Precise Relaxation for Motion Estimation

Jan Lellmann, Emanuel Laude, Thomas Moellenhoff, Daniel Cremers, Michael Moeller, Evgeny Strekalovskiy

SIAM Conference on Imaging Science
Albuquerque, 24 May 2016

MIC
Institute of Mathematics and Image Computing

lellmann@mic.uni-luebeck.de

Variational motion estimation

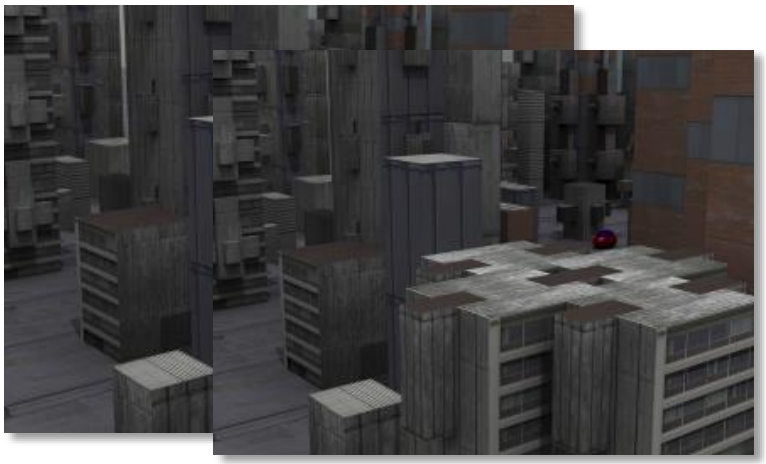

image sequence

optical flow/registration

$$\min_{u:\Omega\to\mathbb{R}^2} \int_\Omega D(I_\sigma^1(x), I_\sigma^2(x+u))dx + \lambda \int_\Omega d\|Du\|$$

u scalar (disparity, stereo) or vector field (optical flow, image registration)

Images: Lellmann, Strekalovskiy, Kötter, Cremers '13

Variational methods

We reconstruct the unknown data *u* from the measurements *b* by minimizing the energy

$$\min_u \{D(T(u); b) + R(u)\}$$

Intuitive (what do we want) and modular (reusable)

In practice: often

$$\min_{u:\Omega \to X} \int_\Omega \rho(x, u(x))dx + \lambda \int_\Omega \sigma(\nabla u)dx$$

Tikhonov '63

Convexity

Convexity assures that every local minimizer of the energy is also a global minimizer

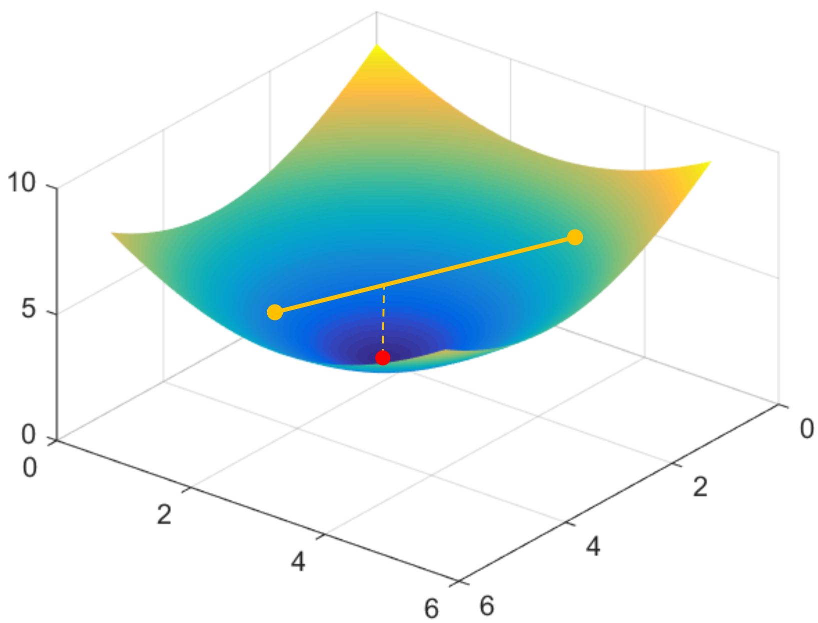

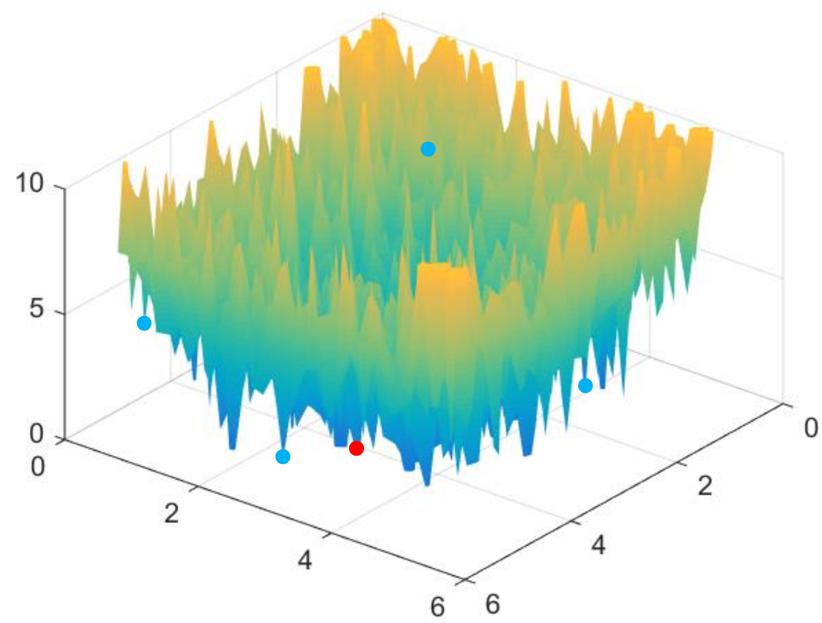

Variational motion estimation

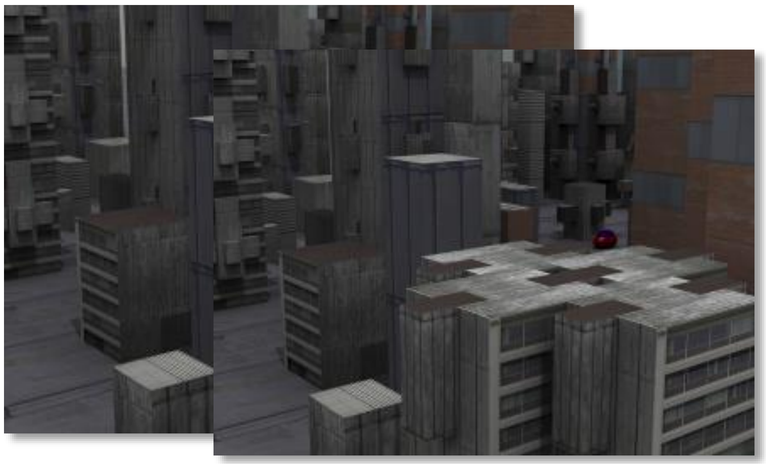

image sequence

optical flow/registration

$$\min_{u:\Omega\to\mathbb{R}^2} \int_\Omega D(I_\sigma^1(x), I_\sigma^2(x+u))dx + \lambda \int_\Omega d\|Du\|$$

u scalar (disparity, stereo) or vector field (optical flow, image registration)

Images: Lellmann, Strekalovskiy, Kötter, Cremers '13

Non-convexity

- Real-world disparity estimation/depth from stereo:

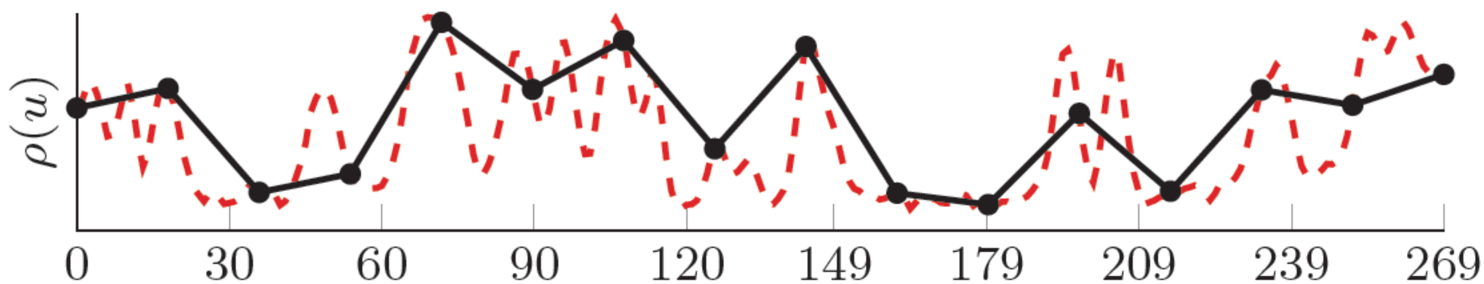

- Data-dependent nonconvexity at every point

Möllenhoff, Laude, Möller, Lellmann, Cremers, CVPR'16

Exact convex relaxation

Exact convex relaxation

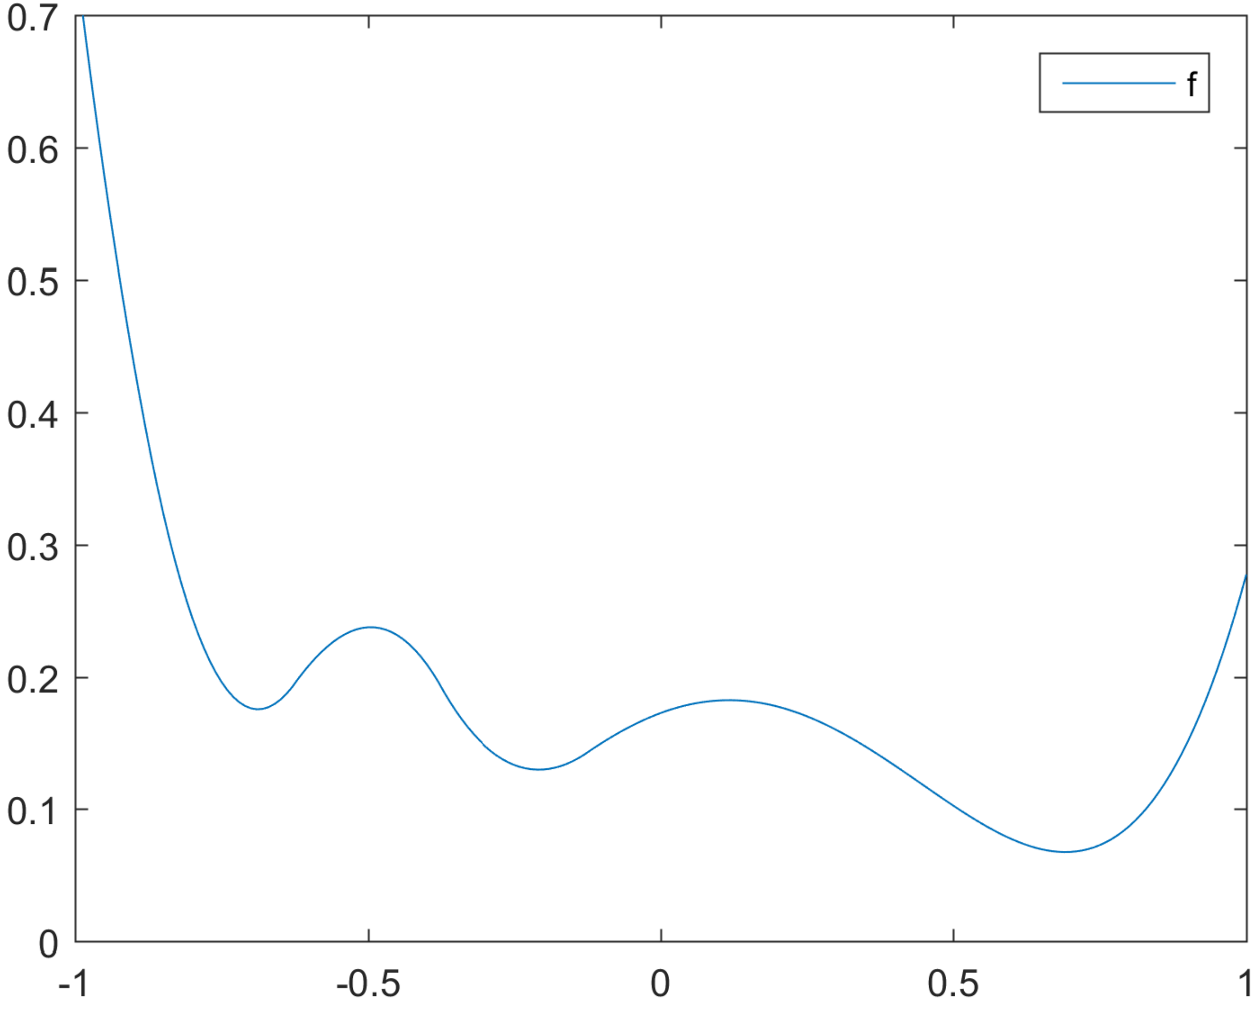

Exact convex relaxation

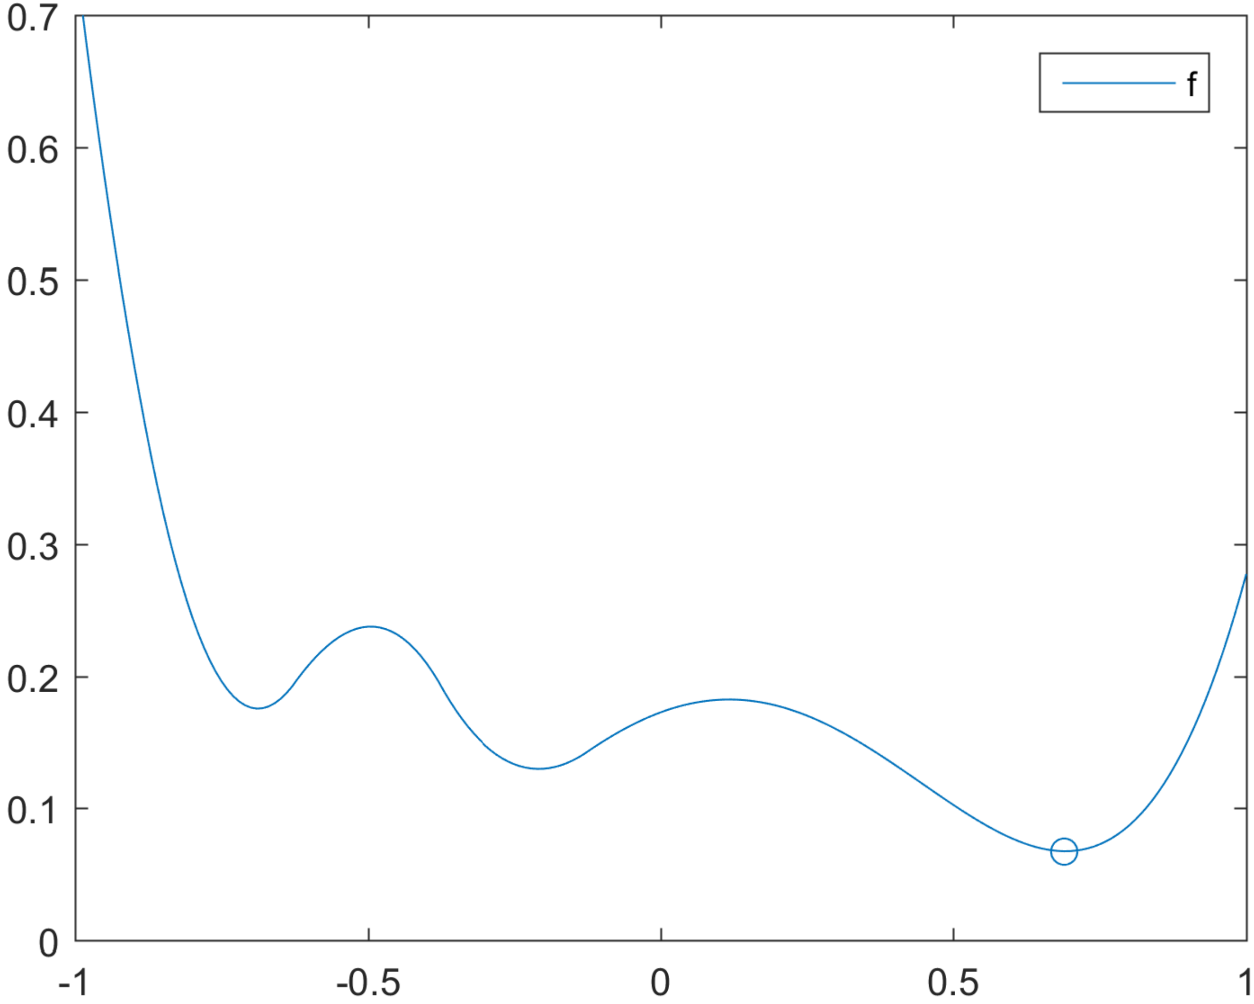

Exact convex relaxation

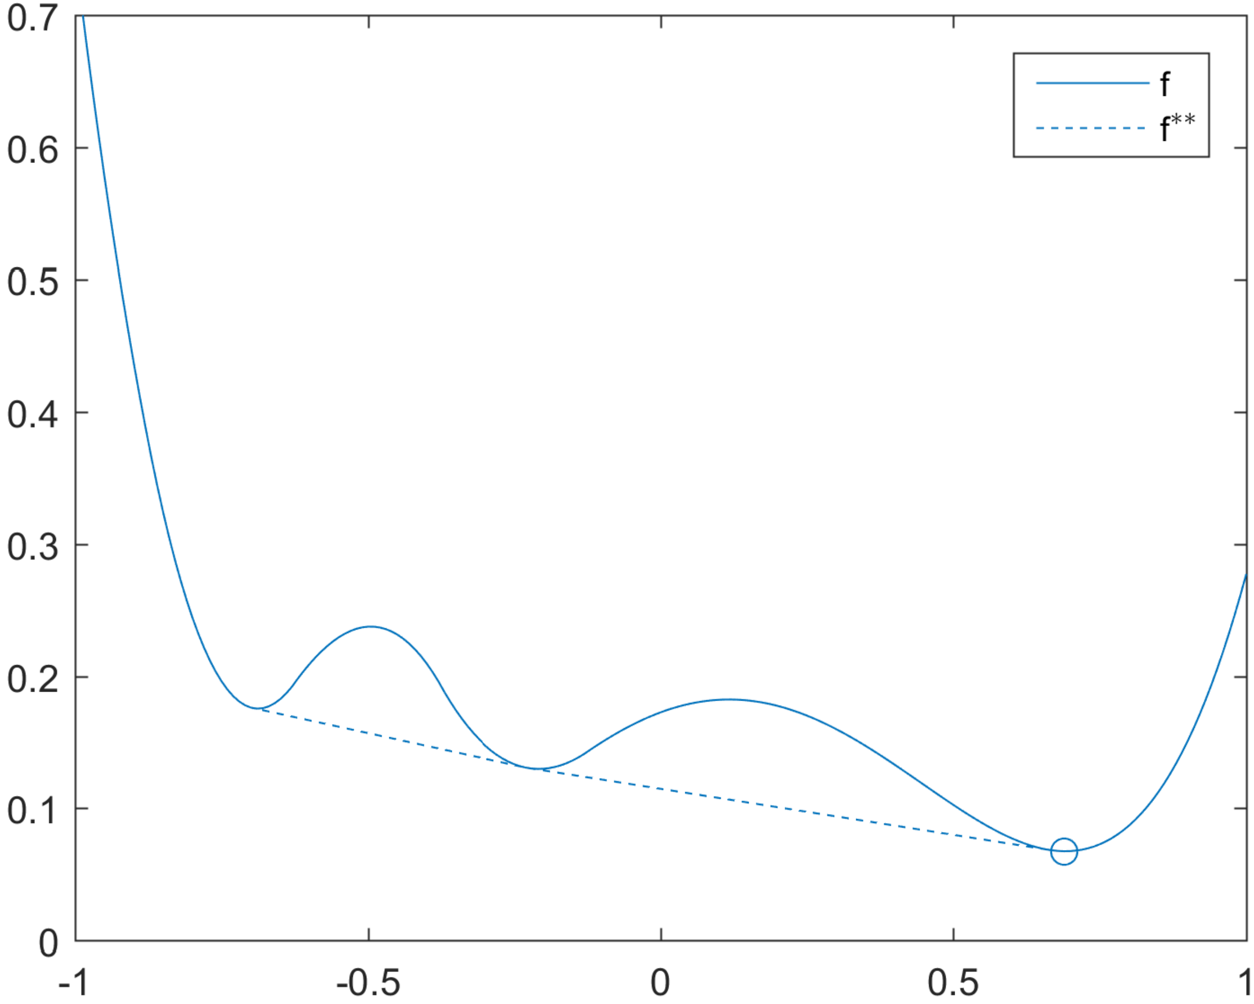

Drawback

This is generally as hard as minimizing the original energy!

Objectives are often sums

$$\min_{u:\Omega\to X} \int_\Omega \rho(x, u(x))dx + \lambda \int_\Omega \sigma(\nabla u)dx$$

