AutoBag: Learning to Open Plastic Bags and Insert Objects

Lawrence Yunliang Chen¹, Baiyu Shi¹, Daniel Seita², Richard Cheng³, Thomas Kollar³, David Held², Ken Goldberg¹

Abstract—Thin plastic bags are ubiquitous in retail stores, healthcare, food handling, recycling, homes, and school lunchrooms. They are challenging both for perception (due to specularities and occlusions) and for manipulation (due to the dynamics of their 3D deformable structure). We formulate the task of "bagging:" manipulating common plastic shopping bags with two handles from an unstructured initial state to an open state where at least one solid object can be inserted into the bag and lifted for transport. We propose a selfsupervised learning framework where a dual-arm robot learns to recognize the handles and rim of plastic bags using UVfluorescent markings; at execution time, the robot does not use UV markings or UV light. We propose the AutoBag algorithm, where the robot uses the learned perception model to open a plastic bag through iterative manipulation. We present novel metrics to evaluate the quality of a bag state and new motion primitives for reorienting and opening bags based on visual observations. In physical experiments, a YuMi robot using AutoBag is able to open bags and achieve a success rate of 16/30 for inserting at least one item across a variety of initial bag configurations. Supplementary material is available at https://sites.google.com/view/autobag.

I. INTRODUCTION

Opening thin plastic bags and then inserting objects for efficient transport is a useful skill for tasks such as grocery shopping, cleaning, recycling, and packing. However, this is very difficult for robots. Deformable thin objects are challenging to manipulate due to their infinite-dimensional state spaces and nonlinear dynamics, and while there has been much prior work on deformable object manipulation, most focus on 1D linear objects such as ropes [28, 53, 63], cables [50], and elastic beams [7, 69], or 2D objects such as fabrics [30, 47, 59, 64], gauze [54], and paper [14, 39]. The 3D structure and elastoplastic material of plastic bags present numerous challenges. Such bags are extremely lightweight, and moving parts of a bag with one gripper often results in the entire bag moving without much change in the opening. Thus, opening a thin plastic bag requires coordination among multiple contacts. Furthermore, many plastic bags are reflective, translucent, or transparent, making perception challenging.

In this work, we formulate "*bagging*"—manipulating a plastic bag from an unstructured initial state so that a robot can open it, insert solid objects into it, and then lift it for transport. We use overhead RGB images for perception and a bimanual robot. We propose a novel pipeline for



Fig. 1: AutoBag. (1) Initial highly unstructured and deformed bag. (2) After a sequence of manipulation steps, the robot orients the bag upward and opens the bag. (3) The robot inserts 2 items into the bag. (4) The robot lifts the bag filled with the inserted items, so it is ready for transport.

bagging that trains a perception model to segment the bag rim and handles through self-supervised data collection. This involves the robot systematically exploring the bag state space by manipulating a bag annotated with ultraviolet (UV) labels [55]. At test time, we deploy the learned segmentation model on bags without UV labels and evaluate the bag opening using novel metrics. We present a novel algorithm, AutoBag, for opening and inserting items into bags. See Figure 1.

This paper makes 6 contributions:

- 1) A novel problem formulation for "bagging;"
- 2) A novel set of primitive actions for manipulating bags including shaking, compressing, flipping, and dilation;
- A self-supervised data collection process where a robot efficiently explores its state space to manipulate bags into diverse configurations to enable recognizing the handles and rim of bags from UV-fluorescent markings;
- Two metrics that quantify the opening of the bag based on the convex hull area and elongation;
- 5) The AutoBag algorithm for bagging;
- 6) An implemented system with experimental results achieving a success rate of 53.3% for inserting at least 1 object over 30 physical trials.

¹The AUTOLab at UC Berkeley (automation.berkeley.edu).

²The Robotics Institute at Carnegie Mellon University.

³Toyota Research Institute, Los Altos, USA.

Correspondence to: yunliang.chen@berkeley.edu

II. RELATED WORK

A. Deformable Object Manipulation

Deformable object manipulation remains challenging for robots [3, 43, 67]. Typical reasons include the complex dynamics and the infinite set of possible configurations. As a rough categorization, deformable object manipulation can be divided into tasks that involve 1D, 2D, or 3D objects.

Manipulation of 1D deformable objects refers to manipulation of items such as cables [25, 37, 38, 50, 68], ropes [36, 58, 66], and other items which can largely be defined by a single linear component. These are used in tasks such as knotting [17, 35] or untangling [12, 57]. Manipulation of 2D objects refers to items such as clothing and fabrics, as studied in recent work on fabric smoothing [4, 15, 18, 27, 29, 40, 48, 59, 61], which often measuring quality using coverage. A smooth fabric with high coverage may make it easier to later do folding, another canonical task explored in prior work [1, 10, 24, 31]. Assistive dressing [8, 9] encompasses a third set of tasks utilizing 2D deformables.

Manipulation of 3D deformable objects adds another dimension. One type of 3D deformable manipulation involves volumetric 3D objects, including plush toys, sponges, and dough. Prior work has studied manipulation of these items to target configurations [26, 32, 41, 52]. A second type of 3D deformable manipulation refers to objects typically held in containers, as in manipulation of liquids [44] and granular media [5, 33, 45], which may require scooping policies [13]. Other references to 3D deformable manipulation refer to thin surfaces arranged in complex 3D patterns, such as plastic bags, which is the main focus of this work.

