



Universiteit
Leiden
The Netherlands

Spectral imaging and tomographic reconstruction methods for industrial applications

Zeegers, M.T.

Citation

Zeegers, M. T. (2023, May 31). *Spectral imaging and tomographic reconstruction methods for industrial applications*. Retrieved from <https://hdl.handle.net/1887/3619550>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/3619550>

Note: To cite this publication please use the final published version (if applicable).

Dictionary-based spectral tomography

Spectral X-ray imaging has the potential to significantly improve the accuracy of material decomposition in tomography. Nevertheless, the problem of volumetric material decomposition in spectral tomography is often highly ill-posed. Therefore, most state-of-the-art methods rely on strict prior information, and are able to provide material decomposition with a limited number of materials, with limited accuracy, or a limited set of spectral configurations.

To resolve these issues, we propose ‘A Dictionary-based Joint reconstruction and Unmixing method for Spectral Tomography’ (ADJUST). Its formulation relies on forming a dictionary of spectral signatures of materials common in CT and prior knowledge of the number of materials present in an object. In particular, we decompose the spectral volume linearly in terms of spatial material maps, a spectral dictionary, and the indicator of materials for the dictionary elements. We propose a memory-efficient accelerated alternating proximal gradient method to find an approximate solution to the resulting bi-convex problem. From numerical demonstrations on several synthetic phantoms, we observe that ADJUST performs exceedingly well compared to other state-of-the-art methods. Additionally, we address the robustness of ADJUST against limited and noisy measurement patterns. The demonstration of the proposed approach on a spectral micro-CT dataset shows its potential for real-world applications.

This chapter is based on:


M. T. Zeegers, A. Kadu, T. van Leeuwen, and K. J. Batenburg. “ADJUST: A Dictionary-based Joint reconstruction and Unmixing method for Spectral Tomography”. *Inverse Problems* 38.12 (2022), p. 125002.

5.1 Introduction


X-ray Computed Tomography (CT) estimates the spatial attenuation map of the object of interest from its measured X-ray projections obtained from different angles. The conventional tomography acquisition setup consists of a polychromatic X-ray source and an X-ray detector that collects the transmitted X-rays. However, these conventional detectors do not discriminate between different incident photon energies and collect the attenuated X-rays solely in one energy bin (Figure 5.1). Images of conventional detectors are reconstructed as grey scale volumes, representing the aggregate attenuation coefficients of the materials. Several materials may correspond to the same grey level, which makes it difficult to determine the material composition of an object. Nonetheless, the attenuations of materials are energy-dependent and their spectral attenuation curves are mutually different. Therefore, by probing multiple energy levels, additional information is obtained for discriminating the materials. A typical approach is dual-energy CT, which allows for more accurate *material decomposition* [40, 274], also known as material unmixing. It uses two polychromatic sources with different peak voltages, and correspondingly two sets of conventional detector panels to measure attenuated X-rays, each from one source. Dual-energy CT is commonly used in clinical settings for separating high-attenuating materials from low-attenuating materials, for example, determining the location of contrast agents [199] such as iodine in the body [142]. However, more accurate material decomposition and concentration determination requires multi-energy CT, enabled by multi-energy X-ray (photon counting) detectors [40, 274].

Multi-energy X-ray detectors allow for collecting X-ray projection data in multiple energy bins with a very high spectral resolution. We can divide these detectors into two classes: (i) detectors that measure all X-ray photons simultaneously and directly categorize these into various spectral channels, and (ii) detectors that indirectly do this. Detectors from the first category are often used for hyperspectral imaging, and some examples include the Hexitec detector [78, 134, 282], Amptek X-123 CdTe detector [236] and the SLcam [213, 247]. On the other hand, the spectroscopic X-ray detectors from the second category indirectly measure many spectral channels by applying threshold scans. Detectors from the Medipix and Pixirad families are examples of this [24, 25]. The number of energy bins that can be recorded simultaneously is limited (usually up to 10), but different energy thresholds, with which X-ray photons with higher energies than that threshold can be detected, can be set between measurements. If the object of interest is static, hyperspectral images can, in principle, be easily obtained using these devices.

The measurements from these detectors, *i.e.* the tomographic projections after their preprocessing and log-correction, are linearly related to the spectral characteristics of the materials present in the object. Although these detectors can now measure the X-ray projections in multiple energy bins (combined referred to as



(a) Data acquisition in conventional X-ray CT



(b) Data acquisition in spectral X-ray CT

Figure 5.1: Comparison of data acquisitions in conventional X-ray CT (a) and spectral X-ray CT (b). In conventional X-ray CT, the energy window is dictated by the source spectrum (grey). None of the detected X-ray photons are distinguished by energy, leading to one sinogram (containing the projections ordered by angle). On the other hand, in spectral CT, the X-ray photons are distinguished according to their energy from multiple windows (colors), yielding multiple sinograms with different characteristics.

spectral projections), it is not straightforward to optimally use the available spectral data for material decomposition. The conventional strategies for reconstruction and material decomposition are not designed to simultaneously process this type of spectral measurements and find accurate material maps (refer to Figure 5.3). Hence, advanced computational techniques are required to infer the material composition from these spectral projections.

5.1.1 Motivation and problem description

Spectral tomographic projections are X-ray measurements at multiple energies that result from the spatial and spectral properties of the materials present in the object. As noted in Section 1.5, *spectral computed tomography* is the process of estimating spatial maps and spectral signatures of materials present in the object from spectral tomographic projections. In general, the spectral CT workflow consists of two steps: (i) *reconstruction*: computing spatial maps from their tomographic projections, (ii) *unmixing*: decomposition of the spectral volume into spatial material maps and their spectral signatures. Mathematically, reconstruction boils down to solving an inverse problem, while decomposition involves matrix factorization, an unsupervised learning approach. These two steps can be performed in a serial fashion to obtain the material maps from the spectral projections (refer to Figure 5.2). However, these two steps can also be combined into a one-step approach, which we refer to as *joint*.

To illustrate the performance of these two-step and one-step approaches, we consider a Shepp-Logan phantom, as described in Figure 5.3. The numerical phantom consists of five materials (vanadium, chromium, manganese, iron and cobalt). For the tomographic projections, a full-view setting is chosen where 180 projections between 0 to π are acquired. Figure 5.3 shows that the two-step approaches do not yield clear material decompositions or spectral signature reconstructions. Although the two-step methods are computationally efficient, these generally do not yield the same solution due to the ill-posedness of the problem. For the same reason, the joint approach is not expected to work well in all cases either. In particular, the current algorithms for joint methods suffer from ill-conditioning of the spectral profiles (refer to Section 5.4.2). The suboptimal performance of these methods motivates us to develop novel reconstruction methods to improve the spatial resolution and precise characterization of materials. To address the ill-posedness, we incorporate spatial and spectral prior information.

5.1.2 Contributions and outline

In this chapter, we propose a new technique called ‘*A Dictionary-based Joint reconstruction and Unmixing method for Spectral Tomography*’ (ADJUST) to reconstruct spatial material maps from their spectral tomographic measurements. ADJUST is a novel bi-convex optimization formulation that incorporates an effective *spatio-spectral* prior. This prior includes (i) *spatial*: the contribution of materials at




Figure 5.2: A schematic overview of various approaches to spectral computed tomography: The spectral sinograms (projections from various angles) obtained at different energy levels make up the input of these algorithms (a). The spectral sinograms can be decomposed first in the projection domain to obtain material sinograms and the spectral signatures (b). The spatial material maps (d) are then obtained by reconstructing each material sinogram separately. Alternatively, reconstruction can be done before the material decomposition by first making a CT reconstruction of every spectral channel using the associated spectral sinograms to obtain the spectral volumes (c). Then, in the image domain, these spectral volumes can be decomposed to obtain the spatial material volumes. These two-step approaches can also be combined into a one-step approach in which the decomposition and reconstruction are carried out as a joint approach.

each location should sum to 1, and (ii) *spectral*: the spectral signature of material should be a linear combination of the elements of a spectral dictionary. To solve the constrained optimization problem, we design a memory-efficient alternating proximal iterative scheme. In particular, we develop numerical methods to compute the required proximal operators quickly. In numerical experiments, we demonstrate that ADJUST performs better than the existing state-of-the-art methods. Furthermore, we show that ADJUST is also applicable to limited-angle problems found in industrial X-ray tomography, optical tomography and electron tomography.

