


Solving Regularized Inverse Problems with ADMM



CSC2529

David Lindell

University of Toronto

cs.toronto.edu/~lindell/teaching/2529

losslandscape.com

*slides adapted from Gordon Wetzstein

Announcements

- HW5 due Wednesday 2/11
- HW6 is out
- No class next week (reading week)
- Proposal due in 2.5 weeks!
- See website for all office hours/problem session dates

Overview

- Brief review of deconvolution with inverse/Wiener filtering
- A Bayesian perspective of inverse problems
- Image priors/regularization and total variation
- The ADMM method
- Image deconvolution with ADMM
- Compressive imaging

Must read: course notes on Deconvolution and Compressive Imaging!

Overview

- Brief review of deconvolution with inverse/Wiener filtering
- A Bayesian perspective of inverse problems
- Image priors/regularization and total variation
- The ADMM method
- Image deconvolution with ADMM
- Compressive imaging

Image Deconvolution – Brief Review



Given: blurry & noisy image

Desired: sharp & noise-free image

Image Deconvolution – Brief Review

- Image formation model:

$$b = c * x + \eta$$

2D measurements

known 2D convolution kernel

2D target image

additive noise




Image Deconvolution – Brief Review

- Image formation model: $b = c * x + \eta$
- Convolution theorem: $b = \mathcal{F}^{-1}\{\mathcal{F}\{c\} \cdot \mathcal{F}\{x\}\} + \eta$
- Inverse filtering: $\tilde{x}_{\text{if}} = \mathcal{F}^{-1}\left\{\frac{\mathcal{F}\{b\}}{\mathcal{F}\{c\}}\right\}$
- Wiener filtering: $\tilde{x}_{\text{wf}} = \mathcal{F}^{-1}\left\{\frac{|\mathcal{F}\{c\}|^2}{|\mathcal{F}\{c\}|^2 + 1/SNR} \cdot \frac{\mathcal{F}\{b\}}{\mathcal{F}\{c\}}\right\}$
- Duality of “signal processing” and “algebraic” interpretation:


$$b = c * x \Leftrightarrow \mathbf{b} = \mathbf{C}\mathbf{x} \qquad \mathbf{C} \in \mathbb{R}^{N \times N}, \quad \mathbf{b}, \mathbf{x} \in \mathbb{R}^N$$

Image Deconvolution – Inverse Filtering


Ground Truth




No Noise



$\sigma=0.1$





$\sigma=1.0$



Measurements

Reconstructions




$$\tilde{x}_{\text{if}} = \mathcal{F}^{-1} \left\{ \frac{\mathcal{F}\{b\}}{\mathcal{F}\{c\}} \right\}$$

Image Deconvolution – Wiener Filtering


Ground Truth




No Noise



$\sigma=0.1$




$\sigma=1.0$



Measurements

Reconstructions



$$\tilde{x}_{\text{wf}} = \mathcal{F}^{-1} \left\{ \frac{|\mathcal{F}\{c\}|^2}{|\mathcal{F}\{c\}|^2 + 1/SNR} \cdot \frac{\mathcal{F}\{b\}}{\mathcal{F}\{c\}} \right\}$$

Image Deconvolution

- Problem: this is an ill-posed inverse problem, i.e., there are infinitely many solutions that satisfy the measurements
- Need some way to determine how “desirable” any one of these feasible solutions is → need an **image prior**

Overview

- Brief review of deconvolution with inverse/Wiener filtering
- A Bayesian perspective of inverse problems
- Image priors/regularization and total variation
- The ADMM method
- Image deconvolution with ADMM
- Compressive imaging

A Bayesian Perspective of Inverse Problems

- Image formation model: $\mathbf{b} = \mathbf{A}\mathbf{x} + \boldsymbol{\eta}$, $\mathbf{b} \in \mathbb{R}^M, \mathbf{x} \in \mathbb{R}^N, \mathbf{A} \in \mathbb{R}^{M \times N}$

- Interpret as random variables:
$$\mathbf{x}_i \sim \mathcal{N}(\mathbf{x}_i, 0), \quad \boldsymbol{\eta}_i \sim \mathcal{N}(0, \sigma^2)$$
$$\mathbf{b}_i \sim \mathcal{N}((\mathbf{A}\mathbf{x})_i, \sigma^2)$$

- Probability of observation i :
$$p(\mathbf{b}_i | \mathbf{x}_i, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(\mathbf{b}_i - (\mathbf{A}\mathbf{x})_i)^2}{2\sigma^2}}$$

- Joint probability of all observations:
$$p(\mathbf{b} | \mathbf{x}, \sigma) = \prod_{i=1}^M p(\mathbf{b}_i | \mathbf{x}_i, \sigma) \propto e^{-\frac{\|\mathbf{b} - \mathbf{A}\mathbf{x}\|_2^2}{2\sigma^2}}$$

A Bayesian Perspective of Inverse Problems

- Bayes' rule:
$$p(\mathbf{x}|\mathbf{b}, \sigma) = \frac{p(\mathbf{b}|\mathbf{x}, \sigma)p(\mathbf{x})}{p(\mathbf{b})} \propto p(\mathbf{b}|\mathbf{x}, \sigma)p(\mathbf{x})$$

\uparrow
posterior

\uparrow
image formation model

\uparrow
prior

- Maximum-a-posterior (MAP) solution:

$$\begin{aligned}\mathbf{x}_{MAP} &= \arg \min_{\mathbf{x}} -\log(p(\mathbf{x}|\mathbf{b}, \sigma)) \\ &= \arg \min_{\mathbf{x}} -\log(p(\mathbf{b}|\mathbf{x}, \sigma)) - \log(p(\mathbf{x})) \\ &= \arg \min_{\mathbf{x}} \frac{1}{2\sigma^2} \|\mathbf{b} - \mathbf{A}\mathbf{x}\|_2^2 + \Psi(\mathbf{x})\end{aligned}$$

A Bayesian Perspective of Inverse Problems

- Terminology: $\Psi(\mathbf{x}) = -\log(p(\mathbf{x}))$


$$\mathbf{x}_{MAP} = \arg \min_{\mathbf{x}} \frac{1}{2\sigma^2} \|\mathbf{b} - \mathbf{A}\mathbf{x}\|_2^2 + \Psi(\mathbf{x})$$

Overview

- Brief review of deconvolution with inverse/Wiener filtering
- A Bayesian perspective of inverse problems
- Image priors/regularization and total variation
- The ADMM method
- Image deconvolution with ADMM
- Compressive imaging

Examples of Image Priors / Regularizers

blurry stuff



stars



“natural” image



Promote smoothness!

$$\Psi(\mathbf{x}) = \|\Delta \mathbf{x}\|_2$$



Laplace operator

Promote sparsity!

$$\Psi(\mathbf{x}) = \|\mathbf{x}\|_1$$

Promote sparse gradients!

$$\Psi(\mathbf{x}) = \text{TV}(\mathbf{x})$$


Total Variation (TV)

express (forward finite difference)
gradient as convolution!

