Supplemental Material for:
On the Distribution of Minima in Intrinsic-Metric Rotation Averaging

Kyle Wilson
Washington College
Department of Math and Computer Science
Chestertown, MD
kwilson24@washcoll.edu

David Bindel
Cornell University
Department of Computer Science
Ithaca, NY
bindel@cs.cornell.edu

1. Connecting the Empirical and Theory Contributions

In the theory sections of this paper we identified the two salient properties that affect local convexity:

Noise: Lower edge noise yields more local convexity

Graph Structure: Greater connectivity yields more local convexity.

The interaction of the two factors is complicated, and we an give exact expression in terms of the spectrum of a particular weighted graph Laplacian (Equation 5). However, we also separate the two factors to seek insight. As shown in Equation 17, algebraic connectivity ($\lambda_2(G)$, the second-smallest eigenvalue of a graph's Laplacian matrix) is a useful measure of the graph structure's contribution to local convexity.

In Figure 1 we look at noise and connectivity on the empirical experiments from Section 4. We visualize three factors:

- Synthetic noise level σ_n
- Graph connectivity $\lambda_2(G)$
- Multiplicity of the most common minimum $\%_{\max}$

For easy problem instances there will be a single dominant local minimum. Nearly 100% of all of the minima we found will be identical.

Notice the following in Figure 1:

- Lower $\lambda_2(G)$ tends to yield lower $\%_{\max}$.
- Higher σ_n tends to yield lower $\%_{\max}$.
- The $\%_{\max}$ is sharply lower below a certain connectivity and below a certain noise level.

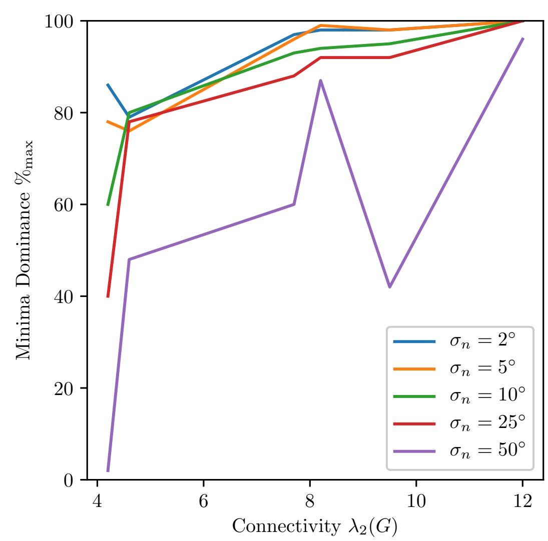

Figure 1. A plot showing how problem difficulty, as quantified by $\%_{\max}$, the multiplicity of the most common minimum, is affected by noise level and problem connectivity.

2. Empirical Experiments with Outliers

The empirical experiments in Section 4 all used a Gaussian-like noise model. However, real problems instances commonly have heavy-tailed noise. Here we present the same types of experiments as in Section 4 but with an inlier/outlier noise model.

Figure 2 is constructed analogously to rows (4) and (6) of Table 1. All problem instances were made with identical $\sigma_n = 5°$ inlier noise and the columns are varying percentages of outlier edges. Outliers are distributed uniformly at random over $SO(3)$.

Notice that the better connected graph (row 2) tolerates more outliers before bad local minima appear everywhere, but by even 5% outliers both problems look very hard.

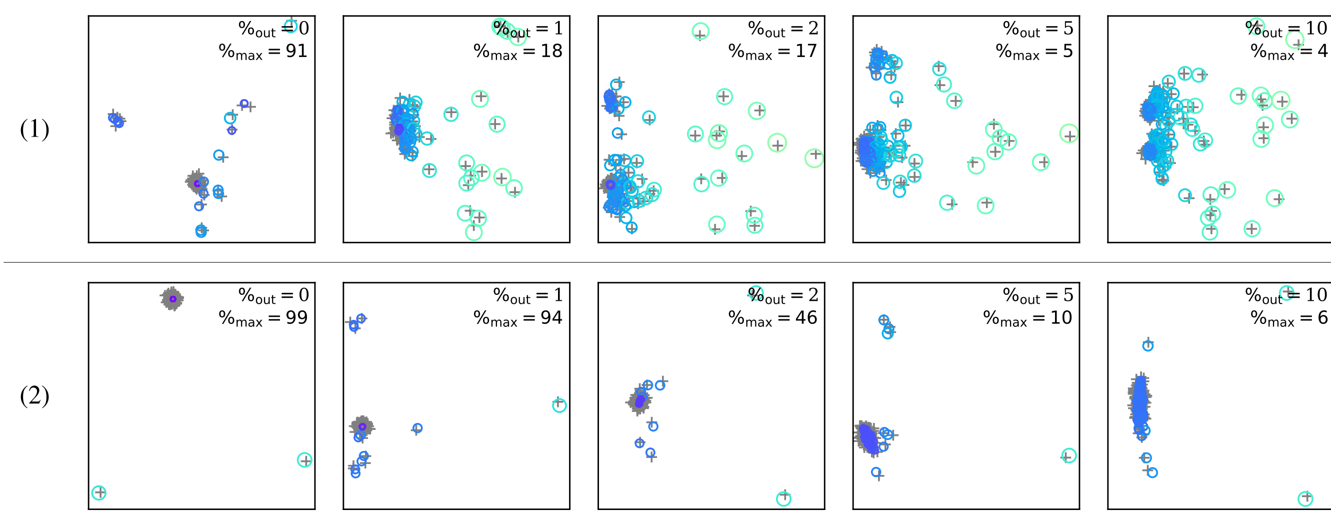

Figure 2. Visualizations in the style of Figure 1 from Section 4 of the paper. All problem instances were constructed from the same G_{nm} random graph instance where $n = 40$. Row (1) has $m = 240$ and row (2) has $m = 400$. All instances have the same 5° inlier noise applied. The columns are in order of increasing outlier percentage. Even a few outliers introduce many bad minima to the cost surface. Row (2), the better connected graph, tolerates more outliers.