Nonprehensile Paper Folding

Danielle White
Advisor: Tomás Lozano-Pérez

***Abstract*—The nonprehensile manipulation of deformable objects is a difficult task that requires the development of a robust controller which can tolerate unknown dynamics. This work will focus specifically on the nonprehensile manipulation of paper, which, in addition to difficult to model dynamics, will also provide little force feedback. As a consequence, the control strategy will need to rely more heavily on vision and estimates of the paper's pose. This proposal describes a high level control strategy that pushes the paper in the normal direction and tracks a position relative to the paper's edge. Various controllers are discusses at varying levels of model complexity and fidelity.**

I. INTRODUCTION

Both nonprehensile manipulation and deformable object manipulation are challenging classes of problems. Nonprehensile manipulation allows the use of simpler and less specialized end effectors, but restricts the control authority the manipulator can exert on the object. "Deformable objects" as a category encompasses a wide range of objects, such as cloth and cables, but these objects present much more complex dynamics than rigid bodies. The design of control strategy that can handle deformable objects using a nonprehensile manipulator would contribute insights into both of these problem.

This thesis aims to tackle one such problem: the nonprehensile manipulation of paper–specifically, folding it. To allow folding without grasping, the paper is fixed on one edge and is left to hang on the other, so the end effector can manipulate the object by pushing it. By disallowing grasping, we force the manipulator to contend with the unknown and difficult to model dynamics of the paper. Additionally, because paper is thin and will exert very small forces on the manipulator, we will likely not be able to rely on force feedback from the end effector and instead will need to use visual feedback.

At a high level, the controller proposed here maintains a particular distance between the edge of the paper and the contact point while continually pushing in the normal direction to eventually fold the paper. A variety of control strategies are proposed that operate on different levels of modeling complexity and prior knowledge about system parameters.

The remainder of this thesis proposal is organized as follows: Section II reviews related work. Section III provides a more detailed description of the task to be executed by the controller. Section IV discusses which of the available approaches for modeling deformable objects is used in this proposal. Section V proposes several control strategies. Section VII presents a timeline for this thesis work.

II. RELATED WORK

The control strategy developed in this thesis work will rely on concepts from compliant control, where the controller exhibits forces and motions in response to the environment. Schumacher et al. provide an overview of the topic of compliant control, and breaks the field down into two subgroups: impedance control and hybrid force-position control [1]. Both strategies are ways to mediate the relationship between forces and motions as a manipulator interacts with its environment: Hogan presents the relationship as an impedance, whereby the controller imposes the physical behavior of some desired physical system (such as a mass-spring-damper system) [2]; Mason decomposes forces and motions into orthogonal direction so that each can be controlled separately [3]. While impedance control may prove a useful strategy at a later point, hybrid-force position control is more directly relevant: the dynamics of contact suggest a decomposition of end-effector dynamics into a tangential, position controlled direction and a normal, force controlled direction. De Schutter and Van Brussel give examples of other problems formulated for hybrid force-position control [4] and implement of hybrid force position control via an external force control loop wrapped around a position controller [5]. Luca and Manes apply hybrid force-position control to a dynamic environment [6].

Other works have also addressed the modeling and manipulation of deformable objects. Jiménez surveys modeling approaches and considerations for deformable objects as well as how those models may be incorporated into planning [7], but restricts scope to prehensile manipulation and works at a higher level of abstraction than this thesis entails. Chang and Padif model a linear deformable object as a series of 15 rigid links connected by spring-damper joints and demonstrate how to calibrate the simulation model to the real object, with the end result being a higher fidelity simulation which can be used to develop further control strategies [8]. Wirnshofer et al. address the parameter uncertainty inherent to the manipulation of deformable objects, and they propose a planner that constructs a search tree in belief space that is robust to inaccuracies in parameters such as friction, stiffness, and damping of manipulated objects [9].

