Sublabel-Accurate Relaxation of Nonconvex Energies

Thomas Möllenhoff[1,*]


Emanuel Laude[1,*]


Michael Moeller[1]


Jan Lellmann[2]


Daniel Cremers[1]

[*]contributed equally [1]TU Munich [2]University of Lübeck

Variational Approaches in Computer Vision

image denoising

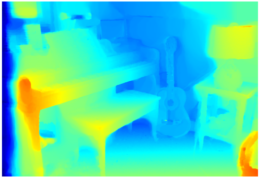

stereo matching

optical flow

Variational Approaches in Computer Vision

image denoising

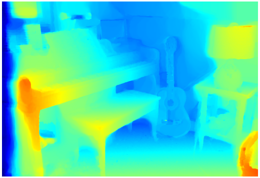

stereo matching

optical flow

Minimize energies of the form

$$\min_{u:\Omega \to \Gamma} \int_{\Omega} \rho\left(x, u(x)\right) + \lambda \cdot \left|\nabla u(x)\right| \mathrm{d}x$$

Variational Approaches in Computer Vision

image denoising

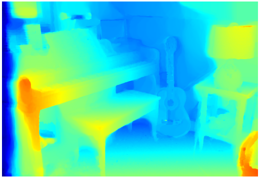

stereo matching

optical flow

Minimize energies of the form

$$\min_{u:\Omega\to\Gamma} \int_{\Omega} \rho\left(x, u(x)\right) + \lambda \cdot \left|\nabla u(x)\right| \mathrm{d}x$$

Challenges:

- Nonconvex data term $\rho : \Omega \times \Gamma \to \mathbb{R}$
- Continuous range $\Gamma = [\gamma_{\mathsf{min}}, \gamma_{\mathsf{max}}] \subset \mathbb{R}$

Discrete Multilabel Optimization [Ishikawa, TPAMI '03]

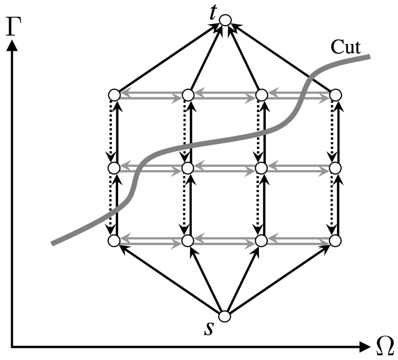

+ optimality guarantees
− discretization of $\Omega \Rightarrow$ grid bias
− discretization of $\Gamma \Rightarrow$ label bias

Continuous Lifting [Pock et al., ECCV '08]

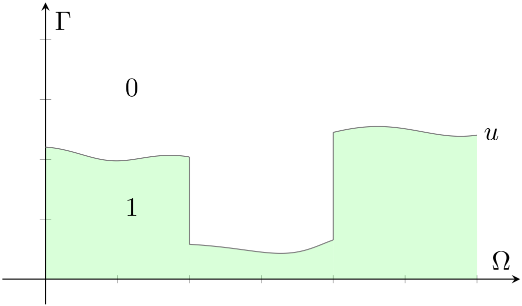

+ optimality guarantees
+ isotropic regularization $\Rightarrow$ no grid bias
− discretization of Γ $\Rightarrow$ label bias

