Visualization of Highly-Dimensional Robot Shape Spaces

Capprin Bass, Brett Stoddard

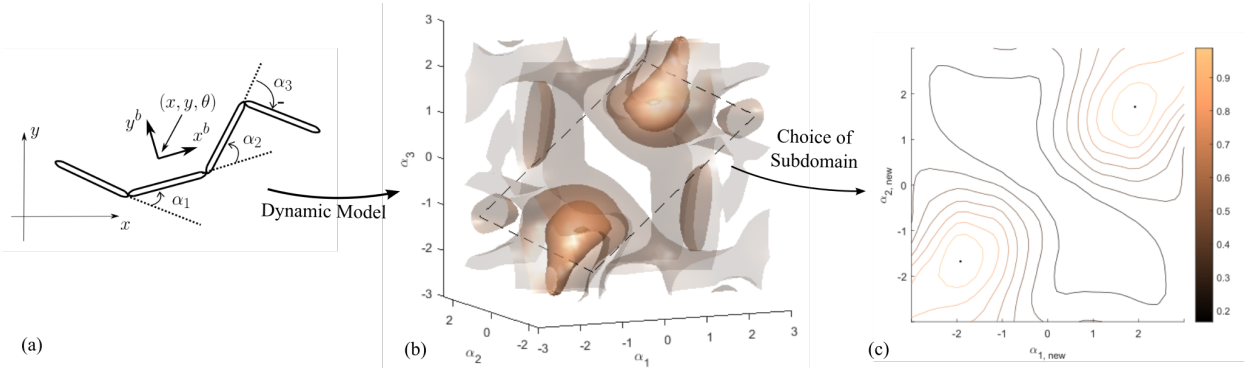

Fig. 1. Visualization and dimensionality reduction for a planar robot system. (a) The four link swimmer resides in the plane, has three shape modes, and position encoded with translation (x,y) and rotation θ. Figure originally published in [1]. (b) The constraint curvature function (CCF) represents magnitude of displacement with respect to configuration $\tilde{\alpha}$. This field represents the magnitude of body rotation θ at any point in the space. (c) By selecting a subdomain of the full shape space, we create a new, reduced shape space. Optimization of the subdomain can preserve the majority of dynamics while reducing the complexity of the system.

Abstract— Visualization of robot behavior is essential; however, current methods are impractical for robots with more than three degrees of freedom. In this work, we summarize and replicate methods for examining the gaits of three and four link swimmers (with 2 and 3 degrees of freedom, respectively) by visualizing their constraint curvature functions. Additionally, we propose and explore a method for selecting an optimal "slice" of the shape space to express and visualize robot dynamics, effectively reducing the dimensionality of the system. We apply this method to two example systems, and discuss applications to higher dimensional systems.

Index Terms—Scalar field visualization, Dimensionality reduction.

1 Introduction

In the field of robotics, systems are designed to move with one or several degrees of freedom. Each degree of freedom may be actuated, or move passively; the dynamics of the system regulate the capabilities of each degree of freedom. Commanded motions along each degree of freedom displace the robot through the world. For kinematic systems (like robot arms), displacements are predictable, and simple controllers can be used; in contrast, dynamic systems require a deliberate control strategy. Force feedback from either the environment or the robot itself destroys the performance of basic controllers. A number of approaches offer improved controllers, including those from the machine learning, geometric mechanics, and optimization communities.

Despite the availability of alternative control strategies, modern robot systems still often suffer from the "curse of dimensionality." Robots with many degrees of freedom, like the four link swimmer in Fig. 1, inherently have a large space of control strategies. The search space is often so large that algorithms are either incomplete or suboptimal, resulting in dissatisfying control. Even if a control strategy is found, it is difficult to explain intuitively, as high-dimensional dynamics are challenging to visualize. If the space is reduced, then control can be achieved with existing approaches, and current visualization techniques may be used to explain the control approach.

