Numerical aspects of artifacts in tomography

Tristan van Leeuwen, Utrecht University

International Workshop on Industrial Tomography 2015

Artifacts in X-ray tomography

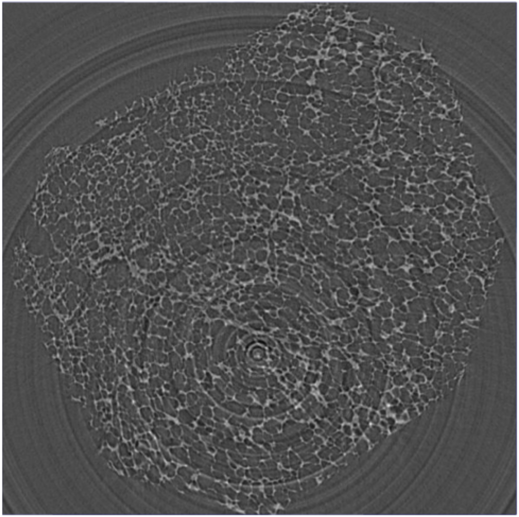

Courtesy of Lucia Mancini

Artifacts in X-ray tomography

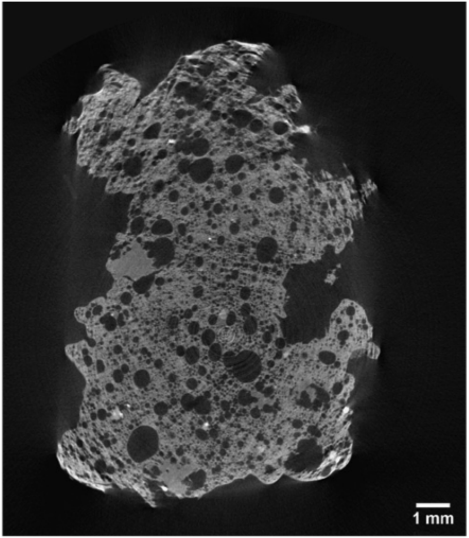

Courtesy of Lucia Mancini

Artifacts in X-ray tomography

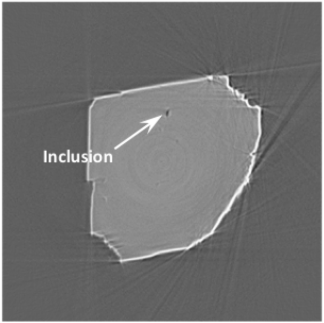

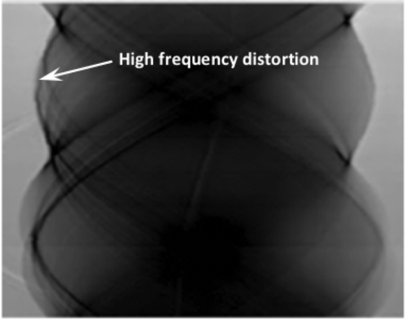

Courtesy of Francesco de Carlo

Artifacts

Physical causes of artifacts

- Noise
- Beam hardening
- Motion blurring
- Scatter
- Limited field of view

Artifacts

Physical causes of artifacts

- Noise
- Beam hardening
- Motion blurring
- Scatter
- Limited field of view

What about the mathematical causes of artefacts?

Overview

- What is an artefact?
- Algebraic reconstruction
- Maximum likelihood estimation
- Examples
- Conclusions

What *is* an artifact

Artefact Something observed in a scientific investigation or experiment that is not naturally present but occurs as a result of the *preparative* or *investigative* procedure.