\mathbf{x}

$$\mathbf{D}_x \mathbf{x} = d_x * x, d_x = \begin{bmatrix} 0 & 0 & 0 \\ 0 & -1 & 1 \\ 0 & 0 & 0 \end{bmatrix}$$

$$\mathbf{D}_y \mathbf{x} = d_y * x, d_y = \begin{bmatrix} 0 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 1 & 0 \end{bmatrix}$$



-0.3 0.3

Total Variation (TV)



\mathbf{x}

better: isotropic

$$\sqrt{(\mathbf{D}_x \mathbf{x})_i^2 + (\mathbf{D}_y \mathbf{x})_i^2}$$

easier: anisotropic

$$\sqrt{(\mathbf{D}_x \mathbf{x})_i^2} + \sqrt{(\mathbf{D}_y \mathbf{x})_i^2}$$



Total Variation (TV)

- Examples are mostly black, indicating that gradient magnitudes are close to 0 \rightarrow natural images have sparse gradients!
- This intuition is well-captured by the TV pseudo-norm, either anisotropic or isotropic:

$$\text{TV}_{\text{anisotropic}}(\mathbf{x}) = \|\mathbf{D}_x \mathbf{x}\|_1 + \|\mathbf{D}_y \mathbf{x}\|_1 = \sum_{i=1}^N |(\mathbf{D}_x \mathbf{x})_i| + |(\mathbf{D}_y \mathbf{x})_i| = \sum_{i=1}^N \sqrt{(\mathbf{D}_x \mathbf{x})_i^2} + \sqrt{(\mathbf{D}_y \mathbf{x})_i^2}$$

$$\text{TV}_{\text{isotropic}}(\mathbf{x}) = \|\mathbf{D} \mathbf{x}\|_{2,1} = \sum_{i=1}^N \left\| \begin{bmatrix} (\mathbf{D}_x \mathbf{x})_i \\ (\mathbf{D}_y \mathbf{x})_i \end{bmatrix} \right\|_2 = \sum_{i=1}^N \sqrt{(\mathbf{D}_x \mathbf{x})_i^2 + (\mathbf{D}_y \mathbf{x})_i^2}$$

Total Variation (TV)

The TV pseudo-norm is one of the most popular regularization schemes for natural images!

Extensions to make it more general or applicable for other data:

- Hyper-Laplacian: Levin et al. 2009, Krishnan & Fergus 2009
- Total generalized variation: Bredies et al. 2009
- Frobenius norm of Hessian: Lefkimmiatis et al. 2003
- ...

How to solve inverse problem that
use these regularizers?

Solving Regularized Inverse Problem

- Objective or “loss” function of general inverse problem:
$$\text{minimize}_x \frac{1}{2} \|\mathbf{b} - \mathbf{A}\mathbf{x}\|_2^2 + \lambda \Psi(\mathbf{x})$$

↑
weight of regularizer
- Practical #1 go-to solution: Adam solver implemented in PyTorch
- 3 simple steps, will explore in problem session & homework:
 1. Implement evaluation of loss function
 2. Set hyperparameters, including learning rate
 3. Run
- The “fine print”: convenient but doesn’t always converge well

Overview

- Brief review of deconvolution with inverse/Wiener filtering
- A Bayesian perspective of inverse problems
- Image priors/regularization and total variation
- The Alternating Direction Method of Multipliers (ADMM)
- Image deconvolution with ADMM
- Compressive imaging

Regularized Image Reconstruction

$$\underset{\{x\}}{\text{minimize}} \quad \frac{1}{2} \|b - Ax\|_2^2 + \lambda \Gamma(x)$$

data fidelity
term

some image prior, such as
 ℓ_1 norm or others

Regularized Image Reconstruction

$$\begin{array}{ll} \underset{\{x\}}{\text{minimize}} & \underbrace{\frac{1}{2} \|b - Ax\|_2^2}_{f(x)} + \underbrace{\lambda \Gamma(z)}_{g(z)} \\ \text{subject to} & Kx - z = 0 \end{array}$$

- split into two parts \rightarrow mathematically equivalent

Regularized Image Reconstruction

$$\begin{array}{ll} \underset{\{x\}}{\text{minimize}} & \underbrace{\frac{1}{2} \|b - Ax\|_2^2}_{f(x)} + \underbrace{\lambda \Gamma(z)}_{g(z)} \\ \text{subject to} & Kx - z = 0 \end{array}$$

$$L(x, z, y) = f(x) + g(z) + y^T (Kx - z)$$

- Lagrangian

Regularized Image Reconstruction

- Lagrangian

$$L_{\rho}(x, z, y) = f(x) + g(z) + y^T (Kx - z)$$

Optimal if all partial derivatives are zero!

$$\nabla_{x,z,y} L = 0$$

Regularized Image Reconstruction

- Lagrangian

$$L_{\rho}(x, z, y) = f(x) + g(z) + y^T (Kx - z)$$

Also implies:

$$\nabla_{x,z,y} L = 0 \Leftrightarrow \nabla_{x,z} f(x) + g(z) = y^T \nabla_{x,z} (Kx - z)$$
$$Kx - z = 0$$

Regularized Image Reconstruction

- Lagrangian

$$L_{\rho}(x, z, y) = f(x) + g(z) + y^T (Kx - z)$$

Also implies:

$$\nabla_{x,z,y} L = 0 \Leftrightarrow \nabla_{x,z} f(x) + g(z) = y^T \nabla_{x,z} (Kx - z)$$

$$\boxed{Kx - z = 0}$$

Regularized Image Reconstruction

- Lagrangian


$$L_{\rho}(x, z, y) = f(x) + g(z) + y^T (Kx - z)$$

Also implies:

$$\nabla_{x,z,y} L = 0 \Leftrightarrow \boxed{\nabla_{x,z} f(x) + g(z) = y^T \nabla_{x,z} (Kx - z)}$$
$$Kx - z = 0$$

Regularized Image Reconstruction

$$\nabla_{x,z,y} L = 0 \Leftrightarrow \nabla_{x,z} f(x) + g(z) = y^T \nabla_{x,z} (Kx - z)$$



Regularized Image Reconstruction

- Augmented Lagrangian

$$L_{\rho}(x, z, y) = f(x) + g(z) + y^T (Kx - z) + \frac{\rho}{2} \|Kx - z\|_2^2$$

Regularized Image Reconstruction

- Augmented Lagrangian

$$L_{\rho}(x, z, y) = f(x) + g(z) + y^T (Kx - z) + \boxed{\frac{\rho}{2} \|Kx - z\|_2^2}$$

We add quadratic penalty: improves convergence properties compared to standard Lagrangian

ρ is called the “penalty parameter” (see Boyd 2011)

Regularized Image Reconstruction

- Scaled dual form of the augmented Lagrangian
 - Given by some algebraic manipulation (easier form to work with)

$$L_{\rho}(x, z, y) = f(x) + g(z) + \frac{\rho}{2} \|Kx - z + u\|^2 + \frac{\rho}{2} \|u\|^2$$

Where $u = y/\rho$ is the scaled dual variable

Regularized Image Reconstruction

repeat until converged

$$x \leftarrow \operatorname{prox}_{\|\cdot\|_2, \rho}(v) = \arg \min_{\{x\}} L_{\rho}(x, z, y)$$

$$z \leftarrow \operatorname{prox}_{\Gamma, \rho}(v) = \arg \min_{\{z\}} L_{\rho}(x, z, y)$$

$$u \leftarrow u + Kx - z$$

- iterative updates - ADMM

Regularized Image Reconstruction

repeat until converged

$$x \leftarrow \operatorname{prox}_{\|\cdot\|_2, \rho}(v) = \arg \min_{\{x\}} L_{\rho}(x, z, y) = \arg \min_{\{x\}} \frac{1}{2} \|Ax - b\|_2^2 + \frac{\rho}{2} \|Kx - v\|, v = z - u$$

$$z \leftarrow \operatorname{prox}_{\Gamma, \rho}(v) = \arg \min_{\{z\}} L_{\rho}(x, z, y) = \arg \min_{\{z\}} \lambda \Gamma(z) + \frac{\rho}{2} \|v - z\|, v = Kx + u$$

$$u \leftarrow u + Kx - z$$

- iterative updates - ADMM

Regularized Image Reconstruction

$$\text{prox}_{\|\cdot\|_2, \rho}(v) = \arg \min_{\{x\}} \frac{1}{2} \|Ax - b\|_2^2 + \frac{\rho}{2} \|Kx - v\|$$

... see notes ...

