Chapter 2
Applications in Image Processing and Computer Vision

Variational Image Processing
Summer School on Inverse Problems 2015

Michael Moeller
Computer Vision
Department of Computer Science
TU München

Variational Methods

$$u(\alpha) \in \arg\min_u H(u, f) + \alpha J(u)$$

Inverse problems perspective:
Variational methods allow to reestablish the continuous dependence on the data. They provide a tool for tackling ill-posed problems.

Variational Methods

$$u(\alpha) \in \arg\min_u H(u,f) + \alpha J(u)$$

Inverse problems perspective:
Variational methods allow to reestablish the continuous dependence on the data. They provide a tool for tackling ill-posed problems.

Modelling perspective:
How can we design $H(u,f)$ and $J(u)$ such that a low value of the resulting energy yields a desired solution in various applications?

Image Deblurring

Seen this morning:

$$u(\alpha) \in \arg\min_u H(u, f) + \alpha J(u)$$

for

$$H(u, f) = \frac{1}{2}\|Au - f\|_2^2, \qquad J(u) = TV(u)$$

where A was a linear blur operator.

Image Deblurring

Seen this morning:

$$u(\alpha) \in \arg\min_u H(u, f) + \alpha J(u)$$

for

$$H(u, f) = \frac{1}{2}\|Au - f\|_2^2, \qquad J(u) = TV(u)$$

where A was a linear blur operator.

By simply changing the meaning of A, we can already tackle many classical image processing problems!

Image inpainting

copyright issues

Image inpainting

How can we unleash the lion?

copyright issues

Input image f

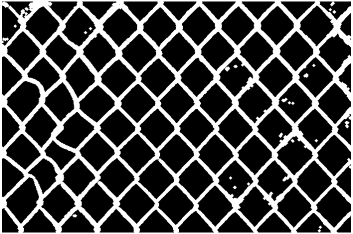

Mask m

Image inpainting

How can we unleash the lion?

copyright issues

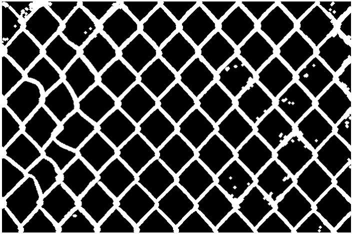

Input image f — Mask m

Joint inpainting and denoising:

We already know how to solve

$$u(\alpha) = \arg\min_u \frac{1}{2}\|Au - f\|_2^2 + \alpha TV(u)$$

Choose

$$Au(x) = \begin{cases} u(x) & \text{if } m(x) = 0, \\ 0 & \text{if } m(x) = 1 \end{cases}$$

Image inpainting

How can we unleash the lion?

copyright issues

Input image f

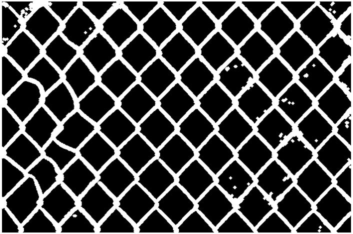

Mask m

Joint inpainting and denoising:

We already know how to solve

$$u(\alpha) = \arg\min_u \frac{1}{2}\|Au - f\|_2^2 + \alpha TV(u)$$

Choose

$$Au(x) = \begin{cases} u(x) & \text{if } m(x) = 0, \\ 0 & \text{if } m(x) = 1 \end{cases}$$

Image inpainting

Exemplar base methods.

image is taken from: *Exemplar-based inpainting from a variational point for view.* Aujol, Ladjal, Masnou.

Nonlocal patch-based image inpainting. Esser, Zhang.

Image inpainting

Exemplar base methods.

Find an image u, coefficients c, based on a dictionary of known patches P via

$$\min_{u,c} \|Pc - u\|_2^2 + \iota_{u_I = f_I}(u) + \iota_{\Delta}(c) + \alpha R(c),$$

such that the regularizer $R(c)$ encourages translations.

image is taken from: *Exemplar-based inpainting from a variational point for view.* Aujol, Ladjal, Masnou.
Nonlocal patch-based image inpainting. Esser, Zhang.

