JHU vision lab

Distributed Consensus Algorithms on Manifolds

Roberto Tron, Bijan Afsari, René Vidal

Johns Hopkins University

Center for Imaging Science

Distributed algorithms in controls

Motion coordination

Formation control

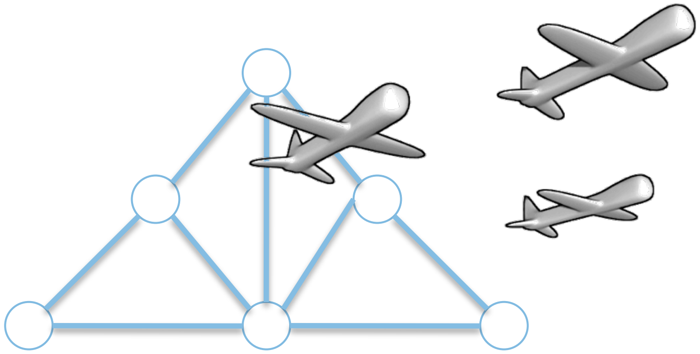

Collaborative mapping

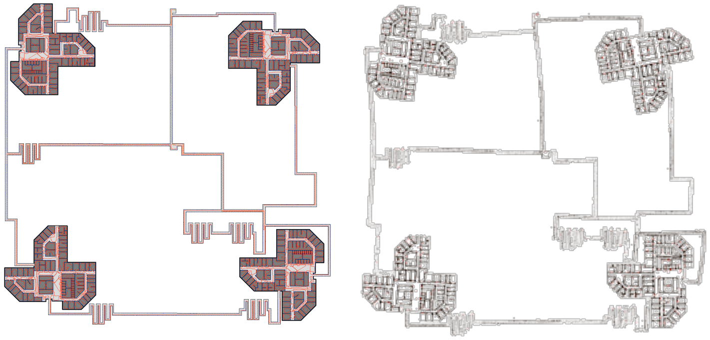

Network optimization

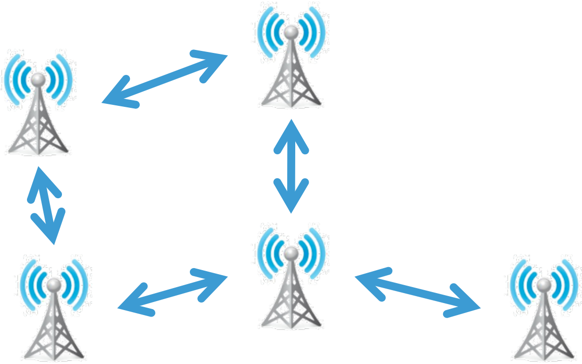

Consensus algorithms

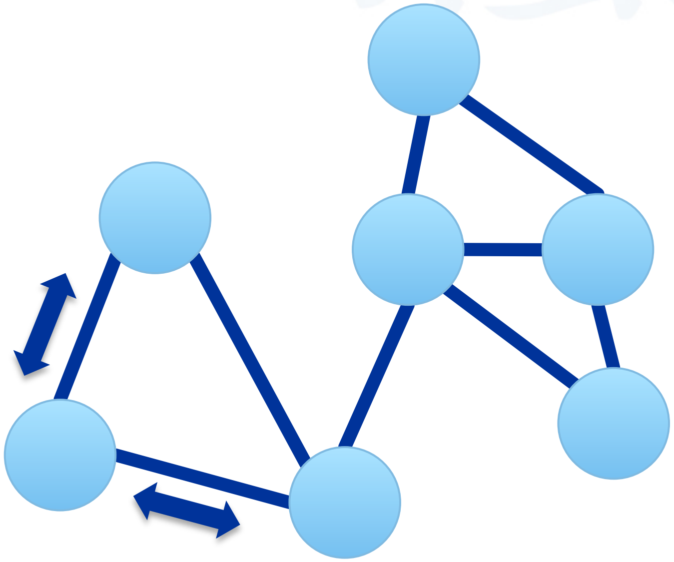

Global agreement from local interactions

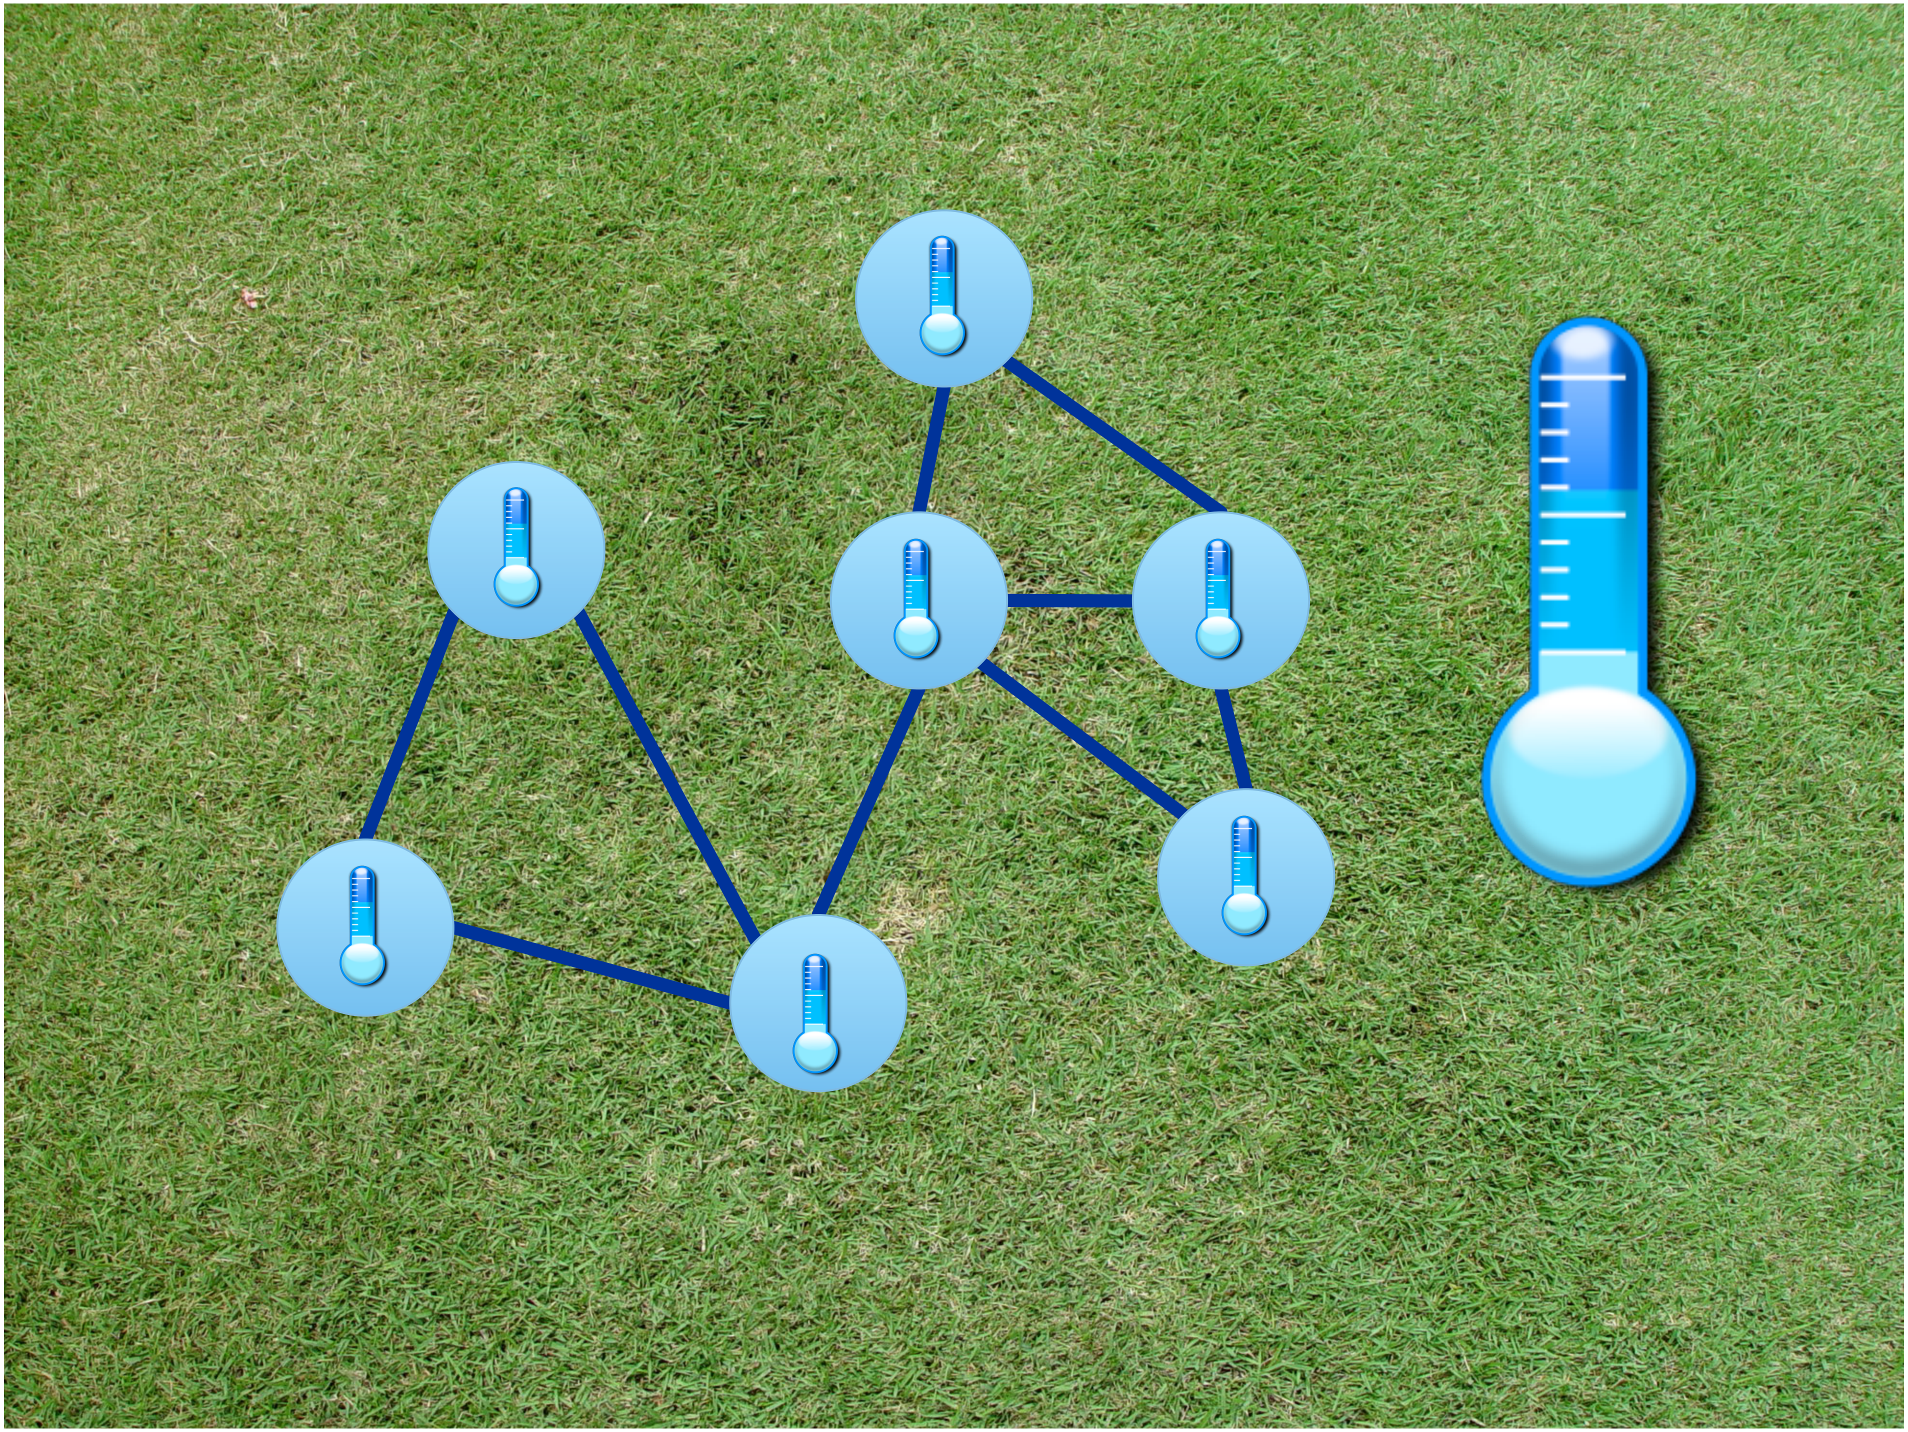

New Settings

Cameras

Today

YouTube 24 hrs/min

30 million cameras
4 bn hrs/week

Tasks	Detection, Segmentation, Motion Estimation, Stabilization, Registration, Categorization	
Application Areas	National Security →	Recreation
End Users	Governments →	Consumer

Camera Networks

- Classical approach
  - One camera wired to one computer
  - Multiple cameras wired to a central processing unit
  - Processing than by a human operator or by a central computer

- Problems
  - Flexibility: Wiring makes it hard to deploy new cameras
  - Robustness: central node failure = entire system failure
  - Scaling: processing and bandwidth requirements do not scale well

Camera Sensor Networks

- Motes
  - Small, wireless devices
  - battery powered
  - limited memory and computing power
- Applications
  - Surveillance
  - Environmental monitoring
  - Smart homes

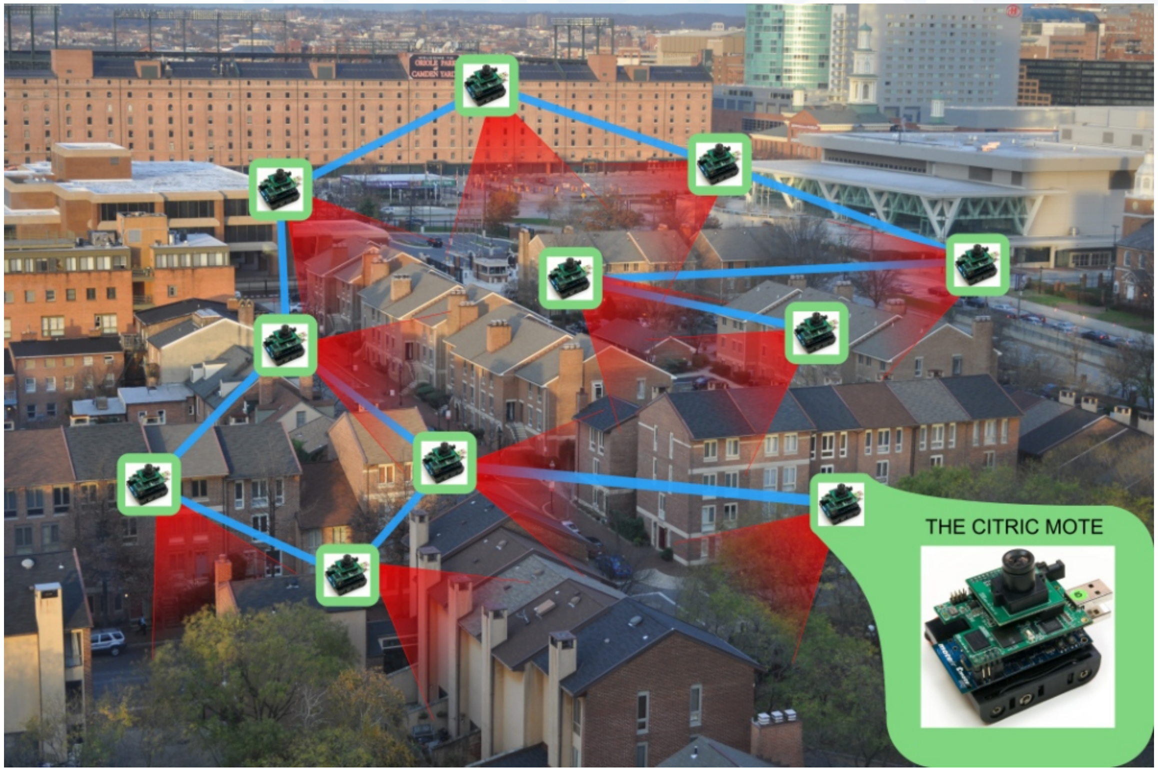

Question

Can we deploy vision algorithms on camera sensor networks?

Challenges to Computer Vision Algorithms

Traditional computer vision algorithms

Existing algorithms are centralized: all images are sent to one node for processing

Computer Vision

- Sensor networks have limited resources
  - Limited processing power and memory
  - Slow wireless channel
  - Nodes can have limited communication range

Challenge

Traditional computer vision algorithms require resources not available in a camera sensor network

Challenges to Sensor Network Algorithms

Traditional distributed algorithms for sensor networks

Existing algorithms have been designed for processing simple scalar measurements

Sensor Networks

- In computer vision applications
  - Measurements (images) are high-dimensional
  - Measurements are corrupted by noise, outliers
  - Estimates are non-Euclidean (e.g. rotations)

Challenge

Traditional sensor network algorithms cannot be directly used for computer vision applications

Toward Distributed Computer Vision Algorithms

- Centralized algorithms
- Considerable complexity

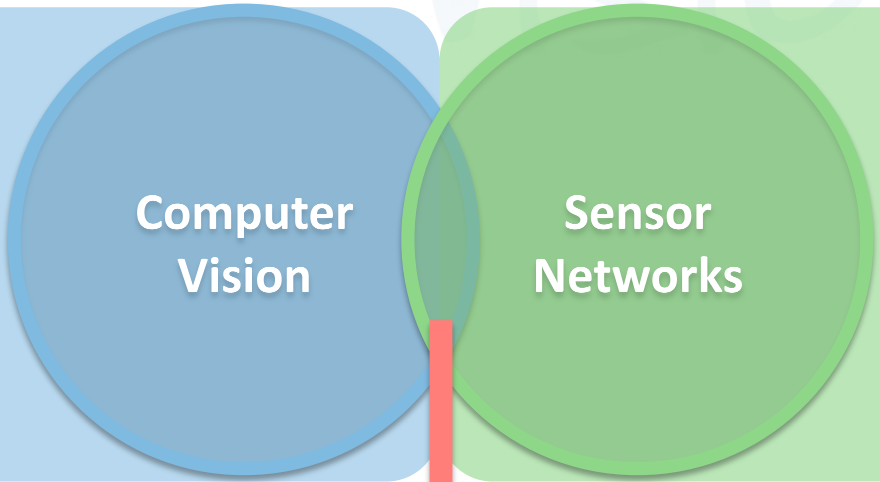

- Distributed algorithms
- Only simple problems

Our goal

- Develop distributed computer vision algorithms that
  - Are efficient: local processing + short communications
  - Converge to the centralized solution
  - Can handle outliers, packet losses, data on manifolds

Consensus on manifolds example

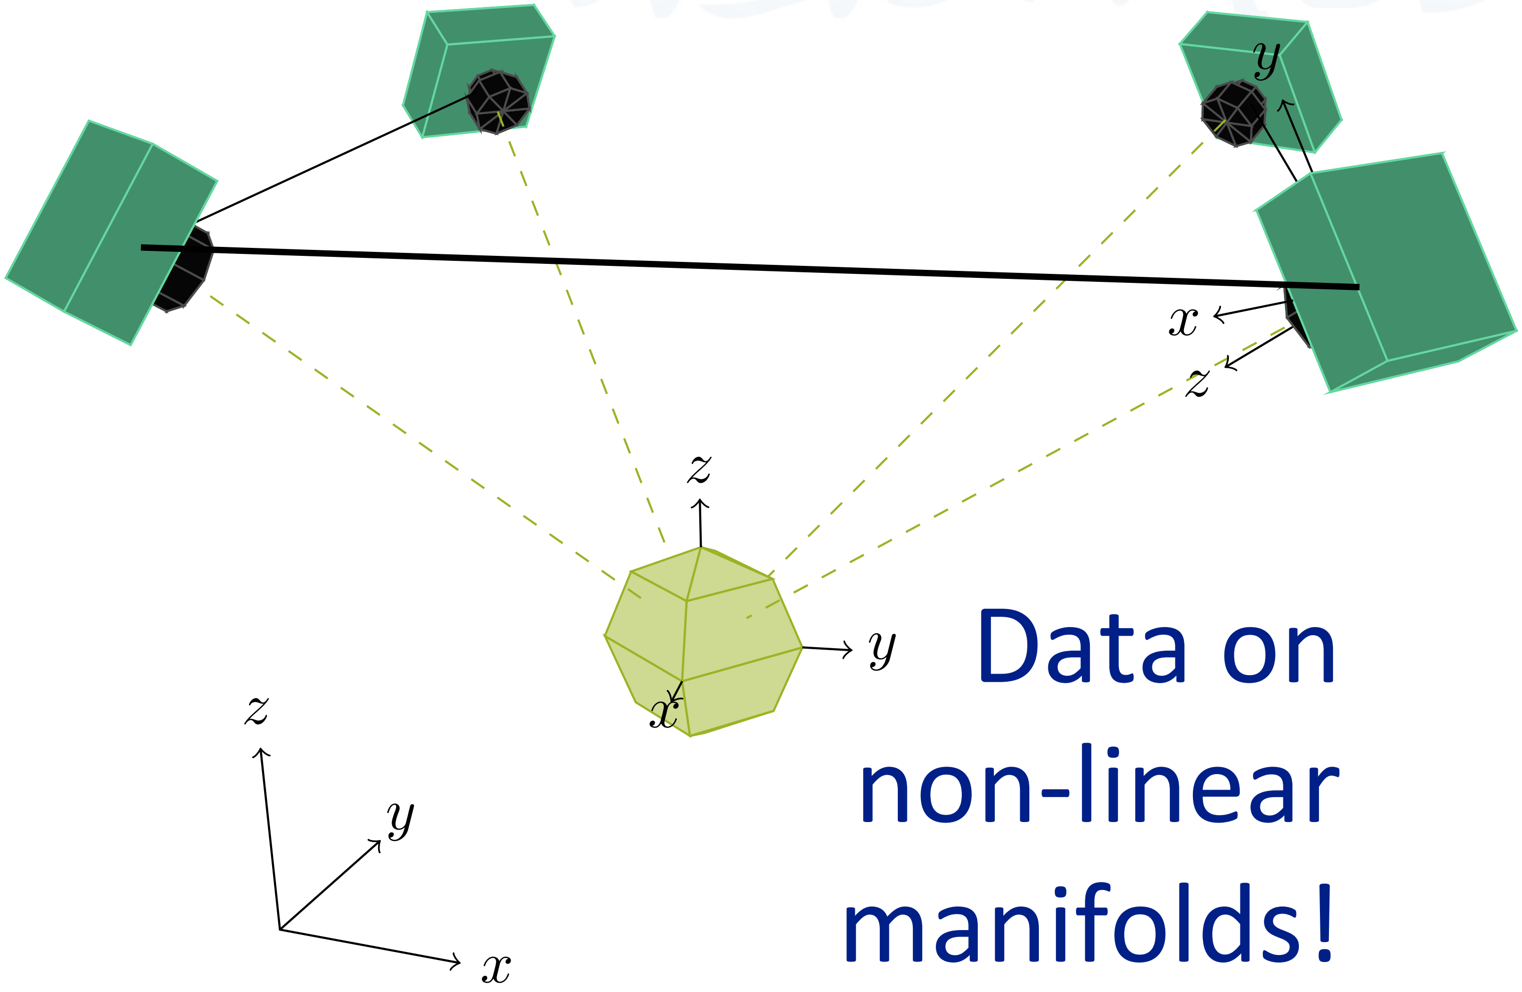

Data on non-linear manifolds!