Sublabel-Accurate Representation

Sublabel-Accurate Representation

Sublabel-Accurate Representation

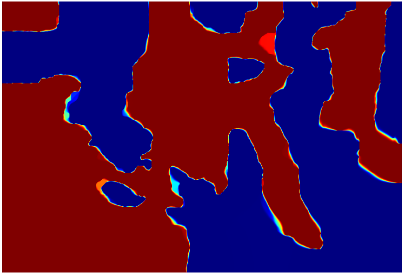

traditional representation

2 labels, 0.07 GB

Sublabel-Accurate Representation

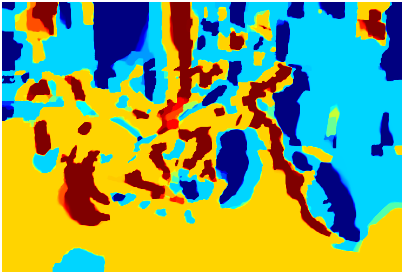

traditional representation

4 labels, 0.14 GB

Sublabel-Accurate Representation

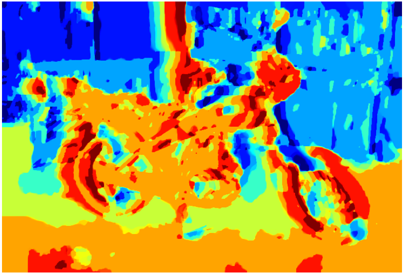

traditional representation

8 labels, 0.27 GB

Sublabel-Accurate Representation

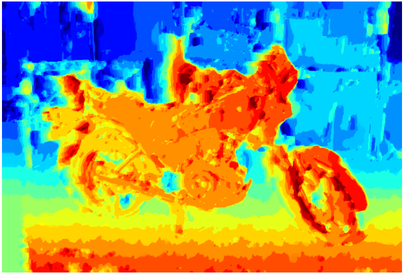

traditional representation

16 labels, 0.54 GB

Sublabel-Accurate Representation

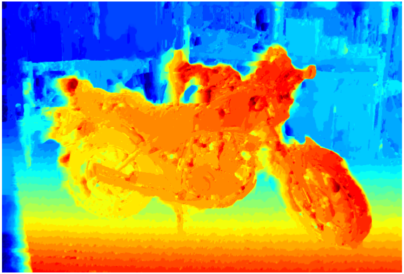

traditional representation

32 labels, 1.09 GB

Sublabel-Accurate Representation

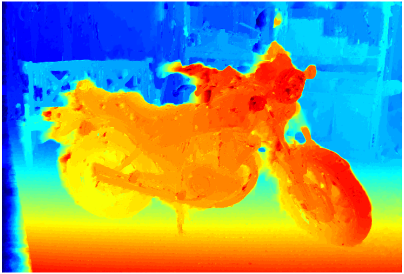

traditional representation

64 labels, 2.17 GB

Sublabel-Accurate Representation

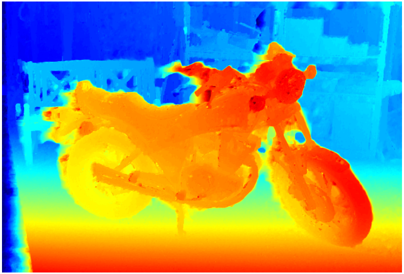

traditional representation

128 labels, 4.34 GB

Sublabel-Accurate Representation

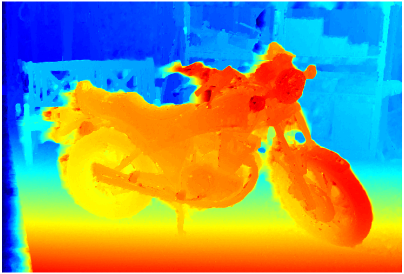

traditional representation

128 labels, 4.34 GB

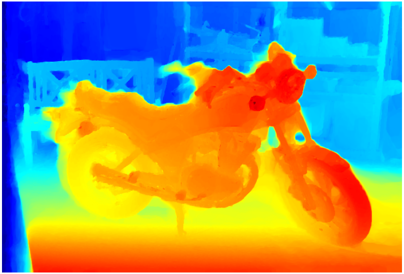

sublabel representation

8 labels, 0.57 GB

Sublabel-Accurate Representation

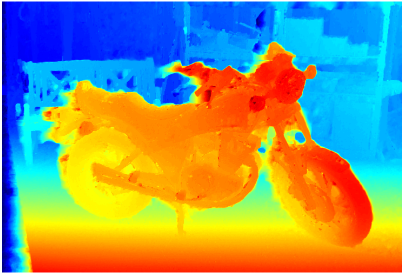

traditional representation

128 labels, 4.34 GB

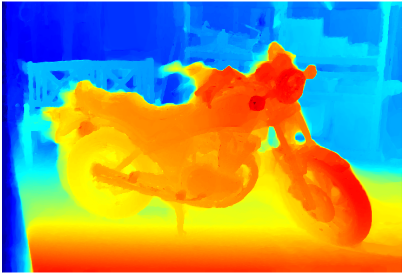

sublabel representation

8 labels, 0.57 GB

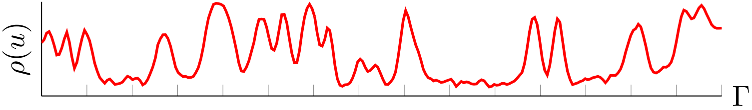

Sublabel-Accurate Representation

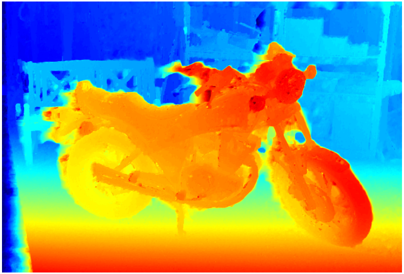

traditional representation

128 labels, 4.34 GB

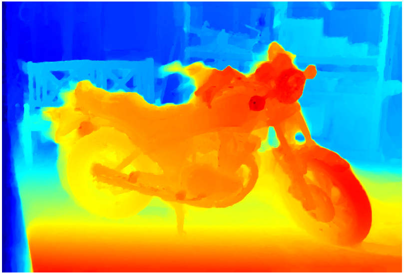

sublabel representation

8 labels, 0.57 GB

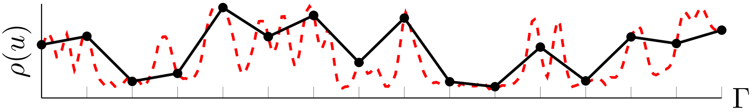

Sublabel-Accurate Representation

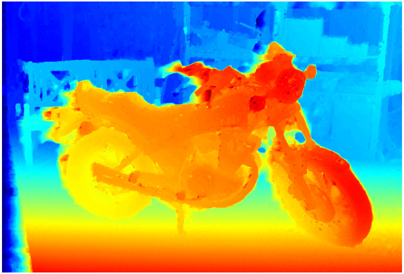

traditional representation

128 labels, 4.34 GB

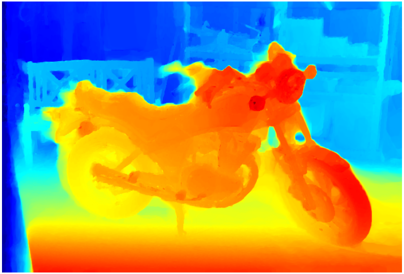

sublabel representation

8 labels, 0.57 GB

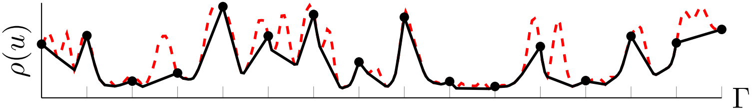

Related Work and Contribution

MRFs with continuous state spaces / continuous graphical models
[Zach, Kohli, ECCV '12], [Fix, Agarwal, ECCV '14]

Related Work and Contribution

MRFs with continuous state spaces / continuous graphical models
[Zach, Kohli, ECCV '12], [Fix, Agarwal, ECCV '14]

Partially sublabel-accurate spatially continuous multilabeling
[Lellmann et al. ICCV '13]

Related Work and Contribution

MRFs with continuous state spaces / continuous graphical models
[Zach, Kohli, ECCV '12], [Fix, Agarwal, ECCV '14]

Partially sublabel-accurate spatially continuous multilabeling
[Lellmann et al. ICCV '13]

Key contributions of this work:

+ First spatially continuous fully sublabel-accurate formulation

Related Work and Contribution

MRFs with continuous state spaces / continuous graphical models
[Zach, Kohli, ECCV '12], [Fix, Agarwal, ECCV '14]

Partially sublabel-accurate spatially continuous multilabeling
[Lellmann et al. ICCV '13]

Key contributions of this work:

+ First spatially continuous fully sublabel-accurate formulation
+ Provably tightest local convex relaxation

Related Work and Contribution

MRFs with continuous state spaces / continuous graphical models
[Zach, Kohli, ECCV '12], [Fix, Agarwal, ECCV '14]

Partially sublabel-accurate spatially continuous multilabeling
[Lellmann et al. ICCV '13]

Key contributions of this work:

+ First spatially continuous fully sublabel-accurate formulation
+ Provably tightest local convex relaxation
+ Unification of lifting and direct convex optimization

Piecewise Linear versus Piecewise Convex Lifting

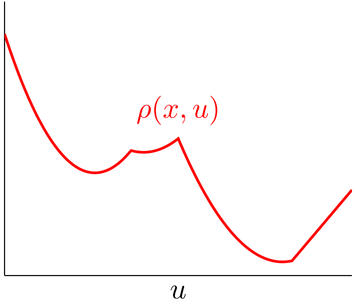

Piecewise Linear versus Piecewise Convex Lifting

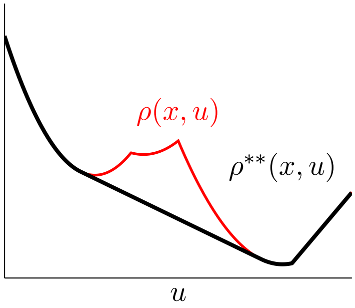

Piecewise Linear versus Piecewise Convex Lifting

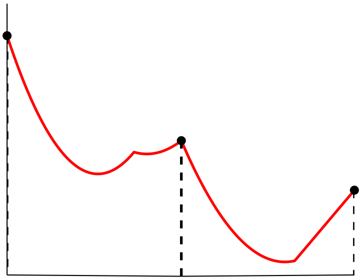

Piecewise Linear versus Piecewise Convex Lifting

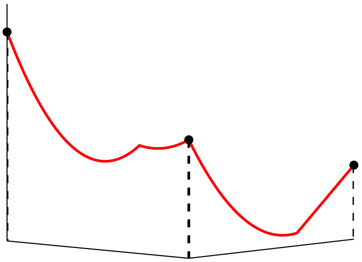

Piecewise Linear versus Piecewise Convex Lifting

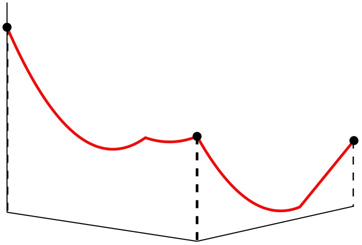

Piecewise Linear versus Piecewise Convex Lifting

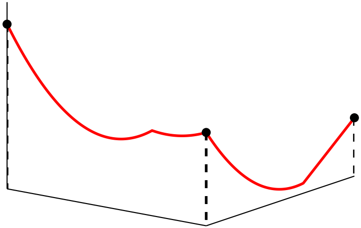

Piecewise Linear versus Piecewise Convex Lifting

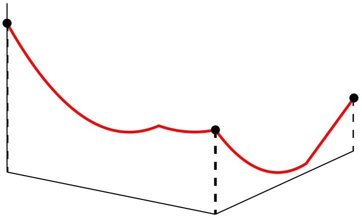

Piecewise Linear versus Piecewise Convex Lifting

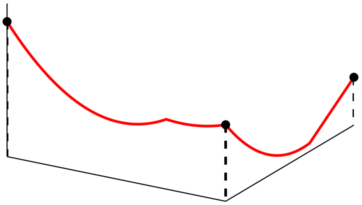

Piecewise Linear versus Piecewise Convex Lifting

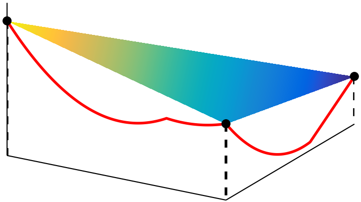

traditional relaxation

Piecewise Linear versus Piecewise Convex Lifting

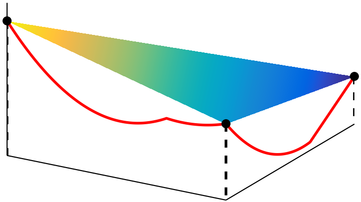

traditional relaxation

- Traditional multilabeling methods [Ishikawa, TPAMI '03], [Pock et al., ECCV '08] specify the cost only at the **labels**

Piecewise Linear versus Piecewise Convex Lifting

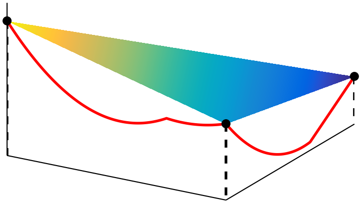

traditional relaxation

- Traditional multilabeling methods [Ishikawa, TPAMI '03], [Pock et al., ECCV '08] specify the cost only at the **labels**
- Leads to a linear relaxation, easy to optimize

Piecewise Linear versus Piecewise Convex Lifting

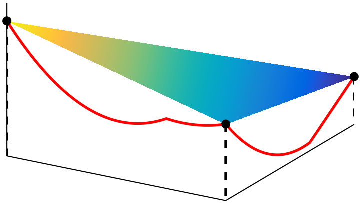

traditional relaxation

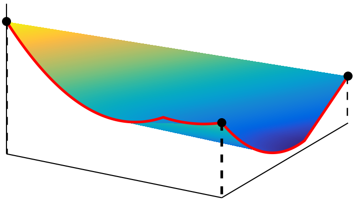

proposed relaxation

- Traditional multilabeling methods [Ishikawa, TPAMI '03], [Pock et al., ECCV '08] specify the cost only at the **labels**
- Leads to a linear relaxation, easy to optimize

