

# On Understanding Disentanglement Puzzles

Xinya Zhang<sup>1</sup>, Paul Kry<sup>2</sup> and Etienne Vouga<sup>3</sup>

<sup>1</sup>Dept. of C.S., The University of Texas at Austin, Texas, USA; xinyazhang@utexas.edu

<sup>2</sup>School of Computer Science, McGill University, Quebec, Canada; kry@cs.mcgill.ca

<sup>3</sup>Dept. of C.S., The University of Texas at Austin, Texas, USA; evouga@cs.utexas.edu

## Abstract

What makes a disentanglement puzzle a puzzle? In this paper, we explore this question from a geometric perspective. We study a set of wire puzzles—disentanglement puzzles consisting of two pieces formed by bending thick metal wire—and wire non-puzzles, contrast their geometric properties both in 3D space and in their 6D configuration space, and propose several qualitative and quantitative measures of how “puzzly” a wire puzzle is.

## Introduction

Rigid disentanglement puzzles are fascinating brainteasers consisting of two or more rigid (usually metal) pieces. The puzzle begins in a tangled state, and the goal is to separate the pieces through a sequence of rigid motions; interest in these puzzles comes from the juxtaposition of the small number of pieces and the complex, unintuitive steps needed to disentangle them. The prototypical rigid disentanglement puzzle is the *alpha puzzle*, shown in Figure 1. Separating its two pieces requires a twist after aligning the gaps in each loop; despite the geometric simplicity of the pieces, the required twist motion is counter-intuitive. The same is not true for, e.g., separating a nut screwed onto a bolt: unscrewing the nut requires a long, twisty screw motion, but few humans would characterize this process as *solving a puzzle*.




Figure 1: Alpha Puzzle




Figure 2: A Collection of Digitized Disentanglement Puzzles.




Figure 3: Example Non-puzzle Problems


What makes the alpha puzzle a *puzzle*, and the nut-and-bolt a non-puzzle? In this paper, we explore this question from a geometric perspective. We study a set of *wire puzzles*—disentanglement puzzles consisting of

two pieces formed by bending thick metal wire, see Figure 2—and a set of wire non-puzzles (Figure 3), contrast their geometric properties both in 3D space and in their 6D configuration space, and propose several qualitative and quantitative measures of how “puzzly” a wire puzzle is. In addition to broadening our understanding of the nature of challenging motion planning tasks like rigid puzzle disassembly, our contributions lay the groundwork for future research on optimizing puzzle difficulty or improving the performance of motion planning algorithms on challenging disassembly tasks.

### Background

Solving rigid disentanglement puzzles is a special case of motion planning, a problem that has received widespread attention across robotics, computer graphics, computational geometry, and mathematics thanks to its applications to robot navigation and manipulation tasks. LaValle’s textbook [5] provides an excellent overview of this area. In motion planning problems, the physical space where the moving object(s) and obstacles reside is called the *workspace* and the space of all transformations of the moving objects (their *configuration* or *state*) is the *configuration space* or C-space. For the wire puzzles we consider in this paper, the workspace is  $\mathbb{R}^3$ . We may assume, without loss of generality, that one of the two wire pieces remains fixed in space while solving the puzzle; the C-space for wire puzzles is therefore  $SE(3) \cong \mathbb{R}^3 \times SO(3)$ , the six-dimensional space of rigid motions of the mover (with three location degrees of freedom, and three degrees of freedom denoting orientation). The C-space can be partitioned into two disjoint subsets,  $C_{\text{free}}$  and  $C_{\text{obs}}$ , the set of collision-free and colliding states respectively, with boundary  $\partial C_{\text{free}}$  between them. Solving a wire puzzle entails finding a path that lies entirely within  $C_{\text{free}}$  and connects a given point (the *start state*) to any point in  $C_{\text{free}}$  representing the two pieces having separated sufficiently far apart (the *goal states*).

Figure 4 illustrates these concepts for a 2D motion planning problem. The workspace, on the left, contains a small red rectangular body (the moving piece) trapped within a larger, fixed cyan body. The C-space of the red rectangle is three-dimensional, with each configuration  $(x, y, \theta)$  specifying its position and orientation.  $C_{\text{free}}$  can be visualized as a volume<sup>1</sup> enclosed by  $\partial C_{\text{free}}$  (on the right), where we have drawn  $\theta$  as the vertical axis. Each of the two “towers” corresponds to the red rectangle moving freely within one of the two empty chambers on the left. The two thin “tunnels” connecting these towers are the configurations where the rectangle, oriented at either  $\theta = 0^\circ$  or  $180^\circ$ , slips through the gap between the chambers. Such tunnels are major challenges for motion planning, as a human or algorithm trying to explore  $C_{\text{free}}$  must first find their entrances, and then navigate through them.