What *is* an artifact

- Things we don't *want* to see

What *is* an artifact

- Things we don't *want* to see

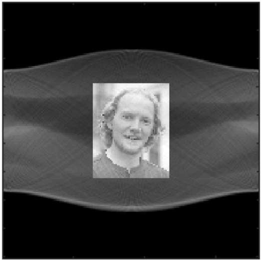

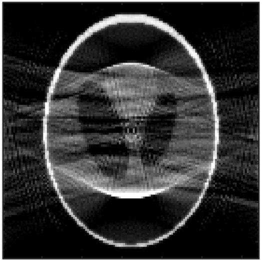

What *is* an artifact

- Things we don't *expect* to see

What *is* an artifact

- Things we don't *expect* to see

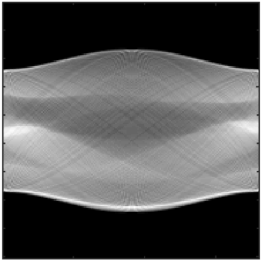

Causes of artifacts

Measurement model

$$\mathbf{p} = W\mathbf{x} + \mathbf{n},$$

Prior model

$$\mathbf{x} = \mathbf{x}_0 + R\mathbf{s},$$

where

W - Projection matrix
$\mathbf{x}$ - image
$\mathbf{p}$ - projection data
$\mathbf{n}$ - noise
$\mathbf{x_0}, R$ - prior information

Causes of artifacts

- Artifacts are caused by making inappropriate assumptions on either the measurement model (W or $\mathbf{n}$) or the prior model (R, $\mathbf{x}_0$ and $\mathbf{s}$).

Causes of artifacts

- Artifacts are caused by making inappropriate assumptions on either the measurement model (W or $\mathbf{n}$) or the prior model (R, $\mathbf{x}_0$ and $\mathbf{s}$).
- We *always* make such assumptions!

Algebraic reconstruction

Find $\mathbf{x}$ such that

$$W\mathbf{x} \approx \mathbf{p} = W\bar{\mathbf{x}} + \mathbf{n}.$$

Algebraic reconstruction

Find $\mathbf{x}$ such that

$$W\mathbf{x} \approx \mathbf{p} = W\overline{\mathbf{x}} + \mathbf{n}.$$

Many algebraic reconstruction techniques can be expressed as

$$\mathbf{x}^{k+1} = \left(I - B^T W\right)\mathbf{x}^k + B^T\mathbf{p},$$

where B is a *preconditioned* version of W:

Algebraic reconstruction

Find $\mathbf{x}$ such that

$$W\mathbf{x} \approx \mathbf{p} = W\overline{\mathbf{x}} + \mathbf{n}.$$

Many algebraic reconstruction techniques can be expressed as

$$\mathbf{x}^{k+1} = \left(I - B^T W\right)\mathbf{x}^k + B^T\mathbf{p},$$

where B is a *preconditioned* version of W:

- ART: $B = HW$,
- SIRT: $B = RWC$,
- Landweber: $B = \lambda W$.

Algebraic reconstruction

Decompose the projection matrix as

$$W = U\Sigma V^T = \sum_i \sigma_i \mathbf{u}_i \mathbf{v}_i^T,$$

and write the pseudo-inverse as

$$W^\dagger = \sum_i \alpha_i \mathbf{v}_i \mathbf{u}_i^T,$$

where $\alpha_i = \sigma_i^{-1}$ when $\sigma_i > 0$ and $\alpha_i = 0$ otherwise.

Algebraic reconstruction

We can express any (algebraic) reconstruction in terms of a *generalized* singular value decomposition

$$\widehat{\mathbf{x}} = \sum_i \alpha_i \widetilde{\mathbf{v}}_i \widetilde{\mathbf{u}}_i^T \mathbf{p},$$

where $U^T R U = I$ and $V^T L V = I$ and $\alpha_i \simeq \sigma_i^{-1}$.

Algebraic reconstruction

- Reconstruction algorithms approximate the pseudo-inverse in various ways.
- Which (if any) null-space elements end up in the final reconstruction depends on the reconstruction algorithm we use.
- To what extend the noise term $\mathbf{n}$ ends up in the reconstruction also depends on the reconstruction algorithm.

