

A convex variational model for learning convolutional image atoms from incomplete data

A. Chambolle, M. Holler and T. Pock

Abstract

A variational model for learning convolutional image atoms from corrupted and/or incomplete data is introduced and analyzed both in function space and numerically. Building on lifting and relaxation strategies, the proposed approach is convex and allows for simultaneous image reconstruction and atom-learning in a general, inverse problems context. Further, motivated by an improved numerical performance, also a semi-convex variant is included in the analysis and the experiments of the paper. For both settings, fundamental analytical properties allowing in particular to ensure well-posedness and stability results for inverse problems are proven in a continuous setting. Exploiting convexity, globally optimal solutions are further computed numerically for applications with incomplete, noisy and blurry data and numerical results are shown.

Mathematical subject classification: 94A08 49M29 65F22 49K30

Keywords: Variational methods, learning approaches, inverse problems, functional lifting, convex relaxation, convolutional lasso, machine learning, texture reconstruction.

1 Introduction

An important task in image processing is to achieve an appropriate regularization or smoothing of images or image related data. In particular, this is indispensable for most application-driven problems in the field, such as denoising, inpainting, reconstruction, segmentation, registration or classification. Also beyond imaging, for general problem settings in the field of inverse problems, an appropriate regularization of unknowns plays a central role as it allows for a stable inversion procedure.

Variational methods and partial-differential-equation-based methods can now be regarded as classical regularization approaches of mathematical image processing (see for instance [42, 3, 40, 30]). An advantage of such methods is the existence of a well-established mathematical theory and, in particular for variational methods, a direct applicability to general inverse problems with provable stability and recovery guarantees [24, 25]. While in particular piecewise smooth images are typically well-described by such methods, their performance for oscillatory or texture-like structures, however, is often limited to pre-described patterns (see for instance [21]).

Data-adaptive methods such as patch-based methods (see for instance [28, 17, 11, 18]) on the other hand are able to exploit redundant structures in images independent of an a-priori description and are, at least for some specific tasks, often superior to variational- and PDE-based methods. In particular machine-learning-based methods have advanced the state-of-the art significantly in many typical imaging applications in the past years. A disadvantage of such methods, however,

is a current lack of mathematical understanding in particular compared to existing results for variational methods in the context of inverse problems. Data-adaptive methods typically build on learning some image atoms in some way or the other and neither the mapping of training data to image atoms nor the application of learned atoms to reconstruction is known to enjoy the same stability and regularization properties as with classical methods. In particular, for typical deep neuronal networks, the training step corresponds to the minimization of a non-convex energy, but neither is this minimization carried out until convergence nor are the properties of local minimizers well understood mathematically.

The goal of this work is to provide a first step towards bridging the gap between data-driven methods and variational methods. Building on a tensorial-lifting approach, we introduce a convex variational method for learning image atoms from noisy and/or incomplete data in an inverse problems context. We further extend this model by a semi-convex variant that improves the performance in some applications. For both settings, we are able to prove well-posedness results in function space and, for the convex version, to compute globally optimal solutions numerically. In particular, classical stability and convergence results for inverse problems such as the ones of [24, 25] are applicable to our model, providing a stable recovery of both learned atoms and images from given, incomplete data.

Our approach is motivated by a sparse, convolutional representation of images via a few image atoms and conceptually allows for a joint learning of image atoms and image reconstruction in a single step. Nevertheless, it can also be regarded purely as an approach for learning image atoms from potentially incomplete data in a training step, after which the learned atoms can be further incorporated in a second step, e.g., for reconstruction or classification. It should also be noted that, while we show some examples where our approach achieves a good practical performance for image reconstruction compared to existing methods, the main purpose of this paper is to provide a mathematical understanding rather than an algorithm that achieves the best performance in practice.

Regarding existing literature in the context of data-adaptive variational learning approaches in imaging, we note that there are many recent approaches that aim to infer either parameter or filters for variational methods from given training data, see e.g. [27, 12, 23]. A continuation of such techniques more towards the architecture of neuronal networks are so called variational networks [26, 2] where not only model parameters but also components of the solution algorithm such as stepsizes or proximal mappings are learned. We also refer to [29] for a recent work on combining variational methods and neuronal networks. While for some of those methods also a function space theory is available, the learning step is still non-convex and one can in general only expect to obtain locally optimal solutions.

1.1 Outline of the paper

In Section 2 we present the main ideas for our approach in a formal setting. This is done from two perspectives, once from the perspective of a convolutional lasso approach and once from the perspective of patch-based methods. In Section 3 we then carry out an analysis of the proposed model in function space, where we motivate our approach via convex relaxation and derive well-posedness results. Section 4 then presents the model in a discrete setting and the numerical solution strategy and Section 5 provides numerical results and a comparison to existing methods. At last, an appendix provides a brief overview on some results for tensor spaces that are used in Section 3. We note that, while the analysis of Section 3 is an important part of our work, the paper is

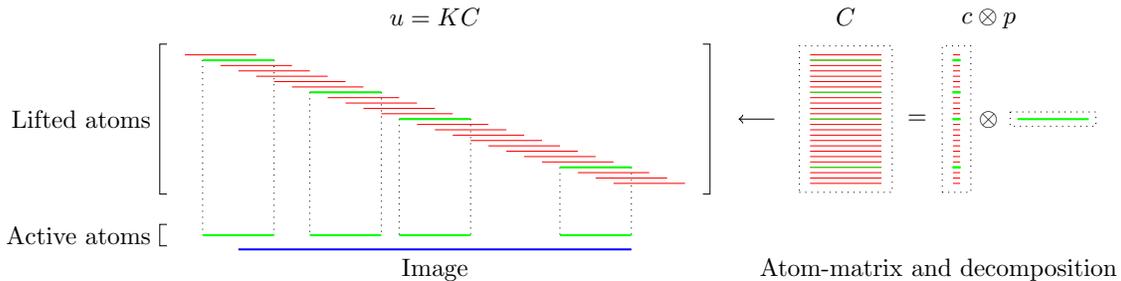


Figure 1: Visualization of the atom-lifting approach for 1D images. The green (thick) lines in the atom-matrix correspond to non-zero (active) atoms and are placed in the image at the corresponding positions.

structured in a way such that readers only interested in the conceptual idea and the realization of our approach can skip Section 3 and focus on Sections 2 and 4.

2 A convex approach to image atoms

In this section we present the proposed approach to image-atom-learning and texture reconstruction, where we focus on explaining the main ideas rather than precise definitions of the involved terms. For the latter, we refer to Section 3 for the continuous model and Section 4 for the discrete setting.

Our starting point is the convolutional lasso problem [44, 15], which aims to decompose a given image u as a sparse linear combination of basic atoms $(p_i)_{i=1}^k$ with coefficient images $(c_i)_{i=1}^k$ by inverting a sum of convolutions as follows

$$\min_{(c_i)_i, (p_i)_i} \sum_{i=1}^k \|c_i\|_1 \quad \text{s.t.} \quad \begin{cases} u = \sum_{i=1}^k c_i * p_i, \\ \|p_i\|_2 \leq 1 \text{ for } i = 1, \dots, k. \end{cases}$$

It is important to note that, by choosing the $(c_i)_i$ to be composed of delta peaks, this allows to place the atoms $(p_i)_i$ and any position in the image. In [44], this model has been used in the context of convolutional neural networks for generating image atoms and other image related tasks. Subsequently, many works have dealt with the algorithmic solution of the resulting optimization problem, where the main difficulty lies in the non-convexity of the atom learning step, and we refer to [22] for a recent review.

Our goal is to obtain a convex relaxation of this model that can be used for both, learning image atoms from potentially noisy data as well as image reconstruction tasks such as inpainting, deblurring or denoising. To this aim, we lift the model to the tensor product space of coefficient images and image atoms, i.e., the space of all tensors $C = \sum_i c_i \otimes p_i$ with $c_i \otimes p_i$ being a rank-1 tensor such that $(c_i \otimes p_i)(x, y) = c_i(x)p_i(y)$. We refer to Figure 1 for a visualization of this lifting in a one-dimensional setting, where both coefficients and image atoms are vectors and $c_i \otimes p_i$ corresponds to a rank-one matrix. Notably, in this tensor product space, the convolution $c_i * p_i$ can be written as linear operator K such that $KC(x) = \sum_i K(c_i \otimes p_i)(x) = \sum_i \int p_i(x - y)c_i(y)$. Exploiting redundancies in the penalization of $(\|c_i\|_1)_i$ and the constraint $\|p_i\|_2 \leq 1, i = 1 \dots, k$

and re-writing the above setting in the lifted tensor space, as discussed in Section 3, we obtain the following minimization problem as convex relaxation of the convolutional lasso approach

$$\min_C \|C\|_{1,2} \quad \text{s.t. } u = KC,$$

where $\|\cdot\|_{1,2}$ takes the 1-norm and 2-norm of C in coefficient and atom direction, respectively. Now while a main feature of the original model was that the number of image atoms was fixed, this is no longer the case in the convex relaxation and would correspond to constraining the rank of the lifted variable C (defined as the minimal number of simple tensors needed to decompose C) to be below a fixed number. As convex surrogate, we add an additional penalization of the nuclear norm of C in the above objective functional (here we refer to the nuclear norm of C in the tensor product space which, in the discretization of our setting, coincides with the classical nuclear norm of a matrix-reshaping of C). Allowing also for additional linear constraints on C via a linear operator M , we arrive at the following convex norm that measures the decomposability of a given image u into a sparse combination of atoms as

$$N_\nu(u) = \min_C \nu \|C\|_{1,2} + (1 - \nu) \|C\|_* \quad \text{such that } u = KC, MC = 0.$$

Interestingly, this provides a convex model for learning image atoms, which for simple images admitting a sparse representation seems quite effective. In addition, this can in principle also be used as a prior for image reconstruction tasks in the context of inverse problems via solving for example

$$\min_u \frac{\lambda}{2} \|Au - u_0\|_2^2 + N_\nu(u),$$

with u_0 given some corrupted data, A a forward operator and $\lambda > 0$ a parameter.

Both the original motivation for our model as well as its convex variant have many similarities with dictionary learning and patch-based methods. The next section strives to clarify similarities and difference and provides a rather interesting, different perspective on our model.

2.1 A dictionary-learning/patch-based methods perspective

In classical dictionary-learning-based approaches, the aim is to represent a resorted matrix of image patches as a sparse combination of dictionary atoms. That is, with $u \in \mathbb{R}^{NM}$ a vectorized version of an image and $D = (D_1, \dots, D_l)^T \in \mathbb{R}^{l \times nm}$ a patch-matrix containing l vectorized (typically overlapping) images patches of size nm , the goal is to obtain a decomposition $D = cd$, where $c \in \mathbb{R}^{l \times k}$ is a coefficient matrix and $p \in \mathbb{R}^{k \times nm}$ is a matrix of k dictionary atoms such that $c_{i,j}$ is the coefficient for the atom $p_{j,\cdot}$ in the representation of the patch D_i . In order to achieve a decomposition in this form, using only a sparse representation of dictionary atoms, a classical approach is to solve

$$\min_{c,p} \frac{\lambda}{2} \|cp - D\|_2^2 + \|c\|_1 + R(p),$$

where R potentially puts additional constraints or cost on the dictionary atoms, e.g. $R(p) = 0$ if $\|p_{j,\cdot}\|_2 \leq 1$ for all j and $R(p) = \infty$ else.

A difficulty with such an approach is again the bilinear and hence non-convex nature of the optimization problem, leading to potentially many non-optimal stationary points and making the approach sensitive to initialization.

As a remedy, one strategy is to consider a convex variant (see for instance [4]). That is, re-writing the above minimization problem (and again using the ambiguity in the product cd to eliminate the L^2 constraint) we arrive at the problem

$$\min_{C:\text{rank}(C)\leq k} \frac{\lambda}{2} \|C - D\|_2^2 + \|C\|_{1,2},$$

where $\|C\|_{1,2} = \sum_i \|C_{i,\cdot}\|_2$. A possible convexification is then given as

$$\min_C \frac{\lambda}{2} \|C - D\|_2^2 + \nu \|C\|_{1,2} + (1 - \nu) \|C\|_*, \quad (1)$$

where $\|\cdot\|_*$ is the nuclear norm of the matrix C .

A disadvantage of such an approach is that the selection of patches is a-priori fixed and that the lifted matrix C has to approximate each patch. In the situation of overlapping patches, this means that individual rows of C have to represent different shifted version of the same patch several times, which inherently contradicts the low-rank assumption.

It is now interesting to see our proposed approach in relation to this dictionary learning methods and the above-described disadvantage: Denote again by K the lifted version of the convolution operator, which in the discrete setting takes a lifted patch-matrix as input and provides an image composed of overlapping patches as output. It is than easy to see that K^* , the adjoint of K , is in fact a patch selection operator and it holds that $KK^* = I$. Now using K^* , the approach in (1) can be re-written as

$$\min_C \frac{\lambda}{2} \|C - K^*u\|_2^2 + \nu \|C\|_{1,2} + (1 - \nu) \|C\|_*, \quad (2)$$

where we remember that u is the original image. Considering the problem of finding an optimal patch-based representation of an image as the problem if inverting K , we can see that the previous approach in fact first applies a right-inverse of K and then decomposes the data. Taking this perspective, however, it seems much more natural to consider instead an adjoint formulation as

$$\min_C \frac{\lambda}{2} \|KC - u\|_2^2 + \nu \|C\|_{1,2} + (1 - \nu) \|C\|_*. \quad (3)$$

Indeed, this means that we do not fix the patch-decomposition of the image a-priori but rather allow the method itself to optimally select the position and size of patches. In particular, given a particular patch at an arbitrary location, this patch can be represented by using only one line of C and the other lines (corresponding to shifted versions) can be left empty. Figure 2 shows the resulting improvement by solving both of the above optimization problems for a particular test image, where the parameters are set such that the data error of both methods, i.e., $\|KC - U\|_2^2$ is approximately the same. As can be seen, solving (2), which we call patch denoising, does not yield meaningful dictionary atoms as the dictionary elements need to represent different, shifted version of the single patch that makes up the image. In contrast to that, solving (3), which we call patch reconstruction, allows to identify the underlying patch of the image and the corresponding patch-matrix is indeed row-sparse.

2.2 The variational model

Now while the proposed model can, in principle, describe any kind of images, in particular its convex relaxation seems best suited for situations where the image can be described by only a

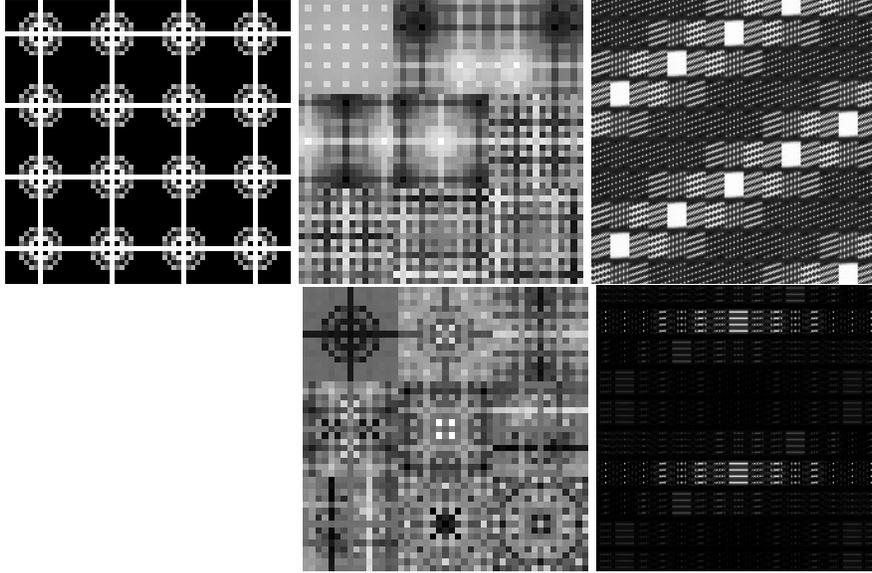


Figure 2: Patch-based representation of test images. Left: Original image, Middle: Nine most important patches for each method (top: patch denoising, bottom: patch reconstruction), Right: Section of the corresponding patch matrices.

few, repeating atoms, as would be for instance the case with texture images. In particular, since we do not include rotations in the model, there are many simple situations, such as u being the characteristic function of a disk, which would in fact require an infinite number of atoms. To overcome this, it seems beneficial to include an additional term which is rotationally invariant and takes care of piecewise smooth parts of the image. Denoting R to be any regularization functional for piecewise smooth data and taking the infimal convolution of this functional with our atom-based norm, we then arrive at the convex model

$$\min_{u,v} \frac{\lambda}{2} \|Au - u_0\|_2^2 + \mu_1 R(u - v) + \mu_2 N_\nu(v),$$

for learning image data and image atoms kernels from potentially noisy or incomplete measurements.