Previous work

Specific Manifolds

- Sphere [Olfati-Saber 2006]
- N-Torus [Kuramoto model, 1975], [Vicksek model, 1995], [Scardovi, Sepulchre 2007]

Extrinsic approach

- Embedding + Projections [Sarlette, Sepulchre 2009], [Hatanaka, Bullo 2010], [Igarashi, Fujita 2010]

Coordination

- Lie groups [Sarlette, Sepulchre 2010]
- Rigid motions [Sarlette, Sepulchre 2009], [Thunberg, Hu 2011], [Bai, Wen], [Igarashi, Spong, 2012], [Smith, Leonard 2001]

Localization

- Planar case [Piovan, Bullo 2008]
- Centralized [Shirmohammadi, Taylor 2010]

Our work

1. Consensus on **any** Riemannian manifold with bounded curvature

2. Distributed image-based localization from images

Distributed optimization

Outline

Introduction

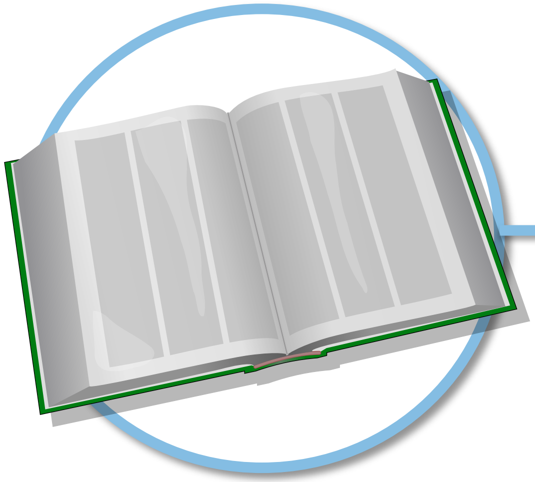

- Consensus in Euclidean spaces
- Riemannian geometry

Riemannian consensus

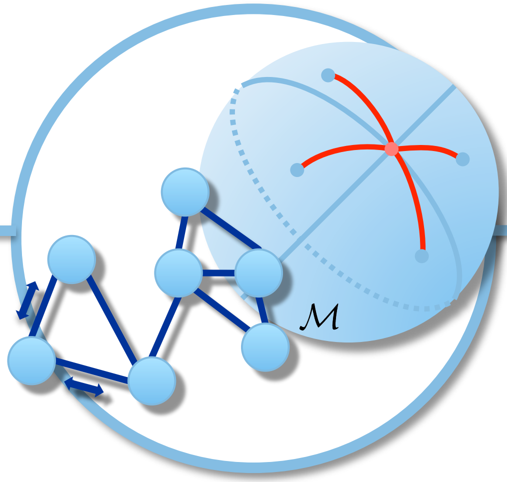

- Choice of step size
- Local convergence results
- Almost-global convergence on SO(3)

Image-based camera network localization

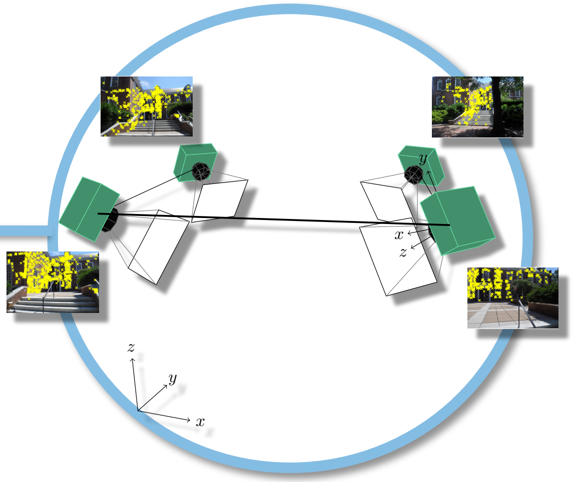

- Setup cost function
- Link with Riemannian Consensus
- Results

JHU vision lab

Review of Euclidean consensus

Center for IMAGING SCIENCE

Notation

$G = (V, E)$ Graph

u_i Measurements

x_i States

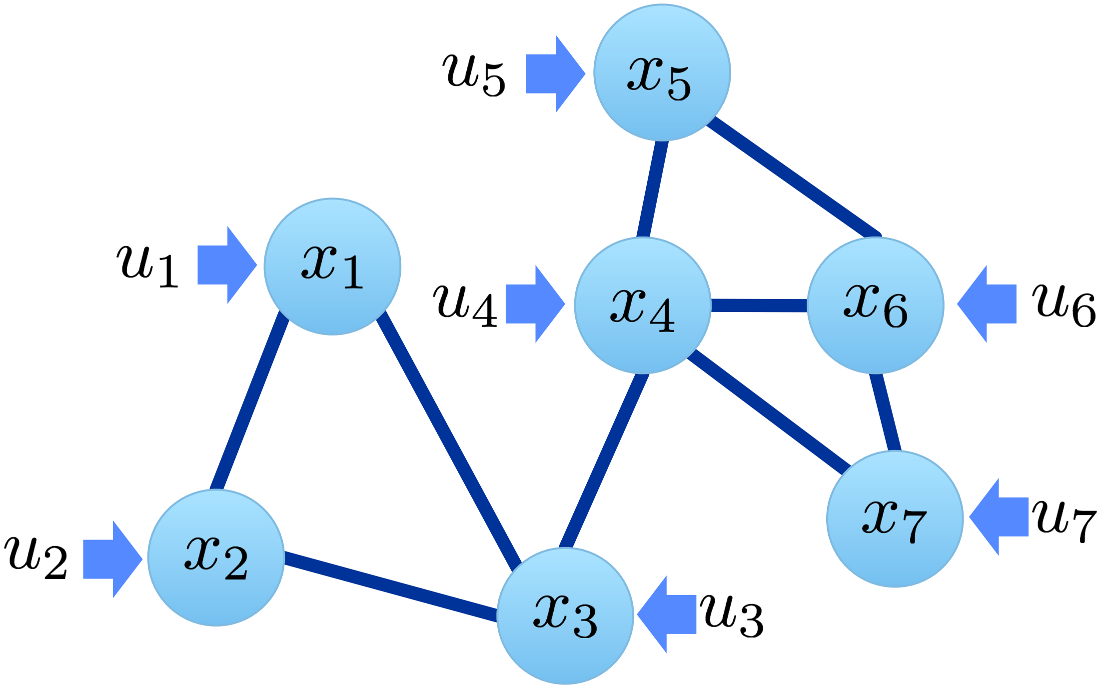

Optimization cost function

$$\varphi(\{x_i\}) = \sum_{(i,j)\in E} \|x_i - x_j\|^2$$

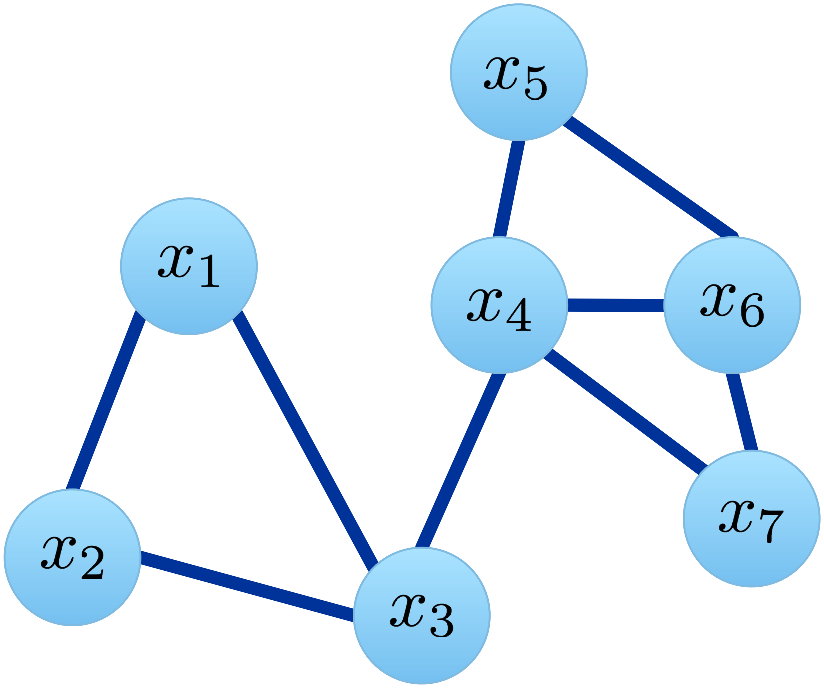

Optimization cost function

$$\varphi(\{x_i\}) = \sum_{(i,j)\in E} \|x_i - x_j\|^2$$

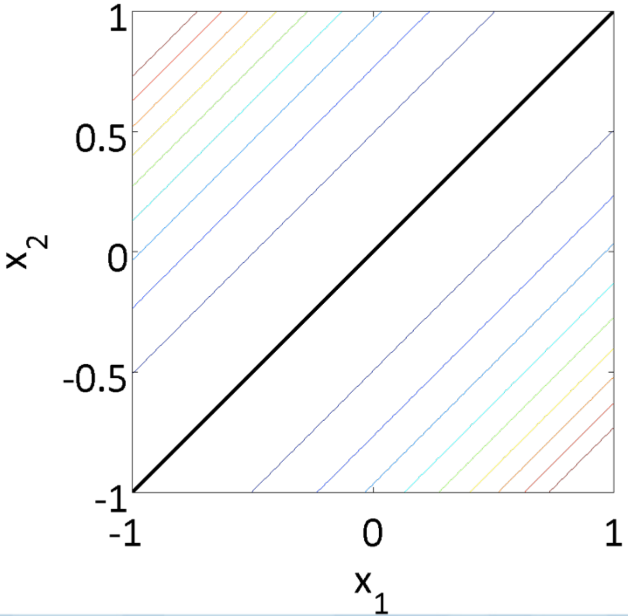

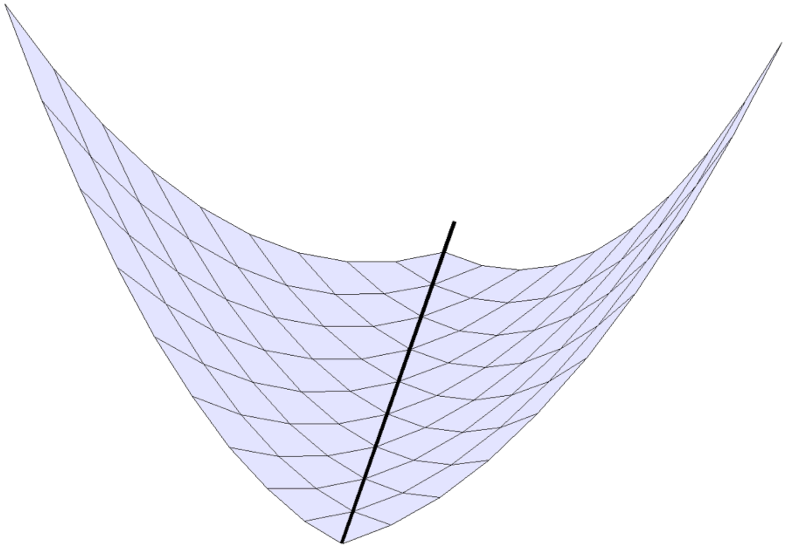

Algorithm = Gradient Descent

$$x_i(0) = u_i$$

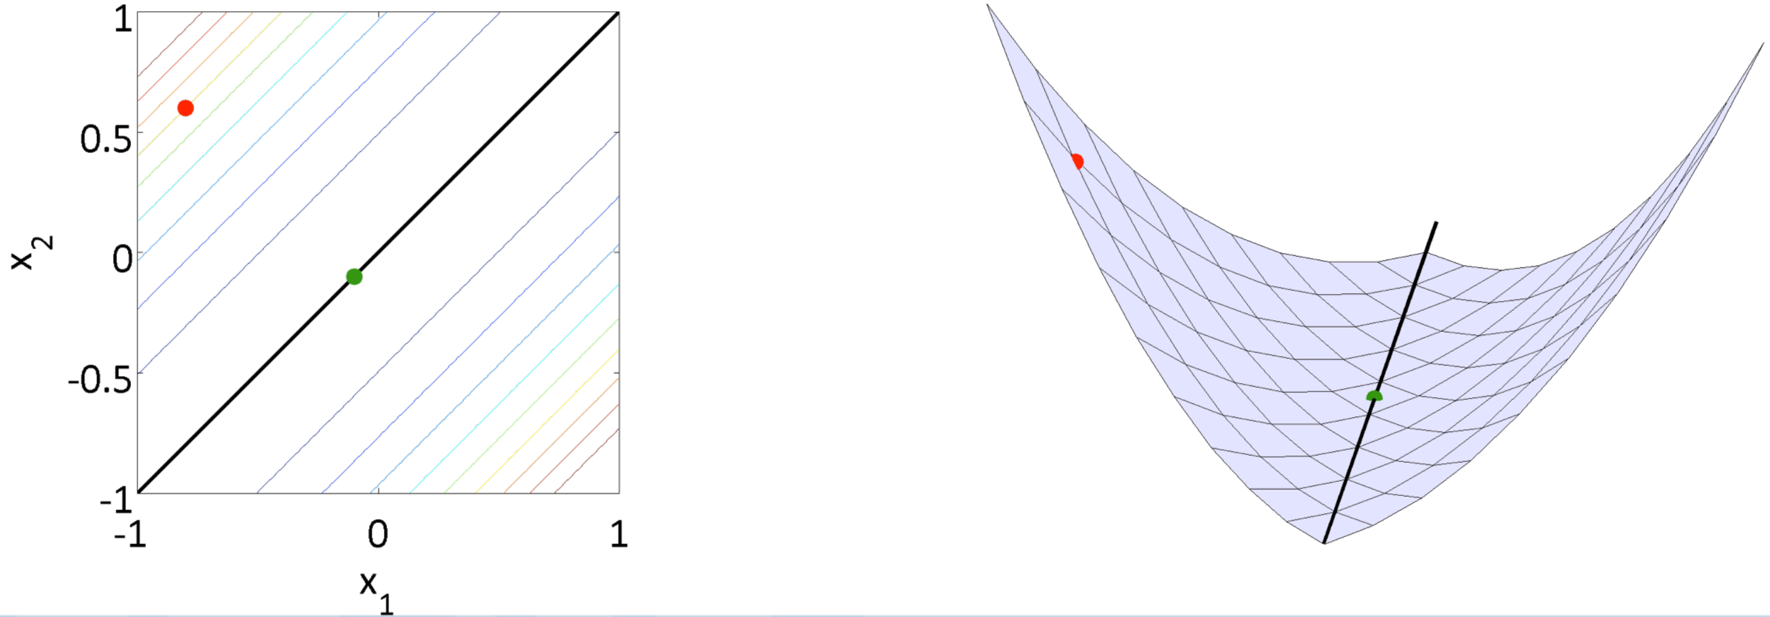

Algorithm = Gradient Descent

$$x_i(0) = u_i$$

$$x_i(t+1) = x_i(t) + \varepsilon \sum_{(i,j)\in E} (x_j(t) - x_i(t))$$

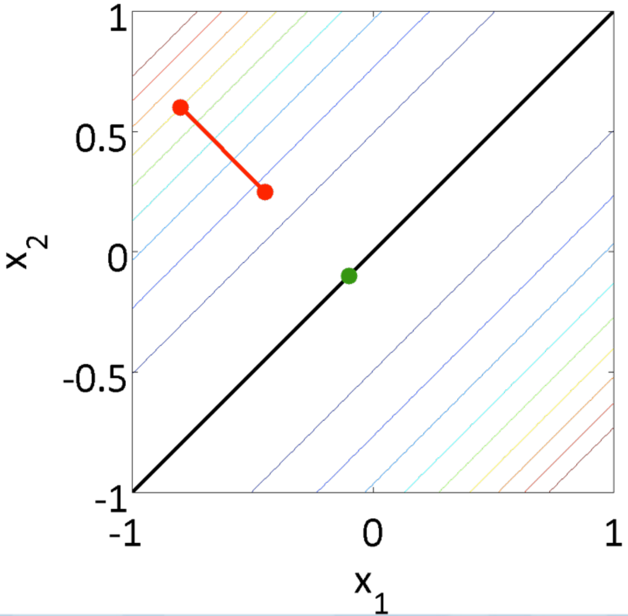

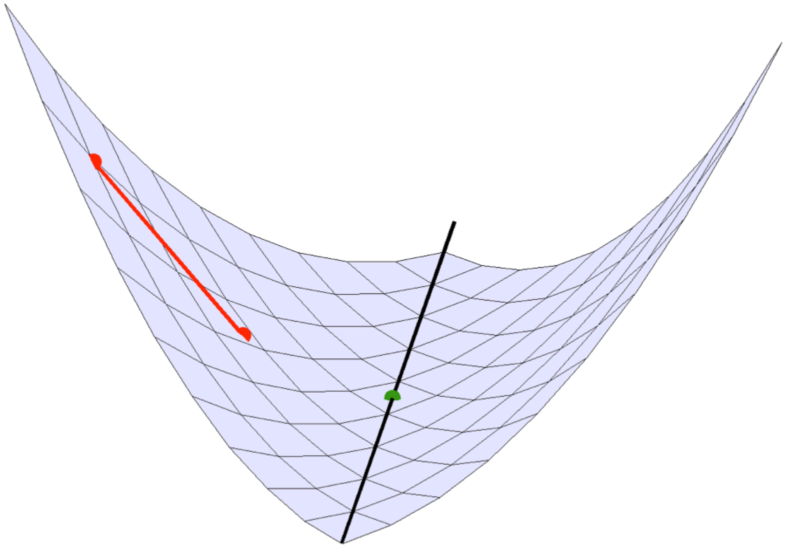

How to choose the step size

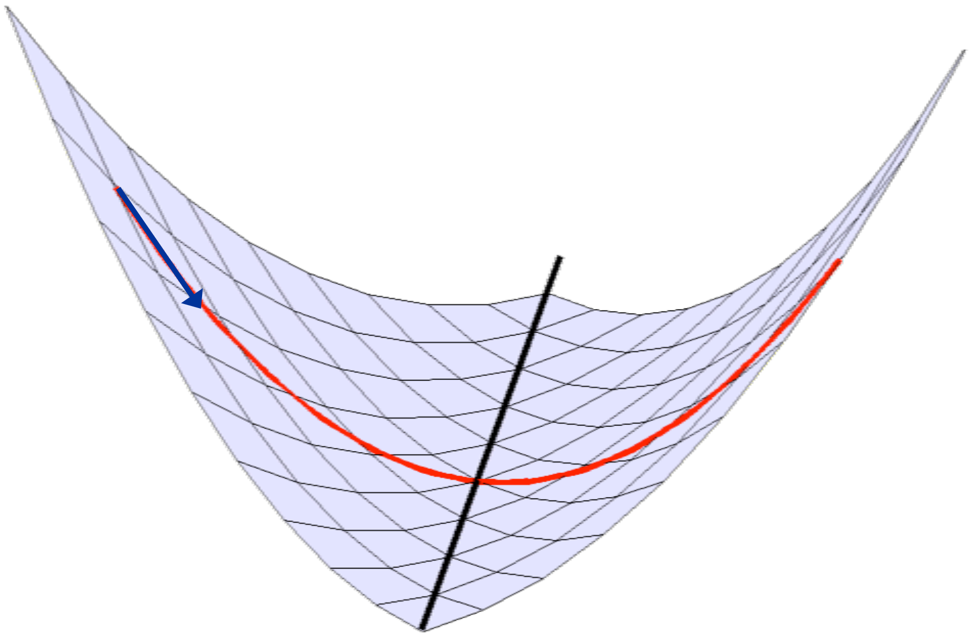

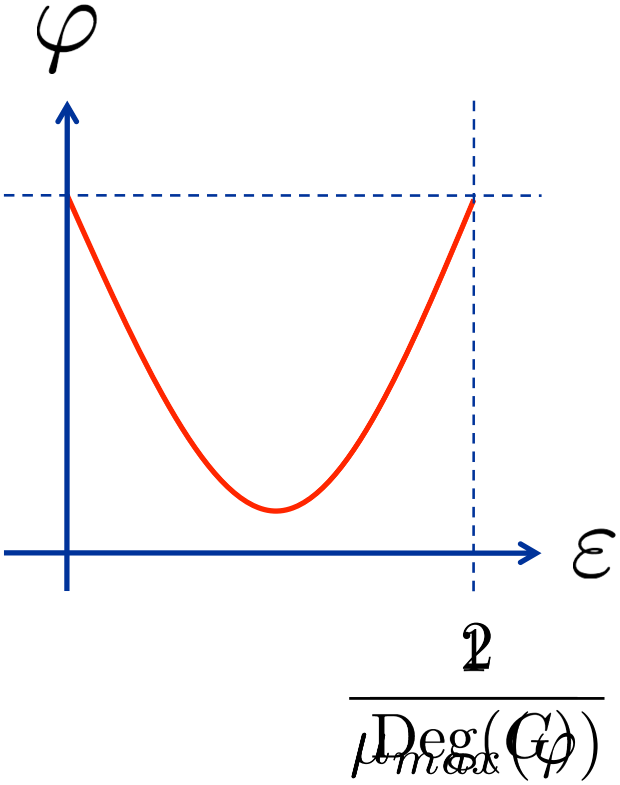

$\mu_{max}(\varphi)$ = bound on max eval of Hessian of φ

$\mathrm{Deg}(G)$ = max # of neighbors

Gradient Descent

$$x_i(0) = u_i$$

$$x_i(t+1) = x_i(t) + \varepsilon \sum_{(i,j)\in E} \left(x_j(t) - x_i(t)\right)$$

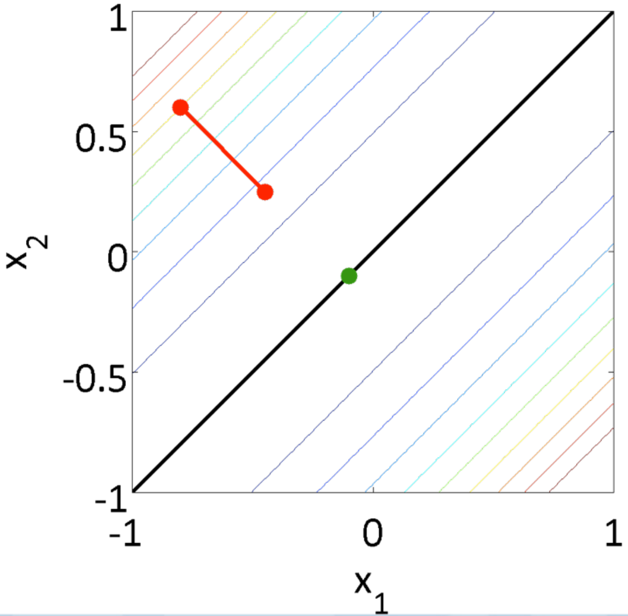

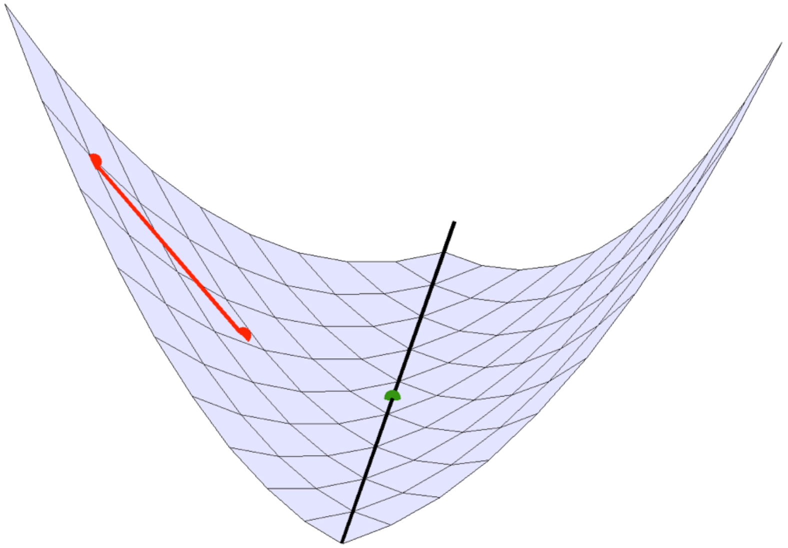

Gradient Descent

$$x_i(0) = u_i$$

$$x_i(t+1) = x_i(t) + \varepsilon \sum_{(i,j)\in E} \left(x_j(t) - x_i(t)\right)$$

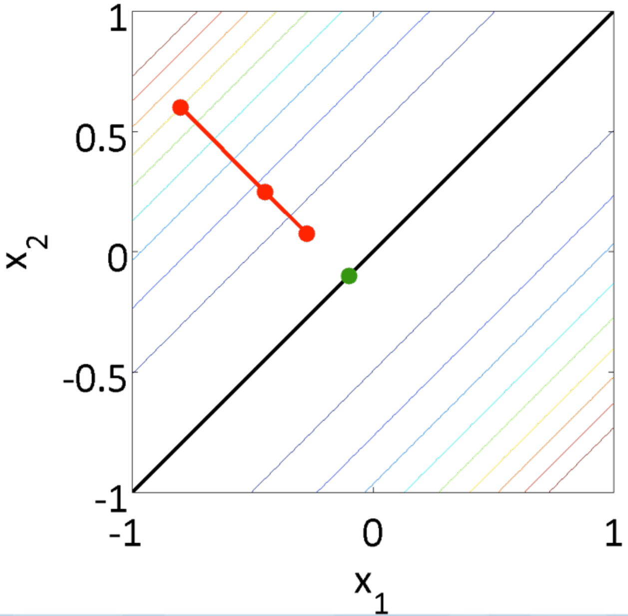

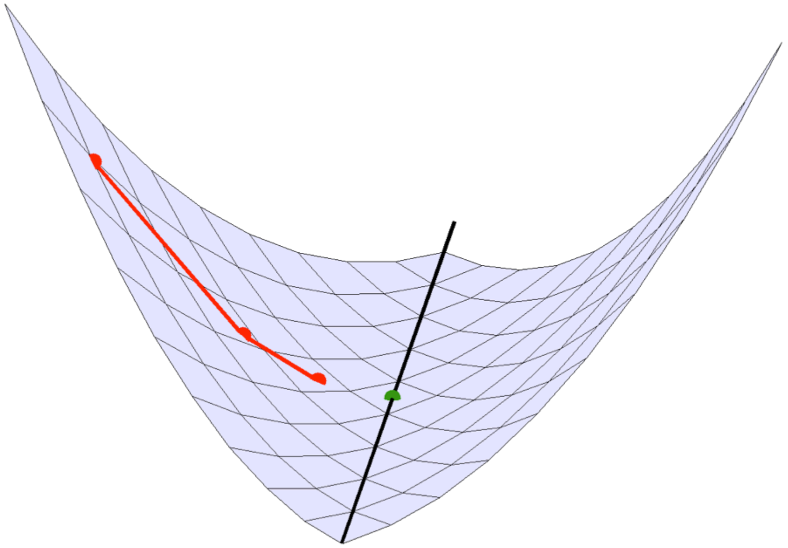

Gradient Descent

$$x_i(0) = u_i$$

$$x_i(t+1) = x_i(t) + \varepsilon \sum_{(i,j)\in E} \left(x_j(t) - x_i(t)\right)$$