The survey by Jiménez also specifically references the manipulation of paper [7]. However, its treatment focuses on the planning of folds (or the topic of foldablity, which is even less relevant to this problem) and abstracts away concerns with paper as a compliant object. Balkcom and Mason construct a paper folding robot that can construct flat origami pieces, but its control strategy does not account for system dynamics because of its specialized design [10]. Elbrechter et al. develop a perception and control system for a set of two dextrous hand manipulators for manipulating paper [11]. Jiang et al. use an underactuated compliant manipulator with advantageous passive dynamics to grasp paper for the task of page turning [12]. All of these examples utilize a specialized manipulator that in some way assists in the task of paper folding.

III. TASK

The end goal of this thesis is to fold a piece of paper without grasping it. To make the task feasible without a gripper, the paper starts on a platform and is fixed on one edge with the other edger hanging off the platform's edge, free to move. Fig. 1 shows the initial and goal configurations of the paper. The initial focus will be only to fold the paper such that the free edge makes contact on the other side, regardless of where the fold occurs, but execution of a more precise fold may be developed further once the controller can execute the initial task.

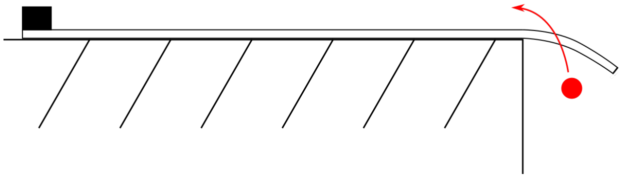

(a) Starting configuration of paper for task.

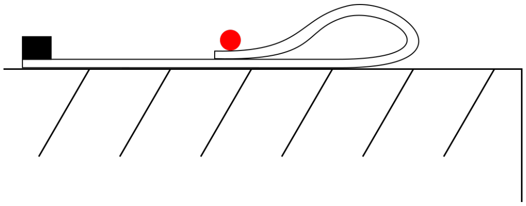

(b) Ending configuration of paper for task.

Fig. 1: Illlustration of this proposal's task. The red circle represents the nonprehensile manipulator, which will also be be abstracted as a sphere in initial simulations, but will be replaced with the Franka Emika Panda.

IV. MODELING APPROACH

High fidelity simulation of the dynamics of deformable objects is computationally intensive, requiring finite element models for anisotropic materials like paper; mass-spring models provide a more practical alternative in exchange for model accuracy [7], [8]. For this proposal, the paper is broken down along its longer edge into a series of rigid links, each of which are joined by a revolute joint that has both stiffness and damping. These parameters are tuned to give approximately realistic behavior based off of the resting configuration of the paper and the settling time of the system.

The controller and paper systems are simulated in Drake, as shown in Fig. 2. The paper's overall dimensions are 8.5 by 11 inches, with a the number of links ranging from 2 to 20 depending on required simulation fidelity. The longer edge of the paper runs parallel to the word's y-axis and the short edge runs parallel to the x-axis. The spherical, nonprehensile manipulator used in initial simulations is confined to moving in the yz-plane and has a radius of 1 centimeter, but eventually a model of the Franka Emika Panda will be used instead. The control inputs to the system are the forces in the y and z direction and torque about the x-axis on the manipulator.

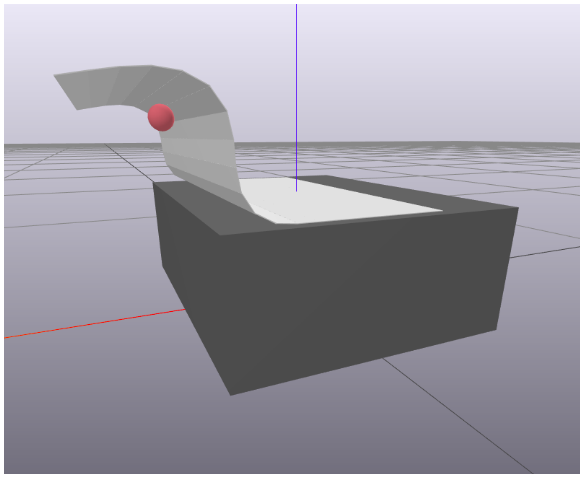

Fig. 2: Screenshot of Drake simulation of system. The gray block shows the platform the paper is affixed to. The red sphere is the manipulator, and the many white links model the paper.