In this paper, we take several steps to reduce the dimensionality of robot shape spaces. We begin by appling existing scientific visualization techniques to identify valuable regions of example shape spaces. These techniques include gradient, contour, and critical point visualizations in 2D, as well as an isosurface visualization in 3D; examples are shown in Fig. 1(b-c). All visualizations take advantage of the constraint curvature function (CCF) [3], a geometric-mechanical tool that encodes important regions of the robot shape space. Visualizations of the CCF are complemented by our own dimensionality reduction approach, which uses a linear subdomain to capture valuable regions of the shape space, as demonstrated in Fig. 1(c). By defining this subdomain, we create a reduced set of robot shape modes over which we can plan robot motion. Our lower-dimensional system maps bijectively to the full shape space, so control policies still apply to the original system. The linear subdomain is also much easier to visualize, lending intuition about system behavior and potentially valuable control policies.

The remainder of this paper is organized as follows. In §2, we describe the details of the underlying geometric-mechanical model, as well as the chosen visualization techniques. In §3, we define our dimensionality reduction technique. In §4, we present CCF visualizations for two example systems, as well as their reduced order models. In §5, we discuss the limits of our method, and comment on future work.

2 Background

Our dimensionality reduction technique depends on results from the geometric mechanics community, as well as established scalar field visualization techniques. Here, we describe the details behind both.

2.1 Geometric Mechanics

Geometric Mechanics refers to a body of research in which system geometry is used to make conclusions about system mechanics. The geometric mechanics community has established a framework and set of tools for analyzing the behavior of dynamic systems; this framework is the basis for work done in this paper.

We assert that robot systems have a shape space R and position space G. The shape space contains all possible ways the system can deform

- *E-mail: basscap@oregonstate.edu.*
- *E-mail: stoddabr@oregonstate.edu.*

itself; shape modes $[\alpha_1, \alpha_2, \cdots, \alpha_n]^T \in R$ refer to the independent degrees of freedom of the system. The position space encodes the possible locations $g \in G$ the robot can be in the world. For planar systems, the special Euclidean group, $G = SE(2)$, is a common choice of position space that supports translation and rotation. Hence, $g = [x, y, \theta]^T$.

Two structures are commonly used to map from shape to position space. The *local connection*, $\mathbf{A}$, is a linear map from shape velocities to body (position) velocities [1] [2]:

$$\dot{g}_b = \mathbf{A}(r)\dot{r}. \tag{1}$$

The local connection is often decomposed into a set of vector fields, describing how velocity $\dot{r}$ at a specific shape r results in x, y, or θ displacement through the world. An alternative structure is the *constraint curvature function*, or CCF, which accounts for local curvature of the shape space with a local Lie bracket [3]:

$$\text{CCF} = d\mathbf{A} + [\mathbf{A}_1, \mathbf{A}_2]. \tag{2}$$

The constraint curvature function is a 2-form, and describes the variation in displacement from infinitesimal shape changes. By taking the magnitude of the CCF, we produce a scalar field that estimates "valuable" regions of the shape space; an example is shown in Fig. 6. High-magnitude regions of this field encode effective collections of shapes for displacement, and are strong regions for motion planning.

2.2 Scalar Field Visualization

Visualization of the CCF-norm allows us to identify high-value regions of the shape space for motion planning. Visualization techniques also have to scale with the dimensionality of the system; we ought to be able to identify conserved structures for up to three shape variables. To effectively visualize the CCF-norm of robot systems, we used color gradient, contour, critical point, and isosurface visualizations.

In the 2D case, gradient, contour, and critical point visualizations tell us about important features in the data. The gradient visualization, when a colorbar is included, provides us intuition about the intermediate regions of the CCF. In contrast, the contour visualization allows us to draw conclusions about the geometry of the CCF, and where critical values (in particular, local maxima) may be found. Critical point detection is also used to compute exactly the highest-value points in the shape space. In this paper, we also combined these visualization techniques, producing gradient-colored contours with rendered maxima.

In the 3D case, visualization is made more challenging because of higher dimensionality and the possibility of occlusions in the field. We chose an isosurface visualization, which can be thought of as the 3D analog to contours. Levelsets in the 3D data are rendered as surfaces, which are colored according to a scalar mapping. To minimize the effect of occlusions, isosurfaces are also given transparency according to their scalar value. This combined approach produces strong intuition about the data and its present features.

3 Method

Our method selects a slice from the shape space such that it can be visualized in two or three dimensions. First, we define a reduced-dimension shape space as a linear subdomain of the full shape space. Then, we optimize the orientation of the new shape space such that it captures as much displacement as possible. Our approach is based on the assumption that an intrinsic, lower dimensional subsystem exists which preserves the most important aspects of the larger system; if this isn't the case, then we destroy valuable structures in the shape space and reduce the capability of the system.