A particular example of this model can be given when choosing $R = \text{TV}$, the total variation function [37]. In this setting, a natural choice for the operator M in the definition of N_ν is to take the point-wise mean of the lifted variable in atom direction, which corresponds to constraining the learned atoms to have zero mean and enforces some kind of orthogonality between the cartoon and the texture part in the spirit of [31]. In our numerical experiments, in order to obtain an even richer model for the cartoon part, we use the second order Total Generalized Variation function (TGV_α^2) [7, 5] as cartoon prior and, in the spirit of a dual TGV_α^2 norm, use M to constrain the 0th and 1st moments of the atoms to be zero.

We also remark that, as shown in the analysis part of Section 3, while an ℓ^1/ℓ^2 -type norm on the lifted variables indeed arises as convex relaxation of the convolutional lasso approach, the addition of the nuclear norm is to some extent arbitrary and in fact in the context of compressed sensing it is known that a summation of two norms is suboptimal for a joint penalization of sparsity

and rank [32] (we refer to Remark 5 below for an extended discussion). Indeed, our numerical experiments also indicate that the performance of our method is to some extent limited by a sub-optimal relaxation of a joint sparsity and rank penalization. To account for that, we also tested with semi-convex potential functions (instead of the identity) for a penalization of the singular values in the nuclear norm. Since this provided a significant improvement in some situations, we also include this more general setting in our analysis and the numerical results.

3 The model in a continuous setting

The goal of this section is to define and analyze the model introduced in Section 2 in a continuous setting. To this aim, we regard images as functions in the Lebesgue spaces $L^q(\Omega)$ with a bounded Lipschitz domain $\Omega \subset \mathbb{R}^d$, $d \in \mathbb{N}$ and $1 < q \leq 2$. Image atoms are regarded as functions in $L^s(\Sigma)$, with $\Sigma \subset \mathbb{R}^d$ a second (smaller) bounded Lipschitz domain (either a circle or a square around the origin) and $s \in [q, \infty]$ an exponent that is a-priori allowed to take any value in $[q, \infty]$, but will be further restricted below. We also refer to the appendix for further notation and results, in particular in the context of tensor product spaces, that will be used in this section.

As described in Section 2 above, the main idea is to synthesize an image via the convolution of a small number of atoms with corresponding coefficient images, where we think of the latter as a sum delta peaks that define the locations where atoms are placed. For this reason, and also due to compactness properties, the coefficient images are modeled as Radon measures in the space $\mathcal{M}(\Omega_\Sigma)$, the dual of $C_0(\Omega_\Sigma)$, where we denote

$$\Omega_\Sigma := \{x \in \mathbb{R}^d \mid \text{there exists } y \in \Sigma \text{ s.t. } x - y \in \Omega\},$$

i.e., the extension of Ω by Σ . The motivation for using this extension of Ω is to allow atoms also to be placed arbitrarily close to the boundary (see Figure 1). We will further use the notation $r' = r/(r-1)$ for an exponent $r \in (1, \infty)$ and denote duality pairings between L^r and $L^{r'}$ and between $\mathcal{M}(\Omega)$ and $C_0(\Omega)$ by (\cdot, \cdot) , while other duality pairings (e.g. between tensor spaces) are denoted by $\langle \cdot, \cdot \rangle$. By $\|\cdot\|_r, \|\cdot\|_{\mathcal{M}}$ we denote standard L^r and Radon norms whenever the domain of definition is clear from the context, otherwise we write $\|\cdot\|_{L^r(\Omega_\Sigma)}, \|\cdot\|_{\mathcal{M}(\Omega_\Sigma)}$ etc.

3.1 The convolutional lasso prior

As first step, we deal with the convolution operator that synthesizes an image from a pair of a coefficient image and an image atom in function space. Formally, we aim to define $K : \mathcal{M}(\Omega_\Sigma) \times L^s(\Sigma) \rightarrow L^q(\Omega)$ as

$$K(c, p)(x) := \int_{\Omega_\Sigma} p(x - y) dc(y),$$

where we extend p by zero outside of Σ . An issue with this definition is that, in general, p is only defined Lebesgue almost everywhere and so we have to give a rigorous meaning to the integration of p with respect to an arbitrary Radon measure. To this aim, we define the convolution operator via duality (see [38]). For $c \in \mathcal{M}(\Omega_\Sigma)$, $p \in L^s(\Sigma)$ we define by $K_{c,p}$ the functional on $C(\bar{\Omega})$ as dense subset of $L^{q'}(\Omega)$ as

$$K_{c,p}(h) := \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \tilde{h}(z + y) \tilde{p}(z) dz d\tilde{c}(y),$$

where \tilde{g} always denotes the zero extension of the function or measure g outside their domain of definition. Now we can estimate with $D > 0$

$$|K_{c,p}(h)| \leq \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} |\tilde{h}(z+y)| |\tilde{p}(z)| \, dz \, d|\tilde{c}|(y) \leq \|\tilde{h}\|_{L^{q'}(\mathbb{R}^d)} \|\tilde{p}\|_{L^q(\mathbb{R}^d)} \|\tilde{c}\|_{\mathcal{M}(\mathbb{R}^d)} \leq D \|p\|_s \|c\|_{\mathcal{M}} \|h\|_{q'}.$$

Hence, by density we can uniquely extend $K_{c,p}$ to a functional in $L^{q'}(\Omega)^* \simeq L^q(\Omega)$ and we denote by $[K_{c,p}]$ the associated function in $L^q(\Omega)$. Now in case p is integrable w.r.t. c and $x \mapsto \int_{\Omega_\Sigma} p(x-y) \, dc(y) \in L^q(\Omega)$, we get by a change of variables and Fubini's theorem that for any $h \in C(\overline{\Omega})$

$$K_{c,p}(h) = \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} \tilde{h}(x) \tilde{p}(x-y) \, dx \, d\tilde{c}(y) = \int_{\Omega} h(x) \left(\int_{\Omega_\Sigma} p(x-y) \, dc(y) \right) \, dx.$$

Hence we get that in this case, $[K_{c,p}](x) = \int_{\Omega_\Sigma} p(x-y) \, dc(y)$ and defining $K : \mathcal{M}(\Omega_\Sigma) \times L^s(\Sigma) \rightarrow L^q(\Omega)$ as

$$K(c, p) := [K_{c,p}]$$

we get that $K(c, p)$ coincides with the convolution of c and p whenever the latter is well defined. Note that K is bilinear and, as the previous estimate shows, there exists $D > 0$ such that $\|K(c, p)\|_q \leq D \|c\|_{\mathcal{M}} \|p\|_s$. Hence $K \in \mathcal{B}(\mathcal{M}(\Omega_\Sigma) \times L^s(\Sigma), L^q(\Omega))$, the space of bounded bilinear operators (see the appendix).

Using the bilinear operator K and denoting by $k \in \mathbb{N}$ a fixed number for kernels, we now define the convolutional lasso prior for an exponent $s \in [q, \infty]$ and for $u \in L^q(\Omega)$ as

$$N_{\text{cl},s}(u) = \inf_{\substack{(c_i)_{i=1}^k \subset \mathcal{M}(\Omega_\Sigma) \\ (p_i)_{i=1}^k \subset L^s(\Sigma)}} \sum_{i=1}^k \|c_i\|_{\mathcal{M}} \quad \text{s.t.} \quad \begin{cases} \|p_i\|_s \leq 1, Mp_i = 0 & i = 1, \dots, k, \\ u = \sum_{i=1}^k K(c_i, p_i) & \text{in } \Omega, \end{cases} \quad (4)$$

and set $N_{\text{cl},s}(u) = \infty$ if the constraint set above is empty. Here, we include an operator $M \in \mathcal{L}(L^s(\Sigma), \mathbb{R}^m)$ in our model that optionally allows to enforce additional constraints on the atoms. A simple example of M that we have in mind is an averaging operator, i.e., $Mp := |\Sigma|^{-1} \int_{\Sigma} p(x) \, dx$, hence the constraint that $Mp = 0$ corresponds to a zero-mean constraint.

3.2 A convex relaxation

Our goal is now to obtain a convex relaxation of the convolutional lasso prior. To this aim, we introduce by \hat{K} and $\hat{M} := I \otimes M$ the lifting of the bilinear operator K and the linear operators I and M , with $I \in \mathcal{L}(\mathcal{M}(\Omega_\Sigma), \mathcal{M}(\Omega_\Sigma))$ being the identity, to the projective tensor product space $X_s := \mathcal{M}(\Omega_\Sigma) \otimes_{\pi} L^s(\Sigma)$ (see the appendix). In this space, we consider a reformulation as

$$N_{\text{cl},s}(u) = \inf_{C \in X_s} \|C\|_{\pi,k,M} \quad \text{s.t.} \quad u = \hat{K}C \text{ in } \Omega, \quad (5)$$

where

$$\|C\|_{\pi,k,M} := \inf \left\{ \sum_{i=1}^k \|c_i\|_{\mathcal{M}} \|p_i\|_s \mid C = \sum_{i=1}^k c_i \otimes p_i \text{ with } Mp_i = 0 \text{ for } i = 1, \dots, k \right\}.$$

Note that this reformulation is indeed equivalent. Next we aim to derive the convex relaxation of $N_{\text{cl},s}$ in this tensor product space. First we consider a relaxation of the functional $\|\cdot\|_{\pi,k,M}$. To this aim, we need an additional assumption on the constraint set $\ker(M)$, which is satisfied for instance if $s = 2$ or for $M = 0$, in particular will be fulfilled by the concrete setting we use later on.

Lemma 1. *Assume that there exists a continuous, linear, norm-one projection onto $\ker(M)$. Then, the convex, lower semi-continuous relaxation of $\|\cdot\|_{\pi,k,M} : X_s \rightarrow \overline{\mathbb{R}}$ is given as*

$$C \mapsto \|C\|_{\pi} + \mathcal{I}_{\ker(\hat{M})}(C),$$

where $\mathcal{I}_{\ker(\hat{M})}(C) = 0$ if $\hat{M}C = 0$ and $\mathcal{I}_{\ker(\hat{M})}(C) = \infty$ else, and $\|\cdot\|_{\pi}$ is the projective norm on X_s given as

$$\|C\|_{\pi} = \inf \left\{ \sum_{i=1}^{\infty} \|c_i\|_{\mathcal{M}} \|p_i\|_s \mid C = \sum_{i=1}^{\infty} c_i \otimes p_i \right\}.$$

Proof. At first note that, for a general function $g : X_s \rightarrow \overline{\mathbb{R}}$, its convex, lower-semi continuous relaxation can be computed as the biconjugate $g^{**} : X_s \rightarrow \overline{\mathbb{R}}$, where $g^*(x^*) = \sup_{x \in X_s} \langle x^*, x \rangle - g(x)$ and $g^{**}(x) = \sup_{x^* \in X_s^*} \langle x^*, x \rangle - g^*(x^*)$. Keeping this in mind, we note that

$$\|\cdot\|_{\pi} + \mathcal{I}_{\ker(\hat{M})} \leq \|\cdot\|_{\pi,k,M} \leq \|\cdot\|_{\pi,1,M},$$

and consequently

$$\|\cdot\|_{\pi} + \mathcal{I}_{\ker(\hat{M})} \leq (\|\cdot\|_{\pi,k,M})^{**} \leq (\|\cdot\|_{\pi,1,M})^{**}.$$

Hence the assertion follows if we show that $(\|\cdot\|_{\pi,1,M})^{**} \leq \|\cdot\|_{\pi} + \mathcal{I}_{\ker(\hat{M})}$. To this aim, we first show that $\|\cdot\|_{\pi,M} \leq \|\cdot\|_{\pi} + \mathcal{I}_{\ker(\hat{M})}$, where we set $\|C\|_{\pi,M} = \|C\|_{\pi,\infty,M}$. Let $C \in X_s$ be such that $\hat{M}C = 0$ and take $(c_i)_i, (p_i)_i$ be such that $\|C\|_{\pi} \geq \sum_{i=1}^{\infty} \|c_i\|_{\mathcal{M}} \|p_i\|_s - \epsilon$ for some $\epsilon > 0$ and $C = \sum_{i=1}^{\infty} c_i \otimes p_i$. Then, with P the projection to $\ker(M)$ as in the assumption, we get that

$$0 = \hat{M}C = \sum_{i=1}^{\infty} c_i \otimes Mp_i = \sum_{i=1}^{\infty} c_i \otimes M(p_i - Pp_i).$$

Now remember that, according to [39, Theorem 2.9], we have $(\mathcal{M}(\Omega_{\Sigma}) \otimes_{\pi} \mathbb{R}^m)^* = \mathcal{B}(\mathcal{M}(\Omega_{\Sigma}) \times \mathbb{R}^m)$ with the norm $\|B\|_{\mathcal{B}} := \sup\{|B(x,y)| \mid \|x\|_{\mathcal{M}} \leq 1, \|y\|_2 \leq 1\}$. Taking arbitrary $\psi \in \mathcal{M}(\Omega)^*$, $\phi \in (\mathbb{R}^m)^*$, we get that $B : (c,p) \mapsto \psi(c)\phi(p) \in \mathcal{B}(\mathcal{M}(\Omega_{\Sigma}) \times \mathbb{R}^m)$ hence

$$\begin{aligned} 0 &= \hat{B}(\hat{M}C) = \sum_{i=1}^{\infty} \hat{B}(c_i \otimes M(p_i - Pp_i)) = \sum_{i=1}^{\infty} \psi(c_i)\phi(M(p_i - Pp_i)) \\ &= \phi \left(\sum_{i=1}^{\infty} \psi(c_i)M(p_i - Pp_i) \right) \\ &= \phi \left(M \left(\sum_{i=1}^{\infty} \psi(c_i)(p_i - Pp_i) \right) \right) \end{aligned}$$

and since ϕ was arbitrary it follows that $M(\sum_{i=1}^{\infty} \psi(c_i)(p_i - Pp_i)) = 0$. Finally, by closedness of $\text{Rg}(I - P)$ we get that $\sum_{i=1}^{\infty} \psi(c_i)(p_i - Pp_i) = 0$ and, since $\mathcal{M}(\Omega_{\Sigma})$ has the approximation property

(see [19, Section VIII.3]), from [39, Proposition 4.6], it follows that $\sum_{i=1}^{\infty} c_i \otimes (p_i - Pp_i) = 0$, hence $C = \sum_{i=1}^{\infty} c_i \otimes Pp_i$ and by assumption $\sum_{i=1}^{\infty} \|c_i\|_{\mathcal{M}} \|p_i\|_s \geq \sum_{i=1}^{\infty} \|c_i\|_{\mathcal{M}} \|Pp_i\|_s$. Consequently, $\|\cdot\|_{\pi} + \mathcal{I}_{\ker(\hat{M})} \geq \|\cdot\|_{\pi, M} - \epsilon$ and since ϵ was arbitrary, the claimed inequality follows.

Now we show that $(\|\cdot\|_{\pi, M})^* \leq (\|\cdot\|_{\pi, 1, M})^*$, from which the claimed assertion follows by the previous estimate and taking the convex conjugate on both sides. To this aim, take $(C_n)_n \subset X_s$ such that

$$(\|\cdot\|_{\pi, M})^*(B) = \sup_{C \in X_s} \langle B, C \rangle - \|C\|_{\pi, M} = \lim_n \langle B, C_n \rangle - \|C_n\|_{\pi, M}$$

and take $(c_i^n)_i, (p_i^n)_i$ such that $Mp_i^n = 0$ for all n, i and

$$C_n = \sum_{i=1}^{\infty} c_i^n \otimes p_i^n \quad \text{and} \quad \sum_{i=1}^{\infty} \|c_i^n\|_{\mathcal{M}} \|p_i^n\|_s \leq \|C_n\|_{\pi, M} + 1/n.$$

We then get

$$\begin{aligned} (\|\cdot\|_{\pi, M})^*(B) &= \lim_n \langle B, C_n \rangle - \|C_n\|_{\pi, M} \leq \lim_n \langle B, \sum_{i=1}^{\infty} c_i^n \otimes p_i^n \rangle - \sum_{i=1}^{\infty} \|c_i^n\|_{\mathcal{M}} \|p_i^n\|_s + 1/n \\ &= \lim_n \lim_m \langle B, \sum_{i=1}^m c_i^n \otimes p_i^n \rangle - \sum_{i=1}^m \|c_i^n\|_{\mathcal{M}} \|p_i^n\|_s + 1/n \\ &\leq \lim_n \sup_m \sup_{\substack{(c_i)_{i=1}^m, (p_i)_{i=1}^m \\ Mp_i=0}} \langle B, \sum_{i=1}^m c_i \otimes p_i \rangle - \sum_{i=1}^m \|c_i\|_{\mathcal{M}} \|p_i\|_s + 1/n \\ &= \sup_m \sup_{\substack{(c_i)_{i=1}^m, (p_i)_{i=1}^m \\ Mp_i=0}} \sum_{i=1}^m B(c_i, p_i) - \|c\|_{\mathcal{M}} \|p\|_s \\ &= \sup_m m \sup_{\substack{c, p \\ Mp=0}} (B(c, p) - \|c\|_{\mathcal{M}} \|p\|_s). \end{aligned}$$

Now it can be easily seen that the last expression equals 0 in case $|B(c, p)| \leq \|c\|_{\mathcal{M}} \|p\|_s$ for all c, p with $Mp = 0$. In the other case, we can pick \hat{c}, \hat{p} with $M\hat{p} = 0$ and $r > 1$ such that $B(\hat{c}, \hat{p}) > r\|\hat{c}\|_{\mathcal{M}}\|\hat{p}\|_s$ and get for any $\lambda > 0$ that

$$\sup_{\substack{c, p \\ Mp=0}} B(c, p) - \|c\|_{\mathcal{M}} \|p\|_s \geq B(\lambda\hat{c}, \hat{p}) - \|\lambda\hat{c}\|_{\mathcal{M}} \|\hat{p}\|_s \geq \lambda(r-1) \rightarrow \infty \text{ as } \lambda \rightarrow \infty.$$

Hence, the last line of the above equation is either 0 or infinity and equals

$$\sup_{\substack{c, p \\ Mp=0}} (B(c, p) - \|c\|_{\mathcal{M}} \|p\|_s) = \sup_C \langle B, C \rangle - \|C\|_{\pi, 1, M} = (\|\cdot\|_{\pi, 1, M})^*. \quad \square$$

This result suggests that the convex, lower semi-continuous relaxation of (5) will be obtained by replacing $\|\cdot\|_{\pi, k, M}$ with the projective tensor norm $\|\cdot\|_{\pi}$ on X_s and the constraint $\hat{M}C = 0$. Our approach to show this will in particular require us ensure lower semi-continuity of this candidate for the relaxation, which in turn requires us to ensure a compactness property of the sublevelsets of the energy appearing in (5) and closedness of the constraints. To this aim, we consider a weak*

topology on X_s and rely on a duality result for tensor product spaces (see the appendix), which states that, under some conditions, the projective tensor product $X_s = \mathcal{M}(\Omega_\Sigma) \otimes_\pi L^s(\Sigma)$ can be identified with the dual of the so called injective tensor product $C_0(\Omega_\Sigma) \otimes_i L^{s'}(\Sigma)$. Different from one what one would expect from the individual spaces, however, this can only be ensured for the case $s < \infty$ which excludes the space $L^\infty(\Sigma)$ for the image atoms. This restriction will also be required later on in order to show well-posedness of a resulting regularization approach for inverse problems, hence we will henceforth always consider the case that $s \in [q, \infty)$ and use the identification $(C_0(\Omega_\Sigma) \otimes_i L^{s'}(\Sigma))^* \hat{=} \mathcal{M}(\Omega_\Sigma) \otimes_\pi L^s(\Sigma)$ (see the appendix).