The remainder of the chapter is organized as follows. Section 5.2 discusses the existing work on spectral CT and material decomposition. Section 5.3 introduces the forward modelling of spectral X-ray tomography. In particular, we derive the linear map from the spatial map of materials to the spectral tomographic measurements. Then, we introduce the inverse problem in Section 5.4 that estimates the spatial material maps and spectral signatures from the spectral tomographic measurements. Here, we also discuss the ill-posedness involved in the inversion process. To reduce this ill-posedness, we introduce ADJUST in Section 5.5. Moreover, we propose an iterative scheme that finds an approximate solution to the resulting bi-convex



Figure 5.3: Comparison of various methods for spectral CT on a five-material spectral Shepp-Logan phantom. The tomographic projections are gathered in 100 equally-sized bins between 5 to 35 keV. The top row shows the true material maps and the spectral signatures of each material. The second and third row show the material maps retrieved from two-step methods, reconstruction-then-unmixing (RU) and unmixing-then-reconstruction (UR), respectively. The fourth row shows the results of the classical joint method (see Section (5.6)). In contrast, the fifth row shows results of the proposed method (see Section 5.5).

formulation. In Section 5.6, we numerically compare ADJUST with other methods on various synthetic phantoms. We also demonstrate the robustness of the method on limited measurement patterns such as sparse-angle tomography and limited view tomography, as well as in situations with limited spectral resolution and higher noise levels. In addition, we apply ADJUST to an experimental X-ray micro-CT dataset. Finally, we discuss the possibilities and limitations of the approach in Section 5.7 and conclude the chapter.

5.2 Related work

For the sake of convenience, we categorize the previous work on spectral CT into (i) (two-step) sequential approaches and (ii) (one-step) joint approaches to reconstruction and material decomposition. Since we focus on multi-spectral CT, we do not discuss the advances in dual-energy CT. However, we refer the reader to the comprehensive review paper [199] that covers dual-energy CT.

For spectral CT, sequential approaches (also known as two-step methods) are mainly (i) reconstruction followed by decomposition [60, 77, 95, 298, 299, 318], and (ii) decomposition followed by reconstruction [125, 190, 248]. In the former category, material decomposition is carried out in the image domain, while in the

latter category, it is carried out in the projection domain (see Figure 5.2). In both approaches, independent methods for material decomposition in the projection domain [75], (multi-channel) spectral reconstruction [312] with various forms of structural or spectral regularization [143, 244, 246], and material decomposition in the image domain [95] can be plugged in. Although these sequential two-step methods are computationally inexpensive, separating the reconstruction and decomposition steps causes information loss [202, 292].

Joint methods (also known as one-step methods) that simultaneously reconstruct multi-channel spectral images and perform material decomposition have been developed to address the issues associated with sequential methods [136, 202]. All current one-step methods are iterative in nature and allow for the incorporation of prior knowledge through regularization. For example, *(i)* structure on material maps is imposed through various penalties [125, 202], *(ii)* material maps are constrained using simplex constraints [95, 298, 299], and *(iii)* structure on spectral signatures is enforced using a spectral dictionary [299, 316]. In particular, when the materials present in the object of interest are precisely known, various techniques improve the quality of the reconstructions. These tailored methods can image contrast agents, such as iodine [318], gold, gadolinium in angiography [190], in the presence of bone and tissue. Sometimes, knowledge about the materials can be used to choose the spectral bins effectively [95]. Moreover, a deep learning approach has been proposed to perform joint decomposition and reconstruction task by generating a training set based on synthetic phantoms [5]. Most existing methods have been developed for, or demonstrated on, a limited number of materials (for example, two [136, 166, 189], three [27, 316], and six [95]), sometimes heavily relying on prior information about the materials, spectral signatures and energy bins. Some of these methods are extendable to more materials [5], but this can be difficult for each additional parameter that may need to be estimated with each new material [27]. However, these methods suffer when *(i)* the number of materials present in the object is larger than 3, *(ii)* the number of projections is smaller than the conventional criterion, *(iii)* the measurements are corrupted with high noise [118].

5.3 Spectral forward model

In this section, we provide the spectral X-ray forward model, by describing how we represent the objects and how spectral X-ray projections are obtained from this. The object is characterised by its attenuation coefficients $\mu(x, E) \in \mathbb{R}_+$, where $x \in \mathbb{R}^d$ is the location and $E > 0$ the X-ray photon energy, with $d \in \{2, 3\}$ being the dimension of our space depending on whether we consider a slice-based or full 3D reconstruction. Given a polychromatic X-ray source with a source spectrum $I_0(E)$ at energy level E , and C energy windows from the set $\mathcal{E} = \{\mathcal{E}_c\}_{c \in \mathcal{C}} = \{[E_c^{\min}, E_c^{\max}]\}_{c \in \mathcal{C}}$ (with $\mathcal{C}$ being its index set with $|\mathcal{C}| = C$) in which a spectral detector captures associated X-ray photons, we model the total X-ray photons captured by a detector pixel in

energy bin $\mathcal{E}_c$ as follows:

$$I(\mathcal{E}_c) = \int_{E_c^{\min}}^{E_c^{\max}} I_0(E) \exp \left(- \int_{\ell} \mu(x, E) dx \right) dE. \quad (5.1)$$

Here, the inner integral is taken over the line ℓ from the X-ray source to a detector pixel. The maximum and minimum energy range depend on detector specifications. We model the energy-dependent attenuation as a linear combination of energy-dependent material attenuations and their spatial contributions. We represent it mathematically as

$$\mu(x, E) = \sum_{m \in \mathcal{M}} \mu_m(E) \alpha_m(x), \quad (5.2)$$

where the material attenuation coefficient μ_m is a function of energy. The material m must be contained in a set of considered materials $\mathcal{M}$ of size M . The proportion of material m at location x is given by α_m . From Equations (5.1) and (5.2), we arrive at the following continuous relationship of measured photons in terms of the material spatial distributions and their attenuations:

$$I(\mathcal{E}_c) = \int_{E_c^{\min}}^{E_c^{\max}} I_0(E) \exp \left(- \sum_{m \in \mathcal{M}} \mu_m(E) \int_{\ell} \alpha_m(x) dx \right) dE.$$

If the spectral bins are sufficiently narrow then the source spectrum and the material-attenuation values can be approximated by their representative (average) values. These values are respectively $\bar{I}_0(\mathcal{E}_c)$ and $\bar{\mu}_m(\mathcal{E}_c)$, and form the following relation for the photon count in energy bin $\mathcal{E}_c$ corresponding to the c^{th} channel:

$$\begin{aligned} I(\mathcal{E}_c) &\approx \bar{I}(\mathcal{E}_c) \\ &= \bar{I}_0(\mathcal{E}_c) \exp \left(- \sum_{m \in \mathcal{M}} \bar{\mu}_m(\mathcal{E}_c) \int_{\ell} \alpha_m(x) dx \right). \end{aligned}$$

In general, the photon count is perturbed by an energy dependent noise distribution. Moreover, the spectral X-ray detector is a photon-counting detector, for which the noise in energy bin $\mathcal{E}_c$ can be modelled using a Poisson distribution with parameter $I(\mathcal{E}_c)$ [221]. If the mean in each energy bin is sufficiently high, a realization $I^{\text{meas}}(\mathcal{E}_c)$ of $I(\mathcal{E}_c)$ measured by a spectral detector can be approximated using Gaussian distribution $\mathcal{N}(0, \sigma^2)$, with variance σ being inversely proportional to $I(\mathcal{E}_c)$. In our experiments, we assume that the average photon count in each bin of detectors are sufficient for this approximation.

5.4 Spectral inverse problem

From the measured spectral X-ray projections, *spectral computed tomography* aims to retrieve the energy-dependent attenuation values $\mu_m(E_c)$ for each material m

and energy channel E_c and the distribution α_m throughout the object for each material m . It boils down to solving the bilinear system, in terms of $\bar{\mu}_m$ and α_m ,

$$\sum_{m \in \mathcal{M}} \bar{\mu}_m(\mathcal{E}_c) \int_{\ell_j} \alpha_m(x) dx = -\ln \left(\frac{I_j^{\text{meas}}(\mathcal{E}_c)}{\bar{I}_0(\mathcal{E}_c)} \right), \quad \forall \ell_j \in L, \mathcal{E}_c \in \mathcal{E}, \quad (5.3)$$

with L and $\mathcal{E}$ denoting the set of rays (with size $|L| = J$) and the set of energy channels, respectively. In tomography, the material distributions are determined by discretizing the object space into the grid of either pixels (2D) or voxels (3D). For now, we consider a three-dimensional scene with N voxels. The proportion of material m in the i^{th} voxel is then given by a_{im} . For each ray over line ℓ_j corresponding to the j^{th} measurement, the quantity w_{ji} determines the contribution of i^{th} voxel to the j^{th} measurement. Usually, the quantity $\bar{I}_0(\mathcal{E}_c)$ can be determined accurately by performing a flatfield measurement (i.e. measurement without object). Therefore, the right-hand side of Equation (5.3) is known. By expressing $-\ln(I_j^{\text{meas}}(\mathcal{E}_c)/\bar{I}_0(\mathcal{E}_c))$ by y_{jc} and $\bar{\mu}_m(\mathcal{E}_c)$ by f_{mc} , we arrive at the following expression:

$$\sum_{m=1}^M \left(\sum_{i=1}^N w_{ji} a_{im} f_{mc} \right) = y_{jc}, \quad j = 1, \dots, J, \text{ and } c = 1, \dots, C. \quad (5.4)$$