**Figure 4:** 2D Workspace and  $\partial C_{\text{free}}$  in 3D C-space

### Related Work

Whereas most motion planning research focuses on the case of many agents navigating around a relatively sparse set of obstacles, disentanglement puzzles consist of only a few moving parts, but their space of motions is carefully designed to be adversarially difficult to navigate. Consequently most practical state-of-the-art motion planning algorithms struggle to solve even simple disentanglement puzzles, such as those shown in Figure 2. The most successful algorithms use heuristics to find and navigate the narrow tunnels. Zhang et al. [11] and D-plan [10] use global workspace features of the piece geometry to predict where tunnels are likely to be.

<sup>1</sup> $C_{\text{free}}$  is periodic along the rotational axis (blue), and hence the top of each “tower” is connected to its bottom.

There has also been recent interest in the inverse problem of *designing* puzzles [9], including sliding puzzles made of interlocking pieces [1], Rubik’s-cube-type twisting puzzles [7], dissection puzzles [8], and 3D jigsaw puzzles [2]. For some of these puzzle types, there is a straightforward quantitative notion of difficulty (such as the number of sliding moves required to solve interlocking burr puzzles) which can be optimized to create challenging puzzles. To our knowledge, however, there is no previous work on designing challenging wire puzzles or analyzing what makes a wire puzzle a puzzle.

## Overview

We study a set of wire puzzles and non-puzzles. The puzzles (originally from Teenitor’s “10-piece Metal Iron Brain Teaser IQ Test Assembly & Disentanglement Puzzle Toy” box set and digitized by Zhang et al. [11]) are shown in Figure 2. We designed several non-puzzles ourselves, shown in Figure 3.

- In **four pipes V1**, one piece starts nestled within the other, but easily slips through any of several gaps.
- In **four pipes V2**, the red piece is larger and now no longer fits through the cyan piece’s gaps. However, disentanglement requires only translating the red piece along a straight path.
- **Curved pipes** is a slightly more complex variant of the previous designs: the red piece must slide on a curved trajectory to escape. The puzzle is still trivial, however, since the red piece is constrained from moving in any way other than along the solution path.
- **Spiral with lid** is a wire-puzzle analogue of a nut and bolt; solving this puzzle requires a long but obvious screw motion to escape the spiral. (**Spiral** shows the same non-puzzle with the endcaps removed, revealing the moving red piece and interior geometry of the spiral.)
- **Wide alpha-z** is a modified version of the true **alpha-z** puzzle: the gap formed by each piece’s loop has been widened to allow the pieces to slip through each other without performing the tricky twisting motion required originally.

### *Indicators of Puzzly-ness*

As discussed above, motion planning problems that involve navigating through regions of  $C_{\text{free}}$  that look like narrow tunnels are particularly challenging, and indeed all of the puzzles shown in Figure 2 require passing through at least one narrow tunnel to solve. But tunnels alone are not sufficient to make a puzzle: **curved pipes**, for example, is one long tunnel (the red piece is constrained to slide along a thin one-dimensional tube in  $C$ -space) but straightforward to solve. Rather, a puzzle is difficult if the solution path requires passing through a sequence of tunnels separated by *bubbles*: regions where  $C_{\text{free}}$  expands to have large diameter in  $C$ -space. These bubbles correspond to configurations where the two puzzle pieces are loosely entangled, and where the large amount of freedom in how to jiggle the pieces obfuscates the next move required (tunnel to enter) to solve the puzzle. We hypothesize that looking for this *tunnel-bubble structure* can effectively discern puzzles from non-puzzles. Second, we observe from the solutions to the puzzles that they all involve twisting moves that entwine segments of the two pieces around each other. We propose that absence of this *twistiness* is a strong indicator of a non-puzzle.

### *Quantitative Approach*

We propose two quantitative metrics for probing existence of tunnel-bubble structure (the *visibility volume* and its anisotropic refinement, the *explained variance ratio*) and one for measuring twistiness (the *Gauss linking integral*), which we describe below. For each puzzle, we first compute a solution to the puzzle (using the algorithm of Zhang et al. [11]), and evaluate each metric at all times along the solution path. The computed solution paths are not necessarily optimal (finding *any* solution is often challenge enough!); instead our approach samples a “typical” path a solver might try (including dead ends and backtracking) and the tunnel-bubble structure and twistiness encountered along the way.