As first step towards the final relaxation result, and also as crucial ingredient for well-posedness results below, we show weak* continuity of the operator \hat{K} on the space X_s .

Lemma 2. *Let $s \in [q, \infty)$. Then the operator $\hat{K} : X_s \rightarrow L^q(\Omega)$ is continuous w.r.t. weak* convergence in X_s and weak convergence in $L^q(\Omega)$. Also, for any $\phi \in C_0(\Omega) \subset L^q(\Omega)$ it follows that $\hat{K}^* \phi \in C_0(\Omega_\Sigma) \otimes_i L^{s'}(\Sigma)$ and, via the identification $C_0(\Omega_\Sigma) \otimes_i L^{s'}(\Sigma) \hat{=} C_0(\Omega_\Sigma, L^{s'}(\Sigma))$ (see the appendix), can be given as $K^* \phi(t) = [x \mapsto \phi(t+x)]$.*

Proof. First we note that for any $\psi \in C_c(\Omega)$, the function $\hat{\psi}$ defined as $\hat{\psi}(t) = [x \mapsto \psi(t+x)]$ (where we extend ψ by 0 to \mathbb{R}^d) is contained in $C_c(\Omega_\Sigma, L^{s'}(\Sigma))$. Indeed, continuity follows since by uniform continuity for any $\epsilon > 0$ there exists a $\delta > 0$ such that for any $r \in \mathbb{R}^d$ with $|r| \leq \delta$ and $t \in \Omega_\Sigma$ with $t+r \in \Omega_\Sigma$

$$\|\hat{\psi}(t+r) - \hat{\psi}(t)\|_{s'} = \left(\int_\Sigma |\psi(t+r+x) - \psi(t+x)|^{s'} dx \right)^{1/s'} \leq \epsilon |\Sigma|^{1/s'}.$$

Also, taking $K \subset\subset \Omega$ to be the support of ψ we get, with K_Σ the extension of K by Σ , for any $t \in \Omega_\Sigma \setminus K_\Sigma$ that $\psi(t+x) = 0$ for any $x \in \Sigma$ and hence $\hat{\psi} = 0$ in $L^{s'}(\Sigma)$ and $\hat{\psi} \in C_c(\Omega_\Sigma, L^{s'}(\Sigma))$.

Now for $\phi \in C_0(\Omega_\Sigma)$, taking $(\phi_n)_n \in C_c(\Omega_\Sigma)$ to be a sequence converging to ϕ , we get that

$$\|\hat{\phi} - \hat{\phi}_n\|_{C_0(\Omega_\Sigma, L^{s'}(\Sigma))} = \sup_t \left(\int_\Sigma |\phi(t+x) - \phi_n(t+x)|^{s'} dx \right)^{1/s'} \leq \|\phi - \phi_n\|_\infty |\Sigma|^{1/s'} \rightarrow 0.$$

Thus $\hat{\phi}$ can be approximated by a sequence of compactly supported functions and hence $\hat{\phi} \in C_0(\Omega_\Sigma, L^{s'}(\Sigma))$. Fixing now $u = c \otimes p \in X_s$ we note that for any $\psi \in C_0(\Omega_\Sigma, L^{s'}(\Sigma))$, the function $t \rightarrow \int_\Sigma \psi(t)(x)p(x) dx$ is continuous, hence we can define the linear functional

$$F_u(\psi) := \int_{\Omega_\Sigma} \int_\Sigma \psi(t)(x)p(x) dx dc(t)$$

and get that F_u is continuous on $C_0(\Omega_\Sigma, L^{s'}(\Sigma))$. Then, since $\hat{\phi} \in C_0(\Omega_\Sigma) \otimes_i L^{s'}(\Sigma)$ it can be approximated by a sequence of simple tensors $(\sum_{i=1}^{m_n} x_i^n \otimes y_i^n)_n$ in the injective norm, which coincides with the norm in $C_0(\Omega_\Sigma, L^{s'}(\Sigma))$ and, using Lemma 21 in the appendix, we get

$$\begin{aligned} \langle u, \hat{\phi} \rangle &= \lim_n \langle u, \sum_{i=1}^{m_n} x_i^n \otimes y_i^n \rangle = \lim_n \sum_{i=1}^{m_n} \langle c, x_i^n \rangle \langle p, y_i^n \rangle = \lim_n \sum_{i=1}^{m_n} \int_{\Omega_\Sigma} \int_\Sigma x_i^n(a) y_i^n(b) p(b) db dc(a) \\ &= \lim_n F_u(\sum_{i=1}^{m_n} x_i^n \otimes y_i^n) = F_u(\hat{\phi}) = \int_{\Omega_\Sigma} \int_\Sigma \phi(a+b)p(b) db dc(a) = (K(c, p), \phi) = (\hat{K}u, \phi) \end{aligned}$$

Now by density of simple tensors in the projective tensor product, it follows that $K^*\phi = \hat{\phi}$. In order to show the continuity assertion, take $(u_n)_n$ weak* converging to some $u \in X_s$. Then by the previous assertion we get for any $\phi \in C_c(\Omega_\Sigma)$ that

$$(Ku_n, \phi) = \langle u_n, K^*\phi \rangle \rightarrow \langle u, K^*\phi \rangle = (Ku, \phi),$$

hence $(Ku_n)_n$ weakly converges to Ku on a dense subset of $L^{s'}(\Omega)$ which, together with boundedness of $(Ku_n)_n$, implies weak convergence. \square

We will also need weak*-to-weak* continuity of \hat{M} , which is shown in the following lemma in a slightly more general situation than needed.

Lemma 3. *Take $s \in [q, \infty)$ and assume that $M \in \mathcal{L}(L^s(\Sigma), Z)$ with Z a reflexive space and define $\hat{M} := I \otimes_\pi M \in \mathcal{L}(X_s, \mathcal{M}(\Omega_\Sigma) \otimes_\pi Z)$, where I is the identity on $\mathcal{M}(\Omega_\Sigma)$. Then \hat{M} is continuous w.r.t. weak star convergence in both spaces.*

Proof. Take $(u_n)_n \in X$ weak* converging to some $u \in X$ and write $u_n = \lim_k \sum_{i=1}^k x_i^n \otimes y_i^n$. We note that, since Z is reflexive, it satisfies in particular the Radon Nikodým property (see the appendix) and hence $C(\Omega_\Sigma) \otimes_i Z^*$ can be regarded as predual of $\mathcal{M}(\Omega_\Sigma) \otimes_\pi Z$ and we test with $\phi \otimes \psi \in C(\Omega_\Sigma) \otimes_i Z^*$. Then

$$\begin{aligned} \langle \hat{M}u_n, \phi \otimes \psi \rangle &= \lim_k \sum_{i=1}^k (x_i^n, \phi)(My_i^n, \psi) = \lim_k \sum_{i=1}^k (x_i^n, \phi)(y_i^n, M^*\psi) \\ &= \langle u_n, \phi \otimes M^*\psi \rangle \rightarrow \langle u, \phi \otimes M^*\psi \rangle = \langle \hat{M}u, \phi \otimes \psi \rangle, \end{aligned}$$

where the convergence follows since $M^*\psi \in L^{s'}(\Sigma)$, the predual of $L^s(\Sigma)$, and hence $\phi \otimes M^*\psi \in C(\Omega_\Sigma) \otimes_i L^{s'}(\Sigma)$. \square

Now we can obtain the convex, lower semi-continuous relaxation of $N_{cl,s}$.

Lemma 4. *With the assumptions of Lemma 1 and $s \in [q, \infty)$, the convex, l.s.c. relaxation of $N_{cl,s}$ is given as*

$$N_{cl,s}^{**}(u) = \inf_{C \in X_s} \|C\|_\pi \quad \text{s.t.} \quad \begin{cases} u = \hat{K}C \text{ in } \Omega, \\ \hat{M}C = 0. \end{cases} \quad (6)$$

Proof. Again we first compute the convex conjugate:

$$\begin{aligned} N_{cl,s}^*(v) &= \sup_u \langle u, v \rangle - N_{cl,s}(u) = \sup_{C \in X_s} \langle \hat{K}C, v \rangle - \|C\|_{\pi,k,M} \\ &= \sup_{C \in X_s} \langle C, \hat{K}^*v \rangle - \|C\|_{\pi,k,M} = \|\hat{K}^*v\|_{\pi,k,M}^*. \end{aligned}$$

Similarly, we see that $N^*(v) = \left(\|\cdot\|_\pi + \mathcal{I}_{\ker(\hat{M})} \right)^* (\hat{K}^*v)$, where

$$N(u) = \inf_{C \in X_s} \|C\|_\pi \quad \text{s.t.} \quad \begin{cases} u = \hat{K}C \text{ in } \Omega \\ \hat{M}C = 0 \end{cases}$$

Now in the proof of Lemma 1, we have in particular shown that $\left(\|\cdot\|_\pi + \mathcal{I}_{\ker(\hat{M})}\right)^* = \|\cdot\|_{\pi,k,M}^*$, hence, if we show that N is convex and lower semi-continuous, the assertion follows from $N(u) = N^{**}(u) = N_{\text{cl},s}^{**}(u)$. To this aim, take a sequence $(u_n)_n$ in $L^q(\Omega)$ converging weakly to some u for which, without loss of generality, we assume that

$$\lim_n N(u_n) = \liminf_n N(u_n) < \infty$$

Now with $(C_n)_n$ such that $\|C_n\|_\pi \leq N(u_n) + n^{-1}$, $\hat{M}C_n = 0$ and $u_n = \hat{K}C_n$ we get that $(\|C_n\|_\pi)_n$ is bounded. Since X_s admits a separable predual (see the appendix), this implies that $(C_n)_n$ admits a subsequence $(C_{n_i})_i$ weak* converging to some C . By weak* continuity of \hat{K} and \hat{M} we get that $u = \hat{K}C$ and $\hat{M}C = 0$, respectively, and by weak* lower semi-continuity of $\|\cdot\|_\pi$ it follows that

$$N(u) \leq \|C\|_\pi \leq \liminf_i \|C_{n_i}\|_\pi \leq \liminf_i N(u_{n_i}) + n_i^{-1} \leq \lim_i N(u_{n_i}) = \liminf N(u_n),$$

which concludes the proof. \square

This relaxation results suggest to use $N(\cdot)$ as in Equation (6) as convex texture prior in the continuous setting. There is, however, an issue with that, namely that such a functional cannot be expected to penalize the number of used atoms at all. Indeed, taking some $C = \sum_{i=1}^l c_i \otimes p_i$ and assume that $\|C\|_\pi = \sum_{i=1}^l \|c_i\|_{\mathcal{M}} \|p_i\|_s$. Now note that we can split any summand $c_{i_0} \otimes p_{i_0}$ as follows: Write $c_{i_0} = c_{i_0}^1 + c_{i_0}^2$ with disjoint support such that $\|c_{i_0}\|_{\mathcal{M}} = \|c_{i_0}^1\|_{\mathcal{M}} + \|c_{i_0}^2\|_{\mathcal{M}}$. Then we can re-write

$$c_{i_0} \otimes p_{i_0} = c_{i_0}^1 \otimes p_{i_0} + c_{i_0}^2 \otimes p_{i_0}$$

which gives a different representation of C by increasing the number of atoms without changing the cost of the projective norm. Hence, in order to maintain the original motivation of the approach to enforce a limited number of atoms, we need to add an additional penalty on C for the lifted texture prior.

3.3 Adding a rank penalization

Considering the discrete setting and the representation of the tensor C as a matrix, the number of used atoms corresponds to the rank of the matrix, for which it is well known that the nuclear norm constitutes a convex relaxation [20]. This construction can in principle also be transferred to general tensor products of Banach spaces via the identification (see Proposition 22 in the appendix)

$$C = \sum_{i=1}^{\infty} x_i \otimes y_i \in X^* \otimes_\pi Y^* \iff T_C \in \mathcal{L}(X, Y^*) \quad \text{where} \quad T_C(x) = \sum_{i=1}^{\infty} x_i(x) y_i$$

and the norm

$$\|C\|_{\text{nuc}} = \|T_C\|_{\text{nuc}} = \inf \left\{ \sum_{i=1}^{\infty} \sigma_i \mid T_C(x) = \sum_{i=1}^{\infty} \sigma_i x_i(x) y_i \text{ s.t. } \|x_i\|_{X^*} \leq 1, \|y_i\|_{Y^*} \leq 1 \right\}.$$

It is important to realize, however, that the nuclear norm of operators depends on the underlying spaces and in fact coincides with the projective norm in the tensor product space (see Proposition 22). Hence adding the nuclear norm in X_s does not change anything and, more generally, whenever

one of the underlying spaces is equipped with and L^1 -type norm we cannot expect a rank-penalizing effect (consider the example of the previous section).

On the other hand, going back to the nuclear norm of a matrix in the discrete setting, we see that it relies on orthogonality and an inner product structure and that the underlying norm is the Euclidean inner product norm. Hence an appropriate generalization of a rank penalizing nuclear norm needs to be built on a Hilbert space setting. Indeed, it is easy to see that any operator between Banach spaces with a finite nuclear norm is compact, and in particular for any $T \in \mathcal{L}(H_1, H_2)$ with finite nuclear norm and H_1, H_2 Hilbert spaces, there are orthonormal systems $(x_i)_i, (y_i)_i$ and uniquely defined singular values $(\sigma_i)_i$ such that

$$Tx = \sum_{i=1}^{\infty} \sigma_i(x, x_i)y_i \quad \text{and in addition} \quad \|T\|_{\text{nuc}} = \sum_{i=1}^{\infty} \sigma_i.$$

Motivated by this, we aim to define an L^2 -based nuclear norm as extended real valued function on X_s as convex surrogate of a rank penalization. To this aim, we consider from now on the case $s = 2$. Remember that the tensor product $X \otimes Y$ of two spaces X, Y is defined as the vector space spanned by linear mappings $x \otimes y$ on the space of bilinear forms on $X \times Y$, which are given as $x \otimes y(B) = B(x, y)$. Now since $L^2(\Omega_\Sigma)$ can be regarded as subspace of $\mathcal{M}(\Omega_\Sigma)$, also $L^2(\Omega_\Sigma) \otimes L^2(\Sigma)$ can be regarded as subspace of $\mathcal{M}(\Omega_\Sigma) \otimes L^2(\Sigma)$. Further, defining for $C \in L^2(\Omega_\Sigma) \otimes L^2(\Sigma)$,

$$\|C\|_{\pi, L^2 \otimes L^2} := \inf \left\{ \sum_{i=1}^n \|x_i\|_2 \|y_i\|_2 \mid C = \sum_{i=1}^n x_i \otimes y_i, x_i \in L^2(\Omega_\Sigma), y_i \in L^2(\Sigma), n \in \mathbb{N} \right\},$$

we get that $\|\cdot\|_\pi \leq B \|\cdot\|_{\pi, L^2 \otimes L^2}$ for a constant $B > 0$, hence also the completion $L^2(\Omega_\Sigma) \otimes_\pi L^2(\Sigma)$ can be regarded as subspace of $\mathcal{M}(\Omega_\Sigma) \otimes_\pi L^2(\Sigma)$. Further, $L^2(\Omega_\Sigma) \otimes_\pi L^2(\Sigma)$ can be identified with the space of nuclear operators $\mathcal{N}(L^2(\Omega_\Sigma), L^2(\Sigma))$ as above such that

$$\|C\|_{\pi, L^2 \otimes L^2} = \sum_{i=1}^{\infty} \sigma_i(T_C) \quad \text{with } (\sigma_i(T_C))_i \text{ the singular values of } T_C.$$

Using this, and introducing a potential function $\phi : [0, \infty) \rightarrow [0, \infty)$, we define for $C \in X_2$,

$$\|C\|_{\text{nuc}, \phi} := \begin{cases} \sum_{i=1}^{\infty} \phi(\sigma_i(T_C)) & \text{if } C \in L^2(\Omega_\Sigma) \otimes_\pi L^2(\Sigma), \\ \infty & \text{else.} \end{cases} \quad (7)$$

We will mostly focus on the case $\phi(x) = x$, in which $\|\cdot\|_{\text{nuc}, \phi}$ coincides with an extension of the nuclear norm and can be interpreted as convex relaxation of the rank. However, since we observed a significant improvement in some cases in practice by choosing ϕ to be a semi-convex potential function, i.e., a function such that $\phi + \tau|\cdot|^2$ is convex for τ sufficiently small, we include the more general situation in the theory.