Here, the total number of measurements for each channel is given by J , and C denotes the total number of (energy) channels. Subsequently, we can write the expression in Equation (5.4) in the following matrix notation

$$\mathbf{W}\mathbf{A}\mathbf{F} = \mathbf{Y}, \quad (5.5)$$

where $\mathbf{Y} \in \mathbb{R}^{J \times C}$ represents tomographic measurements for C number of channels, and $\mathbf{W} \in \mathbb{R}^{J \times N}$ is a projection matrix containing the weights w_{ji} described in Equation (5.4), $\mathbf{A} \in \mathbb{R}^{N \times M}$ consists of M columns of size N , with each column representing a spatial map corresponding to the material present in the object, while $\mathbf{F} \in \mathbb{R}^{M \times C}$ consists of M rows with each row denoting the channel attenuation information of the material. It is important to note that matrices $\mathbf{W}$ and $\mathbf{Y}$ are known and matrices $\mathbf{A}$ and $\mathbf{F}$ are unknown. We formulate the joint spectral tomographic imaging and decomposition problem in a constrained least-squares form as

$$\begin{aligned} \underset{\mathbf{A}, \mathbf{F}}{\text{minimize}} \quad & \mathcal{J}(\mathbf{A}, \mathbf{F}) \triangleq \frac{1}{2} \|\mathbf{Y} - \mathbf{W}\mathbf{A}\mathbf{F}\|_F^2, & (\text{least-squares misfit}) \\ \text{subject to} \quad & \mathbf{A} \geq 0, & (\text{non-negativity of} \\ & & \text{material maps}) \\ & \mathbf{F} \geq 0. & (\text{non-negativity of att-} \\ & & \text{enuation coefficients}) \end{aligned} \quad (5.6)$$

Here, we impose non-negativity constraints on both $\mathbf{A}$ and $\mathbf{F}$. The function $\mathcal{J}(\mathbf{A}, \mathbf{F}) : \mathbb{R}^{N \times M} \times \mathbb{R}^{M \times C} \mapsto \mathbb{R}$ defines the misfit between the true measurements $\mathbf{Y}$ and the simulated measurements $\mathbf{W}\mathbf{A}\mathbf{F}$ using the Frobenius norm. This norm is valid if the noise in the measurements is approximately Gaussian. We note that this joint formulation is a bi-convex optimization problem since the misfit function is bi-convex, and the constraint set is a bi-convex set [103]. We denote the solution set of the joint formulation by

$$\mathcal{B} = \left\{ (\mathbf{A}_{\text{joint}}, \mathbf{F}_{\text{joint}}) \right\} = \arg \min_{\mathbf{A} \geq 0, \mathbf{F} \geq 0} \mathcal{J}(\mathbf{A}, \mathbf{F}).$$

This solution set $\mathcal{B}$ may contain more than one solution if the misfit function $\mathcal{J}$ is not strongly bi-convex. The solution set $\mathcal{B}$ cannot be determined trivially. To find the elements in set $\mathcal{B}$, we need to solve the optimization problem (5.6) using an iterative scheme with an initial estimate of the solution [290].

5.4.1 Practical methods

As outlined in Section 5.2, joint methods aim to simultaneously estimate the spatial material distribution and the energy-dependent attenuation coefficients. However, these one-step methods are not practical due to their high computational cost. In practice, two-step methods are popular, where reconstruction and decomposition are performed separately, because of their modular nature. For each step, tailored solvers are readily available for different platforms. The first category of two-step methods, which we call *RU* (short for Reconstruction-then-Unmixing), estimates a spectral volume from the spectral tomographic measurements, and then decomposes the resulting spectral volume to obtain the material maps and spectral signatures. That is, *RU* solves the following problems in a serial fashion:

$$\begin{aligned} \mathbf{V}_{\text{RU}} &= \arg \min_{\mathbf{V} \geq 0} \frac{1}{2} \|\mathbf{W}\mathbf{V} - \mathbf{Y}\|_F^2 + \lambda \mathcal{R}_1(\mathbf{V}) \\ (\mathbf{A}_{\text{RU}}, \mathbf{F}_{\text{RU}}) &= \arg \min_{\mathbf{A} \geq 0, \mathbf{F} \geq 0} \frac{1}{2} \|\mathbf{A}\mathbf{F} - \mathbf{V}_{\text{RU}}\|_F^2, \end{aligned}$$

where $\mathbf{V}_{\text{RU}} \in \mathbb{R}^{N \times C}$ is a spectral volume, and $\mathbf{A}_{\text{RU}}, \mathbf{F}_{\text{RU}}$ are the material maps and spectral signatures respectively reconstructed by this method. $\mathcal{R}_1 : \mathbb{R}^{N \times C} \mapsto \mathbb{R}$ is a regularization function that incorporates prior information about the spectral volumes and $\lambda \geq 0$ is a regularization parameter. Contrarily, *UR* (short for Unmixing-then-Reconstruction), the other class of two-step methods, separates the spectral tomographic measurements into projections and spectral signatures. These projections then lead to the material maps. *UR* mathematically reads

$$\begin{aligned} (\mathbf{P}_{\text{UR}}, \mathbf{F}_{\text{UR}}) &= \arg \min_{\mathbf{P} \geq 0, \mathbf{F} \geq 0} \frac{1}{2} \|\mathbf{P}\mathbf{F} - \mathbf{Y}\|_F^2 \\ \mathbf{A}_{\text{UR}} &= \arg \min_{\mathbf{A} \geq 0} \frac{1}{2} \|\mathbf{W}\mathbf{A} - \mathbf{P}_{\text{UR}}\|_F^2 + \gamma \mathcal{R}_2(\mathbf{A}) \end{aligned}$$

where $\mathbf{P}_{\text{UR}} \in \mathbb{R}^{J \times M}$ is a material volume, and $\mathbf{F}_{\text{UR}}, \mathbf{A}_{\text{UR}}$ are the material maps and spectral signatures respectively reconstructed by this method. Here, $\mathcal{R}_2 : \mathbb{R}^{N \times M} \mapsto \mathbb{R}$ integrates prior information about the material maps with $\gamma \geq 0$ being the regularization parameter. These practical methods only work well when a complete projection series is available and the object is composed of materials whose spectra are clearly separable. However, the spectra of materials as measured by industrial spectral tomography equipment can be very similar. Although these practical methods do not give accurate solutions, the results can be used as an initial guess for advanced reconstruction methods.

5.4.2 Ill-posedness

The Hadamard conditions to define a *well-posed* problem consist of three criteria: (i) *existence*: There must be an $\mathbf{A}^*$ and $\mathbf{F}^*$ that satisfy $\mathbf{W}\mathbf{A}^*\mathbf{F}^* = \mathbf{Y}$. (ii) *uniqueness*: The solution $\mathbf{A}^*$ and $\mathbf{F}^*$ must be unique. (iii) *stability*: small perturbations in the measurements $\mathbf{Y}$ should not lead to significant deviations in $\mathbf{A}^*$ and $\mathbf{F}^*$. If any of these conditions is violated, we call the problem *ill-posed*. In general, we assume the existence of a solution to the least-squares problem since we use the Euclidean norm to measure the misfit in the discrete setting. However, the uniqueness condition needs to be verified. Moreover, the stability of the solution relies on the conditioning of projection matrix $\mathbf{W}$ and the measurements $\mathbf{Y}$.

In general, the spectral inverse problem has multiple solutions if no prior information is incorporated. To see this, suppose $(\mathbf{A}^*, \mathbf{F}^*)$ is a solution to Equation (5.5), then $(\alpha\mathbf{A}^*, (1/\alpha)\mathbf{F}^*)$ is also a solution to (5.5) for any $\alpha > 0$. Hence, the practical reconstruction methods and classical joint method solve an ill-posed problem. To reduce the non-uniqueness, and hence make the problem less ill-posed, we need to incorporate appropriate spatio-spectral prior information.

