathcal{Y}}}$. Consider $w \in \mathcal{Y}$. We represent the action of Q as $Qw = \sum_{i=1}^k (w, c_i)_{\mathcal{Y}} c_i$. Similarly, P induces a direct sum splitting of \mathcal{X} with respect to $(\cdot, T^*T\cdot)_{\mathcal{X}}$, namely $\mathcal{X} = \mathcal{U} \oplus \mathcal{U}^{\perp(\cdot, T^*T\cdot)_{\mathcal{X}}}$ where $\mathcal{U}^{\perp(\cdot, T^*T\cdot)_{\mathcal{X}}}$ denotes the orthogonal complement of \mathcal{U} with respect to the inner product $(\cdot, T^*T\cdot)_{\mathcal{X}}$. We represent the action of P as $Pv = \sum_{i=1}^k (v, T^*T u_i)_{\mathcal{X}} u_i$. We also can construct these projectors explicitly.

PROPOSITION 5.2. *Let $\mathbf{U} \in \mathcal{X}^k$ be a basis of \mathcal{U} and $\mathbf{C} = T\mathbf{U} \in \mathcal{Y}^k$ be a basis of \mathcal{C} . Then we have*

$$P\cdot = \mathbf{U}(\mathbf{U}, T^*T\mathbf{U})_{\mathcal{X}}^{-1}(\cdot, T^*T\mathbf{U})_{\mathcal{X}} \quad \text{and} \quad Q\cdot = \mathbf{C}(\mathbf{C}, \mathbf{C})_{\mathcal{Y}}^{-1}(\cdot, \mathbf{C})_{\mathcal{Y}}.$$

Proof. This is a standard result, particularly in the setting where $\mathcal{X} = \mathbb{R}^m$ and $\mathcal{Y} = \mathbb{R}^n$, but it is helpful to see why this is true in this more general setting, as well. From [Proposition 5.1](#), we know that $(\mathbf{C}, \mathbf{C}) = (T\mathbf{U}, T\mathbf{U}) = (\mathbf{U}, T^*T\mathbf{U})$ is continuously invertible. Let $\{\hat{u}_{k+1}, \hat{u}_{k+2}, \dots, \hat{u}_{\ell}, \dots\}$ be a basis for $\mathcal{U}^{\perp(\cdot, T^*T\cdot)_{\mathcal{X}}}$. Then for any $v \in \mathcal{X}$, we can write

$$v = \sum_{i=1}^k (v, T^*T u_i)_{\mathcal{X}} u_i + \sum_{i=k+1}^{\infty} (v, T^*T \hat{u}_i)_{\mathcal{X}} \hat{u}_i.$$

Since by definition $(\hat{u}_i, T^*T\mathbf{U})_{\mathcal{X}} = 0$ for all $i > k$, we can write

$$\mathbf{U}(\mathbf{U}, T^*T\mathbf{U})_{\mathcal{X}}^{-1}(v, T^*T\mathbf{U})_{\mathcal{X}} = \sum_{i=1}^k (v, T^*T u_i)_{\mathcal{X}} \mathbf{U}(\mathbf{U}, T^*T\mathbf{U})_{\mathcal{X}}^{-1}(u_i, T^*T\mathbf{U})_{\mathcal{X}}.$$

Observe that, according to the definitions in [\(4.3\)](#) and [\(4.4\)](#), the vector $(u_i, T^*T\mathbf{U})_{\mathcal{X}}$ is the i th column of the matrix $(\mathbf{U}, T^*T\mathbf{U})_{\mathcal{X}}$. Thus,

$$(\mathbf{U}, T^*T\mathbf{U})_{\mathcal{X}}^{-1}(u_i, T^*T\mathbf{U})_{\mathcal{X}} = \mathbf{e}_i \in \mathbb{R}^k$$

the i th Cartesian basis vector. Thus, we have

$$(5.4) \quad \mathbf{U}(\mathbf{U}, T^*T\mathbf{U})_{\mathcal{X}}^{-1}(u_i, T^*T\mathbf{U})_{\mathcal{X}} = \mathbf{U}\mathbf{e}_i = u_i$$

and thus according to [\(5.3\)](#)

$$(5.5) \quad \mathbf{U}(\mathbf{U}, T^*T\mathbf{U})_{\mathcal{X}}^{-1}(v, T^*T\mathbf{U})_{\mathcal{X}} = \sum_{i=1}^k (v, T^*T u_i)_{\mathcal{X}} u_i.$$

As v was chosen arbitrarily, this is true for any element of \mathcal{X} meaning the action of $\mathbf{U}(\mathbf{U}, T^*T\mathbf{U})_{\mathcal{X}}^{-1}(\cdot, T^*T\mathbf{U})_{\mathcal{X}}$ is the action of P . This completes the proof for P . The same line of reasoning yields the proof for Q . \square

We assume that with exact data that [\(1.1\)](#) is consistent. Using the projector P , we decompose the exact true solution,

$$x^{\dagger} = Px^{\dagger} + (I_{\mathcal{X}} - P)x^{\dagger}.$$

With exact data, Px^\dagger would be exactly computable, since we can rewrite

$$(5.6) \quad x_p^\dagger := Px^\dagger = \mathbf{U}(\mathbf{C}, \mathbf{C})_{\mathbf{y}}^{-1} (Tx, \mathbf{C})_{\mathbf{y}} = \mathbf{U}(\mathbf{C}, \mathbf{C})_{\mathbf{y}}^{-1} (y, \mathbf{C})_{\mathbf{y}}.$$

However, we do not have exact data. We have disturbed data y^δ such that $\|y - y^\delta\|_{\mathbf{y}} \leq \delta$. Rewriting

$$(5.7) \quad Px = \mathbf{U}(\mathbf{C}, \mathbf{C})_{\mathbf{y}}^{-1} (y - y^\delta, \mathbf{C})_{\mathbf{y}} + \mathbf{U}(\mathbf{C}, \mathbf{C})_{\mathbf{y}}^{-1} (y^\delta, \mathbf{C})_{\mathbf{y}}$$

we can define

$$(5.8) \quad x_p^\delta := \mathbf{U}(\mathbf{C}, \mathbf{C})_{\mathbf{y}}^{-1} (y^\delta, \mathbf{C})_{\mathbf{y}},$$

which can be computed inexpensively to high precision and may be used as an approximation $x_p^\delta \approx Px^\dagger$. We obtain

PROPOSITION 5.3. *The approximation error between x_p^\dagger and x_p^δ is given by*

$$(5.9) \quad x_p^\dagger - x_p^\delta = \mathbf{U}(\mathbf{C}, \mathbf{C})_{\mathbf{y}}^{-1} (y - y^\delta, \mathbf{C})_{\mathbf{y}},$$

which can be estimated as

$$(5.10) \quad \|x_p^\dagger - x_p^\delta\| \leq \left(\sqrt{\|(\mathbf{U}, \mathbf{U})\|_F \sum_{l=1}^k \|u_l\|^2} \|(\mathbf{C}, \mathbf{C})^{-1}\|_F \right) \delta.$$

Proof. Using (5.6)-(5.8) gives (5.9), which can be estimated as

$$(5.11)$$

$$(5.12) \quad \begin{aligned} \|x_p^\dagger - x_p^\delta\|_{\mathcal{X}} &\leq \left\| \mathbf{U}(\mathbf{C}, \mathbf{C})_{\mathbf{y}}^{-1} (y - y^\delta, \mathbf{C})_{\mathbf{y}} \right\| \\ &\leq \|\mathbf{U}\|_{\mathbb{R}^k \rightarrow \mathcal{X}} \left\| (\mathbf{C}, \mathbf{C})^{-1} \right\|_{\mathbb{R}^k \rightarrow \mathbb{R}^k} \left\| (y - y^\delta, \mathbf{C})_{\mathbf{y}} \right\|_{\mathbb{R}^k} \end{aligned}$$

$$(5.13) \quad \stackrel{(4.2), (4.7)}{\leq} \sqrt{\|(\mathbf{U}, \mathbf{U})\|_F} \left\| (\mathbf{C}, \mathbf{C})^{-1} \right\|_{\mathbb{R}^k \rightarrow \mathbb{R}^k} \|T\| \|y - y^\delta\|_{\mathbf{y}} \sqrt{\sum_{l=1}^k \|u_l\|^2}$$

$$(5.14) \quad = \left(\sqrt{\|(\mathbf{U}, \mathbf{U})\|_F \sum_{l=1}^k \|u_l\|^2} \left\| (\mathbf{C}, \mathbf{C})^{-1} \right\|_F \right) \delta \quad \square$$