Remark 5 (Sparsity and low-rank). *It is important to note that $\|C\|_{\text{nuc}, \phi} < \infty$ restricts C to be contained in the smoother space $L^2(\Omega_\Sigma) \otimes_\pi L^2(\Sigma)$ and in particular does not allow for simple tensors $\sum_{i=1}^k c_i \otimes p_i$ with the c_i 's being composed of delta peaks. Thus we observe some inconsistency of a rank penalization via the nuclear norm and a pointwise sparsity penalty, which is only visible in the continuous setting through the regularity of functions. Nevertheless, such an inconsistency*

has already been observed in the finite dimensional setting in the context of compressed sensing for low-rank AND sparse matrices, manifested via a poor performance of the sum of a nuclear norm and ℓ^1 norm for exact recovery (see [32]). As a result, there exists many literature on improved, convex priors for the recovery of low-rank and sparse matrices, see for instance [36, 16, 35]. While such improved priors can be expected to be highly beneficial for our setting, the question does not seem to be solved in such a way that that can be readily applied in our setting.

One direct way to circumvent this inconsistency would be to include an additional smoothing operator for C as follows: Take $S \in \mathcal{L}(\mathcal{M}(\Omega_\Sigma), \mathcal{M}(\Omega_\Sigma))$ such that $\text{range}(S) \subset L^2(\Omega_\Sigma)$ to be a weak* to weak* continuous linear operator and define the operator $\hat{S} : X_2 \rightarrow X_2$ as $\hat{S} := S \otimes I_{L^2}$, where I_{L^2} denotes the identity in $L^2(\Sigma)$. Then one could alternatively also use $\|SC\|_{\text{nuc}}$ as alternative for penalizing the rank of C while still allowing C to be a general measure. Indeed, in the discrete setting, by choosing S also to be injective, we even obtain the equality $\text{rank}(SC) = \text{rank}(C)$ (where we interpret C and SC as matrices). In practice, however, we did not observe an improvement by including such a smoothing and thus do not include \hat{S} in our model.

3.4 Well-posedness and a cartoon-texture model

Including $\|C\|_{\text{nuc},\phi}$ for $C \in X_2$ as additional penalty in our model, we ultimately arrive at the following variational texture prior in the tensor product space $X_2 := \mathcal{M}(\Omega_\Sigma) \otimes_\pi L^2(\Sigma)$, which is convex whenever ϕ is convex, in particular for $\phi(x) = |x|$.

$$N_\nu(v) = \inf_{C \in X} \nu \|C\|_\pi + (1 - \nu) \|C\|_{\text{nuc},\phi} \quad \text{s.t.} \quad \begin{cases} \hat{M}C = 0, \\ v = \hat{K}C \text{ in } \Omega, \end{cases} \quad (8)$$

where $\nu \in (0, 1)$ is a parameter balancing the sparsity and the rank penalty.

In order to employ N_ν as a regularization term in an inverse problems setting, we need to obtain some lower-semi continuity and coercivity properties. As first step, the following lemma, which is partially inspired by techniques used in [8, Lemma 3.2], shows that, under some weak conditions on ϕ , $\|\cdot\|_{\text{nuc},\phi}$ defines a weak* lower semi-continuous function on X_2 .

Lemma 6. *Assume that $\phi : [0, \infty) \rightarrow [0, \infty)$ is lower semi-continuous, non-decreasing, that*

- $\phi(x) \rightarrow \infty$ for $x \rightarrow \infty$ and that
- there exist $\epsilon, \eta > 0$ such that $\phi(x) \geq \eta x$ for $0 \leq x < \epsilon$.

Then the functional $\|\cdot\|_{\text{nuc},\phi} : X_2 \rightarrow \bar{\mathbb{R}}$ defined as in (7) is lower semi-continuous w.r.t. weak convergence in X_2 .*

Proof. Take $(C_n)_n \subset X_2$ weak* converging to some $C \in X_2$ for which, w.l.o.g., we assume that

$$\liminf_n \|C_n\|_{\text{nuc},\phi} = \lim_n \|C_n\|_{\text{nuc},\phi}.$$

We only need to consider the case that $(\|C_n\|_{\text{nuc},\phi})_n$ is bounded, otherwise the assertion follows trivially. Hence we can write $T_{C_n}(x) = \sum_{i=1}^{\infty} \sigma_i^n(x_i^n, x) y_i^n$ such that $\|C_n\|_{\text{nuc},\phi} = \sum_{i=1}^{\infty} \phi(\sigma_i^n)$. Now we aim to bound $(\|C_n\|_{\pi, L^2 \otimes L^2})_n$ in terms of $(\|C_n\|_{\text{nuc},\phi})_n$. To this aim, first note that the assumptions in ϕ imply that for any $\epsilon' > 0$ there is $\eta' > 0$ such that $\phi(x) \geq \eta' x$ for all $x < \epsilon'$. Also,

$\phi(\sigma_i^n) \leq \|C_n\|_{\text{nuc},\phi}$ for any i, n and via a direct contradiction argument it follows that there exists $\hat{\epsilon} > 0$ such that $\sigma_i^n < \hat{\epsilon}$ for all i, n . Picking $\hat{\eta}$ such that $\phi(x) \geq \hat{\eta}x$ for all $x < \hat{\epsilon}$ we obtain

$$\|C_n\|_{\text{nuc},\phi} = \sum_{i=1}^{\infty} \phi(\sigma_i^n) \geq \hat{\eta} \sum_{i=1}^{\infty} \sigma_i^n = \hat{\eta} \|C_n\|_{\pi, L^2 \otimes L^2},$$

hence $(C_n)_n$ is also bounded as sequence in $L^2(\Omega_\Sigma) \otimes_\pi L^2(\Sigma)$ and admits a (non-re-labeled) subsequence weak* converging to some $\hat{C} \in L^2(\Omega_\Sigma) \otimes_\pi L^2(\Sigma)$, with $L^2(\Omega_\Sigma) \otimes_i L^2(\Sigma)$ being the predual space. By the inclusion $C_0(\Omega_\Sigma) \otimes_i L^2(\Sigma) \subset L^2(\Omega_\Sigma) \otimes_i L^2(\Sigma)$ and uniqueness of the weak* limit, we finally get $\hat{C} = C \in L^2(\Omega_\Sigma) \otimes_i L^2(\Sigma)$ and can write $T_C x = \sum_{i=1}^{\infty} \sigma_i(x_i, x)y_i$ and $\|C\|_{\text{nuc},\phi} = \sum_{i=1}^{\infty} \sigma_i$. By lower semi-continuity of $\|\cdot\|_{\text{nuc}}$ this would suffice to conclude in the case $\phi(x) = x$. For the more general case, we need to show a point-wise lim-inf property of the singular values. To this aim, note that by the Courant-Fischer min-max principle (see for instance [10, Problem 37]) for any compact operator $T \in \mathcal{L}(H_1, H_2)$ with H_1, H_2 Hilbert spaces and λ_k the k -th singular value of T sorted in descending order we have

$$\lambda_k = \sup_{\dim(V)=k} \min_{x \in V, \|x\|=1} \|Tx\|_{H_2}.$$

Now consider $k \in \mathbb{N}$ fixed. For any subspace V with $\dim(V) = k$, the minimum in the equation above is achieved, hence we can denote x_V to be a minimizer and define $F_V(T) := \|Tx_V\|_{H_2}$ such that $\lambda_k = \sup_{\dim(V)=k} F_V(T)$. Since weak* convergence of a sequence (T_n) to T in $L^2(\Omega_\Sigma) \otimes_\pi L^2(\Sigma)$ implies in particular $T_n(x) \rightarrow T(x)$ for all x , by lower semi-continuity of the norm $\|\cdot\|_{H_2}$ it follows that F_V is lower semi-continuous with respect to weak* convergence. Hence this is also true for the function $T \mapsto \lambda_k(T)$ is being the pointwise supremum of a family of lower semi-continuous functional. Consequently, for the sequence $(T_{C_n})_n$ it follows that $\sigma_k \leq \liminf_n \sigma_k^n$. Finally, by monotonicity and lower semi-continuity of ϕ and Fatou's lemma we conclude

$$\|T_C\|_{\text{nuc},\phi} = \sum_k \phi(\sigma_k) \leq \sum_k \phi(\liminf_n \sigma_k^n) = \sum_k \liminf_n \phi(\sigma_k^n) \leq \liminf_n \sum_k \phi(\sigma_k^n) \leq \liminf_n \|T_{C_n}\|_{\text{nuc},\phi}.$$

□

The lemma below now establishes the main properties of N_ν that in particular allow to employ it as regularization term in an inverse problems setting.

Lemma 7. *The infimum in the definition of (8) is attained and $N_\nu : L^q(\Omega) \rightarrow \overline{\mathbb{R}}$ is convex and lower semi-continuous. Further, any sequence $(v_n)_n$ such that $N_\nu(v_n)$ is bounded admits a subsequence converging weakly in $L^q(\Omega)$.*

Proof. The proof is quite standard, but we provide it for the readers convenience. Take $(v_n)_n$ to be a sequence such that $(N_\nu(v_n))_n$ is bounded. Then we can pick a sequence $(C_n)_n$ in X_2 such that $\hat{M}C_n = 0$, $v_n = \hat{K}C_n$ and

$$\nu \|C_n\|_\pi \leq \nu \|C_n\|_\pi + (1 - \nu) \|C_n\|_{\text{nuc},\phi} \leq N_\nu(v_n) + n^{-1}$$

This implies that $(C_n)_n$ admits a subsequence $(C_{n_i})_i$ weak* converging to some $C \in X_2$. Now by continuity of \hat{M} and \hat{K} we have that $\hat{M}C = 0$ and that $(v_{n_i})_i = (\hat{K}C_{n_i})_i$ is bounded. Hence also $(v_{n_i})_i$ admits a (non-re-labeled) subsequence converging weakly to some $v = \hat{K}C$. This already

shows the last assertion. In order to show lower semi-continuity, assume that $(v_n)_n$ converges to some v and, without loss of generality, that

$$\liminf_n N_\nu(v_n) = \lim_n N_\nu(v_n).$$

Now this is a particular case of the argumentation above, hence we can deduce with $(C_n)_n$ as above that

$$\begin{aligned} N_\nu(v) &\leq \nu \|C\|_\pi + (1 - \nu) \|C\|_{\text{nuc}, \phi} \leq \liminf_i \nu \|C_{n_i}\|_\pi + (1 - \nu) \|C_{n_i}\|_{\text{nuc}, \phi} \\ &\leq \liminf_i N_\nu(v_{n_i}) + n_i^{-1} = \liminf_n N_\nu(v_n) \end{aligned}$$

which implies lower semi-continuity. Finally, specializing even more to the case that $(v_n)_n$ is the constant sequence $(v)_n$, also the claimed existence follows. \square

In order to model a large class of natural images and to keep the number of atoms needed in the above texture prior low, we combine it with a second part that models cartoon-like images. Doing so, we arrive at the following model

$$\min_{u, v \in L^q(\Omega)} \lambda D(Au, f_0) + s_1(\mu) R(u - v) + s_2(\mu) N_\nu(v) \quad (\text{P})$$

where we assume R to be a functional that models cartoon images, $D(\cdot, f_0) : Y \rightarrow \overline{\mathbb{R}}$ is a given data discrepancy, $A \in \mathcal{L}(L^q(\Omega), Y)$ a forward model and we define the parameter balancing function

$$s_1(\mu) = 1 - \min(\mu, 0), \quad s_2(\mu) = 1 + \max(\mu, 0). \quad (9)$$

Now we get the following general existence result.

Proposition 8. *Assume that $R : L^q(\Omega) \rightarrow \overline{\mathbb{R}}$ is convex, lower semi-continuous and that there exists a finite dimensional subspace $U \subset L^q(\Omega)$ such that for any $u \in L^q(\Omega)$, $v \in U^\perp$, $w \in U$,*

$$\|v\|_q \leq BR(v), \quad \text{and} \quad R(u + w) = R(u)$$

with $B > 0$ and U^\perp denoting the complement of U in $L^q(\Omega)$. Further assume that $A \in \mathcal{L}(L^q(\Omega), Y)$, $D(\cdot, f_0)$ is convex, lower semi-continuous and coercive on the finite dimensional space $A(U)$ in the sense that for any two sequences $(u_n^1)_n$, $(u_n^2)_n$ such that $(u_n^1)_n \subset U$, $(u_n^2)_n$ is bounded and $(D(A(u_n^1 + u_n^2), f_0))_n$ is bounded, also $(\|Au_n^1\|_q)_n$ is bounded. Then there exists a solution to (P).

Remark 9. *Note that for instance in case D satisfies a triangle inequality, the sequence (u_n^2) in the coercivity assumption is not needed, i.e., can be chosen to be zero.*

Proof. The proof is rather standard and we provide only a short sketch. Take $((u_n, v_n))_n$ a minimizing sequence for (P). From Lemma 7 we get that $(v_n)_n$ admits a (non-re-labeled) weakly convergent subsequence. Now we split $u_n = u_n^1 + u_n^2 \in U + U^\perp$ and $v_n = v_n^1 + v_n^2 \in U + U^\perp$ and by assumption get that $\|u_n^2 - v_n^2\|_q$ is bounded. But since $(\|v_n\|_q)_n$ is bounded, so is $(\|v_n^2\|_q)_n$ and consequently also $(\|u_n^2\|_q)_n$. Now we split again $u_n^1 = u_n^{1,1} + u_n^{1,2} \in \ker(A) \cap U + (\ker(A) \cap U)^\perp$ and note that also $(u_n^{1,2} + u_n^2, v_n)$ is a minimizing sequence for (P). Hence it remains to show that $(u_n^{1,2})_n$ is bounded in order to get a bounded minimizing sequence. To this aim, we note that $(u_n^{1,2})_n \subset (\ker(A) \cap U)^\perp \cap U$ and that A is injective on this finite dimensional space. Hence $\|u_n^{1,2}\|_q \leq B \|Au_n^{1,2}\|_q$ for some $B > 0$ and by the coercivity assumption on the data term we finally get that $(\|u_n^{1,2}\|_q)_n$ is bounded. Hence also $(u_n^{1,2} + u_n^2)_n$ admits a weakly convergent subsequence in $L^q(\Omega)$ and by continuity of A as well as lower semi-continuity of all involved functionals existence of a solution follows. \square

Remark 10 (Choice of regularization). *A particular choice of regularization for R in (P) that we consider in this paper is $R = \text{TGV}_\alpha^2$, with TGV_α^2 the second order total generalized variation functional [7], $q \leq d/(d-1)$ and*

$$Mp := \left(\int_\Sigma p(x) \, dx, \int_\Sigma p(x_1, x_2) x_1 \, dx, \int_\Sigma p(x_1, x_2) x_2 \, dx \right).$$

Since in this case [9, 5]

$$\|u\|_q \leq B \text{TGV}_\alpha^2(u)$$

with $B > 0$ and for all $u \in \mathcal{P}_1(\Omega)^\perp$, the complement of the first order polynomials, and TGV_α^2 is invariant on first order polynomials, the result of Proposition 8 applies.