## Visibility Volumes

In order to locate potential bubble-tunnel structures along a path in C-space, we introduce a metric called *visibility volume*. Given a solution path  $\gamma(t) : [0, 1] \rightarrow C_{\text{free}}$ , the goal of this metric is to measure how much “clearance” there is before hitting  $C_{\text{obs}}$  around each point on the path: large values indicate that the path is moving through the interior of a bubble, and small values indicate a tunnel.

Given a point  $\mathbf{p} = \gamma(t_0) \in C_{\text{free}}$  on the solution path and a tangent vector  $\mathbf{v} = (\mathbf{v}_{\text{trans}}, \mathbf{v}_{\text{rot}})$  to  $SE(3)$  at  $\mathbf{p}$  (representing a screw motion with velocity  $\mathbf{v}_{\text{trans}}$  and angular velocity  $\mathbf{v}_{\text{rot}}$  of the moving piece), we define the *visibility distance* of  $\mathbf{v}$  to be

$$\begin{aligned} \|\mathbf{v}\|_{\text{vis}} &= \min(1, t_{\text{col}}), \\ t_{\text{col}} &= \inf \{t \mid \exp_{\mathbf{p}}(t\mathbf{v}) \in C_{\text{obs}}\}, \end{aligned}$$

where<sup>2</sup>  $\exp_{\mathbf{p}}(t\mathbf{v})$  is the configuration of the moving piece reached by starting at configuration  $\mathbf{p}$ , rotating the piece about the normalized axis  $\widehat{\mathbf{v}}_{\text{rot}}$  (in its body coordinates) by angle  $t\|\mathbf{v}_{\text{rot}}\|$ , and then translating the piece (in global coordinates) by  $t\mathbf{v}_{\text{trans}}$ . The visibility distance measures the first time  $0 \leq t_{\text{col}} \leq 1$  along this screw motion that the moving piece collides with the stationary piece, or 1 if no collision occurs.

Once the puzzle is solved and the two pieces are significantly separated,  $\|\mathbf{v}\|_{\text{vis}}$  is equal to the maximum value of 1 for every direction  $\mathbf{v}$ . Before then,  $\|\mathbf{v}\|_{\text{vis}}$  will be larger within bubbles, and smaller within tunnels (except for very particular directions  $\mathbf{v}$  that happen to align with the direction of the tunnel). This observation motivates the visibility volume,

$$\text{Vis}(\mathbf{p}) = \frac{1}{N} \sum_{i=1}^N \|\mathbf{v}_i\|_{\text{vis}}, \quad (1)$$