Piecewise Linear versus Piecewise Convex Lifting

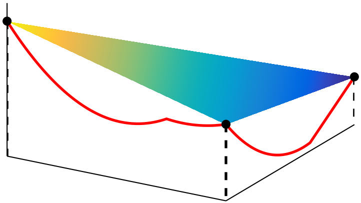

traditional relaxation

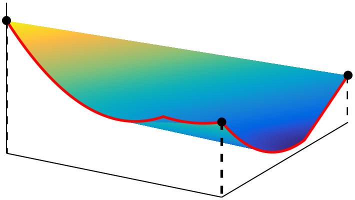

proposed relaxation

- Traditional multilabeling methods [Ishikawa, TPAMI '03], [Pock et al., ECCV '08] specify the cost only at the **labels**
- Leads to a linear relaxation, easy to optimize
- We assign meaningful cost for solutions between the labels

Piecewise Linear versus Piecewise Convex Lifting

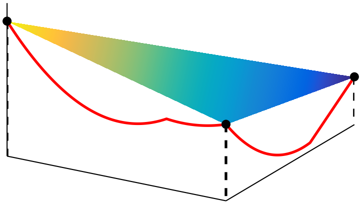

traditional relaxation

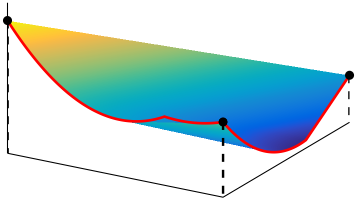

proposed relaxation

- Traditional multilabeling methods [Ishikawa, TPAMI '03], [Pock et al., ECCV '08] specify the cost only at the **labels**
- Leads to a linear relaxation, easy to optimize
- We assign meaningful cost for solutions between the labels
- The proposed relaxation is nonlinear, but still convex!

Tightest Convex Extension

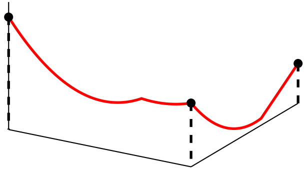

Tightest Convex Extension

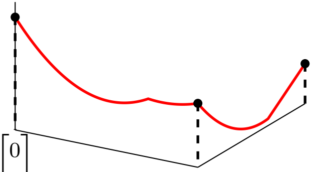

$\mathbf{1}_0 = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$

Tightest Convex Extension

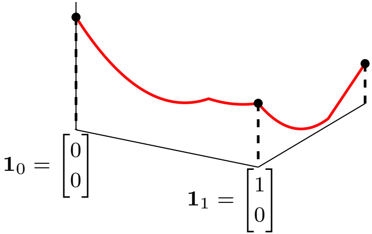

Tightest Convex Extension

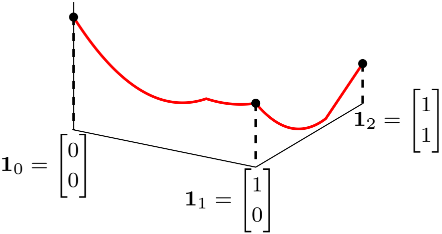

Tightest Convex Extension

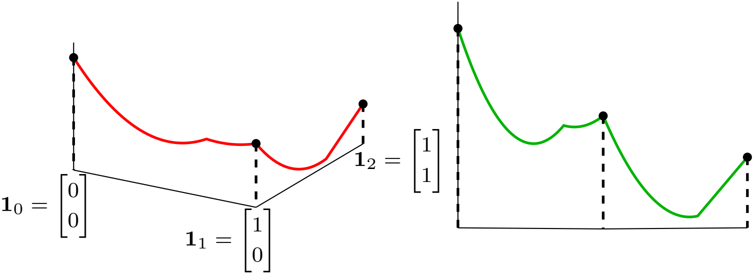

Tightest Convex Extension

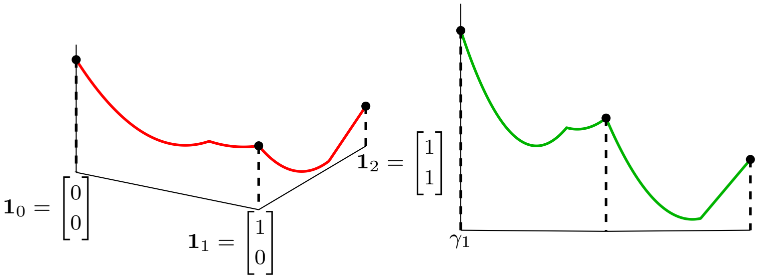

Tightest Convex Extension

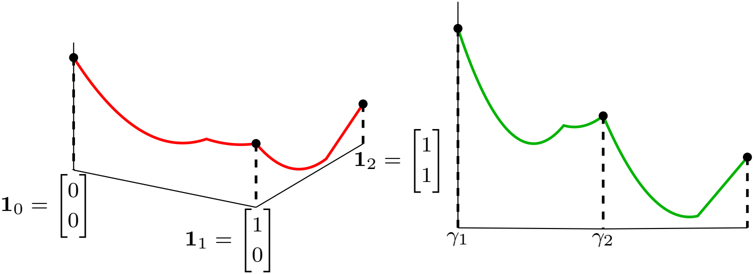

Tightest Convex Extension

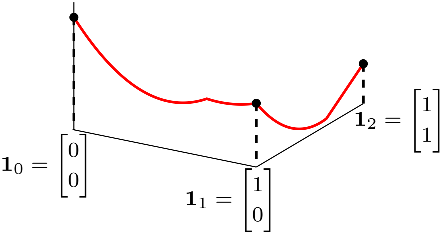

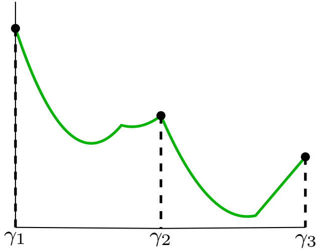

Tightest Convex Extension

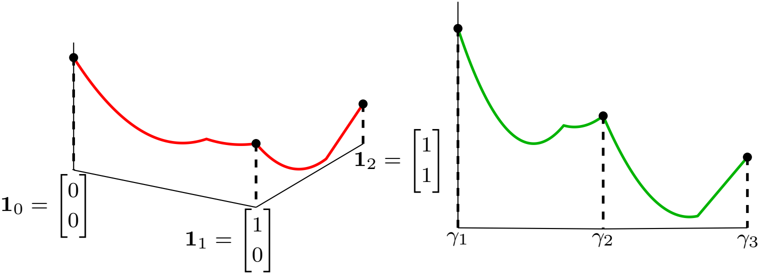

$$\boldsymbol{\rho}(\boldsymbol{u}) = \begin{cases} \rho(\gamma_i + \alpha(\gamma_{i+1} - \gamma_i)), & \text{if } \boldsymbol{u} = \mathbf{1}_{i-1} + \alpha\left(\mathbf{1}_i - \mathbf{1}_{i-1}\right), \\ \infty, & \text{otherwise.} \end{cases}$$