B. Manipulating Deformable Bags

One prior direction in robot manipulation of bags is on the mechanical design of robots suitable for grasping [21] or unloading [22] large sacks. Another direction has provided insights on manipulating knotted bags [19, 20] or bags in highly constrained setups, such as with work on closing ziplock bags [16] or using fully opened, stable paper grocery bags [23]. Recently, DextAIRity [62] used air to efficiently expand bags. They used a setup with three UR5 robots, where two grip the bag and the third manipulated a leaf blower in free space. In this work, we consider a bimanual robot with standard parallel-jaw end-effectors for a highly deformable plastic bag with handles, which the robot has to grasp, open, and then insert items inside for transport.

Seita et al. [46] proposed several deformable object manipulation benchmark tasks in simulation that include a similar problem setup of opening and inserting items into a bag, and then lifting and moving the bag. Transporter Networks [65] are proposed but only evaluated in simulation [6]. Similarly, Weng et al. [60] study modeling and interacting with bags purely in simulation. A large sim-to-real gap in both visual and dynamic properties still needs to be overcome for a policy trained in simulation to transfer to real. Some research with deformable bags [2, 49] assumes that the bag starts with its rim facing upwards and the bag wide open to simplify



Fig. 2: The physical setup with the ABB YuMi. During training, we use 6 UV LED lights and 1 regular LED light. We use a Realsense RGBD camera placed overhead the robot. Left: The workspace under regular lighting. Right: The workspace and a painted bag under UV lighting. The painted rim and handle regions of the bag look normal under regular lighting but glow under UV lights. The UV lights are only used during training, not during execution time.



Fig. 3: Left: 5 plastic bags. The first 4 are used to train the perception module (Sec. IV-A). Bags 1 and 5 (test bag) are used in experiments (Sec. V). **Right**: Bag 3 painted with green UV paint on its handles and red UV paint around its rim, under UV lighting.

item insertion and instead focus on the object rearrangement and the bag lifting steps. In this work, we allow bags to start from unstructured configurations and tackle the challenge of orienting and opening the bag.

In recent work, Gao et al. [11] proposed an algorithm for tying the handles of deformable plastic bags. To facilitate tying, they fill in bags beforehand with items which expand the bags and make their handles more likely to be upright and exposed. Instead of tying filled bags, we study a different problem setting of opening and inserting items into bags, where the bags begin empty and in unstructured states.

III. PROBLEM STATEMENT

We propose *bagging*: autonomously manipulating a thin plastic bag to open it, insert at least one item, and lift it for transport. We consider the broad class of thin plastic shopping bags commonly used in grocery stores that are made of a single sheet of translucent, reflective, and highly flexible thin plastic material cut with two holes for handles from die-cutting [51]. We define the bag "rim" as the edge of the bag around its primary opening as if the plastic handles were cut off. See Figure 3 for an illustration. The term "opening" in this definition refers to the planar surface enclosed by the rim, through which objects are put inside the bag. The orientation of the opening is the direction of the outward-pointing normal vector from the plane formed by the opening. In many configurations, the opening can have an area of zero (e.g., when the rim is fully folded).

We assume the initial bag state is unstructured: deformed and potentially compressed, and resting stably on the surface, in which the bag rim and handles may be partially or fully occluded. We assume that the two sides of the bag do not



Fig. 4: Overview of AutoBag. The robot starts with an unstructured bag with objects, shown to the left. Given this setup, AutoBag follows the procedure shown in the flow chart to first orient the bag upward and then enlarge the bag opening (see Section IV-D for details). If the bag opening metric exceeds a threshold, then AutoBag proceeds to the item insertion stage (Section IV-E). A trial is a full success if the robot lifts the bag with all items in it. *Not shown in the figure: In the implementation, at the initial state, the robot will also check whether there is already an initial opening so it does not unnecessarily perform a **Compress** and **Flip** and potentially ruin a good state.

stick to each other tightly (which occurs in new bags) and that some opening can be created using actions such as shaking, flinging, or pulling. We assume that the bag can be segmented from the workspace (e.g., by color thresholding).

We consider a bimanual robot with parallel-jaw grippers and an overhead calibrated RGB camera above a flat surface. At each timestep t, the robot receives an RGB image observation $I_t \in \mathbb{R}^{W \times H \times 3}$ of the bag configuration s_t , and executes an action \mathbf{a}_t . We assume also a set of small rigid objects \mathcal{O} , placed in known poses for grasping and insertion. The objective is to develop a policy $\pi : I \mapsto \mathbf{a}$ so that the robot opens the bag, places at least one object inside the bag, then lifts the bag off the table while containing the objects for transport.

IV. METHOD

We propose a learned perception module to recognize the bag rim and handles that includes a novel self-supervised data collection process (Section IV-A) for training, during which the robot uses its action primitives (Section IV-B). We propose two metrics for quantifying the bag opening (Section IV-C). We then describe the AutoBag algorithm (visualized in Figure 4) for opening a plastic bag (Section IV-D), followed by item insertion (Section IV-E).

A. Self-Supervised Learning for Perception Module

In this work, we propose to represent bags through semantic segmentation. With this representation, we conjecture that the robot may be able to estimate the bag state from complex and deformed bag configurations. The robot perceives 2 key parts of a bag: **bag handles** and **bag rim** (see Figure 3). We formulate this as an image segmentation task where, given an RGB image, the output is a per-pixel classification among 4 classes: the 2 aforementioned semantic regions, any remaining area of the bag, and the background.