3.1 Linear Subdomain

For a shape space of dimension D (having D shape variables), we reduce the dimension to $D-1$ by taking a linear subdomain of the shape space. For a 2D shape space, this subdomain is a line; for a 3D shape space, it is a plane. The CCF can be projected onto this new domain, where we can search for control policies or visualize system dynamics. Because this domain is embedded in the shape space, the new shape modes are trivially mapped onto the full shape space, and control policies are relevant to the original system.

The linear subdomain requires that we define the size, position, and orientation of the new shape space within the full domain. In our implementation, we chose to maintain the initial range of the shape variables, and fix the position at the origin of the shape space. The orientation of the new subdomain can then be used to select for valuable regions of the shape space, based on some metric for performance. With an optimization approach, we can pick an orientation that maximizes system performance, producing an effective dimensionality reduction.

3.2 Optimization

Our method extracts the optimal subdomain orientation by maximizing the integral of the field enclosed within the subdomain. The objective function maximizes possible displacement within the reduced shape space. In other words, we select the subdomain that satisfies the following optimization problem:

$$\begin{aligned} \max_{R_{D-1}} \quad & \int_{R_{D-1}} \text{CCF}(\vec{\alpha})\, d\alpha_1 \wedge \cdots \wedge d\alpha_D \\ \text{s.t.} \quad & R_{D-1} \in R_D, \end{aligned} \tag{3}$$

where R_{D-1} is a subdomain of the robot's original shape space R_D.

Equation 3 presents a general form of the optimization problem; however, optimizing over all possible subdomains is intractable. As we only consider centered, linear subdomains inside the shape space, R_{D-1} can be parametrized in terms of the angle(s) $\vec{\theta}$ of rotation inside of R_D. By parametrizing the subdomain with a vector of angles $R_{D-1}(\vec{\theta})$, the optimization problem now becomes

$$\begin{aligned} \max_{\vec{\theta}} \quad & \int_{R_{D-1}(\vec{\theta})} \text{CCF}(\vec{\alpha})\, d\alpha_1 \wedge \cdots \wedge d\alpha_D \\ \text{s.t.} \quad & \vec{\theta} \in \mathbb{R}^{\frac{D(D-1)}{2}}. \end{aligned} \tag{4}$$

3.3 Implementation

In our implementation, we calculate the integral in Eqn. 4 using the trapezoidal method, and find the optimal orientation by sampling uniformly across all possible rotations. To explore this in the two and three dimensional cases, we sample from the full domain of angles to find the combination which maximizes the objective function (i.e., integration over the CCF), capturing as much system capability as possible.

By visualizing the result of this parameterization, we observed that this optimization problem is nonconvex but locally smooth, as shown for in Fig. 2. This form has implications for potential future optimization approaches beyond uniform sampling. For instance, it suggests that a gradient ascent method may be useful for finding the best subdomain, but that such an approach is likely to get stuck in a local rather than global maximum. It should also be noted that the objective function may vary wildly for different types of robots; in other words, optimizer hyperparameters may be robot-specific.

4 Results and Visualizations

The dimensionality reduction approach defined in §3 was applied to two example systems: the three link and four link swimmers. A diagram of the latter is available in Fig. 1(a); relevant visualizations and interpretations follow. Further information about these (and other) kinematic systems is available in [1–4].

4.1 Three Link Swimmer

The three link swimmer has 2 degrees of freedom, and its shape space can be represented in the plane. The CCF for the swimmer is represented as a 2D scalar field over this domain.

We first apply a gradient visualization of this field with critical points, as shown in Fig. 3. With this visualization, we can observe two high-value regions of the shape space, in the top-right and bottom-left

[1] In reality, $\dot{g}_b$ is on the Lie algebra of the position space G; in the context of this paper, this is $\mathring{g} \in \mathfrak{se}(2)$. We have simplified its expression here.

