When is Rotations Averaging Hard?

Kyle Wilson
David Bindel
Noah Snavely

kwilson24@washcoll.edu
bindel@cs.cornell.edu
snavely@cs.cornell.edu

The Rotations Averaging Problem

- A key subproblem in global Structure from Motion [2]

Compute: vertex orientations $\mathcal{R} : V \to \mathrm{SO}(3)$

Given: a graph $G = (V, E)$
measures of relative orientation $\widetilde{\mathcal{R}} : E \to \mathrm{SO}(3)$

To Minimize: squared geodesic error $\sum_{(i,j)\in E} d\left(\widetilde{\mathtt{R}}_{ij},\ \mathtt{R}_i\mathtt{R}_j^\top\right)^2$

(this cost is very natural, but there are other common cost functions)

Visualization of the ArtsQuad problem graph (6514 cameras) [1]
Notice the rich, clustered graph topology. Real problems are not G_{np} random graphs.

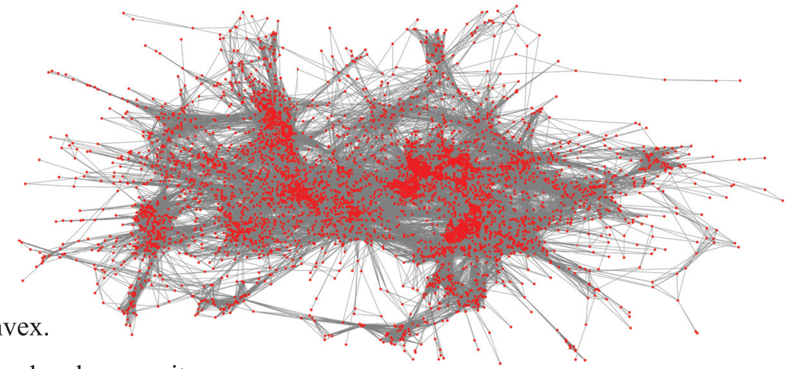

- Good solvers exist [3,4], but sometimes fail. Problem is nonconvex.
- **Contributions**: insights into which problems are easy, bounds on local convexity

Local Convexity

Solver failure ↔ wrong local minima

Why can some problems be reliably solved, but others can't? Which problems have bad local minima?

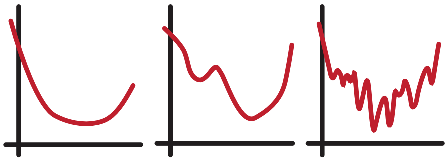

Not all nonconvex problems are equally hard.

Local Convexity: problem's Hessian is positive semi-definite

Locally convex everywhere on a convex domain → convex problem
Locally convex on large neighborhood of a global minima → problem is easier
Sufficient, but not necessary condition.

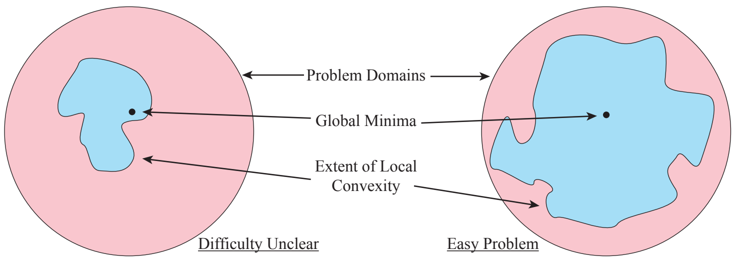

Gauge Ambiguity

Old result: Rotations averaging is locally convex almost nowhere.

New result: Local convexity occurs in many problems of interest.