Remark 11 (Norm-type data terms). *We also note that the result of Proposition 8 in particular applies to $D(w, f_0) := \frac{1}{r} \|w - f_0\|_r^r$ for any $r \in [1, \infty)$ or $D(w, f_0) := \|w - f\|_\infty$, where we extend the norms by infinity to $L^q(\Omega)$ whenever necessary. Indeed, lower semi-continuity of these norms is immediate for both $r \leq q$ and $r > q$ and since the coercivity is only required on a finite dimensional space, it also holds by equivalence of norms.*

Remark 12 (Inpainting). *At last we also remark that the assumptions of Proposition 8 also hold for an inpainting data term defined as*

$$D(w, f_0) := \begin{cases} 0 & \text{if } w = f_0 \text{ a.e. on } \omega \subset \Omega \\ \infty & \text{else,} \end{cases}$$

whenever ω has nonempty interior. Indeed, lower semi-continuity follows from the fact that L^q convergent sequences admit pointwise convergent subsequences and the coercivity follows from finite dimensionality of U and the fact that ω has non-empty interior.

Remark 13 (Regularization in a general setting). *We also note that Lemma 7 provides the basis for employing either N_r directly or its infimal-convolution with a suitable cartoon prior as in Proposition 8 for the regularization of general (potentially non-linear) inverse problems and with multiple data fidelities, see for instance [24, 25] for general results in that direction.*

4 The model in a discrete setting

This section deals with the discretization of the proposed model and its numerical solution. For the sake of brevity, we provide only the main steps and refer to the publicly available source code [13] for all details.

We define $U = \mathbb{R}^{N \times M}$ to be the space of discrete grayscale images, $W = \mathbb{R}^{(N+n-1) \times (M+n-1)}$ to be the space of coefficient images and $Z = \mathbb{R}^{n \times n}$ to be the space of image atoms for which we assume $n < \min\{N, M\}$ and, for simplicity, only consider a square domain for the atoms. The tensor product of a coefficient image $c \in W$ and an atom $p \in Z$ is given as $(c \otimes p)_{i,j,r,s} = c_{i,j} p_{r,s}$ and the lifted tensor space is given as the four-dimensional space $X = \mathbb{R}^{(N+n-1) \times (M+n-1) \times n \times n}$.

Texture norm. The forward operator K being the lifting of the convolution $c * p$ and mapping lifted matrices to the vectorized image space is then given as

$$(KC)_{i,j} = \sum_{r,s=1}^{n,n} C_{i+n-r, j+n-s, r, s}$$

and we refer Figure 1 for a visualization in the one dimensional case. Note that by extending the first two dimensions of the tensor space to $N + n - 1, M + n - 1$ we allow to place an atom at any position where it still effects the image, also partially outside the image boundary.

Also we note that, in order to reduce dimensionality and accelerate the computation, we introduce a stride parameter $\eta \in \mathbb{N}$ in practice which introduces a stride on the possible atom positions. That is, the lifted tensor space and forward operator is reduced in such a way that the grid of possible atom positions in the image is essentially $\{(\eta i, \eta j) \mid i, j \in \mathbb{N}, (\eta i, \eta j) \in \{1, \dots, N\} \times \{1, \dots, M\}\}$. This reduces the dimension of the tensor space by a factor η^{-2} , while for $\eta > 1$ it naturally does not allow for arbitrary atom positions anymore and for $\eta = n$ it corresponds to only allowing non-overlapping atoms. In order to allow for atoms being placed next to each other it is important to choose η to be a divisor of the atom-domain size n and we used $n = 15$ and $\eta = 3$ in all experiments of the paper. In order to avoid extensive indexing and case distinctions, however, we only consider the case $\eta = 1$ here and refer to the source code [13] for the general case.

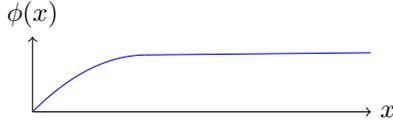
As a straightforward computation shows that, in the discrete lifted tensor space, the projective norm corresponding to discrete $\|\cdot\|_1$ and $\|\cdot\|_2$ norms for the coefficient images and atoms, respectively, is given as a mixed 1-2 norm as

$$\|A\|_\pi = \|A\|_{1,2} = \sum_{i,j=1}^{N,M} \sqrt{\sum_{r,s=1}^{n,n} A_{i,j,r,s}^2}.$$

The nuclear norm for a potential ϕ on the other hand reduces to the evaluation of ϕ on the singular values of a matrix-reshaping of the lifted tensors and is given as

$$\|A\|_{\text{nuc},\phi} = \sum_{i=1}^{nn} \phi(\sigma_i), \quad \text{with } (\sigma_i)_i \text{ the singular values of } B = [A_{(NM,nn)}].$$

where $[A_{(NM,nn)}]$ denotes a reshaping of the tensor A to a matrix of dimensions $NM \times nn$. For the potential function ϕ we consider two choices: Mostly we are interested in $\phi(x) = x$ which yields a convex texture model and enforces sparsity of the singular values. A second choice we consider is $\phi : [0, \infty) \rightarrow [0, \infty)$ given as



$$\phi(x) = \begin{cases} x - \epsilon \delta x^2 & x \in [0, \frac{1}{2\epsilon}] \\ (1 - \delta)x + \frac{\delta}{4\epsilon} & \text{else,} \end{cases} \quad (10)$$

where $\delta < 1, \delta \approx 1$ and $\epsilon > 0$. It is easy to see that ϕ fulfills the assumptions of Lemma 6 and that ϕ is semi-convex, i.e., $\phi + \rho|\cdot|^2$ is convex for $\rho > \delta\epsilon$. While the results of Section 3 hold for this setting even without the semi-convexity assumption, we can in general not expect to obtain an algorithm that provably delivers a globally optimal solution in the semi-convex (or generally non-convex) case. The reason for using a semi-convex potential rather than an arbitrary non-convex one is twofold: First, for a suitably small stepsize τ the proximal mapping

$$\text{prox}_{\tau,\phi}(\hat{u}) = \underset{u}{\text{argmin}} \frac{\|u - \hat{u}\|_2^2}{2\tau} + \phi(u)$$

is well defined and hence proximal-point type algorithms are applicable at least conceptually. Second, since we employ ϕ on the singular values of the lifted matrices A , it will be important for

Algorithm 1 Scheme of implementation for the numerical solution of (DP)

```

1: function CART_TEXT_RECON( $f_0$ )
2:   ( $u, v, C$ )  $\leftarrow$  (0, 0, 0), ( $p, q, d, r, m$ )  $\leftarrow$  (0, 0, 0, 0, 0)
3:   choose  $\sigma, \tau > 0$ 
4:   repeat
5:     Dual updates
6:      $p \leftarrow \text{proj}_{\alpha_1}(p + \sigma(\nabla(\bar{u} - K\bar{C}) - \bar{v}))$ 
7:      $q \leftarrow \text{proj}_{\alpha_0}(q + \sigma\mathcal{E}\bar{v})$ 
8:      $d \leftarrow \text{prox}_{\sigma, (\lambda D_1(\cdot, f_0))^*}(d + \sigma A\bar{u})$ 
9:      $r \leftarrow \text{prox}_{\sigma, (s_2(\mu)\nu\|\cdot\|_{1,2})^*}(r + \sigma\bar{C})$ 
10:     $m \leftarrow (m + \sigma M\bar{C})$ 
11:    Primal updates
12:     $u_+ \leftarrow \text{prox}_{\tau, (\lambda D_2(\cdot, f_0))^*}(u - \tau(\nabla^*p + A^*d))$ 
13:     $v_+ \leftarrow v - \tau(-p + \mathcal{E}^*q)$ 
14:     $C_+ \leftarrow \text{prox}_{\tau, s_2(\mu)(1-\nu)\|\cdot\|_{\text{nuc}, \phi}}(C - \tau(-K^*\nabla^*p + r + M^*m))$ 
15:    Extrapolation and update
16:     $(\bar{u}, \bar{v}, \bar{C}) \leftarrow 2(u_+, v_+, C_+) - (u, v, C)$ ,
17:     $(u, v, C) \leftarrow (u_+, v_+, C_+)$ 
18:  until Stopping criterion fulfilled
19:  return ( $u_+, KC_+, \text{RSV}(C_+)$ )
20: end function

```

Here, the dual variables $(p, q, d, r, m) \in (U^2, U^3, A(U), U, U^3)$ are in the image space of the corresponding operators, $\mathcal{I}_S(z) = 0$ if $z \in S$ and $\mathcal{I}_S(z) = \infty$ else, $\{\|\cdot\|_\infty \leq \delta\} := \{z \mid \|z\|_\infty \leq \delta\}$ with $\|z\|_\infty = \|(z_1, \dots, z_l)\|_\infty = \sup_{i,j} \sqrt{\sum_{s=1}^l (z_{i,j}^s)^2}$ a point-wise infinity norm on $z \in U^l$. The operator B and the functional G are given as

$$B(u, v, C) = (\nabla u - \nabla K C - v, \mathcal{E}v, Au, C, MC), \quad G(x) = G(u, v, C) = \lambda D_2(u, f_0) + s_2(\mu)(1-\nu)\|C\|_{\text{nuc}, \phi}$$

and $F^*(y) = F^*(p, q, d, r, m)$ summarizes all the conjugate functionals as above. Applying the algorithm of [14] to this reformulation yields the numerical scheme as in Algorithm 1.

Note that there, we set either $D_1(\cdot, f_0) \equiv 0$ such that the dual variable d is constant 0 and line 8 of the algorithm can be skipped, or we set $D_2(\cdot, f_0) \equiv 0$ such that the proximal mapping in line 12 reduces to the identity. The concrete choice of D_1, D_2 and the proximal mappings will be given in the corresponding experimental sections. All other proximal mappings can be computed explicitly and reasonably fast: The mappings proj_{α_1} and proj_{α_0} can be computed as pointwise projections to the L^∞ -ball (see for instance [6]) and the mapping $\text{prox}_{\sigma, (s_2(\mu)\nu\|\cdot\|_{1,2})^*}$ is a similar projection given as

$$\text{prox}_{\sigma, (s_2(\mu)\nu\|\cdot\|_{1,2})^*}(R)_{i,j,l,s} = R_{i,j,l,s} / \left(\max \left\{ 1, \left(\sum_{l,s=1}^{n,n} R_{i,j,l,s}^2 \right)^{1/2} / (s_2(\mu)\nu) \right\} \right).$$

Most of the computational effort lies in the computation of $\text{prox}_{\tau, s_2(\mu)(1-\nu)\|\cdot\|_{\text{nuc}, \phi}}$, which, as the following lemma shows, can be computed via an SVD and a proximal mapping on the singular values.

Lemma 14. *Let $\phi : [0, \infty) \rightarrow [0, \infty)$ be a differentiable and increasing function and $\tau, \rho > 0$ be such that $x \mapsto \frac{x^2}{2\tau} + \rho\phi(x)$ is convex on $[0, \infty)$. Then the proximal mapping of $\rho\|\cdot\|_{\text{nuc}, \phi}$ for parameter τ is given as*

$$\text{prox}_{\tau, \rho\|\cdot\|_{\text{nuc}, \phi}}(A) = [(U \text{diag}((\text{prox}_{\tau, \rho\phi}(\sigma_i))_i) V^T)_{((N, M, n, n))}]$$

where $[A_{(NM, nn)}] = U\Sigma V^T$ is the SVD of $[A_{(NM, nn)}]$ and for $x_0 \geq 0$

$$\text{prox}_{\tau, \rho\phi}(x_0) = \min_x \frac{|x - x_0|^2}{2\tau} + \rho\phi(|x|).$$

In particular, in case $\phi(x) = x$ we have

$$\text{prox}_{\tau, \rho\phi}(x_0) = \begin{cases} 0 & \text{if } 0 \leq x_0 \leq \tau\rho, \\ x_0 - \tau\rho & \text{else,} \end{cases}$$

and in case

$$\phi(x) = \begin{cases} x - \epsilon\delta x^2 & \text{if } x \in [0, \frac{1}{2\epsilon}], \\ (1 - \delta)x + \frac{\delta}{4\epsilon} & \text{else,} \end{cases}$$

we have that that $x \mapsto \frac{x^2}{2\tau} + \rho\phi(x)$ is convex whenever $\tau \leq \frac{1}{2\epsilon\delta\rho}$ and in this case

$$\text{prox}_{\tau, \rho\phi}(x_0) = \begin{cases} 0 & \text{if } 0 \leq x_0 \leq \tau\rho, \\ \frac{x_0 - \tau\rho}{1 - 2\epsilon\delta\tau\rho} & \text{if } \tau\rho < x_0 \leq \frac{1}{2\epsilon} + \tau\rho(1 - \delta), \\ x_0 - \tau\rho(1 - \delta) & \text{if } \frac{1}{2\epsilon} + \tau\rho(1 - \delta) < x_0. \end{cases}$$

Proof. At first note that it suffices to consider $\rho\|\cdot\|_{\text{nuc}, \phi}$ as a function on matrices and show the assertion without the reshaping operation. For any matrix A , we denote by $A = U_A \Sigma_A V_A^T$ the SVD of A and $\Sigma_A = \text{diag}((\sigma_i^A)_i)$ contains the singular values sorted in non-increasing order, where Σ_A is uniquely determined by A and U_A, V_A are chosen to be suitable orthonormal matrices.

We first show that $G(A) := \frac{\|A\|_2^2}{2\tau} + \rho\|A\|_{\text{nuc}, \phi}$ is convex. For $\lambda \in [0, 1]$, A, B matrices we get by sub-additivity of the singular values (see for instance [41]) that

$$\begin{aligned} G(\lambda A + (1 - \lambda)B) &= \sum_i \frac{1}{2\tau} (\sigma_i^{\lambda A + (1-\lambda)B})^2 + \rho\phi(\sigma_i^{\lambda A + (1-\lambda)B}) \\ &\leq \sum_i \frac{1}{2\tau} (\lambda\sigma_i^A + (1 - \lambda)\sigma_i^B)^2 + \rho\phi(\lambda\sigma_i^A + (1 - \lambda)\sigma_i^B) \\ &\leq \sum_i \frac{\lambda}{2\tau} (\sigma_i^A)^2 + \frac{1 - \lambda}{2\tau} (\sigma_i^B)^2 + \rho\lambda\phi(\sigma_i^A) + \rho(1 - \lambda)\phi(\sigma_i^B) \\ &\leq \lambda G(A) + (1 - \lambda)G(B). \end{aligned}$$

Now with $H(A) := \frac{\|A - A_0\|_2^2}{2\tau} + \rho\|A\|_{\text{nuc}, \phi}$ we get that $H(A) = G(A) - \frac{1}{2\tau}(2\langle A, A_0 \rangle + \|A_0\|_2^2)$, thus also H is convex. Hence first order optimality conditions are necessary and sufficient and

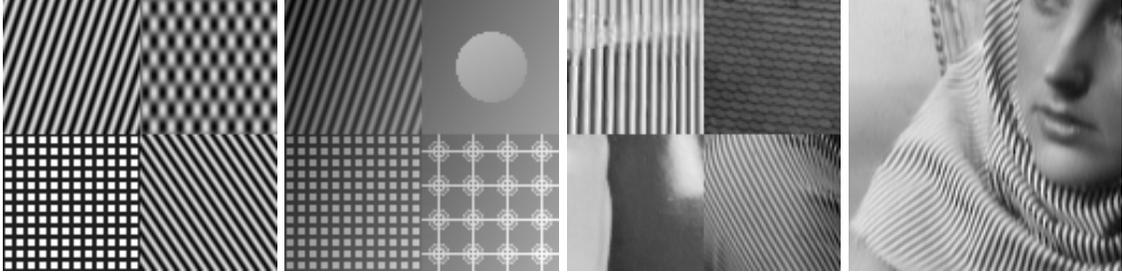


Figure 3: Different test images we will refer to as: Texture, Patches, Mix, Barbara

we get (using the derivative of the singular values as in [33]) with DH the derivative of H that $A = \text{prox}_{\tau, \rho \|\cdot\|_{\text{nuc}, \phi}}(A_0)$ is equivalent to

$$\begin{aligned} 0 &= DH(A) = (A - A_0) + \tau \rho U_A \text{diag}((\phi'(\sigma_i^A))_i) V_A^T \\ &= -A_0 + U_A (\Sigma_A + \tau \rho \text{diag}((\phi'(\sigma_i^A))_i)) V_A^T \end{aligned}$$

and consequently to

$$\sigma_i^{A_0} = \sigma_i^A + \tau \rho \phi'(\sigma_i^A)$$

which is equivalent to

$$\sigma_i^{A_0} = \text{prox}_{\tau, \rho \phi}(\sigma_i^{A_0})$$

as claimed. The other results follow by direct computation. \square

Note also that, in Algorithm 1, KC_+ returns the part of the images that is represented by the atoms (the “texture part”) and $\text{RSV}(C_+)$ stand for right-singular values of $[(C_+)_{(NM, nn)}]$ and returns the image atoms. For the sake of simplicity, we use a rather high, fixed number of iterations in all experiment but note that, alternatively, a duality-gap based stopping criterion (see for instance [6]) could be used.