Image demosaicking

The way a camera records colors:

copyright issues

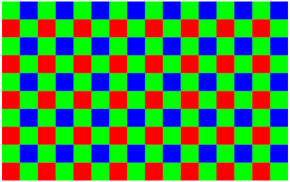

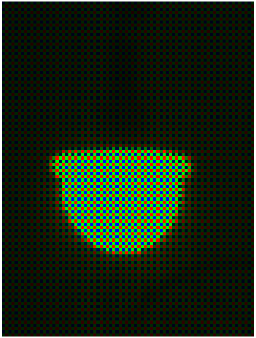

Inpainting problem:

$$u(\alpha) = \arg\min_u \frac{1}{2}\|Au - f\|_2^2 + \alpha R(u)$$

Choose

$$Au(x) = \begin{cases} u(x) & \text{if color known,} \\ 0 & \text{if color not known.} \end{cases}$$

Image Zooming

Input image f Bigger image?

Image Zooming

Input image f

Bilinear interpolation

Image Zooming

Input image f

Bilinear interpolation

High res.

Blurred

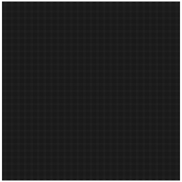

Mask

Input data

Invert this chain!

Image Zooming

Input image f

Bilinear interpolation

We already know how to solve inpainting and deblurring. Combine

$$A = \underbrace{A_1}_{\text{Subsampling operator}} \underbrace{A_2}_{\text{Blur operator}}$$

and solve

$$u(\alpha) = \arg\min_u \frac{1}{2}\|Au - f\|_2^2 + \alpha TV(u).$$

Image Zooming

TV upsampled

Bilinear interpolation

We already know how to solve inpainting and deblurring. Combine

$$A = \underbrace{A_1}_{\text{Subsampling operator}} \underbrace{A_2}_{\text{Blur operator}}$$

and solve

$$u(\alpha) = \arg\min_u \frac{1}{2}\|Au - f\|_2^2 + \alpha TV(u).$$

Medical Imaging

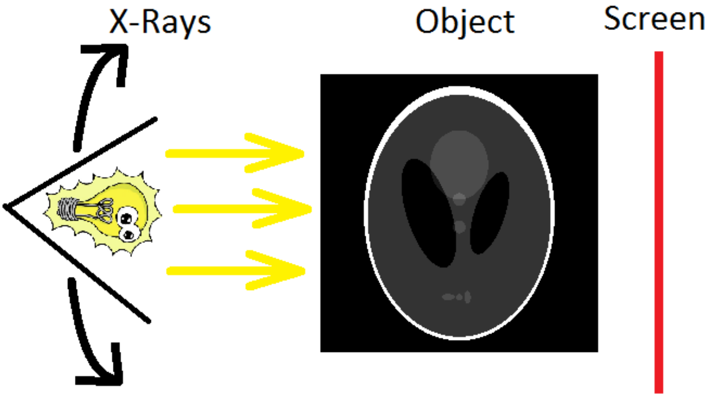

Medical Imaging

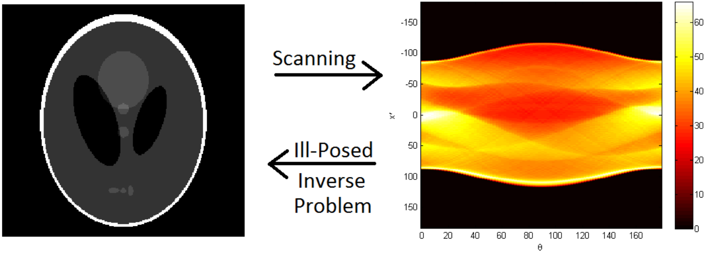

CT Scanning: Reconstruct an image from its line integrals

$$\frac{1}{2}\|Au - f\|^2 + \alpha TV(u)$$

with

$$Au(\Theta, s) = \int_{x \cdot \vec{\Theta} = s} u(x)dx^{\perp}$$

Medical Imaging

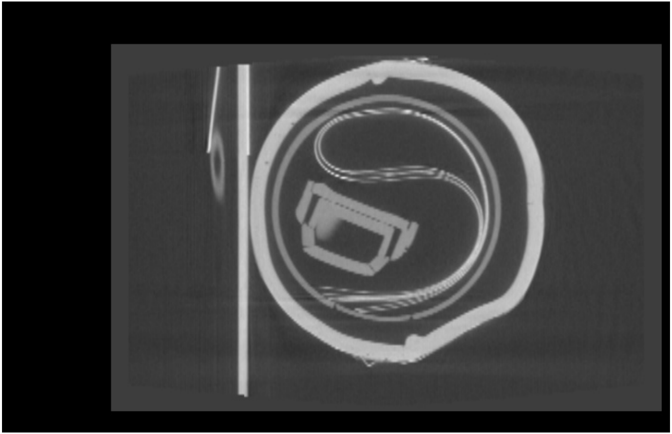

Courtesy of Jahn Müller and Ralf Engbers!