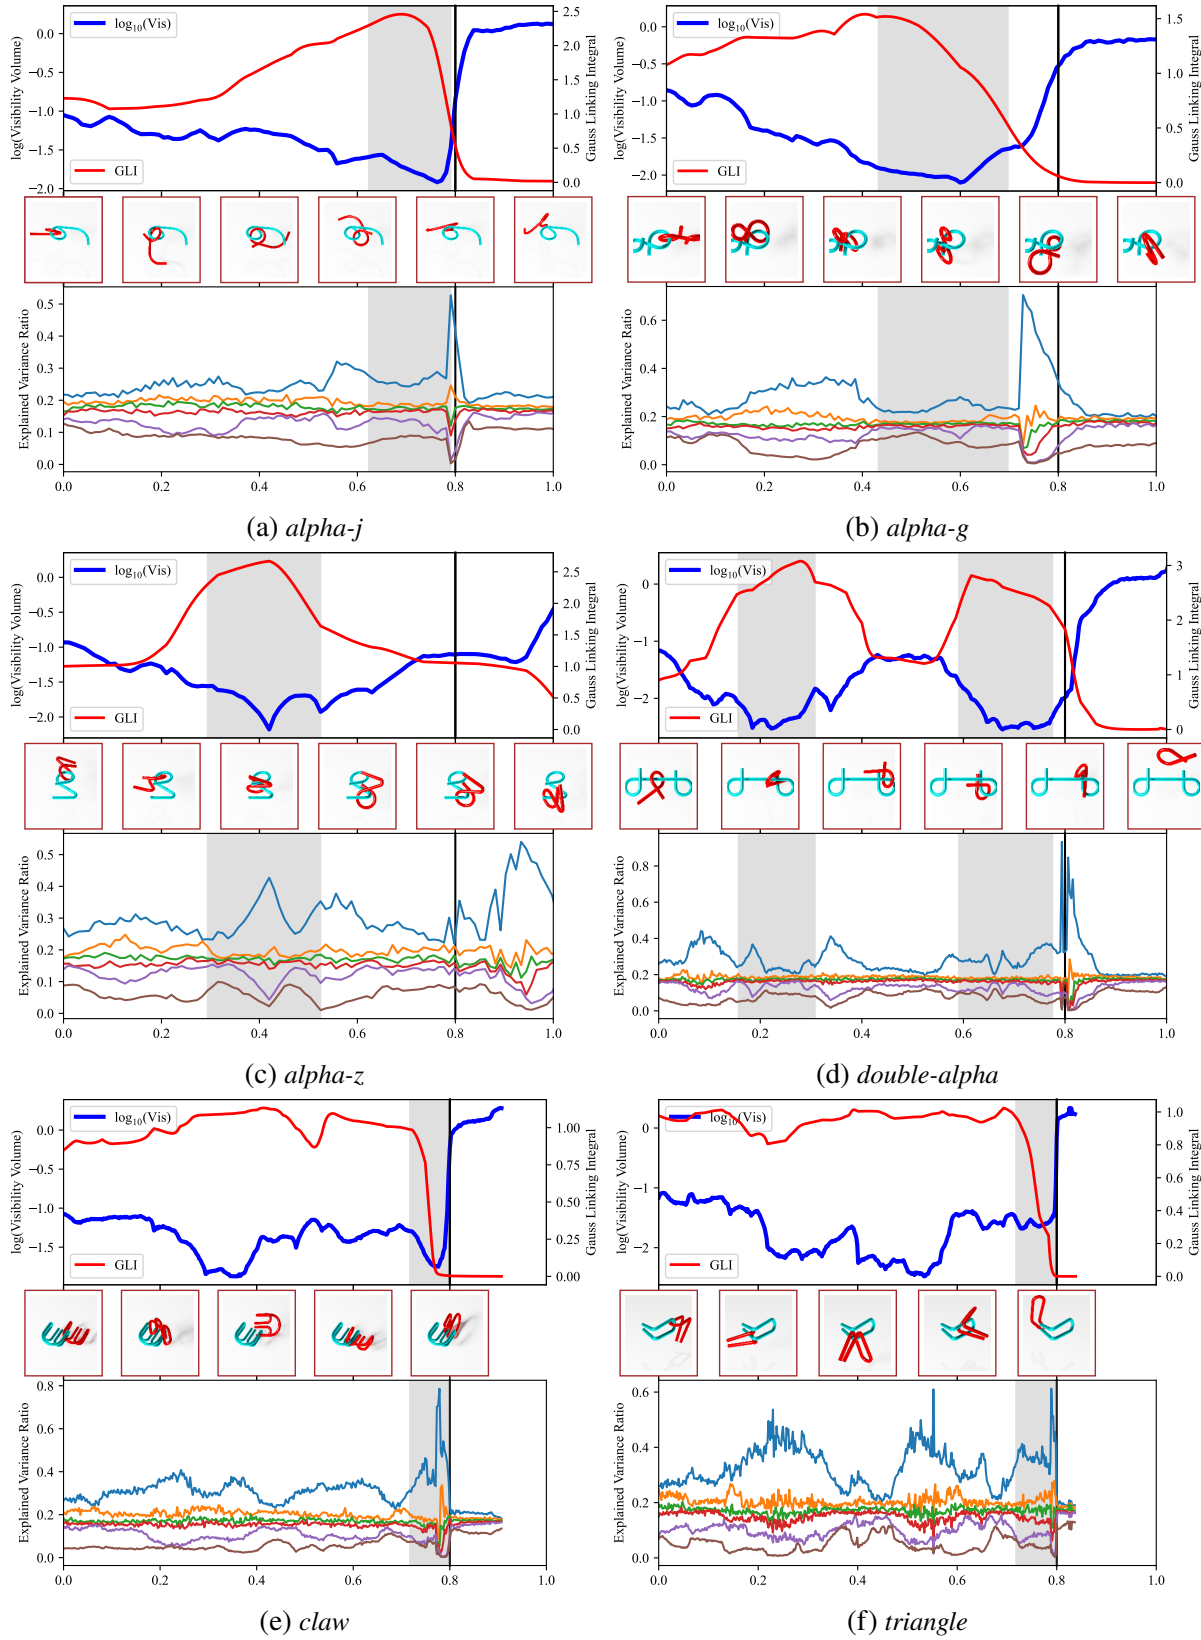
which averages the visibility distance of  $N$  random tangent directions  $\mathbf{v}_i$ . We sample  $\mathbf{v}_i$  by choosing a translation vector uniformly at random from the sphere centered at  $\vec{0}$  of radius  $d_{\text{max}} = 2(r_1 + r_2)$ , where  $r_1, r_2$  are the radii of the bounding spheres around the two pieces, and the rotation axis-angle by sampling a quaternion uniformly from the unit 4-ball. These choices guarantee that the visibility volume increases for larger bubbles, and achieves its maximum value only once the puzzle is solved (since in a tangled state, it is impossible to translate the moving piece by  $d_{\text{max}}$  without causing a collision).