Algebraic reconstruction

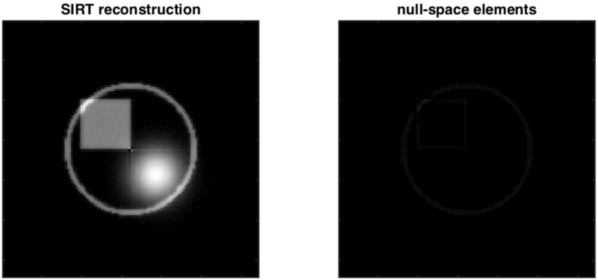

Algebraic reconstruction

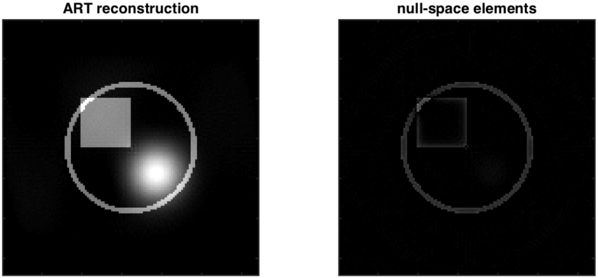

Algebraic reconstruction

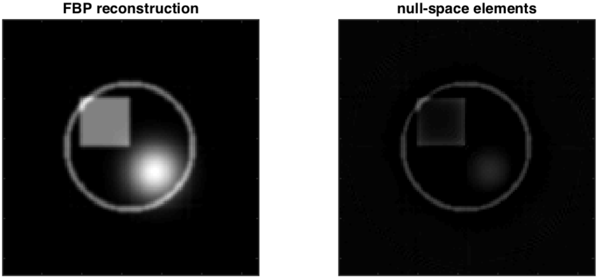

Algebraic reconstruction

Gaussian noise

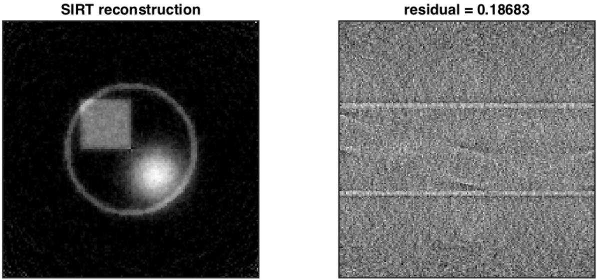

Algebraic reconstruction

Gaussian noise

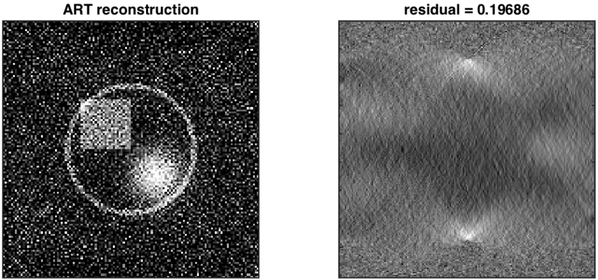

Algebraic reconstruction

Gaussian noise

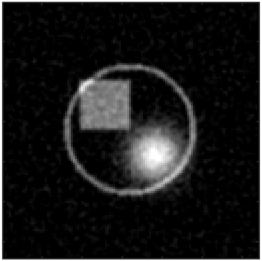

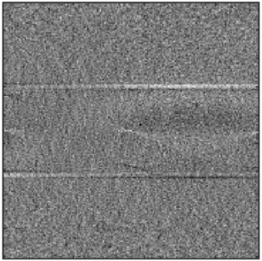

Algebraic reconstruction

Constant offset noise

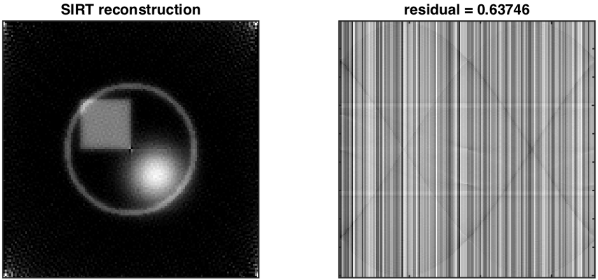

Algebraic reconstruction

Constant offset noise

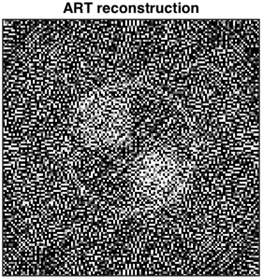

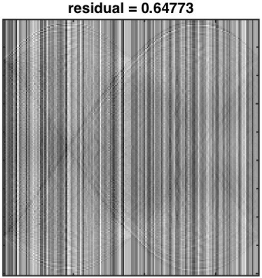

Algebraic reconstruction

Constant offset noise

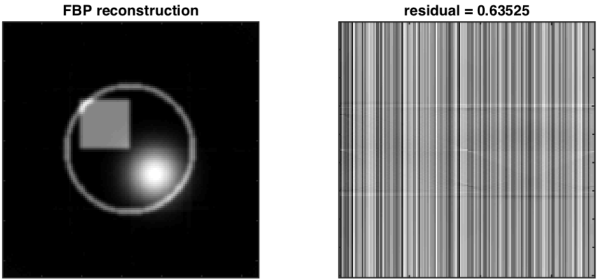

Algebraic reconstruction

Dead pixels

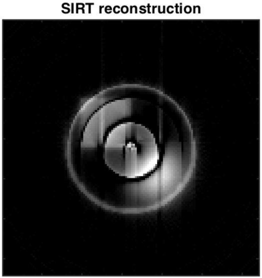

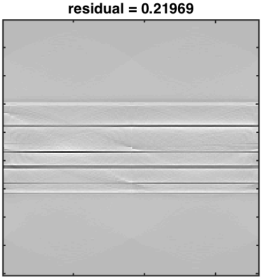

Algebraic reconstruction

Dead pixels

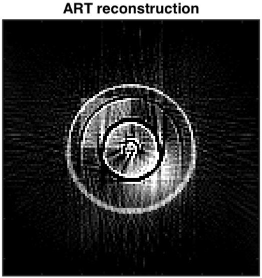

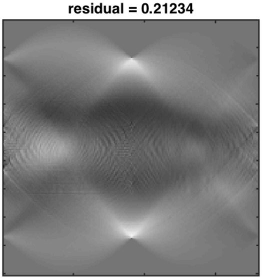

Algebraic reconstruction

Dead pixels

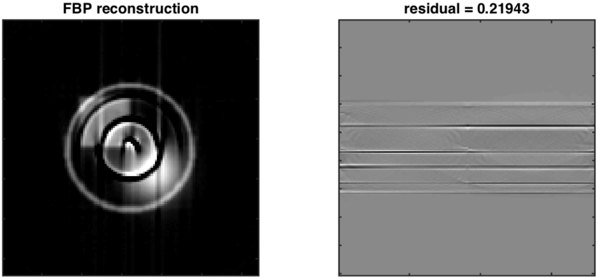

Maximum Likelihood estimation