5 Numerical results

In this section we present numerical results obtained with the proposed method as well as its variants and compare to existing methods. We will mostly focus on the setting of (DP), where $\phi(x) = |x|$ and we use different data terms D . Hence, the regularization term is convex and consists of TGV_α^2 for the cartoon part and a weighted sum of a nuclear norm and $\ell^{1,2}$ norm for the texture part. Besides this choice of regularization (called *CT-cvx*), we will compare to pure TGV_α^2 regularization (called *TGV*), the setting of (DP) with the semi-convex potential ϕ as in (10) (call *CT-scvx*) and the setting of (DP) with TGV replaced by $\mathcal{L}_{\{0\}}$, i.e., only the texture norm is used for regularization, and $\phi(x) = |x|$ (called *TXT*). Further, in the last subsection, we also compare to other methods as specified there. For *CT-cvx* and *CT-scvx* we use the algorithm described in the previous section (where convergence can only be ensured for *CT-cvx*) and for the other variants we use a adaption of the algorithm to the respective special case.

We fix the size of the atom domain to 15×15 pixel and the stride to 3 pixel (see Section 4) for all experiments, and use four different test images (see Figure 3): The first two are synthetic

images of size 120×120 , containing four different blocks of size 60×60 , whose size is a multiple of the chosen atom-domain size. The third and fourth image have size 128×128 (not being a multiple of the atom-domain size) and the third image contains four sections of real images of size 64×64 each (again not a multiple of the atom-domain size). All but the first image contain a mixture of texture and cartoon parts. The first four subsections consider only convex variants of our method ($\phi(x) = |x|$) and the last one considers the improvement obtained with a non-convex potential ϕ and also compares to other approaches.

Regarding the choice of parameters for all methods, we generally aimed to reduce the number of varying parameters for each method as much as possible such that for each method and type of experiment, at most two parameters need to be optimized. Whenever we incorporate the second order TGV functional for the cartoon part, we fix the parameters (α_0, α_1) to $(\sqrt{2}, 1)$. The method *CT-cvx* then essentially depends on the three parameters λ, μ, ν . We experienced that the choice of ν is rather independent of the data and type of experiments, hence we leave it fixed for all experiments with incomplete or corrupted data, leaving our method with two parameters to be adapted: λ defining the tradeoff between data and regularization and μ defining the tradeoff between cartoon and texture regularization. For the semi-convex potential we choose ν as with the convex one, fix $\delta = 0.99$ and use two different choice of ϵ , depending on the type of experiment, hence again leaving two parameters to be adapted. A summary of the parameter choice for all methods is provided in Table 11 below.

We also note that, whenever we tested a range of different parameters for any method presented below, we show the visually best results in the figure. Those are generally not the ones delivering the best result in terms of peak-signal-to-noise ratio, and for the sake of completeness we also provide in Table 10 the best PSNR result obtained with each method and each experiment over the range of tested parameters.

5.1 Image-atom learning and texture separation

As first experiment we test the method *CT-cvx* for learning image atoms and texture separation directly on the ground truth images. To this aim, we use

$$D_1 \equiv 0, \quad D_2(u, f_0) = \mathcal{I}_{\{0\}}(u - f_0).$$

The results can be found in Figure 4, where for the pure texture image we used only the texture norm (i.e. the method *TXT*) without the TGV part for regularization.

It can be observed that the proposed method achieves a good decomposition of cartoon and texture and also is able to learn the most important image structure effectively. While there are some repetitions of shifted structures in the atoms, the different structures are rather well-separated and the first nine atoms corresponding to the nine largest singular values still contain the most important features of the texture parts.

5.2 Inpainting and leaning from incomplete data

This section deals with the task of inpainting a partially available image and learning image atoms from this incomplete data. For reference, we also provide results with pure TGV_α^2 regularization (the method *TGV*). The data fidelity in this case is

$$D_1 \equiv 0, \quad D_2(u, f_0) = \mathcal{I}_{\{v \mid v_{i,j} = (f_0)_{i,j} \text{ for } (i,j) \in \mathcal{M}\}}(u),$$

with \mathcal{M} the index set of known pixels. Again we use only the texture norm for the first image (the method *TXT*) and the cartoon-texture functional for the others.

The results can be found in Figure 5. For the first and third image, 20% of the pixels were given while for the other two, 30% were given. It can be seen that our method is generally still able to identify the underlying pattern of the texture part and to reconstruct it reasonably well. Also the learned atoms are reasonable and are in accordance with the ones learned from the full data as in the previous section. In contrast to that, pure TGV regularization (which assumes piecewise smoothness) has no chance to reconstruct the texture patterns. For the cartoon part, both methods are comparable. It can also be observed that the target-like structure in the bottom right of the second image is not reconstructed well and also not well identified with the atoms (only the 8th one contains parts of this structure). The reason might be that due to the size of the repeating structure there is not enough redundant information available to reconstruct it from the missing data. Concerning the optimal PSNR values of Table 10, we can observe a rather strong improvement with *CT-cvx* compared to *TGV*.

5.3 Learning and separation under noise

In this section we test our method for image-atom-learning and de-noising with data corrupted by Gaussian noise (with standard deviation 0.5 and 0.1 times the image range for the Texture and the other images, respectively). Again we compare to *TGV* regularization in this section (but also to other methods in Section 5.5 below) and use the texture norm for the first image (the method *TXT*). The data fidelity in this case is

$$D_1 \equiv 0, \quad D_2(u, f_0) = \frac{1}{2} \|u - f_0\|_2^2.$$

It can be observed that also under the presence of rather strong noise, our method is able to learn some of the main features of the image within the learned atoms. Also the quality of the reconstructed image is improved compared to *TGV*, in particular for the right-hand side of the *Mix* image, where the top left structure is only visible in the result obtained with *CT-cvx*. On the other hand, the circle of the *Patches* image obtained with *CT-cvx* contains some artifacts of the texture part. Regarding the optimal PSNR values of Table 10, the improvement with *CT-cvx* compared to *TGV* is still rather significant.

5.4 Deconvolution

This section deals with the learning of image features and image reconstruction in an inverse problem setting, where the forward operator is given as a convolution with a Gaussian kernel (standard deviation 0.2, kernel size 9×9 pixels) and the data is degraded by Gaussian noise with standard deviation 0.05 times the image range. The data fidelity in this case is

$$D_1(u, f_0) = \frac{1}{2} \|Au - f_0\|_2^2, \quad D_2 \equiv 0,$$

with A being the convolution operator.

We show results for the *Mix* and the *Barbara* image and compare to *TGV*. It can be seen that the improvement is comparable to the denoising case. In particular, the method is still able to learn reasonable atoms from the given, blurry data and in particular for the texture parts the improvement is quite significant. Regarding the optimal PSNR values, the improvement is roughly around 1 to 1.5 decibel.

5.5 Comparison

This section compares the method $CT-cvx$ to its semi-convex variant $CT-scvx$ and to other methods. At first, we consider the learning of atoms from incomplete data and image inpainting in Figure 8. It can be seen there that for the *Patches* image, the semi-convex variant achieves an almost perfect results: It is able to learn exactly the three atoms that compose the texture part of the image and to inpaint the image very well. For the *Barabara* image, where more atoms are necessary to synthesize the texture part, the two methods yield similar results and also the atoms are similar. These results are also reflected in the PSNR values of Table 10, where $CT-scvx$ is more that 7 decibel better for the *Patches* image and achieve only a slight improvement for *Barbara*.

Next we consider the semi-convex variant $CT-scvx$ for denoising the *Patches* and *Barbara* images of Figure 6. In this setting, also other methods are applicable and we compare to a costume implementation of a variant of the convolutional lasso algorithm (called CL) and to BM3D denoising [17] (called $BM3D$). For the former, we strive to solve the non-convex optimization problem

$$\min_{u, (c_i)_i, (p_i)_i} \text{TV}_\rho(u - \sum_{i=1}^k c_i * p_i) + \sum_{i=1}^k \|c_i\|_1 + \|u - f_0\|_2^2 \quad \text{s.t. } \|p_i\|_2 \leq 1, \int p_i = 0 \text{ for } i = 1, \dots, k$$

where $(c_i)_i$ are coefficient images, p_i are atoms and k is the number of used atoms. Note that we use the same boundary extension, atom-domain-size and stride variable than in the methods $CT-cvx$, $CT-scvx$, and that TV_ρ denotes a discrete TV functional with a slight smoothing of the L^1 norm to make it differentiable (see the source code [13] for details). For the solution we use an adaption of the algorithm of [34]. For BM3D we use the implementation obtained from [1].

Remark 15. *We note that, while we provide the comparison to BM3D in order to have a reference on achievable denoising quality, we do not aim to propose an improved denoising method that is comparable to BM3D. In contrast to BM3D, our method constitutes a variational (convex) approach, that is generally applicable for inverse problems and for which we were able to provide a detailed analysis in function space such that in particular stability and convergence results for vanishing noise can be proven. Furthermore, beyond mere image reconstruction, we regard the ability of simultaneous image-atom-learning and cartoon-texture decomposition as an important feature of our approach.*

Results for the *Patches* and *Barbara* image can be found in Figure 9, where for CL we allowed for three atoms for the *Patches* images an tested 3,5 and 7 atoms for the *Barbara* image, showing the best result that was obtained with 7 atoms. It can be seen that, as with the inpainting results, $CT-scvx$ achieves a very strong improvement compared to $CT-cvx$ for the *Patches* image (obtaining the atoms almost perfectly) and only a slight improvement for the *Barbara* image. The CL method performs similar but slightly worse than $CT-scvx$. While for the *Patches* image also the three main features are identified correctly, they are not centered which leads to artifacts in the reconstruction and might be explained by the method being stuck in a local minimum. For the *Patches* image, the result of $BM3D$ are comparable but slightly smoother than the ones of $CT-scvx$. In particular, the target-like structure in the bottom left is not very well reconstructed with $BM3D$ but suffers from less remaining noise. For the *Barbara* image, $BM3D$ delivers the best and result, but a slight overs-smoothing is visible. Regarding the PSNR values of Table 10, $BM3D$ performs best and $CT-scvx$ second best (better than CL), where in accordance with the visual results the difference of $BM3D$ and $CT-scvx$ is not as high as with *Barbara*.

6 Discussion

Using lifting techniques, we have introduced a (potentially convex) variational approach for learning image atoms from corrupted and/or incomplete data. An important part of our work is the analysis of the proposed model, which shows well-posedness results for the proposed model in function space for a general inverse problem setting. The numerical part shows that indeed our model can effectively learn image atoms from different types of data. While this works well also in a convex setting, moving to a semi-convex setting (which is also captured by our theory) yields a further, significant improvement. While the proposed method can also be regarded solely as image reconstruction method, we believe its main feature is in fact the ability to learn image atoms from incomplete data in a mathematically well understood framework.

Interesting future research questions are for instance the exploration of our method for classification problems or the exploration of similar lifting techniques for a mathematical understanding of deep neural networks.

7 Acknowledgements

MH acknowledges support by the Austrian Science Fund (FWF) (Grant J 4112). TP is supported by the European Research Council under the Horizon 2020 program, ERC starting grant agreement 640156.

A Appendix: Tensor spaces

We recall here some basic results on tensor products of Banach spaces that will be relevant for our work. Most of these results are obtained from [39, 19], to which we refer to for further information and a more complete introduction to the topic.

Throughout this section, let always $(X, \|\cdot\|_X)$, $(Y, \|\cdot\|_Y)$, $(Z, \|\cdot\|_Z)$ be Banach spaces. By X^* we denote the analytic dual of X , i.e., the space of bounded linear functionals from X to \mathbb{R} . By $\mathcal{L}(X, Y)$ and $\mathcal{B}(X \times Y, Z)$ we denote the spaces of bounded linear and bilinear mappings, respectively, where the norm for the latter is given by $\|B\|_{\mathcal{B}} = \sup\{\|B(x, y)\|_Z \mid \|x\|_X \leq 1, \|y\|_Y \leq 1\}$. In case the image space is the reals, we write $\mathcal{L}(X)$ and $\mathcal{B}(X \times Y)$.

Algebraic tensor product. The tensor product $x \otimes y$ of two elements $x \in X$, $y \in Y$ can be defined as a linear mapping on the space of bilinear forms on $X \times Y$ via

$$x \otimes y(A) = A(x, y)$$

The algebraic tensor product $X \otimes Y$ is then defined as the subspace of the space of linear functionals on $\mathcal{B}(X, Y)$ spanned by elements $x \otimes y$ with $x \in X$, $y \in Y$.

Tensor norms. We will use two different tensor norms, the *projective* and the *injective* tensor norm (also known as the largest and smallest reasonable cross norm, respectively). The projective tensor norm on $X \otimes Y$ is defined for $u \in X \otimes Y$ as

$$\|u\|_{\pi} := \inf \left\{ \sum_{i=1}^n \|x_i\|_X \|y_i\|_Y \mid u = \sum_{i=1}^n x_i \otimes y_i, n \in \mathbb{N} \right\}.$$

Note that indeed $\|\cdot\|_\pi$ is a norm and $\|x \otimes y\|_\pi = \|x\|_X \|y\|_Y$ (see [39, Proposition 2.1]). We denote by $X \otimes_\pi Y$ the completion of the space $X \otimes Y$ equipped with this norm. The following result gives a useful representation of elements in $X \otimes_\pi Y$ and their projective norm.

Proposition 16. *For $u \in X \otimes_\pi Y$ and $\epsilon > 0$ there exist bounded sequences $(x_n)_n \subset X$, $(y_n)_n \subset Y$ such that*

$$u = \sum_{i=1}^{\infty} x_n \otimes y_n \quad \text{and} \quad \sum_{i=1}^{\infty} \|x_n\|_X \|y_n\|_Y < \|u\|_\pi + \epsilon.$$

In particular,

$$\|u\|_\pi = \inf \left\{ \sum_{i=1}^{\infty} \|x_i\|_X \|y_i\|_Y \mid u = \sum_{i=1}^{\infty} x_i \otimes y_i \right\}.$$

Now for the injective tensor norm, we note that elements of the tensor product $X \otimes Y$ can be viewed as bounded bilinear forms on $X^* \times Y^*$ by associating to a tensor $u = \sum_{i=1}^n x_i \otimes y_i$ the bilinear form $B_u(\phi, \psi) = \sum_{i=1}^n \phi(x_i)\psi(y_i)$, where this association is unique (see [39, Section 1.3]). Hence $X \otimes Y$ can be regarded as a subspace of $\mathcal{B}(X^* \times Y^*)$ and the injective tensor norm is the norm induced by this space. Thus for $u = \sum_{i=1}^n x_i \otimes y_i$ the injective tensor norm $\|\cdot\|_i$ is given as

$$\|u\|_i = \sup \left\{ \left| \sum_{i=1}^n \phi(x_i)\psi(y_i) \right| \mid \|\phi\|_{X^*} \leq 1, \|\psi\|_{Y^*} \leq 1 \right\}$$

and the injective tensor product $X \otimes_i Y$ is defined as the completion of $X \otimes Y$ with respect to this norm.

Tensor lifting. The next result (see [39, Theorem 2.9]) shows that there is a one-to-one correspondence between bounded bilinear mappings from $X \times Y$ to Z and bounded linear mappings from $X \otimes_\pi Y$ to Z .

Proposition 17. *For $B \in \mathcal{B}(X \times Y, Z)$ there exists a unique linear mapping $\hat{B} : X \otimes_\pi Y \rightarrow Z$ such that $\hat{B}(x \otimes y) = B(x, y)$. Further \hat{B} is bounded and the mapping $B \mapsto \hat{B}$ is an isometric isomorphism between the Banach spaces $\mathcal{B}(X \times Y, Z)$ and $\mathcal{L}(X \otimes_\pi Y, Z)$.*

Using this isometry, for $B \in \mathcal{B}(X \times Y, Z)$ we will always denote by \hat{B} the corresponding linear mapping on the tensor product.

The following result is provided in [39, Proposition 2.3] and deals with the extension of linear operators to the tensor product.

Proposition 18. *Let $S \in \mathcal{L}(X, W)$, $T \in \mathcal{L}(Y, Z)$. Then there exists a unique operator $S \otimes_\pi T : X \otimes_\pi Y \rightarrow W \otimes_\pi Z$ such that $S \otimes_\pi T(x \otimes y) = (Sx) \otimes (Ty)$. Furthermore, $\|S \otimes_\pi T\| = \|S\| \|T\|$.*

Tensor space isometries. The following proposition deals with duality of the injective and the projective tensor products. To this aim, we need the notion of *Radon Nikodým property* and *approximation property*, which we will not define here but rather refer to [39, Sections 4 and 5] and [19]. For our purposes, it is important to note that both properties hold for L^r -spaces with $r \in (1, \infty)$, the *Radon Nikodým property* holds for reflexive spaces, but while we cannot expect the *Radon Nikodým property* to hold for L^∞ and \mathcal{M} , the approximation property does.