In bubbles, Equation (1) approximates the average visibility distance with samples; however in tunnels, as mentioned above, most of the visibility distance is concentrated in a few tangent directions that correspond to motion forward and backward through the tunnel. When calculating visibility volume on a segment of the solution path with tangent vector  $\mathbf{v}_{\text{path}}$ , we augment the random samples in Equation (1) with random perturbations of  $\mathbf{v}_{\text{path}}$ ; each perturbation is a random rotation by up to 15 degrees of the translation direction and rotation axis of  $\mathbf{v}_{\text{path}}$ . We use 4,096 random samples and 512 perturbations of  $\mathbf{v}_{\text{path}}$  for all of our experiments, and weight each group separately to account for their differing density in the tangent space of  $SE(3)$ .

## Explained Variance Ratio

Although the visibility volume is large in bubbles and small in tunnels, it is a crude instrument: the visibility volume will be small even in regions of  $C_{\text{free}}$  that are not one-dimensional tunnels, but where only a few directions of motion are constrained, such as the **spiral** non-puzzle.

<sup>2</sup>The exponential map  $\exp_{\mathbf{p}}(t\mathbf{v})$  computes the rigid motion you get by starting at  $\mathbf{p} \in SE(3)$  and moving along a geodesic (in this case a screw motion) in the direction  $t\mathbf{v}$ .