Objectives are often sums

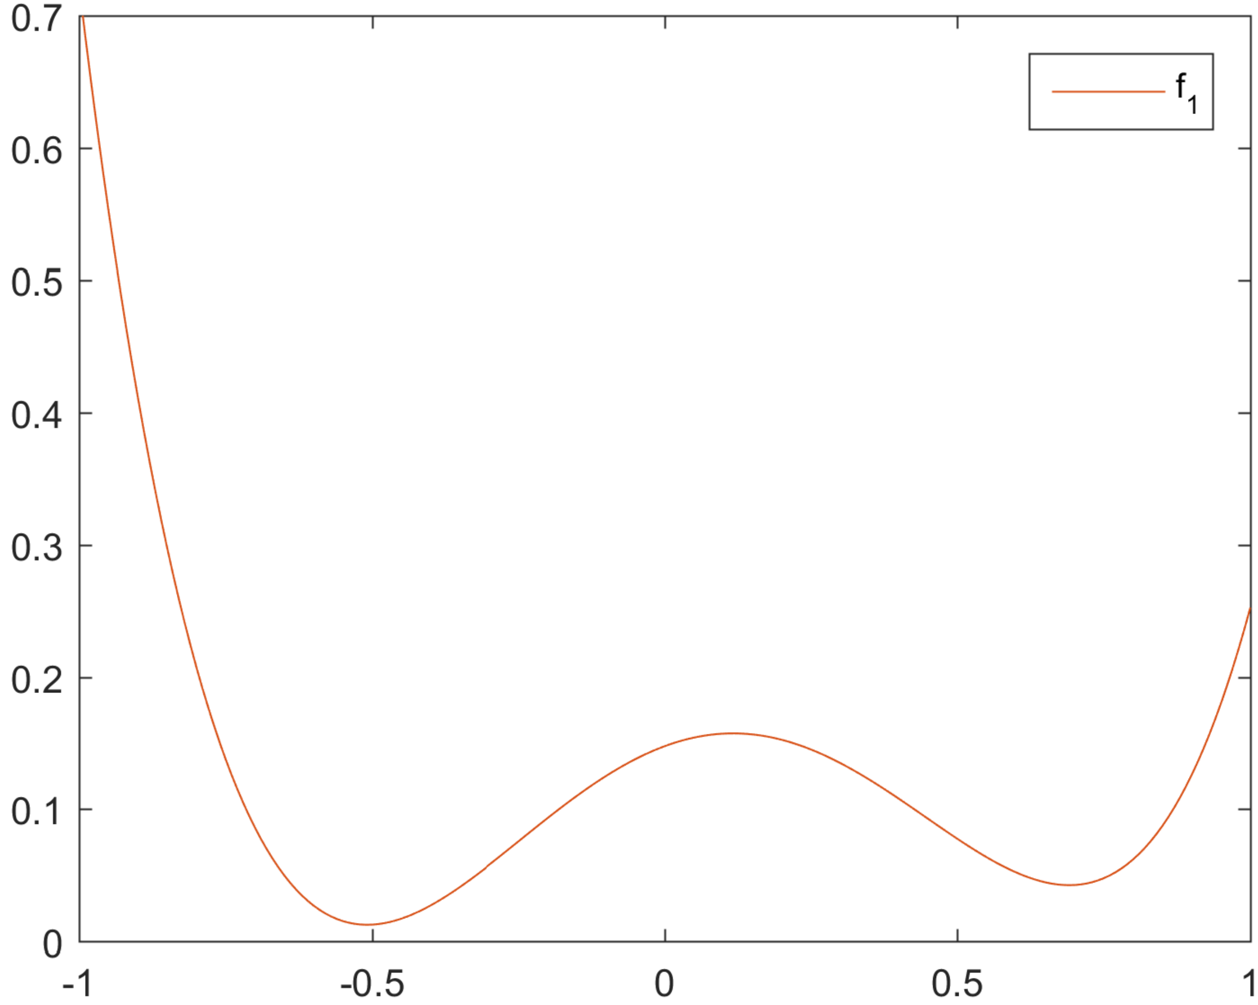

Objectives are often sums

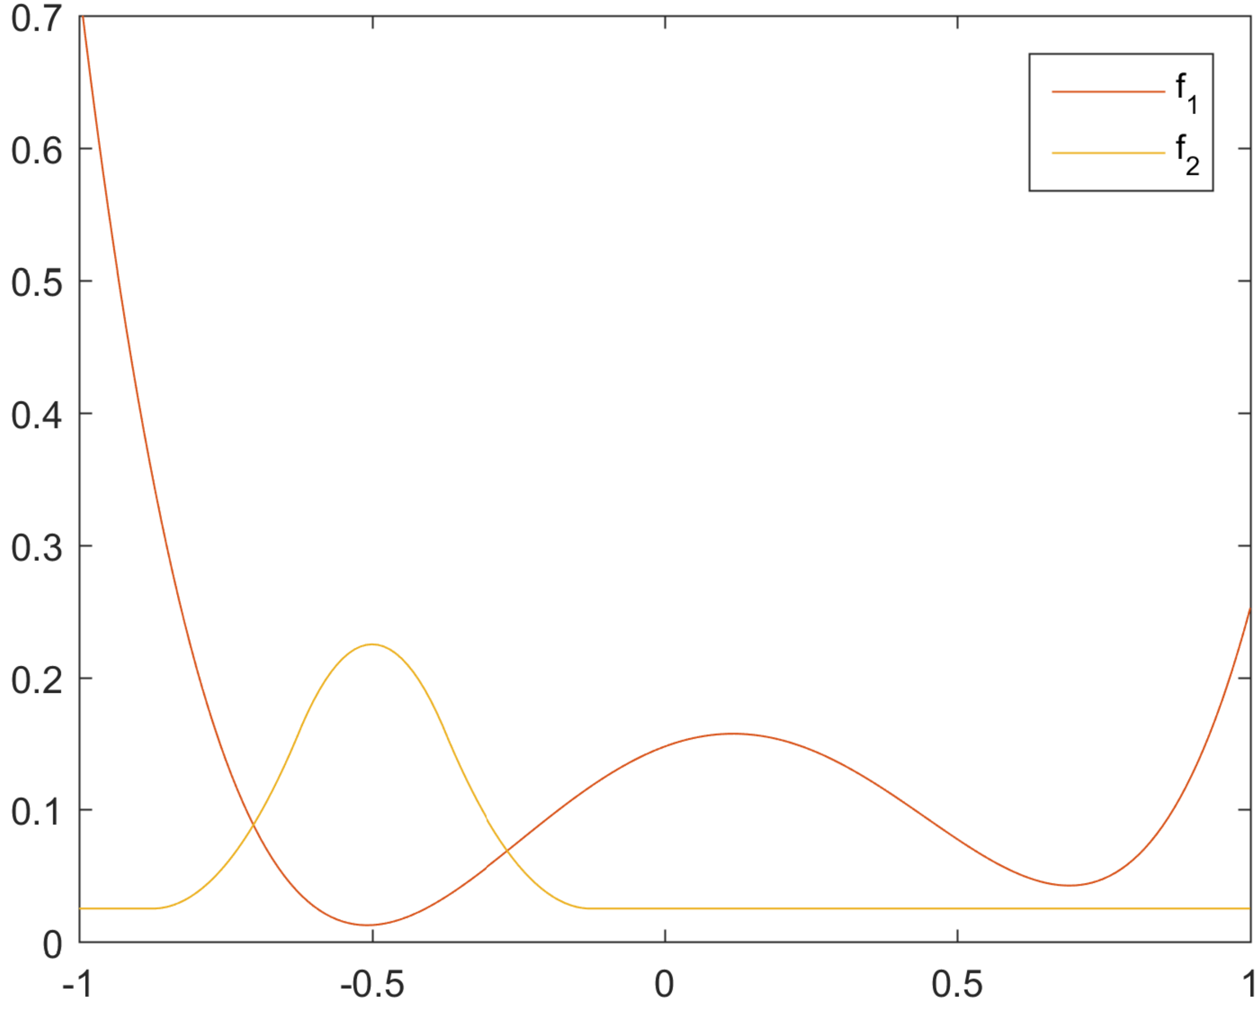

Objectives are often sums

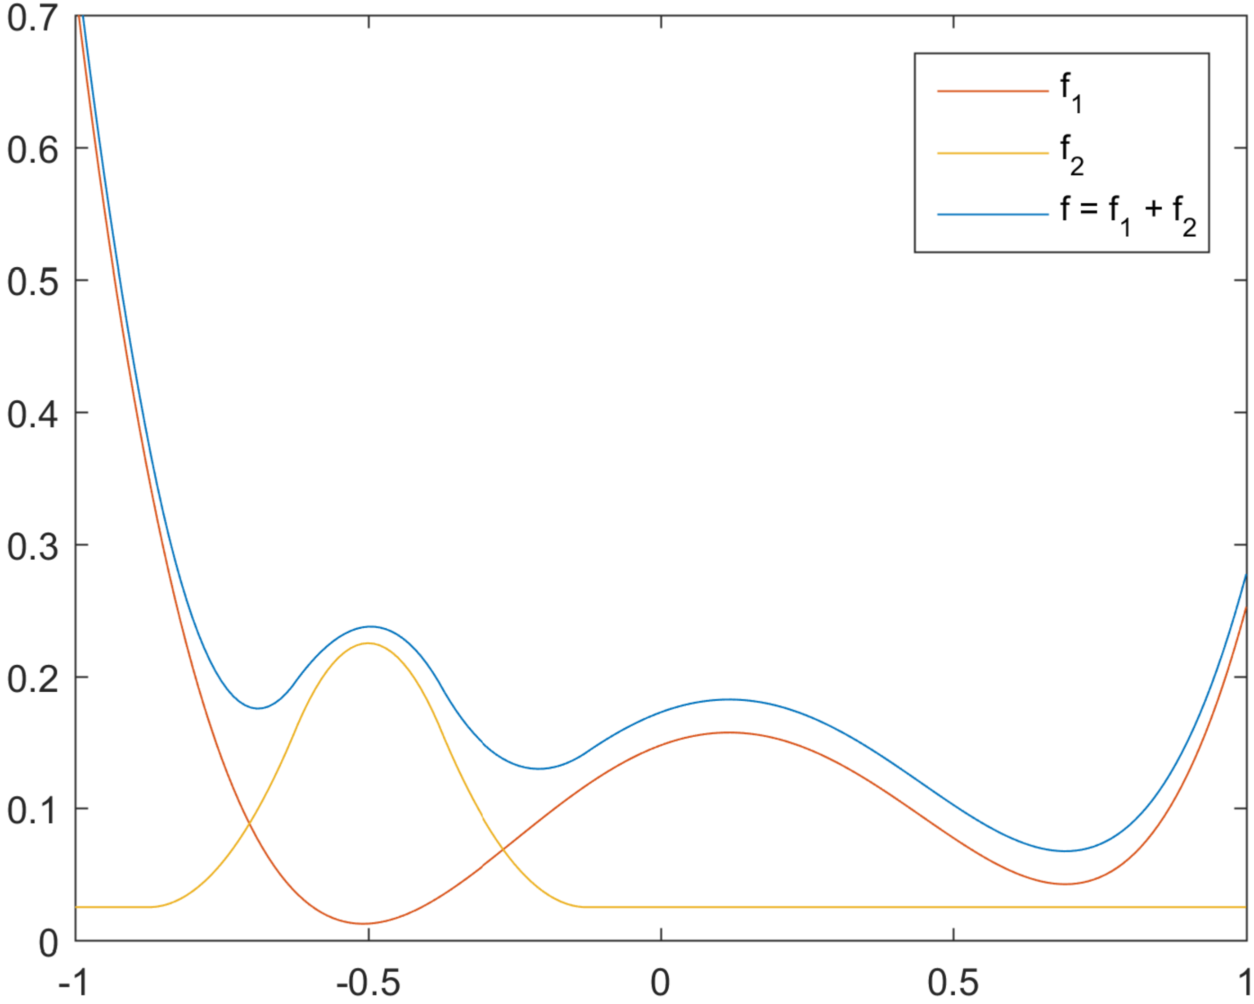

Separate exact relaxation

Separate exact relaxation

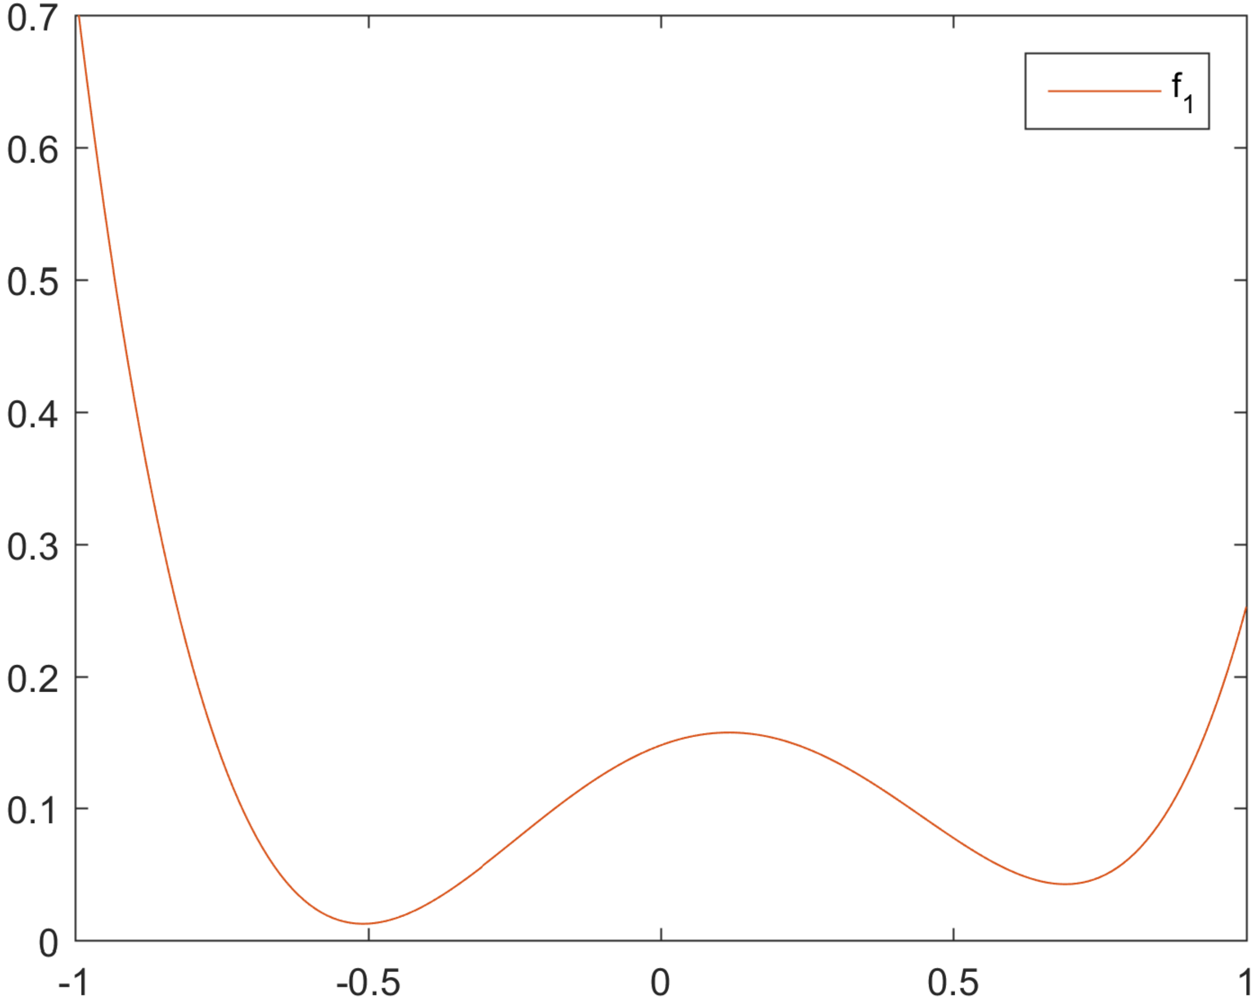

Separate exact relaxation

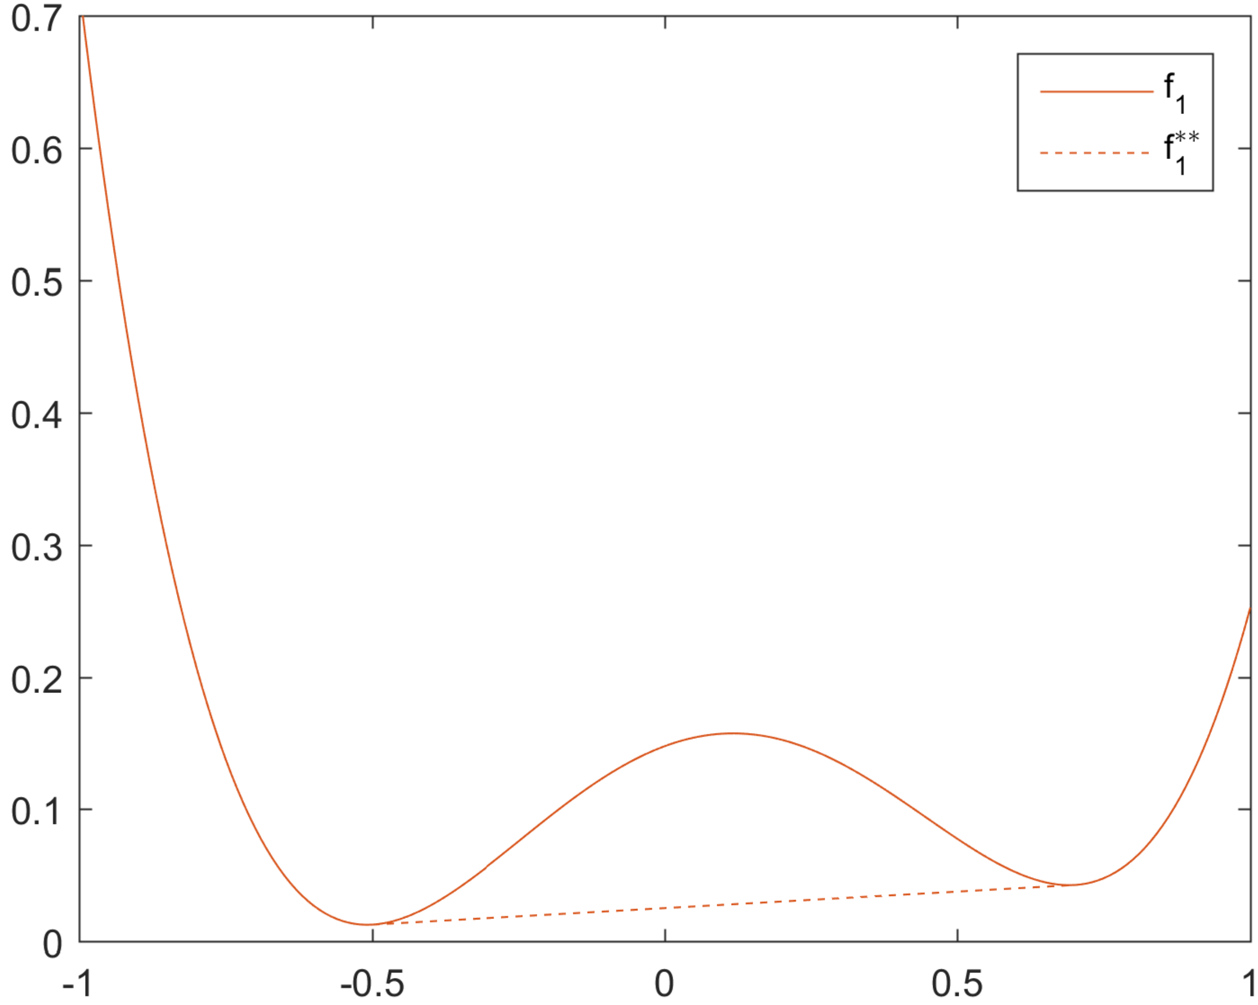

Separate exact relaxation

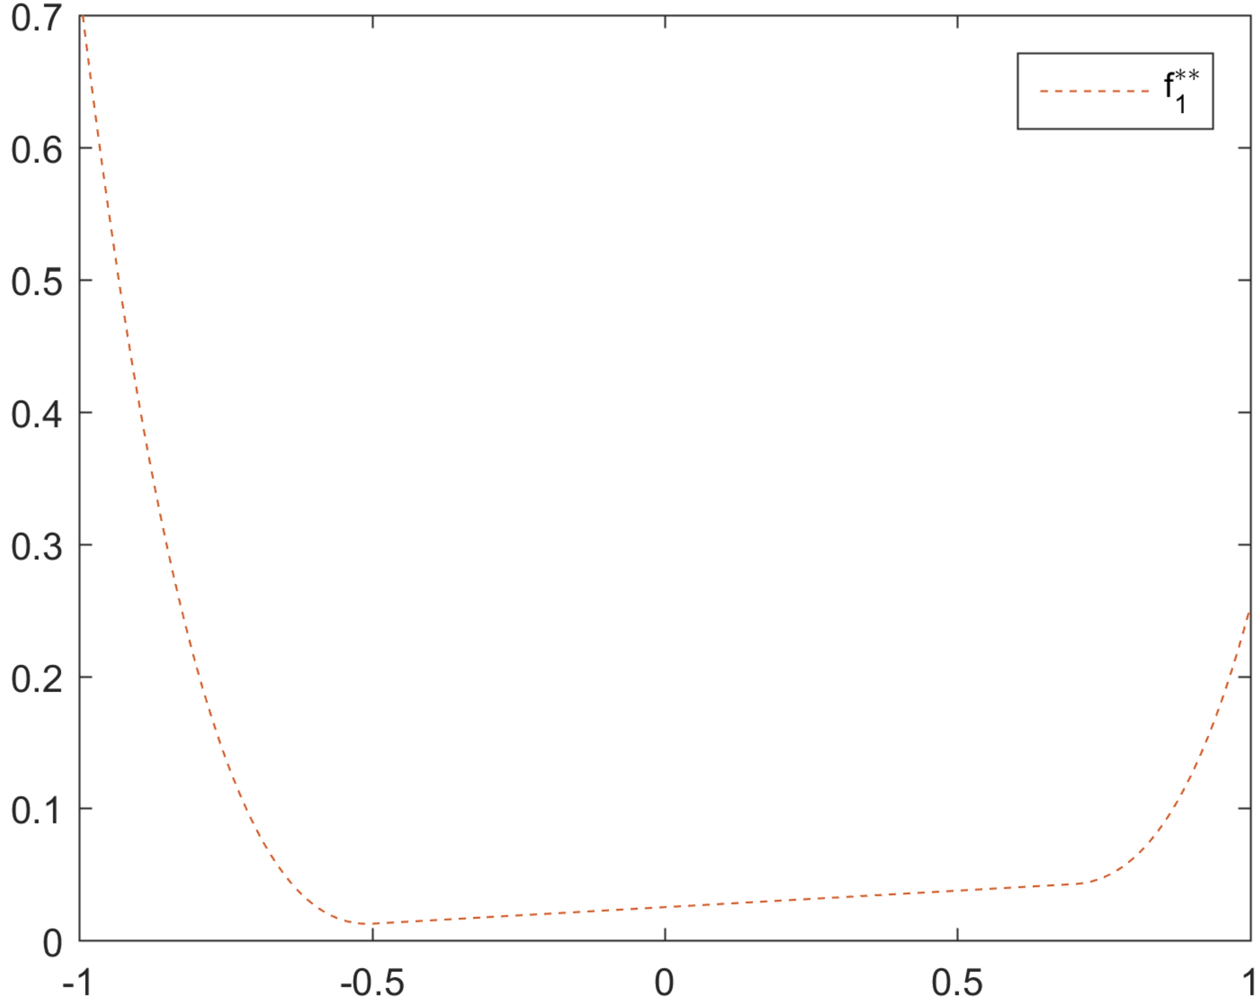

Separate exact relaxation

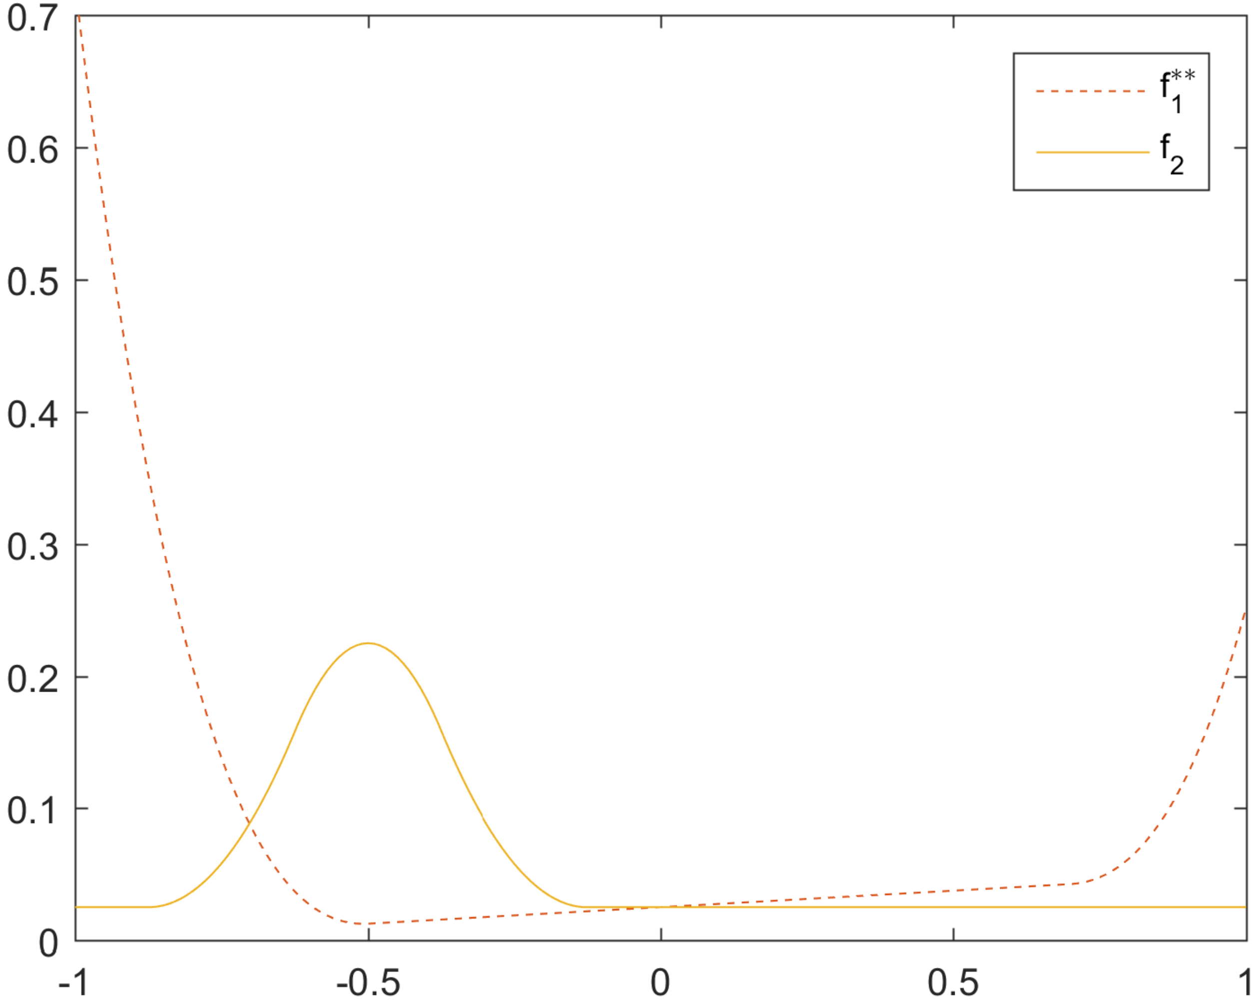

Separate exact relaxation

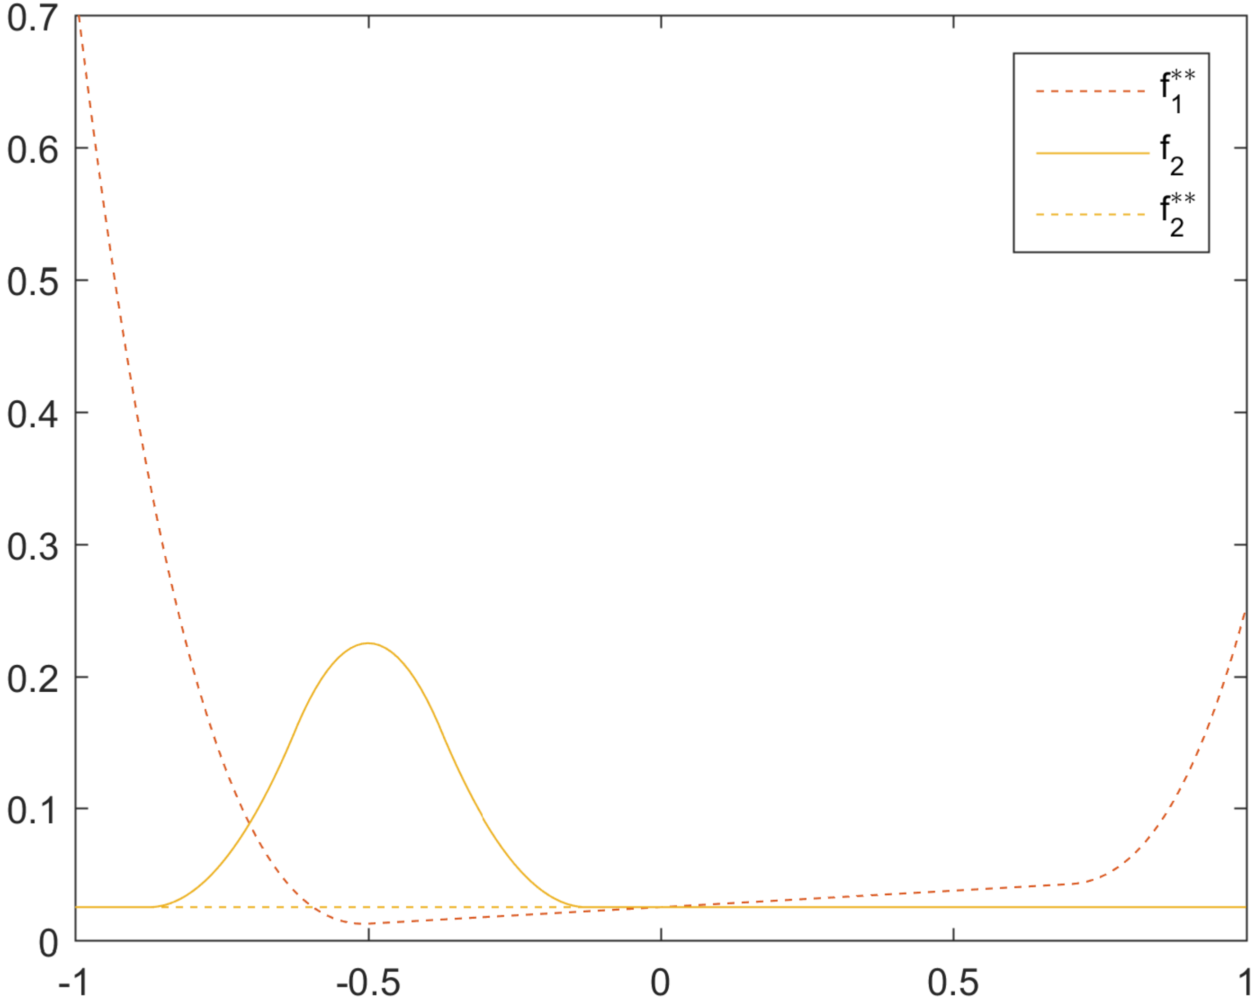

Separate exact relaxation

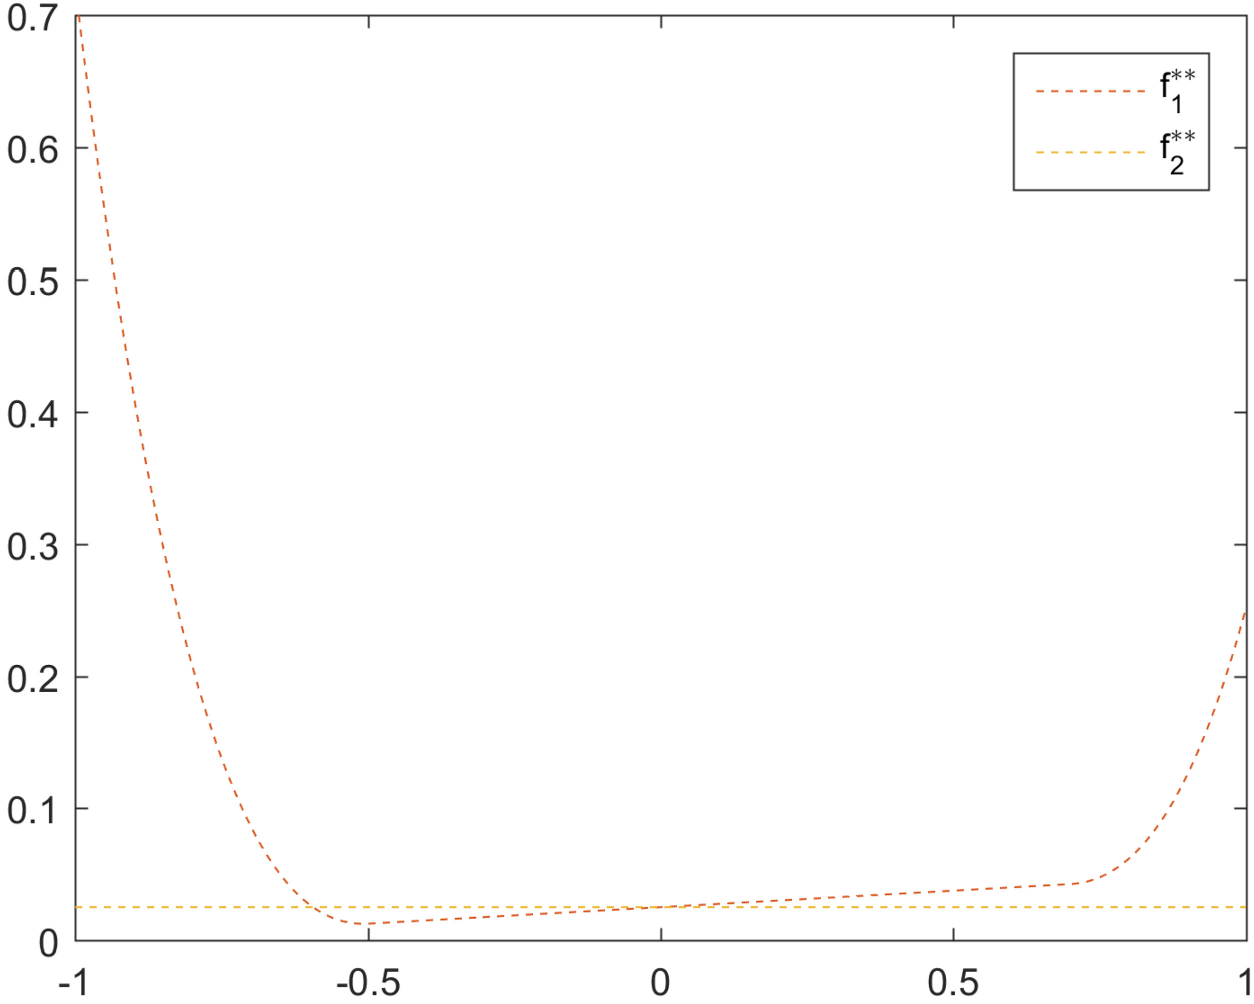

Separate exact relaxation

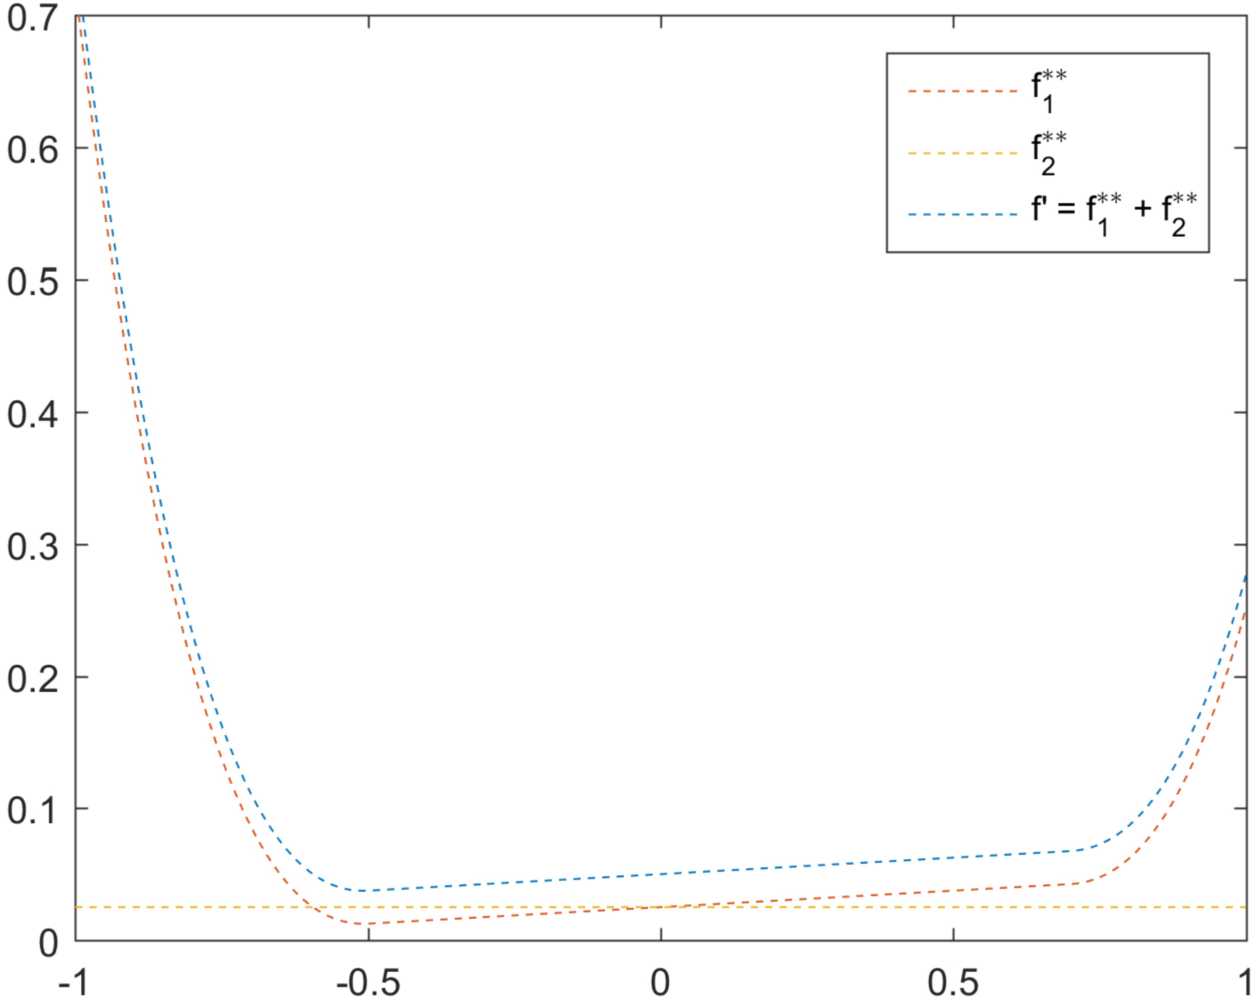

Separate exact relaxation

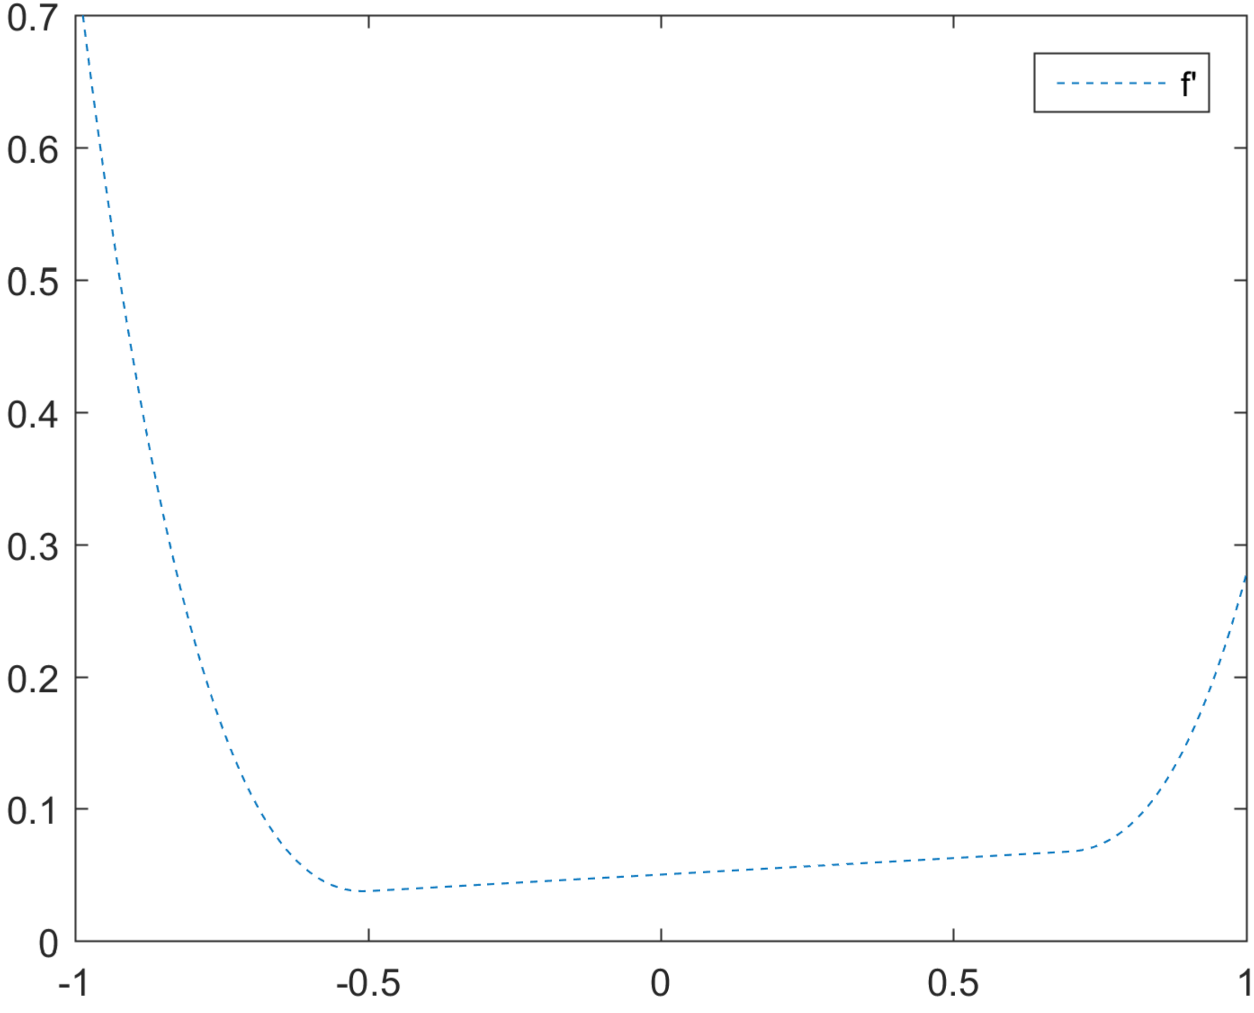

Separate exact relaxation

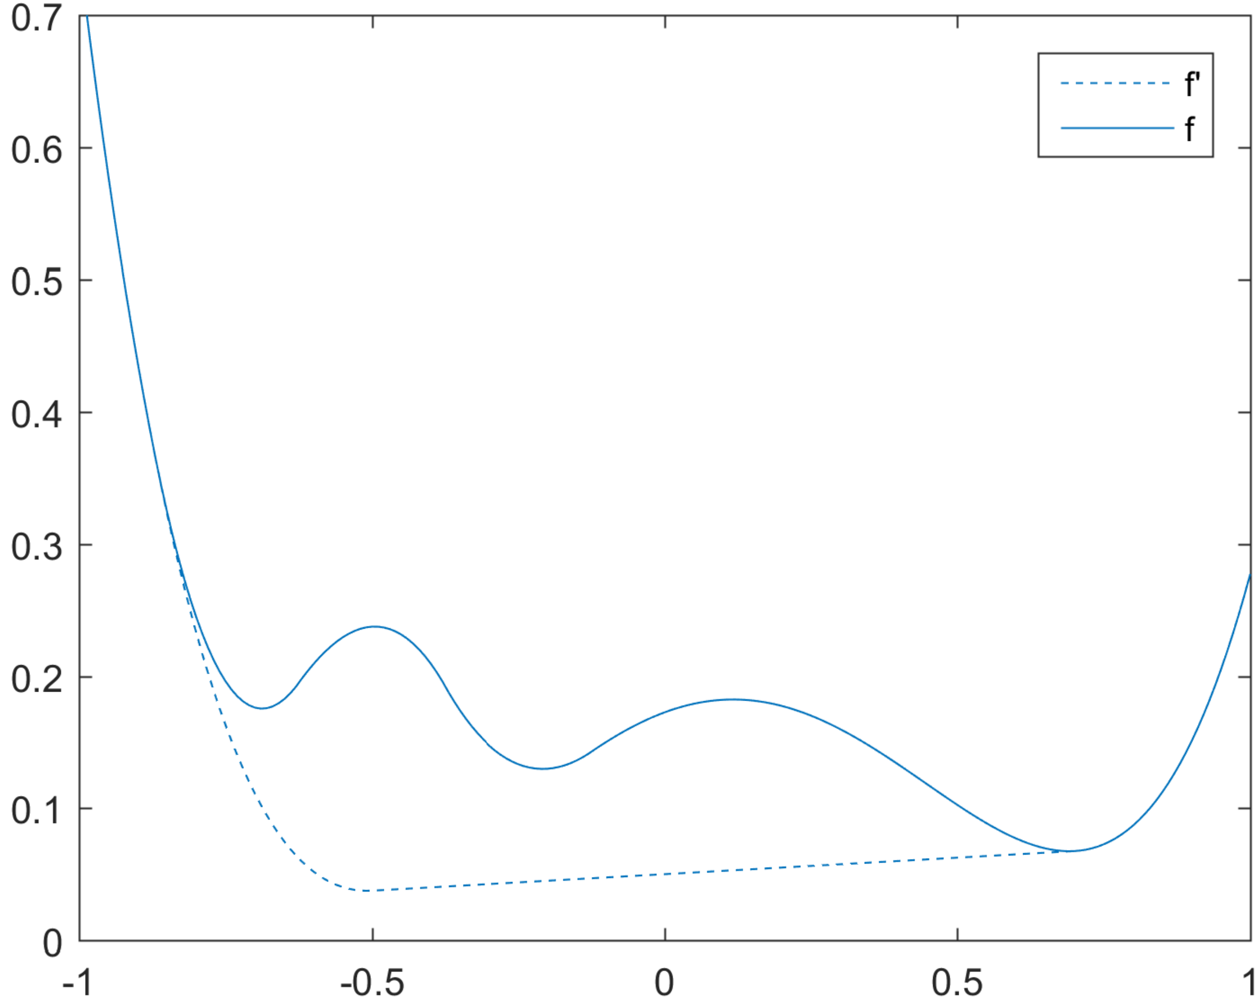

Separate exact relaxation

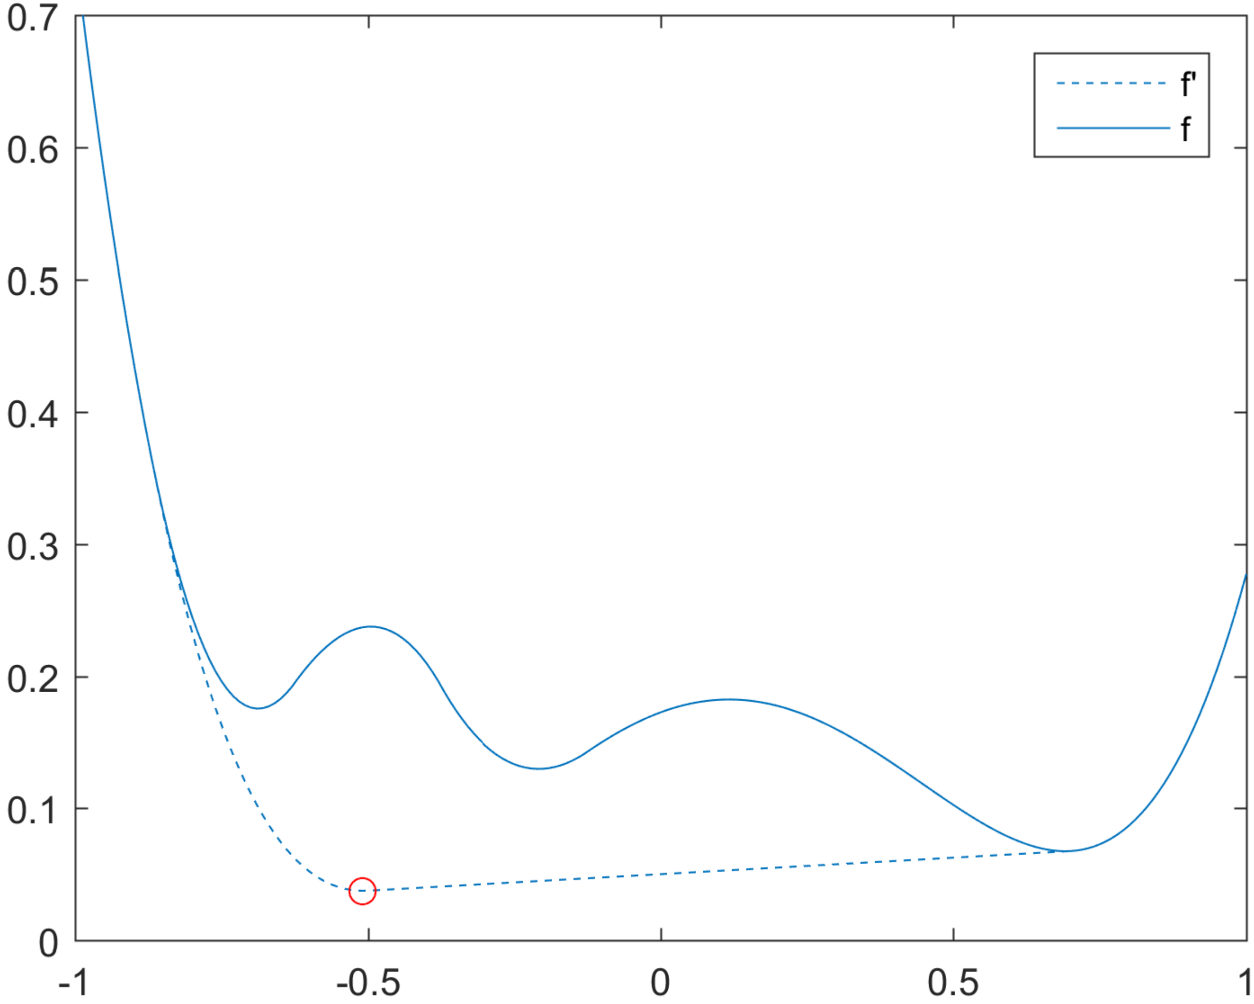

Lifting

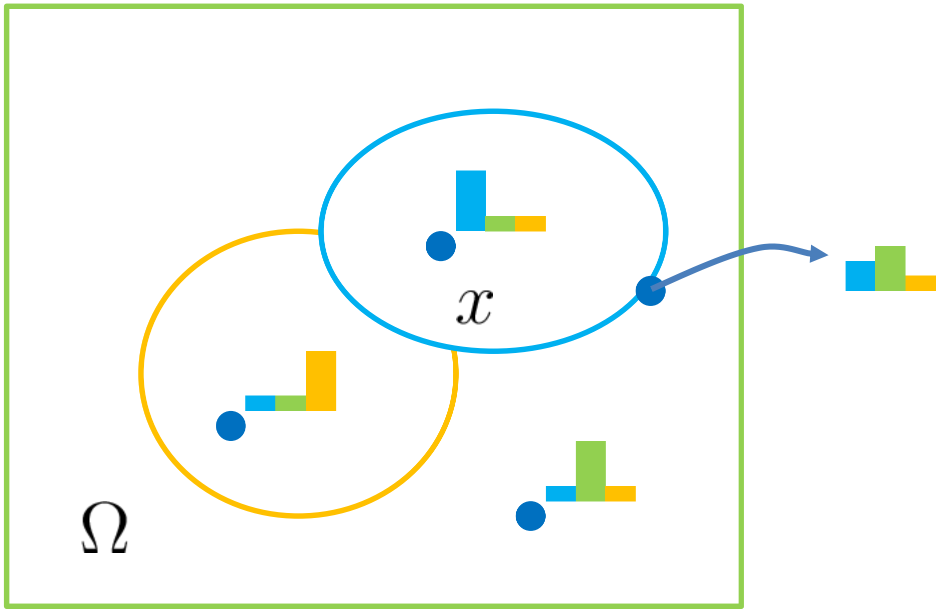

Hard decisions are replaced by soft "probabilities"

Lifting

We extend the problem

$$\min_{u':\Omega\to X} f'(u')$$

to the probability measures:

$$\min_{u:\Omega\to\mathcal{P}(X)} f(u)$$

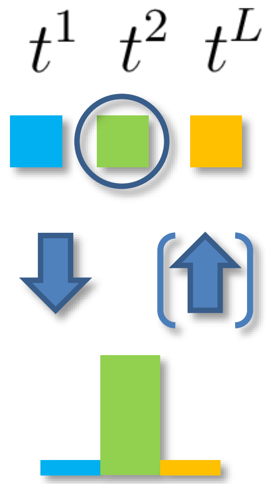

The new energy *f* should agree with *f′* on Dirac measures, and not create artificial minimizers.

Potts '52; Boykov et al. '98, '01; Kleinberg, Tardos '01
Zach et al. '08; Lellmann, Becker, Schnörr '09; Chambolle, Cremers, Pock '11; Yuan, Bae, Tai, Boykov'10
Young measures: Young '37; Currents: Schwartz '51; de Rahm '55; Federer '69
Paired calibrations: Brakke '91; Alberti, Bouchitté, Dal Maso '01

Linear relaxation

- Lifting + relaxation using the biconjugate:

$$\int_\Omega \rho(u(x))dx \rightsquigarrow \int_\Omega \boldsymbol{\rho}^{**}(u(x))dx$$

Linear relaxation

- Lifting + relaxation using the biconjugate:

$$\int_\Omega \rho(u(x))dx \rightsquigarrow \int_\Omega \boldsymbol{\rho}^{**}(u(x))dx$$

- Linear relaxation (*1-sparse* solutions):

$$\boldsymbol{\rho}(z) = \begin{cases} \rho(t^i), & z = e^i,\ i \in \{1, \dots, L\}, \\ +\infty, & \text{otherwise.} \end{cases}$$

Linear relaxation

- Lifting + relaxation using the biconjugate:

$$\int_\Omega \rho(u(x))dx \rightsquigarrow \int_\Omega \boldsymbol{\rho}^{**}(\boldsymbol{u}(x))dx$$