5.5 Proposed method - ADJUST

Since the conventional spectral inverse problem remains ill-posed due to the non-uniqueness of solutions, we propose to incorporate spectral information through a spectral dictionary. The spectral profiles of many materials are readily available [130, 301], and the spectral responses in each spectral channel can easily be computed from these spectral profiles. We model the spectral response for material m as the binary combination of the dictionary elements. That is,

$$\mathbf{f}_m = \hat{r}_{m1}\mathbf{t}_1 + \hat{r}_{m2}\mathbf{t}_2 + \cdots + \hat{r}_{mD}\mathbf{t}_D,$$

where $\mathbf{t}_1, \dots, \mathbf{t}_D$ correspond to the spectral responses of D distinct dictionary materials, and $\hat{r}_{m1}, \hat{r}_{m2}, \dots, \hat{r}_{mD}$ are the coefficients of material m that take the value of either 0 or 1. Suppose the j^{th} material in the dictionary corresponds to the material m , then $\hat{r}_{mj} = 1$, and the other coefficients will be zero. Hence, we

can represent the spectral matrix $\mathbf{F} \in \mathbb{R}^{M \times C}$ as

$$\underbrace{\begin{bmatrix} - & \mathbf{f}_1 & - \\ & \vdots & \\ - & \mathbf{f}_M & - \end{bmatrix}}_{\mathbf{F}} = \underbrace{\begin{bmatrix} - & \hat{\mathbf{r}}_1 & - \\ & \vdots & \\ - & \hat{\mathbf{r}}_M & - \end{bmatrix}}_{\hat{\mathbf{R}}} \underbrace{\begin{bmatrix} | & & | \\ \mathbf{t}_1 & \dots & \mathbf{t}_D \\ | & & | \end{bmatrix}}_{\mathbf{T}},$$

where $\mathbf{T} \in \mathbb{R}^{D \times C}$ is a dictionary of D materials (with $D \geq M$) with spectral information for C channels, and $\hat{\mathbf{R}} \in \{0, 1\}^{M \times D}$ is a spectral coefficient matrix.

Due to the binary constraints, finding such a matrix $\hat{\mathbf{R}}$ jointly with $\mathbf{A}$ is a non-convex problem. To make it convex for fixed $\mathbf{A}$, we relax the binary nature of the variable $\hat{\mathbf{R}}$. Moreover, we apply additional constraints on the material maps to ensure that the total contribution of materials at every voxel does not exceed 1. Thus, the resulting formulation, termed as *A Dictionary-based Joint reconstruction and Unmixing method for Spectral Tomography* (ADJUST), is phrased as

$$\begin{aligned} \underset{\mathbf{A}, \mathbf{R}}{\text{minimize}} \quad & \mathcal{J}(\mathbf{A}, \mathbf{R}) \triangleq \frac{1}{2} \|\mathbf{Y} - \mathbf{WART}\|_F^2, \quad (\text{least-squares misfit}) \\ \text{subject to} \quad & \mathbf{A} \in \mathcal{C}_A, \quad (\text{constraints on spatial map}) \\ & \mathbf{R} \in \mathcal{C}_R, \quad (\text{constraints on dictionary coefficients}) \end{aligned} \tag{5.7}$$

where $\mathbf{Y} \in \mathbb{R}^{J \times C}$ represents spectral tomographic measurements, $\mathbf{W} \in \mathbb{R}^{J \times N}$ is a projection matrix containing the weights w_{ji} described in Equation (5.4), $\mathbf{A} \in \mathbb{R}^{N \times M}$ is the matrix that constitutes the spatial contributions of the materials, $\mathbf{T} \in \mathbb{R}^{D \times C}$ represents the fixed dictionary matrix containing attenuation spectra of many materials, and $\mathbf{R} \in \mathbb{R}^{M \times D}$ is the dictionary coefficient matrix that represents the *continuous* version of $\hat{\mathbf{R}}$. The constraint sets are

$$\begin{aligned} \mathcal{C}_R &\triangleq \left\{ \mathbf{X} \in \mathbb{R}^{M \times D} \mid \underbrace{x_{ij} \geq 0}_{(a)}, \underbrace{\sum_{j=1}^D x_{ij} \leq 1}_{(b)}, \underbrace{\sum_{i=1}^M x_{ij} \leq 1}_{(c)} \right\}, \\ \mathcal{C}_A &\triangleq \left\{ \mathbf{X} \in \mathbb{R}^{N \times M} \mid \underbrace{x_{ij} \geq 0}_{(d)}, \underbrace{\sum_{j=1}^M x_{ij} \leq 1}_{(e)} \right\}. \end{aligned}$$

We provide the details of each constraint below:

- (a) *Non-negativity of $\mathbf{R}$* : Since $\mathbf{R}$ is a dictionary-coefficient matrix that is a convex proxy for $\hat{\mathbf{R}}$, the values must be greater than or equal to 0.
- (b) *Row-sum constraints for $\mathbf{R}$* : In principle, we would like to impose that each material present in the object must be a part of the dictionary. The convex approximation of this condition is that the total contribution of the dictionary elements to represent material should not exceed 1.
- (c) *Column-sum constraints for $\mathbf{R}$* : Each column of $\mathbf{R}$ represents the contribution of the dictionary element to generate materials in the object. Since the number of materials in the object is smaller than the total number of dictionary elements, the contribution of many dictionary elements will be 0. Moreover, the materials present in the object must be distinct. Hence, the contribution of dictionary elements must not exceed 1. Hence, the column-sum constraints impose these conditions.
- (d) *Non-negativity of $\mathbf{A}$* : Each material should have a non-negative contribution to every voxel.
- (e) *Row-sum constraints for $\mathbf{A}$* : The total contribution of materials in each voxel must not exceed 1.

We enumerate the benefits of ADJUST as follows. (i) The ADJUST formulation is less ill-posed when compared with the two-step methods or the Joint formulation given in (5.6). This is due to the fact that the incorporation of the spectral dictionary resolves the scaling issue. (ii) It is a *parameter-free* approach since the constraints are simplex and do not involve any parameters that need to be estimated. The only parameter ADJUST requires is the number of materials present in the object. However, this prior knowledge is generally available to the user. (iii) The optimization problem (5.7) is bi-convex (see Appendix C.2 for the proof). Hence, it can be solved efficiently using the iterative minimization method.

5.5.1 Numerical optimization

To obtain an approximate solution to (5.7), many alternating minimization schemes exist [20, 43, 156]. However, these schemes rely on complete minimization with respect to at least one variable in every step. Moreover, their convergence to a partially optimal solution is slow. Hence, such schemes might not be computationally feasible for large-scale problems. For practical applications, we propose an accelerated variant of Proximal Alternating Linearized Minimization (PALM) [43], by combining it with the acceleration strategy in alternating direction method of multipliers (ADMM) [220]. We term this variant as ‘Alternating Accelerated Proximal Minimization’ (AAPM):

for $k = 0, \dots, K - 1$:

$$\begin{aligned}
 \mathbf{R}_{k+1} &= \text{proj}_{\mathcal{C}_R} \left(\mathbf{R}_k - \alpha \nabla_{\mathbf{R}} \tilde{\mathcal{J}}(\mathbf{A}_k, \mathbf{R}_k, \mathbf{U}_k) \right) && \text{(update dictionary} \\
 &&& \text{coefficients)} \\
 \mathbf{A}_{k+1} &= \text{proj}_{\mathcal{C}_A} \left(\mathbf{A}_k - \beta \nabla_{\mathbf{A}} \tilde{\mathcal{J}}(\mathbf{A}_k, \mathbf{R}_{k+1}, \mathbf{U}_k) \right) && \text{(update material} \\
 &&& \text{maps)} \\
 \mathbf{U}_{k+1} &= \mathbf{U}_k + \rho (\mathbf{W} \mathbf{A}_{k+1} \mathbf{R}_{k+1} \mathbf{T} - \mathbf{Y}) && \text{(running-sum-} \\
 &&& \text{of-errors)}
 \end{aligned} \tag{5.8}$$

where α and β are estimated using a line-search method (*e.g.*, backtracking), ρ is an acceleration parameter chosen from the range $[10^{-3}, 1)$, and $\tilde{\mathcal{J}}(\mathbf{A}, \mathbf{R}, \mathbf{U}) = \mathcal{J}(\mathbf{A}, \mathbf{R}) + \langle \mathbf{U}, \mathbf{Y} - \mathbf{W} \mathbf{A} \mathbf{R} \mathbf{T} \rangle$. The variable $\mathbf{U}$ contains the running sum of errors (*i.e.* residuals):

$$\mathbf{U}_{k+1} = \mathbf{U}_0 + \rho \sum_{i=1}^k \mathbf{R}_i \quad \text{where} \quad \mathbf{R}_i = \mathbf{W} \mathbf{A}_i \mathbf{R}_i \mathbf{T} - \mathbf{Y}.$$

For $\rho = 0$, AAPM is equivalent to PALM. When $\rho > 0$, the error signal $\mathbf{U}$ is driven to zero by feeding back the integral of the error to its input. In Figure 5.4, we plot the residuals versus iterations for various values of ρ . From these results, we conclude that acceleration can be achieved by including the running-sum-of-errors into an alternative iteration scheme. As the ρ values are increased, the residual decreases faster. However, for higher values of ρ , the monotonic decrease of the residual disappears. We consider ρ as a hyperparameter for which a reasonable value can be determined heuristically. The derivation of the method is given in Appendix C.3.