Key Idea: a gauge ambiguity reveals local convexity

$(\mathtt{R}_1, \ldots, \mathtt{R}_n) \equiv (\mathtt{R}_1\mathtt{S}, \ldots \mathtt{R}_n\mathtt{S}) \quad \forall \mathtt{S} \in \mathrm{SO}(3)$

Toy Example:

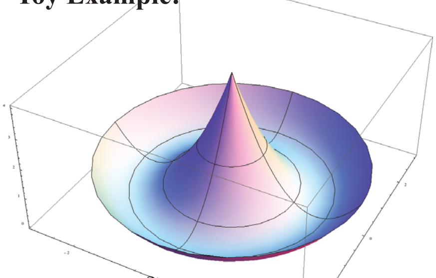

$\min_{r,\theta} (1-r)^2,\ 0 \le r < 2,\ \theta \in [0, 2\pi)$

Note: locally non-convex on $0 \le r < 1$

Fixing the Gauge: remove ambiguity by setting $\theta = 0$ (analagously, $\mathtt{R}_k = \mathbf{I}$).

Results

Hessian of one term of the cost function:

$$\mathbf{H}_{ij} = \begin{bmatrix} \mu\mathbf{I} + (2-\mu)\mathbf{w}\mathbf{w}^\top & -\mu\mathbf{I} - (2-\mu)\mathbf{w}\mathbf{w}^\top \\ -\mu\mathbf{I} - (2-\mu)\mathbf{w}\mathbf{w}^\top & \mu\mathbf{I} + (2-\mu)\mathbf{w}\mathbf{w}^\top \end{bmatrix} + \begin{bmatrix} \mathbf{0} & -\theta[\mathbf{w}]_\times \\ \theta[\mathbf{w}]_\times & \mathbf{0} \end{bmatrix}$$

1. positive semi-definite term (problem structure)
2. indefinite term (arises from curvature of the space)

where:
$\theta =$ residual magnitude
$\mathbf{w} =$ residual (unit) direction
$\mu =$ $\theta \cot(\theta/2)$
$[\cdot]_\times =$ cross product matrix

Fixing the gauge makes the Hessian positive-definite in some parts of the problem domain.

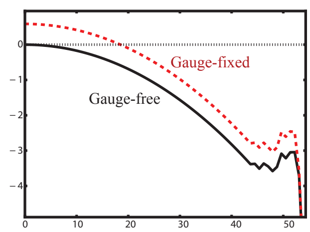

$\lambda_{\min}$ of Hessian along a ray of solutions

Bounds:

Can we get insight by approximating away the directions of residuals?

$\lambda_{\min}\left(\mathbf{L}_{\mathrm{norm}}^{\hat{k}}\right) > 1 \longrightarrow$ local convexity

where $\mathbf{L}_{\mathrm{norm}}^{\hat{k}}$ is a weighted, normalized graph Laplacian with it's k^{th} row and column removed.

And bound by the magnitude of the residuals too?

$\lambda_2(\mathbf{L})/n > \Delta/\mu_{\min} \longrightarrow$ local convexity

structure term, noise term

Application

- Provides insight into good problem instance construction: high algebraic connectivity drives local convexity.
- Bounds are useful for predicting problem instance difficulty

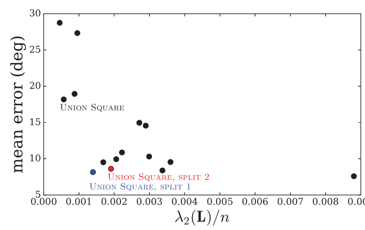

- Bounds can guide a sequence of easier subproblems:

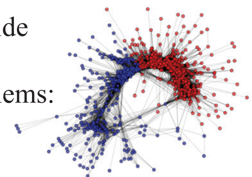

References

[1] Crandall, D., Owens, A., Snavely, N., Huttenlocher, D.: Discrete-continuous optimization for large-scale structure from motion. In: CVPR. (2011)
[2] Govindu, V.M.: Combining two-view constraints for motion estimation. In: CVPR.(2001)
[3] Hartley, R., Trumpf, J., Dai, Y., Li, H.: Rotation averaging. In: IJCV. (2013)
[4] Chatterjee, A., Govindu, V.M.: Efficient and robust large-scale rotation averaging. In: ICCV. (2013)