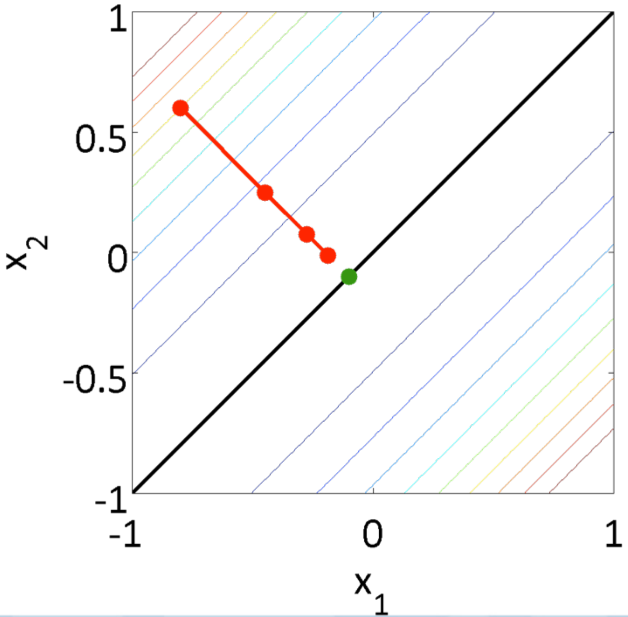

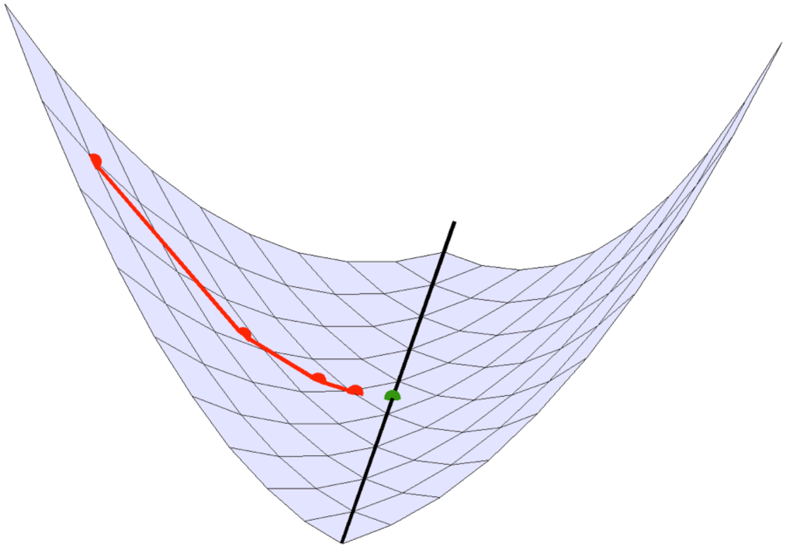

Gradient Descent

$$x_i(0) = u_i$$

$$x_i(t+1) = x_i(t) + \varepsilon \sum_{(i,j)\in E} \left(x_j(t) - x_i(t)\right)$$

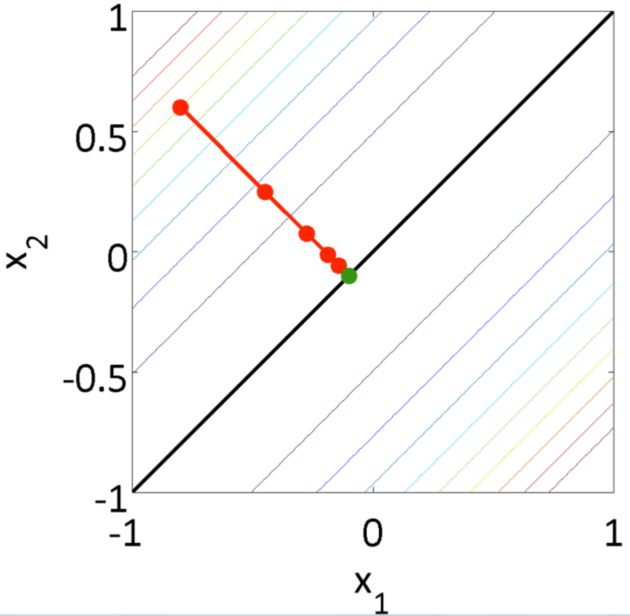

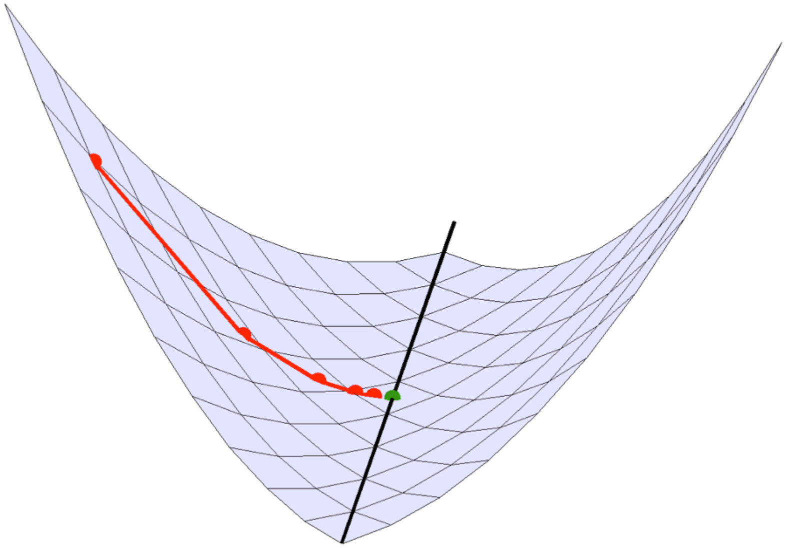

Gradient Descent

$$x_i(0) = u_i$$

$$x_i(t+1) = x_i(t) + \varepsilon \sum_{(i,j)\in E} \left(x_j(t) - x_i(t)\right)$$

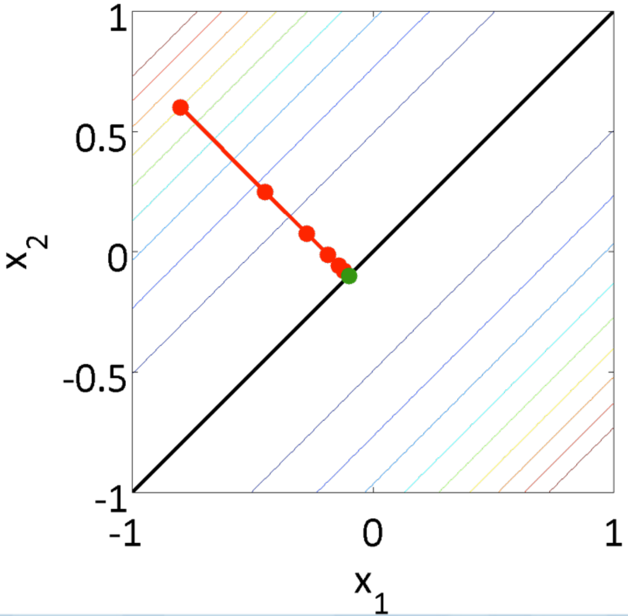

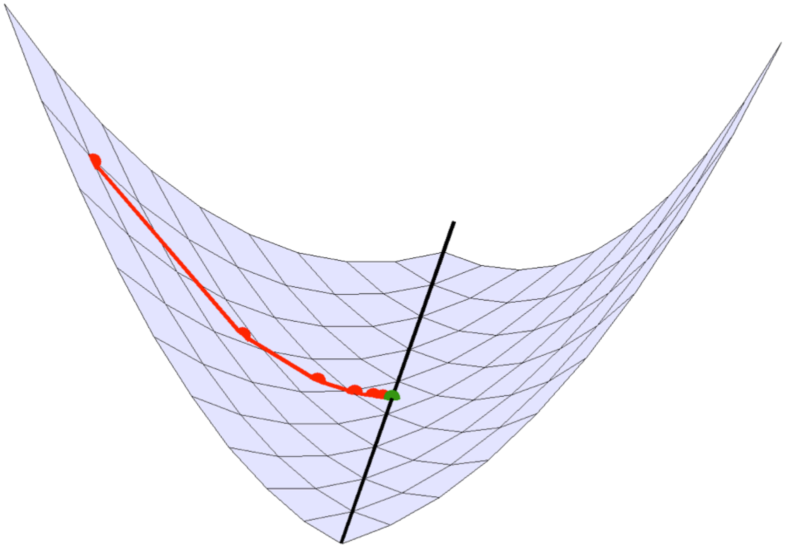

Gradient Descent

$$x_i(0) = u_i$$

$$x_i(t+1) = x_i(t) + \varepsilon \sum_{(i,j)\in E} \big(x_j(t) - x_i(t)\big)$$

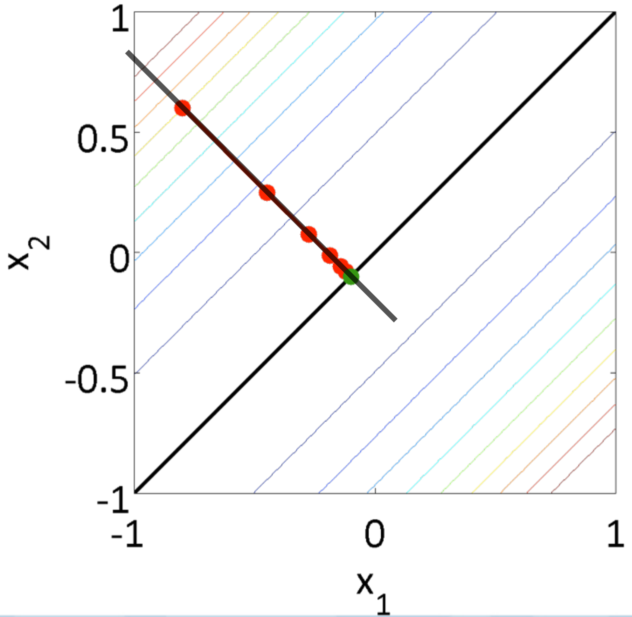

$$\text{mean}\{x_i(t)\} = \text{mean}\{x_i(t+1)\}$$

Convergence to the mean

$$x_1(t_\infty) = \ldots = x_2(t_\infty) = \bar{u} = \frac{1}{N}\sum_i u_i$$

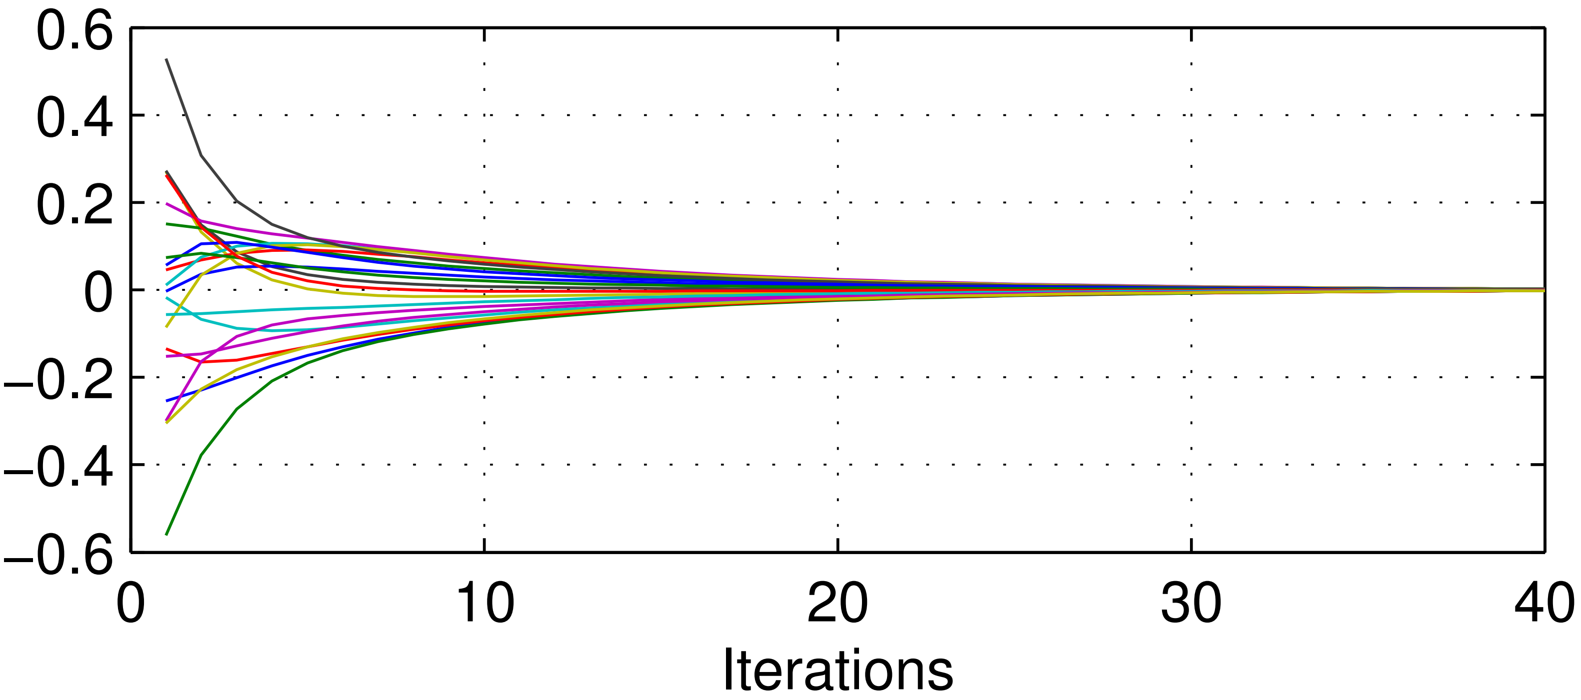

Review of Riemannian geometry

Notation

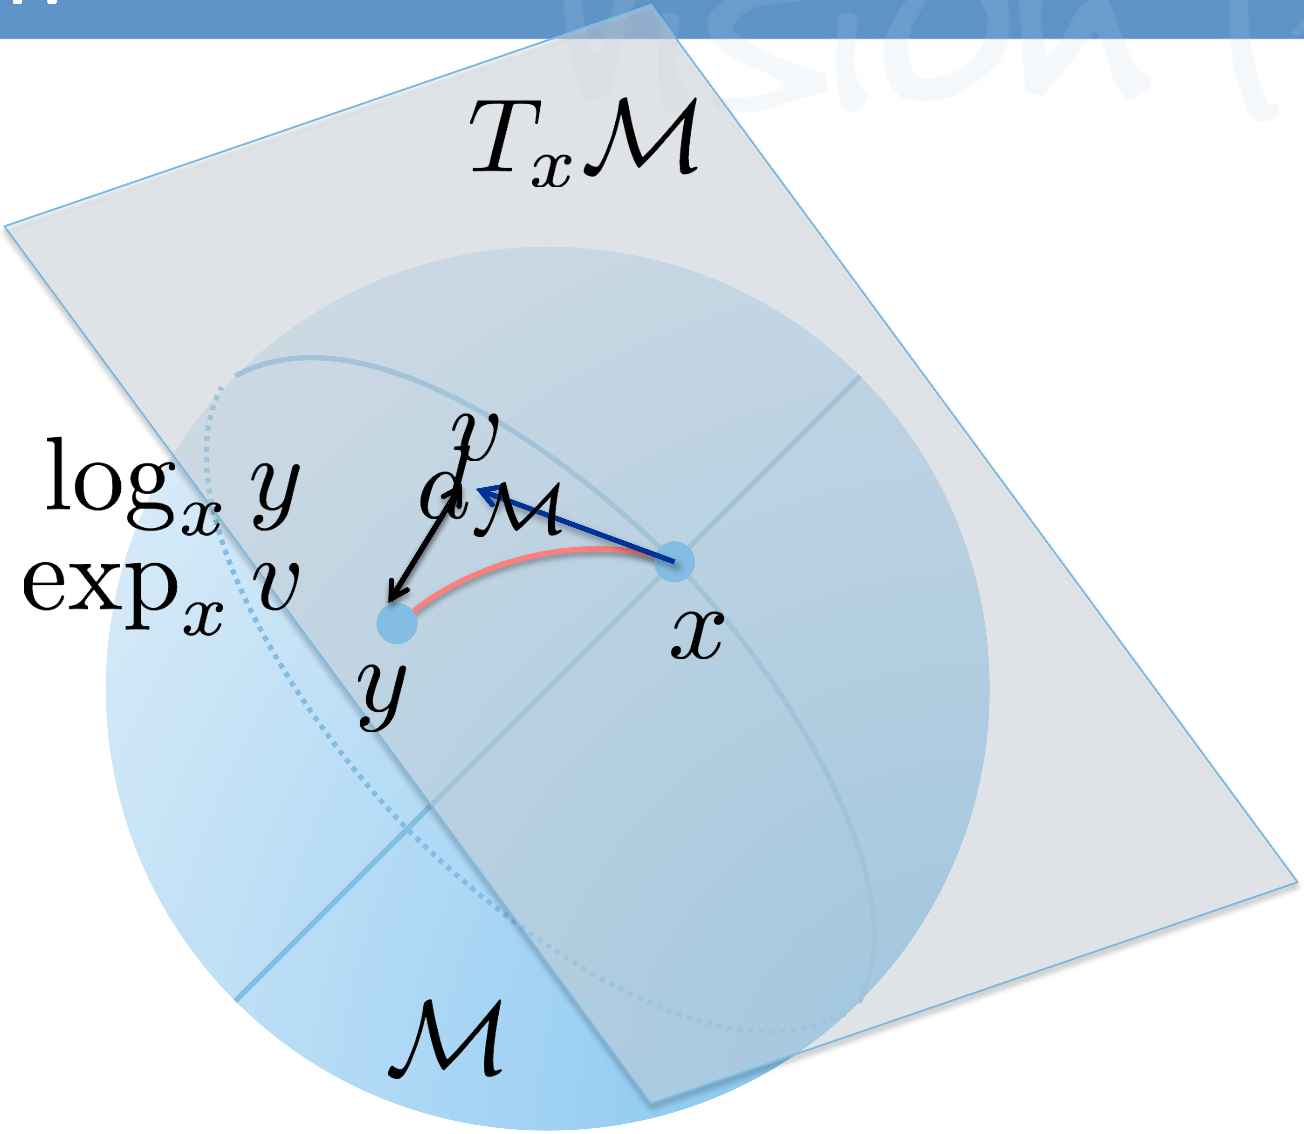

$$\operatorname{grad}_x \frac{1}{2} d^2_{\mathcal{M}}(x, y) = -\log_x y$$

\ma
^2(

JHU vision lab

Riemannian Consensus

Center for IMAGING SCIENCE

Notation

$G = (V, E)$ Graph

u_i Measurements

x_i States

On manifold $\mathcal{M}$

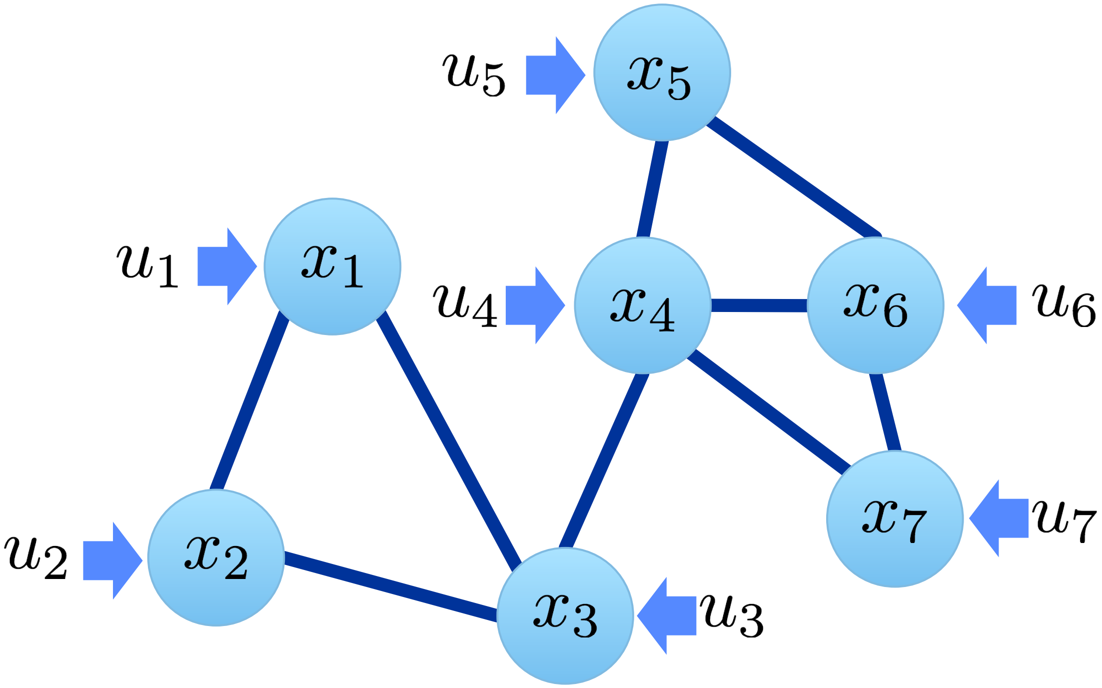

Optimization cost function

$$\varphi(\{x_i\}) = \sum_{(i,j)\in E} d^2_{\mathcal{M}}(x_i, x_j)$$

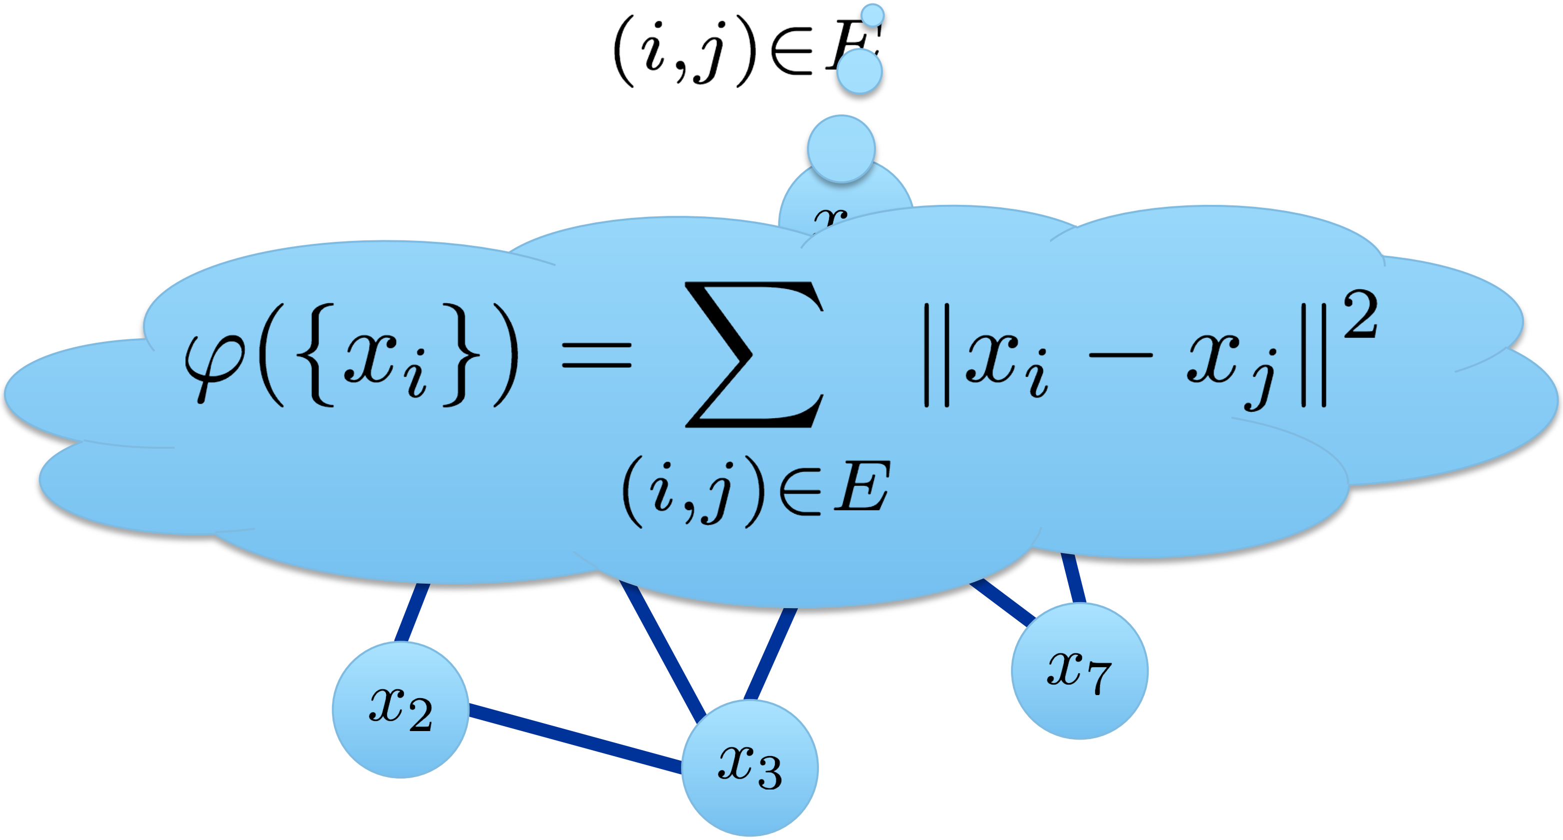

$$\varphi(\{x_i\}) = \sum_{(i,j)\in E} \|x_i - x_j\|^2$$

Optimization cost function

$$\varphi(\{x_i\}) = \sum_{(i,j)\in E} d^2_{\mathcal{M}}(x_i, x_j)$$

Not convex!