It is easy to compute the partial derivatives from basic linear algebra and calculus rules (refer to Appendix C.4). The partial derivatives are

$$\begin{aligned}
 \nabla_{\mathbf{A}} \tilde{\mathcal{J}}(\mathbf{A}, \mathbf{R}, \mathbf{U}) &= \mathbf{W}^T (\mathbf{W} \mathbf{A} \mathbf{R} \mathbf{T} - \mathbf{Y} - \mathbf{U}) \mathbf{T}^T \mathbf{R}^T, \\
 \nabla_{\mathbf{R}} \tilde{\mathcal{J}}(\mathbf{A}, \mathbf{R}, \mathbf{U}) &= \mathbf{A}^T \mathbf{W}^T (\mathbf{W} \mathbf{A} \mathbf{R} \mathbf{T} - \mathbf{Y} - \mathbf{U}) \mathbf{T}^T.
 \end{aligned}$$

The proximal operators are derived in Section 5.5.2. We use the following parameters to determine the stopping criteria:

$$\begin{aligned}
 \epsilon_k^{\text{abs}} &= \|\mathbf{Y} - \mathbf{W} \mathbf{A}_{k+1} \mathbf{R}_{k+1} \mathbf{T}\|_F / \|\mathbf{Y}\|_F \\
 \epsilon_k^{\text{rel}} &= \|\mathbf{A}_{k+1} - \mathbf{A}_k\|_F + \|\mathbf{R}_{k+1} - \mathbf{R}_k\|_F
 \end{aligned}$$

where k is an iteration of the optimization scheme (5.8). The benefits of AAPM are the following:

- *Simple gradient computations*: The gradients have explicit expressions and can be computed using simple matrix-matrix multiplications.

- *Fast proximal operations*: The proximal operations are just orthogonal projections onto convex sets. These operations have either an explicit expression or a simple alternating routine to find the proximal point efficiently.
- *Backtracking line-search*: The improved line-search (*i.e.* finding α, β) makes sure that the progress in the descent direction is appropriate for every iterate.
- *Acceleration through running-sum-of-errors*: The regular update of variable $\mathbf{U}$, which contains the running-sum-of-errors, helps in accelerating the convergence to the partial optimal solution.
- *Memory-efficiency*: The method relies only on forward and adjoint operations with the tomography operator $\mathbf{W}$. Hence, it saves memory to explicitly store the tomography operator in either single or double precision.




Figure 5.4: Numerical demonstration of proposed acceleration scheme on the Shepp-Logan phantom for various ρ values.

5.5.2 Orthogonal projections

In this section, we derive the projections onto the convex sets $\mathcal{C}_R$ and $\mathcal{C}_A$. For $\mathcal{C}_A = \{\mathbf{X} \in \mathbb{R}^{N \times M} \mid \mathbf{X} \geq 0, \mathbf{X}\mathbf{1} \leq \mathbf{1}\}$, its orthogonal projection takes the following form:

$$\text{proj}_{\mathcal{C}_A}(\mathbf{Z}) = \max(\mathbf{Z} - \mathbf{\Lambda}\mathbf{1}, \mathbf{0}), \quad (5.9)$$

where $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_N)$ is the diagonal matrix with weights computed from solving the inequalities

$$\langle \mathbf{1}, \max(\mathbf{z}_{i:}^T - \lambda_i \mathbf{1}, \mathbf{0}) \rangle \leq 1, \quad i = 1, \dots, N, \quad (5.10)$$

with $\mathbf{z}_i$ being the i^{th} row of the matrix $\mathbf{Z}$ [34, Theorem 6.27]. To find the optimal weight λ_i , we carry out bisection on λ_i for which $\max(\langle \mathbf{1}, \max(\mathbf{z}_{i:}^T - \lambda_i \mathbf{1}, \mathbf{0}) \rangle - 1, 0) = 0$ holds, starting with the initial interval $[0, \max(\mathbf{z}_{i:})]$. The function given by $\langle \mathbf{1}, \max(\mathbf{z}_{i:}^T - \lambda_i \mathbf{1}, \mathbf{0}) \rangle = 1$ is piecewise linear, with breakpoints at the values $z_{i1}, \dots, z_{iM}$. Hence, once we have localized λ_i to be between two adjacent values, we can immediately compute the optimal value λ_i^* . Furthermore, the following theorem entails the projection onto set $\mathcal{C}_R$:

Theorem 1. *The convex set $\mathcal{C} = \{\mathbf{X} \in \mathbb{R}^{M \times D} \mid \mathbf{X} \geq 0, \mathbf{X}\mathbf{1} \leq \mathbf{1}, \mathbf{X}^T \mathbf{1} \leq \mathbf{1}\}$ is composed of convex set $\mathcal{C}_1 = \{\mathbf{X} \in \mathbb{R}^{M \times D} \mid \mathbf{X} \geq 0, \mathbf{X}\mathbf{1} \leq \mathbf{1}\}$ and convex set $\mathcal{C}_2 = \{\mathbf{X} \in \mathbb{R}^{M \times D} \mid \mathbf{X} \geq 0, \mathbf{X}^T \mathbf{1} \leq \mathbf{1}\}$. The projection of point $\mathbf{Z} \in \mathbb{R}^{M \times D}$ onto set $\mathcal{C}$ is given by the fixed-point iteration scheme*

$$\mathbf{X}_{t+1} = \text{proj}_{\mathcal{C}_1} \left(\text{proj}_{\mathcal{C}_2} \left(\frac{\mathbf{X}_t + \mathbf{Z}}{2} \right) \right), \quad t = 0, \dots, D \quad (5.11)$$

with $\mathbf{X}_0 = \mathbf{Z}$.

Proof. The proof is given in Appendix C.1. □

For the set $\mathcal{C}_2 = \{\mathbf{X} \in \mathbb{R}^{M \times D} \mid \mathbf{X} \geq 0, \mathbf{X}^T \mathbf{1} \leq \mathbf{1}\}$, the orthogonal projection, derived from [34, Theorem 6.27], takes the form

$$\text{proj}_{\mathcal{C}_2}(\mathbf{Z}) = \max(\mathbf{Z} - \mathbf{1}\mathbf{\Omega}, \mathbf{0}), \quad (5.12)$$

with $\mathbf{\Omega} = \text{diag}(\omega_1, \dots, \omega_D)$ is the diagonal matrix with weights computed from solving the equation

$$\langle \mathbf{1}, \max(\mathbf{z}_i - \omega_i \mathbf{1}, \mathbf{0}) \rangle = 1, \quad i = 1, \dots, D. \quad (5.13)$$

5.6 Experiments and results

This section compares ADJUST with the sequential methods (RU, UR), the classical joint method (cJoint), and five state-of-the-art joint methods on a benchmark synthetic spectral phantom. After this, we compare ADJUST with the sequential methods and the classical joint method on more advanced spectral phantoms (in terms of number of materials and material shapes). Next, we examine the robustness of ADJUST against various limited measurement patterns. Finally, we show the application of ADJUST on spectral X-ray micro-CT dataset. Additional numerical experiments are presented in Appendices C.7, C.8, C.9, and C.10.

5.6.1 Experimental setup

Before presenting the results, we first outline the experimental setup. We describe the phantoms that are used, the settings for the attenuation spectra and the source spectrum, the chosen materials for each phantom and a discussion on the implementation of the algorithms.




Figure 5.5: Visualizations of the numerical phantoms used in the studies.

Spectral phantoms

A number of phantoms are used in our numerical studies, some of which are standardized while others are custom-made. All these phantoms are shown in Figure 5.5.

Mory phantom We use a slightly modified version of the phantom provided in the work by Mory et al. [202] for comparing new one-step methods against the five one-step approaches addressed in their work. The phantom contains three different materials on a 128×128 grid. As opposed to the original phantom, each location contains only one material.

Shepp-Logan phantom This standard phantom is commonly used in tomography for benchmarking. We modify this phantom to have five unique grey values. We discretize it on a 512×512 uniform grid.

Disks phantom In this custom phantom, several disks with different materials are placed on a circle. The phantom is created so that we can place up to 15 different disks on this circle. However, for our numerics, we have taken eight disks and discretized the resulting phantom on a 512×512 pixel grid.

Thorax phantom We use a modified thorax phantom provided in the CONRAD software framework [182]. We created a thorax phantom of 512^3 voxels, took slice $z = 255$, and removed a few ribs. The resulting 512×512 phantom has eight different material candidates, on which we assign five different materials.

Inverse crime refers to the process of using the same forward operator for the generation of synthetic measurements as for the subsequent reconstruction process [139]. To avoid inverse crime in all of our numerical experiments, we generate measurements by increasing the spatial resolution of the spectral phantoms by a factor of 2. For example, the Shepp-Logan phantom is discretized on a grid of 1024×1024 pixels to generate the spectral tomographic measurements. These measurements are, however, acquired on 512 (equally-spaced) detector pixels for 180 projection angles in $[0, \pi)$. The spectral tomographic inversion is then performed on a grid of 512×512 .