Regularized Image Reconstruction

$$\text{prox}_{\|\cdot\|_2, \rho}(v) = \arg \min_{\{x\}} \frac{1}{2} \|Ax - b\|_2^2 + \frac{\rho}{2} \|Kx - v\|$$

... see notes ...

$$\text{prox}_{\|\cdot\|_2, \rho}(v) = \left(\underbrace{A^T A + \rho K^T K}_{\tilde{A}} \right)^{-1} \left(\underbrace{A^T b + \rho K^T v}_{\tilde{b}} \right)$$


- x-update: solve $\tilde{A}x = \tilde{b}$
- symmetric, positive definite matrix \rightarrow conjugate gradient method

Overview

- Brief review of deconvolution with inverse/Wiener filtering
- A Bayesian perspective of inverse problems
- Image priors/regularization and total variation
- The Alternating Direction Method of Multipliers (ADMM)
- Image deconvolution with ADMM
- Compressive imaging

ADMM for Image Deconvolution with TV

Generic: $L_\rho(\mathbf{x}, \mathbf{z}, \mathbf{u}) = f(\mathbf{x}) + g(\mathbf{z}) + \frac{\rho}{2} \|\mathbf{K}\mathbf{x} - \mathbf{z} + \mathbf{u}\|_2^2 + \frac{\rho}{2} \|\mathbf{u}\|_2^2$

Deconv: $L_\rho(\mathbf{x}, \mathbf{z}, \mathbf{u})$ 
 $= \frac{1}{2} \|\mathbf{C}\mathbf{x} - \mathbf{b}\|_2^2 + \lambda \|\mathbf{z}\|_1 + \frac{\rho}{2} \|\mathbf{D}\mathbf{x} - \mathbf{z} + \mathbf{u}\|_2^2 + \frac{\rho}{2} \|\mathbf{u}\|_2^2$

$\mathbf{x} \in \mathbb{R}^N$ unknown sharp image

$\mathbf{C} \in \mathbb{R}^{N \times N}$ circulant convolution matrix for known kernel c

$\mathbf{z}, \mathbf{u} \in \mathbb{R}^{2N}$ slack/dual variable, twice the size of \mathbf{x} !

$\mathbf{D} = \begin{bmatrix} \mathbf{D}_x \\ \mathbf{D}_y \end{bmatrix} \in \mathbb{R}^{2N \times N}$ finite difference gradients, horizontal & vertical

ADMM for Image Deconvolution with TV

$$L_\rho(\mathbf{x}, \mathbf{z}, \mathbf{u}) = \frac{1}{2} \|\mathbf{C}\mathbf{x} - \mathbf{b}\|_2^2 + \lambda \|\mathbf{z}\|_1 + \frac{\rho}{2} \|\mathbf{D}\mathbf{x} - \mathbf{z} + \mathbf{u}\|_2^2 + \frac{\rho}{2} \|\mathbf{u}\|_2^2$$

while not converged:

$$\mathbf{x} \leftarrow \text{prox}_{\|\cdot\|_2, \rho}(\mathbf{v}) = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{C}\mathbf{x} - \mathbf{b}\|_2^2 + \frac{\rho}{2} \|\mathbf{D}\mathbf{x} - \mathbf{v}\|_2^2, \mathbf{v} = \mathbf{z} - \mathbf{u}$$

$$\mathbf{z} \leftarrow \text{prox}_{\|\cdot\|_1, \rho}(\mathbf{v}) = \arg \min_{\mathbf{z}} \lambda \|\mathbf{z}\|_1 + \frac{\rho}{2} \|\mathbf{v} - \mathbf{z}\|_2^2, \mathbf{v} = \mathbf{D}\mathbf{x} + \mathbf{u}$$

ADMM for Image Deconvolution with TV

$$L_\rho(\mathbf{x}, \mathbf{z}, \mathbf{u}) = \frac{1}{2} \|\mathbf{C}\mathbf{x} - \mathbf{b}\|_2^2 + \lambda \|\mathbf{z}\|_1 + \frac{\rho}{2} \|\mathbf{D}\mathbf{x} - \mathbf{z} + \mathbf{u}\|_2^2 + \frac{\rho}{2} \|\mathbf{u}\|_2^2$$

$$\text{prox}_f(v) = \underset{x}{\operatorname{argmin}} f(x) + \frac{1}{2} \|x - v\|_2^2$$

while

$$\mathbf{x} \leftarrow \text{prox}_{\|\cdot\|_2, \rho}(\mathbf{v}) = \arg \min_x \frac{1}{2} \|\mathbf{C}\mathbf{x} - \mathbf{b}\|_2^2 + \frac{\rho}{2} \|\mathbf{D}\mathbf{x} - \mathbf{v}\|_2^2, \mathbf{v} = \mathbf{z} - \mathbf{u}$$

$$\mathbf{z} \leftarrow \text{prox}_{\|\cdot\|_1, \rho}(\mathbf{v}) = \arg \min_z \lambda \|\mathbf{z}\|_1 + \frac{\rho}{2} \|\mathbf{v} - \mathbf{z}\|_2^2, \mathbf{v} = \mathbf{D}\mathbf{x} + \mathbf{u}$$

ADMM for Image Deconvolution with TV

\mathbf{x} - update:

$$\mathbf{x} \leftarrow \text{prox}_{\|\cdot\|_2, \rho}(\mathbf{v}) = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{C}\mathbf{x} - \mathbf{b}\|_2^2 + \frac{\rho}{2} \|\mathbf{D}\mathbf{x} - \mathbf{v}\|_2^2$$

↙ reformulate

$$= \frac{1}{2} (\mathbf{C}\mathbf{x} - \mathbf{b})^T (\mathbf{C}\mathbf{x} - \mathbf{b}) + \frac{\rho}{2} (\mathbf{D}\mathbf{x} - \mathbf{v})^T (\mathbf{D}\mathbf{x} - \mathbf{v})$$

$$= \frac{1}{2} (\mathbf{x}^T \mathbf{C}^T \mathbf{C} \mathbf{x} - 2\mathbf{x}^T \mathbf{C}^T \mathbf{b} + \mathbf{b}^T \mathbf{b}) + \frac{\rho}{2} (\mathbf{x}^T \mathbf{D}^T \mathbf{D} \mathbf{x} - 2\mathbf{x}^T \mathbf{D}^T \mathbf{v} + \mathbf{v}^T \mathbf{v})$$



find solution by setting gradient to 0

$$0 = \nabla_{\mathbf{x}} f(\mathbf{x}) = \mathbf{C}^T \mathbf{C} \mathbf{x} - \mathbf{C}^T \mathbf{b} + \rho \mathbf{D}^T \mathbf{D} \mathbf{x} - \rho \mathbf{D}^T \mathbf{v}$$



closed-form solution

$$\mathbf{x} \leftarrow (\mathbf{C}^T \mathbf{C} + \rho \mathbf{D}^T \mathbf{D})^{-1} (\mathbf{C}^T \mathbf{b} + \rho \mathbf{D}^T \mathbf{v})$$