Lemma 19. *Assume that either X^* or Y^* has the Radon Nikodým property and that either X^* or Y^* has the approximation property. Then*

$$(X \otimes_i Y)^* \hat{=} X^* \otimes_\pi Y^*$$

and for simple tensors $u = \sum_{i=1}^n x_i \otimes y_i \in X \otimes_i Y$ and $u^* = \sum_{i=1}^m x_i^* \otimes y_i^* \in X^* \otimes_\pi Y^*$ the duality pairing is given as

$$\langle u^*, u \rangle = \sum_{i=1}^n \sum_{j=1}^m \langle x_j^*, x_i \rangle \langle y_j^*, y_i \rangle$$

Proof. The identification of the duals is shown in [39, Theorem 5.33]. For the duality pairing, we first note that the action of an element $u^* \in X^* \otimes_\pi Y^*$ on $X \otimes_i Y$ is given as the action of the associated bilinear form B_{u^*} [39, Section 3.4], which for simple tensors $u = \sum_{i=1}^n x_i \otimes y_i$ can be given as

$$\langle B_{u^*}, u \rangle = \sum_{i=1}^n B_{u^*}(x_i, y_i).$$

Now in case also u^* is a simple tensor, i.e., $u^* = \sum_{i=1}^m x_i^* \otimes y_i^*$, the action of this bilinear form can be given more explicitly [39, Section 1.3], which yields

$$\langle B_{u^*}, u \rangle = \sum_{i=1}^n \sum_{i=1}^m \langle x_i^*, x_i \rangle \langle y_i^*, y_i \rangle.$$

□

The duality between the injective and projective tensor product will be used for compactness assertions on subsets of the latter. To this aim, we note in the following lemma that separability of the individual space transfers to the tensor product. As consequence, in case X^* and Y^* satisfy the assumption of Lemma 19 above and both admit a separable predual, also $X^* \otimes_\pi Y^*$ admits a separable predual and hence bounded sets are weakly* compact.

Lemma 20. *Let X, Y be separable spaces. Then both $X \otimes_i Y$ and $X \otimes_\pi Y$ are separable.*

Proof. Take X' and Y' to be dense countable subsets of X and Y , respectively. First note that it suffices to show that any simple tensor $x \otimes y$ can be approximated arbitrarily close by $x' \otimes y'$ with $x' \in X', y' \in Y'$. But this is true since (using [39, Propositions 2.1 and 3.1])

$$\|x \otimes y - x' \otimes y'\| \leq \|x \otimes y - x \otimes y'\| + \|x \otimes y' - x' \otimes y'\| = \|x\| \|y - y'\| + \|y'\| \|x - x'\|,$$

where $\|\cdot\|$ denotes either the projective or the injective norm.

The following result, which can be obtained by direct modification of the result shown at the beginning of [39, Section 3.2], provides an equivalent representation of the injective tensor product in a particular case.

Lemma 21. *Denote by $C_c(\Omega_\Sigma, X)$ the space of compactly supported continuous functions mapping from Ω_Σ to X and denote by $C_0(\Omega_\Sigma, X)$ its completion with respect to the norm $\|\phi\|_\infty := \sup_{t \in \Omega_\Sigma} \|\phi(t)\|_X$. Then we have that*

$$C_0(\Omega_\Sigma) \otimes_i X \hat{=} C_0(\Omega_\Sigma, X)$$

where the isometry is given as the completion of the isometric mapping $J : C_0(\Omega_\Sigma) \otimes X \rightarrow C_0(\Omega_\Sigma, X)$ defined for $u = \sum_{i=1}^n f_i \otimes x_i$ as

$$Ju(t) := \sum_{i=1}^n f_i(t)x_i.$$

Next we consider the identification of tensor products with linear operators which is provided in the following proposition [39, Corollary 4.8].

Proposition 22. *Define the mapping $J : X^* \otimes_\pi Y \rightarrow \mathcal{L}(X, Y)$ as*

$$u = \sum_{i=1}^{\infty} \phi_n \otimes y_n \mapsto L_u : X \rightarrow Y \text{ where } L_u(x) = \sum_{i=1}^{\infty} \phi_n(x)y_n.$$

Then, J is well-defined and has unit norm. Defining $\mathcal{N}(X, Y) \subset \mathcal{L}(X, Y)$ as the range of J , equipped with the norm

$$\|T\|_{\text{nuc}} = \inf \left\{ \sum_{i=1}^{\infty} \|\phi_n\|_{X^*} \|y_n\|_Y \mid T(x) = \sum_{i=1}^{\infty} \phi_n(x)y_n \right\},$$

we get that $\mathcal{N}(X, Y)$ is a Banach space, called the space of nuclear operators. If further either X^* or Y has the approximation property, then J is an isometric isomorphism, that is, we can identify

$$X^* \otimes_\pi Y = \mathcal{N}(X, Y)$$

It is easy to see that nuclear operators are compact that that we can equivalently write

$$\|T\|_{\text{nuc}} = \inf \left\{ \sum_{i=1}^{\infty} \sigma_i \mid T(x) = \sum_{i=1}^{\infty} \sigma_i \phi_i(x)y_i, \|\phi_i\|_{X^*} \leq 1, \|y_i\|_Y \leq 1 \right\}.$$

Also, in a Hilbert space setting (see [43] for details), it is a classical result that for any compact $T \in \mathcal{L}(H_1, H_2)$ with $(H_1, (\cdot, \cdot)), (H_2(\cdot, \cdot))$ Hilbert spaces there exist orthonormal systems $(x_i)_i, (y_i)_i$ and uniquely defined singular values $(\sigma_i)_i := (\sigma_i(T))_i$ such that

$$Tx = \sum_{i=1}^{\infty} \sigma_i(x, x_i)y_i.$$

In addition, in case T has finite nuclear norm, it follows that $\|T\|_{\text{nuc}} = \sum_{i=1}^{\infty} \sigma_i$.

References

- [1] Implementation of BM3D denoising, v2.00 (30 January 2014). Obtained from <http://www.cs.tut.fi/~foi/GCF-BM3D> (October 1, 2018).
- [2] J. Adler and O. Öktem. Learned primal-dual reconstruction. *IEEE transactions on medical imaging*, 37(6):1322–1332, 2018.
- [3] G. Aubert and P. Kornprobst. *Mathematical Problems in Image Processing: Partial Differential Equations and the Calculus of Variations*, volume 147 of *Applied Mathematical Sciences*. Springer, 2006.

- [4] F. Bach, J. Mairal, and J. Ponce. Convex sparse matrix factorizations. *arXiv preprint arXiv:0812.1869*, 2008.
- [5] K. Bredies and M. Holler. Regularization of linear inverse problems with total generalized variation. *Journal of Inverse and Ill-Posed Problems*, 22(6):871–913, 2014.
- [6] K. Bredies and M. Holler. A TGV-based framework for variational image decompression, zooming and reconstruction. Part II: Numerics. *SIAM Journal on Imaging Sciences*, 8(4):2851–2886, 2015.
- [7] K. Bredies, K. Kunisch, and T. Pock. Total generalized variation. *SIAM Journal on Imaging Sciences*, 3(3):492–526, 2010.
- [8] K. Bredies and D. A. Lorenz. Regularization with non-convex separable constraints. *Inverse Problems*, 25(8):085011, 2009.
- [9] K. Bredies and T. Valkonen. Inverse problems with second-order total generalized variation constraints. In *Proceedings of SampTA 2011 - 9th International Conference on Sampling Theory and Applications*, Singapore, 2011.
- [10] H. Brezis. *Functional Analysis, Sobolev Spaces and Partial Differential Equations*. Springer, 2010.
- [11] A. Buades, B. Coll, and J.-M. Morel. A non-local algorithm for image denoising. In *Proc. of the CVPR*, volume 2, pages 60–65, 2005.
- [12] L. Calatroni, C. Cao, J. C. De Los Reyes, C.-B. Schönlieb, and T. Valkonen. Bilevel approaches for learning of variational imaging models. In *Variational Methods in Imaging and Geometric Control, Radon Series on Computational and Applied Mathematics*, volume 18, pages 252–290, 2016.
- [13] A. Chambolle, M. Holler, and T. Pock. Source code for a convex approach to learning convolutional image atoms from incomplete data. Will be available online after acceptance of this manuscript. Retrieved DATE.
- [14] A. Chambolle and T. Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, 2011.
- [15] A. Chambolle and T. Pock. An introduction to continuous optimization for imaging. *Acta Numerica*, 25:161–319, 2016.
- [16] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky. The convex geometry of linear inverse problems. *Foundations of Computational mathematics*, 12(6):805–849, 2012.
- [17] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Trans. Image Process.*, 16(8):2080–2095, 2007.
- [18] J. Delon, A. Desolneux, C. Sutour, and A. Viano. RNLP : Mixing Non-Local and TV-Lp methods to remove impulse noise from images. 2017. MAP5 2016-29, Hal-preprint Nr. hal01381063v2.

- [19] J. Diestel and J. J. Uhl. *Vector Measures*. Mathematical Surveys. American Mathematical Society, 1977.
- [20] M. Fazel. *Matrix rank minimization with applications*. PhD thesis, Stanford University, 2002.
- [21] Y. Gao and K. Bredies. Infimal convolution of oscillation total generalized variation for the recovery of images with structured texture. *SIAM Journal on Imaging Sciences*, 11(3):2021–2063, 2018.
- [22] C. Garcia-Cardona and B. Wohlberg. Convolutional dictionary learning: A comparative review and new algorithms. *IEEE Transactions on Computational Imaging*, 4(3):366–381, 2018.
- [23] M. Hintermüller and C. N. Rautenberg. Optimal selection of the regularization function in a weighted total variation model. part i: Modelling and theory. *Journal of Mathematical Imaging and Vision*, 59(3):498–514, Nov 2017.
- [24] 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(3):987–1010, 2007.
- [25] M. Holler, R. Huber, and F. Knoll. Coupled regularization with multiple data discrepancies. *Inverse Problems*, 34(8):084003, 2018.
- [26] E. Kobler, T. Klatzer, K. Hammernik, and T. Pock. Variational networks: connecting variational methods and deep learning. In *German Conference on Pattern Recognition*, pages 281–293. Springer, 2017.
- [27] K. Kunisch and T. Pock. A bilevel optimization approach for parameter learning in variational models. *SIAM Journal on Imaging Sciences*, 6(2):938–983, 2013.
- [28] M. Lebrun, M. Colom, A. Buades, and J.-M. Morel. Secrets of image denoising cuisine. *Acta Numerica*, 21:475–576, 2012.
- [29] S. Lunz, O. Öktem, and C.-B. Schönlieb. Adversarial regularizers in inverse problems. *arXiv preprint arXiv:1805.11572*, 2018.
- [30] S. Mallat. *A wavelet tour of signal processing - The sparse way, With contributions from Gabriel Peyré*. Elsevier/Academic Press, Amsterdam, 3. edition, 2009.
- [31] Y. Meyer. *Oscillating Patterns in Image Processing and Nonlinear Evolution Equations*. American Mathematical Society, 2001.
- [32] S. Oymak, A. Jalali, M. Fazel, Y. C. Eldar, and B. Hassibi. Simultaneously structured models with application to sparse and low-rank matrices. *IEEE Trans. Information Theory*, 61(5):2886–2908, 2015.
- [33] T. Papadopoulos and M. I. Lourakis. Estimating the jacobian of the singular value decomposition: Theory and applications. In *European Conference on Computer Vision*, pages 554–570. Springer, 2000.

- [34] T. Pock and S. Sabach. Inertial proximal alternating linearized minimization (ipalm) for nonconvex and nonsmooth problems. *SIAM Journal on Imaging Sciences*, 9(4):1756–1787, 2016.
- [35] E. Richard, F. R. Bach, J.-P. Vert, et al. Intersecting singularities for multi-structured estimation. In *ICML (3)*, pages 1157–1165, 2013.
- [36] E. Richard, G. R. Obozinski, and J.-P. Vert. Tight convex relaxations for sparse matrix factorization. In *Advances in neural information processing systems*, pages 3284–3292, 2014.
- [37] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1–4):259–268, 1992.
- [38] W. Rudin. *Fourier analysis on groups*. Courier Dover Publications, 2017.
- [39] R. A. Ryan. *Introduction to tensor products of Banach spaces*. Springer Science & Business Media, 2013.
- [40] O. Scherzer, M. Grasmair, H. Grossauer, M. Haltmeier, and F. Lenzen. *Variational Methods in Imaging*. Springer, 2009.
- [41] R. Thompson. Singular value inequalities for matrix sums and minors. *Linear Algebra and its Applications*, 11(3):251–269, 1975.
- [42] J. Weickert. *Anisotropic Diffusion in Image Processing*. Teubner, Stuttgart, 1998.
- [43] J. Weidmann. *Linear operators in Hilbert spaces*, volume 68. Springer Science & Business Media, 1980.
- [44] M. D. Zeiler, D. Krishnan, G. W. Taylor, and R. Fergus. Deconvolutional networks. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 2528–2535. IEEE, 2010.

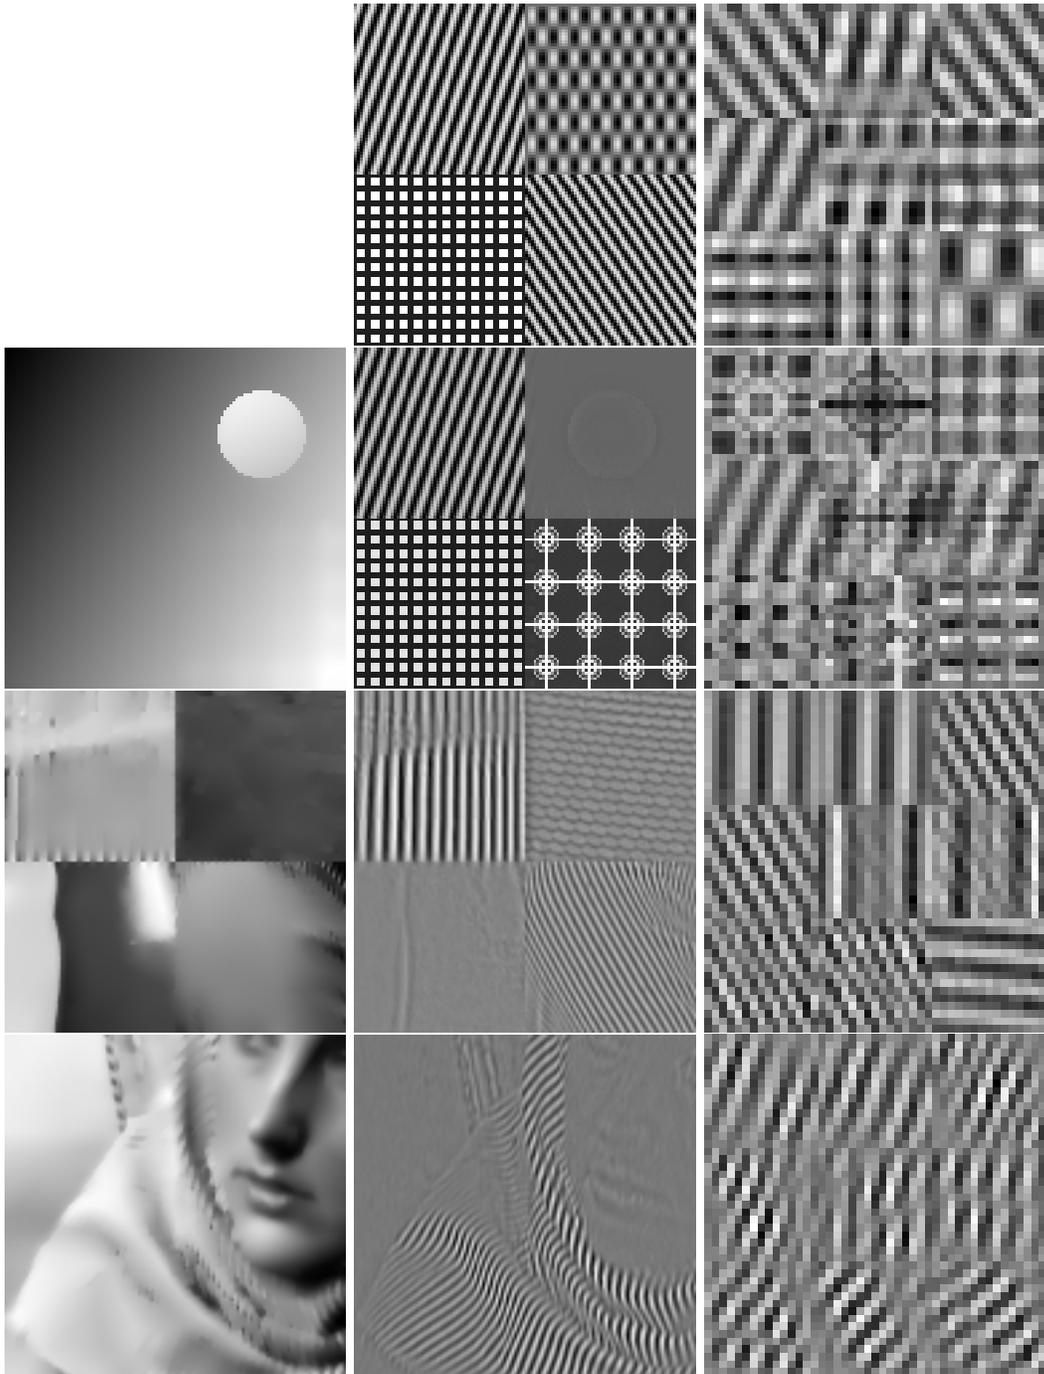


Figure 4: Cartoon-texture decomposition (rows 2-4) and nine most important learned atoms for different test images and the methods *TXT* (row 1) and *CT-cvx* (rows 2-4).