Attenuation spectra, source spectrum and selected materials

To generate the spectral sinograms and construct the dictionary matrix $\mathbf{T}$, we use attenuation spectra provided by the National Institute for Standards and Technology (NIST) [130, 301]. We perform linear interpolation to approximate the real spectra and discretize these such that the energetic centers are located at 100 equidistant values ranging from 5 to 35 keV. We use the corresponding interpolated attenuation values as representative attenuation values in the bins. For experiments with the Thorax phantom, we use a spectral range of [20, 80] keV. Regarding the experiments with the Mory phantom, we use the attenuation spectra provided by Mory et al. in their implementation [100], providing 100 equidistant bins with energetic centers between 20 and 119 keV. In accordance with their data preprocessing procedure, we have scaled both the material maps and the dictionary entries for iodine and gadolinium by their densities and a value of 0.01 to obtain a concentration of 10 mg/ml.

The materials chosen for each phantom (Figure 5.5) are given in Table 5.1. For the up to 42 chosen materials included in the dictionary matrix, we refer to Appendix C.5. For the Thorax phantom, the dictionary consists of the materials that appear in the phantom. This also holds for the Mory phantom.

To generate the source spectrum, we make use of the SpekPy software [47, 230]. With this, we have simulated an X-ray source spectrum from a molybdenum source with a peak voltage of 35kV. This source material provides us with a low energy bound of 5 keV for a positive flux, enabling us to include more material absorption edges (mostly K-edges) in our simulations than with other source materials such as tungsten. However, for experiments with the Thorax phantom, we use a tungsten source with a peak voltage of 80kV. Both source spectra are shown in Figure 5.6. For the Mory phantom, we use the source spectrum provided in their implementation [100], and have changed the matrix for detector response to an identity matrix.

	Shepp-Logan	Disks	Thorax	Mory
1	Vanadium	Arsenic	Bone	Iodine
2	Chromium	Selenium	90% Blood + 10% Iodine	Gadolinium
3	Manganese	Bromine	Soft tissue	Water
4	Iron	Krypton	Blood	-
5	Cobalt	Rubidium	Lung tissue	-
6	-	Strontium	-	-
7	-	Yttrium	-	-
8	-	Zirconium	-	-

Table 5.1: Selection of materials for each phantom.




Figure 5.6: The simulated molybdenum-based and tungsten-based source spectra with respectively peak voltages of 35kVp and 80kVp used in the experiments.

Implementation of the numerical algorithms

We briefly describe the implementation of various algorithms used in our studies below. For the existing one-step methods, we keep a similar naming convention as in the work of Mory et al. [202] by using the last name of the first author of the associated paper. We use the ASTRA toolbox [1, 2] for the simulation of the X-ray projections and implementation of forward and the adjoint operator of tomography [41]. For compatibility with the implementations of the existing methods in [100], we use the AIR Toolbox [113] for the computation of the forward operator and the adjoint instead of ASTRA whenever the Mory phantom is considered. For the last five listed existing one-step methods, we use the implementation and parameter values provided in [100]. Table 5.2 summarizes the prior information used in these methods.

RU For the reconstruction, we solve the Tikhonov-regularized optimization problem with regularization parameter λ set to 10^{-3} . We perform a maximum of 20 conjugate-gradient iterations on the resulting normal equations with the tolerance of 10^{-6} [123]. We use non-negative matrix factorization (NMF) with an alternating least-squares algorithm [39] for 100 maximum iterations in the decomposition step. Since NMF is a non-convex problem, we use ten different initializations to determine the solution.

UR We use the same settings for the decomposition and the reconstruction steps as described in RU.

cJoint We solve the problem using an alternating minimization scheme. The maximum number of iterations is set to 2000 with tolerance, defined as the relative residual, of 10^{-4} . In each iterate, we solve the minimization with a spectral projected gradient scheme [250].

ADJUST We use the AAPM scheme described in (5.8) to find the solution. For all the experiments, we choose ρ value of 10^{-2} and set $\epsilon_k^{\text{abs}}, \epsilon_k^{\text{rel}}$ to 10^{-4} and 10^{-6} , respectively. We run AAPM for maximum 1000 iterations.

Cai This Bayesian reconstruction approach solves a minimization problem with a non-quadratic cost function using a monotone conjugate gradient algorithm with heuristic descent steps [52]. We perform 5000 iterations to find the solution.

Long This is a regularized approach that uses Separable Quadratic Surrogates to minimize Kullback-Leibler cost function with edge-preserving regularization [172]. We run this method with 5000 iterations.

Weidinger This approach is similar to Long with modification in the approximation of regularization function by the Green potential, and leaving out Ordered Subsets that Long uses to speed up convergence [289]. We run the algorithm for 5000 iterations.

Mechlem This approach builds upon Weidinger while replacing the regularization using the Huber function. However, it uses Ordered Subsets and Nesterov acceleration to find the solution [189]. We run the algorithm for 200 iterations.

Barber This approach solves a constrained optimization problem using a primal-dual algorithm where constraints are composition of total-variation and simplex [27]. We run a maximum of 10000 primal-dual iterations.

	Spatial prior information	Spectral prior information
RU	Non-negativity	Non-negativity
UR	Non-negativity	Non-negativity
cJoint	Non-negativity	Non-negativity
Cai	(Modified) TV	Spectral signatures of present materials
Long	Simplex and (modified) TV	Spectral signatures of present materials
Weidinger	(Weak) non-negativity and (modified) TV	Spectral signatures of present materials
Mechlem	(Modified) TV	Spectral signatures of present materials
Barber	Constrained TV	Spectral signatures of present materials
ADJUST	Simplex	Spectral dictionary

Table 5.2: Spatial and spectral information for each method. TV stands for Total-Variation regularization [241]. Note that each method has prior information on the number of materials. The spectral dictionary contains signatures of a large superset of present materials.

5.6.2 Comparison of ADJUST with other methods

Results on the Mory phantom

To compare the ADJUST method with the other approaches listed in Section 5.6.1, we perform numerical studies on the Mory phantom. We take projections from 363 equidistant angles in $[0, \pi)$ using 181 detectors. Moreover, we apply Poisson noise to the resulting projections, with the incident photons being proportional to the source spectrum. We show the reconstruction results in Figure 5.7, and tabulate the performance measures in Table 5.3. We use Mean Squared Error (MSE), Peak Signal-to-Noise Ratio (PSNR), and Structural Similarity Index Measure (SSIM) to assess the reconstruction results with respect to the ground truth phantom (details of the measures are given in Appendix C.6). We list the results as the measures averaged over all material maps. We observe that ADJUST obtains the best values for PSNR, SSIM and MSE.

The two-step methods (RU and UR) find the iodine and the gadolinium locations, but the results appear to be spatially smeared out, which is also reflected in the relatively low PSNR and MSE values. However, the UR method still recovers the shapes reasonably well, as shown in the high value for the SSIM. The cJoint method finds the water location, but fails to recover the locations of iodine and gadolinium. The state-of-the-art one-step methods overall perform better than UR, RU, and cJoint. However, these methods gives rise to edge artefacts (e.g. water edges partially appearing in iodine or gadolinium maps, and blurring near the edges), resulting in suboptimal PSNR and SSIM values. These artefacts hardly




Figure 5.7: Visual comparison of various methods on Mory phantom. GT refers to ground truth.

appear in ADJUST reconstructions, resulting in significantly higher values of the PSNR and SSIM.

RU, UR and cJoint method do not incorporate any prior information about the object other than non-negativity constraints. Hence, these methods struggle to find the optimal solution. However, one-step methods strongly assume the knowledge of materials present in the object. Although they perform better than RU, UR and cJoint, they cannot be applied when the composition of object is not known. In contrast, ADJUST does not know about the material composition of the object. It solely relies on the number of materials present in the object (this number can be estimated through trial and error).

	RU	UR	cJoint	Cai	Weidinger	Long	Mechlem	Barber	ADJUST
MSE	0.0267	0.0790	0.0860	0.0115	0.0117	0.0112	0.0406	0.0115	0.0103
PSNR	20.35	16.44	16.03	19.54	19.33	19.62	14.45	19.41	21.94
SSIM	0.4413	0.6720	0.4119	0.6034	0.6131	0.6187	0.2159	0.6460	0.9616

Table 5.3: Reconstruction error in terms of Mean Squared Error (MSE), Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM) of various methods for the Mory phantom.

Results on the remaining phantoms

We compare the proposed method ADJUST with RU, UR, and Joint method on the Shepp-Logan, Disks and Thorax phantoms. Comparison with the other methods is not possible as those are designed for handling only a limited number of materials. The first two phantoms consist of only hard materials, and hence K-edges are present in the spectra. All these phantoms are more advanced compared to the Mory phantom. The Shepp-Logan and Thorax phantoms are structurally more complicated. On other hand, the Disks phantom contains up to eight hard materials. For the Thorax phantom, we also include soft materials. Nevertheless, we aim to reconstruct the bone, the iodine-blood mixture and the remaining soft materials into three separate classes. For all three phantoms, we measure tomographic projections for 180 equidistant angles from 0 to π . These measurements consist of Poisson noise that is proportional to the incoming photons on the detector. We tabulate the measures on the solutions produced by the RU, UR, cJoint and ADJUST algorithms for the three phantoms in Table 5.4. Moreover, for the Thorax phantom, we show the reconstructed material maps and the recovered spectra in Figure 5.8.