ADMM for Image Deconvolution with TV

\mathbf{x} - update:

$$\mathbf{x} \leftarrow \text{prox}_{\|\cdot\|_2, \rho}(\mathbf{v}) = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{C}\mathbf{x} - \mathbf{b}\|_2^2 + \frac{\rho}{2} \|\mathbf{D}\mathbf{x} - \mathbf{v}\|_2^2$$

$$\mathbf{x} \leftarrow \underbrace{(\mathbf{C}^T \mathbf{C} + \rho \mathbf{D}^T \mathbf{D})}^{-1} \underbrace{(\mathbf{C}^T \mathbf{b} + \rho \mathbf{D}^T \mathbf{v})}$$

Exploit duality of algebraic & signal processing interpretation

$$\mathbf{C}^T \mathbf{C} \Leftrightarrow \mathcal{F}^{-1}\{\mathcal{F}\{c\}^* \cdot \mathcal{F}\{c\}\}$$

$$\mathbf{D}^T \mathbf{z} = \mathbf{D}_x^T \mathbf{v}_1 + \mathbf{D}_y^T \mathbf{v}_2 \Leftrightarrow \mathcal{F}^{-1}\{\mathcal{F}\{d_x\}^* \cdot \mathcal{F}\{v_1\} + \mathcal{F}\{d_y\}^* \cdot \mathcal{F}\{v_2\}\}$$

$$\mathbf{D}^T \mathbf{D} \Leftrightarrow \mathcal{F}^{-1}\{\mathcal{F}\{d_x\}^* \cdot \mathcal{F}\{d_x\} + \mathcal{F}\{d_y\}^* \cdot \mathcal{F}\{d_y\}\}$$

$$\mathbf{C}^T \mathbf{b} \Leftrightarrow \mathcal{F}^{-1}\{\mathcal{F}\{c\}^* \cdot \mathcal{F}\{b\}\}$$

$$\underbrace{\mathbf{C}^T \mathbf{C} + \rho \mathbf{D}^T \mathbf{D}} \Leftrightarrow \mathcal{F}^{-1}\{\mathcal{F}\{c\}^* \cdot \mathcal{F}\{c\} + \rho(\mathcal{F}\{d_x\}^* \cdot \mathcal{F}\{d_x\} + \mathcal{F}\{d_y\}^* \cdot \mathcal{F}\{d_y\})\}$$

$$\underbrace{\mathbf{C}^T \mathbf{b} + \rho \mathbf{D}^T \mathbf{v}} \Leftrightarrow \mathcal{F}^{-1}\{\mathcal{F}\{c\}^* \cdot \mathcal{F}\{b\} + \rho(\mathcal{F}\{d_x\}^* \cdot \mathcal{F}\{v_1\} + \mathcal{F}\{d_y\}^* \cdot \mathcal{F}\{v_2\})\}$$

ADMM for Image Deconvolution with TV

\mathbf{x} - update:

$$\mathbf{x} \leftarrow \text{prox}_{\|\cdot\|_2, \rho}(\mathbf{v}) = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{C}\mathbf{x} - \mathbf{b}\|_2^2 + \frac{\rho}{2} \|\mathbf{D}\mathbf{x} - \mathbf{v}\|_2^2$$

$$\mathbf{x} \leftarrow (\mathbf{C}^T \mathbf{C} + \rho \mathbf{D}^T \mathbf{D})^{-1} (\mathbf{C}^T \mathbf{b} + \rho \mathbf{D}^T \mathbf{v})$$

- Efficient \mathbf{x} -update operates purely on 2D images with FFTs and element-wise multiplications and divisions:

$$\text{prox}_{\|\cdot\|_2, \rho}(\mathbf{z}) = \mathcal{F}^{-1} \left(\frac{\mathcal{F}\{c\}^* \cdot \mathcal{F}\{b\} + \rho(\mathcal{F}\{d_x\}^* \cdot \mathcal{F}\{v_1\} + \mathcal{F}\{d_y\}^* \cdot \mathcal{F}\{v_2\})}{\mathcal{F}\{c\}^* \cdot \mathcal{F}\{c\} + \rho(\mathcal{F}\{d_x\}^* \cdot \mathcal{F}\{d_x\} + \mathcal{F}\{d_y\}^* \cdot \mathcal{F}\{d_y\})} \right)$$



can pre-compute most parts

$$v_1 = \mathbf{v}(1:N), v_2 = \mathbf{v}(N+1:2N)$$

ADMM for Image Deconvolution with TV

\mathbf{z} – update:

$$\mathbf{z} \leftarrow \text{prox}_{\|\cdot\|_1, \rho}(\mathbf{v}) = \arg \min_{\mathbf{z}} \lambda \|\mathbf{z}\|_1 + \frac{\rho}{2} \|\mathbf{v} - \mathbf{z}\|_2^2$$

- Efficient \mathbf{z} –update uses element-wise soft thresholding operator $\mathcal{S}_\kappa(\cdot)$:

$$\text{prox}_{\|\cdot\|_1, \rho}(\mathbf{v}) = \mathcal{S}_\kappa(\mathbf{v}) = \begin{cases} v - \kappa & v > \kappa \\ 0 & |v| \leq \kappa \\ v + \kappa & v < -\kappa \end{cases} = (v - \kappa)_+ - (-v - \kappa)_+$$

$\kappa = \lambda/\rho$

This element-wise soft thresholding is the proximal operator for anisotropic TV, see course notes on block soft thresholding for isotropic TV.

$$\mathbf{v} = \mathbf{D}\mathbf{x} + \mathbf{u}$$

ADMM for Image Deconvolution with Denoiser

\mathbf{x} – update:

$$\mathbf{x} \leftarrow \text{prox}_{\|\cdot\|_2, \rho}(\mathbf{v}) = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{C}\mathbf{x} - \mathbf{b}\|_2^2 + \frac{\rho}{2} \|\mathbf{x} - \mathbf{v}\|_2^2 \quad \mathbf{v} \in \mathbb{R}^N$$

$$\mathbf{x} \leftarrow (\mathbf{C}^T \mathbf{C} + \rho \mathbf{I})^{-1} (\mathbf{C}^T \mathbf{b} + \rho \mathbf{v}) \quad \text{no matrix } \mathbf{D}!$$

- Efficient \mathbf{x} –update operates purely on 2D images with FFTs and element-wise multiplications and divisions:

$$\text{prox}_{\|\cdot\|_2, \rho}(\mathbf{v}) = \mathcal{F}^{-1} \left\{ \frac{\mathcal{F}\{c\}^* \cdot \mathcal{F}\{b\} + \rho \mathcal{F}\{v\}}{\mathcal{F}\{c\}^* \cdot \mathcal{F}\{c\} + \rho} \right\}$$