Remark 5.4. Often with such recycling methods, one takes either \mathcal{U} or \mathcal{C} to be represented with an orthonormal basis, for the purposes of convenience with regard to analysis and implementation. We assume here that \mathcal{C} is represented by an orthonormal basis. If we scale \mathbf{U} such that $\mathbf{C} = T^* \mathbf{U}$ is an orthonormal system spanning \mathcal{C} , we can write (5.10) more compactly as

$$\|Px - x_p^\delta\|_{\mathcal{X}} \leq \|\mathbf{U}\|_{\mathcal{X}} \delta = \mathcal{O}(\delta).$$

Furthermore, P and Q can be expressed more compactly,

$$P \cdot = \mathbf{U}(\cdot, T^* T \mathbf{U})_{\mathcal{X}} \quad \text{and} \quad Q \cdot = \mathbf{C}(\cdot, \mathbf{C})_{\mathbf{y}},$$

due to the fact that $(\mathbf{U}, T^* T \mathbf{U})_{\mathcal{X}} = (\mathbf{C}, \mathbf{C})_{\mathbf{y}} = \mathbf{I}$ when \mathbf{C} is an orthonormal system.

It remains to find a suitable approximation for $(I_{\mathcal{X}} - P)x^\dagger$. We have the following result:

PROPOSITION 5.5. *Under the assumption*

$$(5.15) \quad QT(I - P)x^\dagger = 0$$

holds

$$(5.16) \quad y - y_p = (I - Q)T(I - P)x^\dagger,$$

where

$$(5.17) \quad y_p := Tx_p^\dagger.$$

Proof. We have

$$(5.18) \quad y = Tx^\dagger = QTPx^\dagger + QT(I - P)x^\dagger + (I - Q)TPx^\dagger + (I - Q)T(I - P)x^\dagger.$$

As $x_p^\delta = Px^\dagger \in \mathcal{U}$ we have $y_p = Tx_p^\delta \in T\mathcal{U}$ and therefore

$$(5.19) \quad y_p = Tx_p^\delta = QTPx^\dagger.$$

Further on, as $TPx^\dagger \in T\mathcal{U}$ and $(I - Q)$ projects onto the complement of $T\mathcal{U}$ we obtain

$$(5.20) \quad (I - Q)TPx^\dagger = 0.$$

Inserting (5.19), (5.20) into (5.18) yields

$$(5.21) \quad y - y_p = QT(I - P)x^\dagger + (I - Q)T(I - P)x^\dagger,$$

and Assumption (5.15) yields (5.16) □

Assumption (5.15) transforms to

$$(5.22) \quad \begin{aligned} QT(I - P)x^\dagger = 0 &\iff \langle z, QT(I - P)x^\dagger \rangle = 0 & \forall z \in \mathcal{Y} \\ &\iff \langle T^*Qz, (I - P)x^\dagger \rangle = 0 & \forall z \in \mathcal{Y} \end{aligned}$$

We arrive at

PROPOSITION 5.6. *If*

$$(5.23) \quad T^*\mathcal{C} \perp (I - P)x^\dagger,$$

then (5.15) holds. Particularly, (5.23) is fulfilled if

$$(5.24) \quad T^*T\mathcal{U} \subset \mathcal{U}$$

holds.

Proof. As $Q\mathcal{Y} = \mathcal{C}$, eq. (5.22) transforms to (5.23). As $(I - P)x^\dagger \in \mathcal{U}^\perp$ holds always if we additionally assume that \mathcal{U} is chosen s.t. $T^*\mathcal{C} \perp \mathcal{U}^\perp$, and with $\mathcal{C} = T\mathcal{U}$ this transforms to $T^*T\mathcal{U} \perp \mathcal{U}^\perp$ or, equivalently, (5.24). □

Remark 5.7. Obviously, (5.24) is much stronger than (5.23). However, there are easy examples where (5.23) always holds:

1. If $\mathcal{U} = \text{span}\{u_i : i = 1, \dots, N\}$ and $T^*Tu_i = \sigma_i u_i$ then clearly $T^*T\mathcal{U} \subset \mathcal{U}$.

2. If T is an *unitary* operator, i.e., $\langle Tx_1, Tx_2 \rangle = \langle x_1, x_2 \rangle$ holds for all x_1, x_2 , then we have for $u \in \mathcal{U}$ and $u^\perp \in \mathcal{U}^\perp$

$$\langle T^*Tu, u^\perp \rangle = \langle Tu, Tu^\perp \rangle = \langle u, u^\perp \rangle = 0,$$

i.e., (5.24).

Furthermore, the way we construct our projectors ensures that the assumption holds.

LEMMA 5.8. *The assumption (5.15) holds for any pair of projectors (P, Q) satisfying (3.3).*

Proof. This follows from the fact that the product of complementary projectors is the zero operator, since from (3.3), we have

$$QT(I - P)x^\dagger = Q(I - Q)Tx^\dagger = 0.$$

□

With the above results we can now compute $(I - P)x^\dagger$ as a solution of equation (5.15). However, as always in Inverse Problems, we might not have the exact right hand side $y - y_p$ but some noisy version $y^\delta - y_p^\delta$, where we define $y_p^\delta := Tx_p^\delta$. The data error can be approximated as follows:

PROPOSITION 5.9. *The data error in the left hand side of (5.15) can be estimated as*

$$(5.25) \quad \|(y - y_p) - (y^\delta - y_p^\delta)\| \leq \kappa_{\mathcal{U}} \cdot \delta$$

with

$$(5.26) \quad \kappa_{\mathcal{U}} := 1 + \|T\| \sqrt{\|(\mathbf{U}, \mathbf{U})\|_F \sum_{l=1}^k \|u_l\|^2} \|(\mathbf{C}, \mathbf{C})^{-1}\|_F$$

Proof. The proof is straight forward:

$$(5.27) \quad \|(y - y_p) - (y^\delta - y_p^\delta)\| \leq \|y - y^\delta\| + \|y_p - y_p^\delta\| \leq \delta + \|T\| \|x_p^\dagger - x_p^\delta\|$$

and with (5.10) follows (5.25) □

We are now ready to formulate our solution approach:

Algorithm 1: Augmented Regularization

- 1 Given:
 - data y^δ fulfilling $\|y - y^\delta\| \leq \delta$
 - $\mathcal{U} \subset \mathcal{X}$, \mathcal{U} finite dimensional
 - a regularization method $R_\alpha(y, T)$ for the equation $y = Tx$ with parameter choice rule $\alpha = \alpha(\delta)$
 - 2 Compute x_p^δ according to (5.8) and set $y_p^\delta = Tx_p^\delta$
 - 3 Set $B := (I - Q)T$
 - 4 Compute $\hat{x}_p^\delta := R_\alpha(y^\delta - y_p^\delta, B)$, $\alpha = \alpha(\kappa_{\mathcal{U}}\delta)$
 - 5 Set $x_{\mathcal{U}, \alpha}^\delta := x_p^\delta + \hat{x}_p^\delta$
-

PROPOSITION 5.10. *If for the solution x^\dagger of the equation $Tx = y$ and the chosen augmentation subspace \mathcal{U} eq. (5.15) holds, then the augmented regularization with data y^δ fulfilling $\|y - y^\delta\| \leq \delta$ as described in Algorithm 1 forms a regularization method.*

Proof. As $R_\alpha(y^\delta - y_p^\delta, B)$ is a regularization for (5.16) we have as $\delta \rightarrow 0$, and decomposing $x^\dagger = Px^\dagger + (I - P)x^\dagger = x_p^\dagger + \hat{x}_p^\dagger$ yields

$$\begin{aligned} \|x^\dagger - x_{\mathcal{U},\alpha}^\delta\| &\leq \|x_p^\dagger - x_p^\delta\| + \|\hat{x}_p^\dagger - \hat{x}_p^\delta\| \\ &\stackrel{(5.10)}{\leq} \mathcal{O}(\delta) + \|\hat{x}_p^\dagger - R_{\alpha(\kappa_{\mathcal{U}}\delta)}(y^\delta - y_p^\delta, B)\| \xrightarrow{\delta \rightarrow 0} 0 \end{aligned}$$

which concludes the proof. \square

Remark 5.11. We can thus extend any existing regularization R_α to its augmented method and still obtain a regularization method.