We use a self-supervised data collection procedure for training a segmentation model using UV-fluorescent markings to avoid time-consuming and expensive human annotation [55]. We place 6 programmable UV LED lights overhead and paint the 2 key parts of bags (handles and rim) with transparent UV-fluorescent paints that brightly reflect 2 different colors under UV light. When the UV lights are turned off, the paints are invisible and the bag looks normal under regular lighting. When the regular lights are turned off and the UV lights are turned on, everything is dark except for the regions with UV paints, which glow their unique colors.

With this setup, the robot uses its action primitives (see Section IV-B) to manipulate the bag into different configurations, and by alternating the lighting conditions, the camera collects paired images of the bag in both standard and UV lighting. By extracting the segmentation masks from the image under the UV lights through color thresholding, the system obtains the ground truth segmentation labels corresponding to the image of the bag under regular lighting conditions, which are then used to train the segmentation network. The objective of this process is to manipulate a plastic bag into a diverse set of configurations—both in terms of its volume and its orientation—as the increased data diversity can lead to a higher-quality perception model.

B. Action Primitives

We consider a set of action primitives \mathcal{A} (see Figure 5). Each primitive has a type $m \in \mathcal{M}$ and action-specific parameters ϕ_m : $\mathbf{a} = \langle m, \phi_m \rangle \in \mathcal{A}$. Gripper positions are specified as Cartesian coordinates in pixels, and the grasping height is set to the height of the workspace to ensure successful grasping of the bag. The primitives are:

- 1) **Recenter** (x, y): Grasp the bag at (x, y) from top down, lift it up, and translate it to the workspace center at (0, 0). This prevents the bag from moving off the workspace.
- Rotate (x, y, α, β, γ): Grasp the bag at (x, y) from top down, lift it up, rotate the gripper by Euler angles (α, β, γ), and then directly place the gripper down.
- 3) Shake (x, y, k_s, ℓ, f): Grasp the bag at (x, y) from top down, lift it up, then perform k_s shaking motions of amplitude ℓ and frequency f, where the gripper rotates its wrist side by side, followed by a swing action to lay the bag on the table. This action often expands the surface area of a compressed bag.



Fig. 5: The action primitives in this work. See Section IV-B for details, and Section IV-E for **Pin-Pull** (only used during item insertion). For each primitive, we show a small two-frame overview of it in action. The project website contains full videos of all primitives.

- 4) Fold (x, y, d): Grasp the bag at (x, y) from top down, lift it off the workstation, move the gripper horizontally outward by distance d and then inward and downward to fold the bag and reduce its top-down visible surface area. This action enables the robot to compress and partially reset the bag state. Combined with other actions that expand the bag, we find that this induces greater diversity of bag configurations during self-supervised data collection (Section IV-A). The robot does not use this action when opening the bag at execution time (Section IV-D).
- 5) **Compress** (x, y, k_c) : Grasp the bag at (x, y), lift it up in midair, then press it downwards until contact with the workspace, and repeat for a total of k_c motions. This action changes the side of the bag that is flat after being compressed. Holding the bottom part with the bag opening facing downward and compressing also allow air to inflate the bag opening.
- 6) Flip (x_l, y_l, x_r, y_r, α): Orient both grippers towards each other, each with an angle α with the horizontal plane, grasp opposite ends of the bag at positions (x_l, y_l) and (x_r, y_r). Then, lift both grippers, rotate each gripper by 180 degrees, and then place the bag down. This action tends to change the direction of the bag opening (e.g., from downwards to upwards).
- 7) Dilate (x_l, y_l, x_r, y_r, α, θ, d): Bring both grippers together at positions (x_l, y_l) and (x_r, y_r) and with them pointing downwards with an angle α with the horizontal plane. Then move both grippers away from each other horizontally along direction θ, each moving by distance d. This action can enlarge a small opening. We use this action during bag opening but not during data collection.

During data collection, the robot uses the following policy to select actions to manipulate the bag into diverse states:

- When the bag area (as viewed from above) exceeds a threshold, sample these primitives uniformly at random: Rotate, Shake, Fold, Compress, Flip;
- 2) When the bag area is small, sample these primitives uniformly at random: **Rotate**, **Shake**, **Compress**;
- 3) After Compress, perform a Flip;
- 4) When the bag is off the center, perform Recenter.

C. Bag Opening Metrics

We propose to use two metrics for quantifying the bag opening from a segmented image. Both use the convex hull of the bag rim, denoted as CH, which is the convex region in the image enclosed by the pixels that the perception model (Section IV-A) identifies as the rim.

- 1) Normalized convex hull area A_{CH} : to approximate the size of the bag opening, we compute its convex hull area in 2D pixel space, and divide that by the maximum convex hull of the bag rim. To compute the maximum convex hull value, a human manually manipulates the bag offline to maximize the opening.
- 2) Convex hull elongation E_{CH} : we approximate the bag opening elongation using the ratio of the PCA major and minor axes of the convex hull in 2D pixel space.

See Figure 6 for visualizations. The normalized convex hull area makes the metric bag size-agnostic. In addition, we use convex hull elongation because we observe that for inserting items, a sideways-facing bag with a closed opening is worse than an upward-facing opening which is small but rounded. The normalized convex hull area, however, may give higher values to the former. Consequently, measuring elongation gives extra information about the bag opening.

D. AutoBag Algorithm

The first part of the AutoBag algorithm uses the perception module (Section IV-A) to choose actions (Section IV-B) to open a bag from an unstructured initial state. See Figure 4 for an overview. This consists of the following two stages.