Figure 5.8: Visual comparison of ADJUST with RU, UR, and cJoint method on the Thorax phantom. The top row shows the ground truth and reconstructed material maps and spectral signatures for bone, the second row shows these for iodine and the third row shows (averages of) these for soft materials. We match the colors of the bounding box for material maps with the (recovered) spectral signatures shown on the bottom row.

Phantom		RU	UR	cJoint	ADJUST
Shepp-Logan	MSE	0.0711	0.0598	0.0548	0.0061
	PSNR	16.41	16.66	13.74	23.12
	SSIM	0.2433	0.4497	0.1077	0.9599
Disks	MSE	0.0125	0.0082	0.0063	0.0030
	PSNR	19.50	21.22	23.72	33.32
	SSIM	0.8970	0.8975	0.8882	0.9925
Thorax	MSE	0.0587	0.0622	0.0525	0.0020
	PSNR	19.92	20.96	19.11	36.68
	SSIM	0.6902	0.6094	0.7628	0.9198

Table 5.4: Reconstruction error in terms of MSE, PSNR and SSIM of various methods for various phantoms.

We see that for all measures, ADJUST outperforms the other three methods. The results for the other methods are similar to each other. In general, the Disks phantom leads to very good similarity measures for ADJUST, indicating the capability of dealing with eight materials. For the other two phantoms, where the proportions of materials are different, the measures are slightly worse, but we observe that the reconstructions for the Shepp-Logan and Disks phantoms are visually satisfactory (the visual results for the Disks phantom are given in Appendix C.7, while visual results for the Shepp-Logan phantom are given in Figure 5.3). However, for the Thorax phantom we see a striped pattern in the material map of the soft tissues. A possible reason is that the spectral signature of the combined tissue materials may not be present in the dictionary, and therefore produces visually suboptimal results. On the other hand, it can also be observed that the spectrum of bone is not fully correctly recovered. This may be because it does not have a discontinuity in the chosen spectral range, and is therefore too similar in shape to the tissue spectra. Despite this, ADJUST outperforms the other three methods on all phantoms.

5.6.3 Limited measurement patterns

Through the numerical experiments in this section, we demonstrate that ADJUST is robust. We consider three scenarios: (i) *Sparse-angle tomography*, where the number of measurements is reduced by sampling fewer projections angles, (ii) *Limited-view tomography*, where measurements from a particular range of angles are missing (representing the case of hardware limitations), and (iii) *Sparse channels*, where the number of spectral bins of the detector is limited. We apply these settings to the Shepp-Logan phantom and the Disks phantom, and report the results in Table 5.5.

For the sparse-angle tomography setup, we consider tomographic projections from 10 equidistant angles in the range of 0 to π . For the Shepp-Logan phantom, the spectral signatures are determined correctly and we observed very minor artefacts. For the Disks phantom, the material maps and the spectral signatures are precisely reconstructed, as reflected in the very low MSE and very high PSNR and SSIM values. Therefore, for these phantoms and the selected angles, ADJUST performs well.

For the limited-view tomography setup, we limit the projection angle range from 0 to $2\pi/3$. Restricting the angle range for the projections results in a well-known missing-wedge artefact. We take projections for 60 equidistant projection angles and add Poisson noise. For both phantoms, we observed no missing wedge artefacts when reconstructed with ADJUST. The PSNR and SSIM measures remain very high and the MSE measure remains low. We conclude that for these phantoms, ADJUST can deal well limited-view measurements.

In the sparse channel setting, we reduce the number of spectral bins based on the spectral dictionary. Since we consider a dictionary of 42 hard materials, we

Phantom		Full sampling	Sparse-angle	Limited-view	Sparse channels
Shepp-Logan	MSE	0.0061	0.0113	0.0330	0.0066
	PSNR	23.12	20.27	18.41	23.04
	SSIM	0.9599	0.9435	0.9112	0.9670
Disks	MSE	0.0030	0.0028	0.0057	0.0002
	PSNR	33.32	33.05	26.31	36.76
	SSIM	0.9925	0.9924	0.9807	0.9929


Table 5.5: Reconstruction error in terms of MSE, PSNR, SSIM with ADJUST for various phantoms with limited measurement pattern experiments.

reduce the spectral channels from 100 to 42. These 42 spectral channels are chosen based on the independent columns of the spectral dictionary. For tomography, we choose 60 equidistant angles between 0 and π . We observed that the Shepp-Logan phantom has been reconstructed precisely, but a few artefacts are visible on the edges of the disks for the Disks phantom. These artefacts are reflected in the slightly higher MSE and slightly lower PSNR and SSIM. For the Shepp-Logan phantom there is no obvious decrease of quality in terms of the measurements. So ADJUST appears to be capable of dealing with a sparse channel setting with the given phantoms and the spectral setup.


5.6.4 Limited spectral resolution

In this section, we investigate the performance of the algorithm when each spectral bin spans a wider energy range. This problem is also called *limited spectral resolution*. To this end, we have simulated data with 10 spectral bins instead of 100 bins for the Thorax and Shepp-Logan phantoms over the same energy range as before. The selected number of energy bins is in the same order as the number of bins resulting from the use of multi-spectral X-ray photon-counting detectors, rather than with the hundreds of energy bins of hyperspectral X-ray detectors. This results in a coarser energy resolution for both the source spectra and the attenuation spectra, and the dictionary has been updated accordingly. Apart from the different energy bins, we apply the same settings and configuration as the experiments in Section 5.6.1.


To illustrate the spectral differences, Figure 5.9 shows the attenuation spectra for the materials in the phantoms with the two different energy resolutions. The soft materials, bone and iodine present in the Thorax phantom are easily identifiable since their attenuation spectra are very different from each other in the entire spectrum (*i.e.* 20 – 80 keV). The materials in the Thorax phantom will still be separable with decrease in the spectral resolution. In the Shepp-Logan phantom, however, all included materials have K-edges very close to each other. Hence, differences resulting from the K-edge of the materials in the Shepp-Logan phantom




(a) Thorax - 100 spectral bins



(b) Shepp-Logan - 100 spectral bins



(c) Thorax - 10 spectral bins



(d) Shepp-Logan - 10 spectral bins

Figure 5.9: Attenuation spectra of different materials in the Thorax phantom (left) and the Shepp-Logan phantom right (right), partitioned over 100 energy bins (top) and 10 energy bins (bottom) over the same energy range.

Phantom		RU	UR	cJoint	ADJUST
Shepp-Logan	MSE	0.1446	0.1856	0.1980	0.0976
	PSNR	9.042	9.392	7.259	10.94
	SSIM	0.0213	0.3773	0.0209	0.6228
Thorax	MSE	0.0543	0.0715	0.0533	0.0010
	PSNR	21.22	19.64	17.95	33.05
	SSIM	0.6820	0.6579	0.7608	0.9707

Table 5.6: Reconstruction error in terms of MSE, PSNR, SSIM for various phantoms with 10 times smaller spectral resolution (10 spectral bins in total).

largely disappear. This is also reflected in the performance of various algorithms in Table 5.6, with considerably lower SSIM and PSNR compared to the full spectral resolution experiments for the Shepp-Logan phantom in Table 5.4. On the other hand, for the Thorax phantom the results are similar. Hence, we conclude that the successful application of the proposed algorithm to spectral measurements acquired in bins with larger spectral ranges depends heavily on the complexity of the problem (*i.e.* separability of the materials in the chosen energy range based on spectral responses in the bins).

5.6.5 Noisy measurement patterns

We consider three different noise levels to check the robustness of ADJUST against noise. In particular, we corrupt the spectral tomographic measurements with additive Gaussian noise of strength $\{1, 10, 20\}\%$ followed by Poisson noise with intensity corresponding to the source spectrum. In all three cases, 180 angular projections in the $[0, \pi)$ range are acquired in 100 spectral bins. The effect of noise on the spectral tomographic projections is demonstrated in Figure 5.10. In Table 5.7, we list the performance measures of ADJUST on three different phantoms. We observe a steady decrease in PSNR and SSIM values for the Shepp-Logan and Thorax phantoms. However, the PSNR and SSIM do not suffer from high noise levels (*i.e.* 20% noise) for the Disks phantom. Since the Disks phantom consists of low-rank shapes of small sizes, they can be retrieved well from noisy measurements. However, the performance of ADJUST on noisy datasets may not be extended for complex shapes, as seen from the numerical studies on the Thorax and Shepp-Logan phantom. Hence, these numerical experiments demonstrate that ADJUST is stable against a moderate amount of noise, but may not be reliable against high noise levels.