ADMM for Image Deconvolution with Denoiser

\mathbf{z} – update:

$$\begin{aligned}\mathbf{z} \leftarrow \text{prox}_{\mathcal{D},\rho}(\mathbf{v}) &= \arg \min_{\mathbf{z}} \lambda \Psi(\mathbf{z}) + \frac{\rho}{2} \|\mathbf{v} - \mathbf{z}\|_2^2 \quad \mathbf{v} = \mathbf{x} + \mathbf{u} \\ &= \arg \min_{\mathbf{z}} \Psi(\mathbf{z}) + \frac{\rho}{2\lambda} \|\mathbf{v} - \mathbf{z}\|_2^2\end{aligned}$$

- Efficient \mathbf{z} –update uses arbitrary denoiser $\mathcal{D}(\cdot)$, such as DnCNN and non-local means, using noise variance $\sigma^2 = \frac{\lambda}{\rho}$

$$\text{prox}_{\mathcal{D},\rho}(\mathbf{x}) = \mathcal{D}\left(\mathbf{x}, \sigma^2 = \frac{\lambda}{\rho}\right)$$

Image Deconvolution with ADMM


Target Image



Adam+TV, PSNR 26.1 dB




Measurements, $\sigma=0.1$



ADMM+TV, PSNR 26.3 dB



Wiener Deconv., PSNR 19.5 dB



ADMM+DnCNN, PSNR 26.7 dB




Image Deconvolution with ADMM

ADMM for deconvolution with TV


```
1: initialize  $\rho$  and  $\lambda$ 
2:  $x = \text{zeros}(W, H)$ ;
3:  $z = \text{zeros}(W, H)$ ;
4:  $u = \text{zeros}(W, H)$ ;
5: for  $k = 1$  to  $\text{max\_iters}$  do
6:    $v = z - u$ 
7:    $x = \text{prox}_{\|\cdot\|_2, \rho}(v) = \mathcal{F}^{-1} \left\{ \frac{\mathcal{F}\{c\}^* \cdot \mathcal{F}\{b\} + \rho(\mathcal{F}\{d_x\}^* \cdot \mathcal{F}\{v_1\} + \mathcal{F}\{d_y\}^* \cdot \mathcal{F}\{v_2\})}{\mathcal{F}\{c\}^* \cdot \mathcal{F}\{c\} + \rho(\mathcal{F}\{d_x\}^* \cdot \mathcal{F}\{d_x\} + \mathcal{F}\{d_y\}^* \cdot \mathcal{F}\{d_y\})} \right\}$ 
8:    $v = \mathbf{D}x + u$ 
9:    $z = \text{prox}_{\|\cdot\|_1, \rho}(v) = \mathcal{S}_{\lambda/\rho}(v)$ 
10:   $u = u + \mathbf{D}x - z$ 
11: end for
```

ADMM for deconvolution with denoiser

```
1: initialize  $\rho$  and  $\lambda$ 
2:  $x = \text{zeros}(W, H)$ ;
3:  $z = \text{zeros}(W, H)$ ;
4:  $u = \text{zeros}(W, H)$ ;
5: for  $k = 1$  to  $\text{max\_iters}$  do
6:    $v = z - u$ 
7:    $x = \text{prox}_{\|\cdot\|_2, \rho}(v) = \mathcal{F}^{-1} \left\{ \frac{\mathcal{F}\{c\}^* \cdot \mathcal{F}\{b\} + \rho \mathcal{F}\{v\}}{\mathcal{F}\{c\}^* \cdot \mathcal{F}\{c\} + \rho} \right\}$ 
8:    $v = x + u$ 
9:    $z = \text{prox}_{\mathcal{D}, \rho}(v) = \mathcal{D}\left(v, \sigma^2 = \frac{\lambda}{\rho}\right)$ 
10:   $u = u + x - z$ 
11: end for
```

ADMM - Convergence Criterion

- Run or “unroll” ADMM for K iterations
- Run until change in residual between iterations is $<$ threshold


$$\begin{aligned}v &= z - u \\x &= \mathbf{prox}_{\|\cdot\|_2, \rho}(v) = \mathcal{F}^{-1} \left\{ \frac{\mathcal{F}\{c\}^* \cdot \mathcal{F}\{b\} + \rho \mathcal{F}\{v\}}{\mathcal{F}\{c\}^* \cdot \mathcal{F}\{c\} + \rho} \right\} \\v &= x + u \\z &= \mathbf{prox}_{\mathcal{D}, \rho}(v) = \mathcal{D} \left(v, \sigma^2 = \frac{\lambda}{\rho} \right) \\u &= u + x - z\end{aligned}$$

$$\begin{aligned}v &= z - u \\x &= \mathbf{prox}_{\|\cdot\|_2, \rho}(v) = \mathcal{F}^{-1} \left\{ \frac{\mathcal{F}\{c\}^* \cdot \mathcal{F}\{b\} + \rho \mathcal{F}\{v\}}{\mathcal{F}\{c\}^* \cdot \mathcal{F}\{c\} + \rho} \right\} \\v &= x + u \\z &= \mathbf{prox}_{\mathcal{D}, \rho}(v) = \mathcal{D} \left(v, \sigma^2 = \frac{\lambda}{\rho} \right) \\u &= u + x - z\end{aligned}$$

$$\begin{aligned}v &= z - u \\x &= \mathbf{prox}_{\|\cdot\|_2, \rho}(v) = \mathcal{F}^{-1} \left\{ \frac{\mathcal{F}\{c\}^* \cdot \mathcal{F}\{b\} + \rho \mathcal{F}\{v\}}{\mathcal{F}\{c\}^* \cdot \mathcal{F}\{c\} + \rho} \right\} \\v &= x + u \\z &= \mathbf{prox}_{\mathcal{D}, \rho}(v) = \mathcal{D} \left(v, \sigma^2 = \frac{\lambda}{\rho} \right) \\u &= u + x - z\end{aligned}$$

⋮

Outlook on Unrolled Optimization

- Run or “unroll” ADMM for K iterations
- Interpret as unrolled feedforward network:

$$v = z - u$$

$$x = \mathbf{prox}_{\|\cdot\|_2, \rho}(v) = \mathcal{F}^{-1} \left\{ \frac{\mathcal{F}\{c\}^* \cdot \mathcal{F}\{b\} + \rho \mathcal{F}\{v\}}{\mathcal{F}\{c\}^* \cdot \mathcal{F}\{c\} + \rho} \right\}$$

$$v = x + u$$

$$z = \mathbf{prox}_{\mathcal{D}, \rho}(v) = \mathcal{D} \left(v, \sigma^2 = \frac{\lambda}{\rho} \right)$$

$$u = u + x - z$$

$$v = z - u$$

$$x = \mathbf{prox}_{\|\cdot\|_2, \rho}(v) = \mathcal{F}^{-1} \left\{ \frac{\mathcal{F}\{c\}^* \cdot \mathcal{F}\{b\} + \rho \mathcal{F}\{v\}}{\mathcal{F}\{c\}^* \cdot \mathcal{F}\{c\} + \rho} \right\}$$

$$v = x + u$$

$$z = \mathbf{prox}_{\mathcal{D}, \rho}(v) = \mathcal{D} \left(v, \sigma^2 = \frac{\lambda}{\rho} \right)$$

$$u = u + x - z$$

$$v = z - u$$


$$x = \mathbf{prox}_{\|\cdot\|_2, \rho}(v) = \mathcal{F}^{-1} \left\{ \frac{\mathcal{F}\{c\}^* \cdot \mathcal{F}\{b\} + \rho \mathcal{F}\{v\}}{\mathcal{F}\{c\}^* \cdot \mathcal{F}\{c\} + \rho} \right\}$$

$$v = x + u$$

$$z = \mathbf{prox}_{\mathcal{D}, \rho}(v) = \mathcal{D} \left(v, \sigma^2 = \frac{\lambda}{\rho} \right)$$

$$u = u + x - z$$

⋮




Outlook on Unrolled Optimization

- Run or “unroll” ADMM for K iterations
- Interpret as unrolled feedforward network:

Benefits over unrolled optimization

- Learnable parameters: $\lambda^{(k)}, \rho^{(k)}$, denoiser $\mathcal{D}^{(k)}$
- DenseNet-like skip connections
- Denoiser/regularizer can adapt to matrix \mathbf{C}
- Can train with advanced loss functions (perceptual, adversarial, other network, ...)