Medical Imaging

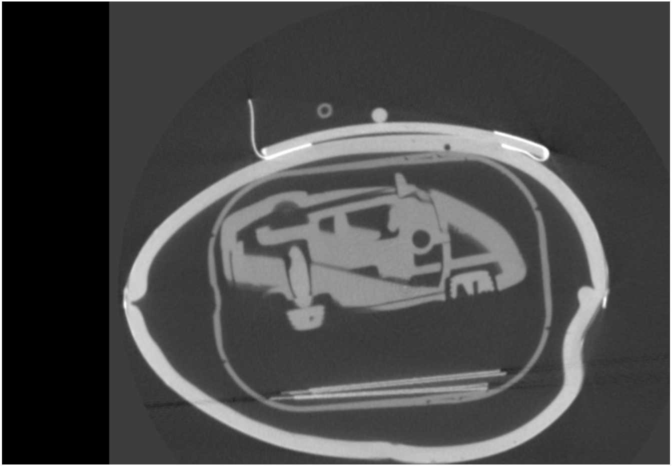

Courtesy of Jahn Müller and Ralf Engbers!

Medical Imaging

Egg interior reconstruction!

Courtesy of Jahn Müller and Ralf Engbers!

Image denoising

Most obvious extension/simplification: Use

$$\frac{1}{2}\|Au - f\|_2^2 + \alpha TV(u)$$

with $A = Id$ for image denoising.

Image denoising

Most obvious extension/simplification: Use

$$\frac{1}{2}\|Au - f\|_2^2 + \alpha TV(u)$$

with $A = Id$ for image denoising.

So far: Modifications of A. But we are not limited to the L^2 distance as a data term!

$$u(\alpha) = \arg\min_u \|u - f\|_1 + \alpha TV(u)$$

Image denoising

Most obvious extension/simplification: Use

$$\frac{1}{2}\|Au - f\|_2^2 + \alpha TV(u)$$

with $A = Id$ for image denoising.

So far: Modifications of A. But we are not limited to the L^2 distance as a data term!

$$u(\alpha) = \arg\min_u \|u - f\|_1 + \alpha TV(u)$$

More robust towards outliers / impulse noise.

Image denoising

$$u(\alpha) = \arg\min_u \|u - f\|_1 + \alpha TV(u)$$

Kodak image database: `http://r0k.us/graphics/kodak/`

Image denoising

$$u(\alpha) = \arg\min_u \|u - f\|_1 + \alpha TV(u)$$

Kodak image database: `http://r0k.us/graphics/kodak/`

Image decompression

How do you usually store an image? Maybe as a .jpg?

What does .jpg do?

Image decompression

How do you usually store an image? Maybe as a .jpg?

What does .jpg do?

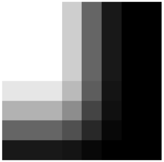

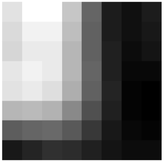

836	657	-61	-129	19	23	4	-1
304	239	-22	-47	7	8	2	0
-166	-130	12	26	-4	-5	-1	0
55	43	-4	-8	1	1	0	0
-14	-11	1	2	0	0	0	0
10	8	-1	-2	0	0	0	0
-6	-4	0	1	0	0	0	0
-1	-1	0	0	0	0	0	0

DCT coef. original patch

840	630	-70	-140	0	0	0	0
280	210	0	-70	0	0	0	0
-140	-140	0	0	0	0	0	0
70	70	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0

Quantized coefficients

Image decompression

Image decompression

- DCT is nothing but an orthonormal transform A.
- Inverse problem to quantization: Coefficients are restricted to certain intervals!