- Linear relaxation (*1-sparse* solutions):

$$\boldsymbol{\rho}(z) = \begin{cases} \rho(t^i), & z = e^i,\ i \in \{1, \dots, L\}, \\ +\infty, & \text{otherwise.} \end{cases}$$

Linear relaxation

- Lifting + relaxation using the biconjugate:

$$\int_\Omega \rho(u(x))dx \rightsquigarrow \int_\Omega \boldsymbol{\rho}^{**}(u(x))dx$$

- Linear relaxation (*1-sparse* solutions):

$$\boldsymbol{\rho}(z) = \begin{cases} \rho(t^i), & z = e^i,\ i \in \{1, \ldots, L\}, \\ +\infty, & \text{otherwise.} \end{cases}$$

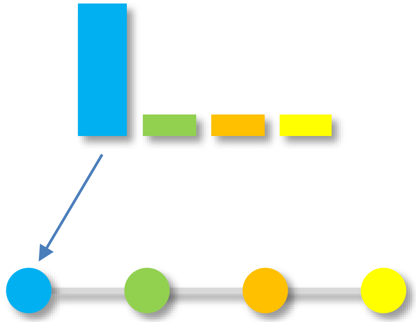

Linear relaxation

- Lifting + relaxation using the biconjugate:

$$\int_\Omega \rho(u(x))dx \rightsquigarrow \int_\Omega \boldsymbol{\rho}^{**}(u(x))dx$$

- Linear relaxation (*1-sparse* solutions):

$$\boldsymbol{\rho}(z) = \begin{cases} \rho(t^i), & z = e^i,\ i \in \{1, \ldots, L\}, \\ +\infty, & \text{otherwise.} \end{cases}$$

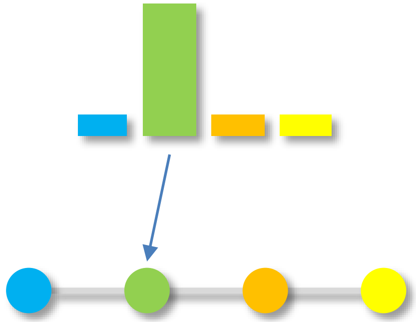

Linear relaxation

- Lifting + relaxation using the biconjugate:

$$\int_\Omega \rho(u(x))dx \rightsquigarrow \int_\Omega \boldsymbol{\rho}^{**}(u(x))dx$$

- Linear relaxation (*1-sparse* solutions):

$$\boldsymbol{\rho}(z) = \begin{cases} \rho(t^i), & z = e^i,\ i \in \{1, \ldots, L\}, \\ +\infty, & \text{otherwise.} \end{cases}$$

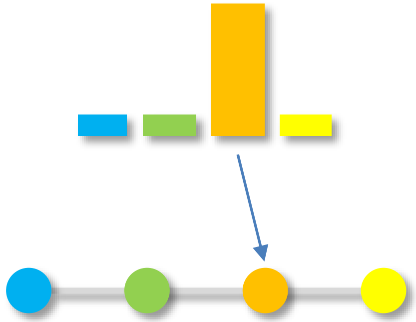

Linear relaxation

- Lifting + relaxation using the biconjugate:

$$\int_\Omega \rho(u(x))dx \rightsquigarrow \int_\Omega \boldsymbol{\rho}^{**}(u(x))dx$$

- Linear relaxation (*1-sparse* solutions):

$$\boldsymbol{\rho}(z) = \begin{cases} \rho(t^i), & z = e^i,\ i \in \{1, \ldots, L\}, \\ +\infty, & \text{otherwise.} \end{cases}$$

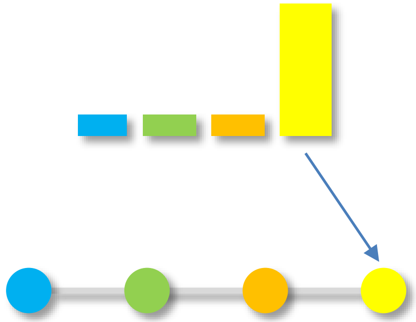

Linear relaxation

- Lifting + relaxation using the biconjugate:

$$\int_\Omega \rho(u(x))dx \rightsquigarrow \int_\Omega \boldsymbol{\rho}^{**}(u(x))dx$$

- Linear relaxation (*1-sparse* solutions):

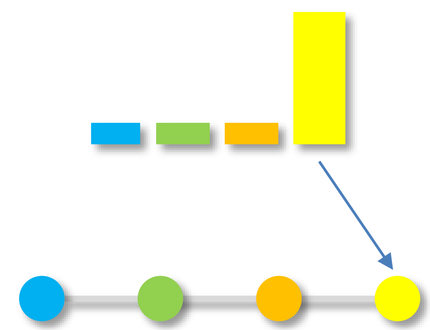

$$\boldsymbol{\rho}(z) = \begin{cases} \rho(t^i), & z = e^i,\ i \in \{1, \dots, L\}, \\ +\infty, & \text{otherwise.} \end{cases}$$

$$\min_{u \in \mathrm{BV}(\Omega, \mathcal{P}(X))} f(u) := \int_\Omega \langle u(x), s(x) \rangle dx + \int_\Omega d\Psi(Du)$$