**Figure 5:** Visibility Volume, Gauss Linking Integral, and Explained Variance Ratio

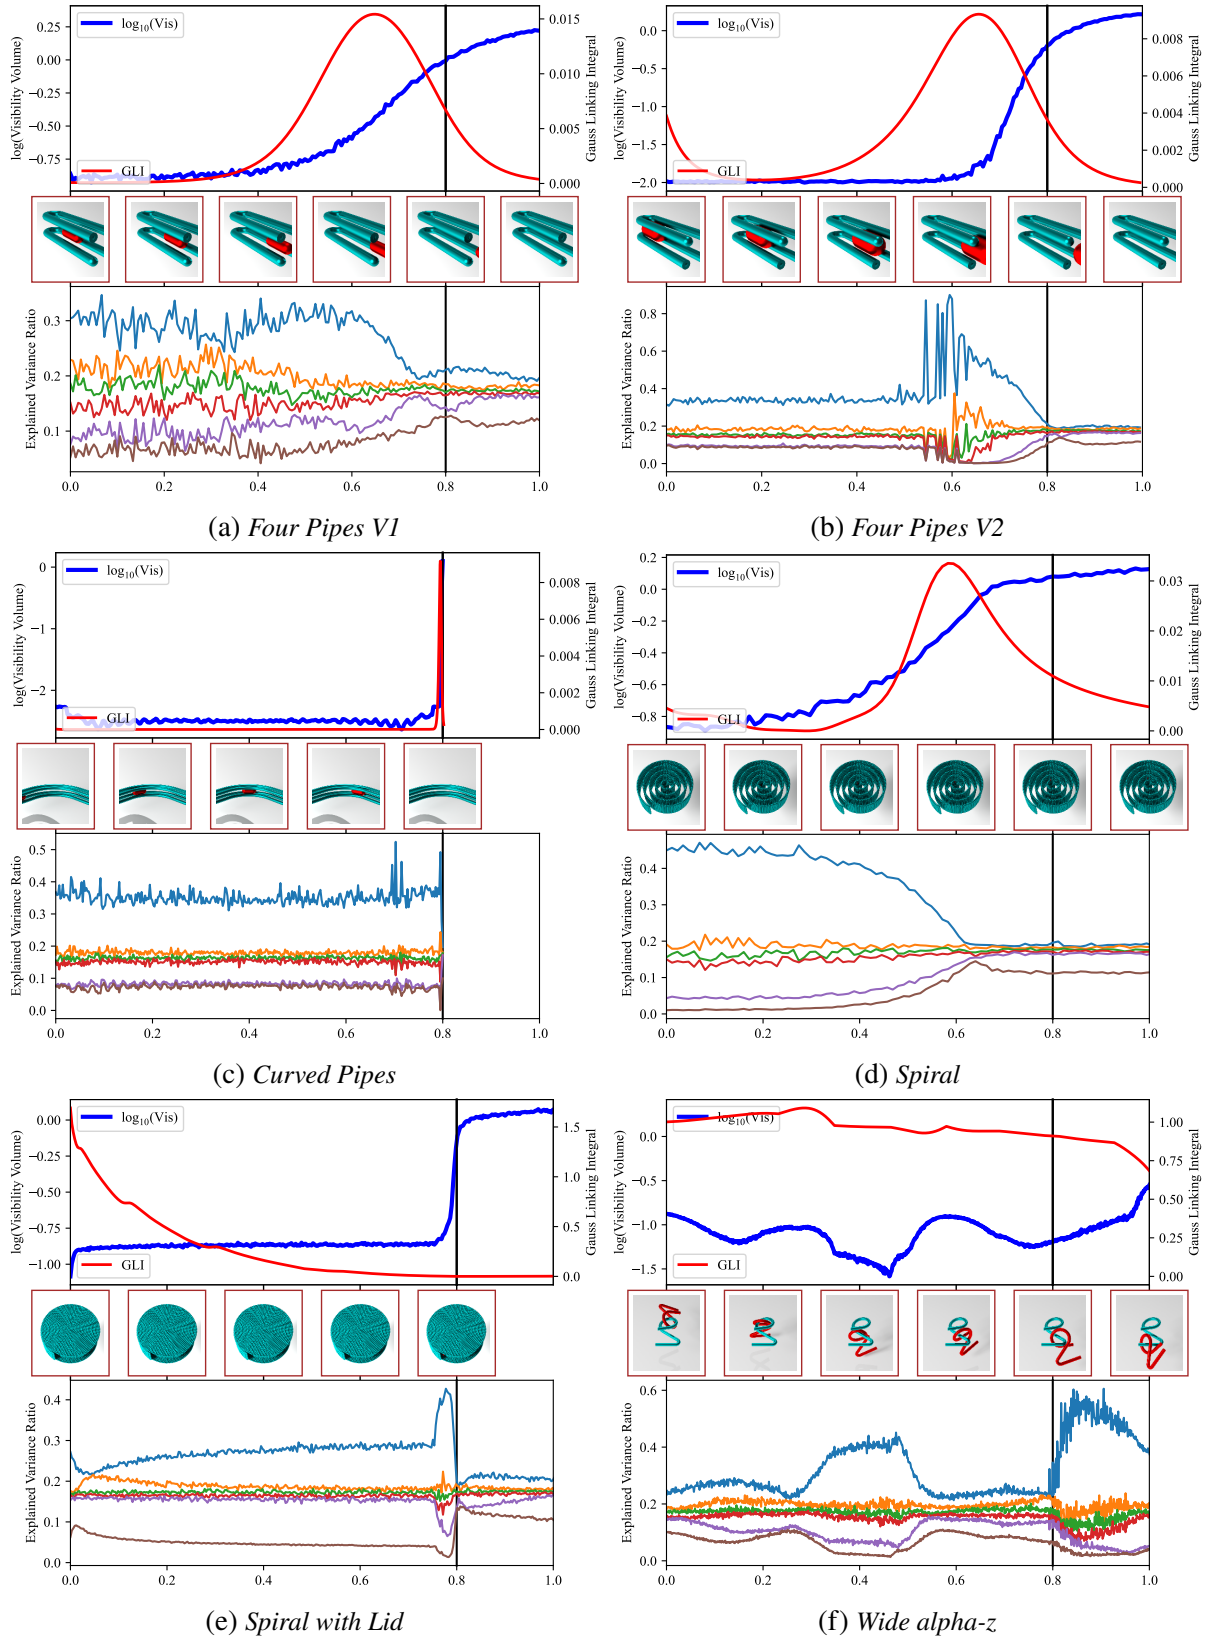


Figure 6: Quantitative Results from Non-Puzzles

We can refine the visibility volume measure  $\text{Vis}(\mathbf{p})$  by examining not just the average visibility distance of tangent vectors at  $\mathbf{p}$ , but also how the visibility distances are clustered in the tangent space. We quantify an anisotropic generalization of Equation (1) by performing *Weighted Tangent Principle Component Analysis* (WtPCA) [6], a.k.a. *Principle Geodesic Analysis* (PGA) [3], on the set of vectors  $\|\mathbf{v}_i\|_{\text{vis}} \mathbf{v}_i$ , where the samples  $\mathbf{v}_i$  are the same as those used to compute  $\text{Vis}(\mathbf{p})$ . We weight the random versus perturbed samples to account for differing tangent space density, as above, and also weight the translation and rotation coordinates of each sample so that the singular values corresponding to the six principal components are identical once the puzzle is solved as the pieces well-separated.

Similar to the classical PCA, the *explained variance* of WtPCA can be defined as  $\sigma_j = \lambda_j^2 / (N - 1)$ , where  $\lambda_j$  are the singular values of the sample covariance matrix. Intuitively, the explained variances describe the lengths of the principal axes of a “best-fit ellipse” to the geometry of  $C_{\text{free}}$  near  $\mathbf{p}$ . Therefore the *explained variance ratios*  $\mu_j = \sigma_j / \sum_{j=1}^6 \sigma_j$  inform us of the anisotropy of  $C_{\text{free}}$  near  $\mathbf{p}$ . In a one-dimensional narrow tunnel, we expect  $\mu_1$  to be large and the other five explained variance ratios to be small. In a bubble, we expect all explained variance ratios to be roughly equal.

### Gauss Linking Integral

Finally, we propose a way to measure how twisted two wire puzzles are about each other. Given two closed curves  $\gamma_1(s)$  and  $\gamma_2(s)$ , the linking number is a topological invariant that measures how many times the curves wind around each other. The linking number can be computed using the Gauss linking integral,