V. PROPOSED CONTROL STRATEGY

To design a controller to complete the task outlined in Section III, we first need to designate which physical quantities the controller will track. Options exist at varying levels of complexity: for example, we could define a desired trajectory for every link in the model. However, generating such trajectories would be intractable, especially considering that these trajectories must be feasible for our nonprehensile manipulator. Ideally, our selection of the tracked quantity would abstract away what parts of the system dynamics are unimportant to the task. In this thesis, we will use the position of the manipulator relative to the paper's edge and the normal velocity of the paper's edge.

Fig. 3 illustrates these values. Here d is defined as the vector from the edge of link (specifically its corner on the side nearer to the manipulator) to the point of contact, and d_T is the projection onto the direction tangent to the link's surface $\hat{T}$, while v_{LN} is the velocity of the link in the normal direction (meaning the contact normal that points from the manipulator towards the link). Tracking d_T ensures that contact is preserved along the surface of the paper, while tracking a positive v_{LN} ensures that eventually, the paper will be folded all the way around.

Tracking v_{LN} and d_T rather than the entire configuration of the paper does simplify our choice of trajectories, but the forces exerted by the deformable object will still complicate their dynamics. Any control strategy must balance the high model fidelity required for precise manipulation of a

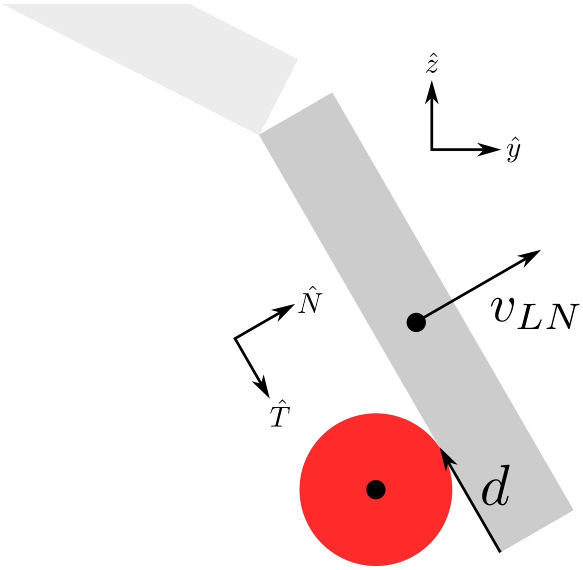

Fig. 3: Illustration of tracked values v_{LN} and d. Note that the $\hat{N}$ vector points in the direction normal from the link's surface away from the manipulator, which the $\hat{T}$ vector points along the surface of the link towards the edge of the paper.

deformable object and the practical concerns of implementing such a strategy. The proposed strategies outlined in this section offer different trade offs between simplicity and modeling accuracy.

A. Feedforward starting point

With the initial goal of simply completing the task in simulation, we implement a fully feedforward controller with full knowledge of the system dynamics. Fig. 4 shows the forces and torques acting on the system, which we can write as follows:

$$a_{LT}m_L = F_{FL} + F_{GT} + F_{OT} \tag{1}$$

$$a_{LN}m_L = F_{GN} + F_{NL} + F_{ON} \tag{2}$$

$$a_{MT}m_M = F_{CT} + F_{FM} \tag{3}$$

$$a_{MN}m_M = F_{CN} + F_{NM} \tag{4}$$

$$I_L\ddot{\theta}_L = -F_{FL}p_{LCN} + F_{NL}p_{LCT} - \frac{F_{ON}w_L}{2} + \tau_O \tag{5}$$