Lifting + linear relaxation

Lifting + linear relaxation

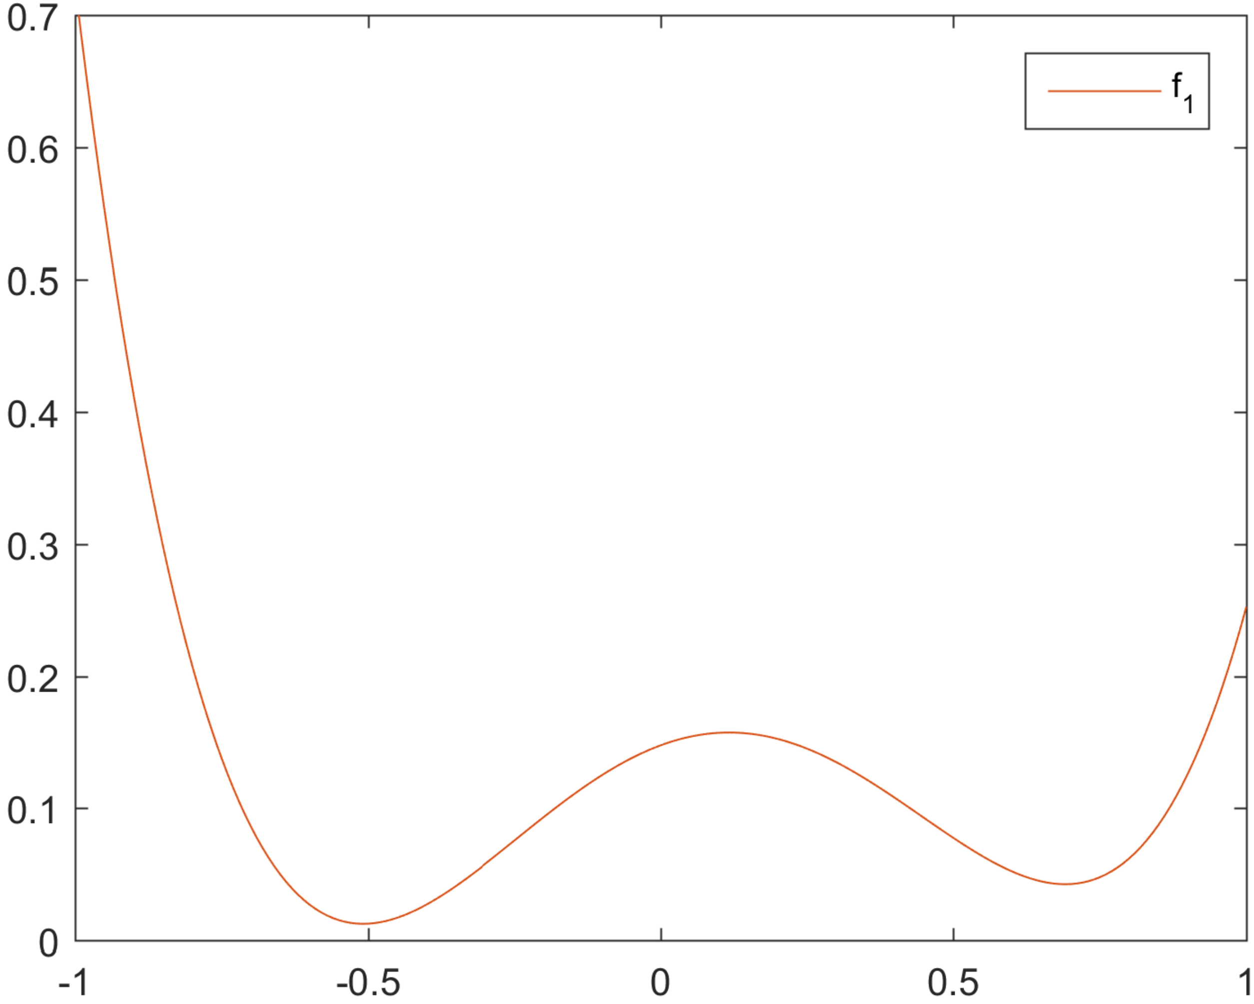

Lifting + linear relaxation

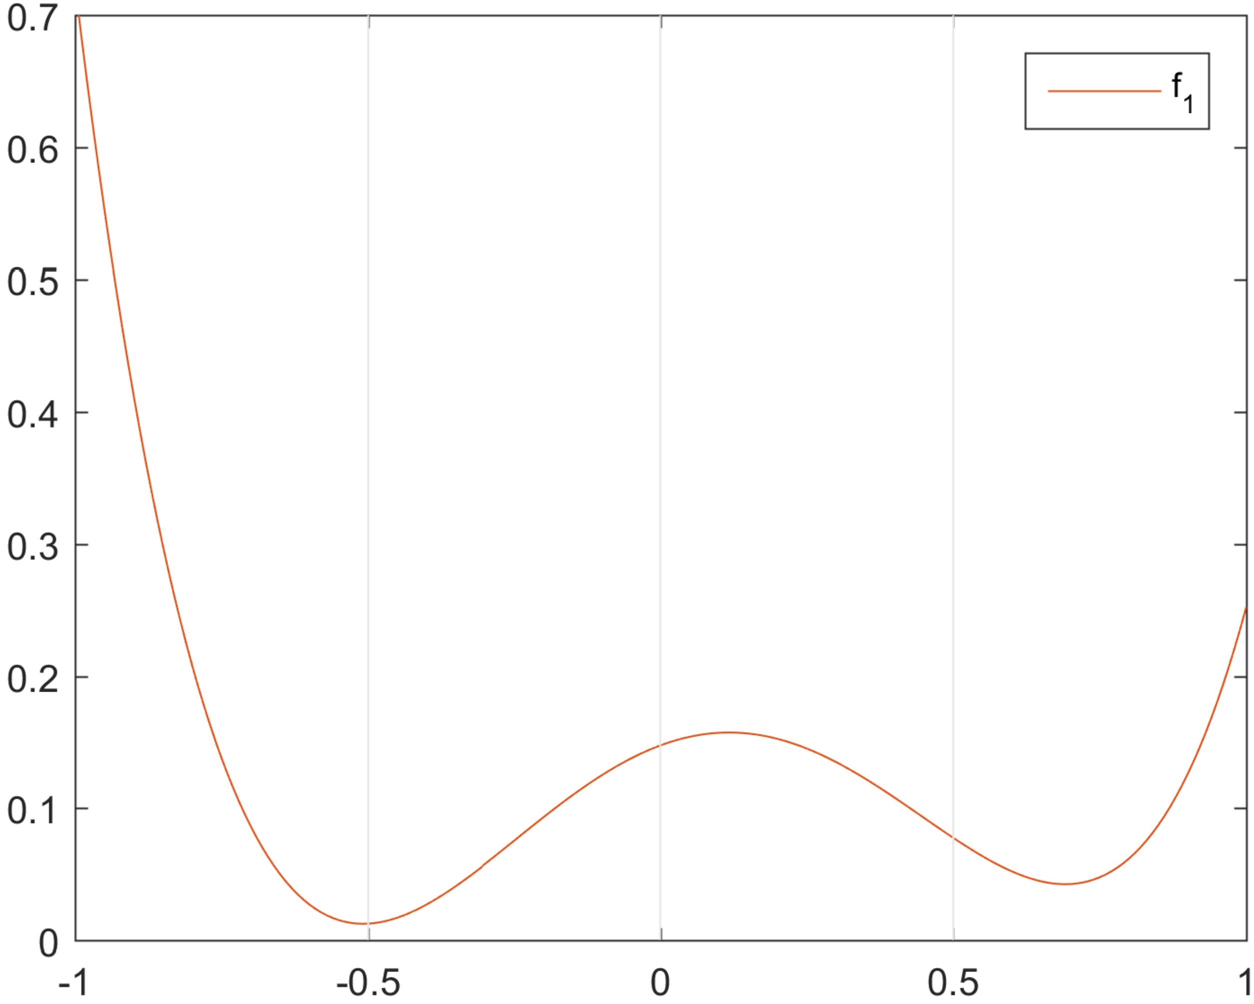

Lifting + linear relaxation

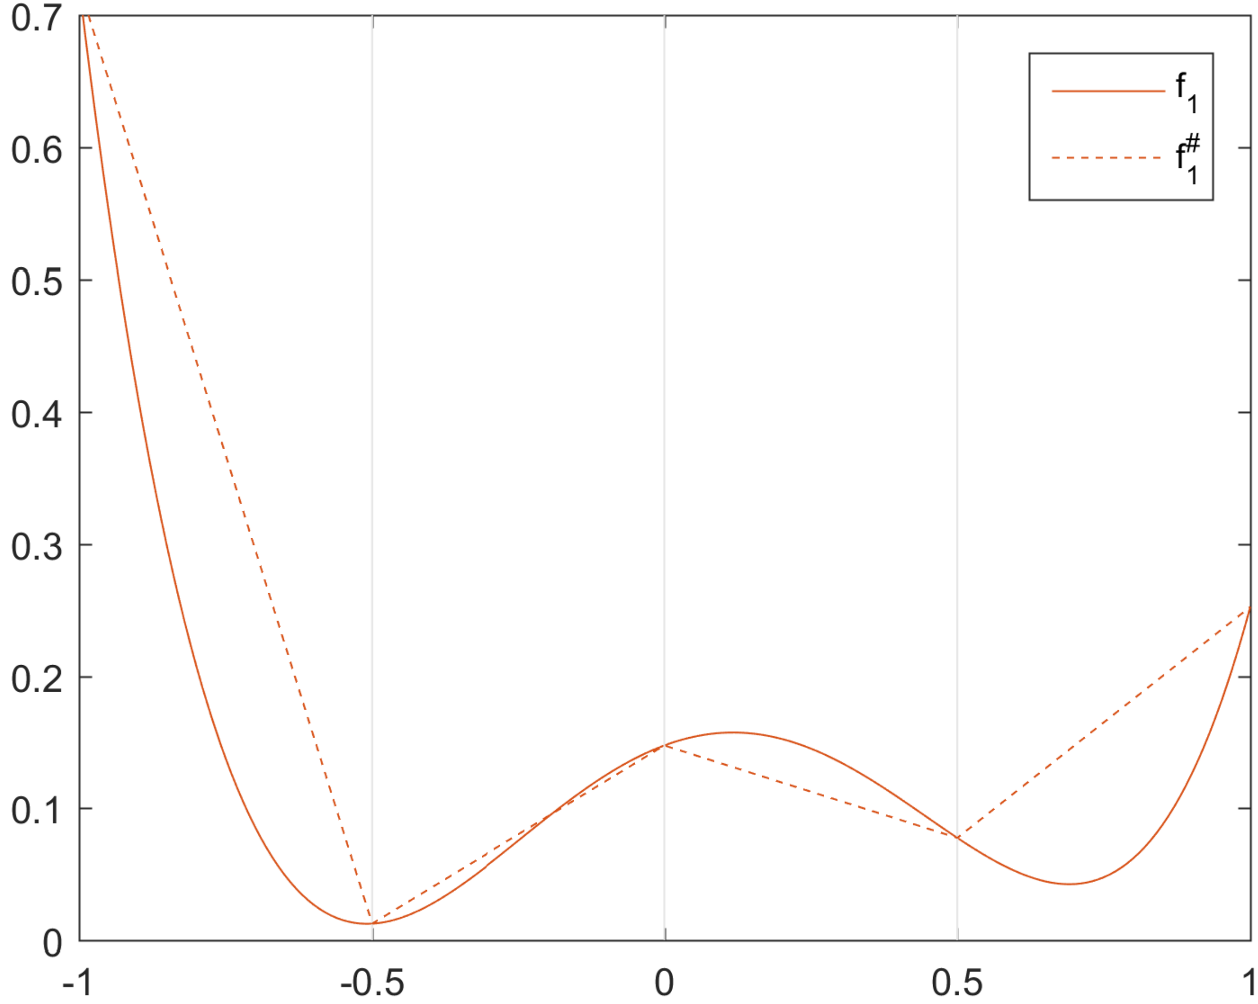

Lifting + linear relaxation

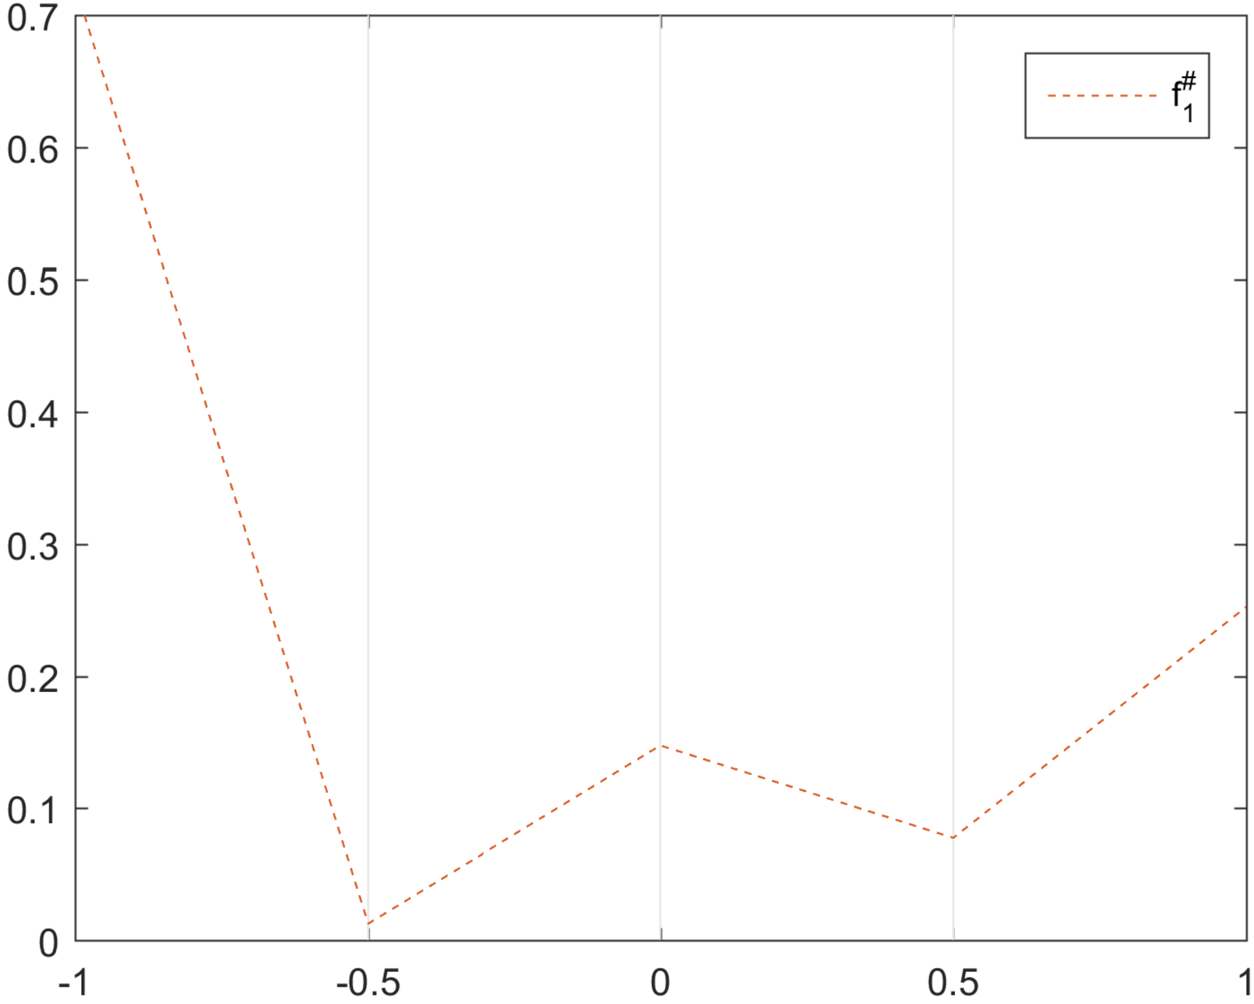

Lifting + linear relaxation

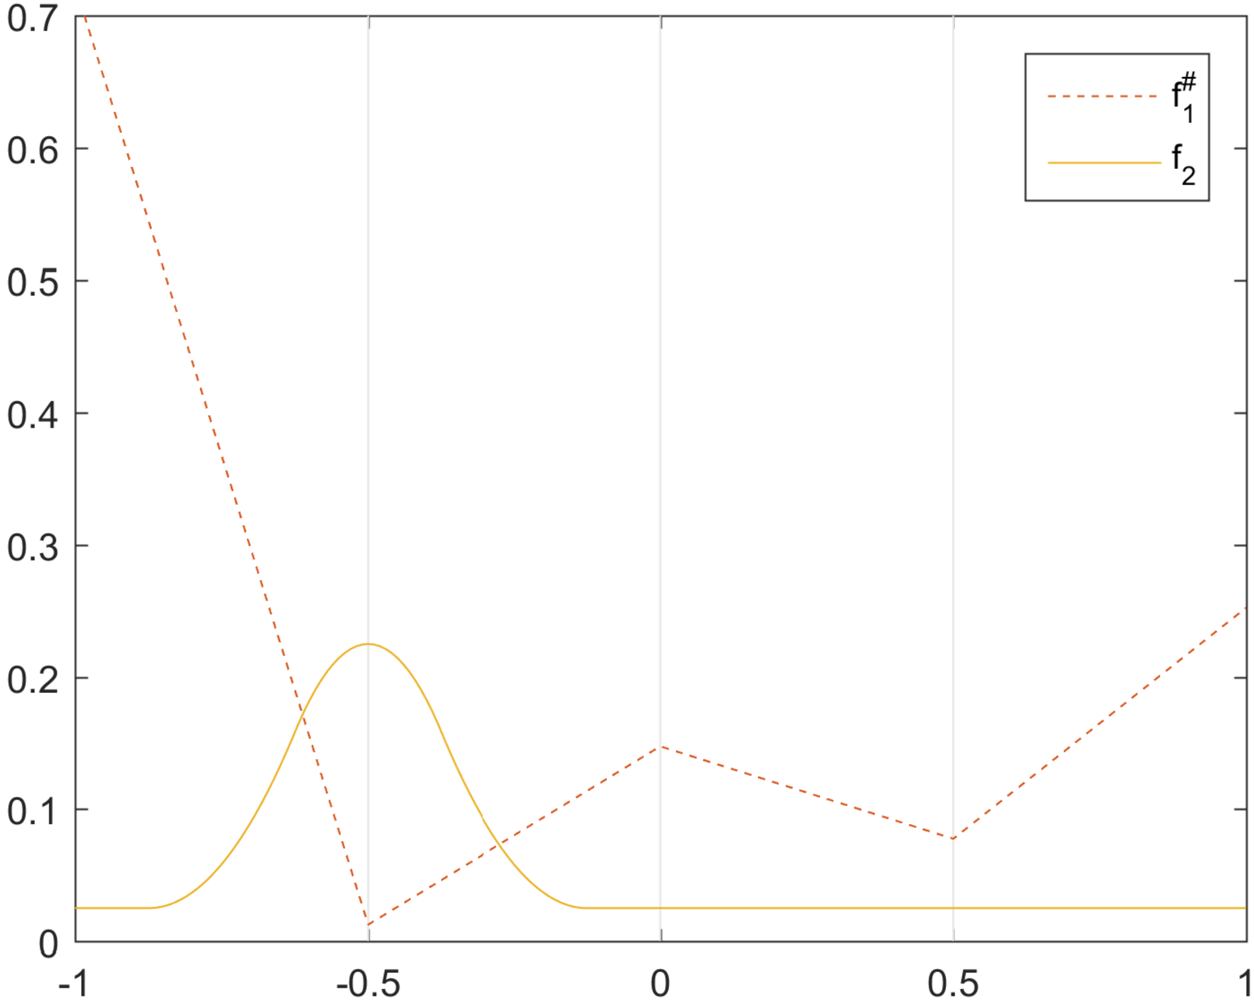

Lifting + linear relaxation

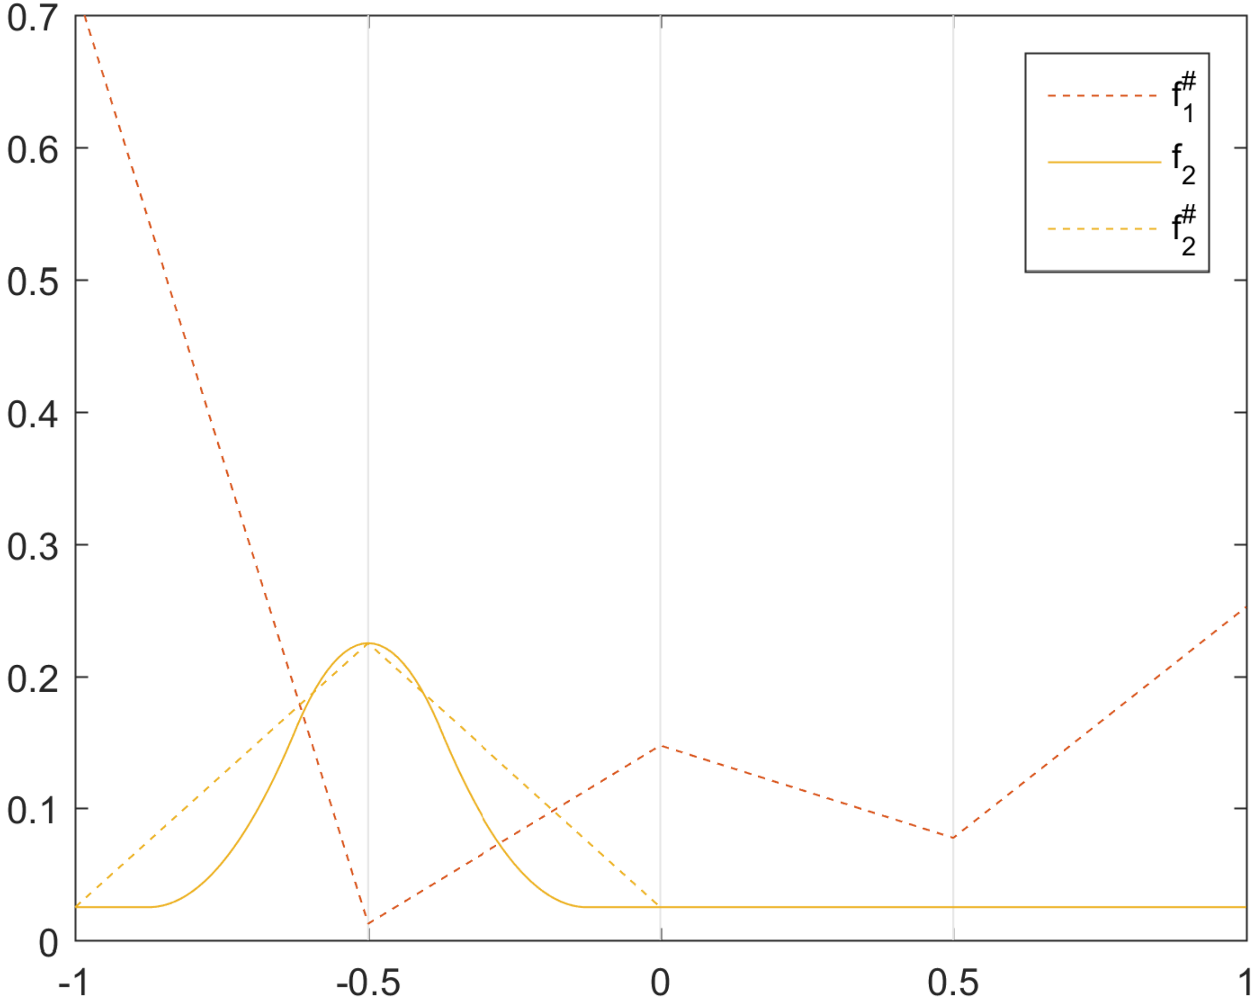

Lifting + linear relaxation

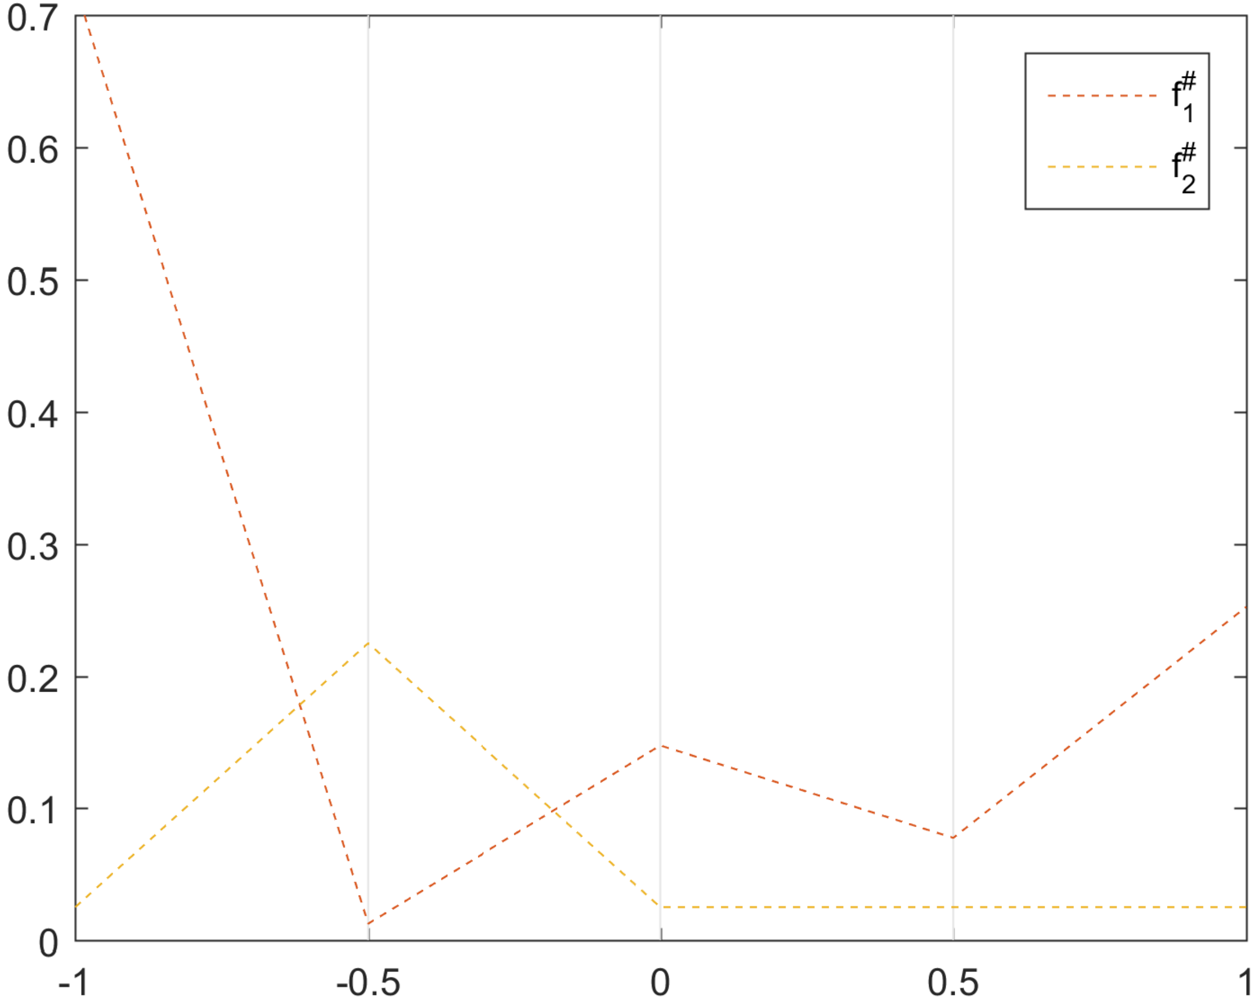

Lifting + linear relaxation

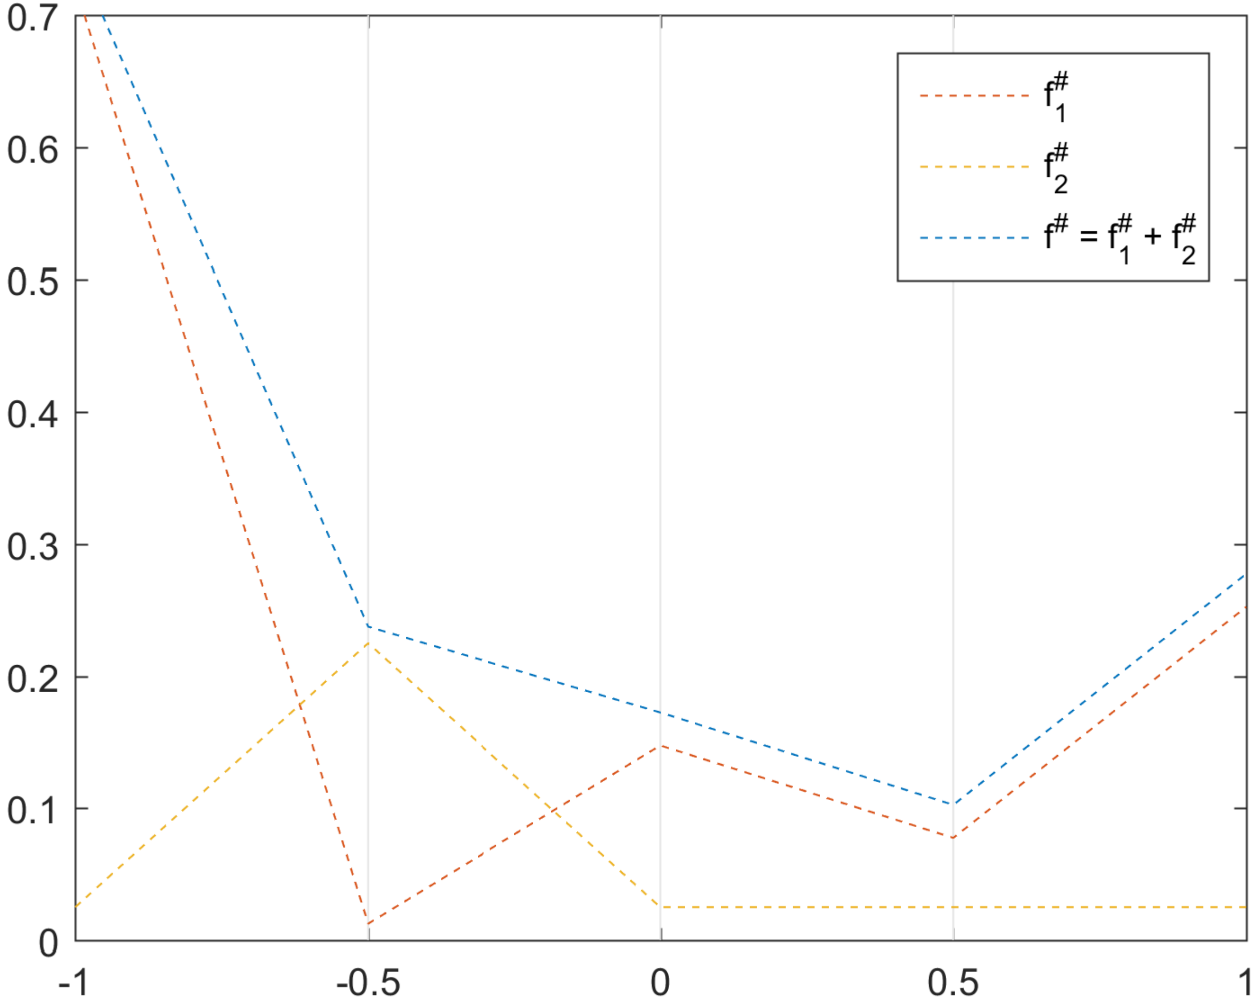

Lifting + linear relaxation

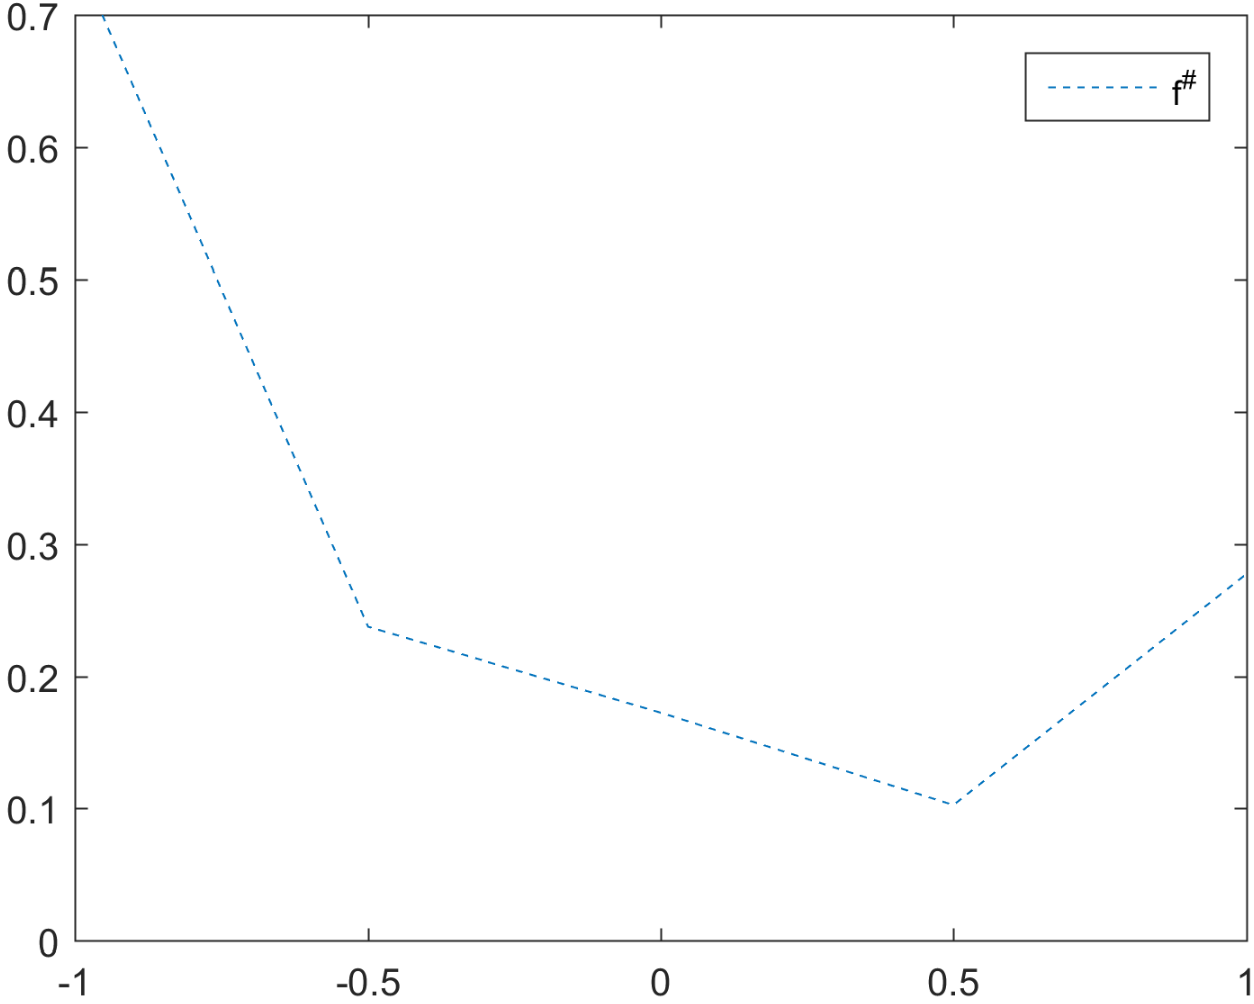

Lifting + linear relaxation

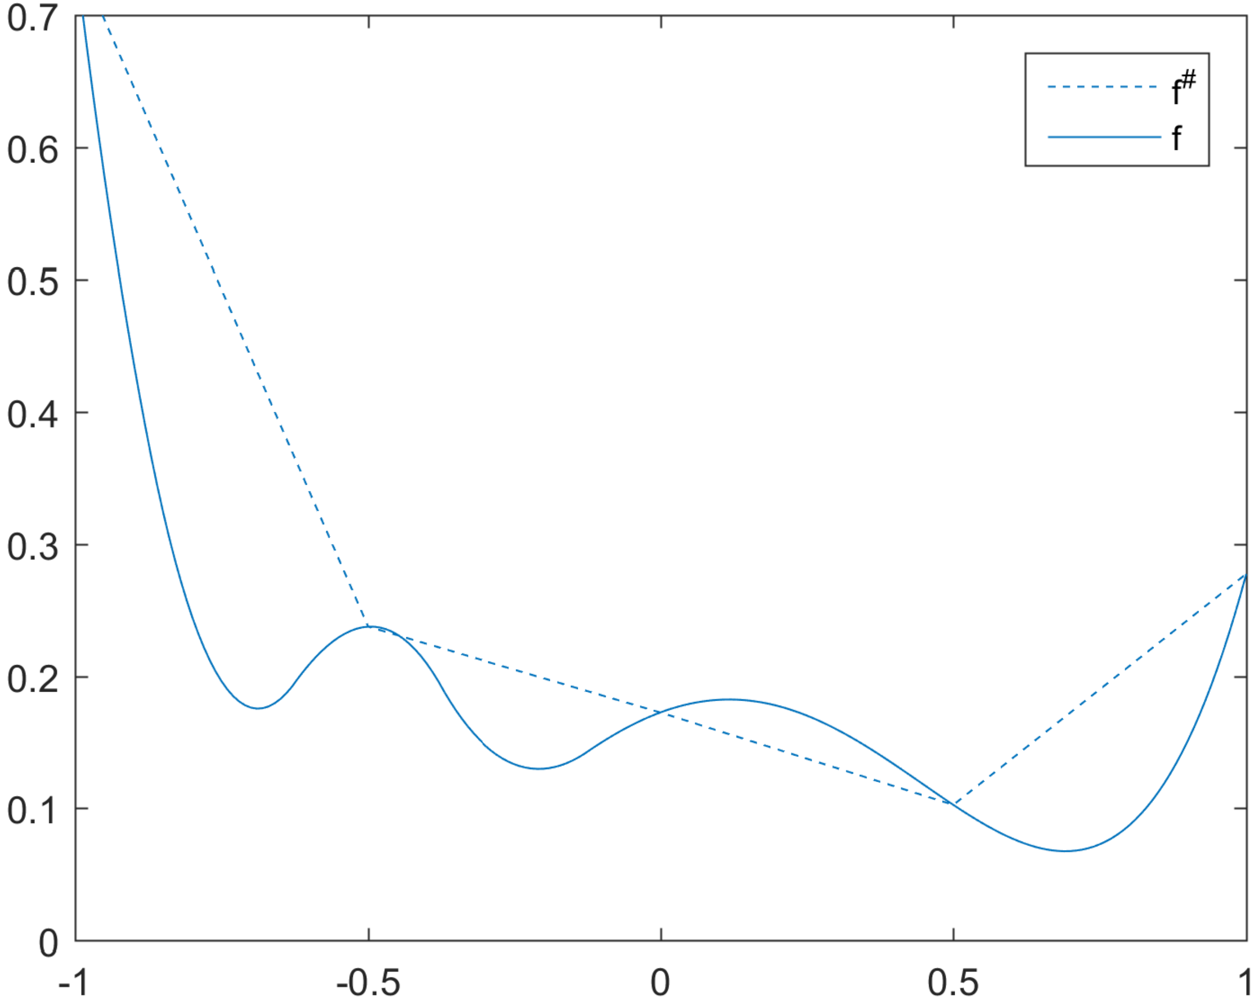

Lifting + linear relaxation

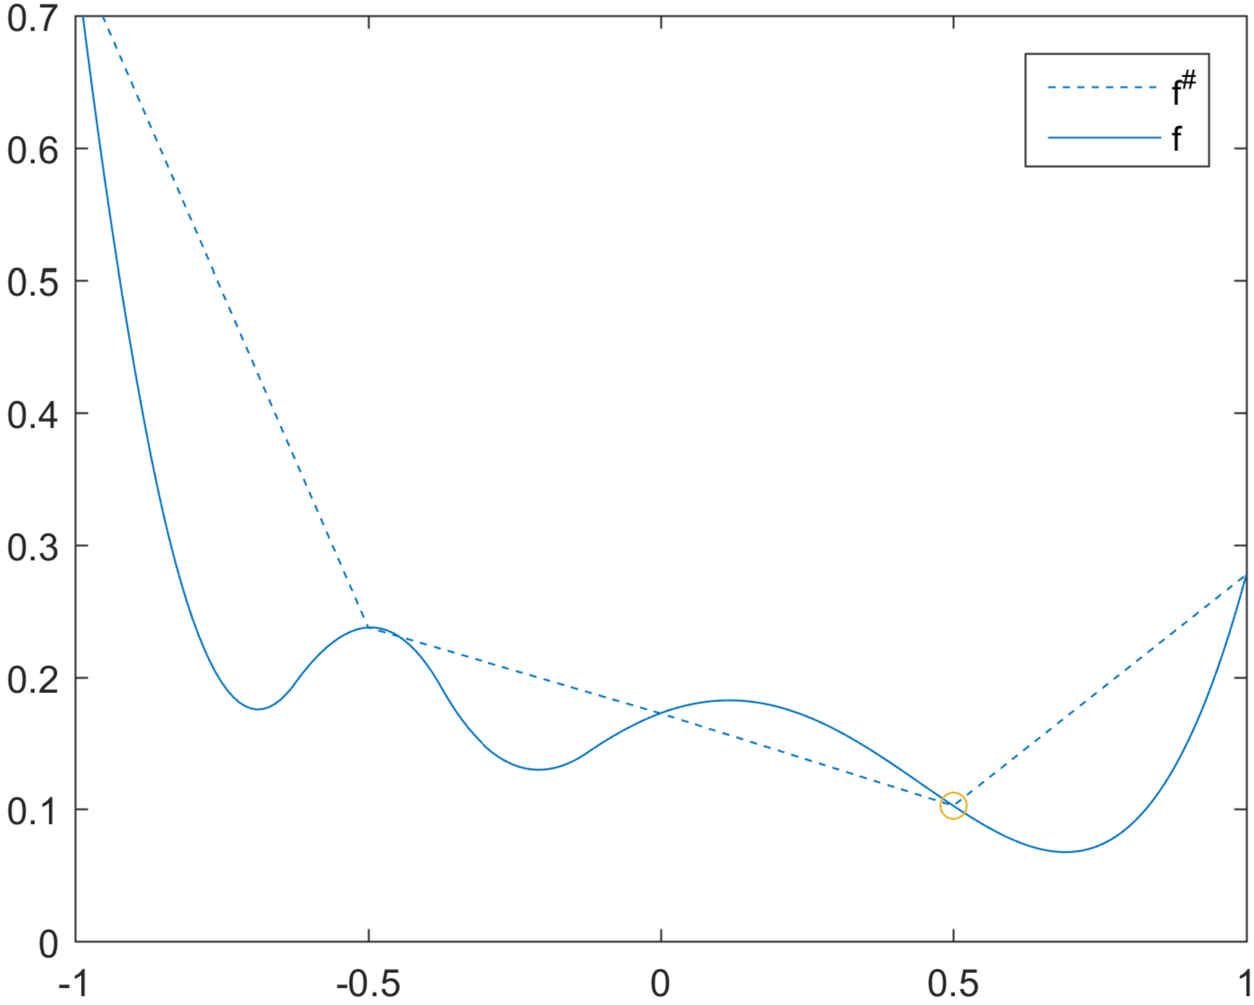

RGB-depth segmentation
(Diebold et al. , SSVM '15)

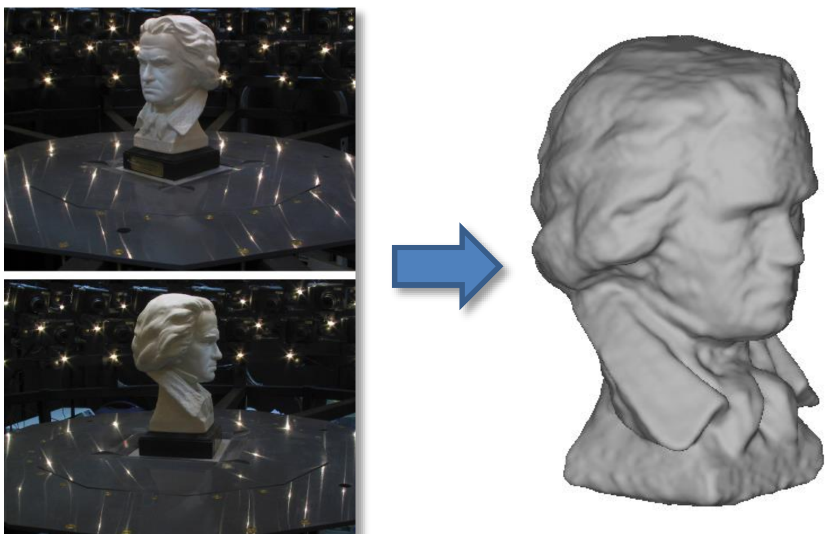

3D reconstruction
(Kolev et al., Int. J. Comp. Vis. '09)

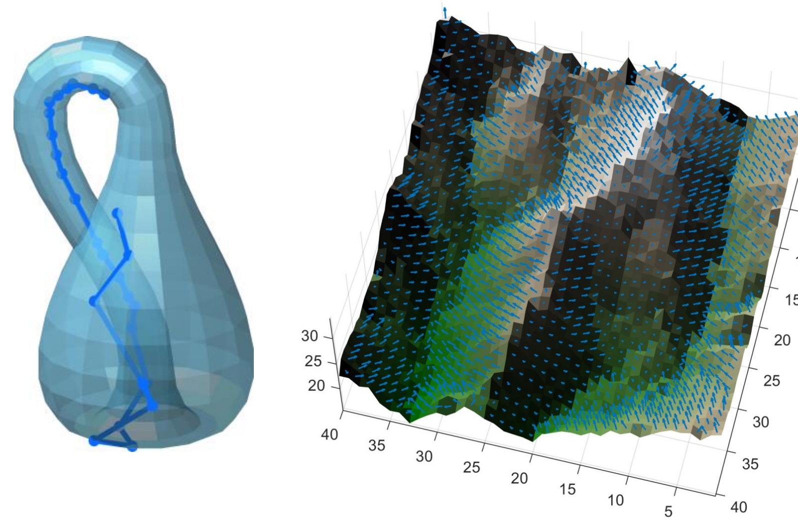

Restoring manifold-valued data
(Cremers, Strekalovskiy '12,
Lellmann et al., ICCV'13)
related: Weinmann, Demart, Storath'14;
Bergman et al.'14

Label bias

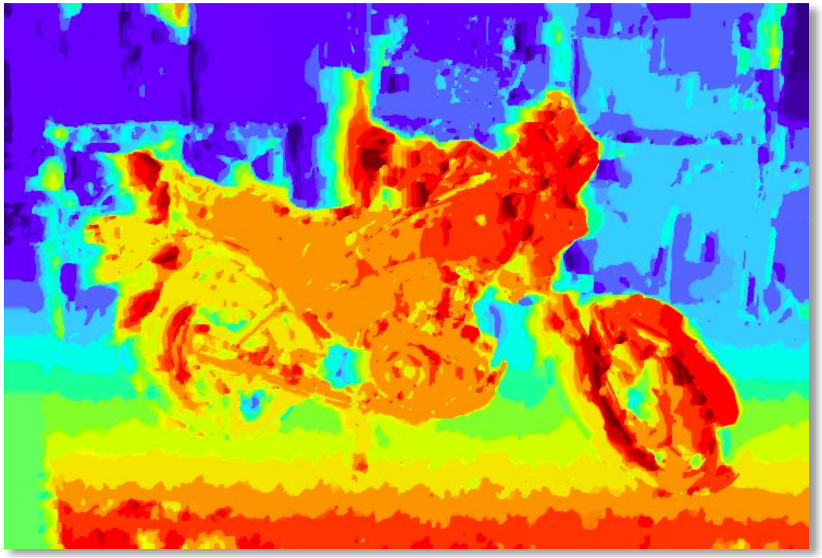

The solution tends strongly towards the chosen labels!

Lifting + linear relaxation

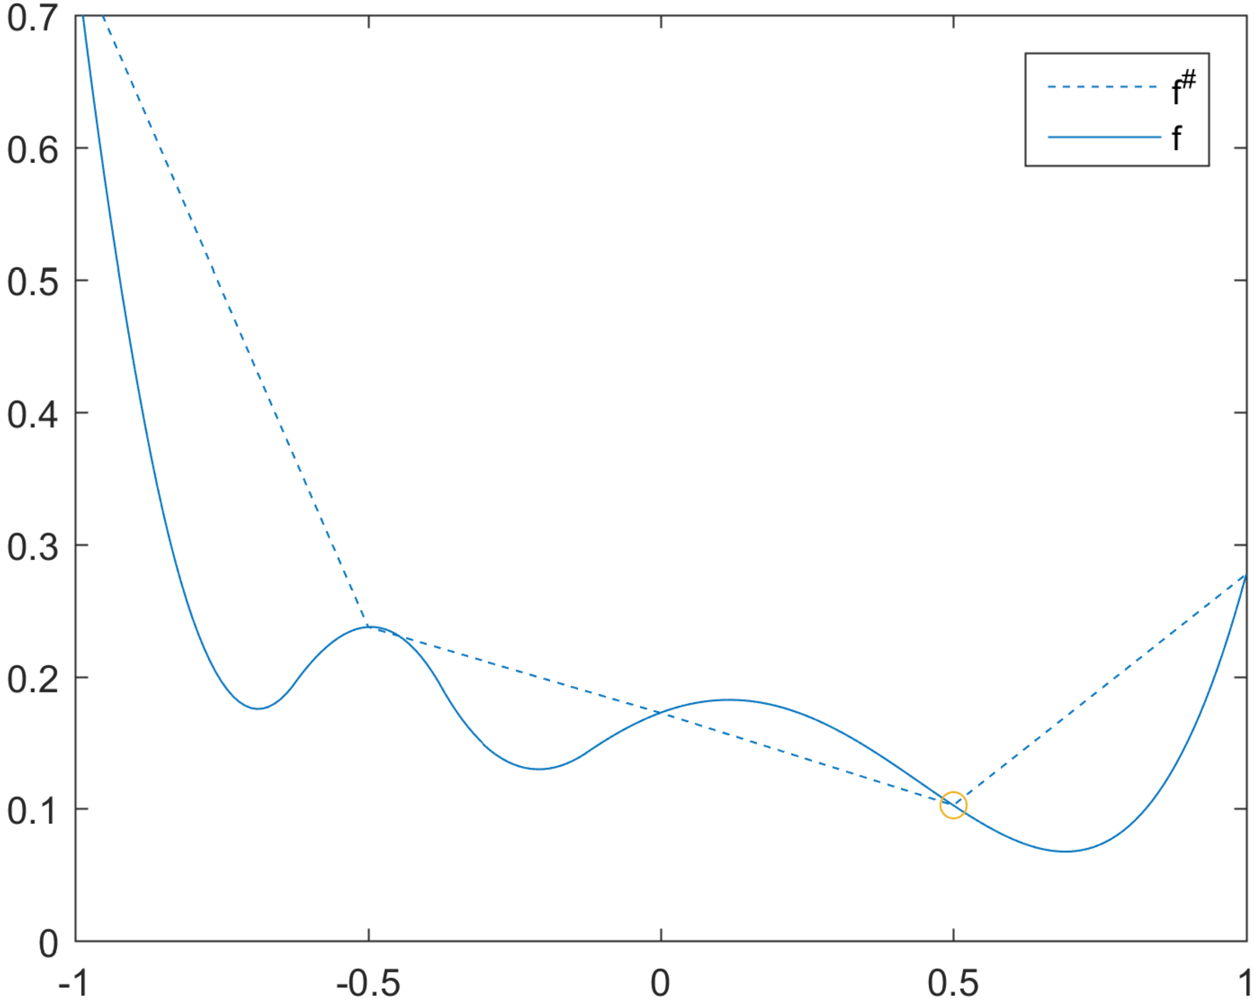

Precise relaxation

- Lifting + relaxation using the biconjugate:

$$\int_\Omega \rho(u(x))dx \rightsquigarrow \int_\Omega \boldsymbol{\rho}^{**}(u(x))dx$$

Möllenhoff, Laude, Möller, Lellmann, Cremers, CVPR'16
related: Zach, Kohli, ECCV'12; Zach, AISTATS'13

Precise relaxation

- Lifting + relaxation using the biconjugate:

$$\int_\Omega \rho(u(x))dx \rightsquigarrow \int_\Omega \boldsymbol{\rho}^{**}(u(x))dx$$

- Linear relaxation (*1-sparse* solutions):

$$\boldsymbol{\rho}(z) = \begin{cases} \rho(t^i), & z = e^i,\ i \in \{1, \dots, L\}, \\ +\infty, & \text{otherwise.} \end{cases}$$

Möllenhoff, Laude, Möller, Lellmann, Cremers, CVPR'16
related: Zach, Kohli, ECCV'12; Zach, AISTATS'13

Precise relaxation

- Lifting + relaxation using the biconjugate:

$$\int_\Omega \rho(u(x))dx \rightsquigarrow \int_\Omega \boldsymbol{\rho}^{**}(u(x))dx$$

- New: Precise relaxation (*2-sparse*!)

$$\boldsymbol{\rho}(z) = \begin{cases} \rho((1-\alpha)t^i + \alpha t^{i+1}), & z = (1-\alpha)e^i + \alpha e^{i+1}, \\ +\infty, & \text{otherwise.} \end{cases}$$

- Linear programs: solution lies on vertex.
- Here: solution between two vertices

Möllenhoff, Laude, Möller, Lellmann, Cremers, CVPR'16
related: Zach, Kohli, ECCV'12; Zach, AISTATS'13

Precise relaxation

- Lifting + relaxation using the biconjugate:

$$\int_{\Omega} \rho(u(x))dx \rightsquigarrow \int_{\Omega} \boldsymbol{\rho}^{**}(u(x))dx$$

- New: Precise relaxation (*2-sparse*!)

$$\boldsymbol{\rho}(z) = \begin{cases} \rho((1-\alpha)t^i + \alpha t^{i+1}), & z = (1-\alpha)e^i + \alpha e^{i+1}, \\ +\infty, & \text{otherwise.} \end{cases}$$

- Linear programs: solution lies on vertex.
- Here: solution between two vertices

Möllenhoff, Laude, Möller, Lellmann, Cremers, CVPR'16
related: Zach, Kohli, ECCV'12; Zach, AISTATS'13

Precise relaxation

- Lifting + relaxation using the biconjugate:

$$\int_\Omega \rho(u(x))dx \rightsquigarrow \int_\Omega \boldsymbol{\rho}^{**}(u(x))dx$$