Tightest Convex Extension

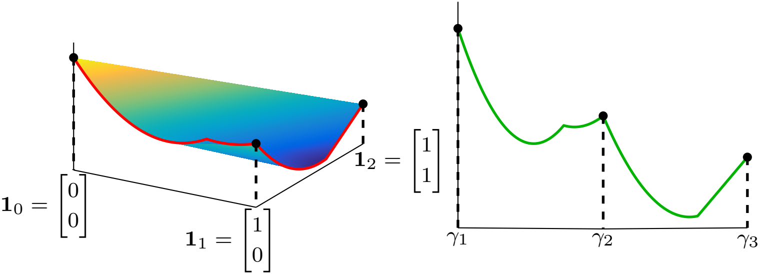

$$\boldsymbol{\rho}(\boldsymbol{u}) = \begin{cases} \rho(\gamma_i + \alpha(\gamma_{i+1} - \gamma_i)), & \text{if } \boldsymbol{u} = \mathbf{1}_{i-1} + \alpha\left(\mathbf{1}_i - \mathbf{1}_{i-1}\right), \\ \infty, & \text{otherwise.} \end{cases}$$

Proposition: The tightest convex extension is given by

$$\boldsymbol{\rho}^{**}(\boldsymbol{u}) = \sup_{\boldsymbol{v} \in \mathcal{C}} \left\langle \begin{bmatrix} \boldsymbol{u} & -1 \end{bmatrix}^\top, \boldsymbol{v} \right\rangle$$

Tightest Convex Extension

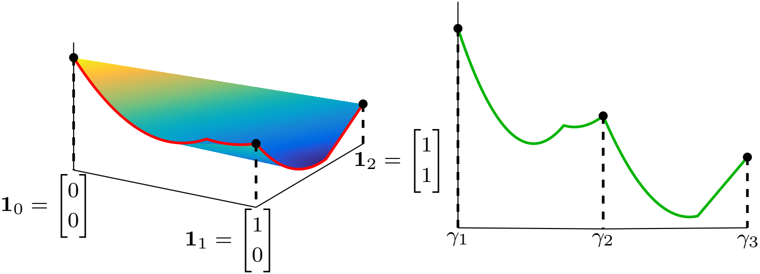

$$\boldsymbol{\rho}(\boldsymbol{u}) = \begin{cases} \rho(\gamma_i + \alpha(\gamma_{i+1} - \gamma_i)), & \text{if } \boldsymbol{u} = \mathbf{1}_{i-1} + \alpha\left(\mathbf{1}_i - \mathbf{1}_{i-1}\right), \\ \infty, & \text{otherwise.} \end{cases}$$

Proposition: The tightest convex extension is given by

$$\boldsymbol{\rho}^{**}(\boldsymbol{u}) = \sup_{\boldsymbol{v} \in \mathcal{C}} \left\langle \begin{bmatrix} \boldsymbol{u} & -1 \end{bmatrix}^\top, \boldsymbol{v} \right\rangle, \; \mathcal{C} = \left\{ \boldsymbol{v} \in \mathbb{R}^L \mid A_i \boldsymbol{v} \in \operatorname{epi}\left(\rho_i^*\right), \forall i \right\}$$

Tightest Convex Extension

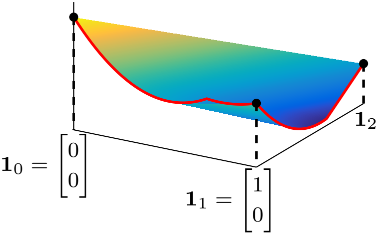

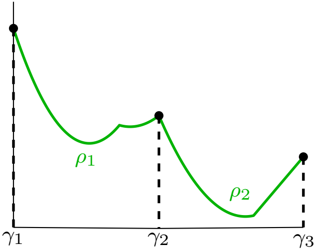

$$\boldsymbol{\rho}(\boldsymbol{u}) = \begin{cases} \rho(\gamma_i + \alpha(\gamma_{i+1} - \gamma_i)), & \text{if } \boldsymbol{u} = \mathbf{1}_{i-1} + \alpha\left(\mathbf{1}_i - \mathbf{1}_{i-1}\right), \\ \infty, & \text{otherwise.} \end{cases}$$

Proposition: The tightest convex extension is given by

$$\boldsymbol{\rho}^{**}(\boldsymbol{u}) = \sup_{\boldsymbol{v} \in \mathcal{C}} \left\langle \begin{bmatrix} \boldsymbol{u} & -1 \end{bmatrix}^\top, \boldsymbol{v} \right\rangle, \; \mathcal{C} = \left\{ \boldsymbol{v} \in \mathbb{R}^L \mid A_i \boldsymbol{v} \in \text{epi}\left(\rho_i^*\right), \forall i \right\}$$

Numerical Optimization

Proposition: Tight local convex extension for lifted regularizer is

$$\int_\Omega |\nabla u| \mathrm{d}x \leftrightarrow \sup_{\mathbf{p}:\Omega \to \mathcal{K}} \langle \mathbf{u}, \mathrm{Div}\, \mathbf{p} \rangle, \; \mathcal{K} = \{\mathbf{p} \mid \|\mathbf{p}_i\| \leq \gamma_{i+1} - \gamma_i, \forall i\}$$