$$\text{GLI}(\gamma_1, \gamma_2) = \frac{1}{4\pi} \int_{\gamma_1} \int_{\gamma_2} \frac{(\gamma_1' \times \gamma_2') \cdot (\gamma_1 - \gamma_2)}{\|\gamma_1 - \gamma_2\|^3}. \quad (2)$$

For two curves that *aren't* closed (such as for instance the centerlines of two entwined wire puzzles), we can still compute the above integral as a (no longer integer) measure of how much the two curves twisting about each other [4]. For puzzles, we expect the GLI to be large in the starting (tangled) state, to change as the two pieces move relative to each other through the tunnel-bubble structure, and to end up zero after the puzzle is solved. A GLI that stays small throughout the solution path suggests a non-puzzle.

### Experiment Results

Figure 5 and 6 list the quantitative analysis results for puzzles and non-puzzles respectively. In each graph, the top image plots the logarithm of the visibility volume and the Gauss linking integral, and the bottom image shows the explained variance ratios of all six singular vectors from the WtPCA. The center image visualizes the poses of pieces at different configurations along the solution path. The three images share the same  $x$ -axis, which is normalized so that  $x = 0.8$  is always considered as “solved” (for puzzles) or “separated” (for non-puzzles). The grey areas in Figure 5 further mark the tunnel region located manually by humans.

Figure 5 shows the visibility volume can locate the bubble-tunnel structures pretty reliably. In all puzzles, the visibility volume starts high and decreases to roughly  $10^{-1}$  times its initial value at known locations of narrow tunnels along the solution path. Furthermore, for the claw puzzle the visibility volume identifies (via a valley in the visibility volume plot around  $t = 0.3$ ) an additional narrow tunnel location that we did not notice at first. A similar additional narrow tunnel identified by the visibility volume for the triangle puzzle reveals a limitation of our analysis: the computed path  $\gamma(t)$  for this puzzle is not optimal, and goes into a dead-end at  $t = 0.3$ , then backs up and goes through another non-essential narrow tunnel ( $t \approx 4.5$ ) before solving the puzzle.