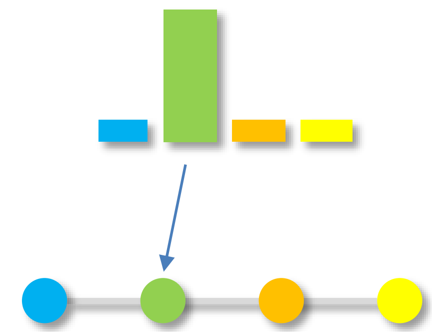

- New: Precise relaxation (*2-sparse*!)

$$\boldsymbol{\rho}(z) = \begin{cases} \rho((1-\alpha)t^i + \alpha t^{i+1}), & z = (1-\alpha)e^i + \alpha e^{i+1}, \\ +\infty, & \text{otherwise.} \end{cases}$$

- Linear programs: solution lies on vertex.
- Here: solution between two vertices

Möllenhoff, Laude, Möller, Lellmann, Cremers, CVPR’16
related: Zach, Kohli, ECCV’12; Zach, AISTATS‘13

Precise relaxation

- Lifting + relaxation using the biconjugate:

$$\int_\Omega \rho(u(x))dx \rightsquigarrow \int_\Omega \boldsymbol{\rho}^{**}(u(x))dx$$

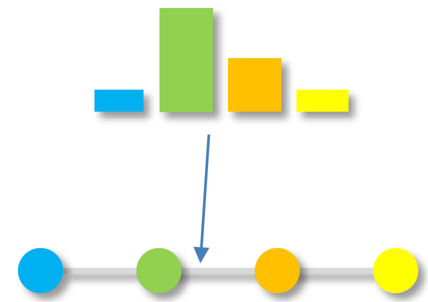

- New: Precise relaxation (*2-sparse*!)

$$\boldsymbol{\rho}(z) = \begin{cases} \rho((1-\alpha)t^i + \alpha t^{i+1}), & z = (1-\alpha)e^i + \alpha e^{i+1}, \\ +\infty, & \text{otherwise.} \end{cases}$$

- Linear programs: solution lies on vertex.
- Here: solution between two vertices

Möllenhoff, Laude, Möller, Lellmann, Cremers, CVPR'16
related: Zach, Kohli, ECCV'12; Zach, AISTATS'13

Precise relaxation

- Lifting + relaxation using the biconjugate:

$$\int_\Omega \rho(u(x))dx \rightsquigarrow \int_\Omega \boldsymbol{\rho}^{**}(u(x))dx$$

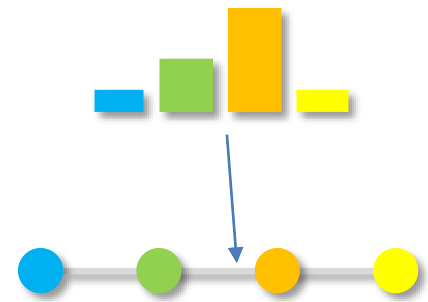

- New: Precise relaxation (*2-sparse*!)

$$\boldsymbol{\rho}(z) = \begin{cases} \rho((1-\alpha)t^i + \alpha t^{i+1}), & z = (1-\alpha)e^i + \alpha e^{i+1}, \\ +\infty, & \text{otherwise.} \end{cases}$$

- Linear programs: solution lies on vertex.
- Here: solution between two vertices

Möllenhoff, Laude, Möller, Lellmann, Cremers, CVPR'16
related: Zach, Kohli, ECCV'12; Zach, AISTATS'13

Precise relaxation

- Lifting + relaxation using the biconjugate:

$$\int_\Omega \rho(u(x))dx \rightsquigarrow \int_\Omega \boldsymbol{\rho}^{**}(u(x))dx$$

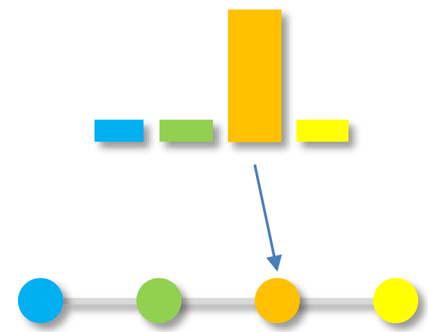

- New: Precise relaxation (*2-sparse*!)

$$\boldsymbol{\rho}(z) = \begin{cases} \rho((1-\alpha)t^i + \alpha t^{i+1}), & z = (1-\alpha)e^i + \alpha e^{i+1}, \\ +\infty, & \text{otherwise.} \end{cases}$$

- Linear programs: solution lies on vertex.
- Here: solution between two vertices

Möllenhoff, Laude, Möller, Lellmann, Cremers, CVPR'16
related: Zach, Kohli, ECCV'12; Zach, AISTATS'13

Relaxing the regularizer

Proposition 4. *The convex envelope of (15) is*

$$\Phi^{**}(\boldsymbol{g}) = \sup_{\boldsymbol{q} \in \mathcal{K}} \langle \boldsymbol{q}, \boldsymbol{g} \rangle, \tag{17}$$

where $\mathcal{K} \subset \mathbb{R}^{k \times d}$ *is given as:*

$$\mathcal{K} = \Big\{ \boldsymbol{q} \in \mathbb{R}^{k \times d} \,\Big|\, \left| \boldsymbol{q}^\top (\mathbf{1}_i^\alpha - \mathbf{1}_j^\beta) \right|_2 \leq \left| \gamma_i^\alpha - \gamma_j^\beta \right|, \quad \forall\, 1 \leq i \leq j \leq k,\ \forall \alpha, \beta \in [0,1] \Big\}. \tag{18}$$

Proposition 5. *In case the labels are ordered, i.e.,* $\gamma_1 < \gamma_2 < \ldots < \gamma_L$, *then the constraint set* $\mathcal{K}$ *from Eq. (36) is equal to*

$$\mathcal{K} = \{ \boldsymbol{q} \in \mathbb{R}^{k \times d} \mid |\boldsymbol{q}_i|_2 \leq \gamma_{i+1} - \gamma_i,\ \forall i \}. \tag{19}$$