Variational model

$$\min_u TV(u) \text{ s.t. } l_{i,j} \leq Au(i,j) \leq b_{i,j}$$

Image decompression

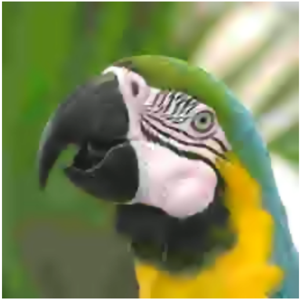

- DCT is nothing but an orthonormal transform A.
- Inverse problem to quantization: Coefficients are restricted to certain intervals!

Variational model

$$\min_u TV(u) \text{ s.t. } l_{i,j} \leq Au(i,j) \leq b_{i,j}$$

With code from: M. Holler, K. Bredis. *A TGV-based framework for variational image decompression, zooming and reconstruction.*

Optical Flow

Given a video, can we estimate the motion in the video?

Middleburry benchmark, *A Database and Evaluation Methodology for Optical Flow*, IJCV. `http://vision.middlebury.edu/flow/`

Optical Flow

Given a video, can we estimate the motion in the video?

Middleburry benchmark, *A Database and Evaluation Methodology for Optical Flow*, IJCV. `http://vision.middlebury.edu/flow/`

Optical Flow

Given a video, can we estimate the motion in the video?

Most common assumption: **color constancy!**

$$v(\alpha) = \arg\min_{v} \frac{1}{2}\int_{\Omega} (f_1(x) - f_2(x + v(x))^2 \, dx + \alpha TV(v)$$

Middleburry benchmark, *A Database and Evaluation Methodology for Optical Flow*, IJCV. `http://vision.middlebury.edu/flow/`

Optical Flow

Given a video, can we estimate the motion in the video?

Difficulty:

$$(A(v))(x) := f_2(x + v(x))$$

is **not a linear operator!**

Middleburry benchmark, *A Database and Evaluation Methodology for Optical Flow*, IJCV. `http://vision.middlebury.edu/flow/`

Nonlinear inverse problems

Form of the optical flow problem:

$$v(\alpha) = \arg\min_{v} \frac{1}{2}\|A(v) - f\|_2^2 + \alpha R(v),$$

for a nonlinear operator A!

Nonlinear inverse problems

Form of the optical flow problem:

$$v(\alpha) = \arg\min_v \frac{1}{2}\|A(v) - f\|_2^2 + \alpha R(v),$$

for a nonlinear operator A!

One strategy for nonlinear inverse problems of the above form:

Iteratively regularized Gauss-Newton-type methods

$$v^{k+1} = \arg\min_v \frac{1}{2}\|A(v^k) + A'(v^k)(v - v^k) - f\|_2^2 + \alpha R(v),$$

Idea:

- Linearize A around current iterate v^k.
- Solve linear inverse problem.
- Repeat until convergence.

Convergence analysis requires smoothness of A.

Optical flow

First order Taylor expansion:

$$f_2(x + v(x)) \approx f_2(x + v^k(x)) + \nabla f_2(x + v^k(x)) \cdot (v(x) - v^k(x))$$

Optical flow

First order Taylor expansion:

$$f_2(x + v(x)) \approx \underbrace{f_2(x + v^k(x))}_{\text{"Warping" } f_2^k} + \nabla \underbrace{f_2(x + v^k(x))}_{\text{"Warping" } f_2^k} \cdot (v(x) - v^k(x))$$

Optical flow

First order Taylor expansion:

$$f_2(x + v(x)) \approx \underbrace{f_2(x + v^k(x))}_{\text{"Warping" } f_2^k} + \nabla \underbrace{f_2(x + v^k(x))}_{\text{"Warping" } f_2^k} \cdot (v(x) - v^k(x))$$

$$v^{k+1} = \arg\min_v \frac{1}{2} \| f_1 - f_2^k - \nabla f_2^k \cdot (v - v^k) \|_2^2 + \alpha TV(v)$$