All reconstruction algorithms make *implicit* assumptions about the noise or the grey-values.

Maximum Likelihood estimation

All reconstruction algorithms make *implicit* assumptions about the noise or the grey-values.

We can make the assumptions explicit by modelling them as random processes, e.g. Gaussian

$$\mathbf{n} \sim \mathcal{N}(\mathbf{n}_0, C_{\mathbf{n}}),$$

$$\mathbf{x} \sim \mathcal{N}(\mathbf{x}_0, C_{\mathbf{x}})$$

Maximum Likelihood estimation

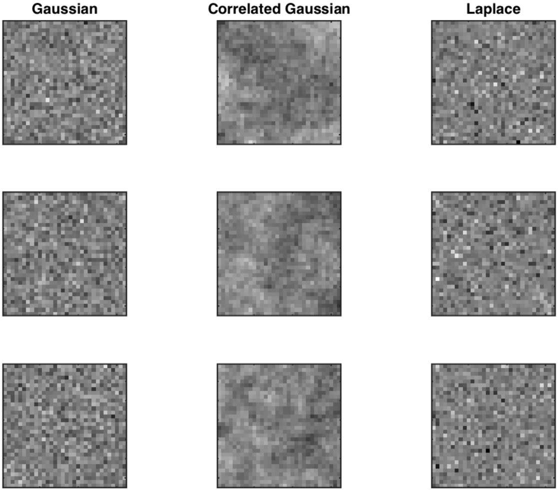

Maximum Likelihood estimation

We can find the most likely $\mathbf{x}$ by maximizing the probability

$$\mathcal{P}(\mathbf{x};\mathbf{p}) \propto \exp\left(-\|W\mathbf{x}-\mathbf{p}-\mathbf{n}_0\|^2_{C_{\mathbf{n}}^{-1}} - \|\mathbf{x}-\mathbf{x}_0\|^2_{C_{\mathbf{x}}^{-1}}\right),$$

or, equivalently, minimizing

$$\mathcal{L}(\mathbf{x};\mathbf{p}) = \|W\mathbf{x}-\mathbf{p}-\mathbf{n}_0\|^2_{C_{\mathbf{n}}^{-1}} + \|\mathbf{x}-\mathbf{x}_0\|^2_{C_{\mathbf{x}}^{-1}}.$$

Maximum Likelihood estimation

For Gaussian noise and correlated Gaussian grey-values

$$\mathcal{L}(\mathbf{x};\mathbf{p}) = \sigma^{-2}\|W\mathbf{x} - \mathbf{p}\|_2^2 + \lambda^2\|D\mathbf{x}\|_2^2,$$

where D is the discretized gradient operator.