$$I_M\ddot{\theta}_M = -F_{FM}p_{MCN} + \tau_{CX} \tag{6}$$

$$F_{NL} = -F_{NM} \tag{7}$$

$$F_{FL} = F_{NL}\mu \tag{8}$$

$$F_{FM} = -F_{FL} \tag{9}$$

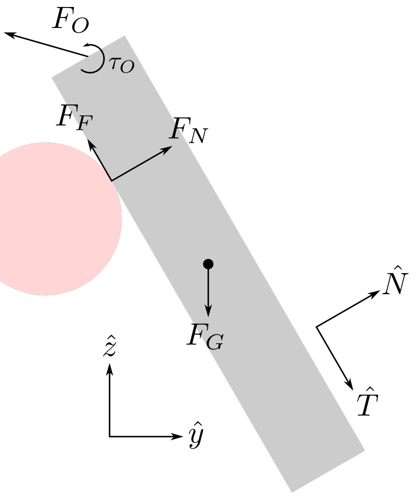

(a) Forces and torques on the last link in the model. F_O and τ_O represent the "object forces," meaning the forces and torque exerted on the last link at the revolute joint by the rest of the object.

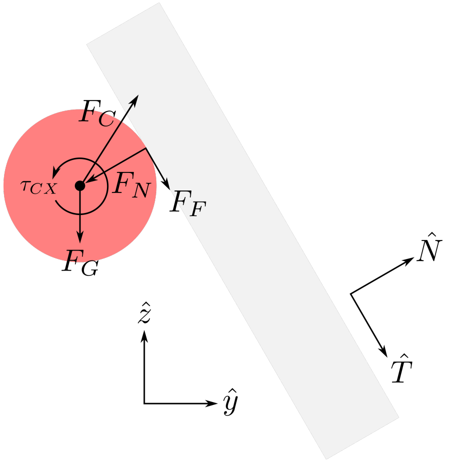

(b) Forces and torques on the manipulator. F_C and τ_{CX} are the forces and torque exerted by the controller.

Fig. 4: Forces and torques on link and manipulator. In both diagrams, F_G is the force due to gravity, and F_N and F_F are the normal and friction forces, respectively.

$$\begin{aligned} a_{MT} - a_{LT} &= \left(\frac{h_L}{2} + r - d_N\right)\ddot{\theta}_L + \ddot{d}_T \\ &\quad - \left(d_T + \frac{w_L}{2}\right)\dot{\theta}_L^2 - 2\dot{\theta}_L\dot{d}_N \end{aligned} \tag{10}$$

$$\begin{aligned} a_{MN} - a_{LN} &= \left(d_T + \frac{w_L}{2}\right)\ddot{\theta}_L + \ddot{d}_N \\ &\quad + \left(\frac{h_L}{2} + r - d_N\right)\dot{\theta}_L^2 + 2\dot{\theta}_L\dot{d}_T \end{aligned} \tag{11}$$

$$\ddot{d}_N = 0 \tag{12}$$

Here the L subscripts indicate accelerations, poses, and so on of the last link; M subscripts indicate those same quantities for the nonprehensile (spherical) manipulator. N and T subscripts are projections on $\hat{N}$ and $\hat{T}$. p_{MC} is the vector from the manipulator's center of mass to the contact point and p_{MC} is the vector from the manipulator's center of mass to the contact point. h_L is the height of the paper and w_L is the lenth of each individual link. mu is the coefficient of friction between the paper and the manipulator, and r is the radius of the spherical manipulator. Theses dynamics describe the sphere manipulator, but the same equations require only slight modifications to model the full robot arm.

These equations reflect the force and moment balance on the link in equations (1), (2), and (5); force and moment balance on the manipulator in equations (3), (4), and (6); friction constraints in equations (7), (8), and (9); differentiating d and relating it to the movements of the link and manipulator in equations (10) and (11); and the rigid body constraint that there is no penetration between the link and the manipulator in equation (12).