Stage 1: If the 2D pixel surface area of the bag is below a threshold $S^{(1)}$, the robot executes a **Shake** to expand the bag, where its grasp points are sampled from the handle region (or anywhere on the bag if handles are not visible). Otherwise, if the area exceeds $S^{(1)}$, the robot grasps the bottom and executes a **Compress**. This flattens the bottom so that when the robot next executes a **Flip**, the bag may stand stably with its opening facing upward. With an updated top-down image, the algorithm uses the perception model to compute the metrics A_{CH} and E_{CH} . When the former is above threshold $A_{CH}^{(1)}$ and the latter is below threshold $E_{CH}^{(1)}$.



Fig. 6: Perception Pipeline. The robot collects images by autonomously manipulating a bag into diverse configurations. It extracts segmentation labels of bag handles and rims through color thresholding from the images under UV light. We use this data to train a semantic segmentation network. During execution time, the perception model takes in an RGB image and predicts the segmentation (white: handles; light gray: rim; dark gray: other parts of the bag; black: background). We fit a convex hull to the predicted rim region (indicated by the red boundary), and compute the area and elongation of the convex hull to quantify the bag opening (Section IV-C). We overlay an ellipse with the major and minor axes computed by PCA on the convex hull.

it is likely that an initial upward-oriented opening exists. If the thresholds are not met, the bag is likely tilted on its side or folded inward, with the opening closed, and the robot resets the state by executing a **Shake** and repeats.

Stage 2: The robot uses **Rotate** and **Dilate** actions to iteratively enlarge the opening. At each iteration, the algorithm queries an overhead image of the bag, estimates its rim positions, and uses PCA to identify the direction of the major and minor axes of the opening. The robot performs a **Rotate** about the *z*-axis so that the minor axis of the bag aligns with the horizontal axis, and then uses **Dilate** to pull the bag opening farther apart from the opening center along its minor axes. This process repeats until the normalized convex hull area reaches a threshold $A_{CH}^{(2)}$ and the elongation metric falls below a threshold $E_{CH}^{(2)}$, suggesting that the opening is large and round enough for object insertion (Section IV-E).

E. AutoBag: Object Insertion and Bag Lifting

The final steps of AutoBag involve inserting the objects into the bag and lifting the bag. To lift the bag, the robot grasps the bag at the workstation height to avoid missed grasps, but this can lead to grasping multiple layers, a common challenge in deformable manipulation [56]. We thus propose the **Pin-Pull** $(x_{pin}, y_{pin}, x_{pull}, y_{pull})$ primitive: one gripper goes to position (x_{pin}, y_{pin}) , and presses down onto the bag ("pinning"). The other gripper goes to position (x_{pull}, y_{pull}) , closes the gripper, and lifts up to a fixed height h or until a torque limit is reached ("pulling"). The purpose of **Pin-Pull** is to stretch the bag after grasping. The additional layer that is accidentally grasped can slip out of the gripper during this process, leading to a higher success rate of grasping a single layer.

Given an overhead image of the bag, the robot estimates the opening by fitting a convex hull on the perceived rim, then divides the space by the number of objects. It grasps each object using known poses and places them in the center of each divided region. Then it identifies the positions of the handles and performs two **Pin-Pull** actions to grasp the left and right handles (or bag boundaries if handles are occluded). Finally, the robot lifts the bag off the table.

V. PHYSICAL EXPERIMENTS

During training and experiments, the bags we use are of size 28-30 cm by 49-54 cm when laid flat (see Figure 3). The flat workspace has dimensions 70 cm by 90 cm.

A. Self-Supervised Training of Perception Module

For data collection, we use 4 training bags (see Figure 3), and collect 500 images for each bag. All images come with automatic labels using the self-supervised procedure. We use a U-Net architecture [42] for the segmentation network, trained with soft DICE loss [34]. We use one NVIDIA Titan Xp GPU, with a batch size of 8, and a learning rate of 5e-4. The trained model achieves a 77% intersection over union (IOU) on the validation set. See the supplement for the learning curve and example predictions.

B. Experiment Protocol

To evaluate AutoBag, we use two bags, one of which is the smallest bag from training. The other, unseen bag has the same size as the largest training bag but has different patterns (see both in Figure 3). Neither has UV paint. The goal is to insert 2 identical 2 oz. spray bottles into each bag. We define a *trial* as an instance of the robot attempting to perform the full end-to-end procedure: to open a bag, insert the items into it, and then lift the bag (with items) off the surface. We allow for up to T = 15 actions (excluding **Recenter**) before the robot must formally lift the bag. If the robot encounters motion planning or kinematic errors during the trial, we reset the robot to the home position and continue the trial. We consider 3 tiers of initial bag configurations:

- <u>Tier 1</u>: The bag starts upward-facing with the rim recognizable but with a small opening. This requires enlarging the bag opening to allow placing objects inside.
- <u>Tier 2</u>: The bag starts at an expanded, slightly wrinkled state lying sideways on the workspace. This requires reorienting the bag upwards and then opening the bag.

Train Bag	Method	Open Bag	#Placed	#Contained	$n \ge 1$ Succ.	n = 2 Succ.
	AB-P	2/6	$0.7 {\pm} 1.0$	0.7 ± 1.0	2/6	2/6
Tier 1	$AB-A_{CH}$	5/6	$1.0 {\pm} 0.9$	0.5 ± 0.5	3/6	0/6
	$AB-E_{CH}$	5/6	0.8 ± 1.0	0.7 ± 1.0	2/6	2/6
	AutoBag	5/6	$1.6{\pm}0.7$	$1.1{\pm}0.8$	4/6	3/6
Tier 2	Side Ins.	1/6	$0.3 {\pm} 0.8$	$0.3 {\pm} 0.8$	1/6	1/6
	Handles	2/6	0.2 ± 0.4	0.2 ± 0.4	1/6	0/6
	AB-P	2/6	0.7 ± 1.0	$0.5 {\pm} 0.8$	2/6	1/6
	AutoBag	4/6	$1.3{\pm}1.0$	0.8±1.0	4/6	2/6
Tier 3	AutoBag	1/6	0.3±0.8	0.3±0.8	1/6	1/6

TABLE I: Physical experiment results of AutoBag, ablations, and baseline methods using a training bag across 3 different tiers.