Maximum Likelihood estimation

For Gaussian noise and correlated Laplace grey-values

$$\mathcal{L}(\mathbf{x};\mathbf{p}) = \sigma^{-2}\|W\mathbf{x}-\mathbf{p}\|_2^2 + \lambda\|D\mathbf{x}\|_1,$$

where $\|\cdot\|_1$ is the ℓ_1 norm.

Maximum Likelihood estimation

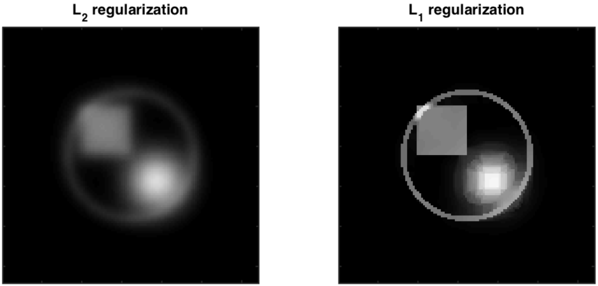

Examples: Students T

- The noise is most likely *not* Gaussian.
- We can use a noise model that more closely reflects the presence of large outliers.

Students T:

$$\mathcal{P}(\mathbf{r}) \propto \prod_i \left(1 + r_i^2\right).$$

Examples: Students T

We now formulate the reconstruction problem as

$$\min_x \rho\left(W\mathbf{x} - \mathbf{p}\right),$$

where

$$\rho(\mathbf{r}) = \sum_i \log(1 + r_i^2).$$

Examples: Students T

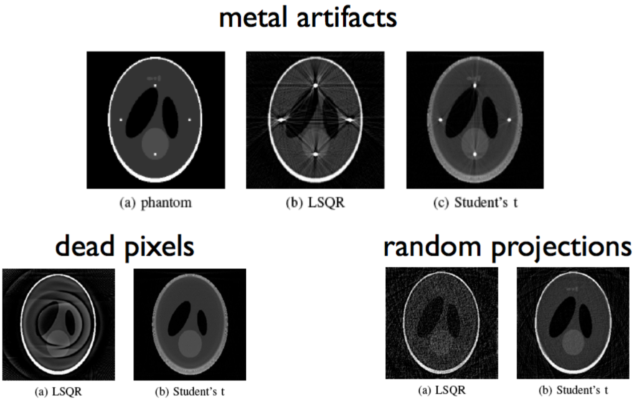

Courtesy of Folkert Bleichrodt

Examples: Alignment

There can also be an uncertainty in W itself, for example due to misalignment.

Examples: Alignment

There can also be an uncertainty in W itself, for example due to misalignment.

We can model this as a small perturbation of the projection matrix: $W = \widetilde{W} + E$

Examples: Alignment

There can also be an uncertainty in W itself, for example due to misalignment.

We can model this as a small perturbation of the projection matrix: $W = \widetilde{W} + E$

The MLE problem is given by

$$\min_{W,\mathbf{x}} \|W\mathbf{x} - \mathbf{p}\|_2^2 + \lambda \|W - \widetilde{W}\|_F^2.$$

Examples: Alignment

A more realistic way to model this is by parametrizing the projection operator as $W(\mathbf{a})$, where $\mathbf{a}$ are the alignment parameters (shifts, angles, rotations).

Examples: Alignment

A more realistic way to model this is by parametrizing the projection operator as $W(\mathbf{a})$, where $\mathbf{a}$ are the alignment parameters (shifts, angles, rotations).

The MLE problem is given by

$$\min_{\mathbf{a},\mathbf{x}} \|W(\mathbf{a})\mathbf{x} - \mathbf{p}\|_2^2,$$

plus any assumptions on $\mathbf{x}$ and $\mathbf{a}$.

Examples: Alignment

We can tackle this problem with an alternating approach

- full reconstruction based on $\mathbf{a}_k$,
- update of alignment based on new reconstruction.