Optical flow

First order Taylor expansion:

$$f_2(x + v(x)) \approx \underbrace{f_2(x + v^k(x))}_{\text{"Warping" } f_2^k} + \nabla \underbrace{f_2(x + v^k(x))}_{\text{"Warping" } f_2^k} \cdot (v(x) - v^k(x))$$

$$v^{k+1} = \arg\min_v \frac{1}{2} \| f_1 - f_2^k - \nabla f_2^k \cdot (v - v^k) \|_2^2 + \alpha TV(v)$$

Common strategies:

- Pyramid scheme, coarse-to-fine

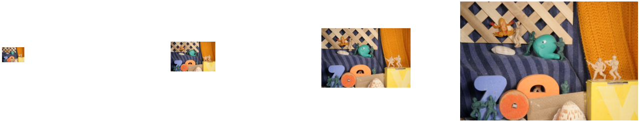

Optical flow

First order Taylor expansion:

$$f_2(x + v(x)) \approx \underbrace{f_2(x + v^k(x))}_{\text{"Warping" } f_2^k} + \nabla \underbrace{f_2(x + v^k(x))}_{\text{"Warping" } f_2^k} \cdot (v(x) - v^k(x))$$

$$v^{k+1} = \arg\min_v \frac{1}{2} \|f_1 - f_2^k - \nabla f_2^k \cdot (v - v^k)\|_2^2 + \alpha TV(v)$$

Common strategies:

- Pyramid scheme, coarse-to-fine

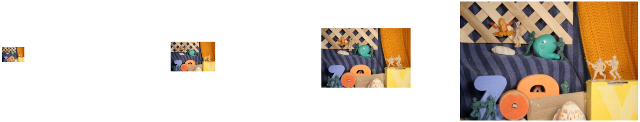

- Different measures of distances:
  - Robustness, e.g. L^1
  - Patch-based fidelities (piecewise rigid motion)

Optical flow

First order Taylor expansion:

$$f_2(x+v(x)) \approx \underbrace{f_2(x+v^k(x))}_{\text{"Warping" } f_2^k} + \nabla \underbrace{f_2(x+v^k(x))}_{\text{"Warping" } f_2^k} \cdot (v(x) - v^k(x))$$

$$v^{k+1} = \arg\min_v \frac{1}{2}\|f_1 - f_2^k - \nabla f_2^k \cdot (v - v^k)\|_2^2 + \alpha TV(v)$$

Common strategies:

- Pyramid scheme, coarse-to-fine

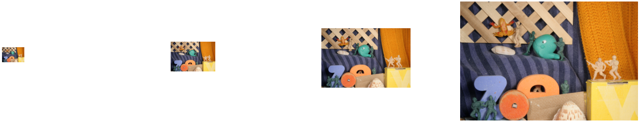

- Different measures of distances:
  - Robustness, e.g. L^1
  - Patch-based fidelities (piecewise rigid motion)
- (Smoothed) brute-force search for good initialization

Optical flow

First order Taylor expansion:

$$f_2(x + v(x)) \approx \underbrace{f_2(x + v^k(x))}_{\text{"Warping" } f_2^k} + \nabla \underbrace{f_2(x + v^k(x))}_{\text{"Warping" } f_2^k} \cdot (v(x) - v^k(x))$$

$$v^{k+1} = \arg\min_v \frac{1}{2} \|f_1 - f_2^k - \nabla f_2^k \cdot (v - v^k)\|_2^2 + \alpha TV(v)$$

Common strategies:

- Pyramid scheme, coarse-to-fine

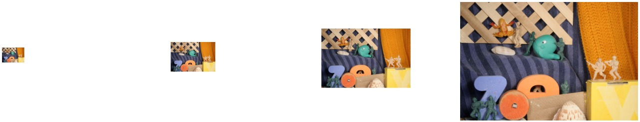

- Different measures of distances:
  - Robustness, e.g. L^1
  - Patch-based fidelities (piecewise rigid motion)
- (Smoothed) brute-force search for good initialization

Example video.

Stereo Matching

For known camera positions, we can relate the optical flow to the 3d coordinates!

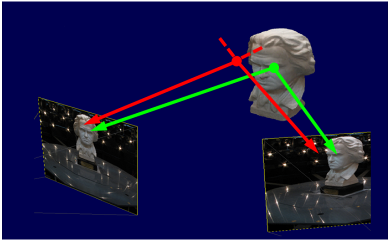

Stereo Matching

For known camera positions, we can relate the optical flow to the 3d coordinates!