In simulation, we can measure the object forces, so if we specify a desired $\ddot{d}_T$, $\ddot{\theta}_M$, and a_{LN}, the dynamics system of equations is fully determined and we can solve for F_{CN}, F_{CT}, and τ_{CX}.

A major drawback of this strategy is that it is not robust to errors in the model or dynamics because it includes no feedback or estimation of the system parameters and it requires the measurement of many quantities that can't be measured in reality.

B. Adaptive control for friction

Adaptive control is used to control systems with a constant, unknown parameter in the system dynamics. Even if the dynamics are nonlinear with respect to the states, they should still be linearly parameterized by the unknown constant. The adaptive controller will then develop an estimate of constant that is sufficient to provide good tracking of the control target.

Adaptive control suits control with an unknown coefficient of friction well, because the coefficient of friction is linearly related to the friction force under Coulomb friction. Furthermore, the coefficient of friction is difficult to measure and unlikely to be known accurately beforehand, so designing a controller which adapts to a different coefficient of friction will be necessary before testing the controller in the real world.

In this thesis work, we use standard adaptive control formulations and use error on d_T to learn the coefficient of friction μ. This also introduces feedback on d_T into the system. The adaptive controller will not include feedback on v_{LN} because the dynamics in the normal direction are not linearly parameterized by μ, but as long as the $\hat{\mu}$ estimated based on error in d_T converges to the true μ (which occurs as long as there is sufficient excitation of the dynamics), that $\hat{\mu}$ can be used for control in directions other than $\hat{T}$.

C. Adaptive control for object forces

The control strategies discussed so far still rely on the measurement of F_O and τ_O, which is impractical for implementing the controller in the real world. With an actual physical system, we will only assume knowledge of the motions of the last link in the chain (ascertained through vision) and the motions of the manipulator. However, if we assume the object forces are a weight sum of known quantities (θ_L, $\dot{\theta}_L^2$, etc), then we can use adaptive control to learn those weights.

D. Observer for object forces

The adaptive control strategy in the previous section does not utilize our model of the paper as a chain of rigid links. To leverage the rigid link model, we would need the positions of all intermediate rigid links, not just the the last link that we assume we can sense. An observer for the system could be used to estimate the positions of these links as hidden states. We can then combine these estimates with adaptive control to model the compliance and damping of the system.

E. Other considerations

Robust control can account for a time-varying bounded disturbance, but may introduce oscillations, especially if the bound is not tight enough. Integrating such a controller may be useful for accounting for errors in the model that can't be represented by known dynamics and unknown constants, but the oscillations introduced may render such a controller impractical.

Another idea that may prove useful at some point is estimating the paper's parameters before the manipulator even makes contact. Specifically, for a rigid link system, the configuration of the system at rest is determined by the compliance of the joints. This could seed the estimate of the system compliance for adaptive control to be closer to the true value before manipulation begins.

VI. PROPOSED EXPERIMENTS

As evidenced by other sections in this proposal, the final controller for folding paper will likely be complex and involve several nested controllers. To integrate such a complex system successfully, we first need to test intermediate versions on simplified prototypes. This section describes how the task will be broken down into different experiment milestones and how those experiments will be implemented.

A. Hinge prototype

When we model the paper as a series of rigid links, the model increases in complexity and accuracy as the number of links increases; conversely, an extremely simple (albeit inaccurate) model of paper is a two link system–in other words, a hinge. When the model is reduced to two links, the object forces can be computed easily given system parameters such as compliance, damping, and mass. Estimation of the entire state of the object is also made easier, as there is only one link to sense.

Friction may still be a difficult force to account for, meaning the adaptive control strategy described in Section V-B will likely still be necessary, but otherwise the feedforward strategy should be effective.