Numerical Optimization

Proposition: Tight local convex extension for lifted regularizer is

$$\int_{\Omega} |\nabla u| \mathrm{d}x \leftrightarrow \sup_{\mathbf{p}:\Omega \to \mathcal{K}} \langle \mathbf{u}, \operatorname{Div} \mathbf{p} \rangle, \; \mathcal{K} = \{\mathbf{p} \mid \|\mathbf{p}_i\| \leq \gamma_{i+1} - \gamma_i, \forall i\}$$

- Leads to convex-concave saddle-point problem

$$\min_{\mathbf{u}:\Omega \to \mathbb{R}^{L-1}} \max_{\substack{\mathbf{v}:\Omega \to \mathcal{C} \\ \mathbf{p}:\Omega \to \mathcal{K}}} \langle \mathbf{u}, \operatorname{Div} \mathbf{p} \rangle + \left\langle \begin{bmatrix} \boldsymbol{u} & -1 \end{bmatrix}^\mathsf{T}, \boldsymbol{v} \right\rangle$$

Numerical Optimization

Proposition: Tight local convex extension for lifted regularizer is

$$\int_{\Omega} |\nabla u| \mathrm{d}x \leftrightarrow \sup_{\mathbf{p}:\Omega \to \mathcal{K}} \langle \mathbf{u}, \operatorname{Div} \mathbf{p} \rangle, \; \mathcal{K} = \{ \mathbf{p} \mid \|\mathbf{p}_i\| \leq \gamma_{i+1} - \gamma_i, \forall i \}$$

- Leads to convex-concave saddle-point problem

$$\min_{\mathbf{u}:\Omega \to \mathbb{R}^{L-1}} \max_{\substack{\mathbf{v}:\Omega \to \mathcal{C} \\ \mathbf{p}:\Omega \to \mathcal{K}}} \langle \mathbf{u}, \operatorname{Div} \mathbf{p} \rangle + \left\langle \begin{bmatrix} \boldsymbol{u} & -1 \end{bmatrix}^{\mathsf{T}}, \boldsymbol{v} \right\rangle$$

- Solved on GPU using a first-order primal-dual algorithm [Pock, Cremers, Bischof, Chambolle, ICCV '09]

