MOCCA: a primal/dual algorithm for nonconvex composite functions with applications to CT imaging

Rina Foygel Barber

Dept. of Statistics, University of Chicago

http://www.stat.uchicago.edu/~rina/mocca.html

Collaborators

- Algorithm & optimization work:
  collaboration with Emil Sidky
- Application to CT:
  collaboration with Emil Sidky, Taly Gilat-Schmidt, & Xiaochuan Pan

Emil Sidky

Xiaochuan Pan

Dept. of Radiology
U. Chicago

Taly Gilat-Schmidt
Dept. Biomedical Eng.
Marquette U.

Computed tomography (CT) imaging

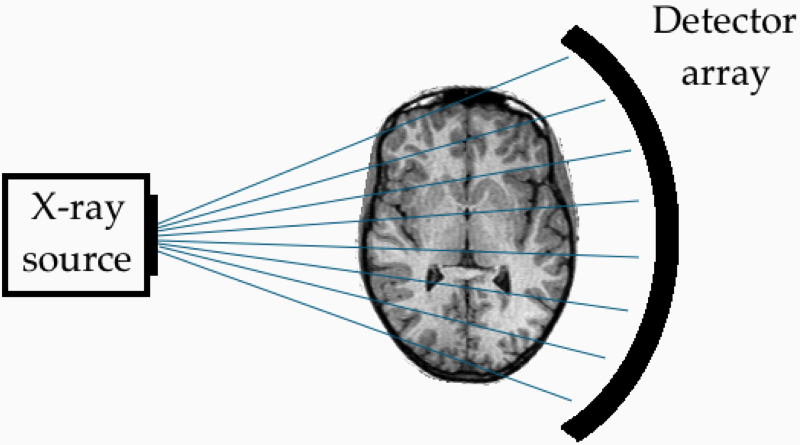

Computed tomography (CT) imaging

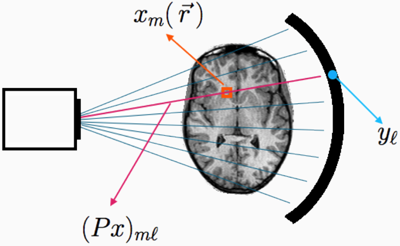

- Measure: y_ℓ = number of photons detected along ray ℓ
- Want to estimate the materials at each point inside the object:

$$x_m(\vec{r}) = \text{density of material } m \text{ at location } \vec{r}$$

- Distribution of y is $\approx$ determined by projections of x:

$$(Px)_{m\ell} = \text{amount of material } m \text{ along ray } \ell$$

Computed tomography (CT) imaging

If the X-ray beam is monochromatic,

for each ray ℓ the number of photons detected is

$$y_\ell \approx \mathsf{Poisson}\left(I_{\mathsf{total}} \cdot \exp\left\{ -\sum_m \mu_m \cdot \underbrace{(Px)_{m\ell}}_{\substack{\text{amount of material } m \\ \text{along ray } \ell}} \right\} \right)$$

μ_m = attenuation coefficient for material m

I_{total} = total intensity of X-ray spectrum / detector sensitivity

Computed tomography (CT) imaging

X-ray beam used in CT is polychromatic:

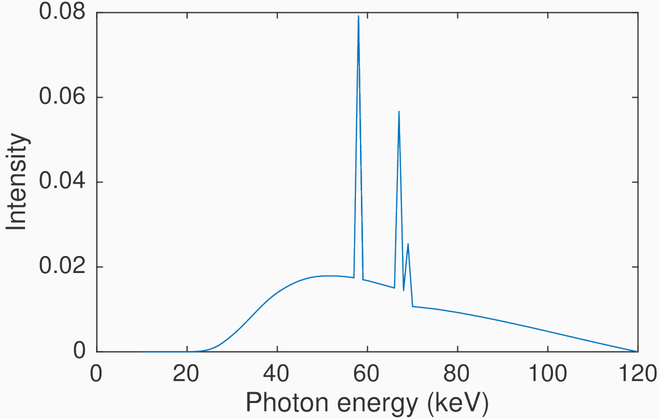

Computed tomography (CT) imaging

For polychromatic X-ray beam:

$$y_\ell \approx \text{Poisson}\left(I_{\text{total}} \int_E S(E) \cdot \exp\left\{ -\sum_m \underbrace{\mu_m(E)}_{\substack{\text{attenuation coefficient for} \\ \text{material } m \text{ at energy } E}} \cdot (Px)_{m\ell} \right\} \, \text{d}E \right)$$