EXAMPLE 1. *Let us illustrate this for the Landweber method, which has been defined in (2.6). If the iteration is stopped by the discrepancy principle, i.e., if the stopping index k_* is determined as the first index such that*

$$(5.28) \quad \|y^\delta - Tx_{k_*}^\delta\| \leq \tau\delta \quad \tau > 1$$

holds, then it is a regularization method, see e.g., [8]. For given data y^δ , the augmented Landweber method would read as

$$(5.29) \quad x_{k_*}^\delta = x_p^\delta + \hat{x}_{k_*}^\delta$$

where x_p^δ is defined in (5.8) and $\hat{x}_{k_}^\delta$ is the k_* th iterate of the Landweber iteration applied to the equation*

$$(5.30) \quad (I - Q)Tx = y^\delta - y_p^\delta, \quad y_p^\delta = Tx_p^\delta.$$

Using

$$(5.31) \quad y_p^\delta = Tx_p^\delta = T\mathbf{U}(\mathbf{C}, \mathbf{C})_{\mathbf{y}}^{-1}(y^\delta, \mathbf{C})_{\mathbf{y}} = \mathbf{C}(\mathbf{C}, \mathbf{C})_{\mathbf{y}}^{-1}(y^\delta, \mathbf{C})_{\mathbf{y}} = Qy^\delta$$

the iteration for \hat{x}_k^δ reads as

$$\begin{aligned} \hat{x}_{k+1}^\delta &= \hat{x}_k^\delta + T^*(I - Q)(y^\delta - y_p^\delta - (I - Q)T\hat{x}_k^\delta) \\ (5.32) \quad &= \hat{x}_k^\delta + T^*(I - Q)(y^\delta - T\hat{x}_k^\delta). \end{aligned}$$

Usually such an iteration is started with $\hat{x}_0^\delta = 0$. Instead, we can incorporate x_p^δ directly into the iteration by setting $\hat{x}_0^\delta = x_p^\delta$.

5.1. What to recycle. One question we have not addressed thus far concerns what should the subspace \mathcal{U} encode? This is a question that hinges very much on the application. Suffice it to say, a detailed review of recycling strategies is beyond the scope of this paper, but we refer the reader to [36, Section 6] for a more detailed discussion. Some common choices for augmentation vectors are approximate solutions or solutions to related problems, approximate or (when available) exact eigenvectors/singular vectors, and vectors satisfying convergence model optimality properties.

The use of approximate eigenvectors as a recycling strategy was first suggested for recycling between cycles of GMRES applied to one linear system to mitigate the effects on convergence speed of restarting [28]. The algorithmic choices in that paper necessitated that the approximate eigenvectors be harmonic Ritz vectors. Ritz-type eigenvector approximations with respect to a square matrix \mathbf{A} are obtained applying

a Galerkin or Petrov-Galerkin approximation strategy to the eigenvalue problem of the form

$$\text{select } \mathbf{v} \in \mathcal{V} \quad \text{such that} \quad \mathbf{A}\mathbf{v} - \lambda\mathbf{v} \perp \mathcal{W},$$

for some $\lambda \in \mathbb{R}$. The spaces \mathcal{V} and \mathcal{W} are usually spaces which have been generated during the iteration, so that the solution to the Ritz problem reduces to a (generalized) eigenvalue problem. This strategy was extended and combined with GCRO-type optimal methods in a way that did not require augmentation by harmonic Ritz vectors, and this enabled recycling between multiple linear systems [29]. This has been further extended to the case of short-recurrence schemes not requiring restart; see, e.g., [5, 23, 38]. In that case, one stores and updates a running recycle *window* but not to use in the current iteration. Rather, it is being built to be used in subsequent systems.

It should be noted that in the context of discrete ill-posed problems, a poor or noisy choice of \mathcal{U} may induce a projected problem (3.4) which is even more ill-posed. However, it has been demonstrated that there are many good choices for \mathcal{U} ; see, e.g., [3, 4, 9]. More recently, a recycled LSQR algorithm has been proposed where approximate eigenvectors are used as well as other data which become available in real-time [5]. As LSQR is a short-recurrence method, the authors employ a windowed recycling strategy, where the recycled subspace is updated every p iterations so that only a window of vectors generated by the short recurrence iteration must be stored. They propose a few different recycling approaches appropriate for ill-posed problems: truncated singular vectors, solution- or sparsity-oriented compression, and reduced basis decomposition.

In the next section, we use the recycling method framework from [36] to develop augmented versions of steepest descent and Landweber methods. In Section 7, we combine this method with recycling strategies appropriate to each individual application, building on strategies from the literature but adjusted to fit the particular applications.

6. Augmented gradient descent methods. We now delve further into the case that we apply a gradient descent-type method (i.e., Landweber) to the projected problem, in order to produce a practical implementation of this algorithm. In particular, we first explore augmenting the steepest descent method, which can be understood as a Landweber-type method with a step-length that is dynamically chosen to minimize the residual norm. We have shown in Proposition 3.1 that this is a projection method and thus can be augmented in the above-discussed framework. Augmented methods such as this rely on certain identities to achieve algorithmic advantages.

LEMMA 6.1. *The residual produced by the augmented iterative method with approximation $\hat{x}_k^\delta = \mathbf{U}(\mathbf{C}, \mathbf{C})_{\mathcal{Y}}^{-1} (y^\delta, \mathbf{C})_{\mathcal{Y}} + (I_{\mathcal{X}} - P) \hat{t}_k^\delta$ defined as in (3.5) where \hat{t}_k^δ is the approximate solution of (3.4) coincides with the residual associated to \hat{t}_k^δ for the projected problem (3.4), i.e.,*

$$y^\delta - T\hat{x}_k^\delta = (I_{\mathcal{Y}} - Q) (y^\delta - T\hat{t}_k^\delta) \perp \mathcal{C}.$$

Proof. This result has been shown in a number of papers for finite-dimensional, well-posed problems; see, [7] and [6, 15, 16, 18]. We show it here for completeness by

direct calculation, by first writing

$$\begin{aligned} y^\delta - T\hat{x}_j^\delta &= y^\delta - T \left(\mathbf{U}(\mathbf{C}, \mathbf{C})_y^{-1} (y^\delta, \mathbf{C})_y + (I_{\mathcal{X}} - P) \hat{t}_j^\delta \right) \\ y^\delta - T\hat{x}_j^\delta &= y^\delta - \left(\mathbf{C}(\mathbf{C}, \mathbf{C})_y^{-1} (y^\delta, \mathbf{C})_y + T(I_{\mathcal{X}} - P) \hat{t}_j^\delta \right). \end{aligned}$$

Then we observe that it follows from [Proposition 5.2](#) and [\(3.3\)](#) that this is equivalent to

$$\begin{aligned} y^\delta - T\hat{x}_j^\delta &= y^\delta - (Qy^\delta + (I_{\mathcal{X}} - Q)T\hat{t}_j^\delta) \\ &= (I_{\mathcal{X}} - Q)y^\delta - (I_{\mathcal{X}} - Q)T\hat{t}_j^\delta. \end{aligned}$$

It follows directly that $y^\delta - T\hat{x}_j^\delta \perp \mathcal{C}$. \square

Note that this proof holds in the presence or absence of noise and whether or not the problem being treated is consistent. In the case that the right-hand side is noise-free and the problem is consistent, we have from [\(5.7\)](#) that

$$(6.1) \quad \mathbf{U}(\mathbf{C}, \mathbf{C})_y^{-1} (y, \mathbf{C})_y = Px^\dagger = x_p^\dagger,$$

meaning it is the best approximation of x^\dagger in \mathcal{U} with respect to the T^*T -norm. In the presence of noise, we previously characterized and bounded the error introduced into [\(6.1\)](#) in [Proposition 5.3](#).

It directly follows that the construction of the gradient descent direction for Landweber applied to [\(3.4\)](#) can be simplified.

COROLLARY 6.2. *The gradient descent direction for Landweber applied to [\(3.4\)](#) admits the simplification,*

$$T^*(I_{\mathcal{Y}} - Q)(y^\delta - T\hat{t}_j^\delta) = T^*(y^\delta - T\hat{x}_j^\delta);$$

i.e., the descent direction is the same as that for the full problem [\(1.1\)](#).