Examples: Alignment

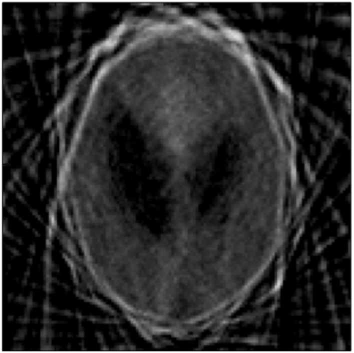

Examples: Alignment

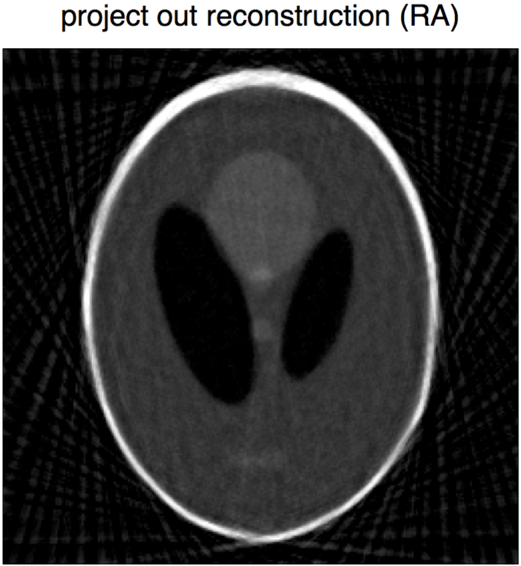

Conclusions

- Artifacts are a combination of null-space elements and reconstructed noise.
- All reconstruction algorithms make *implicit* assumptions about the null-space elements and the noise.
- We make these assumptions *explicit* by casting reconstruction as a maximum likelihood estimation problem.
- If the assumptions are violated, such approaches usually worse results than classical methods.

Conclusions

- Ideally we want to test various noise models and priors to investigate the uncertainty

Conclusions

- Ideally we want to test various noise models and priors to investigate the uncertainty
- A naive approach would require many iterative reconstructions

Conclusions

- Ideally we want to test various noise models and priors to investigate the uncertainty
- A naive approach would require many iterative reconstructions
- Can we approximate the action of advanced reconstruction algorithms efficiently?

Conclusions

- Ideally we want to test various noise models and priors to investigate the uncertainty
- A naive approach would require many iterative reconstructions
- Can we approximate the action of advanced reconstruction algorithms efficiently?
- If we can do this, how do we compare different results and visualize the uncertainty?

Algebraic reconstruction

Null space Vectors $\mathbf{z}$ for which $W\mathbf{z} = 0$

- Some components in the image do not affect the data
- This cause non-uniqueness of the solution

Algebraic reconstruction

Column space Vectors $\mathbf{y}$ for which we can find a vector $\mathbf{x}$ such that $\mathbf{y} = W\mathbf{x}$

- Some noise components can be explained by the model.
- This causes artefacts in the reconstruction

Algebraic reconstruction

Consider the measurement model only

$$W\mathbf{x} = \mathbf{p} + \mathbf{n}.$$

Split the true image $\bar{\mathbf{x}} = \bar{\mathbf{x}}_r + \bar{\mathbf{x}}_n$ with $W\bar{\mathbf{x}}_n = 0$.

Algebraic reconstruction

Consider the measurement model only

$$W\mathbf{x} = \mathbf{p} + \mathbf{n}.$$

Split the true image $\bar{\mathbf{x}} = \bar{\mathbf{x}}_r + \bar{\mathbf{x}}_n$ with $W\bar{\mathbf{x}}_n = 0$.

Express the data as

$$\mathbf{p} = W\bar{\mathbf{x}} = W\bar{\mathbf{x}}_r.$$

Algebraic reconstruction

Consider the measurement model only

$$W\mathbf{x} = \mathbf{p} + \mathbf{n}.$$

Split the true image $\bar{\mathbf{x}} = \bar{\mathbf{x}}_r + \bar{\mathbf{x}}_n$ with $W\bar{\mathbf{x}}_n = 0$.

Express the data as

$$\mathbf{p} = W\bar{\mathbf{x}} = W\bar{\mathbf{x}}_r.$$

We interpret $\mathbf{n}$ as everything else.