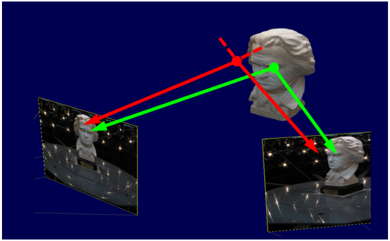

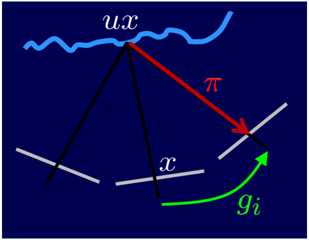

Stereo Matching

For known camera positions, we can relate the optical flow to the 3d coordinates!

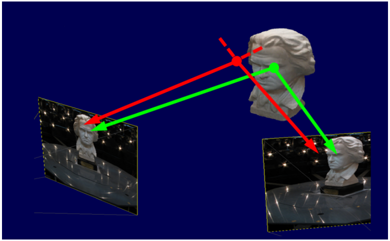

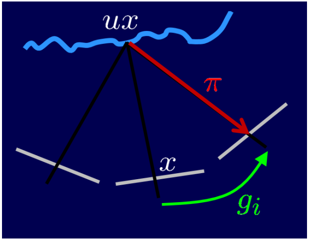

Optimize photo-consistency to find depth:

$$\sum_i \frac{1}{2} \int_\Omega |f_1(x) - f_i(\pi(g_i(u \cdot x)))| \, dx + \alpha TV(u)$$

3d Reconstruction

Full 3d reconstruction (segmentation problem in 3d):

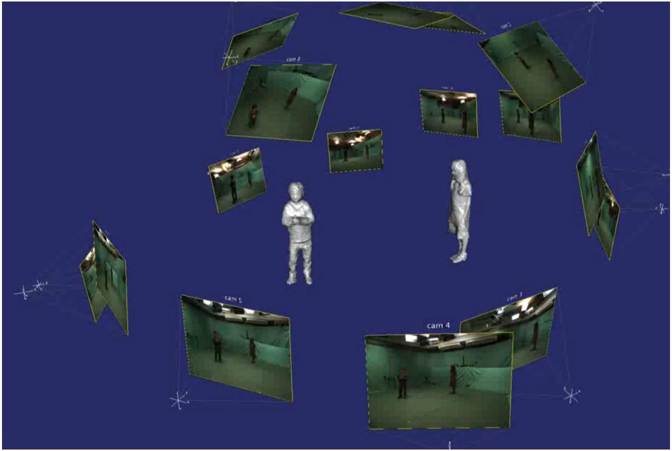

A Convex Relaxation Approach to Space Time Multi-view 3D Reconstruction. Oswald, Cremers.

Stereo Matching

Simpler version of the optical flow: Stereo matching!

Stereo images after rectification

From: Wikipedia, `https://de.wikipedia.org/wiki/Stereokamera`
Scharstein, Szeliski. *A taxonomy and evaluation of dense two-frame stereo correspondence algorithms.* `http://vision.middlebury.edu/stereo/`

Stereo Matching

$$v(\alpha) = \arg\min_{v} \frac{1}{2} \int_{\Omega} (f_1(x_1, x_2) - f_2(x_1, x_2 + v(x_1, x_2)))^2 \, dx + TV(v)$$

Stereo Matching

$$v(\alpha) = \arg\min_v \frac{1}{2} \int_\Omega \underbrace{(f_1(x_1, x_2) - f_2(x_1, x_2 + v(x_1, x_2)))^2}_{=\rho(x,v) \text{ nonconvex, but 1d!}} dx + TV(v)$$

Stereo Matching

$$v(\alpha) = \arg\min_v \frac{1}{2}\int_\Omega \underbrace{(f_1(x_1,x_2) - f_2(x_1,x_2+v(x_1,x_2)))^2}_{=\rho(x,v) \text{ nonconvex, but 1d!}} dx + TV(v)$$

First option: Same as optical flow:

$$f_2(x_1,x_2+v(x_1,x_2)) \approx f_2^k + (v - v^k)\,\partial_x f_2^k$$

Stereo Matching

$$v(\alpha) = \arg\min_v \frac{1}{2} \int_\Omega \underbrace{(f_1(x_1,x_2) - f_2(x_1, x_2 + v(x_1,x_2)))^2}_{=\rho(x,v) \text{ nonconvex, but 1d!}} dx + TV(v)$$

First option: Same as optical flow:

$$f_2(x_1, x_2 + v(x_1,x_2)) \approx f_2^k + (v - v^k)\, \partial_x f_2^k$$

Second option: Can we convexify the data term at each point?