With these simplifications, we now show that indeed the above-described fully augmented Landweber construction is indeed updating the approximation to the solution to the full problem [\(1.1\)](#) in the optimal direction at each iteration. Furthermore, if when applying Landweber to [\(3.4\)](#), we choose the step length dynamically such that we are applying the method of *steepest descent* to [\(3.4\)](#), this will be shown in the following theorem to be equivalent to minimizing the residual over the sum of subspaces $\mathcal{S}_{j+1} = \mathcal{U} + \text{span}\{T^*(y^\delta - Tx_j)\}$.

As this should be a steepest descent-type method which is a \mathcal{Y} -norm minimization method, the correct constraint space is the image of the correction space under the action of the operator, namely $\mathcal{L}_{j+1} = T\mathcal{U} + \text{span}\{TT^*(y^\delta - Tx_j)\}$. At step $j+1$, we have $x_{j+1} = x_j + \mathbf{U}\mathbf{u}_{j+1} + \alpha_{j+1}T^*(y^\delta - Tx_j)$ where $\mathbf{u}_{j+1} \in \mathbb{R}^k$ and $\alpha_{j+1} \in \mathbb{R}$ are determined by the minimization constraint. The following is an adaption of a more general result, proven in [\[7\]](#).

THEOREM 6.3. *The $(j+1)$ st approximate solution to the augmented steepest descent problem minimization,*

$$\begin{aligned} &\text{select } \mathbf{U}\mathbf{u}_{j+1} + \alpha_{j+1}v_j \in \mathcal{U} + \text{span}\{v_j\} \quad \text{where } v_j = T^*(y^\delta - Tx_j) \\ &\text{such that } (\mathbf{u}_{j+1}, \alpha_{j+1}) = \underset{\substack{\mathbf{u} \in \mathbb{R}^k \\ \alpha \in \mathbb{R}}}{\operatorname{argmin}} \|y - T(x_j + \alpha v_j + \mathbf{U}\mathbf{u})\|_{\mathcal{Y}} \end{aligned}$$

satisfies

1. $\mathbf{u}_{j+1} = \mathbf{u}_{j+1}^{(1)} + \mathbf{u}_{j+1}^{(2)}$ with

$$(6.2) \quad \mathbf{U}\mathbf{u}_{j+1}^{(1)} = \begin{cases} \operatorname{argmin}_{x \in \mathcal{U}} \|e_0^\dagger - x\|_{(T^*T, \cdot)_{\mathcal{X}}} - u_y & \text{if } j = 0 \\ \mathbf{0} & \text{otherwise} \end{cases}$$

$$(6.3) \quad \mathbf{U}\mathbf{u}_{j+1}^{(2)} = -\operatorname{argmin}_{x \in \mathcal{U}} \|x - \alpha_{j+1}v_{j+1}\|_{(T^*T, \cdot)_{\mathcal{X}}}$$

where $u_y = \mathbf{U}(\mathbf{C}, \mathbf{C})_{\mathcal{Y}}^{-1} (y - y^\delta, \mathbf{C})_{\mathcal{Y}}$,

2. and α_{j+1} satisfies

$$(6.4) \quad \alpha_{j+1}v_{j+1} = \operatorname{argmin}_{x \in \operatorname{span}\{v_j\}} \|(I_{\mathcal{Y}} - Q)r_j - (I_{\mathcal{Y}} - Q)Tx\|_{\mathcal{Y}},$$

i.e., α_{j+1} is obtained via applying an iteration of steepest descent to the projected problem, where $v_{j+1} = T^*r_j$ and $e_0^\dagger = x^\dagger - x_0$.

Proof. As this is a minimization problem, we prove the result by studying the structure of the associated orthogonal projector. We can rewrite this minimization as

$$(\mathbf{u}_{j+1}, \alpha_{j+1}) = \operatorname{argmin}_{\substack{\mathbf{u} \in \mathbb{R}^k \\ \alpha \in \mathbb{R}}} \|r_j - T(\alpha v_j + \mathbf{U}\mathbf{u})\|_{\mathcal{Y}}.$$

Thus, we seek the pair $(\mathbf{u}_{j+1}, \alpha_{j+1})$ giving the \mathcal{Y} -norm best approximation of r_j in the space $T\mathcal{S}_{j+1}$. This is equivalent to computing the \mathcal{Y} -orthogonal projection of r_j into that space, which we denote $Q_{T\mathcal{S}_{j+1}}$. Recalling that $\mathbf{C} = T\mathbf{U}$, from [Proposition 5.2](#), we can express this as

$$\begin{aligned} T(\alpha v_j + \mathbf{U}\mathbf{u}) &= Q_{T\mathcal{S}_{j+1}}r_j \\ &= [\mathbf{C} \quad Tv_{j+1}] ([\mathbf{C} \quad Tv_{j+1}], [\mathbf{C} \quad Tv_{j+1}])_{\mathcal{Y}}^{-1} (r_j, [\mathbf{C} \quad Tv_{j+1}])_{\mathcal{Y}}, \end{aligned}$$

where $[\mathbf{C} \quad Tv_{j+1}] \in \mathcal{L}(\mathbb{R}^{(k+1)}, \mathcal{Y})$. This means that $(\mathbf{u}_{j+1}, \alpha_{j+1})$ form the solution of

$$\begin{bmatrix} (\mathbf{C}, \mathbf{C})_{\mathcal{Y}} & (Tv_{j+1}, \mathbf{C})_{\mathcal{Y}} \\ (\mathbf{C}, Tv_{j+1})_{\mathcal{Y}} & (Tv_{j+1}, Tv_{j+1})_{\mathcal{Y}} \end{bmatrix} \begin{bmatrix} \mathbf{u}_{j+1} \\ \alpha_{j+1} \end{bmatrix} = \begin{bmatrix} (r_j, \mathbf{C})_{\mathcal{Y}} \\ (r_j, Tv_{j+1})_{\mathcal{Y}} \end{bmatrix}$$

To prove the statement of the theorem, we eliminate \mathbf{u}_{j+1} from the second block of equations (following [\[30\]](#)) yielding the two sets of equations,

$$\begin{aligned} (\mathbf{C}, \mathbf{C})_{\mathcal{Y}} \mathbf{u}_{j+1} + \alpha_{j+1} (Tv_{j+1}, \mathbf{C})_{\mathcal{Y}} &= (r_j, \mathbf{C})_{\mathcal{Y}} \\ \alpha_{j+1} \left[(Tv_{j+1}, Tv_{j+1})_{\mathcal{Y}} - (\mathbf{C}, Tv_{j+1})_{\mathcal{Y}} (\mathbf{C}, \mathbf{C})_{\mathcal{Y}}^{-1} (Tv_{j+1}, \mathbf{C})_{\mathcal{Y}} \right] &= (r_j, Tv_{j+1})_{\mathcal{Y}} \\ &\quad - (\mathbf{C}, Tv_{j+1})_{\mathcal{Y}} (\mathbf{C}, \mathbf{C})_{\mathcal{Y}}^{-1} (r_j, \mathbf{C})_{\mathcal{Y}}. \end{aligned}$$

We can consolidate the second set of equations, yielding

$$\begin{aligned} \alpha_{j+1} \left(Tv_{j+1} - \mathbf{C} (\mathbf{C}, \mathbf{C})_{\mathcal{Y}}^{-1} (Tv_{j+1}, \mathbf{C})_{\mathcal{Y}}, Tv_{j+1} \right)_{\mathcal{Y}} &= \\ \left(r_j - \mathbf{C} (\mathbf{C}, \mathbf{C})_{\mathcal{Y}}^{-1} (r_j, \mathbf{C})_{\mathcal{Y}}, Tv_{j+1} \right)_{\mathcal{Y}}. \end{aligned}$$

Recalling the definition of the projector Q , we substitute and then take advantage of the self-adjointness of orthogonal projectors,

$$\begin{aligned}\alpha_{j+1} ((I_Y - Q) T v_{j+1}, T v_{j+1})_{\mathcal{Y}} &= ((I_Y - Q) r_j, T v_{j+1})_{\mathcal{Y}} \\ \alpha_{j+1} \left((I_Y - Q)^2 T v_{j+1}, T v_{j+1} \right)_{\mathcal{Y}} &= \left((I_Y - Q)^2 r_j, T v_{j+1} \right)_{\mathcal{Y}} \\ \alpha_{j+1} ((I_Y - Q) T v_{j+1}, (I_Y - Q) T v_{j+1})_{\mathcal{Y}} &= ((I_Y - Q) r_j, (I_Y - Q) T v_{j+1})_{\mathcal{Y}}.\end{aligned}$$