$S(E) =$ distribution of X-ray spectrum intensity / detector sensitivity across energies E

Computed tomography (CT) imaging

Existing algorithms for CT treat the measurements
as a log linear function of the image:

$$\log\left(\mathbb{E}\left[y\right]\right) \approx \text{Linear function of } Px$$

- Filtered back projection (FBP) — used in clinical CT

Computed tomography (CT) imaging

$$\log\left(\mathbb{E}\left[\frac{y_\ell}{I_{\mathsf{total}}}\right]\right)$$
$$= \log\left(\int_E S(E)\cdot\exp\left\{-\sum_m \mu_m(E) \ \cdot\ (Px)_{m\ell}\right\} \mathsf{d}E\right)$$

If we swap $\log(\cdot)$ with averaging:

$$\approx -\sum_m \left[\int_E S(E)\cdot\mu_m(E)\,\mathsf{d}E\right]\cdot\ (Px)_{m\ell}$$

Computed tomography (CT) imaging

Ignoring the X-ray spectrum leads to beam hardening:

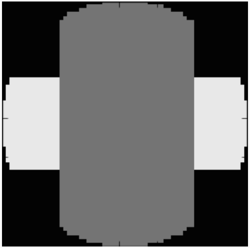

true object

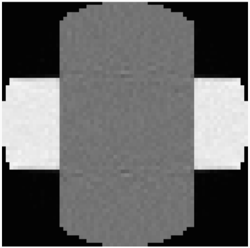

full Poisson model

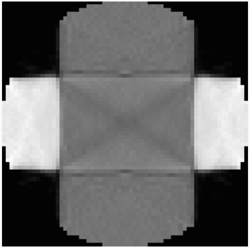

log-linear
Poisson model

Computed tomography (CT) imaging

Beam hardening in practice:

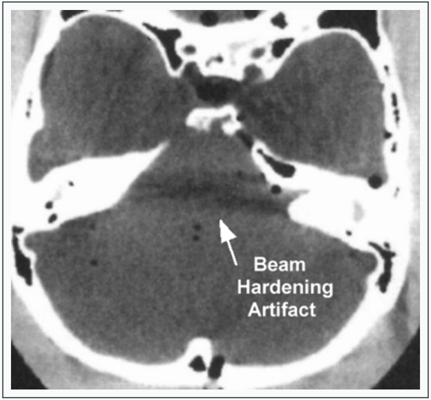

Goldman, J. Nucl. Med. Technol., 2007

CT optimization problem

After discretization into pixels, want to minimize

$$\sum_{\text{rays } \ell} \underbrace{\mathcal{L}\Bigg(y_\ell; \sum_{\text{energy } i} s_{\ell i} \cdot \exp\Big\{ -(\mu^\top P x)_{\ell i} \Big\}\Bigg)}_{\text{Poisson negative log-likelihood}} \;+\; \left(\begin{array}{c} \text{Total variation} \\ \text{constraints, etc} \end{array} \right)$$

Vector $x =$ discretized materials map

CT optimization problem

Spectral CT: photon detection is split into multiple energy "windows" (bands):

$$\sum_{\substack{\text{windows } w\\ \text{rays } \ell}} \underbrace{\mathcal{L}\Bigg(y_{w\ell};\ \sum_{\text{energy } i} s_{w\ell i}\cdot\exp\Big\{-(\mu^\top P x)_{\ell i}\Big\}\Bigg)}_{\text{Poisson negative log-likelihood}} + \begin{pmatrix}\text{Total variation}\\ \text{constraints, etc}\end{pmatrix}$$

Vector x = discretized materials map

Optimization problem

General problem:

Want to minimize

$$\mathsf{F}(Kx) + \mathsf{G}(x)$$

where F and G might be nonconvex and/or nondifferentiable

Optimization problem: differentiable case

If F is differentiable & G has an easy proximal map:

- Proximal gradient descent:

$$\begin{cases} \widetilde{x}_{t+1} = x_t - \frac{1}{\eta} K^\top \nabla \mathsf{F}(K x_t), \\ x_{t+1} = \arg\min \left\{ \frac{1}{2} \left\| x - \widetilde{x}_{t+1} \right\|_2^2 + \frac{1}{\eta} \mathsf{G}(x) \right\} \end{cases}$$

- Accelerated version: add an extrapolation step,

$$x_{t+1} \leftarrow x_{t+1} + \theta(x_{t+1} - x_t)$$

Convex: Beck & Teboulle 2009

Nonconvex: Loh & Wainwright 2013; Ochs et al 2014

Optimization problem: convex case

If F, G are convex:

ADMM (alternating direction method of multipliers)

- Rewrite optimization:

$$\min_{x,w} \max_{u} \left\{ \mathsf{F}(w) + \mathsf{G}(x) + \langle u, Kx - w \rangle + \frac{\sigma}{2} \|Kx - w\|_2^2 \right\}$$

- Alternate between minimizing over x and w, and updating u

Boyd et al, *FnTML*, 2011

Optimization problem: convex case

CP (Chambolle-Pock algorithm)

Saddle point problem $$\min_x \max_y \left\{ \underbrace{\langle Kx, y\rangle - \overbrace{\mathsf{F}^*(y)}^{\text{Fenchel conjugate of F}}}_{\mathsf{F}(Kx) = \text{ max over } y} + \mathsf{G}(x) \right\}$$

Iterate:

$$x_{t+1} = \arg\min \left\{ \langle Kx, y_t\rangle + \mathsf{G}(x) + \frac{1}{2\tau} \|x - x_t\|_2^2 \right\}$$
$$y_{t+1} = \arg\max \left\{ \langle K\bar{x}_{t+1}, y\rangle - \mathsf{F}^*(y) - \frac{1}{2\sigma} \|y - y_t\|_2^2 \right\}$$

extrapolation $x_{t+1} + \theta(x_{t+1} - x_t)$

- Equivalent to ADMM with an added preconditioning step

Chambolle & Pock, *J. Math. Imaging Vision*, 2011

Optimization problem: convex case

Can we run CP or ADMM if F & G are nonconvex?

- Example: $x \mapsto \mathsf{F}(Kx) + \mathsf{G}(x)$ is convex,

  but F is strongly concave in some directions

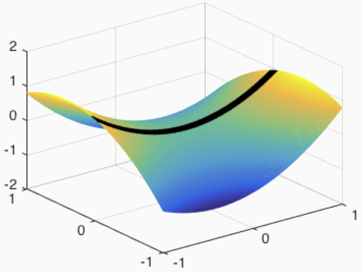

- ADMM / CP may diverge immediately
- CP may converge to the wrong solution because $\mathsf{F}^{**} \neq \mathsf{F}$

MOCCA algorithm

Main idea:

1. Take local convex approximations to F and G
2. Take one step (or a few steps) of the CP algorithm
3. Repeat until convergence

MOCCA $\approx$ majorization/minimization + primal/dual updates

Main question:

How should we construct the local convex approximations?

MOCCA algorithm

- Split F & G into convex + differentiable components:

$$\mathsf{F} = \mathsf{F}_{\mathsf{cvx}} + \mathsf{F}_{\mathsf{diff}}, \quad \mathsf{G} = \mathsf{G}_{\mathsf{cvx}} + \mathsf{G}_{\mathsf{diff}}$$

- Convex approximations at step t:

$$\mathsf{F}_t(w) = \mathsf{F}_{\mathsf{cvx}}(w) + \left[\mathsf{F}_{\mathsf{diff}}(z^t_{\mathsf{F}}) + \langle w - z^t_{\mathsf{F}}, \nabla \mathsf{F}_{\mathsf{diff}}(z^t_{\mathsf{F}})\rangle\right]$$
$$\mathsf{G}_t(x) = \mathsf{G}_{\mathsf{cvx}}(x) + \left[\mathsf{G}_{\mathsf{diff}}(z^t_{\mathsf{G}}) + \langle x - z^t_{\mathsf{G}}, \nabla \mathsf{G}_{\mathsf{diff}}(z^t_{\mathsf{G}})\rangle\right]$$

MOCCA algorithm

- How do we pick expansion points z^t_{F} and z^t_{G}?

$$\underbrace{\mathsf{F}(Kx)}_{\text{dual variable } y} + \underbrace{\mathsf{G}(x)}_{\text{primal variable } x}$$

- z^t_{G} = primal variable x_t
- z^t_{F} = primal point that <u>mirrors</u> the dual variable y_t

MOCCA algorithm

Iterate:

$$x_{t+1} = \arg\min \left\{ \langle Kx, y_t \rangle + \mathsf{G}_t(x) + \frac{1}{2\tau} \|x - x_t\|_2^2 \right\}$$

$$y_{t+1} = \arg\max \left\{ \langle K\bar{x}_{t+1}, y \rangle - \mathsf{F}_t^*(y) - \frac{1}{2\sigma} \|y - y_t\|_2^2 \right\}$$

$$z_{\mathsf{F}}^{t+1} = \frac{1}{\sigma}(y_t - y_{t+1}) + K\bar{x}_{t+1}, \qquad z_{\mathsf{G}}^{t+1} = x_{t+1}$$

- Step sizes σ, τ should satisfy $\sigma\tau \|K\|^2 < 1$.