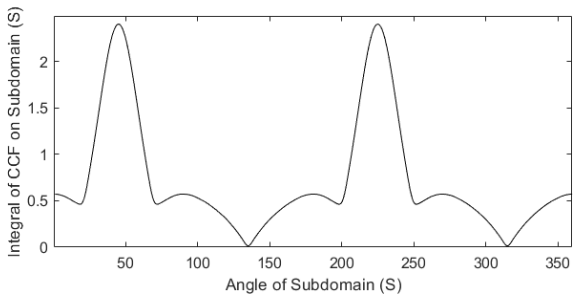

Fig. 2. Change in the integrated CCF based on subdomain angle, for a three link swimmer. Here, the subdomain is denoted as S. This figure represents the objective function for selecting the optimal subdomain. The nonconvex objective function makes optimization nontrivial.

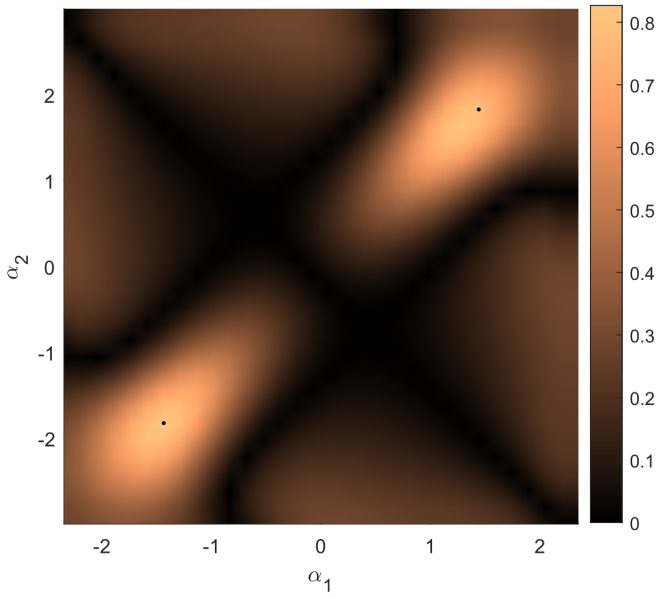

Fig. 3. CCF for the rotation of the three link swimmer, rendered with a gradient color map. Maxima are represented with black points. The two maximal corner regions in the shape space are particularly effective for rotation.

corners, respectively. Other, suboptimal regions also appear along the edges of the space, which could be used for constrained locomotion.

We also apply a contour visualization of the field in Fig. 4, with gradient coloring of the contours and a visualization of maxima. With this approach, we observe the same structures present in the gradient visualization, but can make stronger conclusions about the geometry of the CCF. In addition, this visualization contains the chosen "slice" for our optimized subdomain, which bisects the two high-value regions of the shape space. The cross-section of this subdomain (and corresponding reduction of the shape space) is available in Fig. 5; note the corresponding range and conserved high-value geometry.

4.2 Four Link Swimmer

The four link swimmer has 3 degrees of freedom, and so its shape space is represented as a volume. The CCF is a corresponding 3D scalar field over the domain, requiring a more careful visualization method.

For the four link swimmer, we apply the isosurface visualization discussed in §2; a projection of this visualization is available in Fig. 6. Although the 2D projection of the 3D geometry into the page reduces the effectiveness of the visualization, we can still identify important structures. In the extrema of the domain, we observe two high-value regions, analogous to those observed in the three link swimmer. Indeed,

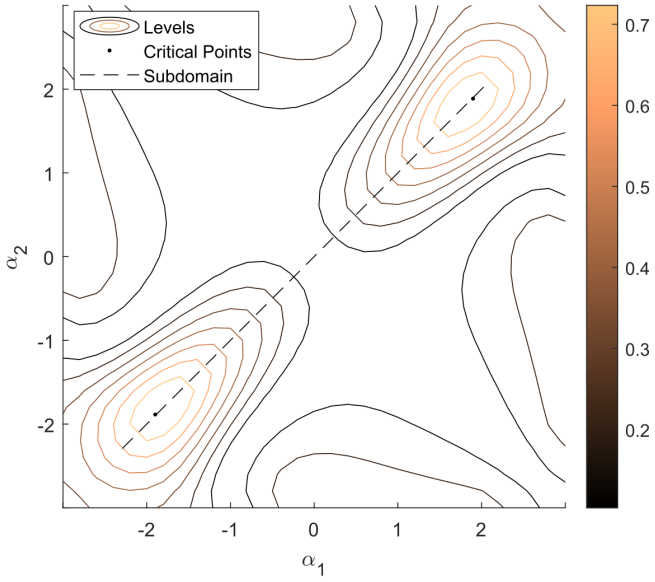

Fig. 4. CCF for the rotation of the three link swimmer, rendered with colored contours, and including critical points. An optimized subdomain is also shown, which captures as much area as possible in a one-dimensional shape space.

many structures are conserved between the swimmers, including the suboptimal extrema along the edges of the shape space.