Thus, the minimizer α_{j+1} satisfies

$$\left((I_Y - Q) r_j - (I_Y - Q) T (\alpha_{j+1} v_{j+1}), (I_Y - Q) T v_{j+1} \right)_{\mathcal{Y}} = 0.$$

As $v_{j+1} = T^* r_j$, then we see that α_{j+1} is the solution of

$$\begin{aligned}\text{select } \alpha_{j+1} v_{j+1} &\in \text{span} \{ T^* (I_Y - Q) r_j \} \\ \text{such that } (I_Y - Q) r_j - (I_Y - Q) T (\alpha_{j+1} v_{j+1}) &\perp \text{span} \{ (I_Y - Q) T T^* (I_Y - Q) r_j \}.\end{aligned}$$

From [Proposition 3.1](#), we then see that this is the residual constraint formulation of the $(j+1)$ st step of steepest descent applied to [\(6.4\)](#). We can then solve for \mathbf{u}_{j+1} ,

$$\mathbf{u}_{j+1} = (\mathbf{C}, \mathbf{C})_{\mathcal{Y}}^{-1} [(r_j, \mathbf{C})_{\mathcal{Y}} - \alpha_{j+1} (T v_{j+1}, \mathbf{C})_{\mathcal{Y}}].$$

Expanding to obtain the contribution from \mathcal{U} yields

$$\mathbf{U} \mathbf{u}_{j+1} = \mathbf{U} (\mathbf{C}, \mathbf{C})_{\mathcal{Y}}^{-1} (r_j, \mathbf{C})_{\mathcal{Y}} - \alpha_{j+1} \mathbf{U} (\mathbf{C}, \mathbf{C})_{\mathcal{Y}}^{-1} (T v_{j+1}, \mathbf{C})_{\mathcal{Y}},$$

and we observe that second term in this sum is [\(6.3\)](#). For the first term, we must consider two cases. If $j = 0$, then $x_0 = 0$, and $r_0^\delta = y^\delta$. Thus, we can write

$$\mathbf{U} (\mathbf{C}, \mathbf{C})_{\mathcal{Y}}^{-1} (r_0, \mathbf{C})_{\mathcal{Y}} = \mathbf{U} (\mathbf{C}, \mathbf{C})_{\mathcal{Y}}^{-1} (y^\delta, \mathbf{C})_{\mathcal{Y}} = x_p^\delta.$$

From [Proposition 5.3](#), we have that $x_p^\delta = x_p^\dagger - \mathbf{U} (\mathbf{C}, \mathbf{C})_{\mathcal{Y}}^{-1} (y - y^\delta, \mathbf{C})_{\mathcal{Y}}$. It is established in [\(6.1\)](#) that x_p^\dagger is the minimizer in [\(6.2\)](#). Lastly, if $j > 0$, [Lemma 6.1](#) tells us that $r_j \perp \mathcal{C}$; thus $(r_j, \mathbf{C})_{\mathcal{Y}} = \mathbf{0}$. \square

COROLLARY 6.4. *The full augmented steepest descent approximation can be represented as*

$$x_{j+1} = \begin{cases} x_p^\delta + \alpha_1 T^* y^\delta - \alpha_i P(T^* y^\delta) & \text{if } j = 0 \\ x_j + \alpha_i T^* r_j - \alpha_i P(T^* r_j) & \text{otherwise} \end{cases}$$

with optimal step length

$$\alpha_i = \frac{\|T^* r_i\|_{\mathcal{Y}}^2}{\|(I_Y - Q) T T^* r_i\|_{\mathcal{Y}}^2}.$$

We note that $x_p^\delta \in \mathcal{U}$ can be computed one time initially, as it comes only from initial data. We encapsulate all of this in the case of augmented steepest descent with dynamically determined optimal step length as [Algorithm 2](#).

Note: the reader may notice that Line 3 of [Algorithm 2](#) asks that the user compute a QR-factorization of $T\mathbf{U}$, which essentially has a finite number of elements from

an infinite-dimensional space \mathcal{Y} as its “columns”. By QR-factorization in this context, we mean performing steps of the Gram-Schmidt process and storing the orthogonalization and normalization coefficients in an upper-triangular matrix \mathbf{R} . Right-multiplication of $\mathbf{C} \in \mathcal{L}(\mathbb{R}^k, \mathcal{Y})$ by a matrix should be understood as applying the definition in (4.2) for each column of \mathbf{R} , thereby recovering the relationship between $T\mathbf{U}$ and \mathbf{C} arising from the Gram-Schmidt process. This is a well-defined procedure and process which is discussed, e.g., in [37, Lecture 7].

Algorithm 2: Augmented steepest descent for the normal equations

```

1 Given:  $\mathbf{U} \in \mathcal{L}(\mathbb{R}^k, \mathcal{X})$  representing  $\mathcal{U}$ 
2 Set  $r_0 = y - Tx_0$ 
3 Compute QR-factorization  $T\mathbf{U} = \mathbf{C}\mathbf{R}$ 
4  $\mathbf{U} \leftarrow \mathbf{U}\mathbf{R}^{-1}$ 
5  $\mathbf{u}^{(1)} = (r_0, \mathbf{C})_{\mathcal{Y}}$ 
6  $x \leftarrow x_0 + \mathbf{U}\mathbf{u}^{(1)}$ 
7  $r \leftarrow r_0 - \mathbf{C}\mathbf{u}^{(1)}$ 
8 while STOPPING-CRITERIA do
9    $\alpha_i = \frac{\|T^*r_i\|_{\mathcal{Y}}^2}{\|(I_{\mathcal{Y}} - Q)TT^*r_i\|_{\mathcal{Y}}^2}$ 
10   $\hat{\mathbf{w}} \leftarrow (TT^*r, \mathbf{C})_{\mathcal{Y}}$ 
11   $x \leftarrow x + \alpha_i T^*r - \alpha_i \mathbf{U}\hat{\mathbf{w}}$ 
12   $r \leftarrow r - \alpha_i TT^*r + \alpha_i \mathbf{C}\hat{\mathbf{w}}$ 
13 end
```

Furthermore, it is then clear that if we no longer choose α_i dynamically according to a minimization criteria and instead fix $\alpha_i \equiv \alpha$ that we can propose an augmented Landweber method, which we encapsulate as Algorithm 3.

Algorithm 3: Augmented Landweber for the normal equations

```

1 Given:  $\mathbf{U} \in \mathcal{L}(\mathbb{R}^k, \mathcal{X})$  representing  $\mathcal{U}$ ,  $\alpha > 0$ 
2 Set  $r_0 = y - Tx_0$ 
3 Compute QR-factorization  $T\mathbf{U} = \mathbf{C}\mathbf{R}$ 
4  $\mathbf{U} \leftarrow \mathbf{U}\mathbf{R}^{-1}$ 
5  $\mathbf{u}^{(1)} = (r_0, \mathbf{C})_{\mathcal{Y}}$ 
6  $x \leftarrow x_0 + \mathbf{U}\mathbf{u}^{(1)}$ 
7  $r \leftarrow r_0 - \mathbf{C}\mathbf{u}^{(1)}$ 
8 while STOPPING-CRITERIA do
9    $\hat{\mathbf{w}} \leftarrow (TT^*r, \mathbf{C})_{\mathcal{Y}}$ 
10   $x \leftarrow x + \alpha T^*r - \alpha \mathbf{U}\hat{\mathbf{w}}$ 
11   $r \leftarrow r - \alpha TT^*r + \alpha \mathbf{C}\hat{\mathbf{w}}$ 
12 end
```

7. Numerical experiments. We present here experiments from two subproblems from the field of adaptive optics for large ground-based telescopes as well as for a toy blurring problem presented as a proof-of-concept for an alternative use of these methods.