Figure 5.10: Clean and noisy sinograms (*i.e.* tomographic projections) of Shepp-Logan phantom for energy of 5.6 keV (top-row) and 7.7 keV (bottom-row).

Phantom		1% noise	10% noise	20% noise
Shepp-Logan	MSE	0.0032	0.0067	0.0187
	PSNR	25.68	22.53	18.68
	SSIM	0.9738	0.9012	0.7076
Disks	MSE	0.0001	0.0034	0.0003
	PSNR	39.20	32.44	35.53
	SSIM	0.9989	0.9779	0.9785
Thorax	MSE	0.0061	0.0129	0.0123
	PSNR	29.70	25.93	23.33
	SSIM	0.8716	0.8474	0.8464

Table 5.7: Reconstruction error in terms of MSE, PSNR, SSIM for various phantoms with three different noise levels.

5.6.6 Experiment on micro-CT data

In this subsection, we test the performance of ADJUST on a publicly available X-ray microtomography dataset [262] generated by a conventional laboratory-based CT scanner equipped with a photon-counting line detector (TESCAN Polydet) containing a semiconductor crystal (CdTe) [263]. The particle mixture sample contains pure gold, tungsten and lead, along with quartz. For a total of 128 energy bins with a spectral range of ca. 20 to 160 keV, tomographic projections are acquired on 255 equally-spaced detector pixels for 600 angles in $[0, 2\pi)$. We perform a reconstruction using the simultaneous iterative reconstruction algorithm (SIRT) [140] for each channel and determine the location of each material manually (refer to Figure 5.11). We plot the gold, lead, and tungsten spectrum obtained through channel-wise reconstructions along with the corresponding NIST spectra in Figure 5.12.

For a fair comparison, we apply RU, UR, and cJoint along with ADJUST on this spectral microtomography dataset. We use the same settings as mentioned in Section 5.6.1. For ADJUST, we use the spectral dictionary with four materials (*i.e.* gold, lead, tungsten and quartz) with their spectrum obtained from NIST (the spectrum of quartz is obtained from the reconstructed channels). Since it is evident from Figure 5.12 that the NIST spectra with appropriate scaling match closely to the spectra in the sample in the range of $[53.35, 127.05]$ keV, we reduce the spectral range of the dataset from 53.35 to 127.05 keV (amounting to 64 spectral bins in total). The reconstruction results are demonstrated in Figure 5.13. From these results, it is clear that RU, UR, and cJoint cannot precisely reconstruct the spectral signatures of materials. These results, furthermore, suggest the need of stronger spectral regularization to separate the materials. Out of all the spatial maps reconstructed with these classical methods, only the spatial map of tungsten (red) produced by the RU method approximately matches the expected material map (shown in Figure 5.11). On the other hand, the spatial maps recovered from




Figure 5.11: CT reconstruction of the 7th energy bin (a) of the spectral micro-CT dataset, in which the locations of the tungsten (orange), gold (green) and lead (blue) particles are highlighted (b).




Figure 5.12: Attenuation spectra of the materials in the sample, showing both the spectra from NIST (scaled by 0.015) and the spectra extracted from the channel-wise CT reconstructions, including quartz. The spectral range is from 20.35 keV to 161.15 keV (white region), with the bins having a spectral resolution of 1.1 keV.




Figure 5.13: Reconstruction results of methods RU, UR, cJoint, and ADJUST on the microtomography dataset of particle mixtures. The zoomed sections demonstrate the separation capabilities of ADJUST compared to RU, UR and cJoint on particles of three different materials (tungsten, gold and lead). We match the color of zoomed sections with the spectral plots for improved readability.




Figure 5.14: Detailed reconstruction results of ADJUST on X-ray microtomography dataset of particle mixtures. Reconstructed material maps for lead, tungsten and gold are shown in high resolution (left column). In addition, material maps with indications of identified particles are shown (right column). Green circles indicate correctly identified particles. None of the maps contain false negatives, and the remaining identified particles are false positives (denoted by red circles).

ADJUST match almost precisely with the expected material maps. Figure 5.14 shows that all particles are identified and that there are no false negatives. False positives are mostly small and faint, and many of these occur at the borders of other particles, mainly between gold and tungsten. As reference, Figure 5.14 also shows the reconstruction images for each material in high resolution. Finally, we note that the NIST spectra do not overlap entirely with the spectra recovered through channel-wise reconstruction. Therefore, the spectral dictionary must be calibrated for the spectral detector setup for various materials. We expect improvement in reconstruction results of ADJUST with a calibrated spectral dictionary.

5.7 Conclusions and discussion

Spectral imaging is an emerging topic in X-ray tomography since it adds an additional dimension to the measurements, which can be exploited to retrieve the material composition of the object of interest. Recently, joint approaches (also known as one-step methods) have emerged as a promising technique for solving the spectral imaging problem by incorporating all the prior information in a single step. These joint approaches reduce the ill-posedness of the spectral imaging problem. However, the spectral signatures of many materials are very similar, making the joint approaches likely to fail when many materials are involved. To tackle the problems with the joint approach, we propose the ADJUST framework that integrates the structure of spectral signatures by creating a dictionary of all the known materials and uses this to jointly reconstruct and carry out material decomposition in a single step. Since the resulting formulation is a bi-convex optimization problem, we propose an Alternating Accelerated Proximal Minimization (AAPM) scheme to find a solution. Through numerical experiments, we show that ADJUST performs better over practical methods as well as state-of-the-art joint approaches on various simulated phantoms.

Obtaining projections from all directions requires high experimental time. Moreover, X-ray machines may not allow for sampling in all directions. Hence, practical methods do not help in determining the material composition of structurally more complicated objects (either because of intricate structures or a wider variety of materials) with most X-ray configurations. However, ADJUST is robust against limited tomography measurement patterns on phantoms that are more complicated.

A natural question is to check if the utilization of a spectral dictionary (*i.e.* representing $\mathbf{F}$ as $\mathbf{RT}$) in two-step methods can produce optimal results. For example, in UR method, we can first decompose spectral measurements $\mathbf{Y} = \mathbf{ZRT}$, where $\mathbf{Z}$ incorporates the projections per material. Later, the spatial maps can be extracted by solving $\mathbf{WA} = \mathbf{Z}$. However, in the decomposition step, we allow for unrealistic projections due to lack of knowledge of tomography operator $\mathbf{W}$ and spatial properties of $\mathbf{A}$, especially when dealing with limited measurements. Moreover, allowing unrealistic projections leads to unrealistic materials since simplex

constraints on $\mathbf{R}$ give rise to convex combination of (many) dictionary elements, and not of only a single element. A similar argument holds for RU method combined with the spectral dictionary. These unrealistic solutions are, however, penalized by proper *spatio-spectral regularization* in the ADJUST framework. Hence, ADJUST will always perform better than spectral dictionary versions of two-step methods even when complete measurements are available.

There are some limitations to the ADJUST framework. It can only separate hard materials from each other and separate hard materials from soft materials. Moreover, if neither the spectral signature of the material is present in the dictionary, nor it can be composed as a linear combination of the elements from the dictionary, ADJUST will fail in recovering that material. Although we tested ADJUST against moderate Gaussian and Poisson noise, it is not straightforward to assume that ADJUST will behave stably against real (extremely) noisy datasets that are common in energy-dispersive X-ray tomography.

The experiments conducted on the micro-CT dataset suggest the potential of ADJUST for application to experimental spectral CT data. Apart from the noise level on experimental data, the successful application of ADJUST depends on several factors. As shown in the experiments with the limited spectral resolution, the number and range of the energy bins, combined with the problem's complexity, determines the degree of possible material decomposition. If the spectral resolution is too small such that materials cannot be distinguished anymore, ADJUST will not yield proper material decompositions (and neither will other methods). Moreover, for successful reconstruction and decomposition, there is a limit to the number of (differentiable) materials that can be included in the dictionary for a given spectral resolution.

For an uncalibrated spectral detector, the spectral dictionary can be measured, if it is not available through the manufacturer. Additionally, when working with real data, an interesting consideration for future work would be to estimate the source spectrum along with the spatial material maps and their signatures. Although our framework is based on the assumption that the spectral tomographic measurements consist of additive white noise, we can extend it to tackle Poisson noise by replacing the least-squares loss by Kullback-Leibler function. However, the bi-convexity can no longer be guaranteed and the solution obtained through AAPM may not be partially optimal. We leave this extension for future work.

Code and data availability

The source code of ADJUST, along with the RU, UR, and cJoint algorithms, are available on <https://github.com/mzeegers/ADJUST> [138]. These MATLAB codes make use of open-source toolboxes, in particular the ASTRA toolbox [1, 2], Spot Operator toolbox [37], MinConf optimization package [250]. The scripts for

testing the mentioned algorithms on the phantoms and the data are also made open-source in the ADJUST Github repository. The micro-CT spectral dataset is available at <https://rodare.hzdr.de/record/1627> [262].