Overview


- Brief review of deconvolution with inverse/Wiener filtering
- A Bayesian perspective of inverse problems
- Image priors/regularization and total variation
- The Alternating Direction Method of Multipliers (ADMM)
- Image deconvolution with ADMM
- Compressive imaging

Compressive Imaging

Single-pixel Imaging



Single-pixel Imaging



original



10%

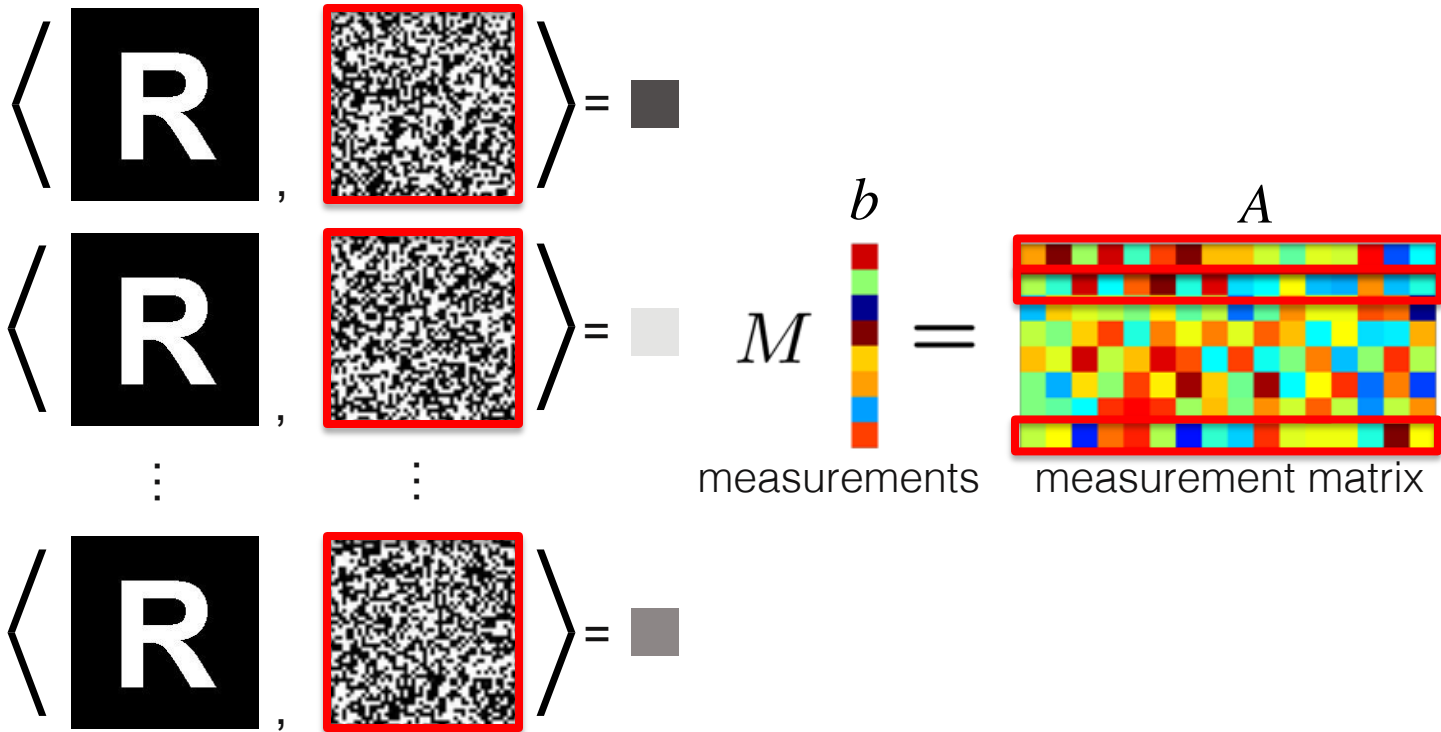


5%



2%

Single-pixel Imaging



Under-determined Inverse Problems

- Image formation model: $\mathbf{b} = \mathbf{A}\mathbf{x} + \boldsymbol{\eta}$, $\mathbf{b} \in \mathbb{R}^M, \mathbf{x} \in \mathbb{R}^N, \mathbf{A} \in \mathbb{R}^{M \times N}$
- What makes it under-determined (or a compressive imaging problem):
 $M < N$
- Problem: infinitely many solutions satisfy the observations!
Same problem as ill-posed problems! \rightarrow need **image priors**


Under-determined Inverse Problems

- Image formation model: $\mathbf{b} = \mathbf{A}\mathbf{x} + \boldsymbol{\eta}$, $\mathbf{b} \in \mathbb{R}^M, \mathbf{x} \in \mathbb{R}^N, \mathbf{A} \in \mathbb{R}^{M \times N}$
- Standard approach – the least-norm solution: $\tilde{\mathbf{x}}_{\text{ln}} = \mathbf{A}^T(\mathbf{A}\mathbf{A}^T)\mathbf{b}$
- This is the solution of optimization problem
$$\begin{array}{ll}\text{minimize}_{\mathbf{x}} & \|\mathbf{x}\|_2 \\ \text{subject to} & \mathbf{A}\mathbf{x} = \mathbf{b}\end{array}$$

Note: among the infinitely many solutions satisfying the observations, the least-norm solution is the one with the smallest L2 norm, thus equivalent to $\|\cdot\|_2$ regularizer


Under-determined Inverse Problems

- Image formation model: $\mathbf{b} = \mathbf{A}\mathbf{x} + \boldsymbol{\eta}$, $\mathbf{b} \in \mathbb{R}^M$, $\mathbf{x} \in \mathbb{R}^N$, $\mathbf{A} \in \mathbb{R}^{M \times N}$
- Standard approach – the least-norm solution: $\tilde{\mathbf{x}}_{\text{ln}} = \mathbf{A}^T(\mathbf{A}\mathbf{A}^T)\mathbf{b}$
- Results (not great):




Other Inverse Problems in Imaging

Images: Wikipedia






Computed Tomography (CT)



Magnetic Resonance Imaging (MRI)

Fourier domain



Hyperspectral Imaging

- Computational photography
- Light field imaging
- Thermal imaging
- ...