Consensus configurations are global minima

Optimization cost function

$$\varphi(\{x_i\}) = \sum_{(i,j)\in E} d^2_{\mathcal{M}}(x_i, x_j)$$

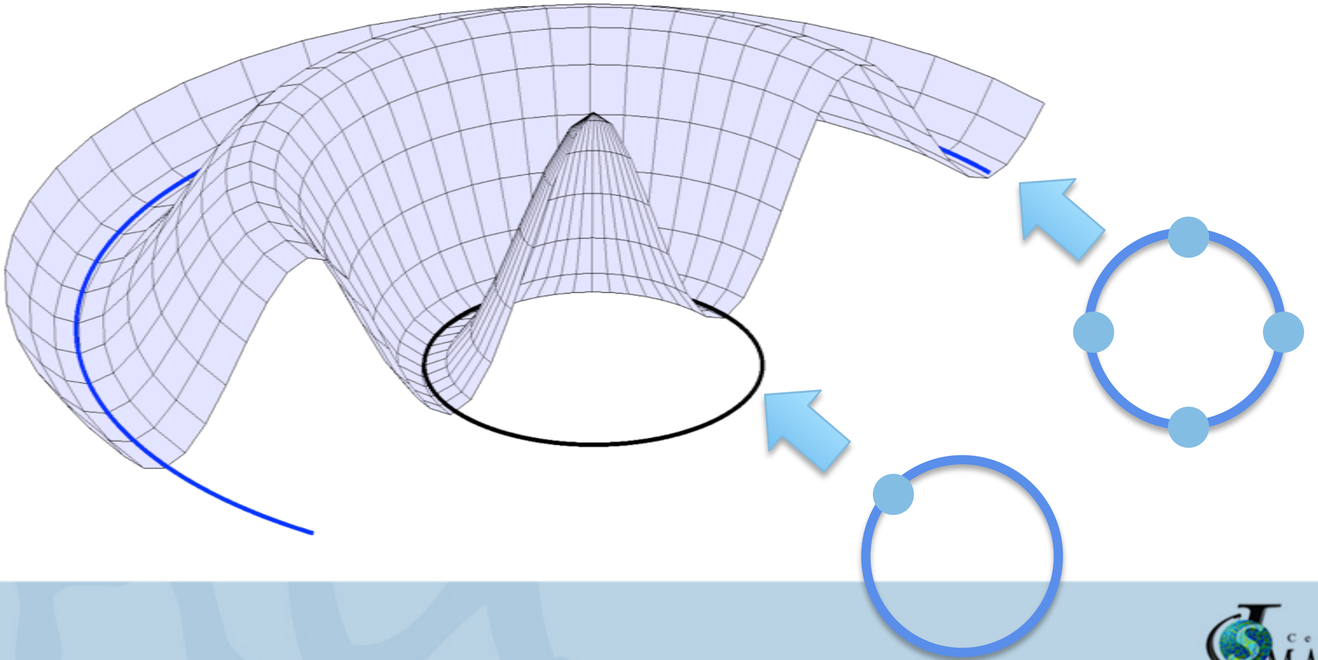

Gradient descent

$$x_i(0) = u_i$$

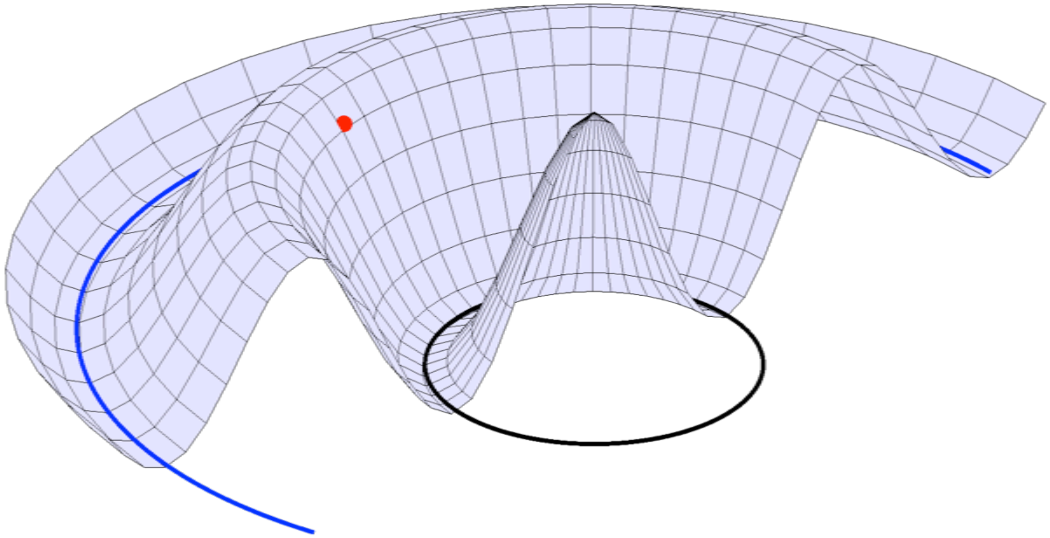

Gradient descent

$$x_i(0) = u_i$$

$$x_i(t+1) = \exp_{x_i(t)} \sum_{(i,j)\in E} \left(\varepsilon \log_{x_i(t)} x_j(t)\right)$$

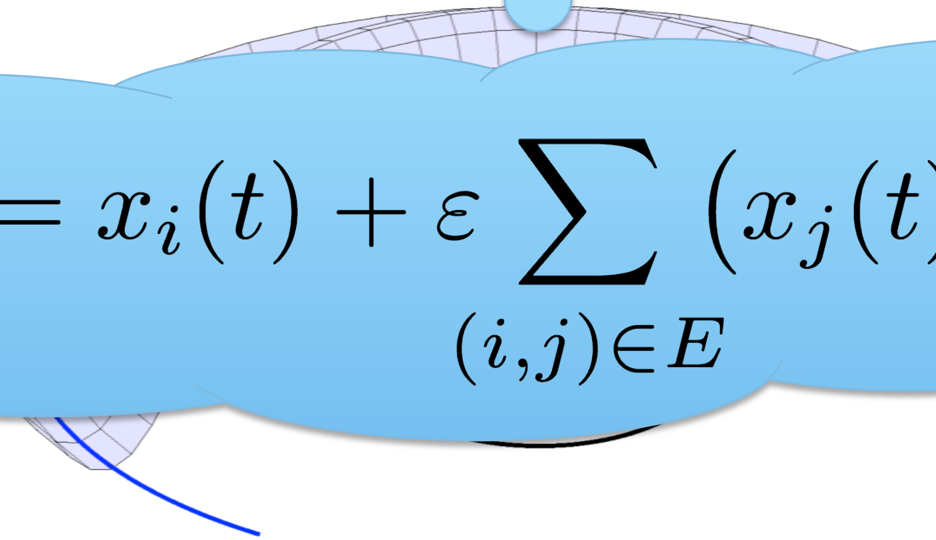

$$x_i(t+1) = x_i(t) + \varepsilon \sum_{(i,j)\in E} \left(x_j(t) - x_i(t)\right)$$

How to choose the step size

Theorem

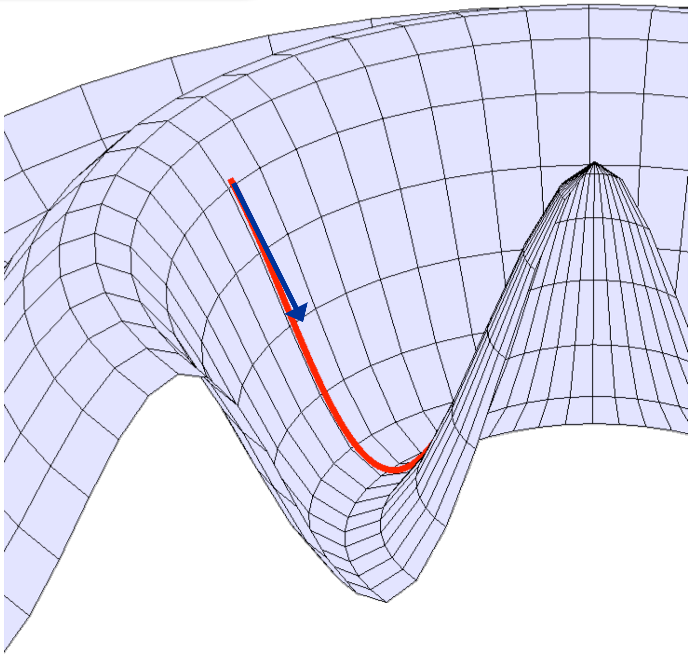

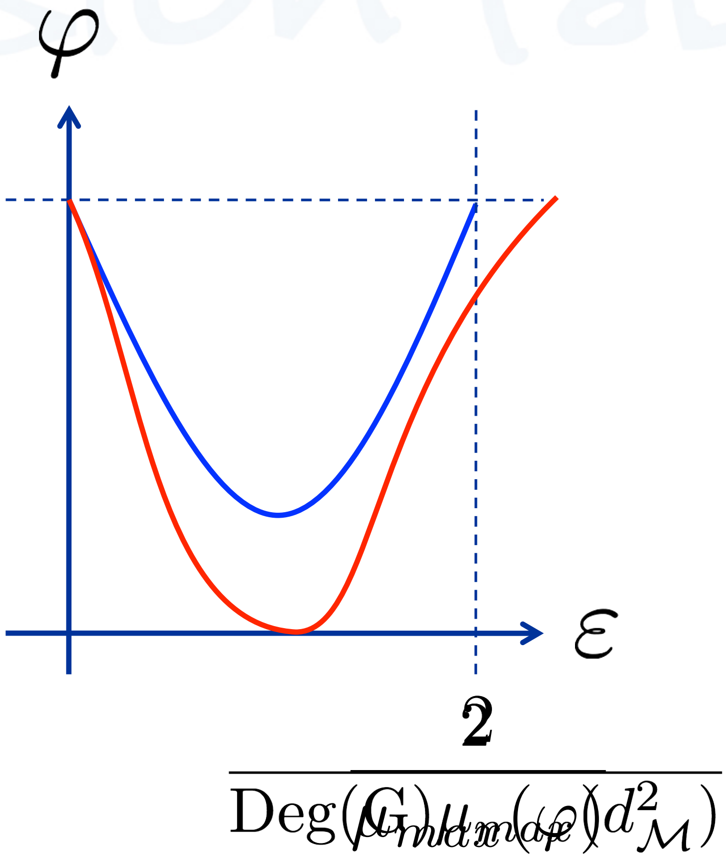

$\mathrm{Deg}(G)$ = max # of neighbors

$\mu_{max}(d^2_{\mathcal{M}})$ = bound on max eval of Hessian of $d^2_{\mathcal{M}}$

Tron, Afsari, Vidal - “Riemannian Consensus for Manifolds with Bounded Curvature”
to appear in Transactions on Automatic Control, 2012

Spaces of non-negative constant curvature

$$\mu_{max}(d^2_{\mathcal{M}}) = 2$$

$$\varepsilon < \frac{1}{\mathrm{Deg(G)}}$$

All other spaces

$\mu_{max}(d^2_{\mathcal{M}})$ might depend on max distance between neighboring states

Gradient Descent

$$x_i(0) = u_i$$

$$x_i(t+1) = \exp_{x_i(t)} \sum_{(i,j)\in E} \left(\varepsilon \log_{x_i(t)} x_j(t)\right)$$

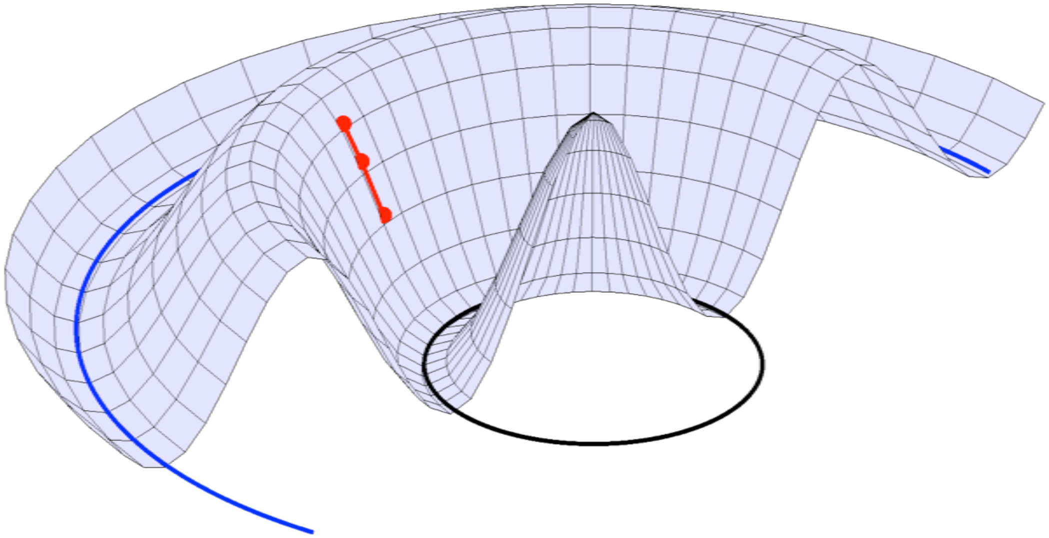

Gradient Descent

$$x_i(0) = u_i$$

$$x_i(t+1) = \exp_{x_i(t)} \sum_{(i,j)\in E} \left(\varepsilon \log_{x_i(t)} x_j(t)\right)$$

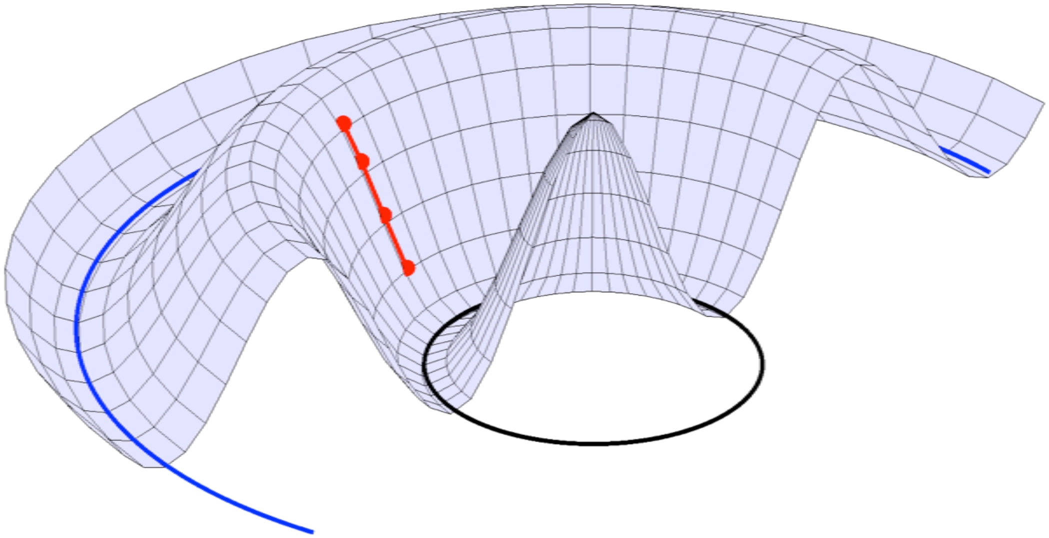

Gradient Descent

$$x_i(0) = u_i$$

$$x_i(t+1) = \exp_{x_i(t)} \sum_{(i,j)\in E} \left(\varepsilon \log_{x_i(t)} x_j(t)\right)$$

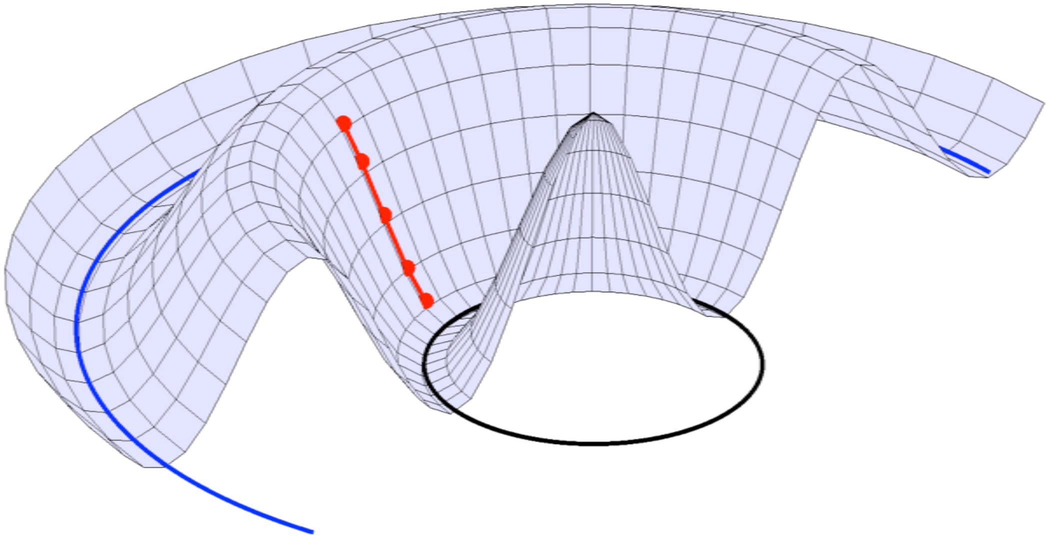

Gradient Descent

$$x_i(0) = u_i$$

$$x_i(t+1) = \exp_{x_i(t)} \sum_{(i,j)\in E} \left(\varepsilon \log_{x_i(t)} x_j(t)\right)$$

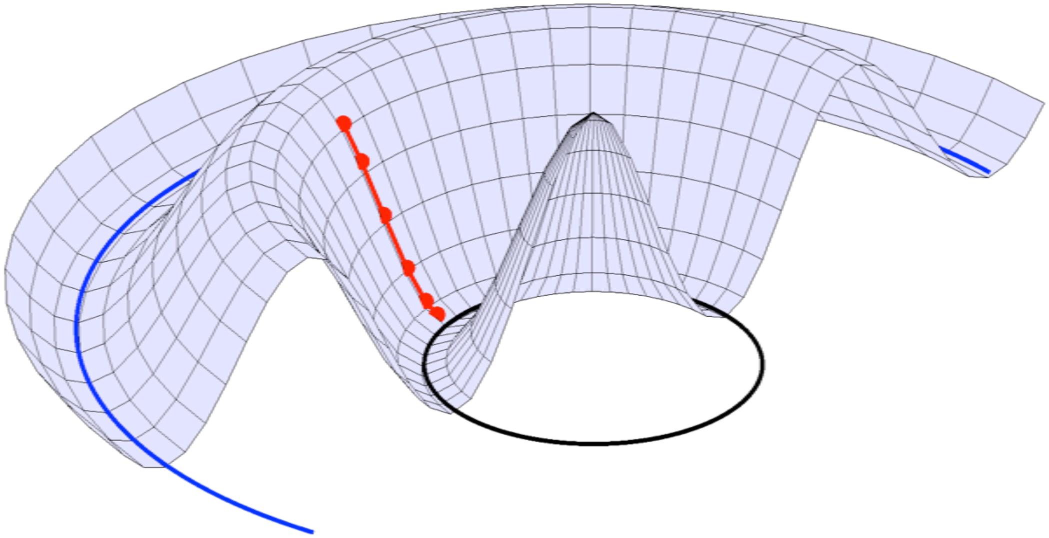

Convergence

Theorem

- Measurements not too dispersed
- Step size small enough

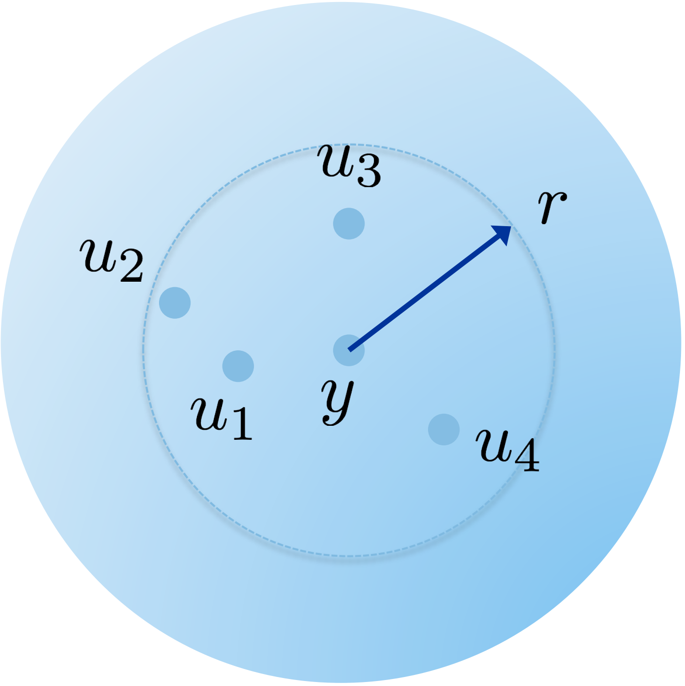

Convergence to consensus configurations

Tron, Afsari, Vidal - “Riemannian Consensus for Manifolds with Bounded Curvature” to appear in Transactions on Automatic Control, 2012

1. Given a set containing all global minimizers

$$S = \{(x_1, \ldots, x_N) \in \mathcal{M}^N : \exists x_0 \in \mathcal{M} \quad \text{s.t.} \quad \max_i d_{\mathcal{M}}(x_i, x_0) < r^*\}$$

$$r^* = \frac{1}{2} \min \left\{ \text{inj } \mathcal{M}, \frac{\pi}{\sqrt{\Delta}} \right\}$$

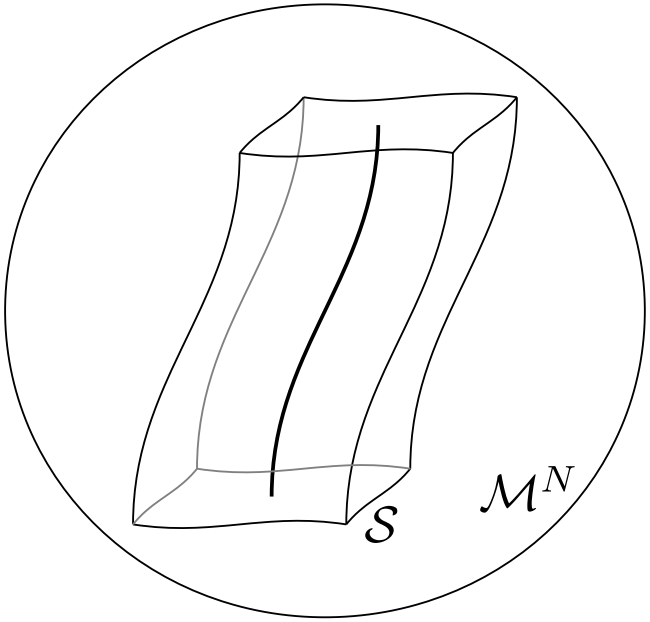

2. Guarantee iterates do not leave this set

Tron, Afsari, Vidal - “Average Consensus on Riemannian Manifolds with Bounded Curvature”
IEEE Conference on Decision and Control 2011