Convex Conjugate

Convex conjugate

Let $F : X \to \mathbb{R} \cup \{\infty\}$. We define $F^* : X^* \to \mathbb{R} \cup \{\infty\}$ by

$$F^*(p) = \sup_{u \in X} \langle u, p \rangle - F(u)$$

Convex Conjugate

Convex conjugate

Let $F : X \to \mathbb{R} \cup \{\infty\}$. We define $F^* : X^* \to \mathbb{R} \cup \{\infty\}$ by

$$F^*(p) = \sup_{u \in X} \langle u, p \rangle - F(u)$$

Note that

$$F(u) + F^*(p) \geq \langle u, p \rangle$$

for all u, p.

Convex Conjugate

Convex conjugate

Let $F : X \to \mathbb{R} \cup \{\infty\}$. We define $F^* : X^* \to \mathbb{R} \cup \{\infty\}$ by

$$F^*(p) = \sup_{u \in X} \langle u, p \rangle - F(u)$$

Note that

$$F(u) + F^*(p) \geq \langle u, p \rangle$$

for all u, p.

Note that $F^*(p)$ is convex.

Convex Conjugate

Convex conjugate

Let $F : X \to \mathbb{R} \cup \{\infty\}$. We define $F^* : X^* \to \mathbb{R} \cup \{\infty\}$ by

$$F^*(p) = \sup_{u \in X} \langle u, p \rangle - F(u)$$

Note that

$$F(u) + F^*(p) \geq \langle u, p \rangle$$

for all u, p.

Note that $F^*(p)$ is convex.

Note that

$$F(u) \geq F^{**}(u)$$

Convex Conjugate

Biconjugate

The biconjugate F^{**} of F is the largest lower semi-continuous convex underapproximation of F, i.e. $F^{**} \leq F$.

If F is a proper, lower-semi continuous, convex function, then $F^{**} = F$.

Functional lifting

Discuss (=try to draw) some convex relaxations!

Functional lifting

Discuss (=try to draw) some convex relaxations!

May work very well and may come at the risk to loose some information.

Functional lifting

Discuss (=try to draw) some convex relaxations!

May work very well and may come at the risk to loose some information.

Idea to have a convex problem but approximate the original energy more closely:
Move to a higher dimensional space!

Functional lifting

Discuss (=try to draw) some convex relaxations!

May work very well and may come at the risk to loose some information.

Idea to have a convex problem but approximate the original energy more closely:
Move to a higher dimensional space!

Consider $E : \mathbb{R} \to \mathbb{R}$ nonconvex and try to find $\min_v E(v)$.

- Idea: Discretize the possible values of v: $v_1, ..., v_l$.
- Move to higher dimensions: $u \in \mathbb{R}^l$ with

$$u = e_i \qquad \text{means} \qquad v = v_i$$

– Functional lifting –

Functional lifting

Discuss (=try to draw) some convex relaxations!

May work very well and may come at the risk to loose some information.

Idea to have a convex problem but approximate the original energy more closely:
Move to a higher dimensional space!

Consider $E : \mathbb{R} \to \mathbb{R}$ nonconvex and try to find $\min_v E(v)$.

- Idea: Discretize the possible values of v: $v_1, ..., v_l$.
- Move to higher dimensions: $u \in \mathbb{R}^l$ with

$$u = e_i \qquad \text{means} \qquad v = v_i$$

– Functional lifting –

Reformulate energy

$$\bar{E}(u) = \begin{cases} E(v_i) & \text{if } u = e_i \\ \infty & \text{else.} \end{cases}$$

Convex relaxation

The functional

$$\bar{E}(u) = \begin{cases} E(v_i) & \text{if } u = e_i \\ \infty & \text{else.} \end{cases}$$

is (still) not convex.