Other Inverse Problems in Imaging

- All these inverse problems have important applications and are very different
- Yet, they all boil down to the same inverse problem, each with a different matrix \mathbf{A} :
$$\text{minimize}_{\mathbf{x}} \frac{1}{2} \|\mathbf{b} - \mathbf{Ax}\|_2^2 + \lambda \Psi(\mathbf{x})$$
- The methods derived here also apply to all those problems and applications; single-pixel imaging is a great example problem \rightarrow “if you can solve this, you can solve anything”

Review of ADMM for General Inverse Problems

- Objective or “loss” function of general inverse problem:
$$\underset{\mathbf{x}}{\text{minimize}} \quad \underbrace{\frac{1}{2} \|\mathbf{b} - \mathbf{A}\mathbf{x}\|_2^2}_{f(\mathbf{x})} + \underbrace{\lambda \Psi(\mathbf{x})}_{g(\mathbf{x})}$$

↑
weight of regularizer

- Reformulate as:
$$\underset{\{\mathbf{x}, \mathbf{z}\}}{\text{minimize}} \quad \underbrace{\frac{1}{2} \|\mathbf{b} - \mathbf{A}\mathbf{x}\|_2^2}_{f(\mathbf{x})} + \underbrace{\lambda \Psi(\mathbf{z})}_{g(\mathbf{z})}$$

subject to $\mathbf{D}\mathbf{x} - \mathbf{z} = \mathbf{0}$

- Remove constraints using augmented Lagrangian
$$L_\rho(\mathbf{x}, \mathbf{z}, \mathbf{u}) = f(\mathbf{x}) + g(\mathbf{z}) + \frac{\rho}{2} \|\mathbf{D}\mathbf{x} - \mathbf{z} + \mathbf{u}\|_2^2 + \|\mathbf{u}\|_2^2$$

Review of ADMM for General Inverse Problems

$$L_\rho(\mathbf{x}, \mathbf{z}, \mathbf{u}) = f(\mathbf{x}) + g(\mathbf{z}) + \frac{\rho}{2} \|\mathbf{D}\mathbf{x} - \mathbf{z} + \mathbf{u}\|_2^2 + \|\mathbf{u}\|_2^2$$

- Alternating gradient descent approach to solving penalty formulation leads to following iterative algorithm:

while not converged:

$$\mathbf{x} \leftarrow \text{prox}_{f,\rho}(\mathbf{v}) = \arg \min_{\mathbf{x}} L_\rho(\mathbf{x}, \mathbf{z}, \mathbf{u}) = \arg \min_{\mathbf{x}} f(\mathbf{x}) + \frac{\rho}{2} \|\mathbf{D}\mathbf{x} - \mathbf{v}\|_2^2$$

$$\mathbf{z} \leftarrow \text{prox}_{g,\rho}(\mathbf{v}) = \arg \min_{\mathbf{z}} L_\rho(\mathbf{x}, \mathbf{z}, \mathbf{u}) = \arg \min_{\mathbf{z}} g(\mathbf{z}) + \frac{\rho}{2} \|\mathbf{v} - \mathbf{z}\|_2^2$$

Review of ADMM for General Inverse Problems

\mathbf{x} - update:

$$\mathbf{x} \leftarrow \text{prox}_{\|\cdot\|_2, \rho}(\mathbf{v}) = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \frac{\rho}{2} \|\mathbf{D}\mathbf{x} - \mathbf{v}\|_2^2$$

$$\mathbf{x} \leftarrow \underbrace{(\mathbf{A}^T \mathbf{A} + \rho \mathbf{D}^T \mathbf{D})}^{\tilde{\mathbf{A}}}^{-1} \underbrace{(\mathbf{A}^T \mathbf{b} + \rho \mathbf{D}^T \mathbf{v})}_{\tilde{\mathbf{b}}}$$

- For general inverse problems, we don't necessarily have an efficient closed-form solution for this problem, like we did for the deconvolution problem
- Use matrix-free iterative solver, such as the conjugate gradient method, to solve $\tilde{\mathbf{A}}\mathbf{x} = \tilde{\mathbf{b}}$ (e.g., `scipy.sparse.linalg.cg`)

Review of ADMM for General Inverse Problems


\mathbf{z} – update for TV regularizer in closed form:

$$\mathbf{z} \leftarrow \text{prox}_{\|\cdot\|_1, \rho}(\mathbf{v}) = \arg \min_{\mathbf{z}} \lambda \|\mathbf{z}\|_1 + \frac{\rho}{2} \|\mathbf{v} - \mathbf{z}\|_2^2 = \mathcal{S}_{\kappa}(\mathbf{v})$$

\mathbf{z} – update for denoising-based regularizer in closed form:

$$\mathbf{z} \leftarrow \text{prox}_{\mathcal{D}, \rho}(\mathbf{v}) = \arg \min_{\mathbf{z}} \lambda \Psi(\mathbf{z}) + \frac{\rho}{2} \|\mathbf{v} - \mathbf{z}\|_2^2 = \mathcal{D}\left(\mathbf{v}, \sigma^2 = \frac{\lambda}{\rho}\right)$$

ADMM – Results



Short tangent on Half Quadratic Splitting (HQS)
(Another solver for constrained optimization problems)

The Half-quadratic Splitting (HQS) Method

- Objective or “loss” function of general inverse problem:
$$\underset{x}{\text{minimize}} \quad \underbrace{\frac{1}{2} \|\mathbf{b} - \mathbf{A}\mathbf{x}\|_2^2}_{f(x)} + \underbrace{\lambda \Psi(\mathbf{x})}_{g(\mathbf{x})}$$

↑
weight of regularizer

- Reformulate as:
$$\underset{\{x, z\}}{\text{minimize}} \quad \underbrace{\frac{1}{2} \|\mathbf{b} - \mathbf{A}\mathbf{x}\|_2^2}_{f(x)} + \underbrace{\lambda \Psi(\mathbf{z})}_{g(\mathbf{z})}$$

subject to $\mathbf{D}\mathbf{x} - \mathbf{z} = \mathbf{0}$

- Remove constraints using penalty term (equivalent for large ρ):
$$L_\rho(\mathbf{x}, \mathbf{z}) = f(\mathbf{x}) + g(\mathbf{z}) + \underbrace{\frac{\rho}{2} \|\mathbf{D}\mathbf{x} - \mathbf{z}\|_2^2}_{\text{penalty term}}$$

The Half-quadratic Splitting (HQS) Method

$$L_\rho(\mathbf{x}, \mathbf{z}) = f(\mathbf{x}) + g(\mathbf{z}) + \frac{\rho}{2} \|\mathbf{D}\mathbf{x} - \mathbf{z}\|_2^2$$

- Alternating gradient descent approach to solving penalty formulation leads to following iterative algorithm:

while not converged:

$$\mathbf{x} \leftarrow \text{prox}_{f,\rho}(\mathbf{z}) = \arg \min_{\mathbf{x}} L_\rho(\mathbf{x}, \mathbf{z}) = \arg \min_{\mathbf{x}} f(\mathbf{x}) + \frac{\rho}{2} \|\mathbf{D}\mathbf{x} - \mathbf{z}\|_2^2$$

$$\mathbf{z} \leftarrow \text{prox}_{g,\rho}(\mathbf{D}\mathbf{x}) = \arg \min_{\mathbf{z}} L_\rho(\mathbf{x}, \mathbf{z}) = \arg \min_{\mathbf{z}} g(\mathbf{z}) + \frac{\rho}{2} \|\mathbf{D}\mathbf{x} - \mathbf{z}\|_2^2$$