(As in Chambolle & Pock 2011)

- Can use a preconditioning step to avoid computing $\|K\|$

(Pock & Chambolle 2011)

Case study: nonconvex total variation

The problem:

- True signal $x^\star \in \mathbb{R}^d$ has total-variation sparsity

  (nearby pixels often have identical values)

- Problem: minimize loss $\mathcal{L}(x)$ subject to sparsity in $\nabla_{\mathrm{2d}} x$

  2-dim. gradient operator

- Common approach: penalize $\|\nabla_{\mathrm{2d}} x\|_1$

  $\rightsquigarrow$ bias due to shrinkage on large gradient values

Case study: nonconvex total variation

Use a nonconvex TV penalty to reduce bias from shrinkage:

$$\mathsf{logTV}_\beta(x) = \sum_i \beta \cdot \log\left(1 + |(\nabla_{2\mathrm{d}} x)_i|/\beta\right)$$

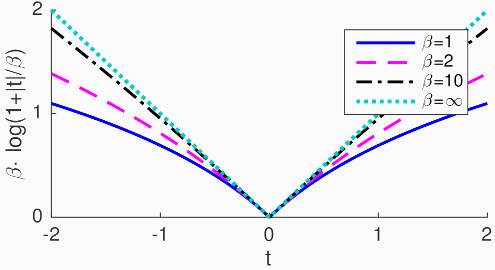

Equivalent to $\|x\|_{\mathsf{TV}} = \|\nabla_{2\mathrm{d}} x\|_1$ when $\beta = \infty$.

Parekh & Selesnick (2015)
Related to reweighted ℓ_1 sparsity, Candès et al (2008)

Case study: nonconvex total variation

Optimization problem for least squares loss:

$$\text{minimize } \frac{1}{2}\|b - Ax\|_2^2 + \nu \cdot \mathsf{logTV}_\beta(x)$$

$$\mathsf{logTV}_\beta(x) = \underbrace{\|\nabla_{2\mathrm{d}} x\|_1}_{\text{convex}} + \underbrace{\left[\beta \log(1 + |\nabla_{2\mathrm{d}} x|/\beta) - \|\nabla_{2\mathrm{d}} x\|_1\right]}_{h(\nabla_{2\mathrm{d}} x) = \text{differentiable}}$$

Define:

$$\begin{aligned} \mathsf{F}_{\mathsf{cvx}}(w) &= \nu \cdot \|w\|_1 & \qquad \mathsf{G}_{\mathsf{cvx}}(x) &= \tfrac{1}{2}\|b - Ax\|_2^2 \\ \mathsf{F}_{\mathsf{diff}}(w) &= \nu \cdot h(w) & \qquad \mathsf{G}_{\mathsf{diff}}(x) &\equiv 0 \end{aligned}$$

Case study: nonconvex total variation

MOCCA for least squares + nonconvex TV

$$x_{t+1} = \left(\mathbf{I} + \tau A^\top A\right)^{-1} \left(x_t + \tau A^\top b - \tau \nabla_{2\mathrm{d}}^\top y_t\right)$$

$$y_{t+1} = \mathsf{Truncate}_\nu \left(y_t + \sigma \nabla_{2\mathrm{d}} \bar{x}_{t+1} - \lambda \nabla h(z_{\mathsf{F}}^t)\right) + \lambda \nabla h(z_{\mathsf{F}}^t)$$

$$z_{\mathsf{F}}^{t+1} = \tfrac{1}{\sigma}(y_t - y_{t+1}) + K \bar{x}_{t+1}$$

Case study: nonconvex total variation

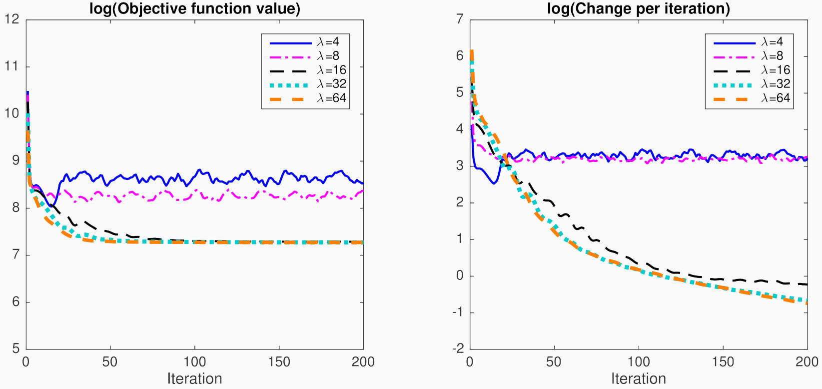

Problem size: $x \in \mathbb{R}^{25\times 25}$ with block structure; 200 measurements
Tuning parameter λ: $\sigma = \frac{\lambda}{2}$, $\tau = \frac{1}{2\lambda}$