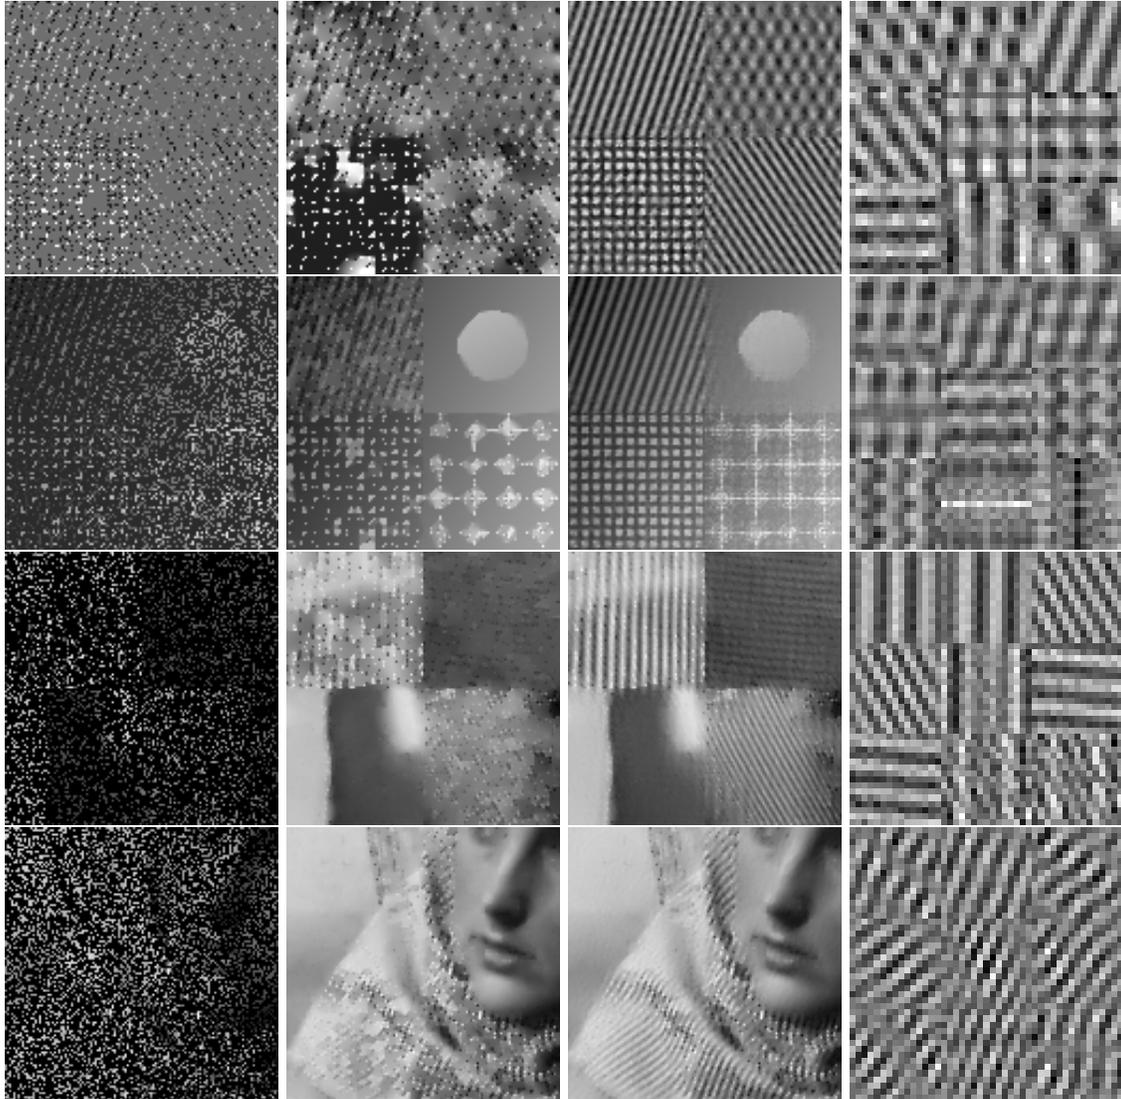


Figure 5: Image inpainting from incomplete data. From left to right: Data, TGV-based reconstruction, proposed method (only TXT in first row), nine most important learned atoms. Rows 1,3: 20% of pixels, rows 2,4: 30% of pixels.

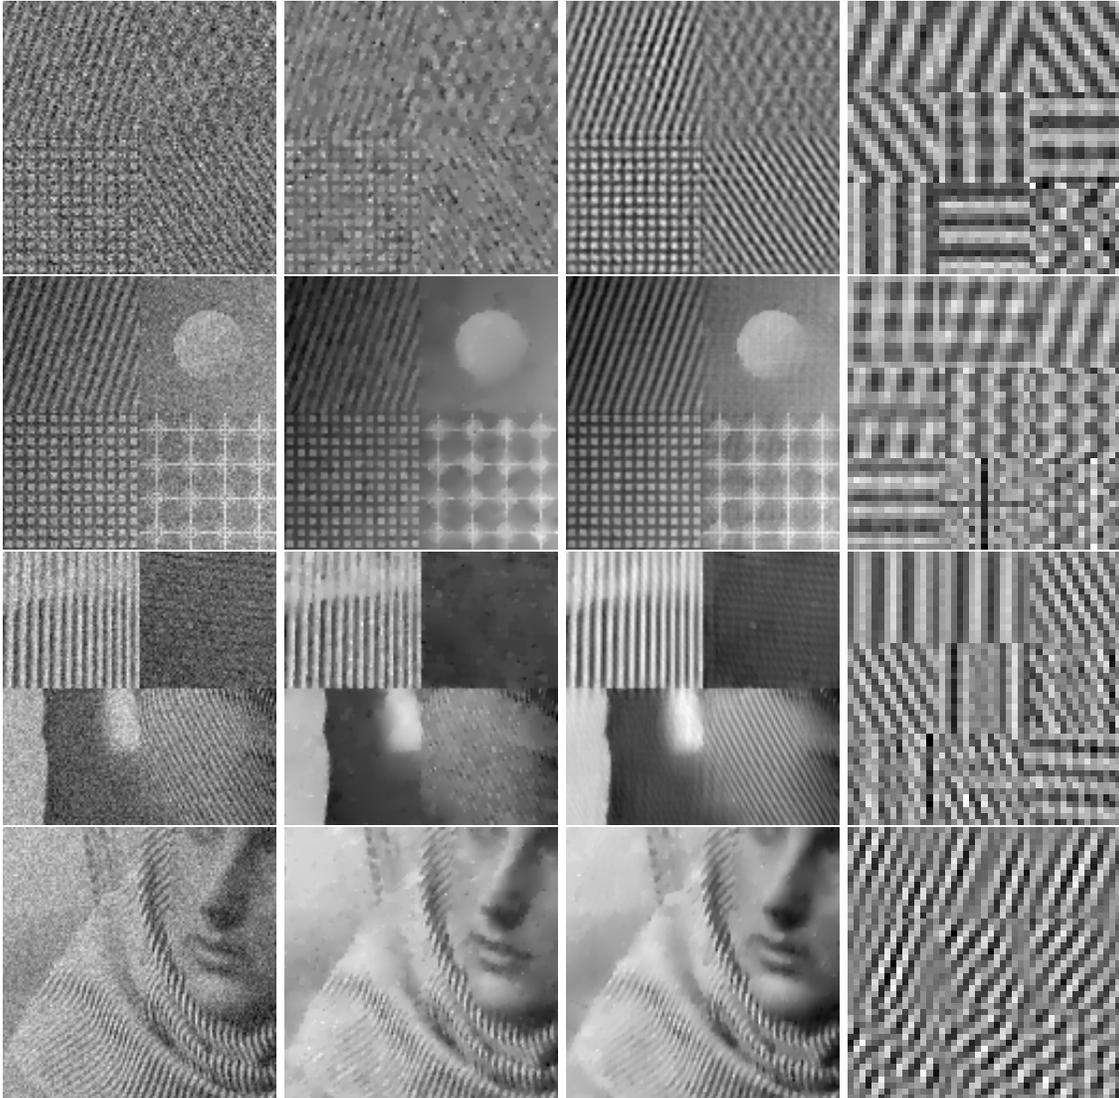


Figure 6: Denoising and atom-learning. From left to right: Noisy data, TGV-based reconstruction, proposed method (only *TXT* for the fist image), nine most important learned atoms.

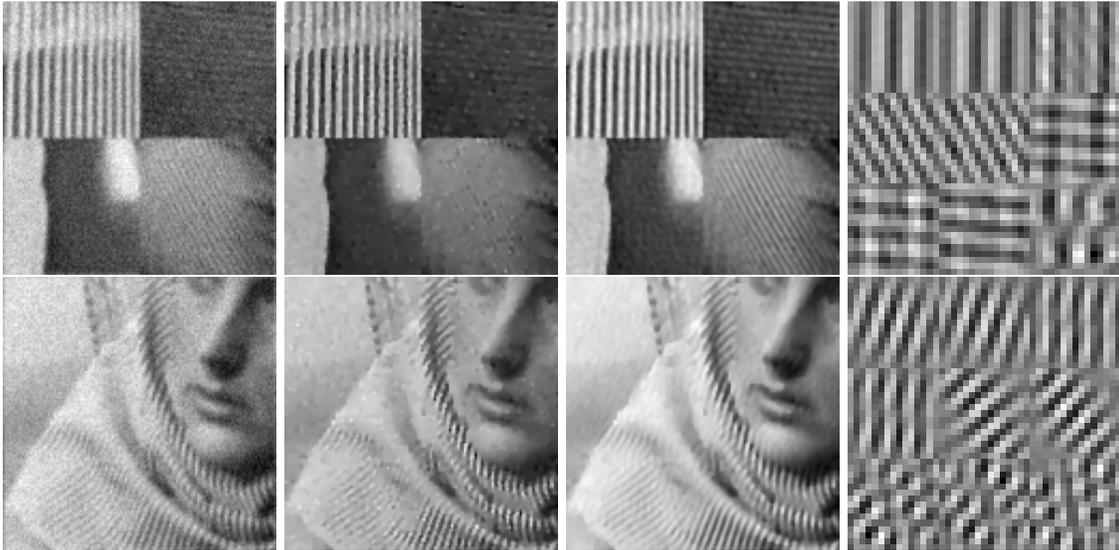


Figure 7: Reconstruction from blurry and noisy data. From left to right: Data, TGV, proposed, learned atoms.

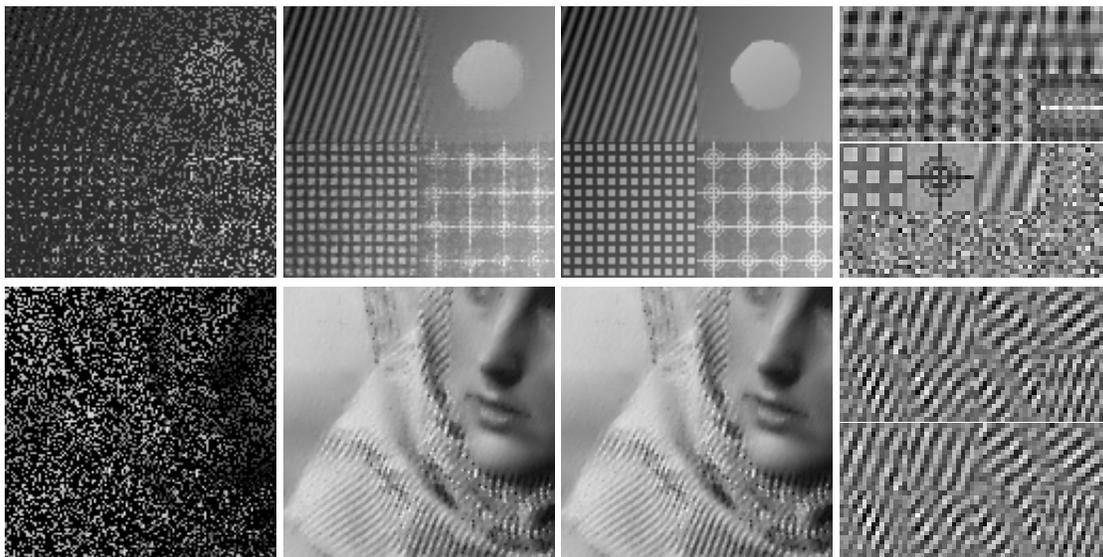


Figure 8: Comparison of $CT-cvx$ and $CT-scvx$ for inpainting with 30% of the pixels given. From left to right: Data, convex, semi-convex, convex atoms (top), semi-convex atoms (bottom).

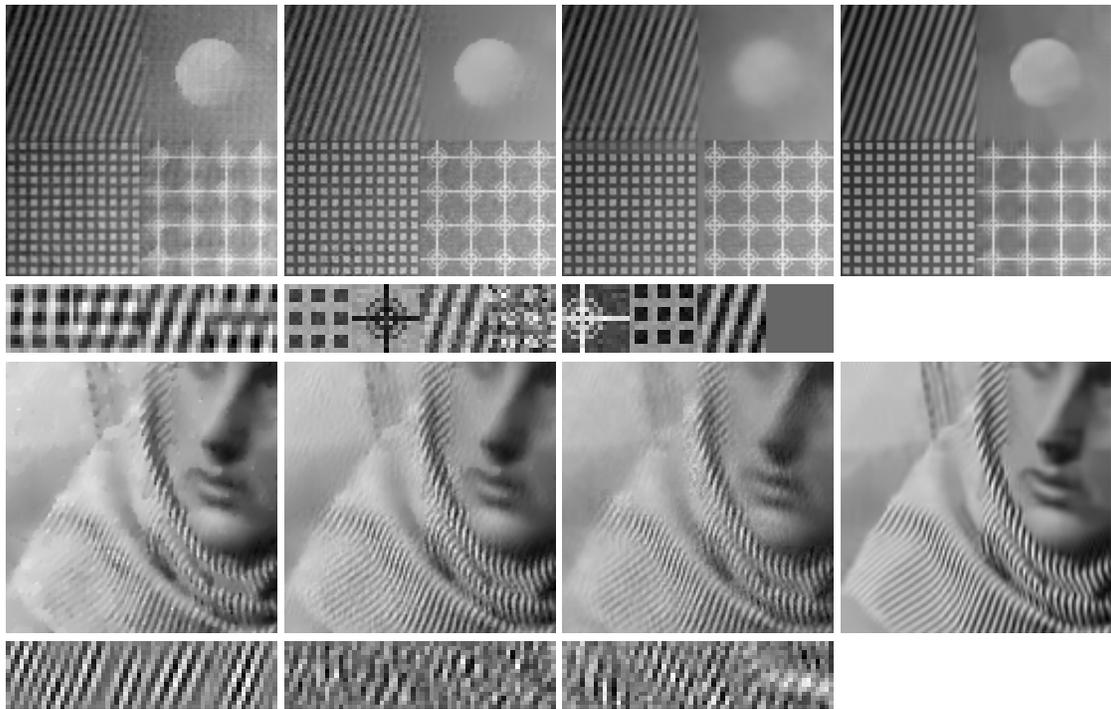


Figure 9: Comparison of different methods for denoising the *Patches* and *Barbara* images from Figure 6. From left to right: *CT-cvx*, *CT-scvx*, *CL* and *BM3D*. The four most important learned atoms are shown below the images.

	Texture	Patches	Mix	Barbara
Inpainting				
<i>TGV</i>	10.32	19.28	20.19	20.58
<i>TXT/ CT-cvx</i>	17.59	25.55	23.38	23.48
<i>CT-scvx</i>		32.74		23.6
Denoising				
<i>TGV</i>	11.83	23.96	23.74	23.99
<i>TXT/ CT-cvx</i>	16.06	25.91	26.07	25.0
<i>CT-scvx</i>		29.4		25.56
<i>CL</i>		29.09		25.14
<i>BM3D</i>		30.82		28.15
Deconvolution				
<i>TGV</i>			23.58	22.86
<i>CT-cvx</i>			25.19	23.93

Figure 10: Best PSNR result achieved with each method for the parameter test range as specified in Table 11. The best achieve result for each experiment is written in bold.

	<i>CT-cvx</i>			<i>CT-scvx</i>				<i>TGV</i>	<i>TXT</i>			<i>BM3D</i>	<i>CL</i>	
	λ	μ	ν	λ	μ	ν	ϵ	λ	λ	ν	λ	λ	μ	
Decomp.	-	opt	0.95						-	0.75				
Inp.	-	opt	0.975	-	opt	0.975	0.1	-	-	0.975				
Den.	opt	opt	0.975	opt	opt	0.975	2.0	opt	opt	0.975	opt	opt	opt	
Deconv.	opt	opt	0.975					opt						

Figure 11: Parameter choice for all methods and experiments used in the paper. Here, λ always defines the trade-off between data fidelity and regularization, μ defined the trade-off between cartoon and texture, ν defined the trade-off between the 1/2 norm and the penalization of singular values and ϵ defines the degree of non-convexity for the semi-convex potential. Whenever a parameter was optimized over a certain range for each experiment, we write *opt*.