FIG. 1. *Residual and error convergence for large-scale Gaussian blur problem with $\sigma = 6$ and augmentation space generated using vectors resulting from applying 2 iterations steepest descent to the noisy (dashed curve) and noise-free (solid black curve) data using Gaussian blur operators for five smaller standard deviations chosen equally-spaced from the interval $[0.5, 1.5]$.*

7.1. Gaussian blurring model. Our first experiment concerns an academic problem which serves as a proof of concept. However, we still execute a large-scale version of this experiment using an implicitly-defined, matrix free problem generated by the IR-Tools software package [17]. The IR-tools package does provide an out-of-the-box generator function for the Gaussian blurring problem, `PRBlurGauss()`; however, it does not allow one to specify precise standard deviation (i.e., the spread) of the point-spread function (PSF). Looking under the hood, though, one sees that there is a helper function called `psfGauss()` which generates a PSF function for a Gaussian with specified spread, which can then be fed into the general purpose `PRBlur()` problem generator, which outputs a matrix-free operator which convolves images with the PSF as well as the true solution and noisy right-hand side. A variety of solution images are available; we choose a predefined geometric pattern for this experiment. To make the problem sufficiently large-scale, we choose the image size to be 500×500 , meaning the implicit side of the matrix induced by the PSF is 250000×250000 . The IR-tools package is careful to not generate exact right-hand side data unless it is specified. We rerun `PRBlur()` with the option `CommitCrime='true'` to obtain \mathbf{b}_{true} , but this is only used to understand the effects of noise in generating the recycling vectors. In these experiments, we generated the problem for standard deviation $\sigma = 6$. To construct an augmentation subspace \mathcal{U} , we selected vectors resulting from applying two iterations of steepest descent to the right-hand side data but for different Gaussian blurring operators with five smaller standard deviations chosen equally-spaced from the interval $[0.5, 1.5]$. An orthonormal basis for the span of these (nine) vectors was obtained and taken as \mathcal{U} . We also generated a set of “clean” recycled vectors using the noise-free \mathbf{b}_{true} , to see what effect noise has on the recycling process itself. It is demonstrated that such vectors do indeed offer better performance as recycled vectors, meaning that noise does effect the recycling process.

In Figure 1, we see that for both problems, this strategy yields an improvement in convergence, both for the noise-free and noise-contaminated problems. This is a promising numerical result and a proof of concept, but it requires further analysis and experimentation to understand it in the proper concept to guide the use of this strategy for real-world, non-academic problems.

7.2. Adaptive optics: wavefront reconstruction. Turbulences in the atmosphere have a significant impact on the imaging quality of modern ground based

FIG. 2. *The right-hand side and the best reconstruction produced by steepest descent for each, respectively.*





FIG. 3. *The true image and reconstructed images using augmented steepest descent. We show them both as images in the top row and three-dimensional surfaces in the bottom row.*



telescopes. In order to correct for the impact of the atmosphere, *Adaptive Optics* (AO) systems are utilized. In these systems, the incoming light from bright guide stars is measured by wavefront sensors in order to obtain information of the actual turbulence in the atmosphere. In a next step, deformable mirrors use this information for an improvement of the obtained scientific images. For a review on the different mathematical aspects of AO systems we refer to [11]. E.g., *Single Conjugate Adaptive Optics* (SCAO) systems use one sensor and one mirror for the correction and are able to obtain a good image quality close to the direction of the guide star. Figure 4 shows a setup for a SCAO system. For the instruments of the new generation of telescopes, e.g., the Extremely Large Telescope (ELT) of the European Southern Observatory (ESO), the Pyramid sensor (P-WFS) will be used frequently to obtain the wavefront. A linear approximation of the connection between the incoming wavefront φ and the sensor measurements (s_x, s_y) - assuming a non-modulated sensor - can be described as

$$(7.1) \quad s_x(x, y) = \frac{1}{\pi} \int_{-B(y)}^{B(y)} \frac{\Phi(x', y)}{x - x'} dx' ,$$

analogous for s_y . An overview of existing reconstruction methods for the P-WFS can be found in [22, 20, 19], and a throughout analysis of the various P-WFS models is given in [21, 21].




FIG. 4. Principle of SCAO

The linear models can be used in case of small wavefront aberrations which is the case in a *closed loop* setting, i.e., if the wavefront sensor is optically located *behind* the deformable mirror. Due to the rapidly changing atmosphere, the shape of the deformable mirror has to be readjusted every 1-2 milliseconds over the whole imaging process, which might last several minutes. Therefore, the reconstructions of the wavefronts have to be accurate and have to be obtained in real-time. As a quality measure we use the Strehl ratio (SR), which relates the imaging quality of the

telescope without disturbing atmosphere ($SR = 1$) to the corrected imaging quality, i.e., $SR \in [0, 1]$. In the beginning of the observation the Strehl ratio is low, as the deformable mirror is not yet adjusted to the turbulence conditions. After a few time steps, the mirror is in a state where only small adjustments are necessary - *the loop has been closed*. This happens in our case when $SR \in [0.6, 0.7]$. In this state, the linear modes for the P-WFS are valid. We show the utility of the recycled steepest descent method for multiple steps of an closed-loop iteration for a wavefront reconstruction for the instrument METIS of the ELT. The recycling subspaces were formed by the previous 1 – 3 reconstructions, see Figure 5 for the results. In the long run, both steepest descent with recycling and warm restart (i.e, where the last reconstruction is used as a starting value for the new reconstruction) yield the same Strehl ratio. However, the method with recycling is able to close the loop much faster, which is important for a stable and efficient imaging of the scientific object.




FIG. 5. Monitoring of wavefront reconstruction accuracy over time.

7.3. Adaptive optics: Image reconstruction. Our final example concerns another problem related to astronomical imaging. As for any optical system, the design of the system determines the obtainable imaging quality: Instead of the "true" image one always obtains a blurred image. The extend of blurring is described mathematically by the point spread function (PSF) of the optical system, and the measured image can be modeled as the convolution of the true image and the PSF:

$$(7.2) \quad I(t) = \int I_{true}(s) \cdot PSF(t - s) ds.$$

Thus, the original image can be recovered - assuming the PSF is known - by solving the linear ill-posed problem (7.2). In this example we want to reconstruct an astronomical image from the measured image obtained by a telescope using the approximate PSF

of the telescope. It is well known that the PSF of a telescope without a turbulent atmosphere above the telescope (e.g., a telescope in space) is given by

$$(7.3) \quad PSF_0(x) = |\mathcal{F}\{P\}(x)|^2$$

where P is the characteristic function of the aperture of the main mirror of the telescope (usually an annulus) and $\mathcal{F}(\cdot)$ represents the Fourier transform of the argument. Please note that the recovery of the original image from the measured one results in most cases in an ill posed problem, at least for smooth PSF, which is, e.g., the case for (7.3). For ground based telescopes, however, the turbulent atmosphere above the telescope has to be taken into account. The related PSF can be modeled as

$$(7.4) \quad PSF_\phi(x, t) = \left| \mathcal{F}\{P(\cdot)e^{i\phi(\cdot, t)}\}(x) \right|^2.$$

Here, $\phi(r, t)$ describes the turbulence above the telescope. As we have seen above, an AO systems aims to reduce the impact of the turbulence to the imaging process - see Section 7.2. For various reasons, e.g., the time delay between measuring the turbulences and their correction or the limited resolution of wavefront sensors and deformable mirrors, such an correction will never be perfect and therefore there will always remain a residual (uncorrected) turbulence. Based on the data from the AO system it is, however, possible to estimate those uncorrected turbulences and obtain an approximation of PSF_ϕ of the observation. A full description of this problem, i.e. the reconstruction of the PSF from AO data of the scientific observation, can be found, e.g., in [10, Section 5.4] and references therein. As astronomical images are acquired over a period of time, one actually recovers a time-averaged PSF, denoted $\langle PSF_\phi(\cdot, t) \rangle$.


In our experiments, the time averaged PSF_ϕ with residual turbulences taken into account is simulated by OCTOPUS [26, 25], the simulation tool of the European Southern Observatory. In addition, we use a representation of $PSF(x)$, the PSF of the telescope without turbulences. Both PSFs are represented by matrices, i.e., as the functions spatially discretized for images of a certain size. Please note that the recovery of the original image from the convolution data remains ill-posed in both cases.