AutoBag Test Bag	Open Bag	#Placed	#Contained	$n \ge 1$ Succ.	n = 2 Succ.
Tier 1 Tier 2	4/6 3/6	$1.2{\pm}1.0 \\ 0.9{\pm}1.0$	$1.0 {\pm} 0.9 \\ 0.7 {\pm} 0.8$	4/6 3/6	2/6 1/6

TABLE II: Results of AutoBag on the test bag, across the first two tiers. We report similar evaluation metrics as in Table I.

• <u>Tier 3</u>: Any other, more complex configuration, such as when the bag is compressed with no visible handles.

At the start of each trial, we initialize the bag by compressing and then adjusting it to make it fall into one of the tiers. A trial is an " $n \ge 1$ success" if the robot can contain at least one bottle in the bag when lifting. A trial is an "n = 2success" when the robot lifts the bag off the surface while containing both bottles. Additionally, we report the number of times the robot successfully opens the bag ("Open Bag") (i.e., proceeds to item insertion), and the number of objects that the robot correctly places in the bag opening before lifting ("#Placed"). Objects may fall out when the robot lifts the bag, meaning that the number of objects contained ("#Contained") is upper-bounded by #Placed.

C. AutoBag and Baseline Methods

We compare AutoBag to baselines and ablations, which we summarize here and defer details to the supplement.

To evaluate the proposed perception module, we perform ablation studies in Tier 1 and Tier 2 that remove the perception module ("AB-P"). Instead of detecting the rim and using the rim area and elongation as the metrics for determining whether the bag is sufficiently opened, it uses the area and elongation of the bag as an approximation. In addition, for **Dilate**, instead of dilating from the opening center, it uses the bag center. To evaluate the proposed metrics for quantifying the bag opening, we perform ablation studies in Tier 1. "AB- A_{CH} " uses only the convex hull elongation metric, and "AB- E_{CH} " uses only the convex hull area metric.

Additionally, for Tier 2, since the bag opening faces sideways, we design 2 different baselines. "Side Ins." attempts to grasp the top layer of the bag to open it and then insert objects from the side. "Handles" grasps the two handles, performs horizontal and vertical flinging to open the bag, and then releases one gripper to insert the objects while the other gripper still holds the bag.

VI. RESULTS AND LIMITATIONS

We report results on the training bag in Table I, where we run 6 trials per experiment setting. The results suggest that AutoBag achieves an n = 2 success rate of 3/6, 2/6, and 1/6 on the three respective tiers and outperforms baselines and ablations in all metrics. The low success rate of "AB-P" for opening the bag suggests the perception module is important for determining whether a bag is actually opened and for placing the gripper inside the opening for **Dilate**. Using only the opening area metric achieves a good $n \ge 1$ success rate but fails to put both objects in, as the opening is not sufficiently wide. On the other hand, using only the elongation metric ensures the bag opening is round but does not guarantee that the opening is large enough. For sideways insertion, the main challenge is grasping the top layer only to create a sideways opening. We find that it often either misses the grasp or grasps two layers, due to the tight space between the two layers. The main failure reason for the "Handles" baseline is that the bag collapses or the opening closes as soon as one gripper releases the handle to grasp the inserted object, leaving no room for the robot to insert the objects.

On the test bag, we conduct experiments on Tier 1 and Tier 2 only, due to the extra challenge from the test bag perception. The results in Table II suggest that AutoBag can attain 4/6 and 3/6 partial successes.

We summarize the failure modes of AutoBag:

- (A) Number of trials reaches the maximum limit (38%).
- (B) Bag gets pushed out of the workspace (3%).
- (C) Objects are not placed inside the opening (13%).
- (D) Objects falls out when lifting the bag (23%).
- (E) Fail to lift the bag successfully (23%).

(A) is often caused by instability of the perception model prediction, especially in complex bag configurations and around decision boundaries, leading to unsuitable/ineffective actions. (B) occurs due to robot arm motion planning. (C) is due to inaccuracy in the rim recognition, especially for the unseen test bag. (D) is a major bottleneck for achieving *full* success, due to non-ideal grasp position for lifting. (E) is caused by slipped grasps. More details, including failure reasons for baselines, are in the supplementary material.

VII. CONCLUSION AND FUTURE WORK

In this paper, we propose a novel problem and an algorithm, AutoBag, for opening a thin plastic bag and inserting items. In future work, we will study better manipulation techniques to increase the success rate and speed up AutoBag.

ACKNOWLEDGEMENTS

This research was performed at the AUTOLAB at UC Berkeley in affiliation with the Berkeley AI Research (BAIR) Lab, and the CITRIS "People and Robots" (CPAR) Initiative. The authors were supported in part by donations from Toyota Research Institute. Lawrence Yunliang Chen is supported by the National Science Foundation (NSF) Graduate Research Fellowship Program under Grant No. 2146752. Daniel Seita and David Held are supported by NSF CAREER grant IIS-2046491. We thank Justin Kerr, Roy Lin, and the reviewers for their valuable feedback.