New: lifting + precise relaxation

New: lifting + precise relaxation

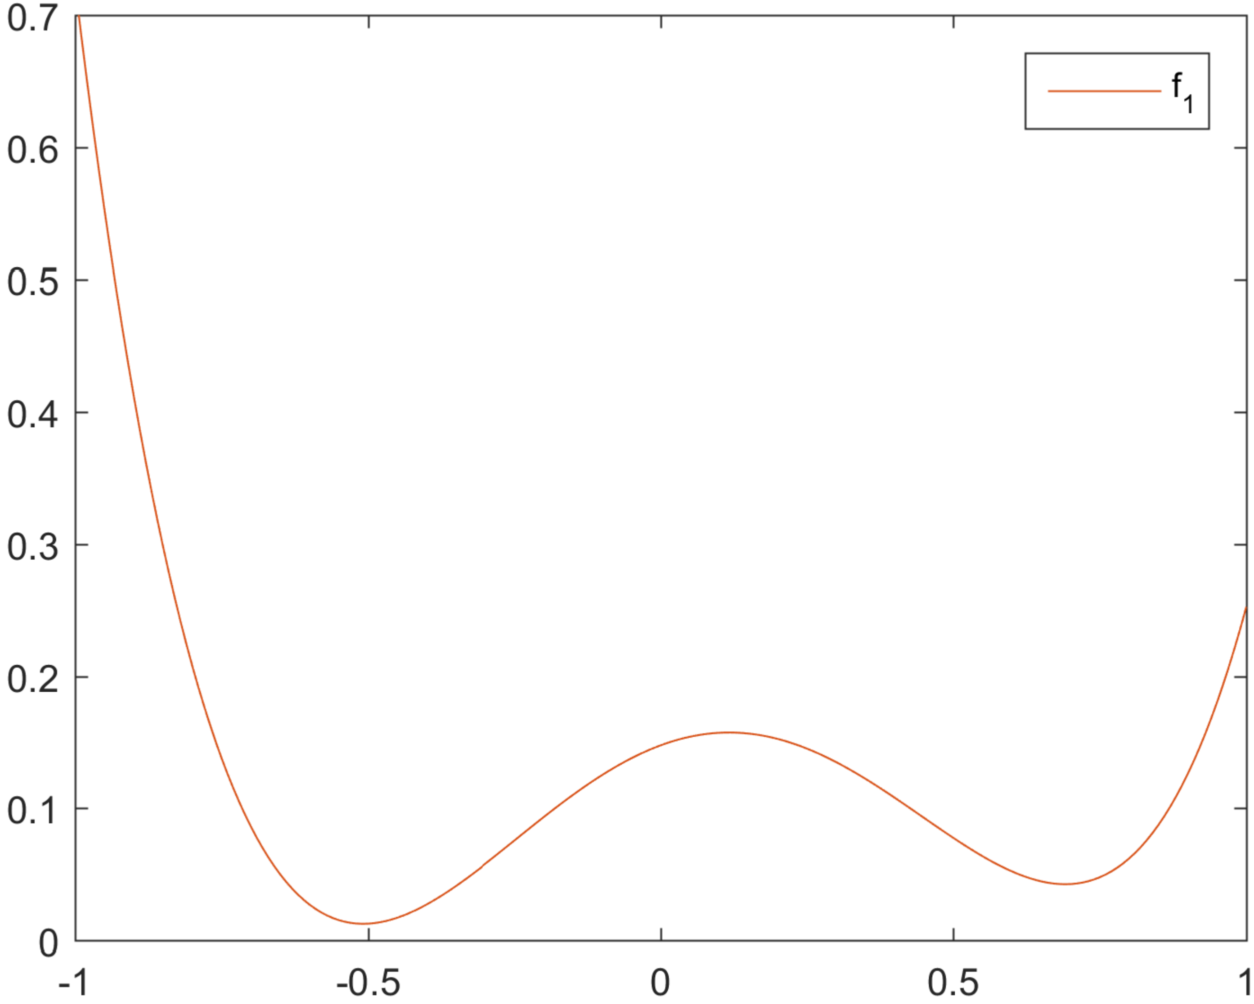

New: lifting + precise relaxation

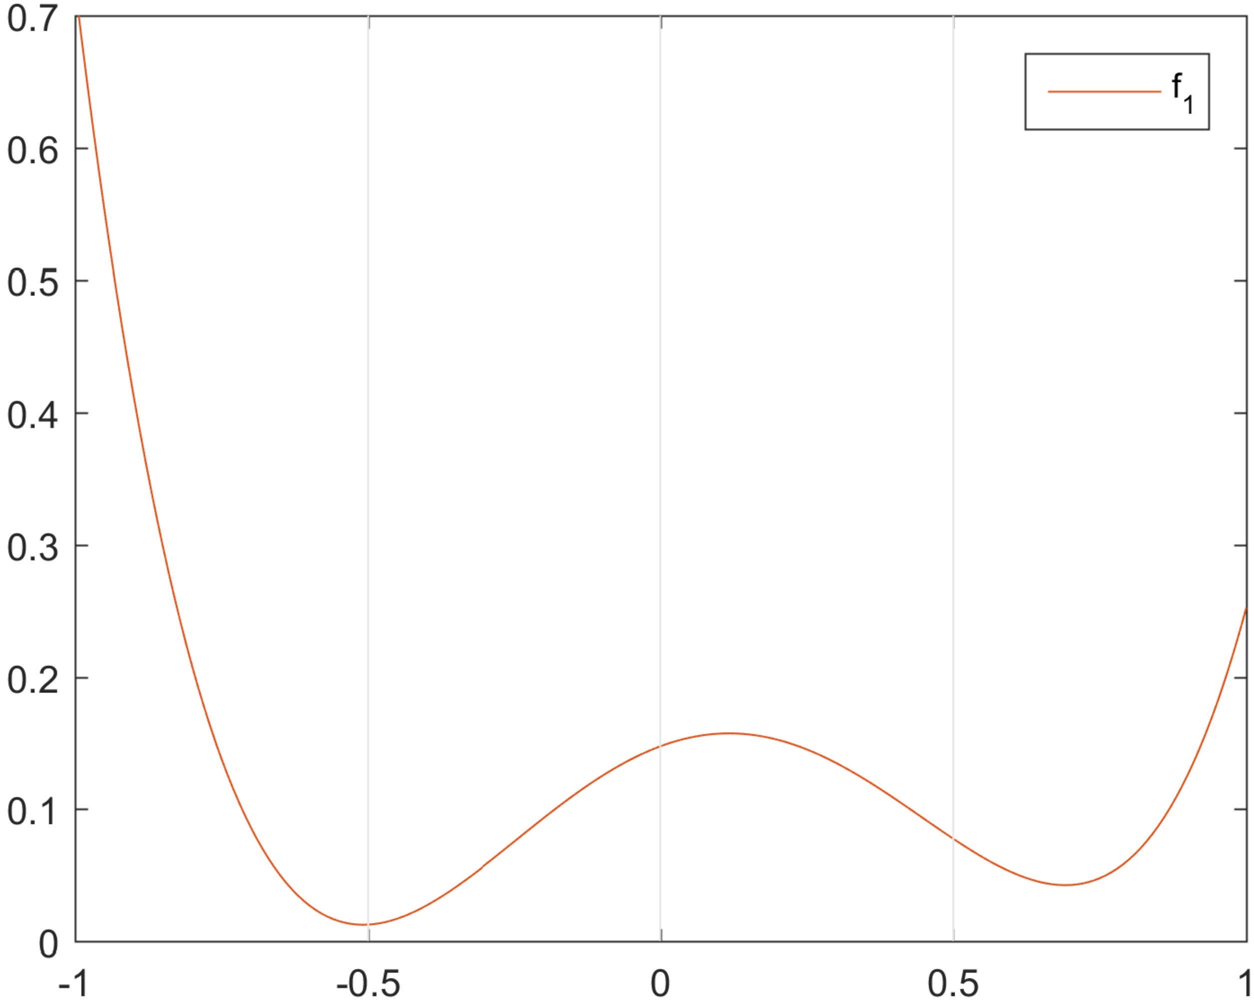

New: lifting + precise relaxation

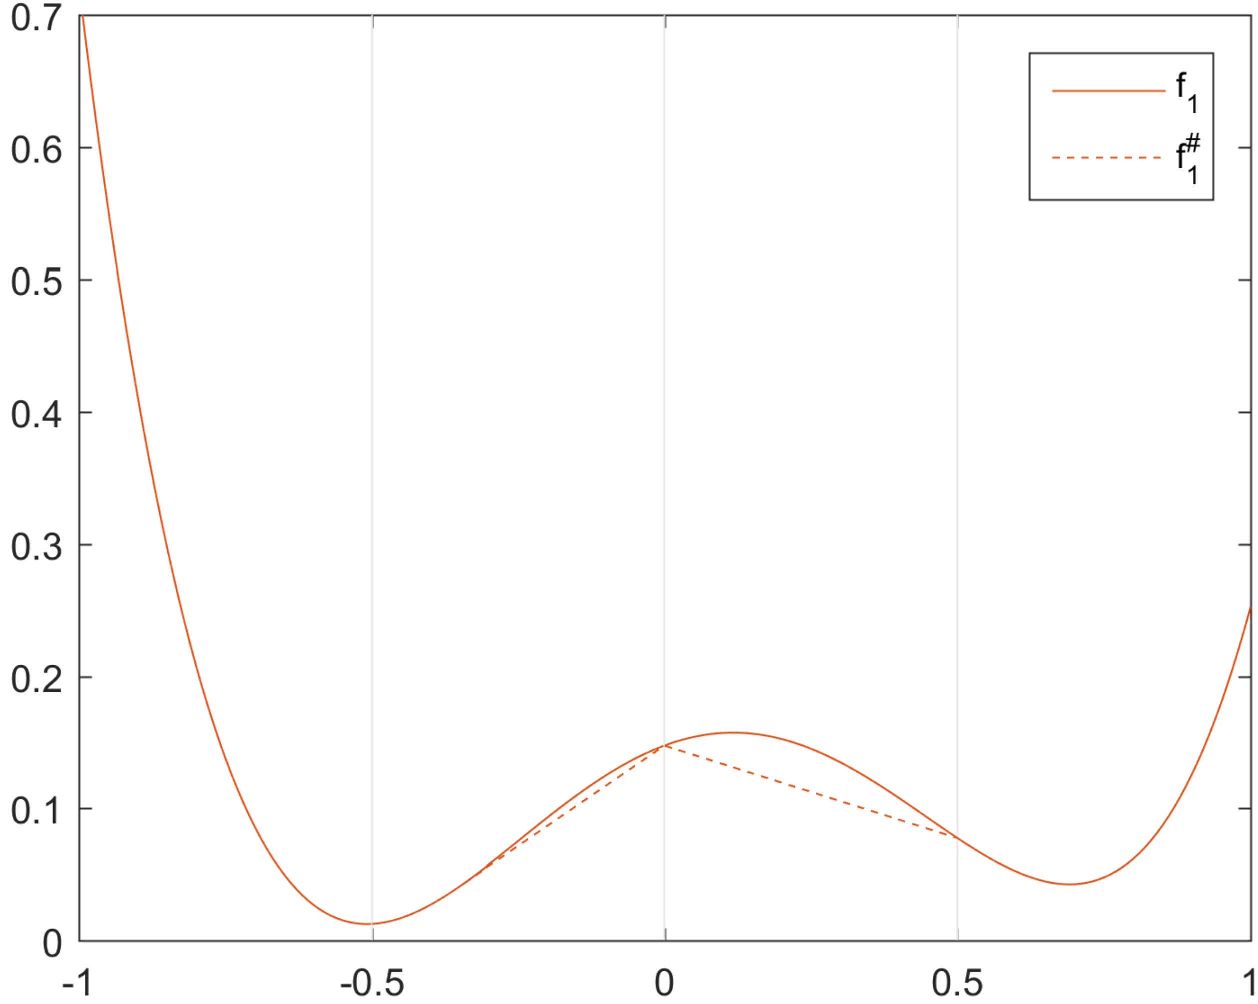

New: lifting + precise relaxation

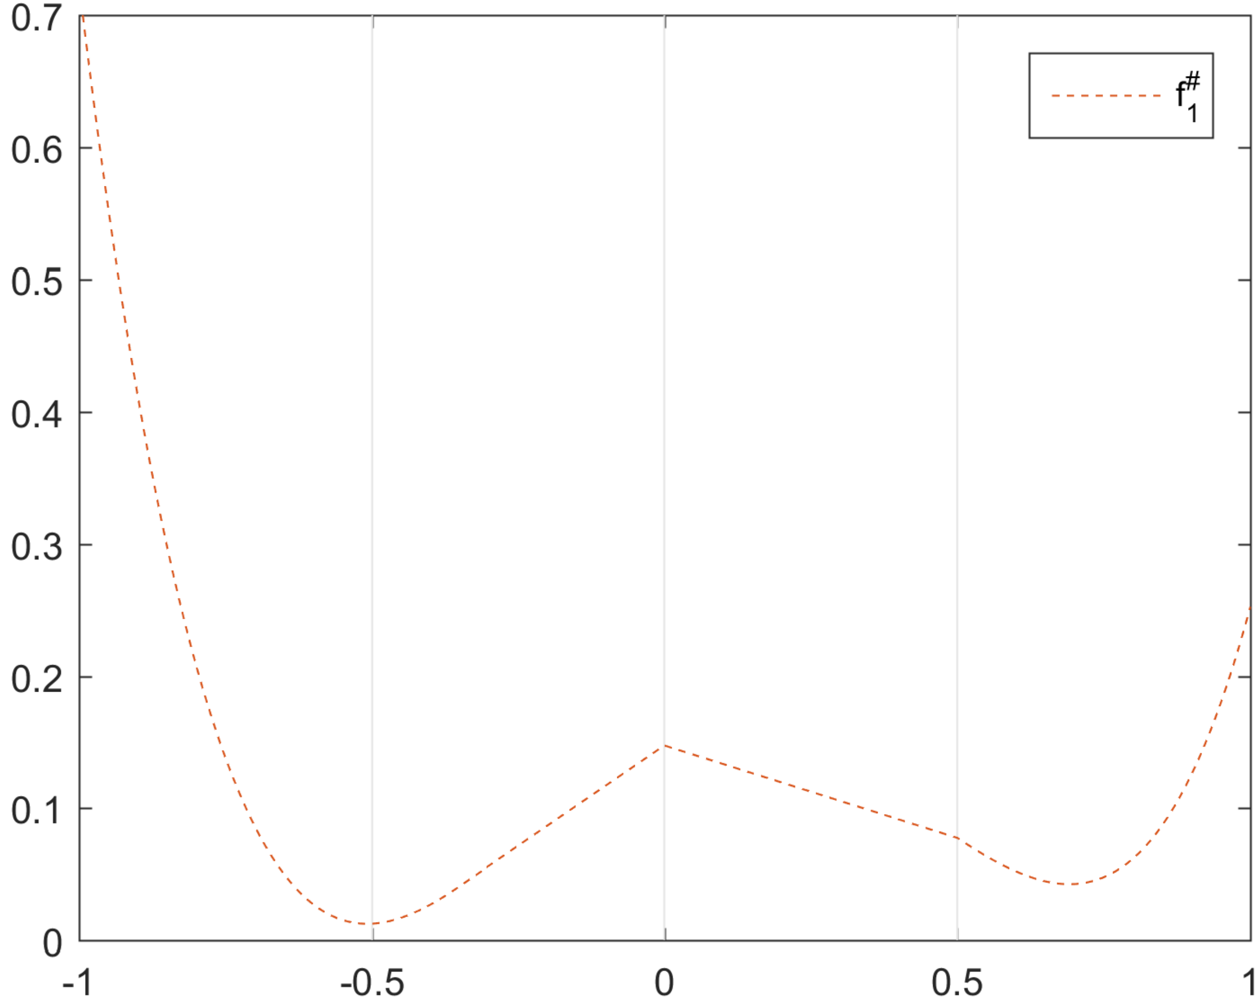

New: lifting + precise relaxation

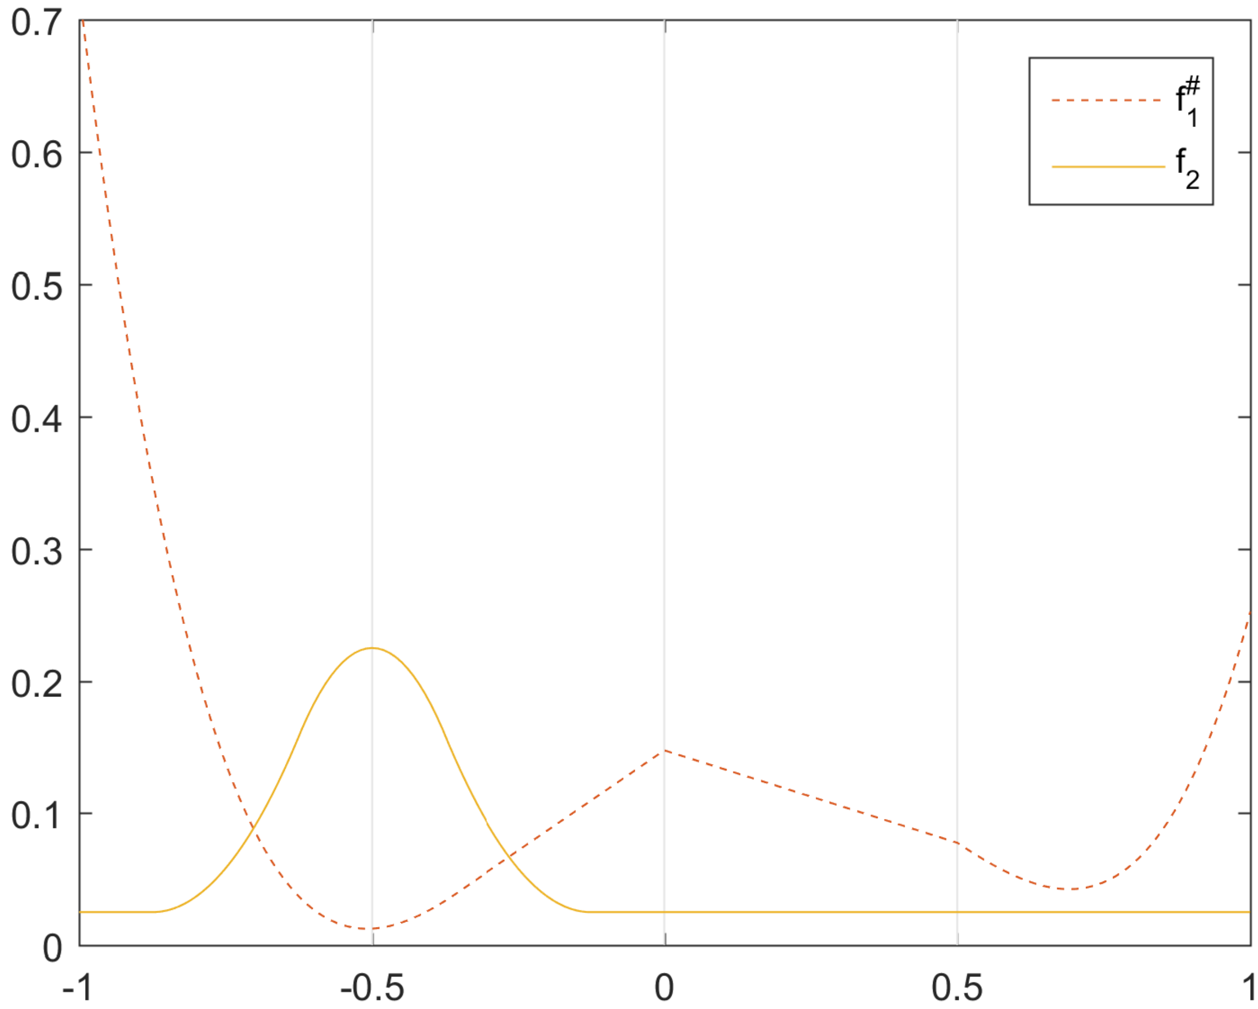

New: lifting + precise relaxation

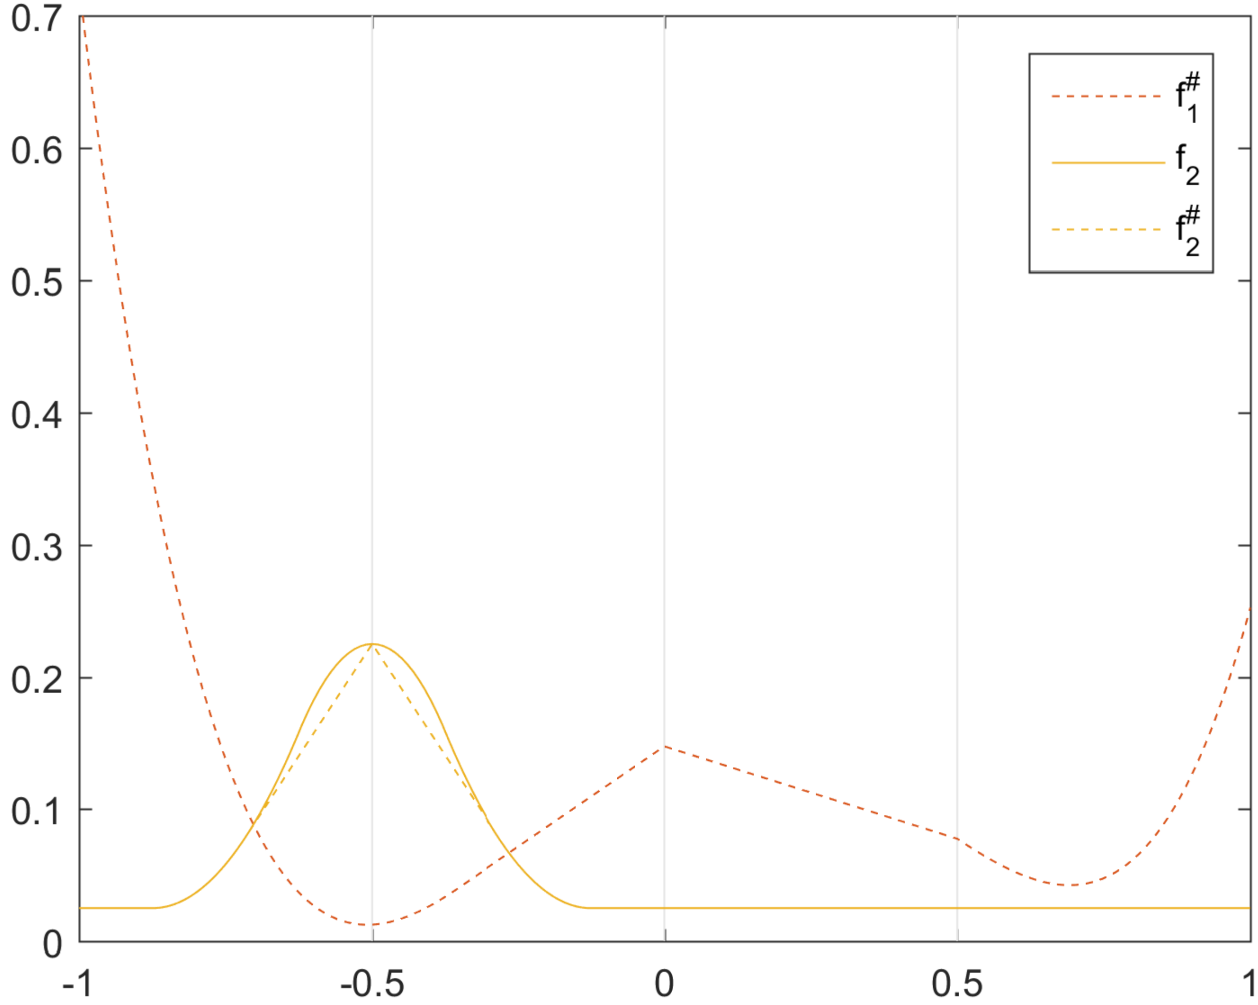

New: lifting + precise relaxation

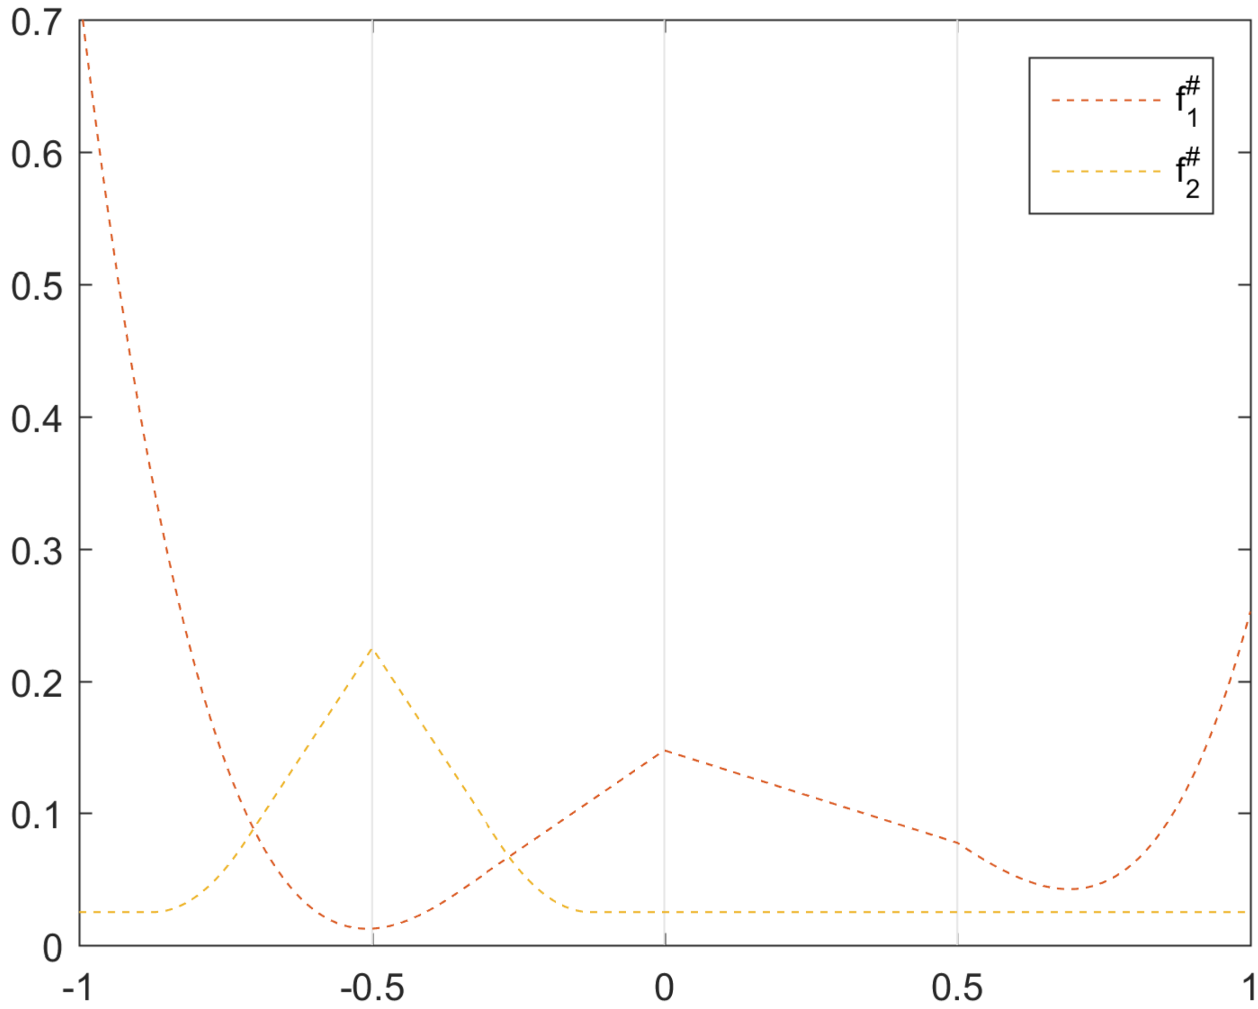

New: lifting + precise relaxation

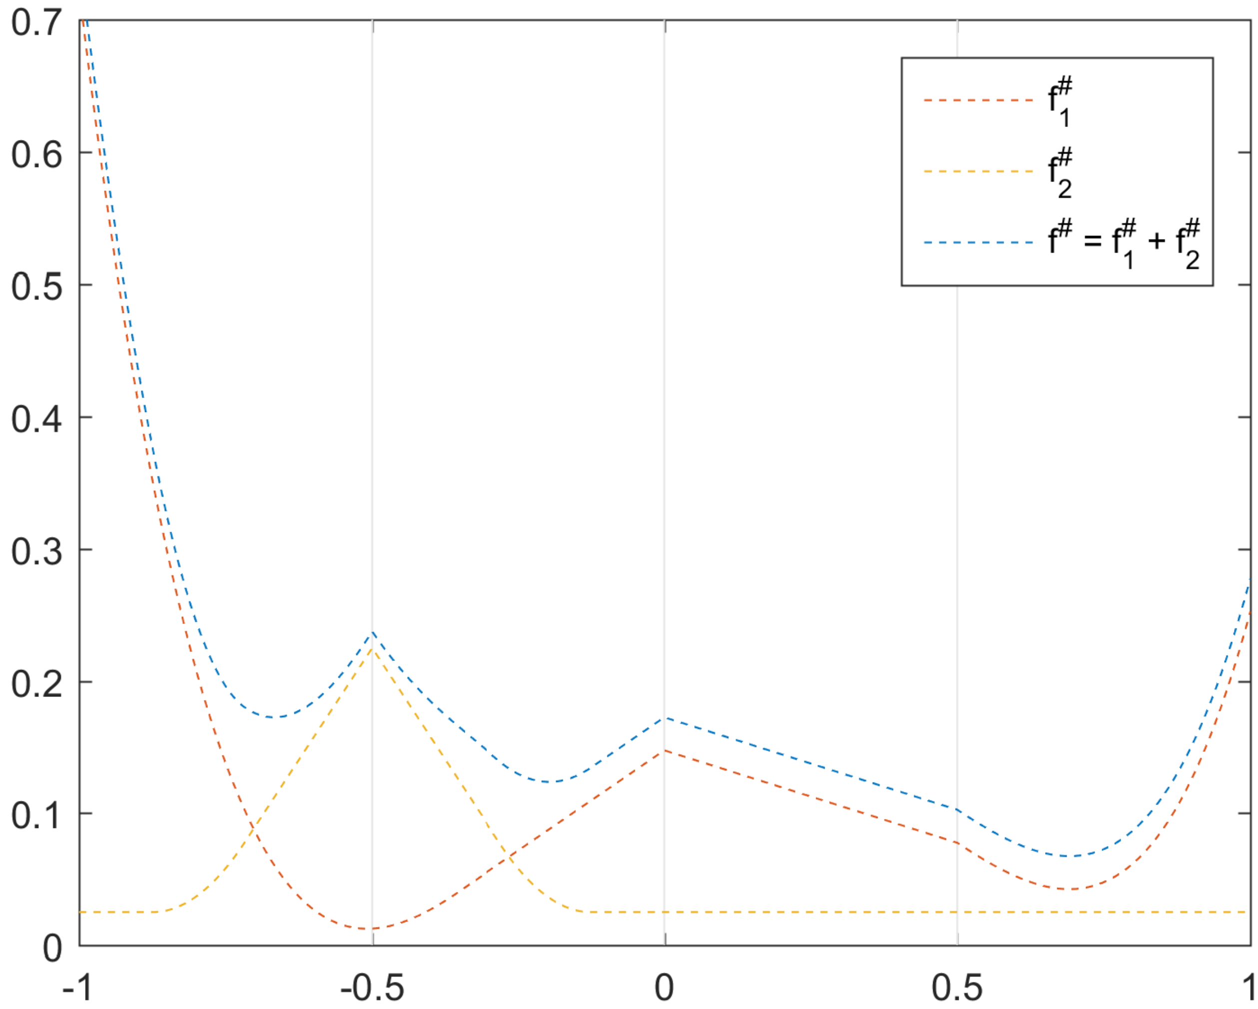

New: lifting + precise relaxation

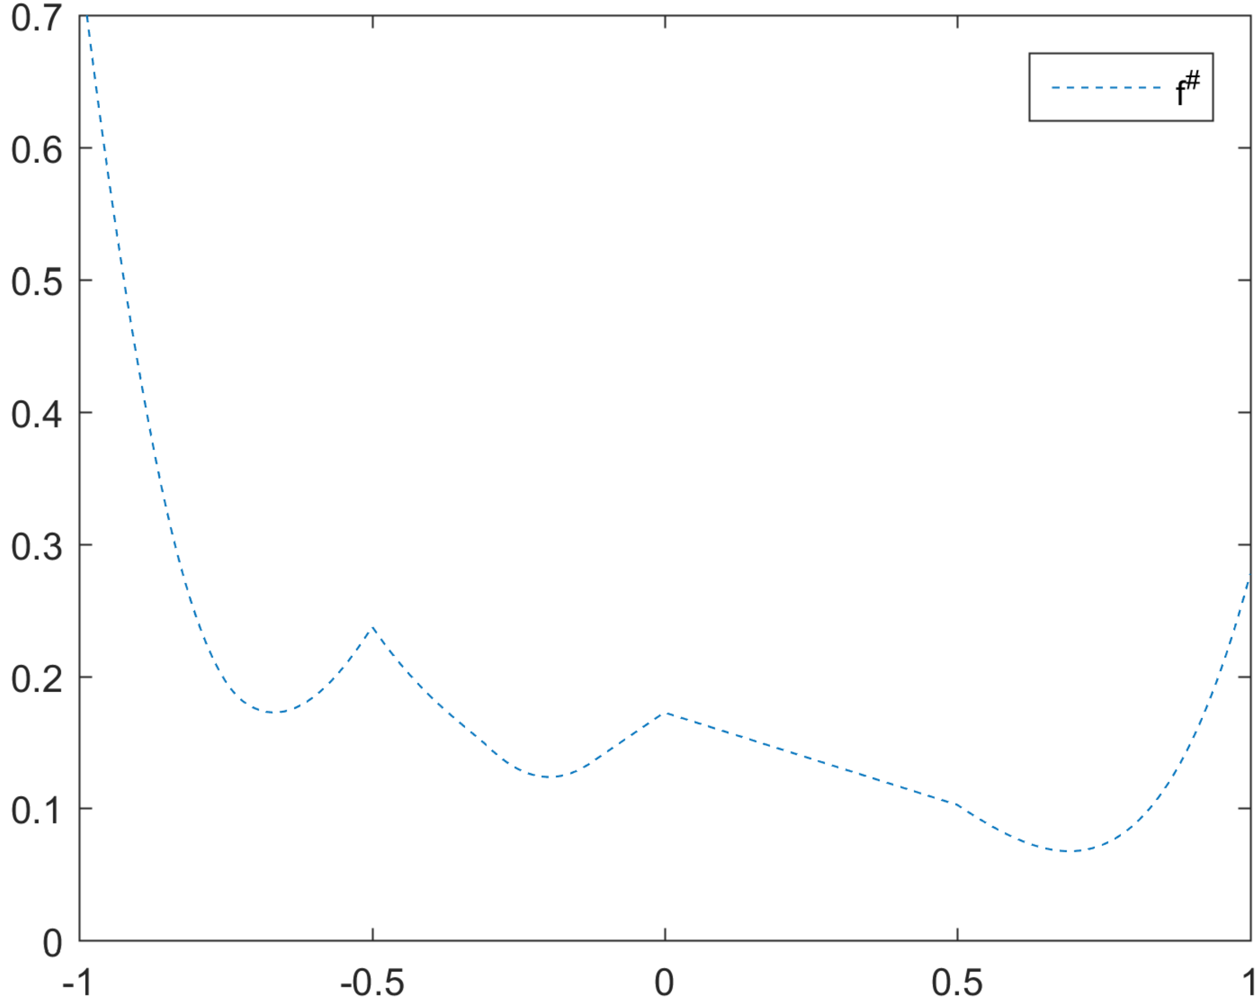

New: lifting + precise relaxation

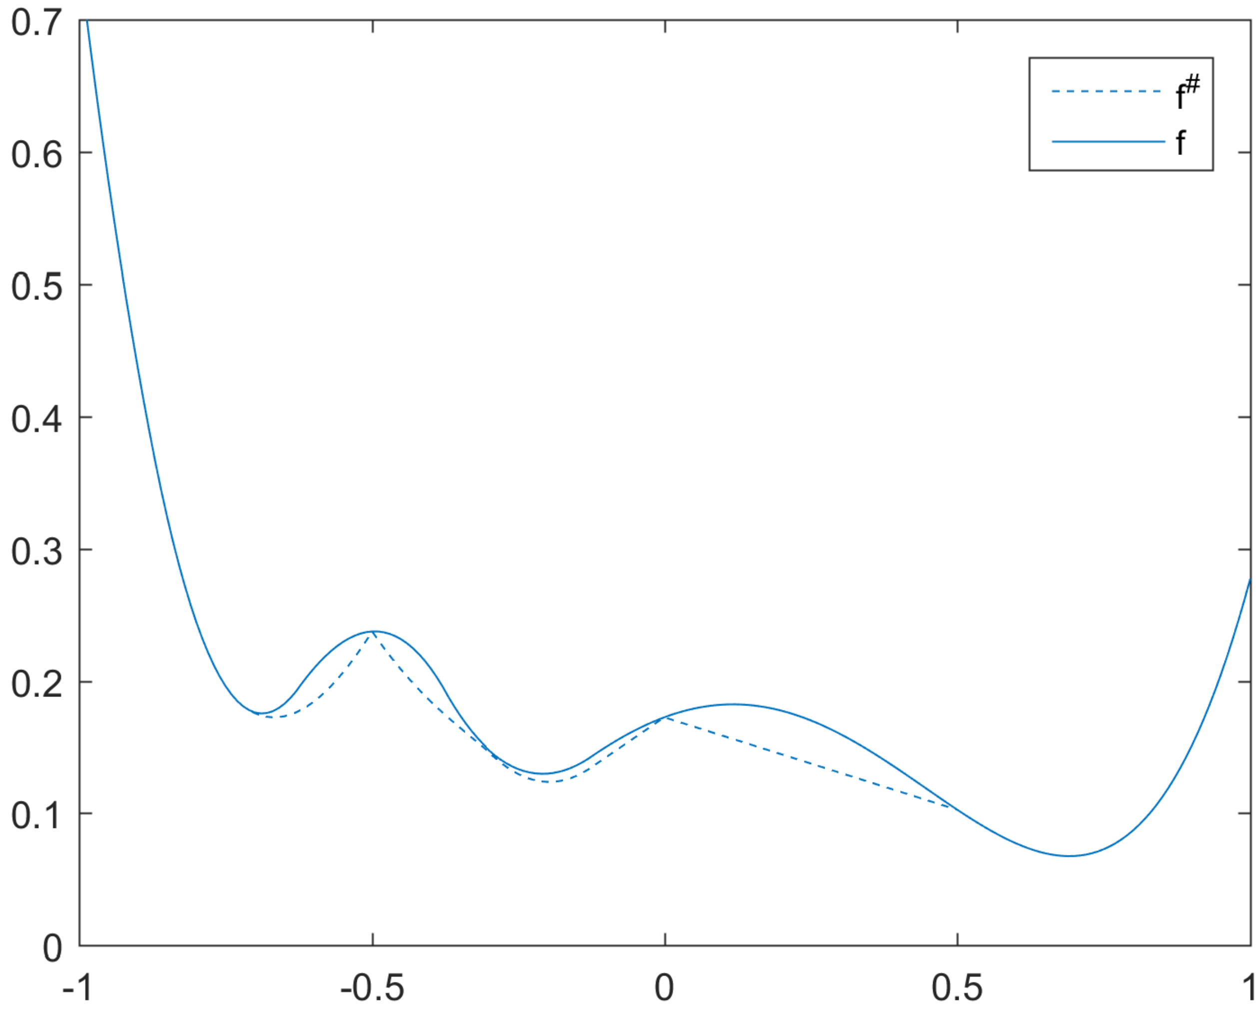

New: lifting + precise relaxation

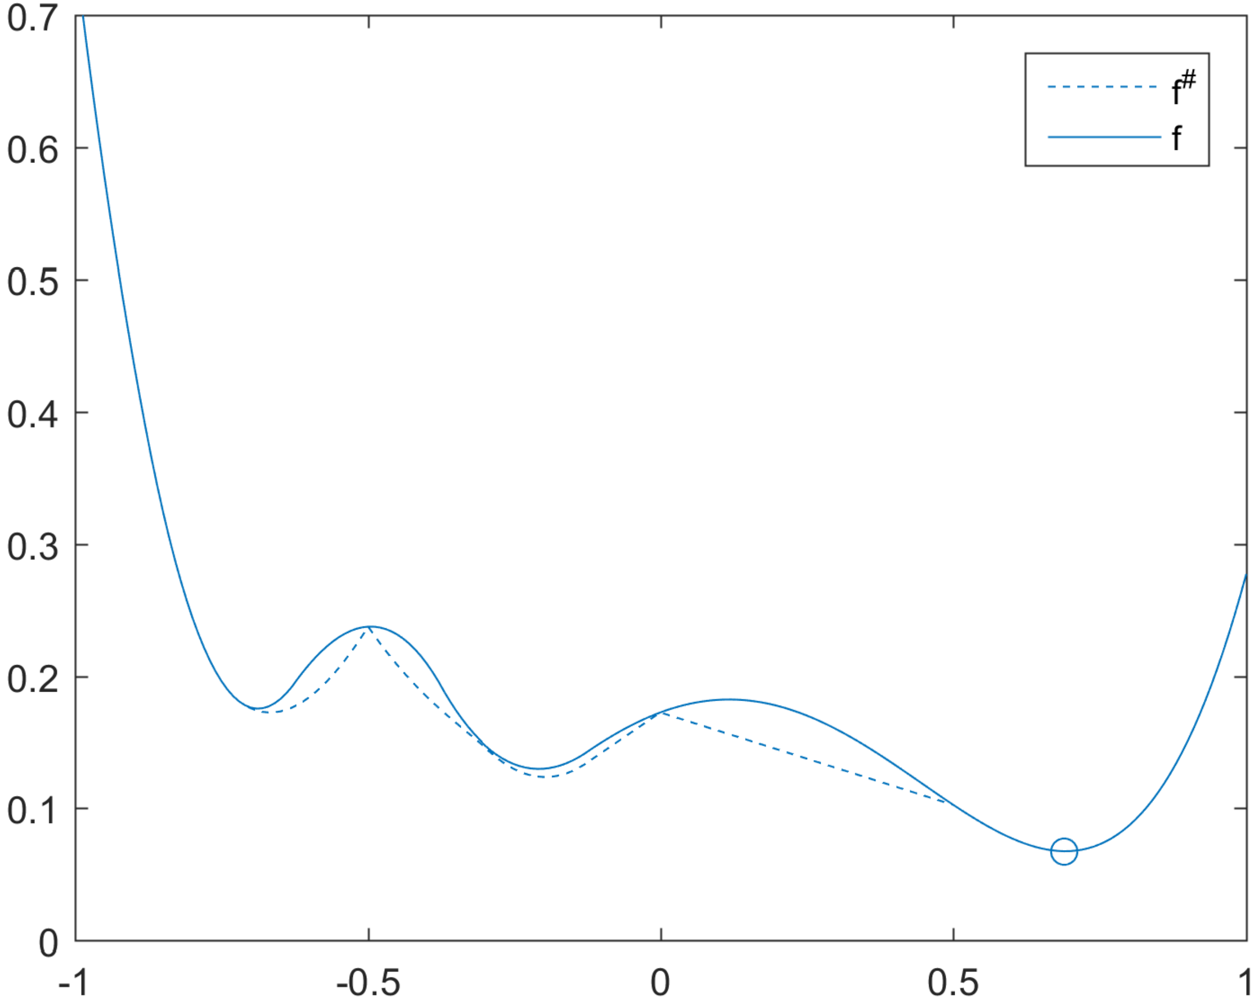

Results – disparity estimation

Möllenhoff, Laude, Möller, Lellmann, Cremers, CVPR’16

Results – disparity estimation

linear lifting

$L = 2$

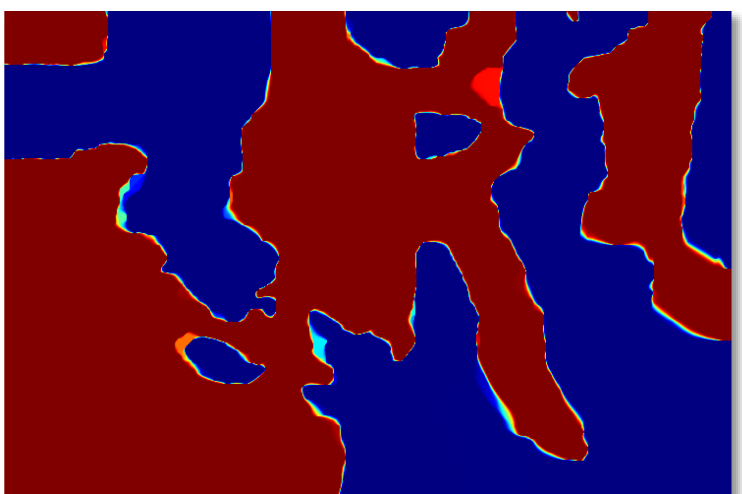

Möllenhoff, Laude, Möller, Lellmann, Cremers, CVPR'16

Results – disparity estimation

linear lifting

$L = 4$

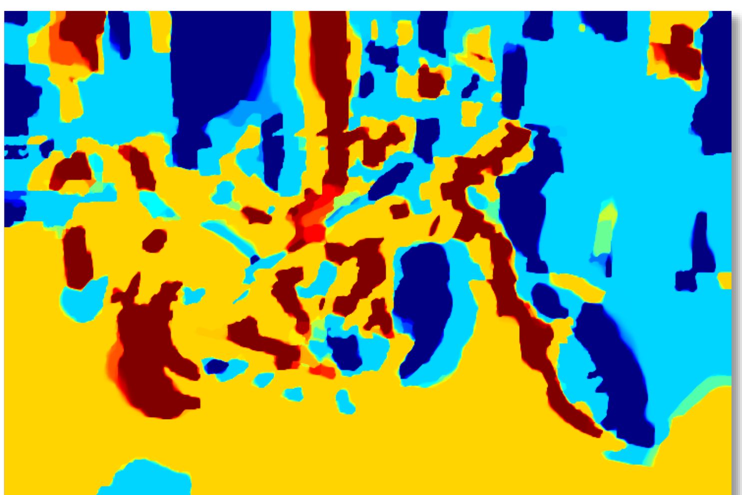

Möllenhoff, Laude, Möller, Lellmann, Cremers, CVPR'16

Results – disparity estimation

linear lifting

$L = 8$

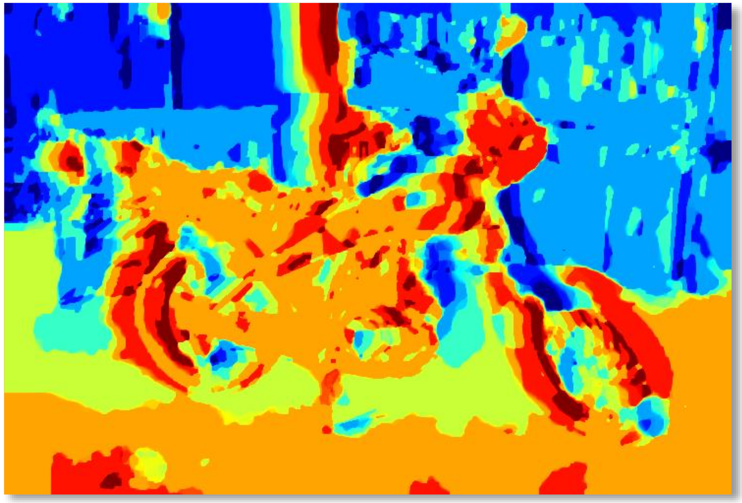

Möllenhoff, Laude, Möller, Lellmann, Cremers, CVPR'16

Results – disparity estimation

linear lifting

$L = 16$

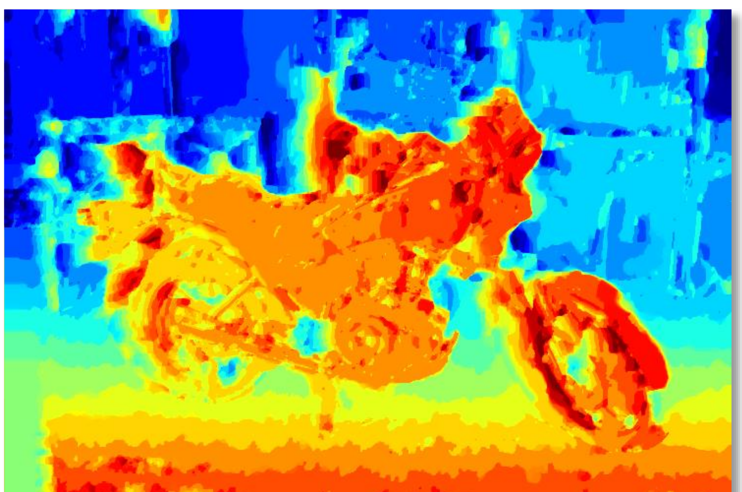

Möllenhoff, Laude, Möller, Lellmann, Cremers, CVPR’16

Results – disparity estimation

linear lifting

$$L = 32$$

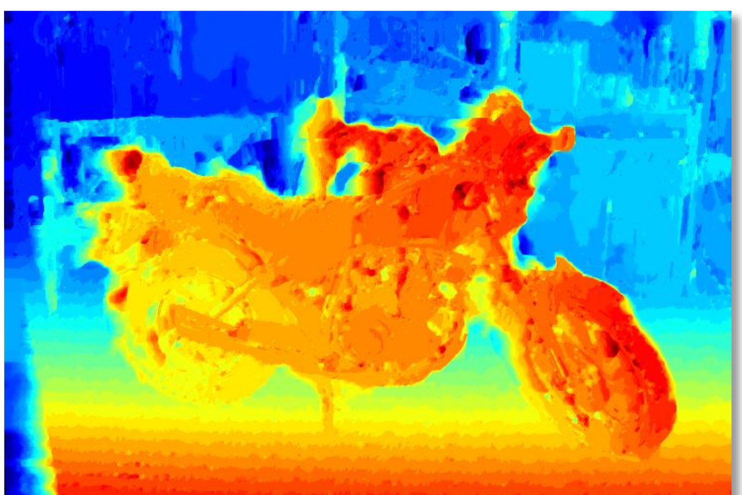

Möllenhoff, Laude, Möller, Lellmann, Cremers, CVPR'16

Results – disparity estimation

linear lifting

$L = 32$

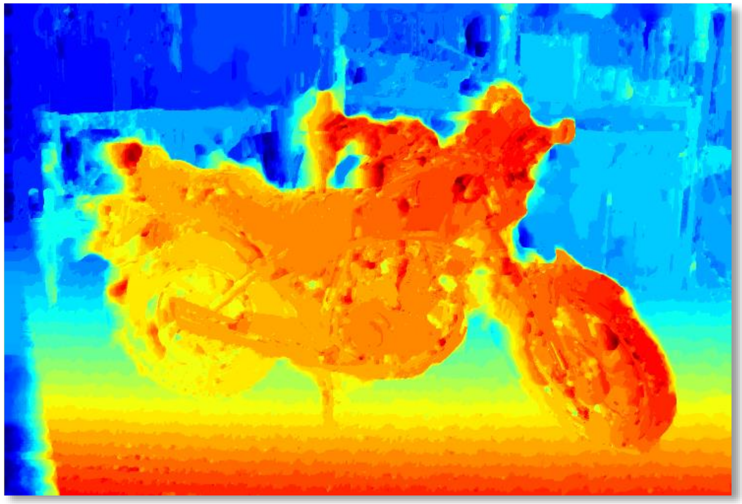

precise lifting

$L = 2$

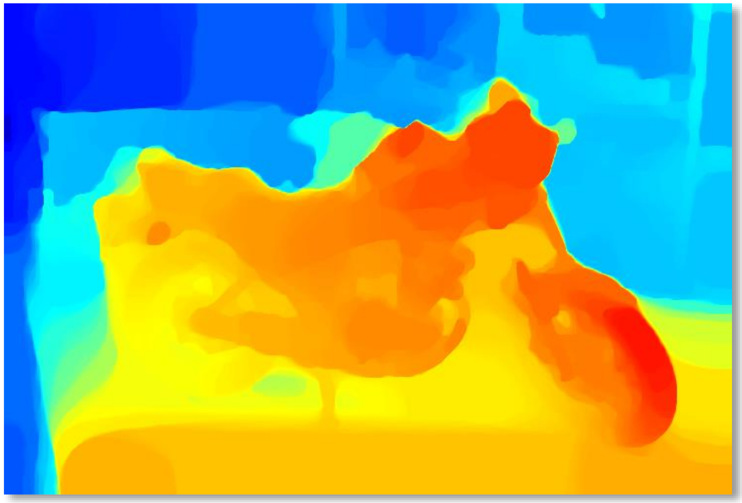

Möllenhoff, Laude, Möller, Lellmann, Cremers, CVPR'16

Results – disparity estimation

linear lifting

$L = 32$

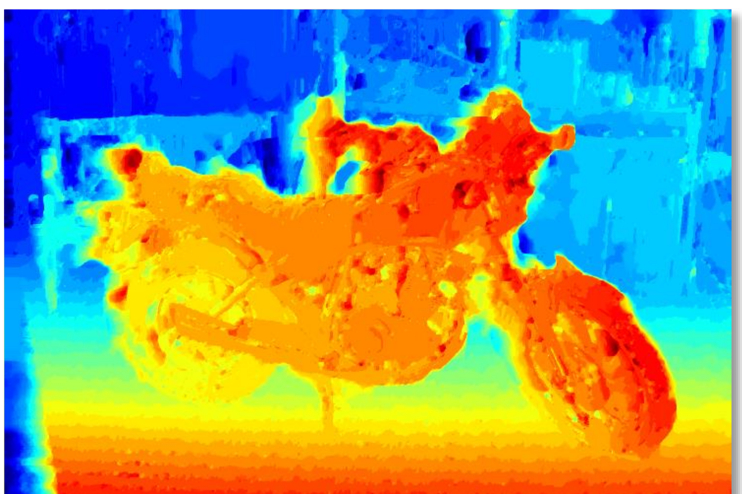

precise lifting

$L = 4$

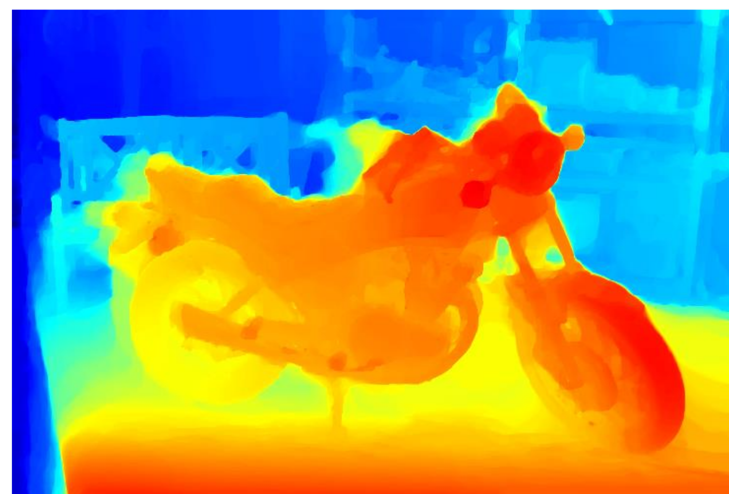

Möllenhoff, Laude, Möller, Lellmann, Cremers, CVPR'16

Results – disparity estimation

linear lifting

$L = 32$

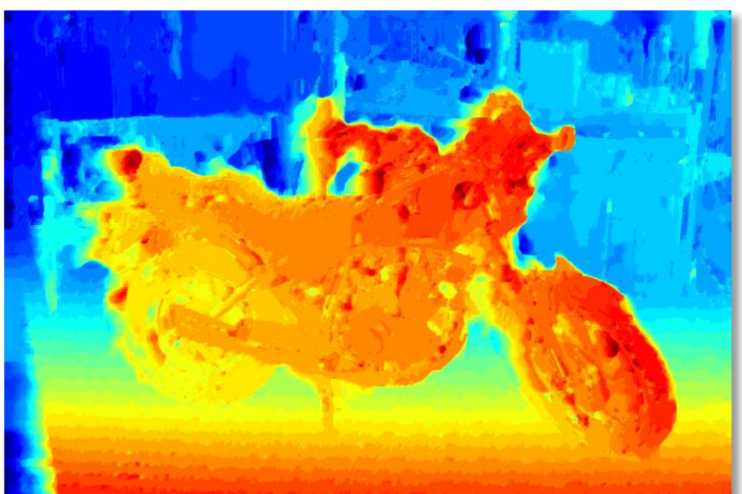

precise lifting

$L = 8$

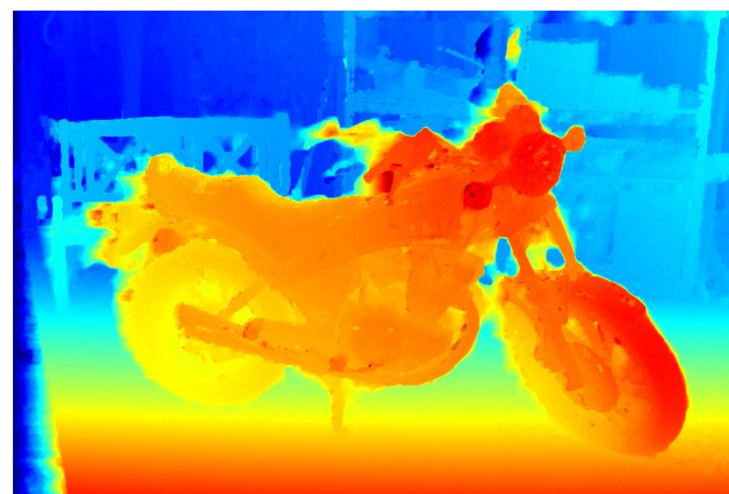