Convex relaxation

The functional

$$\bar{E}(u) = \begin{cases} E(v_i) & \text{if } u = e_i \\ \infty & \text{else.} \end{cases}$$

is (still) not convex.

Compute the best possible convex underapproximation (E^{**}):

$$\bar{E}^{**}(u) = \begin{cases} \sum_i u_i(x) E(v_i) & \text{if } u_i \geq 0, \sum_i u_i = 1 \\ \infty & \text{else.} \end{cases}$$

Convex relaxation

The functional

$$\bar{E}(u) = \begin{cases} E(v_i) & \text{if } u = e_i \\ \infty & \text{else.} \end{cases}$$

is (still) not convex.

Compute the best possible convex underapproximation (E^{**}):

$$\bar{E}^{**}(u) = \begin{cases} \sum_i u_i(x) E(v_i) & \text{if } u_i \geq 0, \sum_i u_i = 1 \\ \infty & \text{else.} \end{cases}$$

Or for $\rho_i := E(v_i)$

$$\bar{E}^{**}(u) = \begin{cases} \langle u, \rho \rangle & \text{if } u_i \geq 0, \sum_i u_i = 1 \\ \infty & \text{else.} \end{cases}$$

Back to Stereo Matching

Stereo matching energy:

$$v(\alpha) = \arg\min_{v} \frac{1}{2} \underbrace{\int_{\Omega} (f_1(x_1, x_2) - f_2(x_1, x_2 + v(x_1, x_2)))^2 \; dx}_{=\rho(x,v) \text{ nonconvex, but 1d!}} + TV(v)$$

Back to Stereo Matching

Stereo matching energy:

$$v(\alpha) = \arg\min_v \frac{1}{2} \underbrace{\int_\Omega (f_1(x_1,x_2) - f_2(x_1, x_2 + v(x_1,x_2)))^2 \, dx}_{=\rho(x,v) \text{ nonconvex, but 1d!}} + TV(v)$$

Introduce a variable u that is lifted for each x, i.e. $u(x) \in \mathbb{R}^n$.

Back to Stereo Matching

Stereo matching energy:

$$v(\alpha) = \arg\min_{v} \underbrace{\frac{1}{2}\int_{\Omega} (f_1(x_1,x_2) - f_2(x_1, x_2 + v(x_1,x_2)))^2 \, dx}_{=\rho(x,v) \text{ nonconvex, but 1d!}} + TV(v)$$

Introduce a variable u that is lifted for each x, i.e. $u(x) \in \mathbb{R}^n$.

Total energy:

$$E(u) = \frac{1}{2}\int_{\Omega} \langle u(x), \rho(x) \rangle \, dx + \alpha R(u) + \mathfrak{i}_{\Delta}(u)$$

for a suitable regularization R which mimics $TV(v)$.

Back to Stereo Matching

Stereo matching energy:

$$v(\alpha) = \arg\min_{v} \underbrace{\frac{1}{2}\int_{\Omega} (f_1(x_1,x_2) - f_2(x_1, x_2 + v(x_1,x_2)))^2 \, dx}_{=\rho(x,v) \text{ nonconvex, but 1d!}} + TV(v)$$

Introduce a variable u that is lifted for each x, i.e. $u(x) \in \mathbb{R}^n$.

Total energy:

$$E(u) = \frac{1}{2}\int_{\Omega} \langle u(x), \rho(x) \rangle \, dx + \alpha R(u) + \imath_{\Delta}(u)$$

for a suitable regularization R which mimics $TV(v)$.

Further readings:

- Convex Relaxation of Vectorial Problems with Coupled Regularization. Strekalovskiy, Chambolle, Cremers. References therein.
- Nonsmooth Convex Variational Approaches to Image Analysis. PhD thesis Jan Lellmann.
- High-dimension multi-label problems: convex or non convex relaxation? Papadakis, Yildizoglu, Aujol, Caselles.

Can you do all this yourself?

We have seen a great number of applications, all based on

$$u(\alpha) = \arg\min_u E_\alpha(u)$$

(mostly for convex energies E).

Can you do all this yourself?

We have seen a great number of applications, all based on

$$u(\alpha) = \arg\min_u E_\alpha(u)$$

(mostly for convex energies E).

Next lecture:

How can we solve these problems in practice?