REFERENCES

- Y. Avigal, L. Berscheid, T. Asfour, T. Kröger, and K. Goldberg, "SpeedFolding: Learning Efficient Bimanual Folding of Garments," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems* (IROS), 2022.
- [2] A. Bahety, S. Jain, H. Ha, N. Hager, B. Burchfiel, E. Cousineau, S. Feng, and S. Song, "Bag all you need: Learning a generalizable bagging strategy for heterogeneous objects," *arXiv preprint arXiv:2210.09997*, 2022.
- [3] J. Borras, G. Alenya, and C. Torras, "A Grasping-centered Analysis for Cloth Manipulation," arXiv preprint arXiv:1906.08202, 2019.
- [4] L. Y. Chen, H. Huang, E. Novoseller, D. Seita, J. Ichnowski, M. Laskey, R. Cheng, T. Kollar, and K. Goldberg, "Efficiently Learning Single-Arm Fling Motions to Smooth Garments," in *Int. S. Robotics Research (ISRR)*, 2022.
- [5] S. Clarke, T. Rhodes, C. Atkeson, and O. Kroemer, "Learning Audio Feedback for Estimating Amount and Flow of Granular Material," in *Conf. on Robot Learning (CoRL)*, 2018.
- [6] E. Coumans and Y. Bai, PyBullet, a Python Module for Physics Simulation for Games, Robotics and Machine Learning, http: //pybullet.org, 2021.
- [7] S. Duenser, J. M. Bern, R. Poranne, and S. Coros, "Interactive Robotic Manipulation of Elastic Objects," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2018.
- [8] Z. Erickson, M. Collier, A. Kapusta, and C. Kemp, "Tracking Human Pose During Robot-Assisted Dressing using Single-Axis Capacitive Proximity Sensing," in *IEEE Robotics and Automation Letters*, 2018.
- [9] Z. Erickson, V. Gangaram, A. Kapusta, C. K. Liu, and C. C. Kemp, "Assistive Gym: A Physics Simulation Framework for Assistive Robotics," in *Proc. IEEE Int. Conf. Robotics and Automation* (*ICRA*), 2020.
- [10] A. Ganapathi, P. Sundaresan, B. Thananjeyan, A. Balakrishna, D. Seita, J. Grannen, M. Hwang, R. Hoque, J. Gonzalez, N. Jamali, K. Yamane, S. Iba, and K. Goldberg, "Learning Dense Visual Correspondences in Simulation to Smooth and Fold Real Fabrics," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2021.
- [11] C. Gao, Z. Li, H. Gao, and F. Chen, "Iterative Interactive Modeling for Knotting Plastic Bags," in *Conf. on Robot Learning (CoRL)*, 2022.
- [12] J. Grannen, P. Sundaresan, B. Thananjeyan, J. Ichnowski, A. Balakrishna, M. Hwang, V. Viswanath, M. Laskey, J. E. Gonzalez, and K. Goldberg, "Untangling Dense Knots by Learning Task-Relevant Keypoints," in *Conf. on Robot Learning (CoRL)*, 2020.
- [13] J. Grannen, Y. Wu, S. Belkhale, and D. Sadigh, "Learning Bimanual Scooping Policies for Food Acquisition," in *Conf. on Robot Learning (CoRL)*, 2022.
- [14] Y. Guo, X. Jiang, and Y. Liu, "Deformation Control of a Deformable Object Based on Visual and Tactile Feedback," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2021.
- [15] H. Ha and S. Song, "FlingBot: The Unreasonable Effectiveness of Dynamic Manipulation for Cloth Unfolding," in *Conf. on Robot Learning (CoRL)*, 2021.
- [16] R. B. Hellman, C. Tekin, M. van der Schaar, and V. J. Santos, "Functional Contour-following via Haptic Perception and Reinforcement Learning," in *IEEE Transactions on Haptics*, 2018.
- [17] J. Hopcroft, J. Kearney, and D. Krafft, "A Case Study of Flexible Object Manipulation," in *Int. Journal of Robotics Research (IJRR)*, 1991.
- [18] R. Hoque, D. Seita, A. Balakrishna, A. Ganapathi, A. Tanwani, N. Jamali, K. Yamane, S. Iba, and K. Goldberg, "VisuoSpatial Foresight for Multi-Step, Multi-Task Fabric Manipulation," in *Proc. Robotics: Science and Systems (RSS)*, 2020.
- [19] A. Howard and G. Bekey, "Prototype System for Automated Sorting and Removal of Bags of Hazardous Waste," in SPIE, Intelligent Robots and Computer Vision, 1996.
- [20] A. Howard and G. Bekey, "Intelligent Learning for Deformable Object Manipulation," in *Autonomous Robtoics*, 2000.
- [21] H. Kazerooni and C. Foley, "A Robot Mechanism for Grapsing Sacks," in *IEEE Trans. Automation Science and Engineering*, 2005.
- [22] A. Kirchheim, M. Burwinkel, and W. Echelmeyer, "Automatic Unloading of Heavy Sacks From Containers," in *IEEE International Conference on Automation and Logistics*, 2008.
- [23] E. Klingbeil, D. Rao, B. Carpenter, V. Ganapathi, A. Y. Ng, and O. Khatib, "Grasping With Application to an Autonomous Checkout

Robot," in Proc. IEEE Int. Conf. Robotics and Automation (ICRA), 2011.