In addition, the optimized subdomain was computed for the four link swimmer. The "slice" plane is visible in Fig. 6, and the projection of the CCF onto the plane is shown in Fig. 7. Here, the similarity between the swimmers is reinforced; as before, we see two high-value regions of the shape space. Given a relaxation of constraints on the subdomain (a larger plane, for example), the two domains may be even more similar.

Both Fig. 6 and Fig. 7 are larger reprints of Fig. 1(b-c), and are included for clarity and easier reading.

5 Discussion

Our dimensionality reduction approach, combined with effective visualization techniques, demonstrates interesting qualities of the example systems. Interpretation is done here, as well as speculation about future work and possible improvements to this approach.

5.1 Interpretation of Results

Immediately, the visualizations presented in this paper imply some amount of conserved behavior for N-link swimmers. Not only are similar structures present in the full shape spaces, but our basic dimensionality reduction implies a correspondence between the three and four link swimmers. At the moment, this claim is built only on intuition lended by our visualization techniques; however, an improved method of subdomain selection and optimization could allow for numeric validation through either field comparison or test gait execution.

This approach allows us to think of lower-dimensional systems as "approximations" on their higher-dimensional counterparts. Higher dimensional systems have the most shape modes available to them; however, a reduced order system may still be able to mostly replicate the same behavior. The extent that a simpler system is effective is application-dependent, but certain robot domains may be able to simplify existing designs or planning approaches, based on analysis and results using this technique. Alternatively, the general behavior of systems may be *visualized* in lower-dimensional space, but planning could occur in the full shape space, taking advantage of all shape modes while allowing the behavior to still be explained.

Future improvements on this method will either reinforce the fact that similar systems have conserved dynamics, or will highlight emergent artifacts of lowering the "shape resolution" of systems.

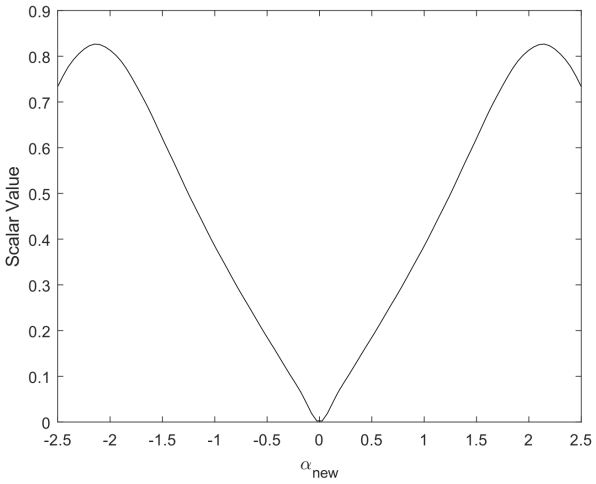

Fig. 5. Slice of the CCF for the rotation of the three link swimmer, corresponding to the optimized subdomain. This new subdomain defines a trival, one-dimensional shape space.

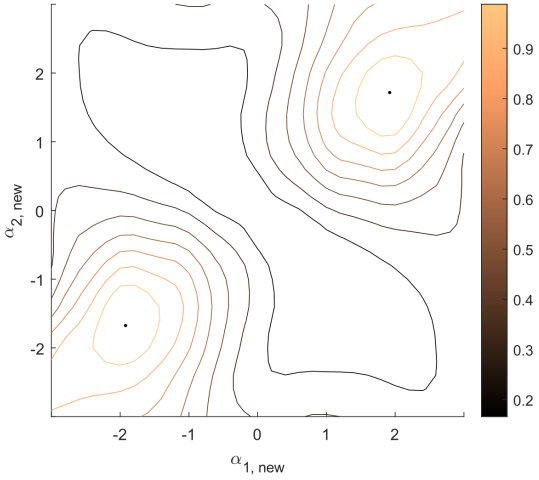

Fig. 7. Slice of the CCF for the rotation of the four link swimmer, corresponding to an optimized subdomain that defines a new, two-dimensional shape space. Note that the subdomain is analogous to that of the three link swimmer, which has equivalent dimensionality.