Backprojecting the Lifted Solution

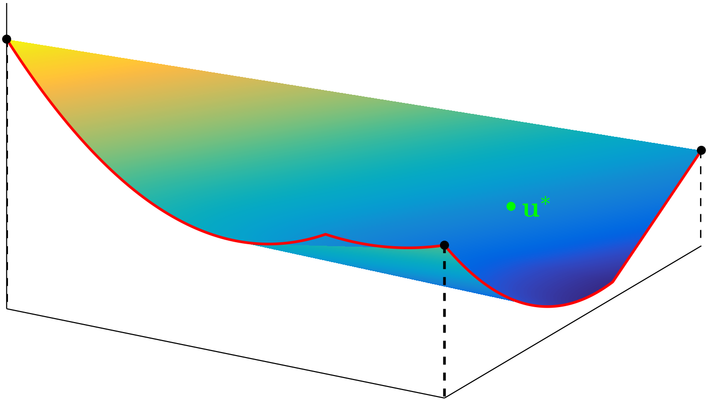

Backprojecting the Lifted Solution

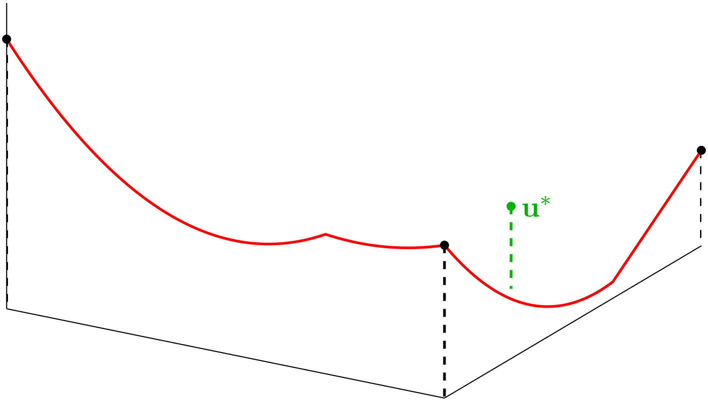

Backprojecting the Lifted Solution

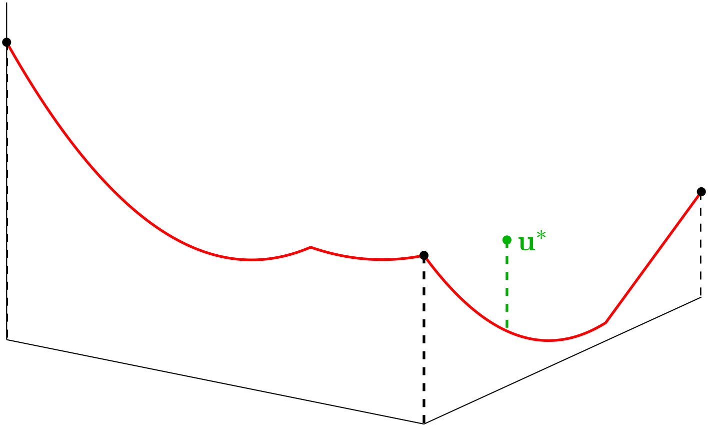

Backprojecting the Lifted Solution

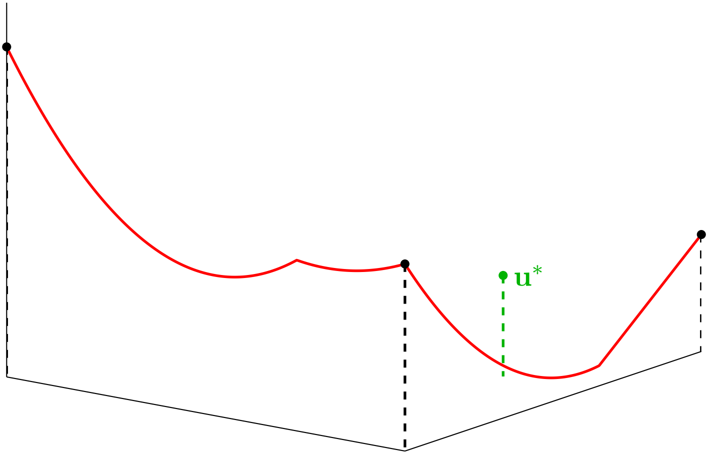

Backprojecting the Lifted Solution

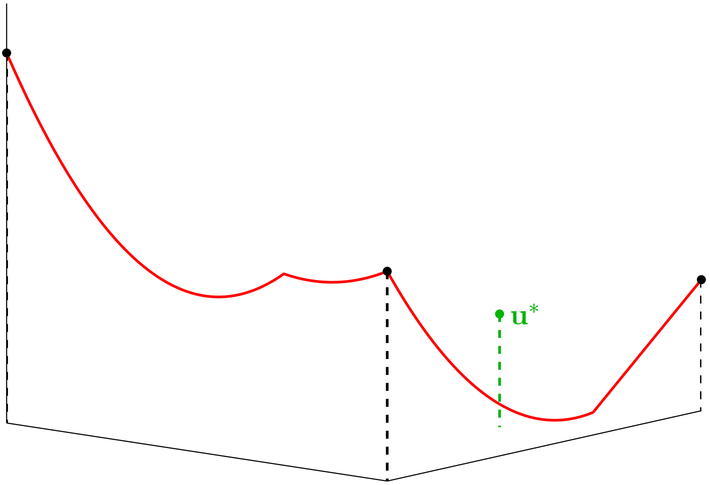

Backprojecting the Lifted Solution

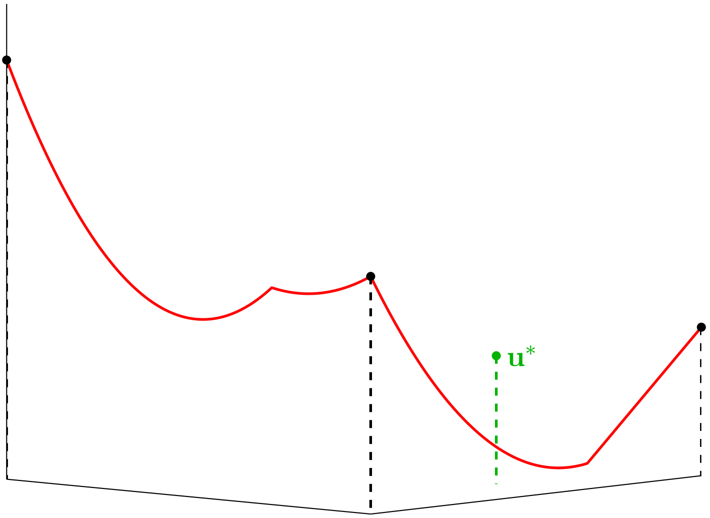

Backprojecting the Lifted Solution

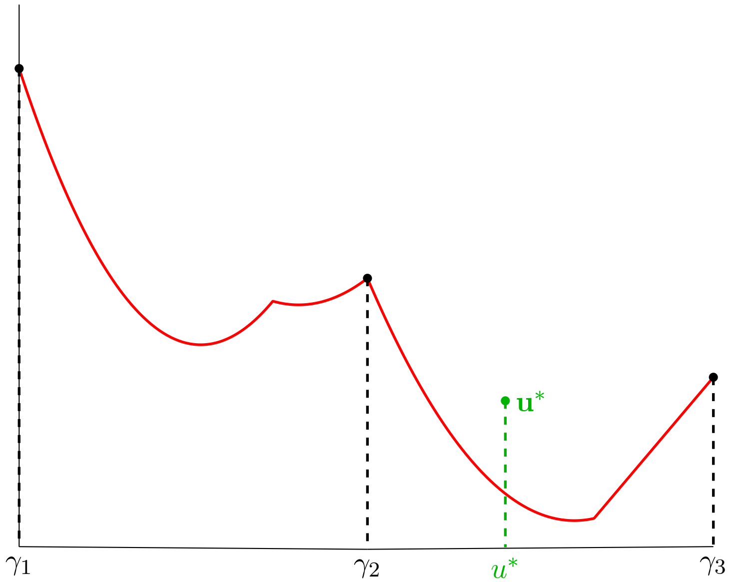

Convex Case: $\rho\left(x, u(x)\right) = \left(u(x) - f(x)\right)^2$

direct, no labels
$0.6s$, 11.78 MB

Convex Case: $\rho\left(x, u(x)\right) = \left(u(x) - f(x)\right)^2$

direct, no labels
$0.6s$, 11.78 MB

sublabel, 2 labels,
$1s$, 27 MB

Convex Case: $\rho\left(x, u(x)\right) = \left(u(x) - f(x)\right)^2$

direct, no labels
$0.6s$, 11.78 MB

sublabel, 2 labels,
$1s$, 27 MB

sublabel, 10 labels,
$15s$, 211 MB

Convex Case: $\rho\left(x, u(x)\right) = \left(u(x) - f(x)\right)^2$

direct, no labels
$0.6s$, 11.78 MB

sublabel, 2 labels,
$1s$, 27 MB

sublabel, 10 labels,
$15s$, 211 MB

traditional, 8 labels,
suboptimal, 113 MB

Convex Case: $\rho\left(x, u(x)\right) = \left(u(x) - f(x)\right)^2$

direct, no labels
$0.6s$, 11.78 MB

sublabel, 2 labels,
$1s$, 27 MB

sublabel, 10 labels,
$15s$, 211 MB

traditional, 8 labels,
suboptimal, 113 MB

traditional, 16 labels,
suboptimal, 226 MB

Convex Case: $\rho\left(x, u(x)\right) = \left(u(x) - f(x)\right)^2$

direct, no labels
$0.6s$, 11.78 MB

sublabel, 2 labels,
$1s$, 27 MB

sublabel, 10 labels,
$15s$, 211 MB

traditional, 8 labels,
suboptimal, 113 MB

traditional, 16 labels,
suboptimal, 226 MB

traditional, 256 labels,
suboptimal, 3619 MB

Stereo Matching, $\rho(x, u(x)) = \|I_1(x) - I_2(x_1 + u(x), x_2)\|$

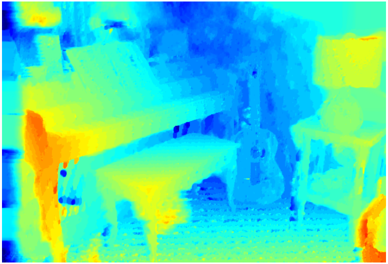

traditional, 48 labels, 52s, 1.49 GB

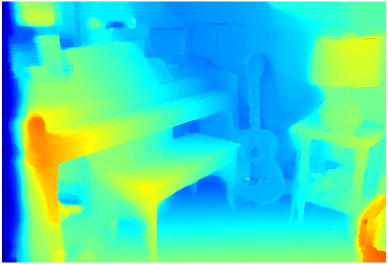

sublabel, 8 labels, 30s, 0.49 GB

Stereo Matching, $\rho(x, u(x)) = \|I_1(x) - I_2(x_1 + u(x), x_2)\|$

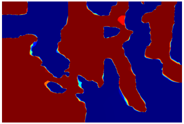

traditional, 2 labels

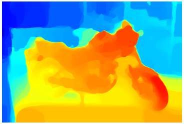

sublabel, 2 labels

Stereo Matching, $\rho(x, u(x)) = \|I_1(x) - I_2(x_1 + u(x), x_2)\|$

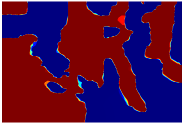

traditional, 2 labels

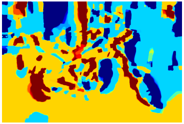

traditional, 4 labels

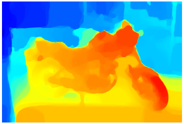

sublabel, 2 labels

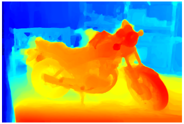

sublabel, 4 labels

Stereo Matching, $\rho(x, u(x)) = \|I_1(x) - I_2(x_1 + u(x), x_2)\|$

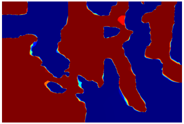

traditional, 2 labels