The Half-quadratic Splitting (HQS) Method

$$L_{\rho}(\mathbf{x}, \mathbf{z}) = f(\mathbf{x}) + g(\mathbf{z}) + \frac{\rho}{2} \|\mathbf{D}\mathbf{x} - \mathbf{z}\|_2^2$$

- Alternating gradient descent approach to solving penalty form
- Steps not tied together with dual variable

- Can be very sensitive to the penalty parameter, requiring more tuning than ADMM (technically, penalty needs to go to infinity)


$$\mathbf{z} \leftarrow \text{prox}_{g,\rho}(\mathbf{D}\mathbf{x}) = \arg \min_{\mathbf{z}} L_{\rho}(\mathbf{x}, \mathbf{z}) = \arg \min_{\mathbf{z}} g(\mathbf{z}) + \frac{\rho}{2} \|\mathbf{D}\mathbf{x} - \mathbf{z}\|_2^2$$

Applications of Compressive Imaging


Compressive Medical Imaging

- reduce acquisition time, radiation exposure, or allow for more patients in same time, ...
- examples: x-ray computed tomography and MRI

Computed Tomography (CT)




Reconstructed 2D Slices




Computed Tomography – Fourier Slice Theorem

- measurements = Fourier slices
- compressive CT: e.g. fewer slices




frequency domain





primal domain

Magnetic Resonance Imaging

- measurements = (random) Fourier coefficients
- compressive MRI: fewer Fourier coefficients



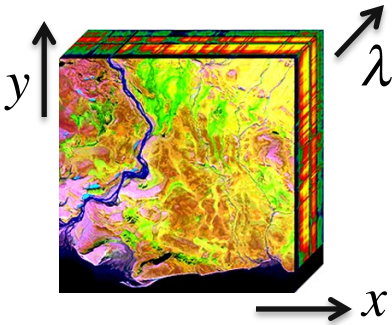
frequency domain



Compressive Imaging: CT & MRI


- people in bio-medical imaging often hesitant about priors:
 - few guarantees for success
 - if reconstruction breaks, not clear how exactly
 - is that feature a reconstruction artifact or the thing I'm looking for?

Compressive Hyperspectral Imaging

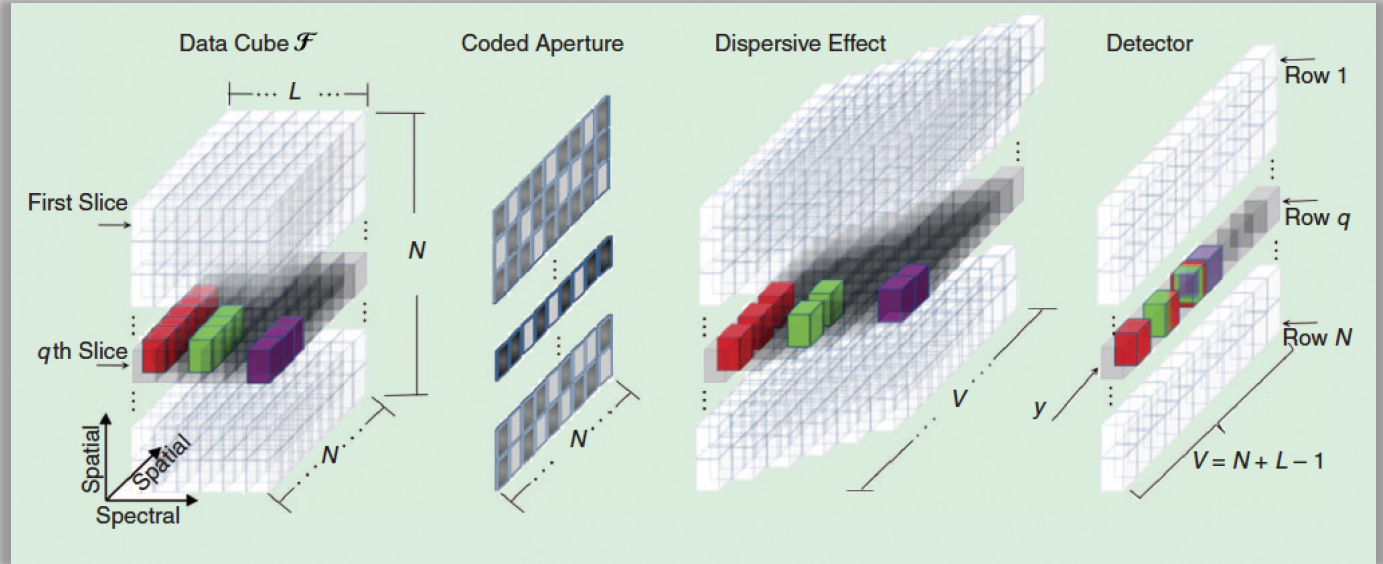
- motivation: 

- conventional: either scan over xy or over λ !
- idea: capture hyperspectral datacube with a single, coded image
 - use compressive sensing to reconstruct
- first approach: CASSI (coded aperture snapshot spectral imager), Wagadarikar 2008

Compressive Hyperspectral Imaging



Compressive Hyperspectral Imaging



Compressive Hyperspectral Imaging

- moderate quality for snapshot, but good quality for coded multi-shot
- applications: remote sensing, cultural heritage, ...



Compressive Imaging Everywhere

- metamaterials
- THz imaging
- x-ray imaging
- thermal IR
- ultra-fast imaging
- not as much on compressive coherent imaging (could be interesting for course projects: OCT, holography, ...)
- ...

Notes

- compressive imaging is an exploding area: check COSI, ICCP, CVPR, ICCV conferences, other optics journals and conferences
- most variants of compressive imaging problems can be implemented with ADMM
- check lecture notes online to help with homework
- Increasingly we want to learn the sensing matrices, reconstruction using neural networks and datasets...

References and Further Reading

Must read: course notes on Image Deconvolution with ADMM & course notes on compressive imaging

Adam

- D. Kingma, J. Ba “Adam: A method for stochastic optimization”, ICLR 2015

ADMM

- S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein “Distributed optimization and statistical learning via the alternating direction method of multipliers”, Foundation and Trends in Machine Learning, 2001

Single-pixel Imaging

- M. Duarte, M. Davenport, D. Takhar, J. Laska, T. Sun, K. Kelly, R. Baraniuk “Single-pixel imaging via compressive sampling”, IEEE Signal Processing Magazine 2008

TV Prior and Extensions

- L. Rudin, S. Osher, E. Fatemi “Nonlinear total variation-based noise removal algorithm”, Physica D, 1992
- A. Levin, Y. Weiss, F. Durand, W. Freeman “Understanding and evaluating blind deconvolution algorithms”, CVPR 2009
- D. Krishnan, R. Fergus “Fast Image Deconvolution using Hyper-Laplacian Priors”, NIPS 2009
- K. Bredies, K. Kunisch, T. Pock “Total Generalized Variation”, Technical Report 2009
- S. Lefkimmiatis, J. Ward, M. Unser “Hessian Schatten-Norm Regularization for Linear Inverse Problems”, IEEE Transactions on Image Processing 2003

Unrolled Optimization

- S. Diamond, V. Sitzmann, F. Heide, G. Wetzstein “Unrolled optimization with deep priors”, arxiv, 2017