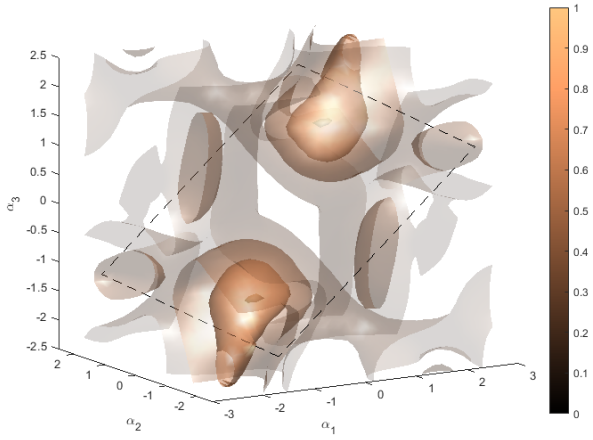

Fig. 6. CCF for the rotation of the four link swimmer, represented as an isosurface. Two high-value regions are clear, in the bottom-left and top-right corners of the shape space. An optimized choice of planar subdomain is also shown, which defines a reduced shape space.

5.2 Future Work

This work is a first step towards effective dimensionality reduction and visualization of robot system behavior. However, many simplifying assumptions reduce the effectiveness of our method; future work will remove these assumptions and improve results.

This work only addresses reduction of systems with up to three shape modes; however, this approach would be most useful for systems with at least four shape modes. The behavior of these high-dimensional systems is difficult to visualize with existing techniques, but approximations of their behavior (through subdomain optimization) could make complex systems more approachable. Extension to higher dimensions requires a more robust mathematical formulation of the embedded subdomain and associated optimization strategy.

The systems used in this paper do not have flat shape spaces, despite the cartesian coordinates used throughout this paper to visualize system dynamics. Independently actuated, rotational joints give the swimmers *toroidal* shape spaces; this can be seen in the conserved structures present in Fig. 4 and Fig. 6, cut off by the bounds of the visualization. By failing to take this underlying topology into account, we potentially miss interesting structures, and certainly fail to show all critical points. Future work ought to take into account the topology of the space to reveal these structures and perform robust critical point detection.

In a similar vein, this work also does not account for the cost of motion in the shape space. Previous work [1,4] defines cost metrics that further deform the shape space according to physical actuator limits. By incorporating the metric into this method, visualizations will better reflect behavior that the system is physically capable of performing.

Finally, this work can be improved significantly with a broader choice of subdomain, and an efficient optimization technique. The linear subdomain produces a reasonable dimensionality reduction; however, we unrealistically constrain this domain, and there are likely curved embeddings that more effectively describe system behavior, which would require more intelligent optimization approaches to find. As an example, the gradient of the field could be used to deform and improve the subdomain over a set of iterations.

References

[1] R. L. Hatton, Z. Brock, S. Chen, H. Choset, H. Faraji, R. Fu, N. Justus, and S. Ramasamy. The geometry of optimal gaits for inertia-dominated kinematic systems, 2021.

[2] R. L. Hatton and H. Choset. Connection vector fields for underactuated systems. In *Proceedings of the IEEE BioRobotics Conference*, pp. 451–456, October 2008.

[3] R. L. Hatton and H. Choset. Nonconservativity and noncommutativity in locomotion. *European Physical Journal Special Topics: Dynamics of Animal Systems*, 224(17–18):3141–3174, 2015.

[4] S. Ramasamy and R. L. Hatton. Soap-bubble optimization of gaits. In *Decision and Control (CDC), 2016 IEEE 55th Conference on*, pp. 1056–1062. IEEE, 2016.

A Division of Tasks

- **Capprin Bass** (domain expert): provided data for the project, programmed 2- and 3-D visualizations, defined objective function and updated dimensionality reduction approach, and wrote majority of paper (roughly 70%).

- **Brett Stoddard** (visualizationer): experimented with statistical dimensionality reduction - this was our original, planned approach, but details of problem formulation rendered this impossible. With change to dimensionality reduction algorithm, programmed subdomain optimizer, and wrote minority of paper (roughly 30%).

B Code

All source code and data is available for download at:
`https://github.com/Capprin/cs553-final`.