- [24] R. Lee, D. Ward, A. Cosgun, V. Dasagi, P. Corke, and J. Leitner, "Learning Arbitrary-Goal Fabric Folding with One Hour of Real Robot Experience," in *Conf. on Robot Learning (CoRL)*, 2020.
- [25] V. Lim, H. Huang, L. Y. Chen, J. Wang, J. Ichnowski, D. Seita, M. Laskey, and K. Goldberg, "Planar Robot Casting with Real2Sim2Real Self-Supervised Learning," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2022.
- [26] X. Lin, C. Qi, Y. Zhang, Y. Li, Z. Huang, C. G. Katerina Fragkiadaki, and D. Held, "Planning with Spatial-Temporal Abstraction from Point Clouds for Deformable Object Manipulation," in *Conf. on Robot Learning (CoRL)*, 2022.
- [27] X. Lin, Y. Wang, Z. Huang, and D. Held, "Learning Visible Connectivity Dynamics for Cloth Smoothing," in *Conf. on Robot Learning (CoRL)*, 2021.
- [28] W. H. Lui and A. Saxena, "Tangled: Learning to Untangle Ropes with RGB-D Perception," in *Proc. IEEE/RSJ Int. Conf. on Intelli*gent Robots and Systems (IROS), 2013.
- [29] X. Ma, D. Hsu, and W. S. Lee, "Learning Latent Graph Dynamics for Deformable Object Manipulation," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2022.
- [30] J. Maitin-Shepard, M. Cusumano-Towner, J. Lei, and P. Abbeel, "Cloth Grasp Point Detection Based on Multiple-View Geometric Cues with Application to Robotic Towel Folding," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2010.
- [31] J. Matas, S. James, and A. J. Davison, "Sim-to-Real Reinforcement Learning for Deformable Object Manipulation," *Conf. on Robot Learning (CoRL)*, 2018.
- [32] C. Matl and R. Bajcsy, "Deformable Elasto-Plastic Object Shaping using an Elastic Hand and Model-Based Reinforcement Learning," in Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS), 2021.
- [33] C. Matl, Y. Narang, R. Bajcsy, F. Ramos, and D. Fox, "Inferring the Material Properties of Granular Media for Robotic Tasks," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2020.
- [34] F. Milletari, N. Navab, and S.-A. Ahmadi, "V-net: Fully convolutional neural networks for volumetric medical image segmentation," in 2016 fourth international conference on 3D vision (3DV), IEEE, 2016, pp. 565–571.
- [35] T. Morita, J. Takamatsu, K. Ogawara, H. Kimura, and K. Ikeuchi, "Knot Planning from Observation," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2003.
- [36] A. Nair, D. Chen, P. Agrawal, P. Isola, P. Abbeel, J. Malik, and S. Levine, "Combining Self-Supervised Learning and Imitation for Vision-Based Rope Manipulation," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2017.
- [37] H. Nakagaki, K. Kitagi, T. Ogasawara, and H. Tsukune, "Study of Insertion Task of a Flexible Wire Into a Hole by Using Visual Tracking Observed by Stereo Vision," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 1996.
- [38] H. Nakagaki, K. Kitagi, T. Ogasawara, and H. Tsukune, "Study of Deformation and Insertion Tasks of Flexible Wire," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 1997.
- [39] A. Namiki and S. Yokosawa, "Robotic Origami Folding with Dynamic Movement Primitives," in Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS), 2015.
- [40] K. Puthuveetil, C. C. Kemp, and Z. Erickson, "Bodies Uncovered: Learning to Manipulate Real Blankets Around People via Physics Simulations," in *IEEE Robotics and Automation Letters*, 2022.
- [41] C. Qi, X. Lin, and D. Held, "Learning Closed-Loop Dough Manipulation Using a Differentiable Reset Module," in *IEEE Robotics* and Automation Letters, 2022.
- [42] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*, Springer, 2015, pp. 234–241.
- [43] J. Sanchez, J.-A. Corrales, B.-C. Bouzgarrou, and Y. Mezouar, "Robotic Manipulation and Sensing of Deformable Objects in Domestic and Industrial Applications: a Survey," in *Int. Journal* of *Robotics Research (IJRR)*, 2018.
- [44] C. Schenck and D. Fox, "Visual Closed-Loop Control for Pouring Liquids," in Proc. IEEE Int. Conf. Robotics and Automation (ICRA), 2017.