Möllenhoff, Laude, Möller, Lellmann, Cremers, CVPR'16

Results – disparity estimation

linear lifting

$$L = 270$$

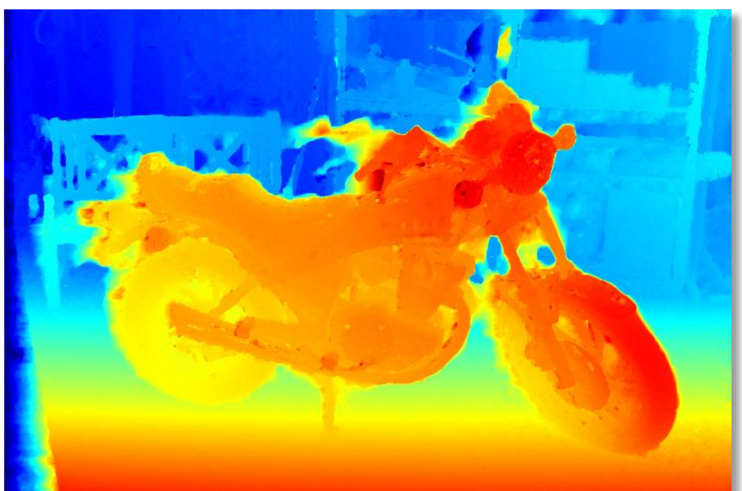

precise lifting

$$L = 8$$

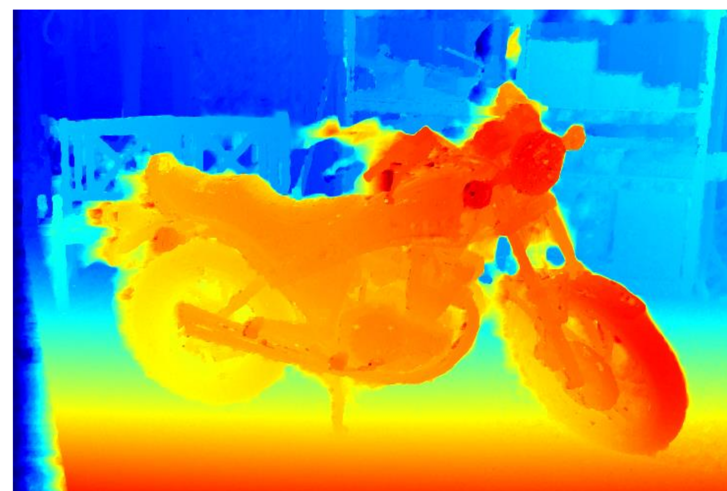

Möllenhoff, Laude, Möller, Lellmann, Cremers, CVPR'16

Related work

- Zach, Kohli'12; Zach'13; Fix, Agarwal'14
  - piecewise convex
  - different relaxation
  - MRF-based – not isotropic yet

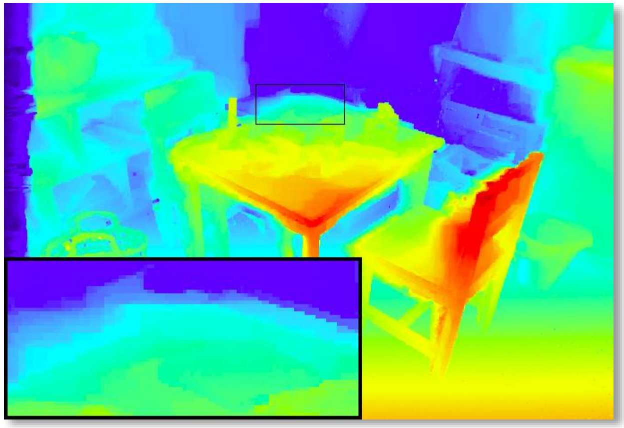

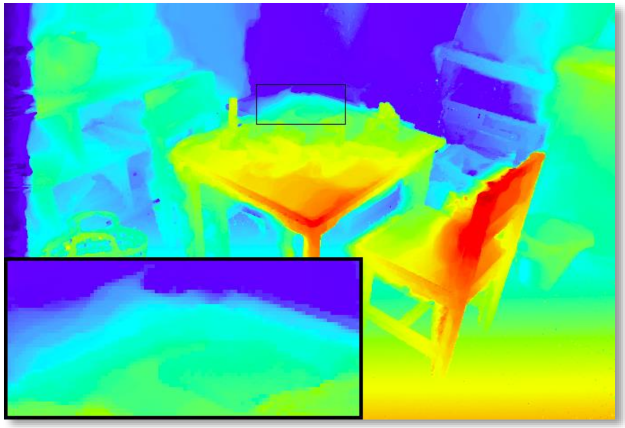

Outlook – multiple dimensions

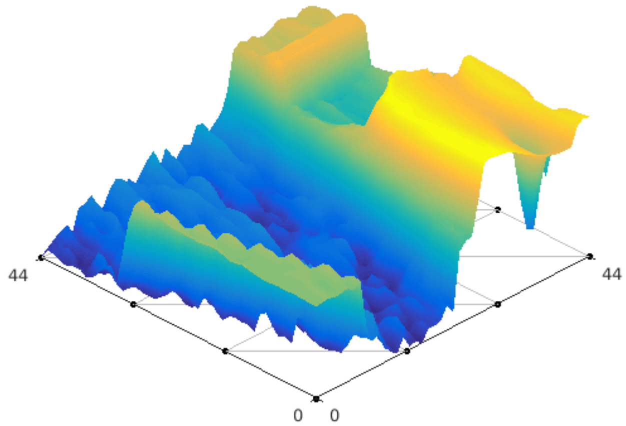

Outlook – multiple dimensions

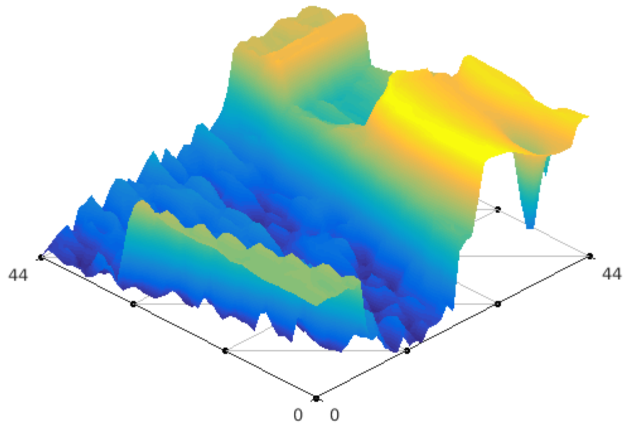

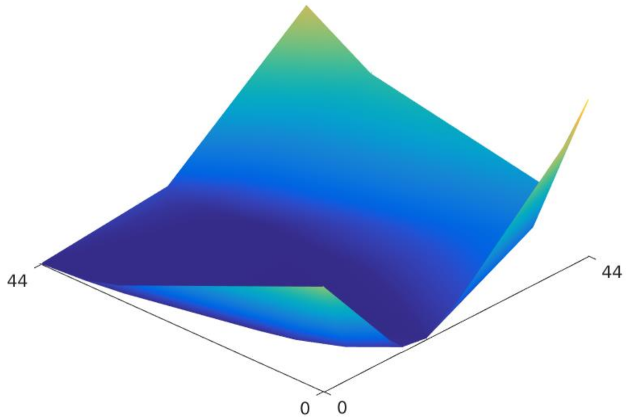

Outlook – multiple dimensions

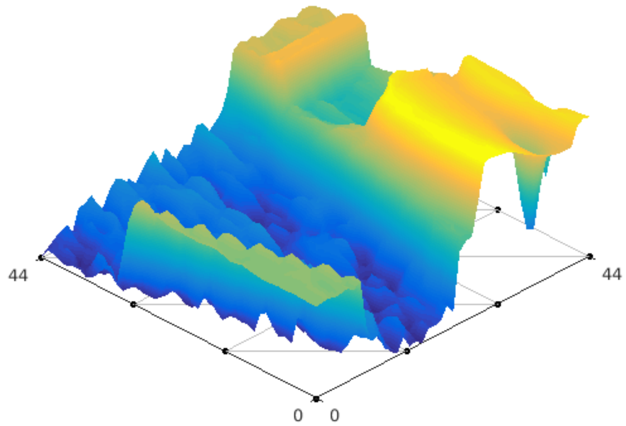

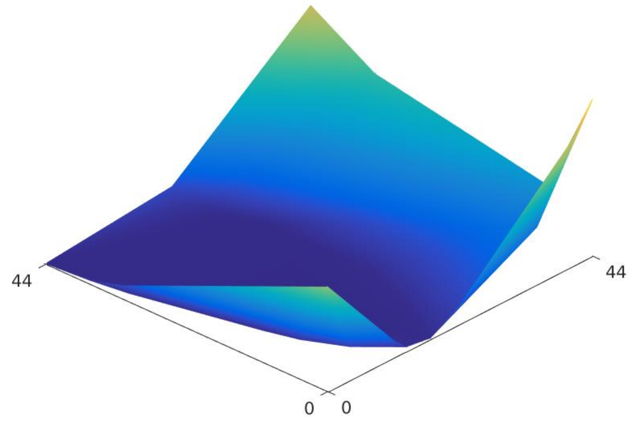

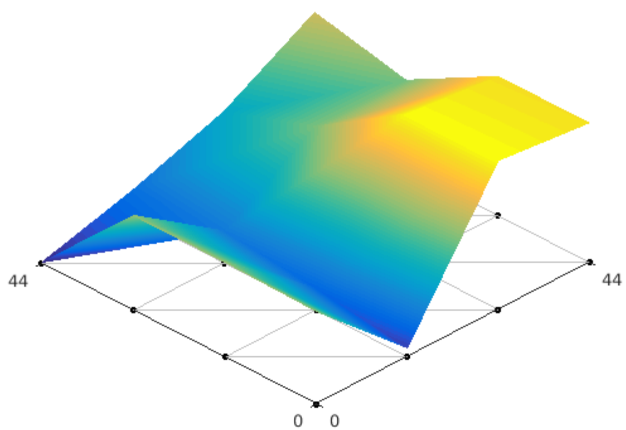

Outlook – multiple dimensions

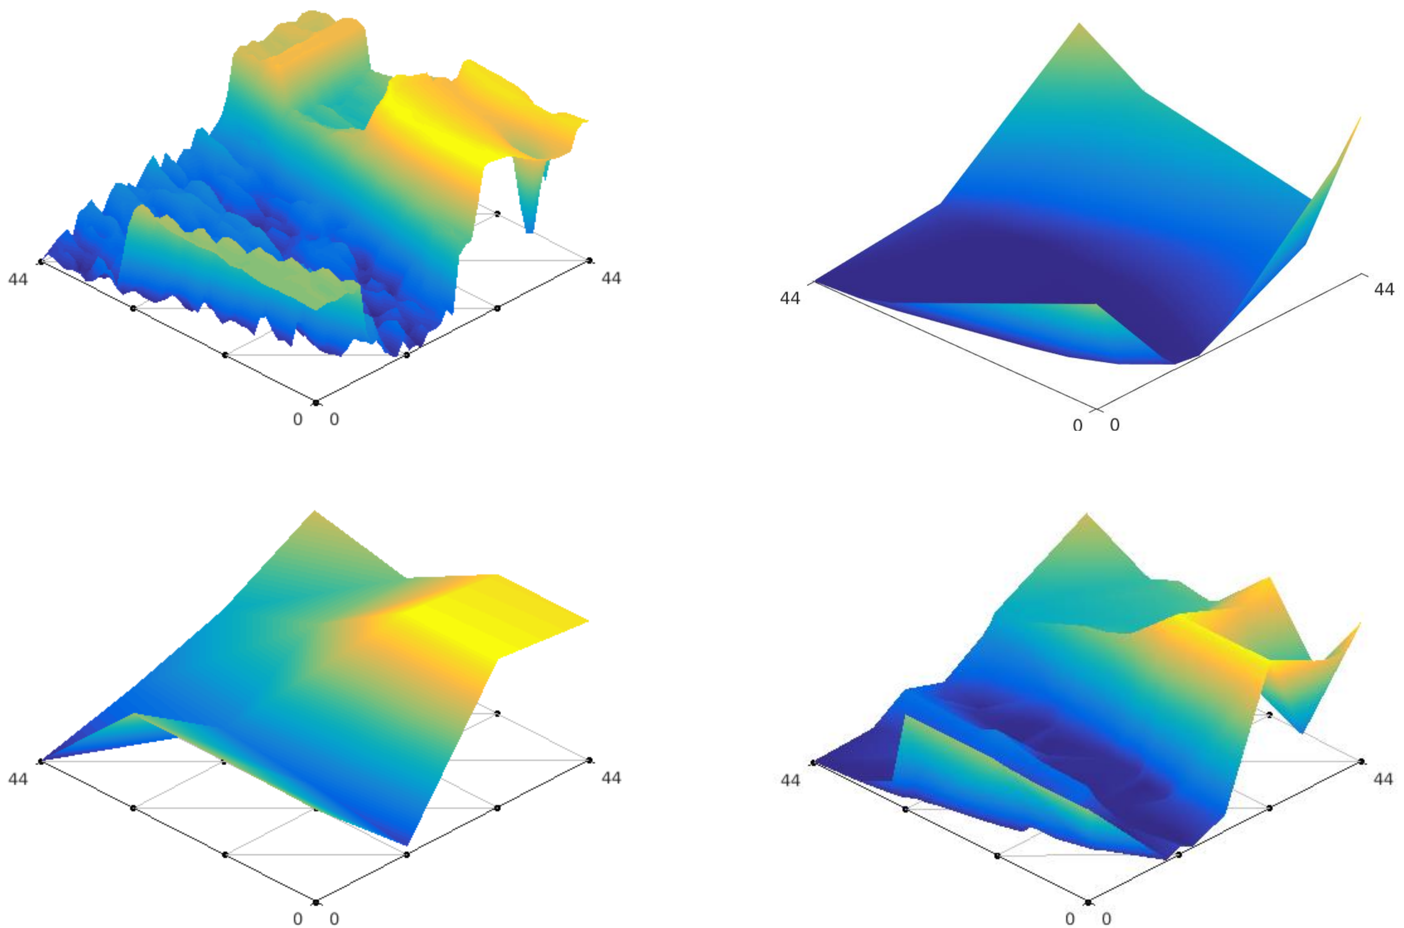

Large displacement optical flow

input sequence
(Middlebury)

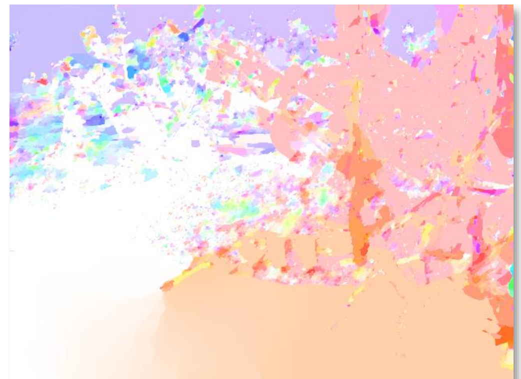

LP-style relaxation (L. et al.'13)
7x7, 5.2GB, 33min, aep 2.65

ours, 2x2 labels
0.63GB, 17min, aep1.28

ground truth

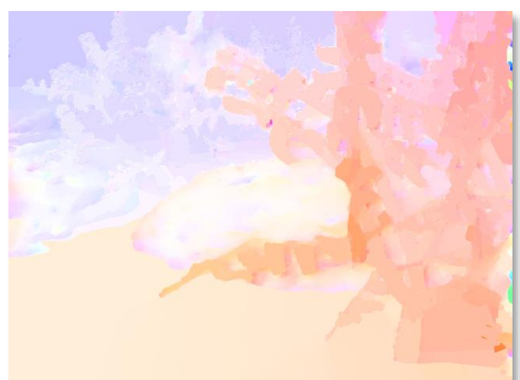

product spaces (Goldluecke et al.'13),
28x28, 9.3GB, 60min, aep 1.39

ours, 6x6 labels
10.1GB, 56min, aep 0.9

Large displacement optical flow

input sequence

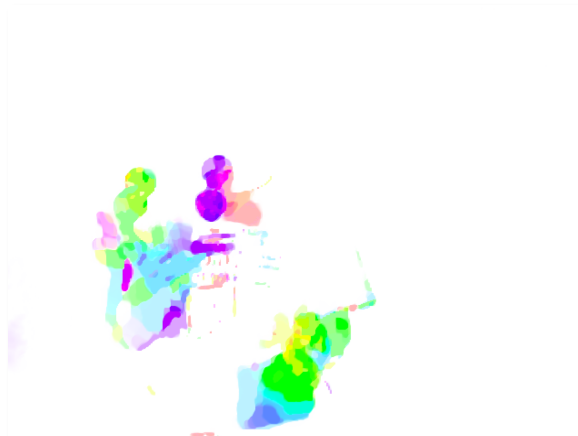

MRF-based, 81x81 labels

ours, 4 labels

References:

- T. Möllenhoff, E. Laude, M. Möller, J. Lellmann, D. Cremers: *Sublabel–Accurate Relaxation of Nonconvex Energies*. arXiv Tech. Rep., 2015
- E. Laude, T. Möllenhoff, M. Möller, J. Lellmann, D. Cremers: *Sublabel-Accurate Convex Relaxation of Vectorial Multilabel Energies*. arXiv Tech. Rep., 2016

Take-home

- Goal: global minimizer of nonconvex energies
- Lift into larger space
- Relax piecewise convex
- Much smaller problems, often 2-4 labels enough

MIC
Institute of Mathematics and Image Computing
lellmann@mic.uni-luebeck.de
Möllenhoff, Laude, Möller, Lellmann, Cremers'16

Relaxation Methods in Variational Image Processing

Jan Lellmann, Emanuel Laude, Thomas Moellenhoff,
Daniel Cremers, Michael Moeller, Evgeny Strekalovskiy

37th Northern German Colloquium on
Applied Analysis and Numerical Mathematics
23 April 2016

MIC
Institute of Mathematics and
Image Computing

lellmann@mic.uni-luebeck.de