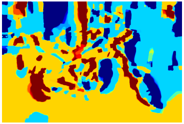

traditional, 4 labels

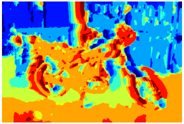

traditional, 8 labels

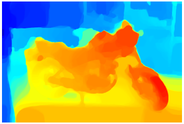

sublabel, 2 labels

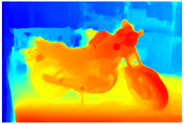

sublabel, 4 labels

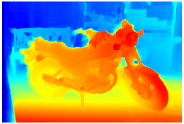

sublabel, 8 labels

Stereo Matching, $\rho(x, u(x)) = \|I_1(x) - I_2(x_1 + u(x), x_2)\|$

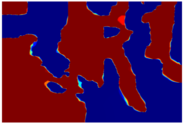
traditional, 2 labels

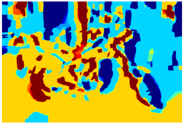
traditional, 4 labels

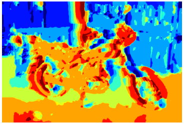
traditional, 8 labels

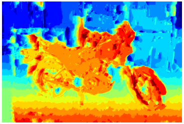
traditional, 16 labels

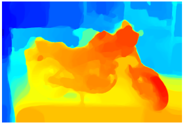
sublabel, 2 labels

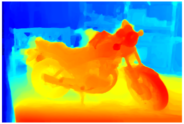
sublabel, 4 labels

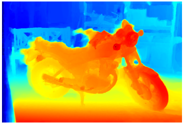
sublabel, 8 labels

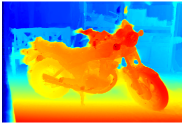
sublabel, 16 labels

Conclusion

- We proposed a sublabel-accurate relaxation for a certain class of nonconvex energies

Conclusion

- We proposed a sublabel-accurate relaxation for a certain class of nonconvex energies
- Requires far fewer labels than traditional lifting techniques

Conclusion

- We proposed a sublabel-accurate relaxation for a certain class of nonconvex energies
- Requires far fewer labels than traditional lifting techniques
- Leads to substantial improvements in runtime and memory

Conclusion

- We proposed a sublabel-accurate relaxation for a certain class of nonconvex energies
- Requires far fewer labels than traditional lifting techniques
- Leads to substantial improvements in runtime and memory
- It generalizes traditional lifting methods from piecewise linear to piecewise convex approximations

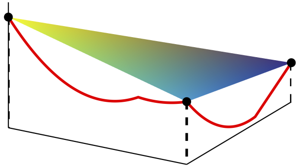

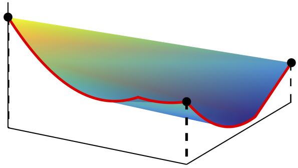

Conclusion

- We proposed a sublabel-accurate relaxation for a certain class of nonconvex energies
- Requires far fewer labels than traditional lifting techniques
- Leads to substantial improvements in runtime and memory
- It generalizes traditional lifting methods from piecewise linear to piecewise convex approximations

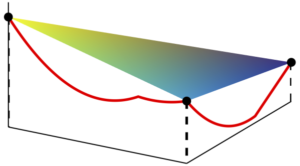
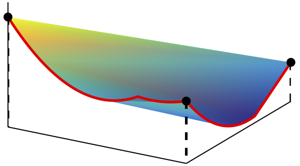

https://github.com/tum-vision/sublabel_relax

Comparison with [Zach, Kohli, ECCV '12]

- Denoising with robust data term

$$\rho(x, u) = (\alpha/2) \min \left\{ \nu, (u - f(x))^2 \right\}$$

- Special case of our method: anisotropic regularizer $\|\nabla u\|_1$
- Our relaxation uses only **half** the number of variables

proposed tight relaxation,
33 labels, **Energy: 194836**

[Zach, Kohli, ECCV '12], DC-MRF,
33 labels, Energy: 194845