- [45] C. Schenck, J. Tompson, D. Fox, and S. Levine, "Learning Robotic Manipulation of Granular Media," in *Conf. on Robot Learning* (*CoRL*), 2017.
- [46] D. Seita, P. Florence, J. Tompson, E. Coumans, V. Sindhwani, K. Goldberg, and A. Zeng, "Learning to Rearrange Deformable Cables, Fabrics, and Bags with Goal-Conditioned Transporter Networks," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2021.
- [47] D. Seita, A. Ganapathi, R. Hoque, M. Hwang, E. Cen, A. Tanwani, A. Balakrishna, B. Thananjeyan, J. Ichnowski, N. Jamali, K. Yamane, S. Iba, J. Canny, and K. Goldberg, "Deep Imitation Learning of Sequential Fabric Smoothing From an Algorithmic Supervisor," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems* (IROS), 2020.
- [48] D. Seita, N. Jamali, M. Laskey, R. Berenstein, A. K. Tanwani, P. Baskaran, S. Iba, J. Canny, and K. Goldberg, "Deep Transfer Learning of Pick Points on Fabric for Robot Bed-Making," in *Int. S. Robotics Research (ISRR)*, 2019.
- [49] D. Seita, J. Kerr, J. Canny, and K. Goldberg, "Initial Results on Grasping and Lifting Physical Deformable Bags with a Bimanual Robot," in *IROS Workshop on Robotic Manipulation of Deformable Objects in Real-world Applications*, 2021.
- [50] Y. She, S. Dong, S. Wang, N. Sunil, A. Rodriguez, and E. Adelson, "Cable Manipulation with a Tactile-Reactive Gripper," in *Proc. Robotics: Science and Systems (RSS)*, 2020.
- [51] Shearing (manufacturing) Wikipedia, the free encyclopedia, [Online; accessed 15-September-2022].
- [52] B. Shen, Z. Jiang, C. Choy, L. J. Guibas, S. Savarese, A. Anandkumar, and Y. Zhu, "ACID: Action-Conditional Implicit Visual Dynamics for Deformable Object Manipulation," in *Proc. Robotics: Science and Systems (RSS)*, 2022.
- [53] P. Sundaresan, J. Grannen, B. Thananjeyan, A. Balakrishna, M. Laskey, K. Stone, J. E. Gonzalez, and K. Goldberg, "Learning Rope Manipulation Policies Using Dense Object Descriptors Trained on Synthetic Depth Data," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2020.
- [54] B. Thananjeyan, A. Garg, S. Krishnan, C. Chen, L. Miller, and K. Goldberg, "Multilateral Surgical Pattern Cutting in 2D Orthotropic Gauze with Deep Reinforcement Learning Policies for Tensioning," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2017.
- [55] B. Thananjeyan, J. Kerr, H. Huang, J. E. Gonzalez, and K. Goldberg, "All You Need is LUV: Unsupervised Collection of Labeled Images using Invisible UV Fluorescent Indicators," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2022.
- [56] S. Tirumala, T. Weng, D. Seita, O. Kroemer, Z. Temel, and D. Held, "Learning to Singulate Layers of Cloth using Tactile Feedback," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2022.
- [57] V. Viswanath, K. Shivakumar, J. Kerr, B. Thananjeyan, E. Novoseller, J. Ichnowski, A. Escontrela, M. Laskey, J. E. Gonzalez, and K. Goldberg, "Autonomously Untangling Long Cables," in *Proc. Robotics: Science and Systems (RSS)*, 2022.
- [58] A. Wang, T. Kurutach, K. Liu, P. Abbeel, and A. Tamar, "Learning Robotic Manipulation through Visual Planning and Acting," in *Proc. Robotics: Science and Systems (RSS)*, 2019.
- [59] T. Weng, S. Bajracharya, Y. Wang, K. Agrawal, and D. Held, "FabricFlowNet: Bimanual Cloth Manipulation with a Flow-based Policy," in *Conf. on Robot Learning (CoRL)*, 2021.
- [60] Z. Weng, F. Paus, A. Varava, H. Yin, T. Asfour, and D. Kragic, "Graph-based Task-specific Prediction Models for Interactions between Deformable and Rigid Objects," in *Proc. IEEE/RSJ Int. Conf.* on Intelligent Robots and Systems (IROS), 2021.
- [61] Y. Wu, W. Yan, T. Kurutach, L. Pinto, and P. Abbeel, "Learning to Manipulate Deformable Objects without Demonstrations," in *Proc. Robotics: Science and Systems (RSS)*, 2020.
- [62] Z. Xu, C. Chi, B. Burchfiel, E. Cousineau, S. Feng, and S. Song, "DextAIRity: Deformable Manipulation Can be a Breeze," in *Proc. Robotics: Science and Systems (RSS)*, 2022.
- [63] M. Yan, Y. Zhu, N. Jin, and J. Bohg, "Self-Supervised Learning of State Estimation for Manipulating Deformable Linear Objects," in *IEEE Robotics and Automation Letters*, 2020.
- [64] W. Yan, A. Vangipuram, P. Abbeel, and L. Pinto, "Learning Predictive Representations for Deformable Objects Using Contrastive Estimation," in *Conf. on Robot Learning (CoRL)*, 2020.

- [65] A. Zeng, P. Florence, J. Tompson, S. Welker, J. Chien, M. Attarian, T. Armstrong, I. Krasin, D. Duong, V. Sindhwani, and J. Lee, "Transporter Networks: Rearranging the Visual World for Robotic Manipulation," in *Conf. on Robot Learning (CoRL)*, 2020.
- [66] H. Zhang, J. Ichnowski, D. Seita, J. Wang, H. Huang, and K. Goldberg, "Robots of the Lost Arc: Self-Supervised Learning to Dynamically Manipulate Fixed-Endpoint Ropes and Cables," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2021.
- [67] J. Zhu, A. Cherubini, C. Dune, D. Navarro-Alarcon, F. Alambeigi, D. Berenson, F. Ficuciello, K. Harada, J. Kober, X. Li, J. Pan, W. Yuan, and M. Gienger, "Challenges and Outlook in Robotic Manipulation of Deformable Objects," *arXiv preprint arXiv:2105.01767*, 2021.
- [68] J. Zhu, B. Navarro, R. Passama, P. Fraisse, A. Crosnier, and A. Cherubini, "Robotic Manipulation Planning for Shaping Deformable Linear Objects with Environmental Contacts," in *IEEE Robotics and Automation Letters*, 2019.
- [69] S. Zimmermann, R. Poranne, and S. Coros, "Dynamic Manipulation of Deformable Objects with Implicit Integration," in *IEEE Robotics* and Automation Letters, 2021.