Application to spectral CT

Simulated CT measurements from object with 2 materials:

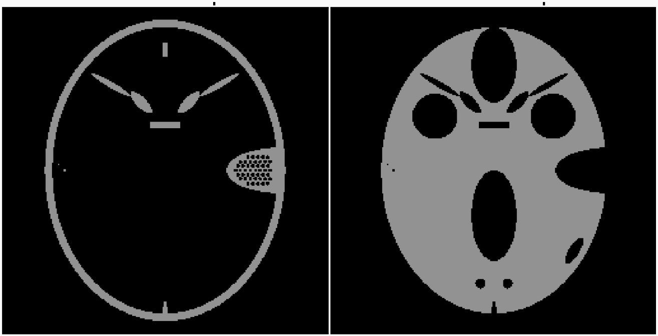

Bone | Brain

FORBILD head phantom (Lauritsch & Bruder 2001)

Application to spectral CT

Minimize:

$$\sum_{\substack{\text{windows } w \\ \text{rays } \ell}} \underbrace{\mathcal{L}\left(y_{w\ell};\ \sum_{\text{energy } i} s_{w\ell i} \cdot \exp\left\{-(\mu^\top P x)_{\ell i}\right\}\right)}_{\text{Poisson negative log-likelihood}} + \left(\begin{array}{c}\text{Total variation} \\ \text{constraints, etc}\end{array}\right)$$

Application to spectral CT

Minimize:

$$\underbrace{\mathcal{L}(\mu^\top P \cdot x)}_{\text{Poisson negative log-likelihood}} \quad + \quad \underbrace{\delta \begin{pmatrix} \|x_{\text{bone}}\|_{\text{TV}} \leq \gamma_{\text{bone}} \\ \& \\ \|x_{\text{brain}}\|_{\text{TV}} \leq \gamma_{\text{brain}} \end{pmatrix}}_{\text{convex indicator function}}$$

Application to spectral CT

MOCCA setup: minimize $\mathsf{F}(Kx) + \mathsf{G}(x)$

$$w = \begin{pmatrix} w_1 \\ w_2 \\ w_3 \end{pmatrix} = \begin{pmatrix} \mu^\top P \\ \nabla_{\text{bone}} \\ \nabla_{\text{brain}} \end{pmatrix} \cdot x = Kx$$

$$\begin{cases} \mathsf{F}(w) = \begin{pmatrix} \text{local convex/concave} \\ \text{quadratic approx. to } \mathcal{L}(w_1) \end{pmatrix} + \delta \begin{pmatrix} \|w_2\|_1 \leq \gamma_{\text{bone}} \\ \& \\ \|w_3\|_1 \leq \gamma_{\text{brain}} \end{pmatrix} \\ \\ \mathsf{G}(x) \equiv 0 \end{cases}$$

Application to spectral CT

Algorithm:

1. Take one step of the MOCCA algorithm
2. Update local convex/concave quadratic approximation to $\mathcal{L}(\cdot)$
3. Update step sizes
4. Repeat until convergence

Application to spectral CT

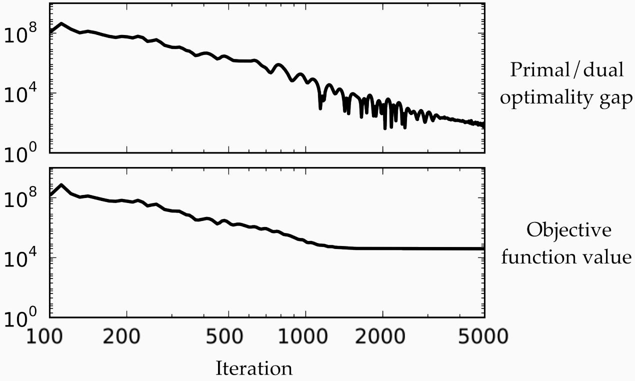

Application to spectral CT

Using the Poisson likelihood vs. a least squares loss:

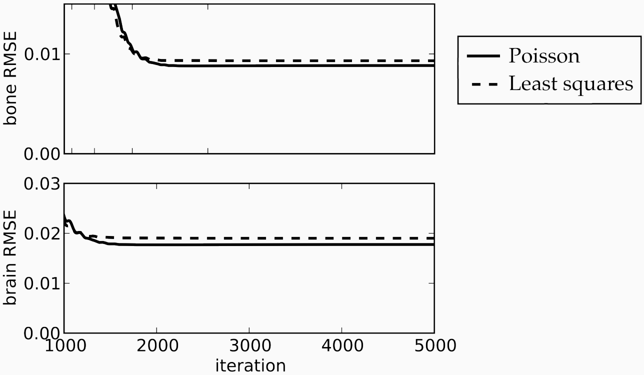

Application to spectral CT

How critical is the choice of TV constraints γ_{bone} & γ_{brain}?

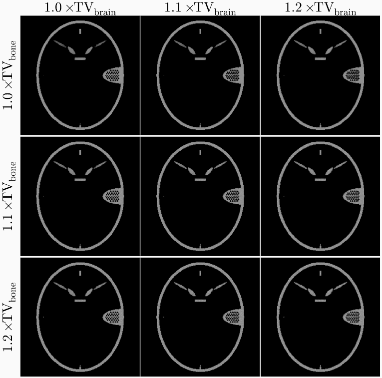

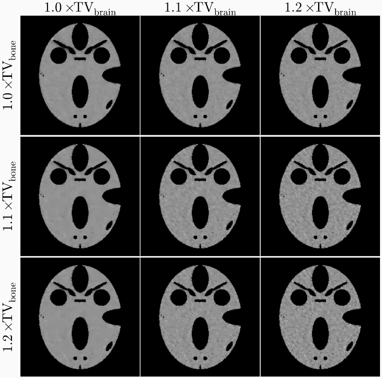

Theoretical results

Question 1:
If MOCCA converges, has it converged to the right solution?

Theoretical results

Theorem 1: critical points

If the MOCCA algorithm converges with

$$(x_t, y_t, z_t) \to (\widehat{x}, \widehat{y}, \widehat{z})$$

then $\widehat{x}$ is a critical point of the optimization problem,

$$0 \in K^\top \partial \mathsf{F}_{\mathsf{cvx}}(K\widehat{x}) + K^\top \nabla \mathsf{F}_{\mathsf{diff}}(K\widehat{x}) + \partial \mathsf{G}_{\mathsf{cvx}}(\widehat{x}) + \nabla \mathsf{G}_{\mathsf{diff}}(\widehat{x})$$

Theoretical results

Question 2:
Is MOCCA guaranteed to converge (& at what rate)?

Theoretical results

Stable MOCCA algorithm (with "inner loop")

At stage t,

1. Run the "inner loop": fixing expansion points $(z_{\mathsf{F}}^t, z_{\mathsf{G}}^t)$,
   update (x, y) variables L_{t+1} times
2. Update (x, y) variables by averaging over stage t:

$$(x_{t+1}, y_{t+1}) = \frac{1}{L_{t+1}} \sum_{\ell=1}^{L_{t+1}} (x_{t+1;\ell}, y_{t+1;\ell})$$

3. Update expansion points by averaging over stage t:

$$\begin{cases} z_{\mathsf{F}}^{t+1} = \frac{1}{L_{t+1}} \sum_{\ell=1}^{L_{t+1}} \frac{1}{\sigma}(y_{t+1;\ell-1} - y_{t+1;\ell}) + K\bar{x}_{t+1;\ell} \\ z_{\mathsf{G}}^{t+1} = \frac{1}{L_{t+1}} \sum_{\ell=1}^{L_{t+1}} x_{t+1;\ell} \end{cases}$$

Theoretical results

Background—restricted strong convexity:

- Definition: a loss function $\mathcal{L}(x)$ satisfies RSC if

$$\langle x - x', \partial\mathcal{L}(x) - \partial\mathcal{L}(x')\rangle \gtrsim \left\|x - x'\right\|_2^2 - \tfrac{\log(d)}{n}\left\|x - x'\right\|_1^2$$

- Convex: accurate recovery of sparse/structured signals in high-dimensional statistics

  Negahban et al 2009

- Nonconvex: local minima guaranteed to be near global min for (differentiable loss) $+$ (sparsity penalty)

  Loh & Wainwright 2013

Theoretical results

Restricted convexity/smoothness assumptions for MOCCA:

- $\mathsf{F}_{\mathsf{cvx}}$ is Λ_{F}-convex and $\mathsf{F}_{\mathsf{diff}}$ is Θ_{F}-smooth
- $\mathsf{G}_{\mathsf{cvx}}$ is Λ_{G}-convex and $\mathsf{G}_{\mathsf{diff}}$ is Θ_{G}-smooth

- The overall optimization problem is nearly convex:

$$\underbrace{(Kx)^\top(\Lambda_{\mathsf{F}} - \Theta_{\mathsf{F}})(Kx)}_{\text{Convexity of F}} + \underbrace{x^\top(\Lambda_{\mathsf{G}} - \Theta_{\mathsf{G}})x}_{\text{Convexity of G}} \succeq C_{\mathsf{cvx}} \|x\|_2^2 - \tau^2 \|x\|_{\mathsf{restrict}}^2 \ .$$

ℓ_1 norm / any structured norm

- Optimization is over bounded region $\{x : \|x\|_{\mathsf{restrict}} \leq R\}$

Theoretical results

Theorem 2: convergence guarantee

For the stable form of the MOCCA algorithm with $L_t \sim C^t$,

$$\|x_t - x^\star\|_2 \lesssim C^{-t/2} + \tau R,$$

for any critical point $x^\star$ with $\|x^\star\|_{\text{restrict}} \leq R$.

Number of iterations to calculate x_t is $L_1 + \cdots + L_t \sim C^t$

$$\rightsquigarrow \quad \|x_t - x^\star\|_2 \sim \frac{1}{\sqrt{(\text{computational cost})}} + \tau R$$

Theoretical results

Main ingredient: contraction property

Consider two convex approximations:

$$\left\{\begin{aligned} &\mathsf{F}_z(Kx) + \mathsf{G}_z(x) \\ &\mathsf{F}_{z'}(Kx) + \mathsf{G}_{z'}(x) \end{aligned}\right. \qquad \text{with minimizers} \qquad \left.\begin{aligned} &x^\star_z \\ &x^\star_{z'} \end{aligned}\right\}$$

Then

$$\left\| \begin{pmatrix} x^\star_z - x^\star_{z'} \\ Kx^\star_z - Kx^\star_{z'} \end{pmatrix} \right\| \leq (1-\epsilon) \left\| \begin{pmatrix} z_{\mathsf{G}} - z'_{\mathsf{G}} \\ z_{\mathsf{F}} - z'_{\mathsf{F}} \end{pmatrix} \right\| + C \cdot \tau R$$

for some $\epsilon > 0$ and $C < \infty$.