Convergence

Theorem

Spaces of constant, non-negative curvature

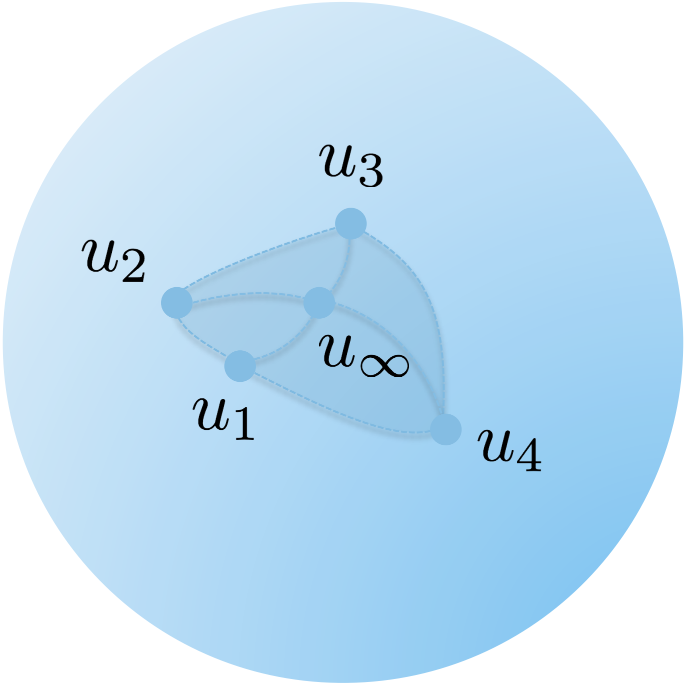

- Converge to a point inside convex hull
- Point need not be the global Frechet mean

Tron, Afsari, Vidal - “Riemannian Consensus for Manifolds with Bounded Curvature” to appear in Transactions on Automatic Control, 2012

Convergence

- General result
  - If initial measurements are inside Sconv
  - And step size is small enough
  - Then, Riemannian Consensus converges to a set in Sconv

- Constant non-negative curvature
  - Convergence set is enlarged to Sr
  - The convergence result is to a single consensus state instead of to a set
  - The consensus state is shown to lie in the convex hull of the initial measurements

Experiments – Sphere(7)

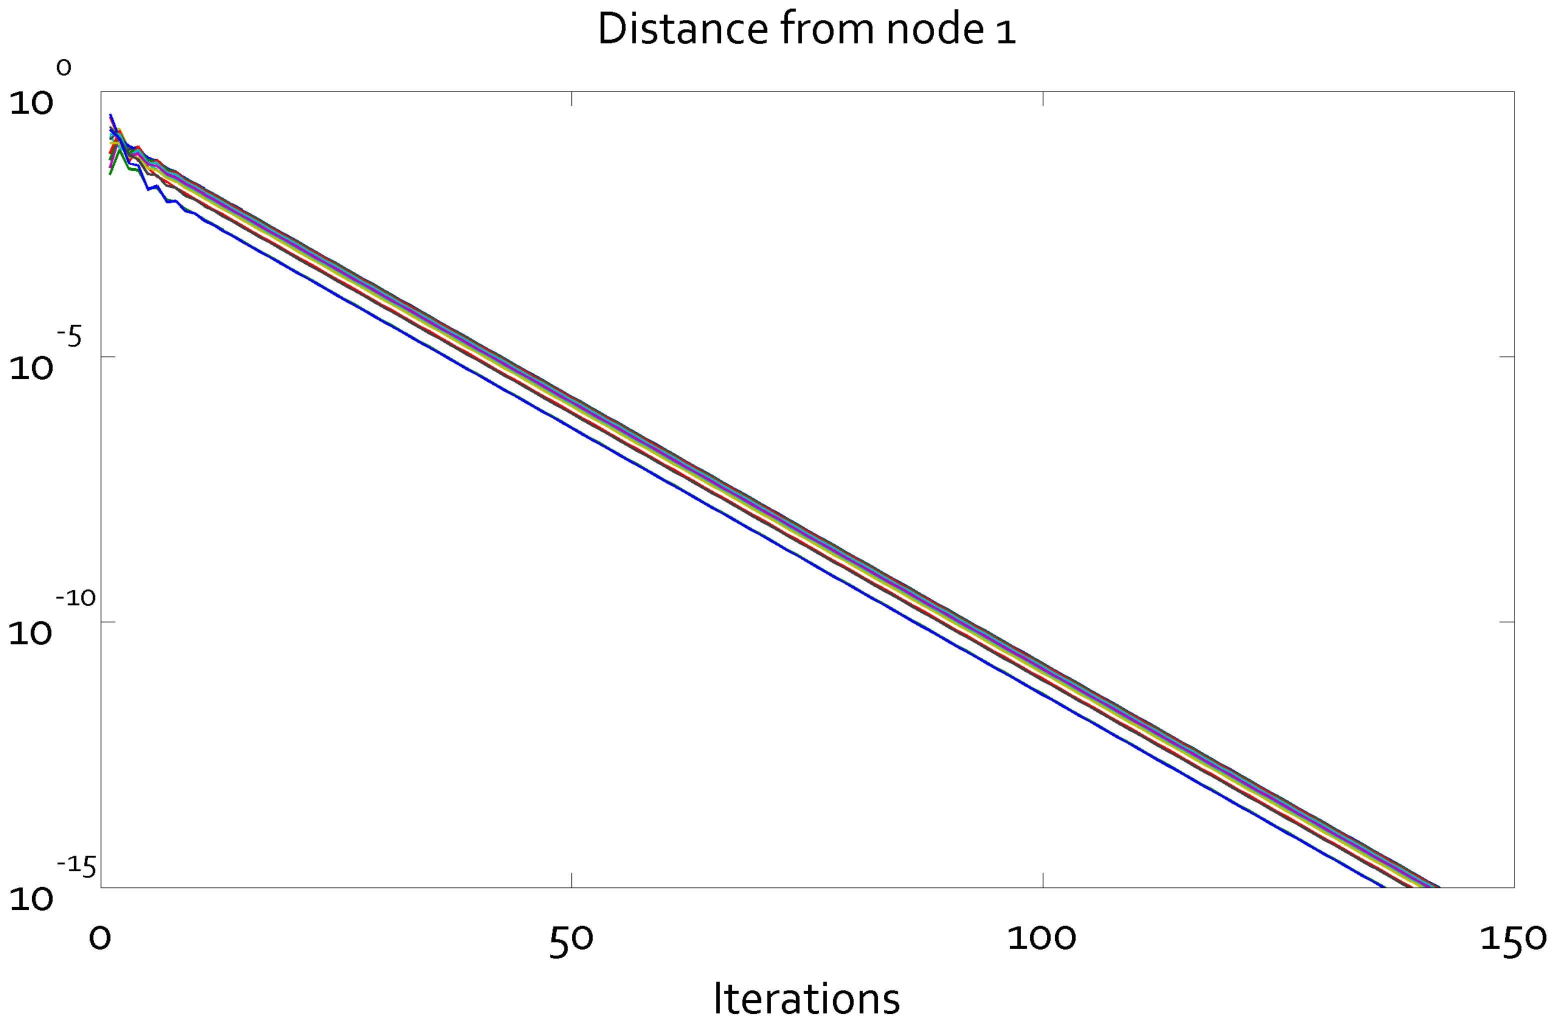

N=15 nodes, 4-regular graph

Experiments – SO(7)

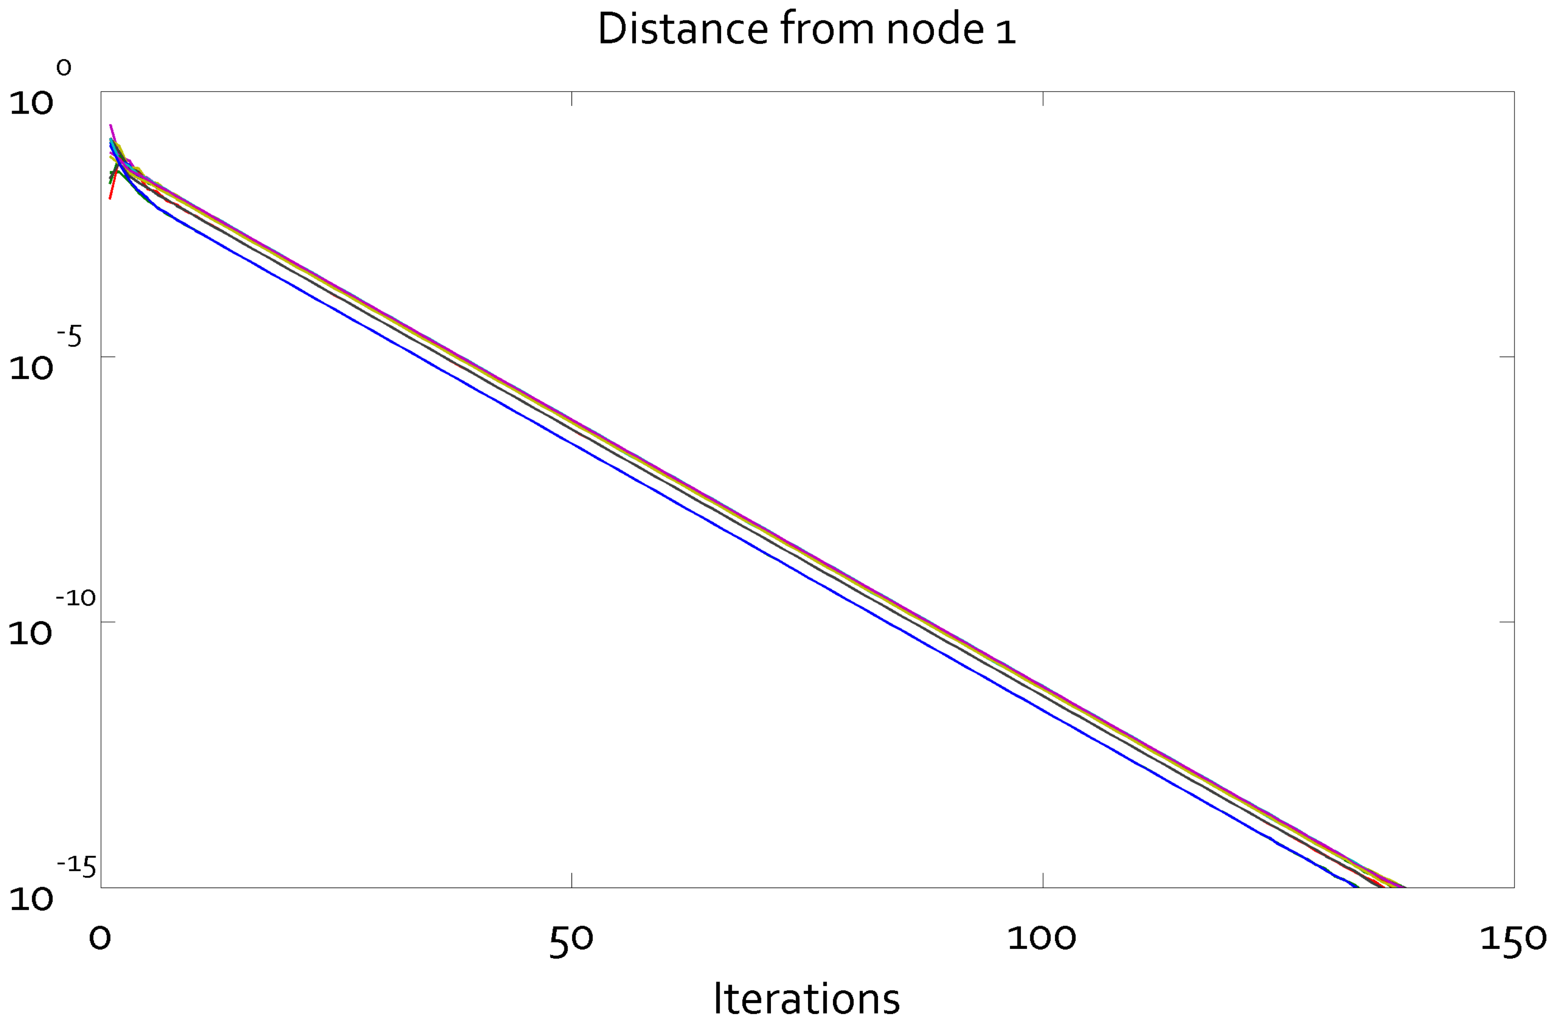

N=15 nodes, 4-regular graph

Experiments – SO(3) – Closed geodesic

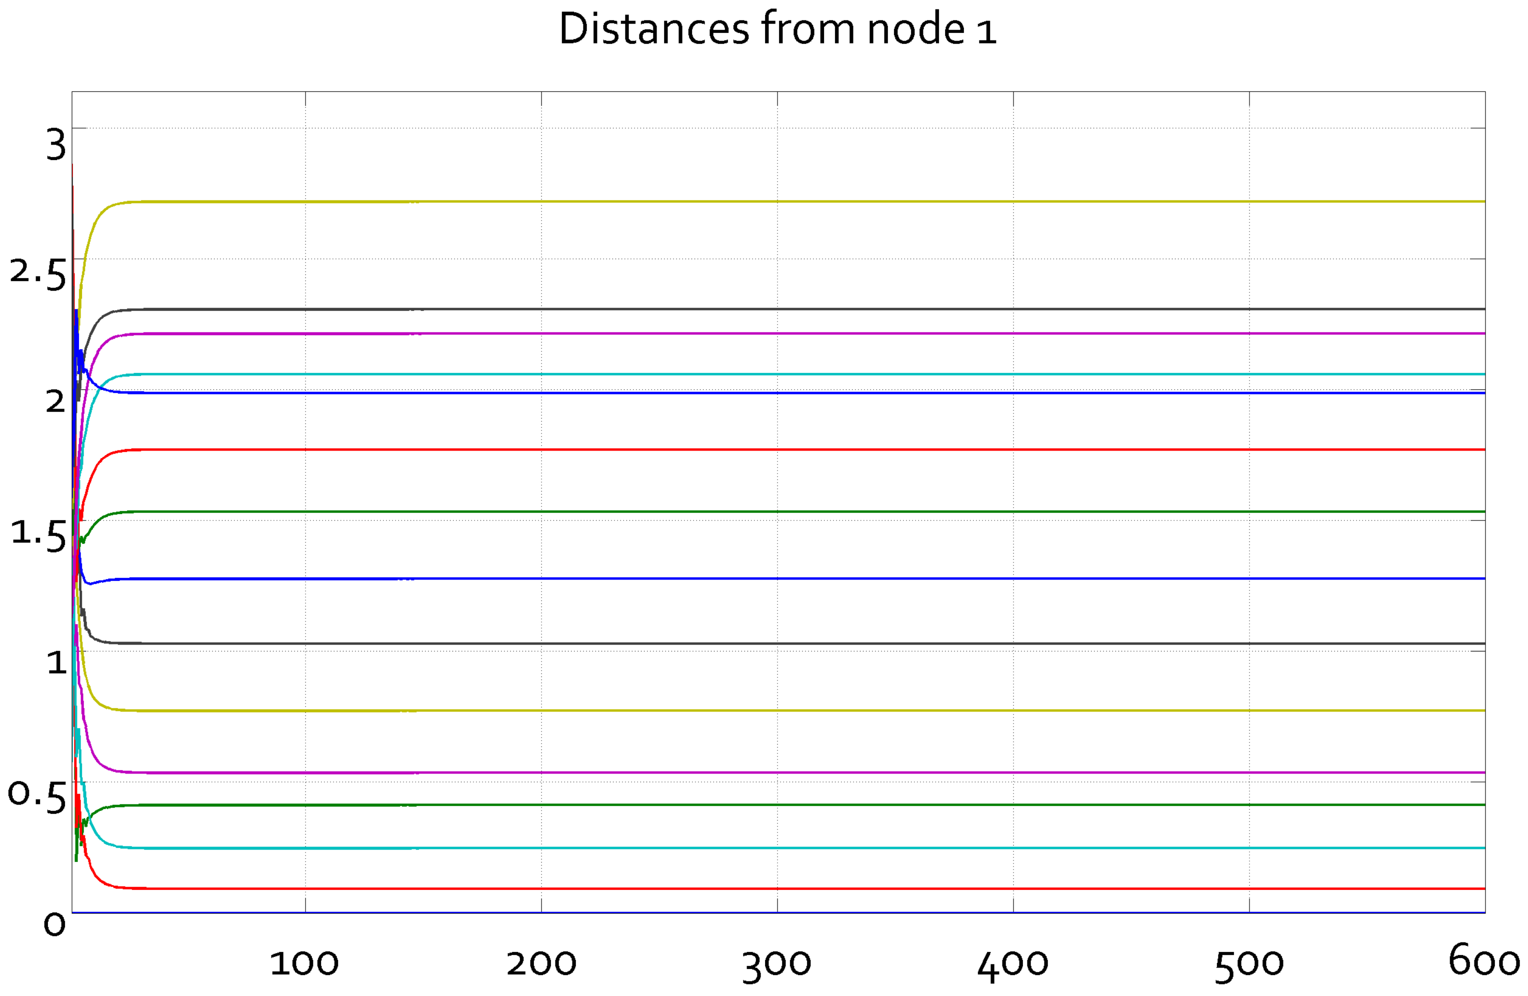

N=15 nodes, 4-regular graph, closed geodesic

$$\varphi(\{x_i\}) = \sum_{(i,j)\in E} d^2_{\mathcal{M}}(x_i, x_j)$$

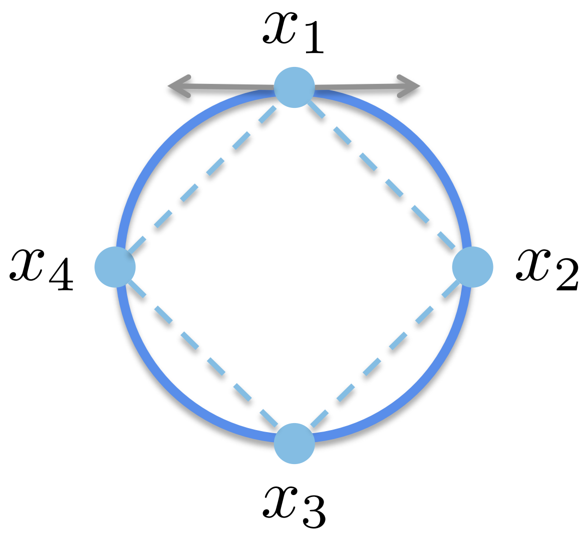

Local minimum

$$\varphi(\{x_i\}) = \sum_{(i,j)\in E} d^2_{\mathcal{M}}(x_i, x_j)$$

$$d^2_{\mathcal{M}} =$$

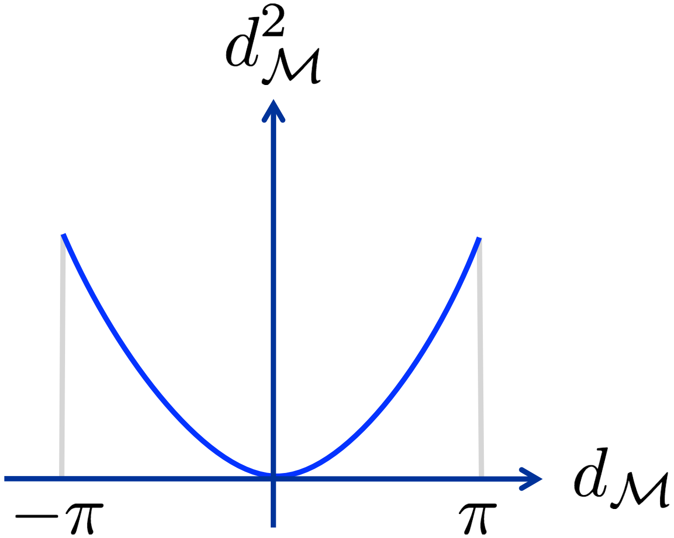

$$\varphi(\{x_i\}) = \sum_{(i,j)\in E} \xcancel{d^2_{\mathcal{M}}}(x_i, x_j)$$

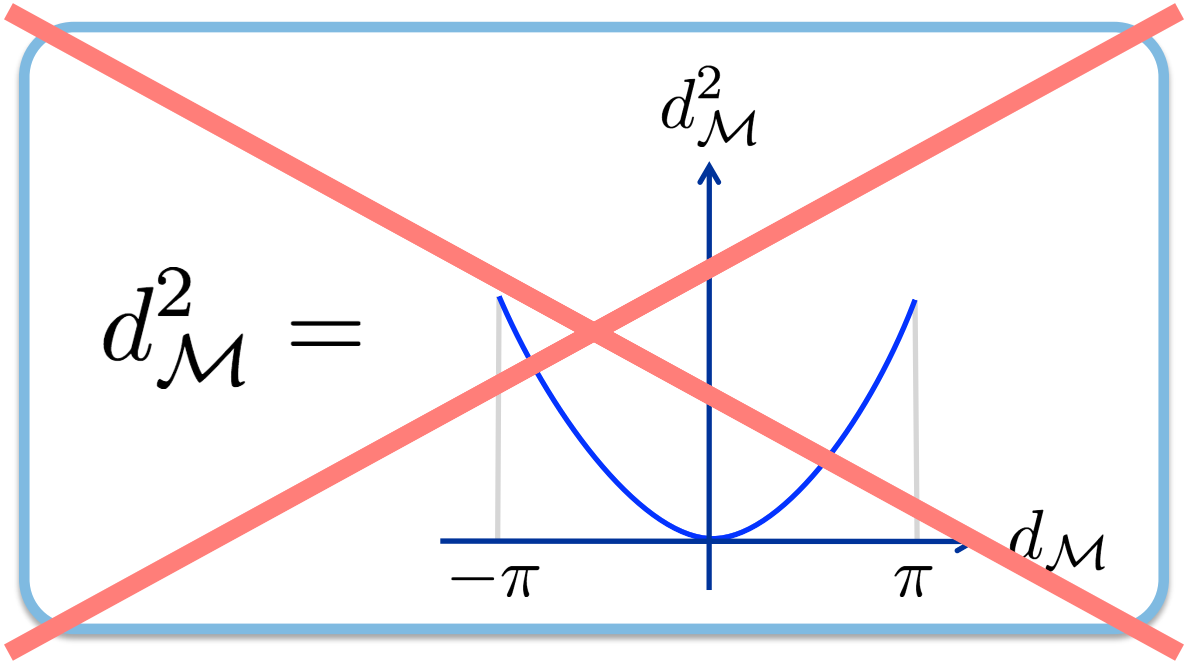

d

Tron, Afsari, Vidal – “Intrinsic consensus on SO(3) with almost-global convergence”
to appear in Conference on Decision and Control 2012

$$\varphi(\{x_i\}) = \sum_{(i,j)\in E} f_{SO(3)}\left(d_{\mathcal{M}}(x_i, x_j)\right)$$

$$f_{SO(3)} =$$

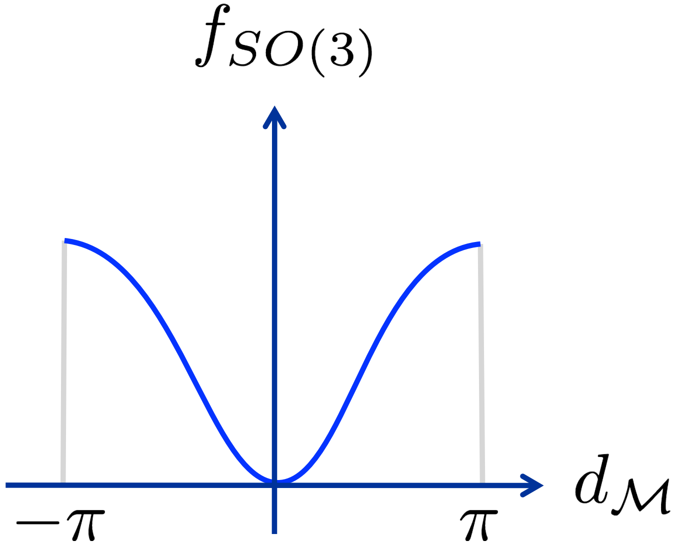

f_

Tron, Afsari, Vidal – “Intrinsic consensus on SO(3) with almost-global convergence”
to appear in Conference on Decision and Control 2012

$$\varphi(\{x_i\}) = \sum_{(i,j)\in E} f_{SO(3)}\left(d_{\mathcal{M}}(x_i, x_j)\right)$$