Convolution of images with PSFs was performed using the function `psfMatrix()` available, .e.g., as a part of the `Matlab` toolbox `IRTtools` [17]. We use the `Star Cluster` image, which can be obtained from the test problems of the `Matlab` package `Restore Tools` [1]. The provided artificial `Star Cluster` image is 257×257 ; so for these experiments, the discretized PSFs are resized (i.e., downscaled) to also be 257×257 using the `imresize()` function of `Matlab`. A `psfMatrix` object is then instantiated which can be applied via multiplication to an image to perform a convolution, i.e.,

$$\text{Apsf} = \text{psfMatrix}(\text{psf_257}, \text{'periodic'}, [C \ C], [N \ N])$$

where $N = 257$ and $C = 129$, the center coordinate of the image, and we use periodic boundary conditions. Entry (C, C) of the PSF indicates where the value of the PSF at $(0, 0)$ is located, where we assume the square image is centered at the origin. The PSF inducing the adjoint operator can be created via the relation $PSF_{adj}(x, y) = PSF(-x, -y)$ which, along with the image being centered at the origin, can be used to construct the discrete PSF inducing the discretized adjoint

FIG. 6. The true and blurred **Star Cluster** image with generated stars of various brightnesses, with the true image shown in log-plot to clearly show all the generated stars clearly



operator [39]. In Figure 6, we show log plots of the true **Star Cluster** image and its blurred counterpart convolved with the PSF representing the atmospheric distortion. The blurred image was perturbed with random noise generated with uniform distribution using Matlab's `rand()` function, which was then scaled to have norm $10^{-1} \|b\|_F$, where b represents the convolved true image shown on the right in Figure 6.

In this study, we demonstrate that subspace recycling works well for the considered deconvolution problem. As for a telescope with a good AO system PSF_0 is at least a coarse approximation of PSF_ϕ , we propose to form the recycling subspace from a few eigenvectors of the matrix representing PSF_0 . Here, we take advantage of the fact that although PSF_ϕ changes for every observation, PSF_0 remains constant. This enables us to pre-calculate eigenvectors of its induced convolution operator ahead of time offline.

We calculated the telescope eigenvectors via the matrix-free `eigs()` routine of Matlab which uses a Krylov-based iteration to compute eigenvalues and eigenvectors for a few of the largest eigenvalues of the operator. In this experiment, we chose to compute the top 200 but had to perform post-processing step to eliminate the spurious complex eigenvalues/-vectors that are created in pairs, leaving us with 198 eigenvalues. Ordering the eigenvalues thereof in descending order, we augment with different collections of these pre-computed eigenvectors, starting with just the first, then the first two, and so on, up to all 198. In Figure 7, we show the residual and error curves for all experiments together. One sees that augmenting with these eigenvectors is effective in increasing the speed of convergence of the steepest descent method, with acceleration increasing as we add more vectors, but with diminishing returns.

It is important to investigate how long it takes to achieve some manner of semi-convergence. Thus, for each experiment, we track at which iteration semiconvergence is achieved and also at which iteration the residual-based discrepancy principle is achieved. For the operator induced by the telescope PSF, we also monitor jumps in eigenvalue magnitude, and these are plotted separately. This is shown in Figure 8. We also show for completeness the first nine eigenvectors (i.e., eigenimages) of the operator induced by the telescope PSF in Figure 9.

Lastly, for the case in which we augmented with the first 37 eigenvectors, we show the actual image produced at the iteration in which semiconvergence was reached. We show both the raw image produced by augmented steepest descent, as well as one post-processed by thresholding pixels with values less than 10 to be zero. This is shown (again in log plot) in Figure 10. Augmentation of additional eigenvectors beyond the

FIG. 7. *Residual and error convergence of steepest descent (shown in red) and augmented steepest descent (shown in shades of gray) with augmentation spaces being eigenvectors of the operator induced by the telescope PSF, where we used the eigenvectors associated to the largest eigenvalues. The curve's darkness is determined by how many eigenvectors were used in the augmentation. More vectors produces a darker curve.*


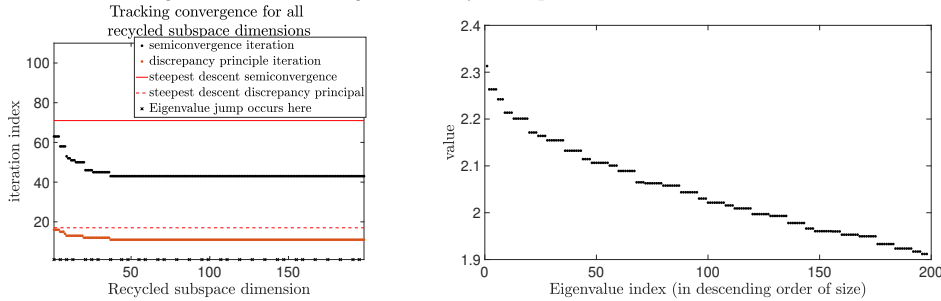


FIG. 8. *We show (left figure) for each recycled subspace dimension from this experiment at what iteration semiconvergence and the discrepancy principle are achieved, and we also mark where jumps in the eigenvalues of the operator induced by the PSF (the Airy function) of the telescope occur. On the right, we show the eigenvalues of this operator.*



first 37 produced no additional improvement for this problem.

8. Conclusions. In this paper, we have shown that under basic assumptions, the subspace recycling scheme can be combined with any known regularization scheme, and the resulting augmented scheme is still a regularization. This opens many possibilities for schemes for developing new augmented regularization schemes. We have further demonstrated this by proposing an augmented gradient descent scheme, and we show how useful recycling easily-computed information within the augmented gradient descent method can be for accelerating semi-convergence, both in an academic problem and in problems arising from applications in astronomy.

Acknowledgments. Victoria Hutterer and Ronny Ramlau are partially supported by the SFB "Tomography Across the Scales" (Austrian Science Fund Project F68-N36). The work in this paper was initiated during a visit by the second author to RICAM. The authors would like to thank Roland Wagner for providing sample PSFs for the last numerical experiment. Furthermore, the third author thanks Roland Wagner for his help in using the PSF to induce the adjoint operator needed for applying Landweber-type iterations. Lastly, the authors thank the referees for their constructive comments.

REFERENCES

FIG. 9. The first nine eigenfunctions (with eigenvalues in descending order) of the operator induced by the PSF, i.e., (7.4), of the telescope. Note that the first function is actually a constant up to roundoff noise.





FIG. 10. The image produced by augmented steepest descent at the semiconvergence iteration, raw and with pixel thresholding (all pixels of value less than 10 thresholded to zero), for the experiment augmenting the first 37 eigenvectors.



- [1] *Restoretools: An object oriented matlab package for image restoration*, 2012, <http://www.mathcs.emory.edu/~nagy/RestoreTools/>. Accessed on 28. June 2020.
- [2] K. AHUJA, E. DE STURLER, S. GUGERCIN, AND E. R. CHANG, *Recycling BiCG with an application to model reduction*, SIAM J. Sci. Comput., 34 (2012), pp. A1925–A1949, <https://doi.org/10.1137/100801500>, <http://dx.doi.org/10.1137/100801500>.
- [3] J. BAGLAMA AND L. REICHEL, *Augmented GMRES-type methods*, Numerical Linear Algebra with Applications, 14 (2007), pp. 337–350, <https://doi.org/10.1002/nla.518>, <http://dx.doi.org/10.1002/nla.518>.
- [4] J. BAGLAMA AND L. REICHEL, *Decomposition methods for large linear discrete ill-posed prob-*