The hinge prototype will be implemented using a 3D printed hinge. AprilTag fiducials [13] will be used to locate the link, and the Franka Emika Panda will be used to execute the strategy. Although this robot has a prehensile manipulator, the fingers will remain held together to create a nonprehensile manipulator.

Once the controller successfully folds the hinge prototype, the test setup will be expanded to manipulate multi-link systems instead, constructed similarly to the hinge prototype but with multiple links. This will allow us to test the controller with more complex object forces but still avoid prematurely introducing all the complexities of a compliant object like paper.

The final experiment will be to repeat this process with actual paper.

VII. TIMELINE

This table shows my proposed timeline for completing my thesis work:

September	Hinge controller simulation
October	Hinge controller experiments
November	Multi-link adaptive controller simulation
December	Multi-link adaptive controller experiments
January	Multi-link observer controller simulation
February	Multi-link observer controller simulation
March	Experiments with real paper
April	Controller improvements for real paper
May	Thesis writing

REFERENCES

[1] M. Schumacher, J. Wojtusch, P. Beckerle, and O. von Stryk, "An introductory review of active compliant control," *Robotics and Autonomous Systems*, vol. 119, pp. 185–200, 2019. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0921889018307772

[2] N. Hogan, "Impedance Control: An Approach to Manipulation: Part I—Theory," *Journal of Dynamic Systems, Measurement, and Control*, vol. 107, no. 1, pp. 1–7, 03 1985. [Online]. Available: https://doi.org/10.1115/1.3140702

[3] M. T. Mason, "Compliance and force control for computer controlled manipulators," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 11, no. 6, pp. 418–432, 1981.

[4] J. D. Schutter and H. V. Brussel, "Compliant robot motion i. a formalism for specifying compliant motion tasks," *The International Journal of Robotics Research*, vol. 7, no. 4, pp. 3–17, 1988. [Online]. Available: https://doi.org/10.1177/027836498800700401

[5] ——, "Compliant robot motion ii. a control approach based on external control loops," *The International Journal of Robotics Research*, vol. 7, no. 4, pp. 18–33, 1988. [Online]. Available: https://doi.org/10.1177/027836498800700402

[6] A. De Luca and C. Manes, "Hybrid force-position control for robots in contact with dynamic environments," *IFAC Proceedings Volumes*, vol. 24, no. 9, pp. 177–182, 1991, 3rd IFAC Symposium on Robot Control 1991 (SYROCO'91), Vienna, Austria, 16-18 September 1991. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1474667017510526

[7] P. Jiménez, "Survey on model-based manipulation planning of deformable objects," *Robotics and Computer-Integrated Manufacturing*, vol. 28, no. 2, pp. 154–163, 2012. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0736584511000986

[8] P. Chang and T. Padif, "Sim2real2sim: Bridging the gap between simulation and real-world in flexible object manipulation," in *2020 Fourth IEEE International Conference on Robotic Computing (IRC)*, 2020, pp. 56–62.

[9] F. Wirnshofer, P. S. Schmitt, P. Meister, G. v. Wichert, and W. Burgard, "Robust, compliant assembly with elastic parts and model uncertainty," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 6044–6051.

[10] D. J. Balkcom and M. T. Mason, "Robotic origami folding," *The International Journal of Robotics Research*, vol. 27, no. 5, pp. 613–627, 2008. [Online]. Available: https://doi.org/10.1177/0278364908090235

[11] C. Elbrechter, R. Haschke, and H. Ritter, "Bi-manual robotic paper manipulation based on real-time marker tracking and physical modelling," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2011, pp. 1427–1432.

[12] C. Jiang, A. Nazir, G. Abbasnejad, and J. Seo, "Dynamic flex-and-flip manipulation of deformable linear objects," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 3158–3163.

[13] E. Olson, "Apriltag: A robust and flexible visual fiducial system," in *2011 IEEE International Conference on Robotics and Automation*, 2011, pp. 3400–3407.