Algebraic reconstruction

The (least-squares) solution is given by

$$\widehat{\mathbf{x}} = \mathbf{W}^{\dagger}\mathbf{p} + \widehat{\mathbf{x}}_n + \mathbf{W}^{\dagger}\mathbf{n},$$

where $W^{\dagger}$ is the *pseudo-inverse* and $W\widehat{\mathbf{x}}_n = 0$.

Algebraic reconstruction

The (least-squares) solution is given by

$$\widehat{\mathbf{x}} = \mathbf{W}^{\dagger}\mathbf{p} + \widehat{\mathbf{x}}_n + \mathbf{W}^{\dagger}\mathbf{n},$$

where $W^{\dagger}$ is the *pseudo-inverse* and $W\widehat{\mathbf{x}}_n = 0$.

Ideally, we have

$$\widehat{\mathbf{x}} = \bar{\mathbf{x}}_r + \widehat{\mathbf{x}}_n + \mathbf{W}^{\dagger}\mathbf{n}.$$

So, if $\|\mathbf{n}\|_2$ is small, $\widehat{\mathbf{x}}_r \approx \bar{\mathbf{x}}_r$ but in general $\widehat{\mathbf{x}}_n \not\approx \bar{\mathbf{x}}_n$.

Algebraic reconstruction

Gallery of singular vectors

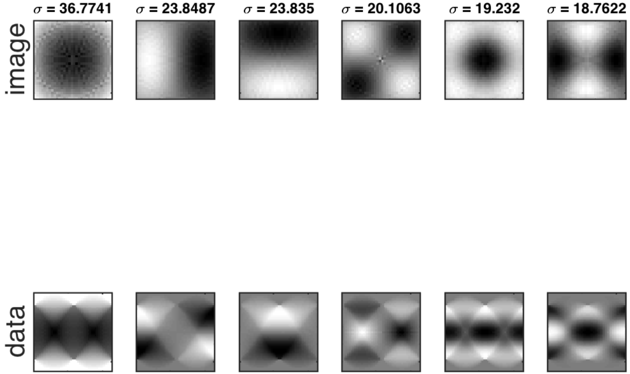

Linear algebra

Gallery of singular vectors

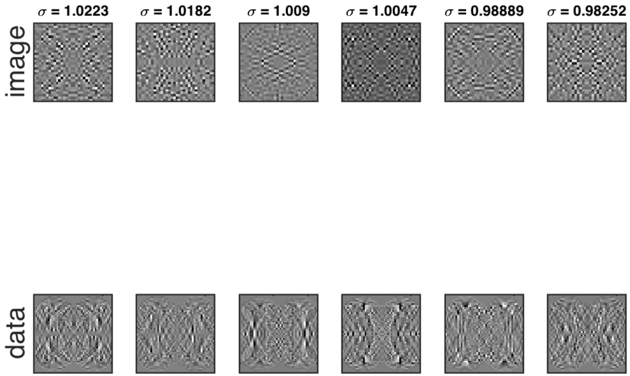

Linear algebra

Gallery of singular vectors

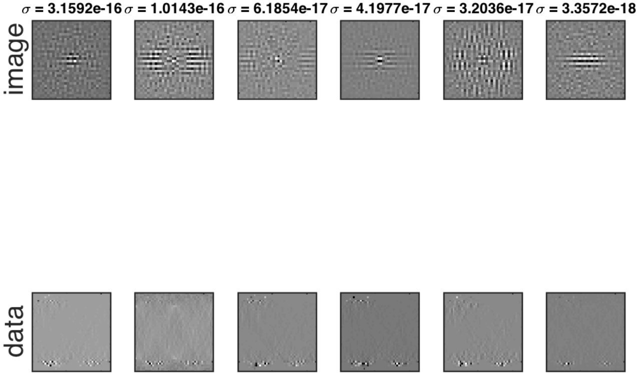

Algebraic reconstruction

We can think of these as minimizing the residual in a weighted norm

$$\min_{\mathbf{x}} \|W\mathbf{x} - \mathbf{p}\|_Q^2.$$

Algebraic reconstruction

We can think of these as minimizing the residual in a weighted norm

$$\min_{\mathbf{x}} \|W\mathbf{x} - \mathbf{p}\|_Q^2.$$

- ART: $Q = H$,
- SIRT: $Q = R$.