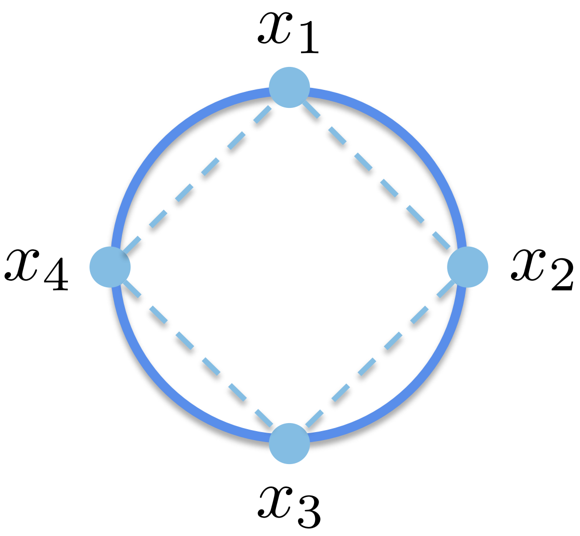

Saddle point

$$\varphi(\{x_i\}) = \sum_{(i,j)\in E} f_{SO(3)}\big(d_{\mathcal{M}}(x_i, x_j)\big)$$

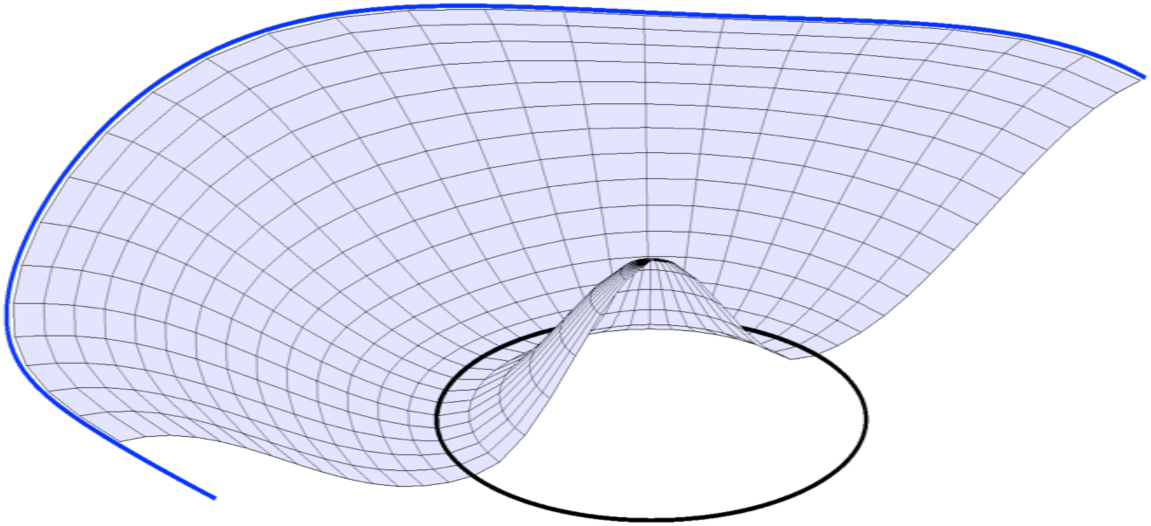

Almost-global convergence

Theorem

On SO(3), the only set of stable equilibria is the set of global minimizers

$$f_{SO(3)} = \frac{a}{b} - (\frac{1}{b} + \theta)\exp(-b\theta)$$

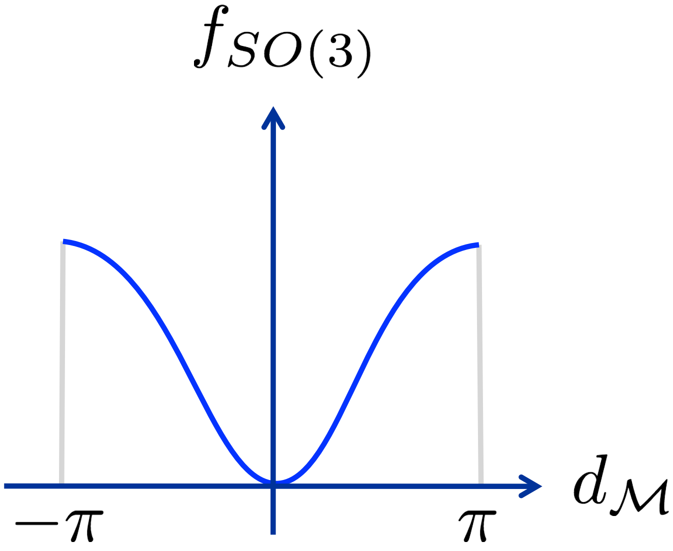

Tron, Afsari, Vidal – "Intrinsic consensus on SO(3) with almost-global convergence" to appear in Conference on Decision and Control 2012