Theoretical results

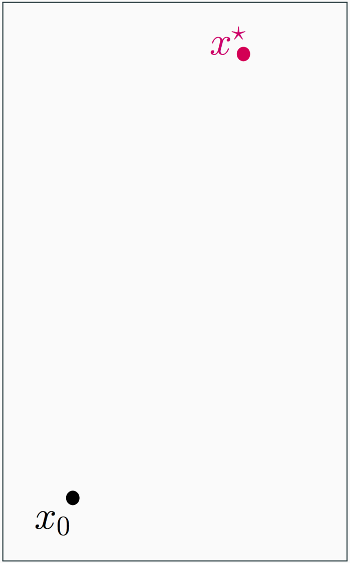

Theoretical results

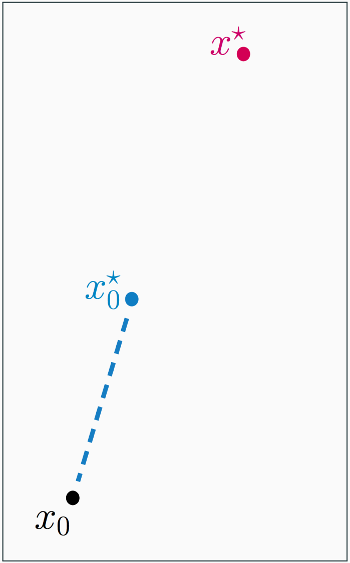

Theoretical results

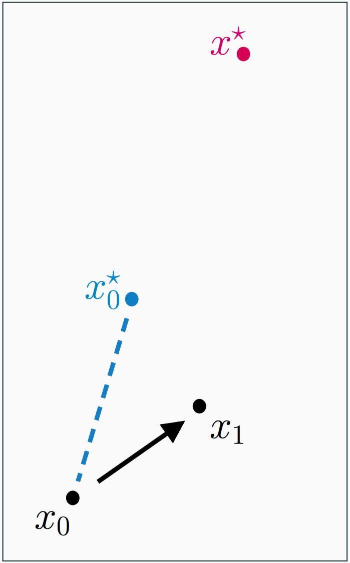

Theoretical results

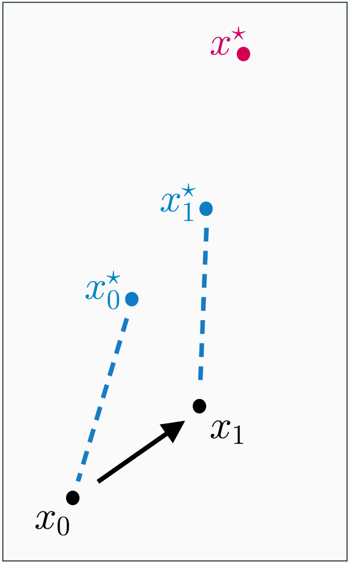

Theoretical results

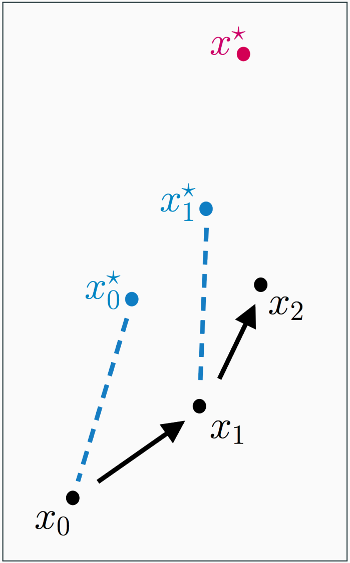

Theoretical results

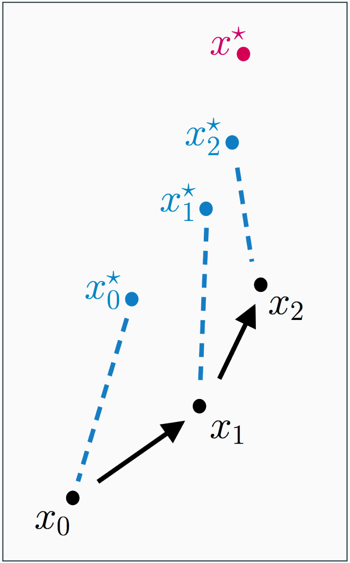

Theoretical results

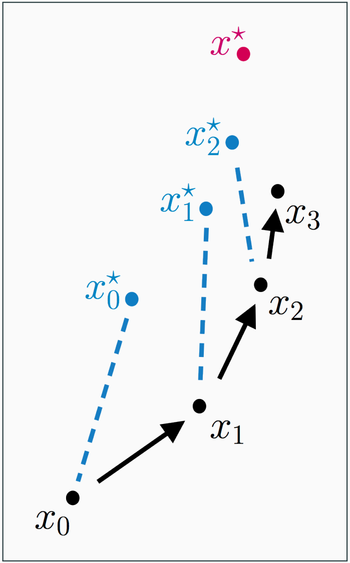

Theoretical results

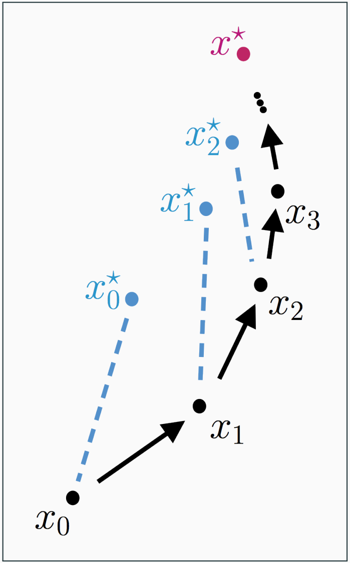

Future directions

Optimization & theory:

- Is the stable "inner loop" version of MOCCA necessary?
- Without RSC, guarantee convergence to stationary point?
- Adaptive step sizes for faster convergence?

CT imaging:

- Detector sensitivity is not known exactly & may vary over detector cells $\rightsquigarrow$ data-adaptive calibration?
- Apply MOCCA directly to Poisson likelihood, without quadratic approximation?

Thank you!

Website: `http://www.stat.uchicago.edu/~rina/mocca.html`

Funding: partially supported by NIH Grants R21EB015094, CA158446, CA182264, and EB018102. The contents of this presentation are solely the responsibility of the authors and do not necessarily represent the official views of the National Institutes of Health.