The explained variance ratio of alpha-z, double-alpha, claw, and triangle successfully locates low-dimensional regions in  $C_{\text{free}}$  at  $t = 0.4$ ,  $t = 0.2, 0.7$ ,  $t = 0.35, 0.75$ , and  $t = 0.75$  respectively. The alpha-g

puzzle at  $t = 0.725$  is an interesting case: the smaller ring of the red moving piece slips through the larger ring of the teal fixed piece. The pieces lose degrees of freedom while their rings are nestled, as detected by the explained variance ratios, though not enough to form a one-dimensional narrow tunnel. We also observed that the first explained variance ratio increases for most puzzles when the moving pieces enter a large open region in  $C_{\text{free}}$ . The exit of a tunnel forms a mushroom-like shape that allows the moving piece to move into the open space along several directions in  $SE(3)$  while limiting its movement in others. The explained variances along different principle components of samples of this “mushroom” have an uneven distribution.

Comparing the puzzle plots to the corresponding plots for non-puzzles in Figure 6 partially verifies that tunnel-bubble structure (especially as measured by visibility volume) and twistiness (as measured by GLI) define a puzzle. The visibility volume, explained variance ratio, and GLI plots for most non-puzzles lack the peaks and valleys seen in puzzles. One exception is the wide alpha-z: none of our metrics can differentiate this non-puzzle from the original alpha-z puzzle. At  $t = 0.4$  the red piece slips through a narrow gap between the wires of the teal piece, without the unintuitive “twist”; visibility volume labels this region a tunnel, suggesting additional criteria are needed to fully distinguish puzzles from non-puzzles.

### Discussion and Future Directions

Our interest in distinguishing puzzles from non-puzzles stems from our desire to design new puzzles. We believe that geometric features of the puzzle shapes and their C-space, such as visibility volume, visibility volume shape, and GLI can be useful as objective functions in an optimization. One might also consider puzzles with more than two pieces, requiring an analysis in a higher dimensional C-space. Finally, our current measures cannot detect all differences between a puzzle and non-puzzle. In the future, more measures, like the curvatures of the tunnel medial axis, could be tested to augment our current suite.

### References

- [1] R. Chen, Z. Wang, P. Song, and B. Bickel. “Computational Design of High-level Interlocking Puzzles.” *ACM Transactions on Graphics (SIGGRAPH 2022)*, vol. 41, no. 4, 2022, pp. 150:1 – 150:15.
- [2] G. Elber and M.-S. Kim. “Synthesis of 3D Jigsaw Puzzles over Freeform 2-manifolds.” *Computers & Graphics*, vol. 102, 2022, pp. 339–348.
- [3] P. T. Fletcher, C. Lu, S. M. Pizer, and S. Joshi. “Principal geodesic analysis for the study of nonlinear statistics of shape.” *IEEE transactions on medical imaging*, vol. 23, no. 8, 2004, pp. 995–1005.
- [4] E. S. Ho and T. Komura. “Character Motion Synthesis by Topology Coordinates.” *Computer Graphics Forum*, vol. 28, no. 2. Wiley Online Library, 2009, pp. 299–308.
- [5] S. M. LaValle. *Planning Algorithms*. Cambridge University Press, 2006.
- [6] N. Miolane, N. Guigui, A. Le Brigant, J. Mathe, B. Hou, Y. Thanwerdas, S. Heyder, O. Peltre, N. Koep, H. Zaatiti *et al.* “Geomstats: a Python Package for Riemannian Geometry in Machine Learning.” *The Journal of Machine Learning Research*, vol. 21, no. 1, 2020, pp. 9203–9211.
- [7] T. Sun and C. Zheng. “Computational Design of Twisty Joints and Puzzles.” *ACM Trans. Graph.*, vol. 34, no. 4, Jul. 2015.
- [8] K. Tang, P. Song, X. Wang, B. Deng, C.-W. Fu, and L. Liu. “Computational Design of Steady 3D Dissection Puzzles.” *Computer Graphics Forum*, vol. 38, no. 2, 2019, pp. 291–303.
- [9] O. van Deventer *et al.* “That Is Not Art, It Is a Puzzle!” *Bridges 2019 Conference Proceedings*. Tessellations Publishing, 2019, pp. 1–8.
- [10] L. Zhang, X. Huang, Y. J. Kim, and D. Manocha. “D-Plan: Efficient Collision-Free Path Computation for Part Removal and Disassembly.” *Computer-Aided Design and Applications*, vol. 5, no. 6, 2008, pp. 774–786.
- [11] X. Zhang, R. Belfer, P. G. Kry, and E. Vouga. “C-Space Tunnel Discovery for Puzzle Path Planning.” *ACM Trans. Graph.*, vol. 39, no. 4, aug 2020.