Center for IMAGING SCIENCE

Experiments – Squared distance

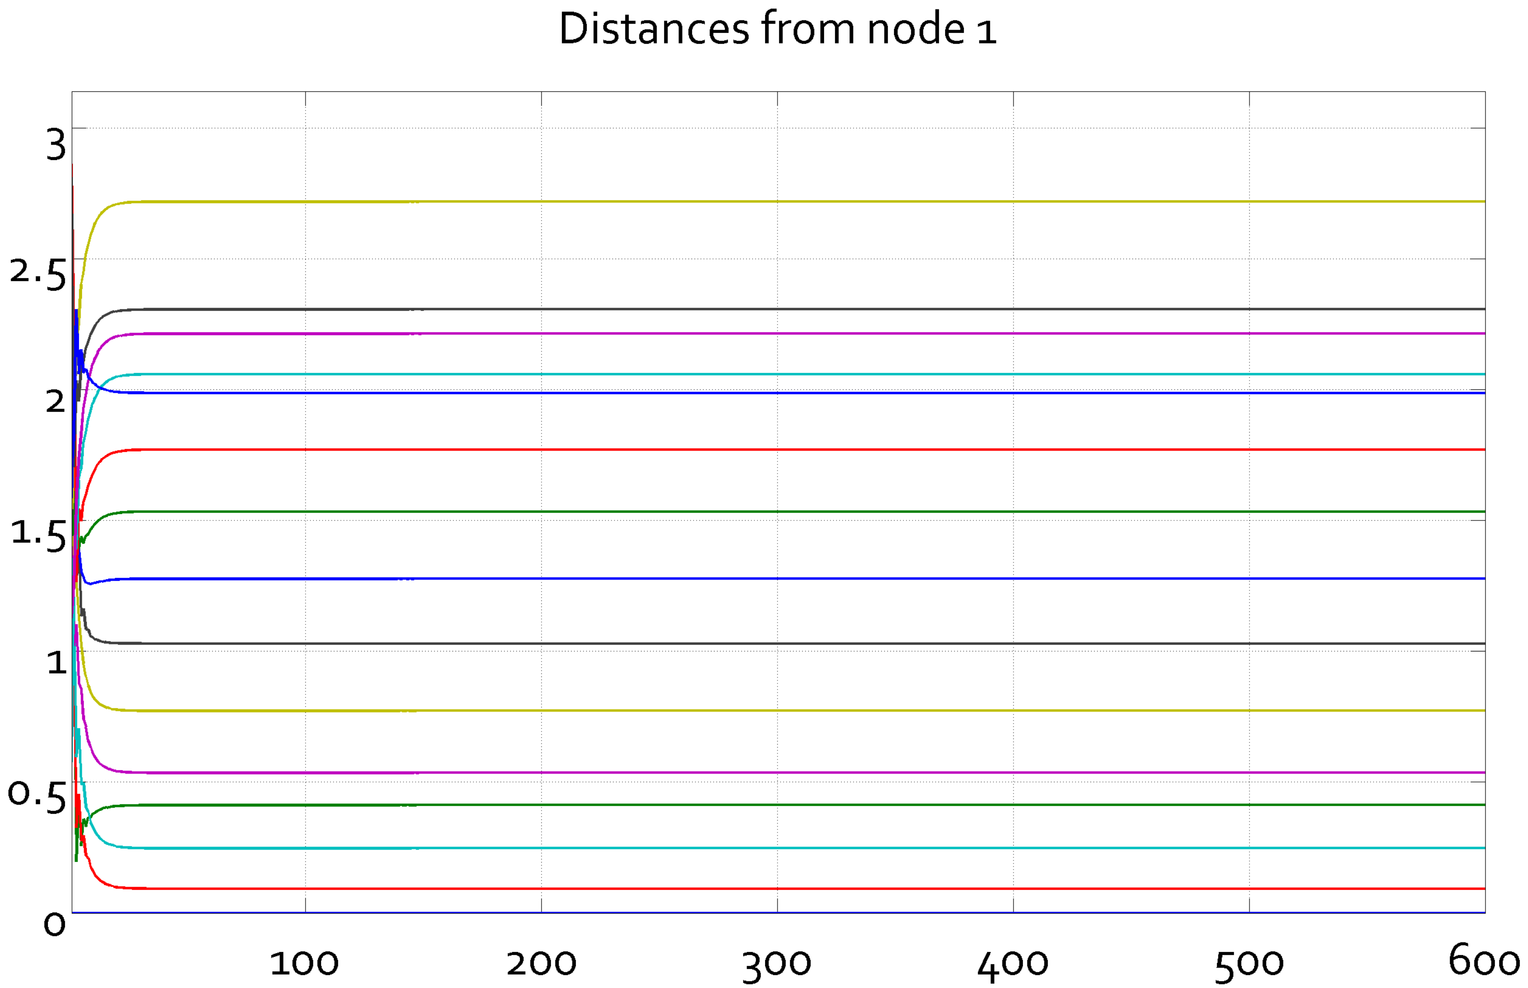

N=15 nodes, 4-regular graph, closed geodesic

Experiments – Reshaped distance

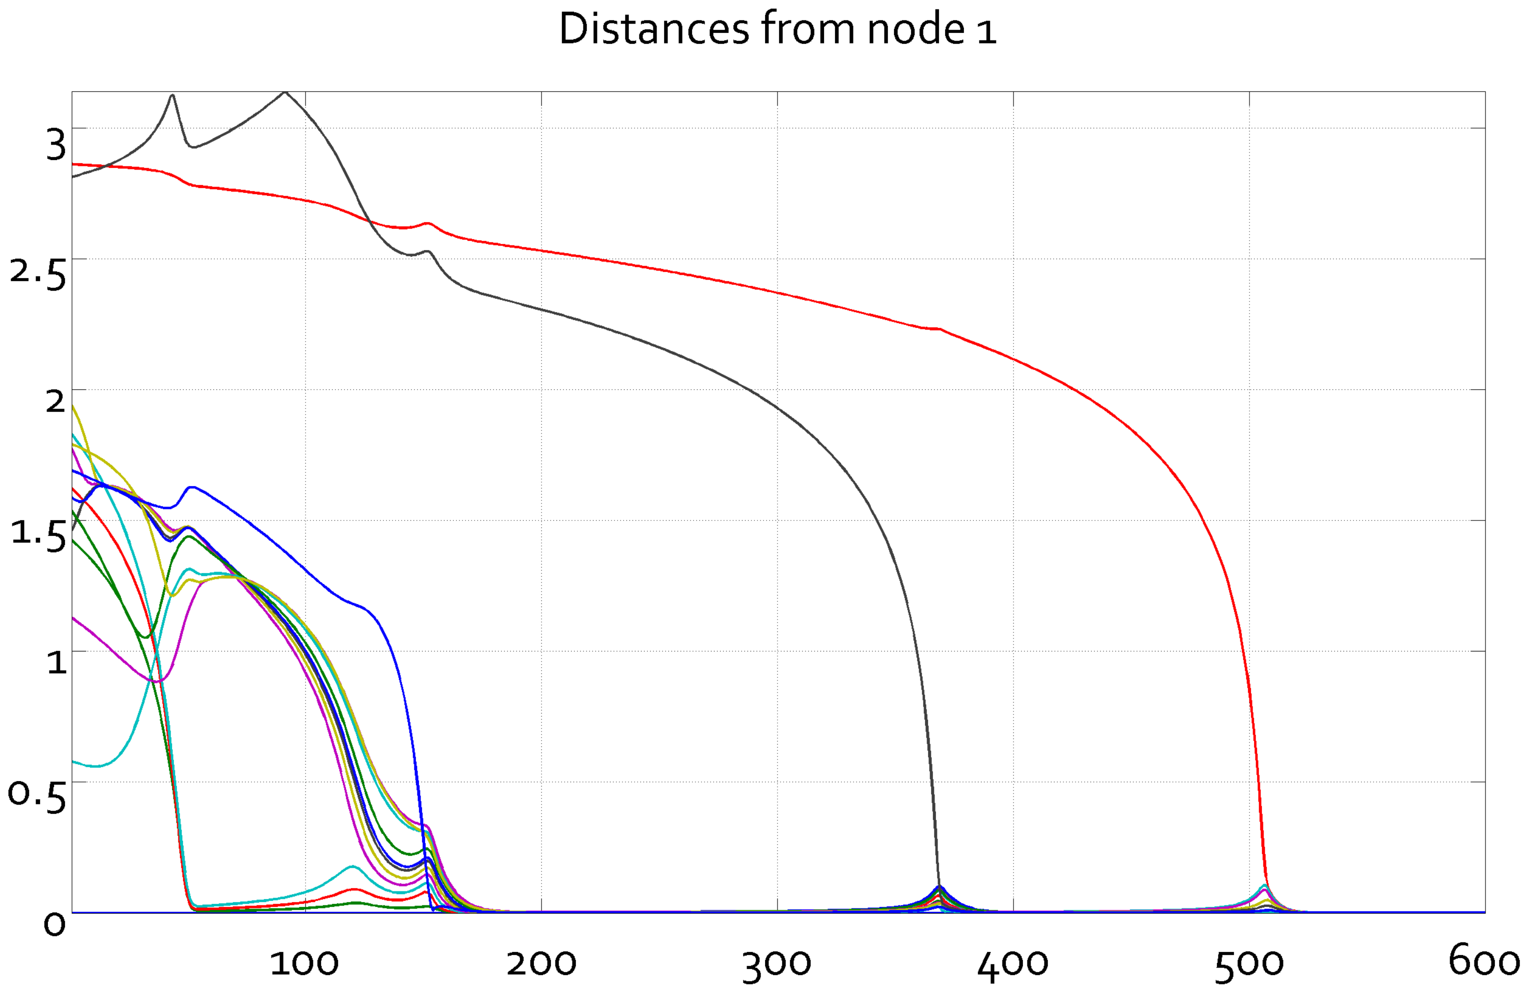

Average number of iterations for a K-regular network with N nodes

N	50	100	200	500	100
K	4	8	16	40	4
Avg. It.	2376	2593	2651	2932	4157

JHU vision lab

Image-based camera network localization

Center for IMAGING SCIENCE

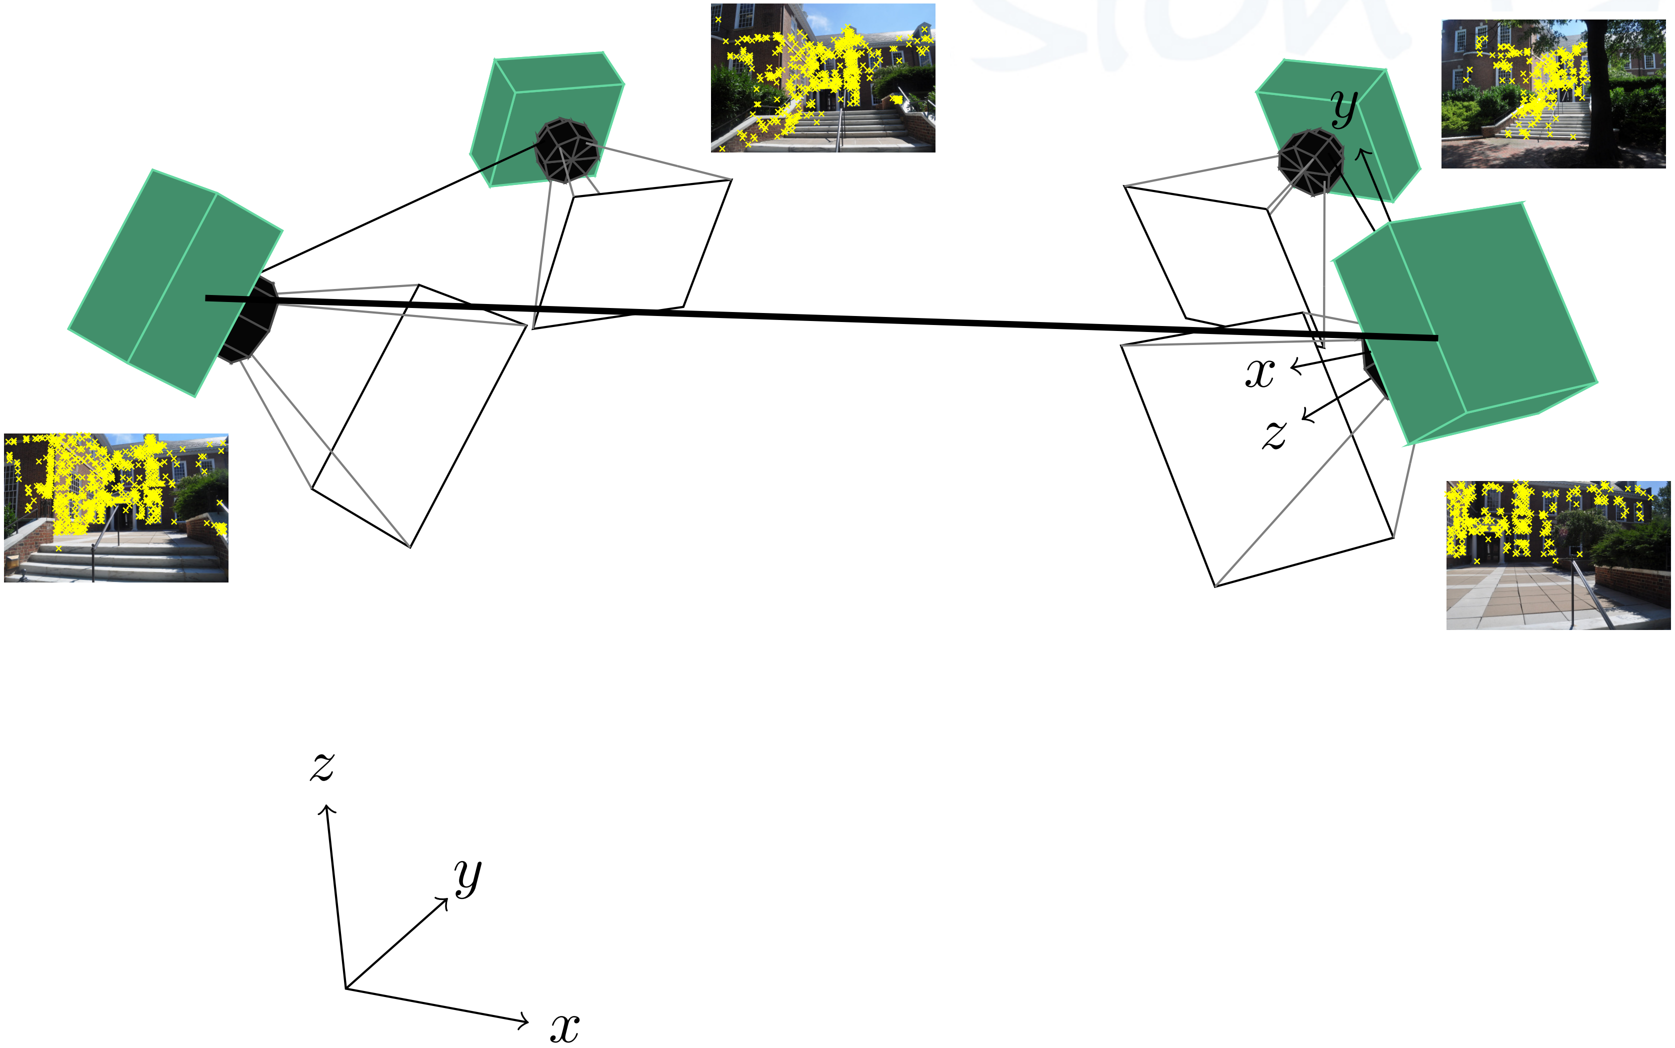
y
x
z
z
y
x

Vision graph

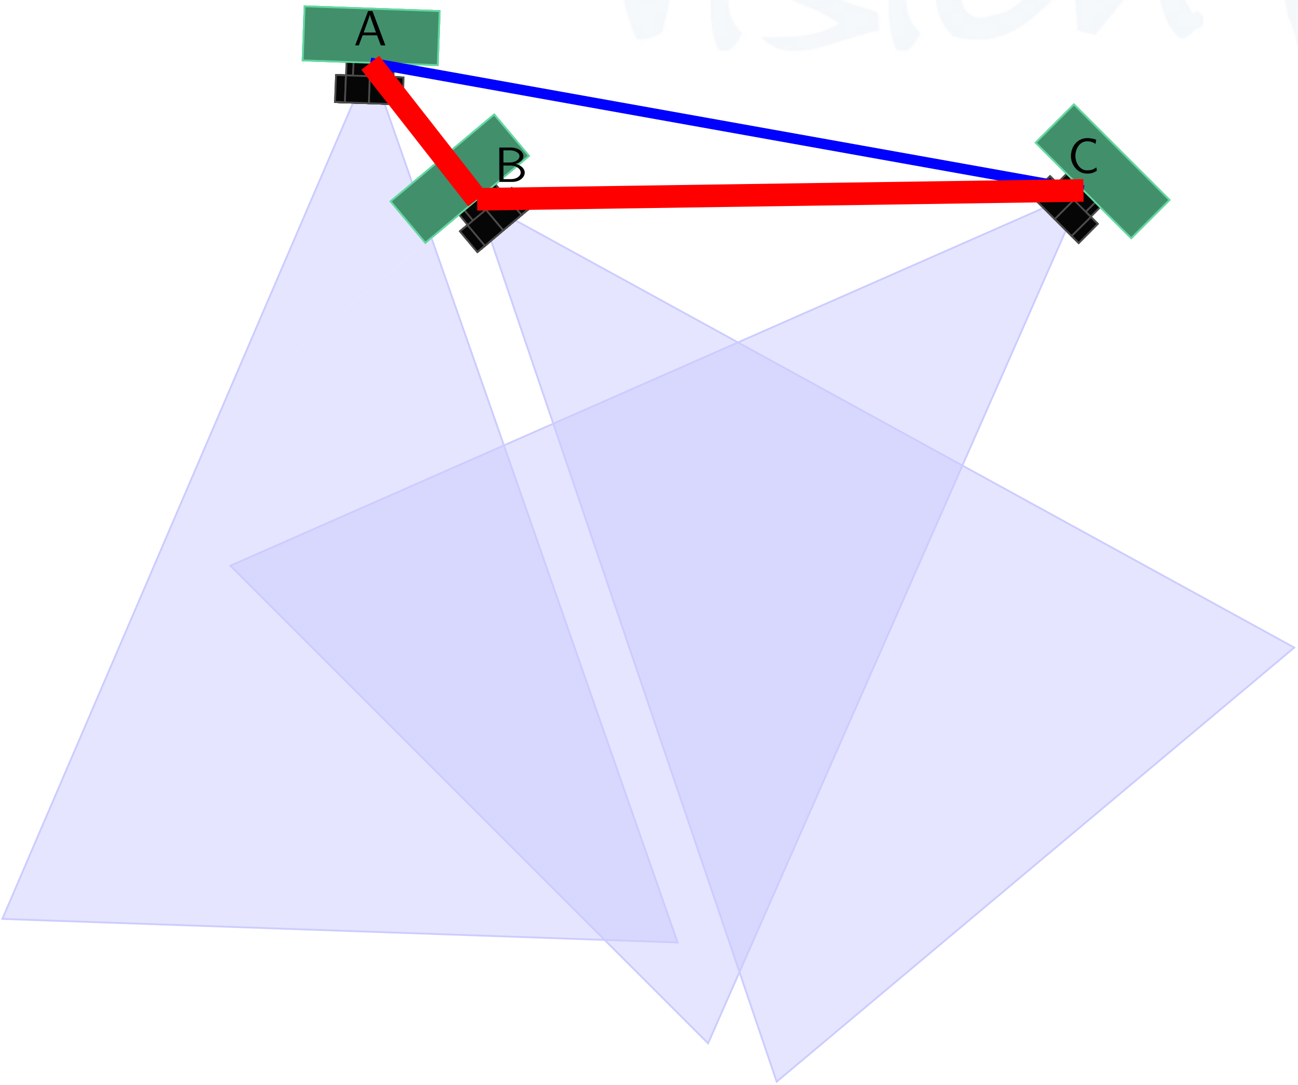

Input

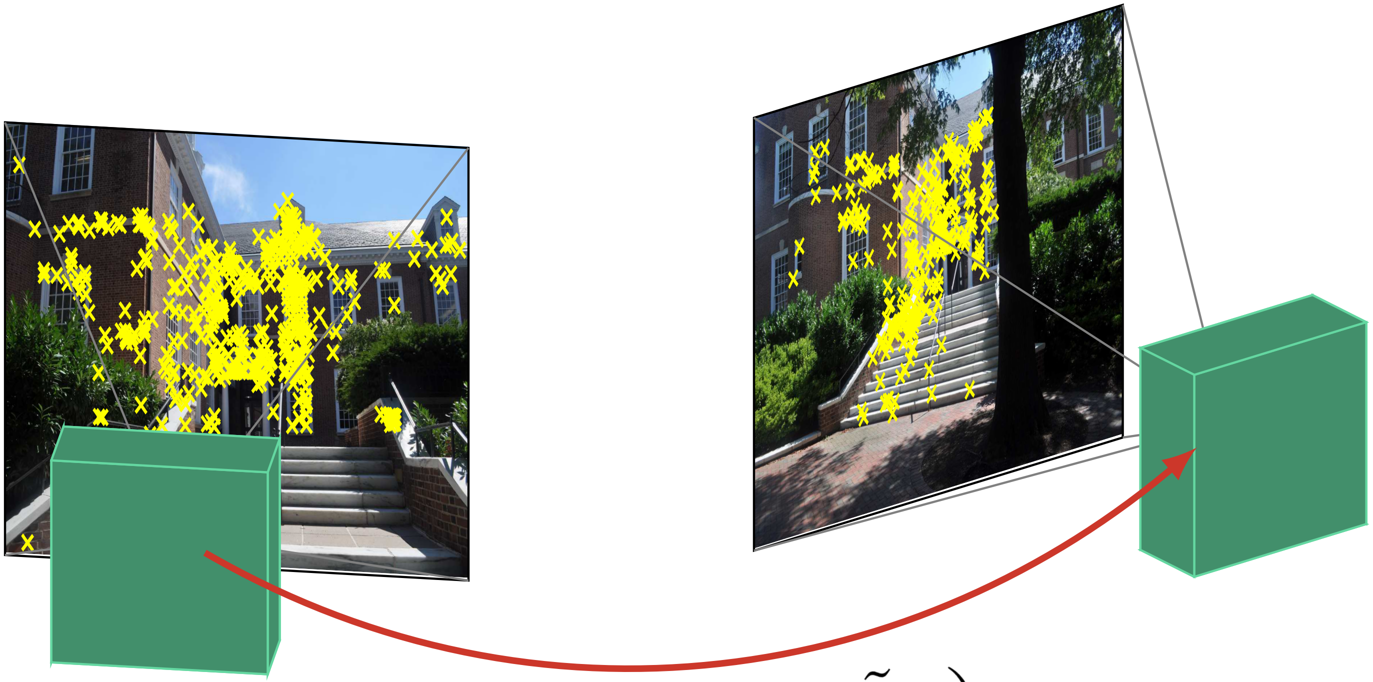

$$(\tilde{R}_{ij}, \tilde{t}_{ij})$$

(\ti

Output

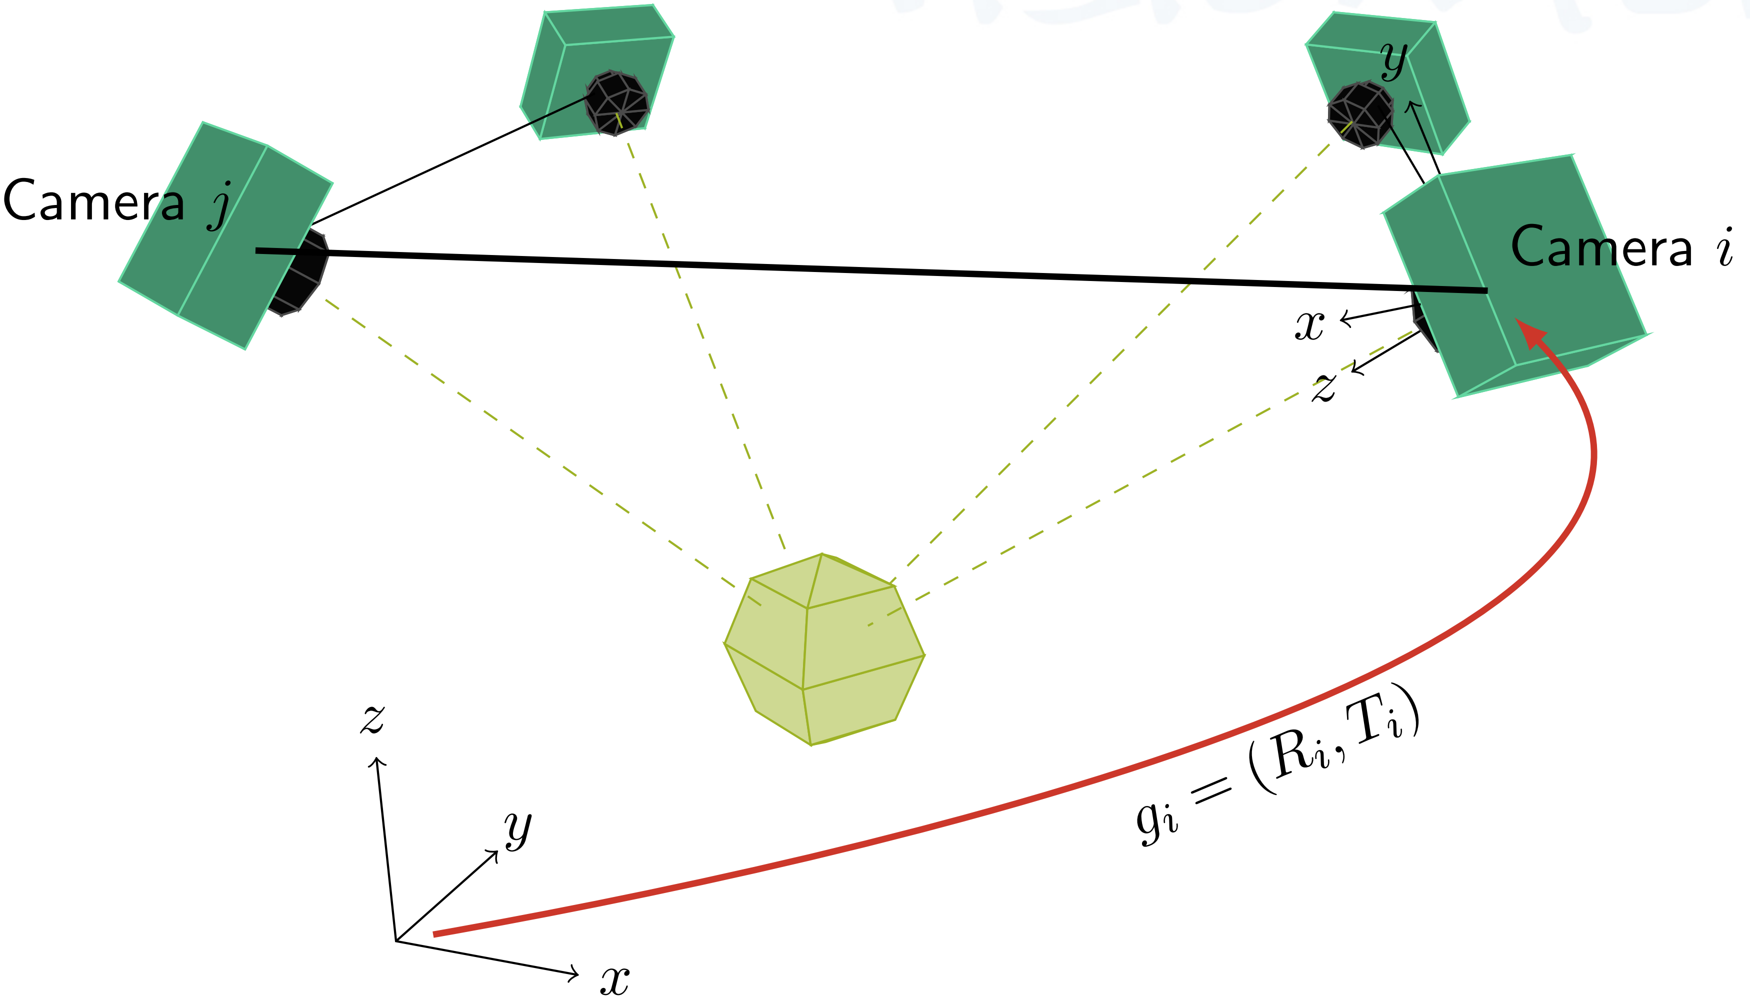

\r
(3

$$\min_{\{g_{ij}\}} \sum_{(i,j)\in E} \boxed{f\left(d_{SE(3)}\left(\boxed{g_{ij}}, \boxed{\tilde{g}_{ij}}\right)\right)}$$

d^

Challenge 1: Inconsistent Transf.

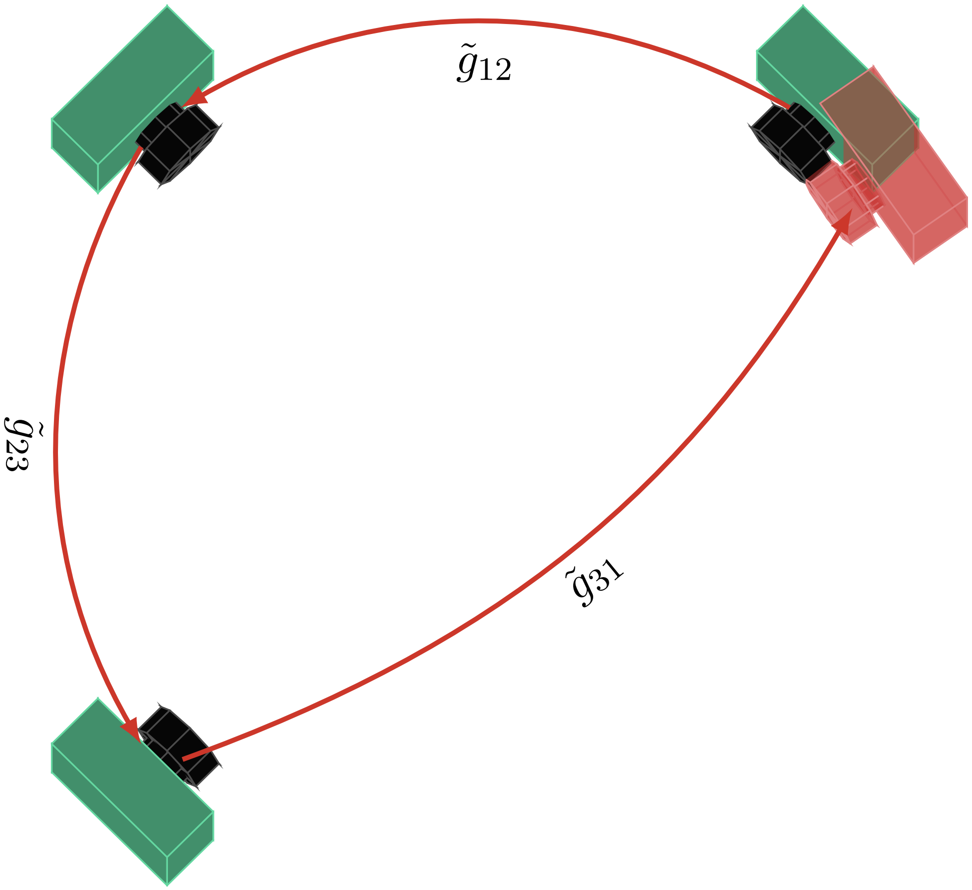

Optimization problem

\r
(3

$$\min_{\{g_{ij}\}} \sum_{(i,j)\in E} f\left(d_{SE(3)}(g_{ij}, \tilde{g}_{ij})\right)$$

d^

s.t. $\{g_{ij}\}$ are consistent

Pose reparametrization

Theorem

$\{g_{ij}\}$ are consistent

$$\Updownarrow$$

$$g_{ij} = g_i^{-1} g_j$$

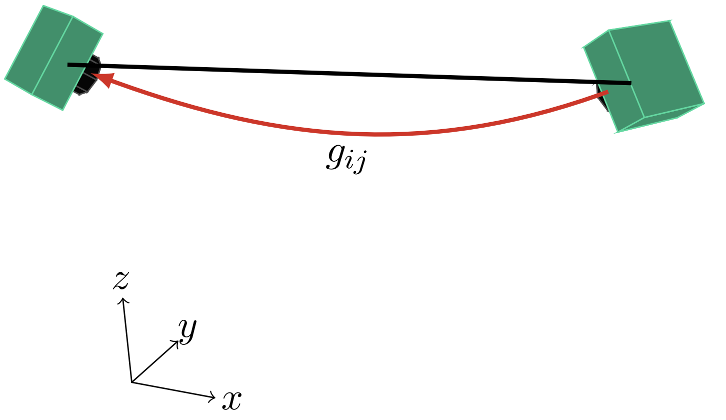

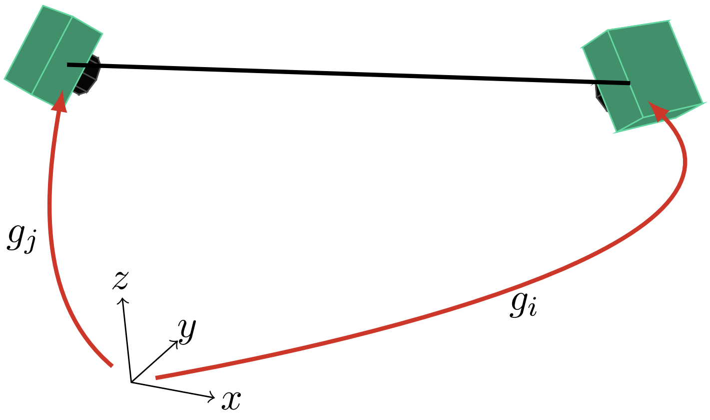

Tron, Vidal – "Distributed image-based 3-D localization in camera sensor networks"
Conference on Decision and Control 2009

Optimization problem

$$g_{ij} = g_i^{-1} g_j$$

$$\min_{\{g_{ij}\}} \sum_{(i,j)\in E} f\left(d_{SE(3)}(g_{ij}, \tilde{g}_{ij})\right)$$

\r
(3

d^

Optimization problem

$$f\left(d_{SE(3)}(\cdot,\cdot)\right) = f_{SO(3)}(\cdot,\cdot) + f_{\mathbb{R}^3}(\cdot,\cdot)$$

$$\min_{\{g_{ij}\}} \sum_{(i,j)\in E} f\left(d_{SE(3)}(g_i^{-1}g_j, \tilde{g}_{ij})\right)$$

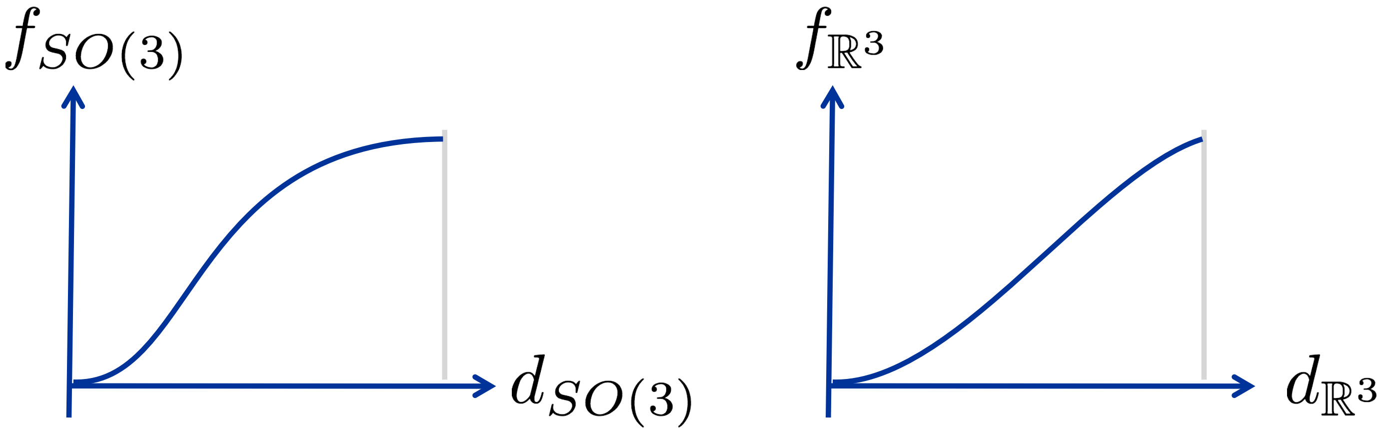

\r
(3

f\b

f

d_

d

Optimization problem

$$\min_{\{R_i,T_i\}} \sum_{(i,j)\in E} f_{SO(3)}(R_i^T R_j, \tilde{R}_{ij})$$

$$+ \sum_{(i,j)\in E} f_{\mathbb{R}^3}(R_i^T (T_j - T_i), \tilde{T}_{ij})$$

$$\tilde{T}_{ij} = \lambda_{ij} \tilde{t}_{ij}$$

Optimization problem

$$\min_{\{R_i,T_i,\lambda_{ij}\}} \sum_{(i,j)\in E} f_{SO(3)}(R_i^T R_j, \tilde{R}_{ij})$$

$$+ \sum_{(i,j)\in E} f_{\mathbb{R}^3}(R_i^T(T_j - T_i), \lambda_{ij}\tilde{t}_{ij})$$

Challenge 2: Scale Ambiguity

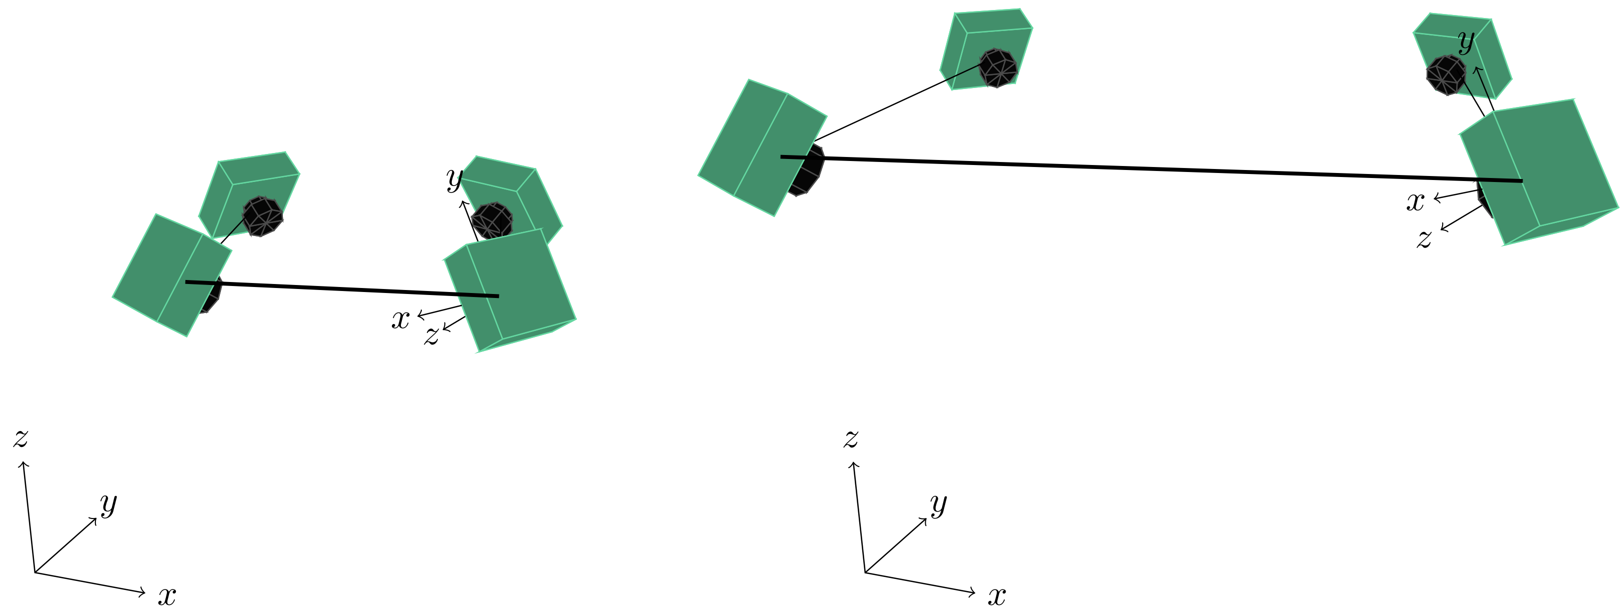

Fix shortest edge length

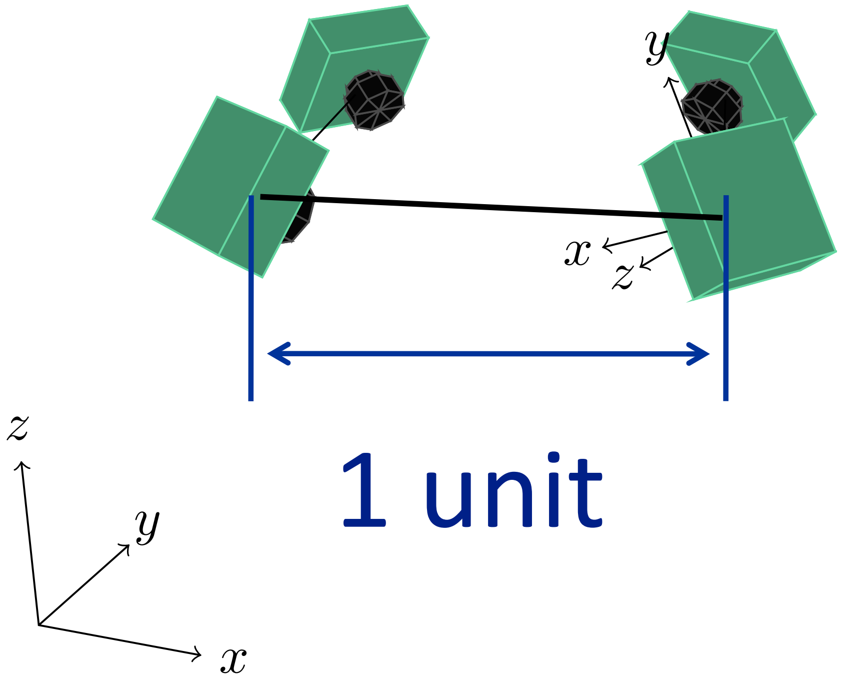

1 unit

Optimization problem

$$\min_{\{R_i,T_i,\lambda_{ij}\}} \boxed{\sum_{(i,j)\in E} f_{SO(3)}(R_i^T R_j, \tilde{R}_{ij})}$$

$$\boxed{\begin{aligned} &+ \sum_{(i,j)\in E} f_{\mathbb{R}^3}(R_i^T(T_j - T_i), \lambda_{ij}\tilde{t}_{ij}) \\ &\text{s.t. } \lambda_{ij} \geq 1 \end{aligned}}$$

How to get global minimizer?