- lems, *Journal of Computational and Applied Mathematics*, 198 (2007), pp. 332–343, <https://doi.org/10.1016/j.cam.2005.09.025>, <http://dx.doi.org/10.1016/j.cam.2005.09.025>.
- [5] J. CHUNG, E. DE STURLER, AND J. JIANG, *Hybrid projection methods with recycling for inverse problems*, (2020), <https://arxiv.org/abs/2007.00207>.
 - [6] E. DE STURLER, *Nested Krylov methods based on GCR*, *Journal of Computational and Applied Mathematics*, 67 (1996), pp. 15–41, [https://doi.org/10.1016/0377-0427\(94\)00123-5](https://doi.org/10.1016/0377-0427(94)00123-5), [http://dx.doi.org/10.1016/0377-0427\(94\)00123-5](http://dx.doi.org/10.1016/0377-0427(94)00123-5).
 - [7] E. DE STURLER, M. KILMER, AND K. M. SOODHALTER, *Krylov subspace augmentation for the solution of shifted systems: a review*, in preparation.
 - [8] M. DEFRISE AND C. D. MOL, *A note on stopping rules for iterative methods and filtered svd*, in *Inverse Problems: An Interdisciplinary Study*, P. C. Sabatier, ed., Academic Press, 1987, pp. 261–268.
 - [9] Y. DONG, H. GARDE, AND P. C. HANSEN, *R³GMRES: including prior information in GMRES-type methods for discrete inverse problems*, *Electron. Trans. Numer. Anal.*, 42 (2014), pp. 136–146.
 - [10] L. DYKES, R. RAMLAU, L. REICHEL, K. M. SOODHALTER, AND R. WAGNER, *Lanczos-based fast blind deconvolution methods*, *Journal of Computational and Applied Mathematics*, (To Appear).
 - [11] B. ELLERBROEK AND C. VOGEL, *Inverse problems in astronomical adaptive optics*, *Inverse Problems*, 25 (2009), p. 063001.
 - [12] H. ENGL AND R. RAMLAU, *Regularization of Inverse Problems*, Springer, New York, 2015, pp. 1233–1241.
 - [13] H. W. ENGL, M. HANKE, AND A. NEUBAUER, *Regularization of inverse problems.*, Dordrecht: Kluwer Academic Publishers, 1996.
 - [14] Y. A. ERLANGGA AND R. NABBEN, *Deflation and balancing preconditioners for Krylov subspace methods applied to nonsymmetric matrices*, *SIAM Journal on Matrix Analysis and Applications*, 30 (2008), pp. 684–699.
 - [15] A. GAUL, *Recycling Krylov subspace methods for sequences of linear systems: Analysis and applications*, PhD thesis, Technischen Universität Berlin, 2014.
 - [16] A. GAUL, M. H. GUTKNECHT, J. LIESEN, AND R. NABBEN, *A framework for deflated and augmented Krylov subspace methods*, *SIAM Journal on Matrix Analysis and Applications*, 34 (2013), pp. 495–518, <https://doi.org/10.1137/110820713>, <http://dx.doi.org/10.1137/110820713>.
 - [17] S. GAZZOLA, P. C. HANSEN, AND J. G. NAGY, *IR Tools: a MATLAB package of iterative regularization methods and large-scale test problems*, *Numer. Algorithms*, 81 (2019), pp. 773–811, <https://doi.org/10.1007/s11075-018-0570-7>, <https://doi-org.libproxy.temple.edu/10.1007/s11075-018-0570-7>.
 - [18] M. H. GUTKNECHT, *Deflated and augmented Krylov subspace methods: a framework for deflated BiCG and related solvers*, *SIAM J. Matrix Anal. Appl.*, 35 (2014), pp. 1444–1466, <https://doi.org/10.1137/130923087>, <https://doi-org.libproxy.temple.edu/10.1137/130923087>.
 - [19] V. HUTTERER AND R. RAMLAU, *Nonlinear wavefront reconstruction methods for pyramid sensors using Landweber and Landweber-Kaczmarz iteration*, *Applied Optics*, 57 (2018), pp. 8790–8804.
 - [20] V. HUTTERER AND R. RAMLAU, *Wavefront reconstruction from non-modulated pyramid wavefront sensor data using a singular value type expansion*, *Inverse Problems*, 34 (2018), p. 035002.
 - [21] V. HUTTERER, R. RAMLAU, AND I. SHATOKHINA, *Real-time adaptive optics with pyramid wavefront sensors: part I. a theoretical analysis of the pyramid sensor model*, *Inverse Problems*, 35 (2019), p. 045007, <https://doi.org/10.1088/1361-6420/ab0656>, <https://doi.org/10.1088%2F1361-6420%2Fab0656>.
 - [22] R. I. SHATOKHINA, V. HUTTERER, *Review on methods for wavefront reconstruction from pyramid wavefront sensor data*, *Journal of Astronomical Telescopes, Instruments and Systems*, 6 (2020), p. 010901, <https://doi.org/10.1117/1.JATIS.6.1.010901>.
 - [23] M. E. KILMER AND E. DE STURLER, *Recycling subspace information for diffuse optical tomography*, *SIAM Journal on Scientific Computing*, 27 (2006), pp. 2140–2166, <https://doi.org/10.1137/040610271>, <http://dx.doi.org/10.1137/040610271>.
 - [24] L. LANDWEBER, *An Iteration Formula for Fredholm Integral Equations of the First Kind*, *American Journal of Mathematics*, 73 (1951), pp. 615–624.
 - [25] M. LE LOUARN, C. VÉRINAUD, V. KORKIAKOSKI, AND E. FEDRIGO, *Parallel simulation tools for AO on ELTs*, in *Advancements in Adaptive Optics*, Proc. SPIE 5490, 2004, pp. 705–712.
 - [26] M. LE LOUARN, C. VÉRINAUD, V. KORKIAKOSKI, N. HUBIN, AND E. MARCHETTI, *Adaptive optics simulations for the European Extremely Large Telescope - art. no. 627234*, in *Advances*

- in Adaptive Optics II, Prs 1-3, vol. 6272, 2006, pp. U1048–U1056.
- [27] A. K. LOUIS, *Inverse und schlecht gestellte Probleme*, Teubner Studienbücher Mathematik, Vieweg+Teubner Verlag, 1989.
 - [28] R. B. MORGAN, *GMRES with deflated restarting*, SIAM J. Sci. Comput., 24 (2002), pp. 20–37, <https://doi.org/10.1137/S1064827599364659>, <https://doi-org.libproxy.temple.edu/10.1137/S1064827599364659>.
 - [29] M. L. PARKS, E. DE STURLER, G. MACKEY, D. D. JOHNSON, AND S. MAITI, *Recycling Krylov subspaces for sequences of linear systems*, SIAM Journal on Scientific Computing, 28 (2006), pp. 1651–1674, <https://doi.org/10.1137/040607277>, <http://dx.doi.org/10.1137/040607277>.
 - [30] M. L. PARKS, K. M. SOODHALTER, AND D. B. SZYLD, *A block recycled gmres method with investigations into aspects of solver performance*, (2016), <https://arxiv.org/abs/1604.01713>.
 - [31] R. RAMLAU AND M. ROSENSTEINER, *An efficient solution to the atmospheric turbulence tomography problem using Kaczmarz iteration*, Inverse Problems, 28 (2012), p. 095004.
 - [32] R. RAMLAU AND B. STADLER, *An augmented wavelet reconstructor for atmospheric tomography*, 2020, <https://arxiv.org/abs/2011.06842>.
 - [33] Y. SAAD, *Krylov subspace methods for solving large unsymmetric linear systems*, Mathematics of Computation, 37 (1981), pp. 105–126, <https://doi.org/10.2307/2007504>, <http://dx.doi.org/10.2307/2007504>.
 - [34] Y. SAAD, *Iterative Methods for Sparse Linear Systems*, SIAM, Philadelphia, Second ed., 2003.
 - [35] O. SCHERZER, M. GRASMAIR, H. GROSSAUER, M. HALTMEIER, AND F. LENZEN, *Variational Methods in Imaging*, Springer, New York, 2009.
 - [36] K. M. SOODHALTER, E. DE STURLER, AND M. KILMER, *A survey of subspace recycling iterative methods*, GAMM Mitteilungen, (2020). Applied and numerical linear algebra topical issue.
 - [37] L. N. TREFETHEN AND D. BAU, III, *Numerical linear algebra*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1997, <https://doi.org/10.1137/1.9780898719574>, <https://doi-org.libproxy.temple.edu/10.1137/1.9780898719574>.
 - [38] S. WANG, E. DE STURLER, AND G. H. PAULINO, *Large-scale topology optimization using preconditioned Krylov subspace methods with recycling*, International Journal for Numerical Methods in Engineering, 69 (2007), pp. 2441–2468, <https://doi.org/10.1002/nme.1798>, <http://dx.doi.org/10.1002/nme.1798>.
 - [39] R. WANGER. Personal Communication, Mar. 2020.
 - [40] M. YUDYTSKIY, T. HELIN, AND R. RAMLAU, *Finite element-wavelet hybrid algorithm for atmospheric tomography*, J. Opt. Soc. Am. A, 31 (2014), pp. 550–560, <https://doi.org/10.1364/JOSAA.31.000550>, <http://josaa.osa.org/abstract.cfm?URI=josaa-31-3-550>.