Tron, Vidal - “Distributed Image-Based 3-D Localization of Camera Sensor Networks” in Conference on Decision and Control 2009

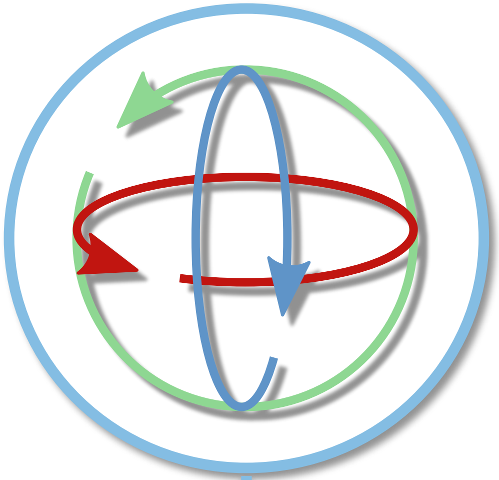

Initialization of rotations

$$\min_{\{R_i\}} \sum_{(i,j)\in E} f_{SO(3)}(R_i^T R_j, \tilde{R}_{ij})$$

\r
(I

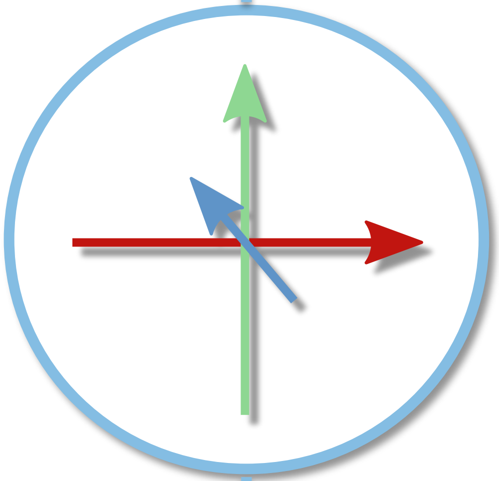

Initialization of translations

$$\min_{\{T_i,\lambda_{ij}\}} \sum_{(i,j)\in E} f_{\mathbb{R}^3}(R_i^T(T_j - T_i), \lambda_{ij}\tilde{t}_{ij})$$

$$\text{s.t. } \lambda_{ij} \geq 1$$

\r
R
\t

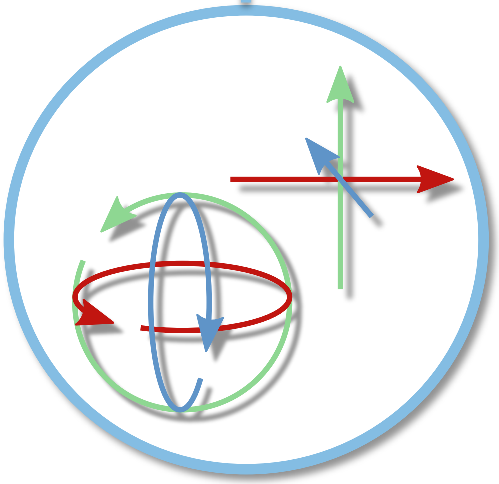

Complete optimization

$$\min_{\{R_i, T_i, \lambda_{ij}\}} \varphi(\{R_i, T_i, \lambda_{ij}\})$$

$$\text{s.t. } \lambda_{ij} \geq 1$$

\min
{R_i
\lam

Gradient descent

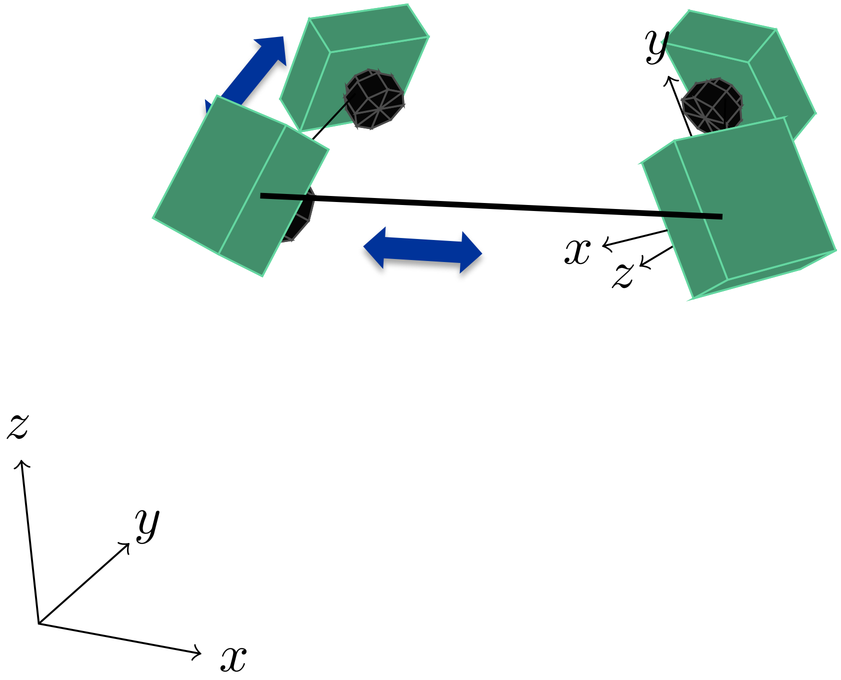

$$R_i(t+1) = R_i(t)\exp\Big(\varepsilon \sum_{(i,j)\in E} \log(R_i^\top R_j \tilde{R}_{ij}^\top) + \ldots\Big)$$

How to choose the step size

Theorem

$$\varphi(\{R_i, T_i, \lambda_{ij}\}) = \sum_{(i,j)\in E} \varphi_{ij}(R_i, R_j, T_i, T_j, \lambda_{ij})$$

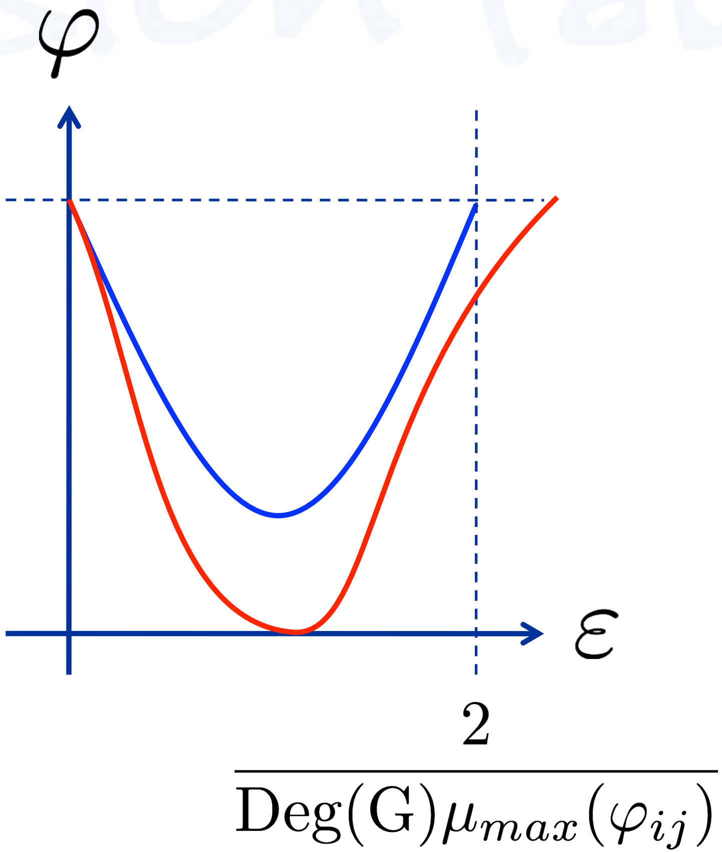

$\mathrm{Deg}(G)$ = max # of neighbors

$\mu_{max}(\varphi_{ij})$ = bound on max eval of Hessian of φ_{ij}

Noiseless case

Theorem

In noiseless case, equivalent to Riemannian consensus

$$\min_{\{R_i\}} \sum_{(i,j)\in E} f_{SO(3)}(R_i^T R_j, \tilde{R}_{ij})$$

$$\Updownarrow$$

$$\min_{\{R_i\}} \sum_{(i,j)\in E} f_{SO(3)}(R_i', R_j')$$

Noiseless case

Theorem

In noiseless case, non-convex, but no local minima!

$$\min_{\{T_i,\lambda_{ij}\}} \sum_{(i,j)\in E} f_{\mathbb{R}^3}(R_i^T(T_j - T_i), \lambda_{ij}\tilde{t}_{ij})$$

$$\text{s.t. } \lambda_{ij} \geq 1$$

Experiments on synthetic data

Experiment: 8 cameras, 30 scene points, 4-regular graph

1. Camera poses and points
2. Images + Noise
3. Relative poses
4. Localization

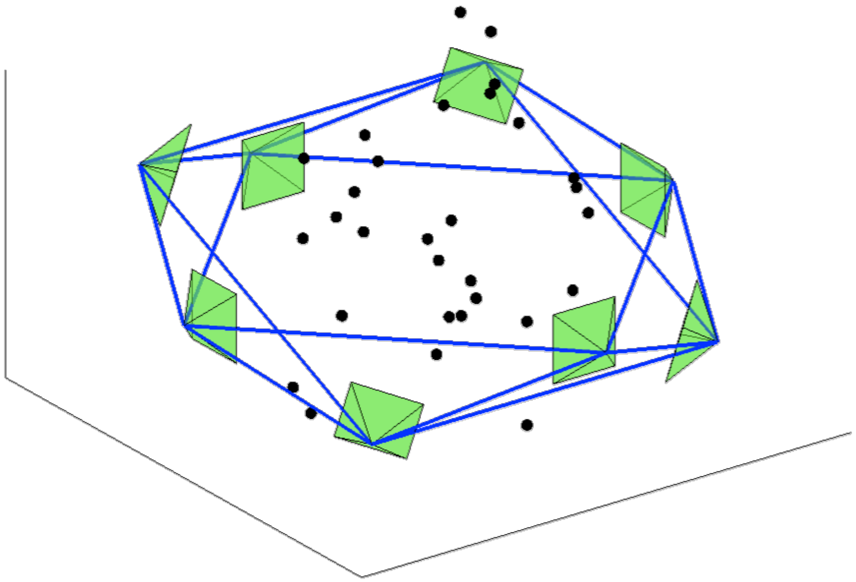

Experiments on synthetic data

Rotation errors [degrees]

Noise	0 pixels	1 pixels	3 pixels
Initial	0.00 ±0.00	2.77 ±0.58	4.80 ±1.74
Final	0.00 ±0.00	0.13 ±0.00	0.39 ±0.03

Translation direction errors [degrees]

Noise	0 pixels	1 pixels	3 pixels
Initial	0.00 ±0.00	0.11 ±0.00	0.33 ±0.05
Final	0.00 ±0.00	0.09 ±0.00	0.29 ±0.03

Geometric variance of scale ratios

Noise	0 pixels	1 pixels	3 pixels
Final	1.000	1.002	1.005

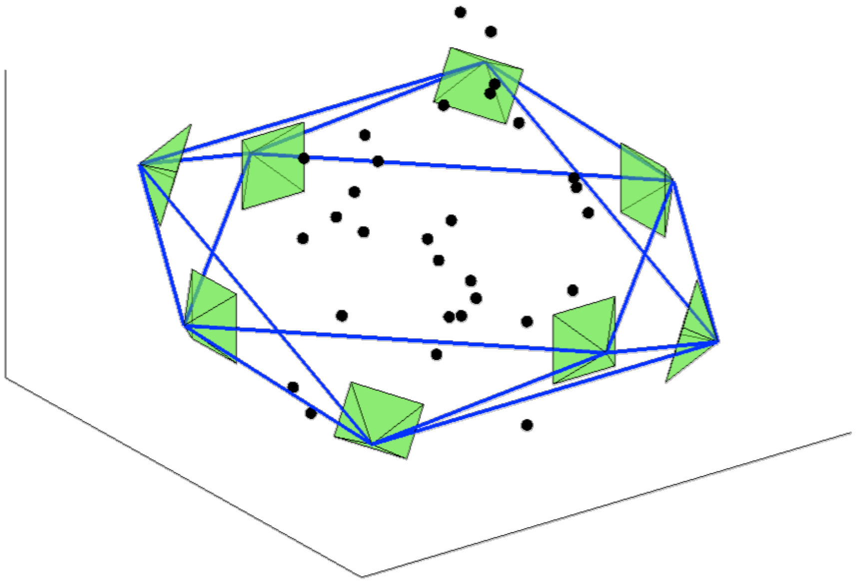

Experiments on real data

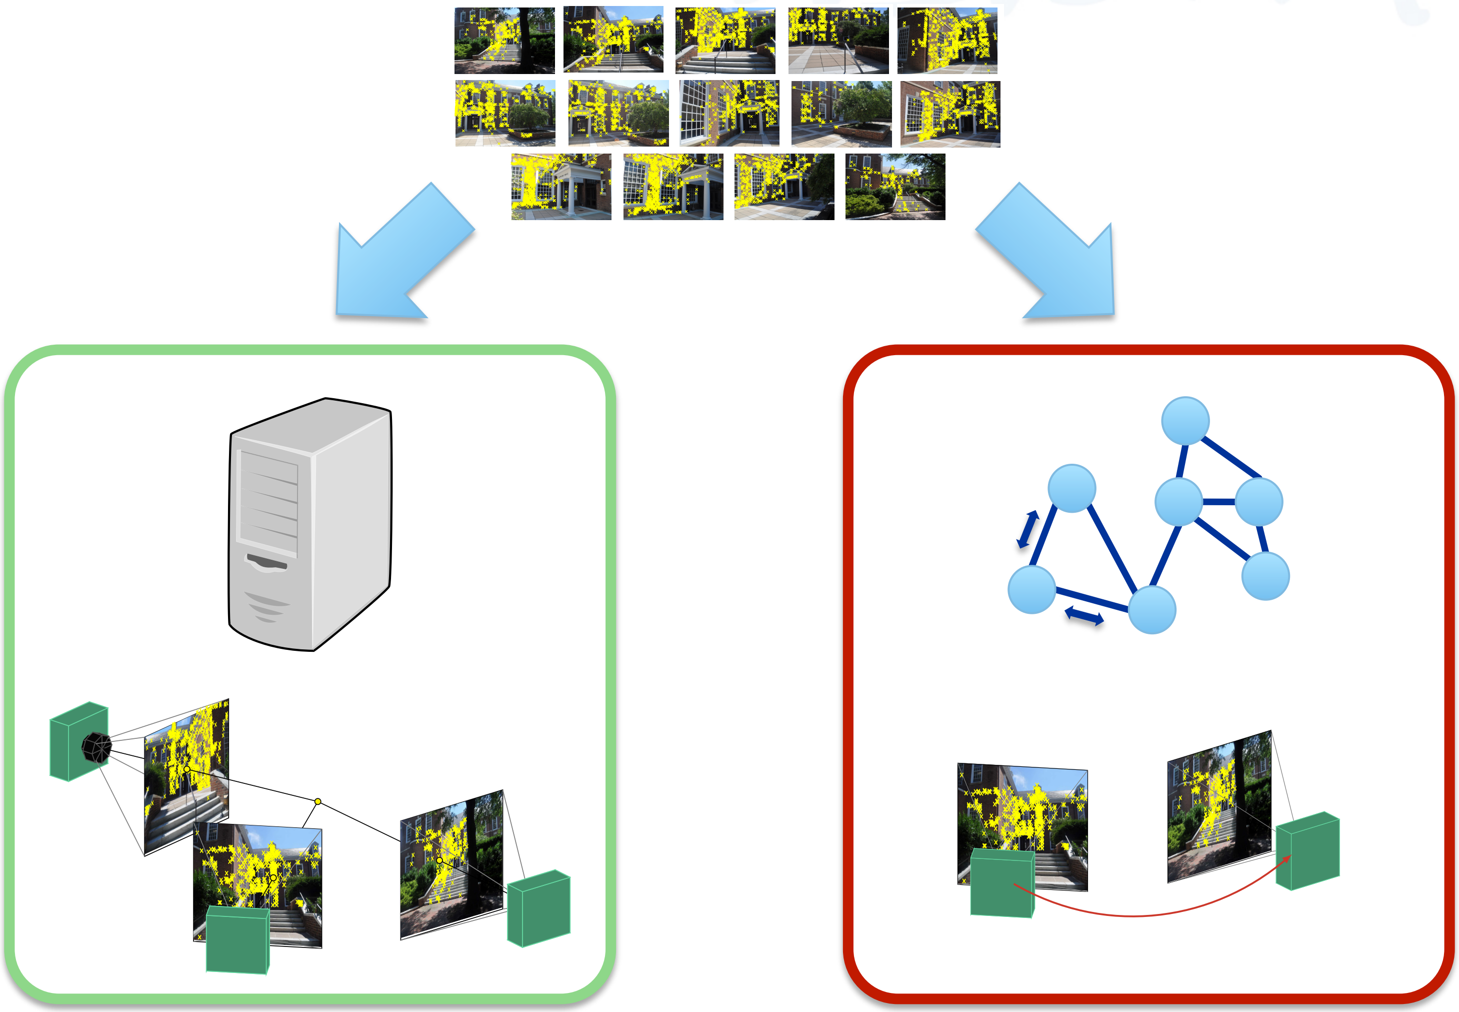

Snavely, Seitz, Szeliski. “Photo Tourism: Exploring image collections in 3D” in ACM Transactions on Graphics

Experiments on real data

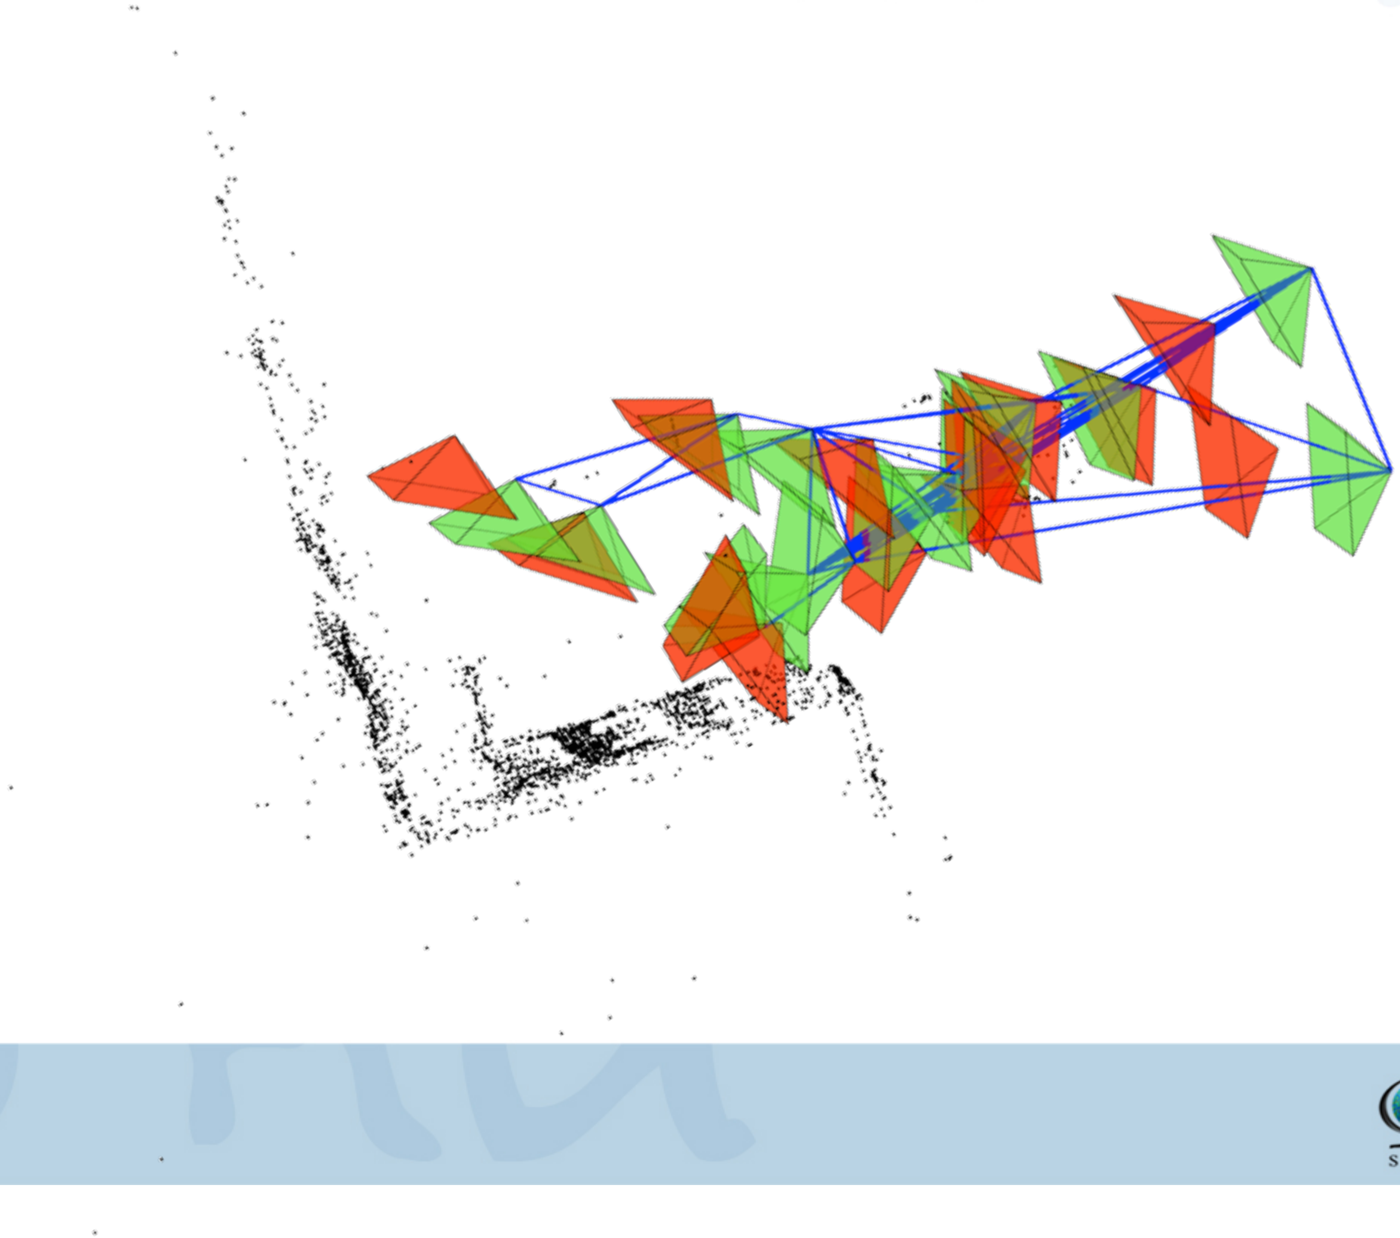

Experiments on real data

Rotation errors per edge

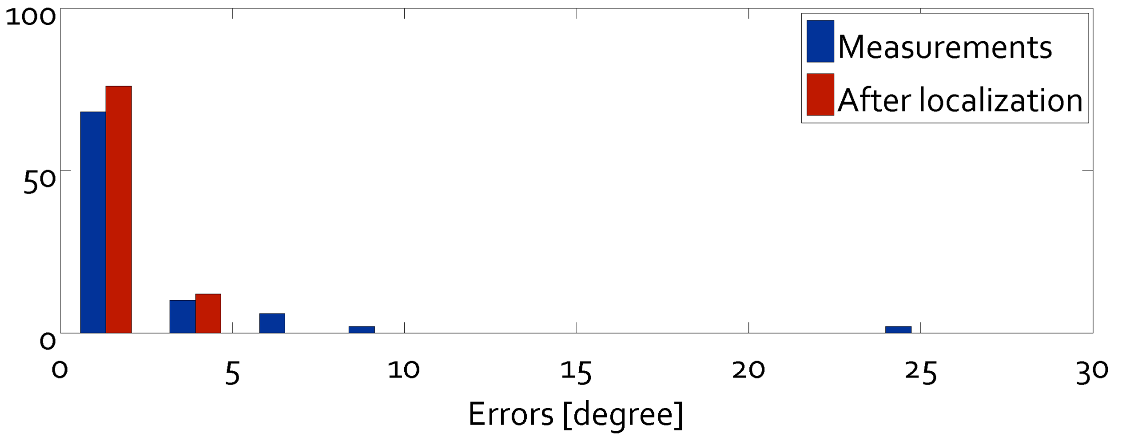

Translation direction errors per edge

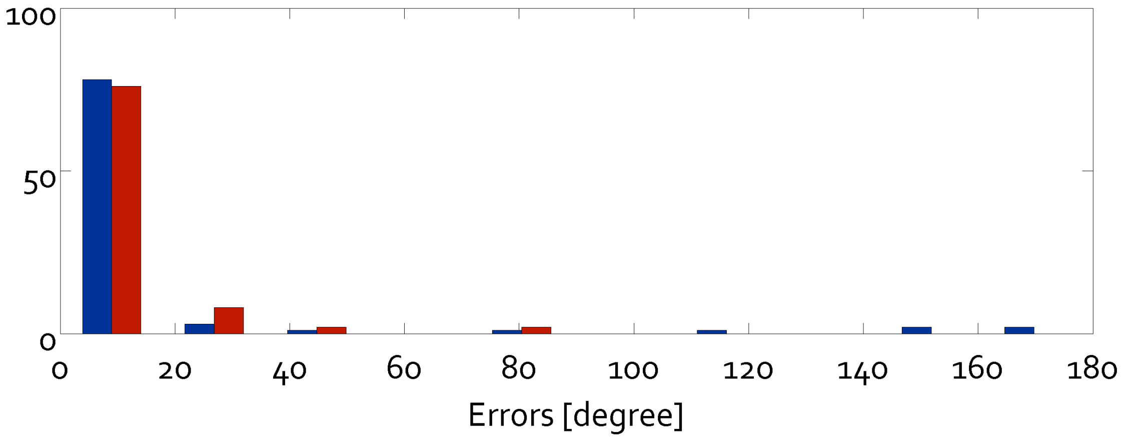

Conclusion

Riemannian Consensus

Theory

Sufficient convergence on any Riemannian manifold

Almost-global convergence on SO(3)

Applications

Image-based camera network localization

Future work

Convergence rate

Incorporate uncertainties

Dynamic case

- Time-varying measurements
- Pose + velocity

Improved localization

Thanks!