

Right-Side-Out: Learning Zero-Shot Sim-to-Real Garment Reversal

Chang Yu^{*1}, Siyu Ma^{*1}, Wenxin Du¹, Zeshun Zong¹, Han Xue³, Wendi Chen³,
Cewu Lu³, Yin Yang⁴, Xuchen Han², Joseph Masterjohn², Alejandro Castro², Chenfanfu Jiang¹

Abstract—Turning garments right-side out is a challenging manipulation task: it is highly dynamic, entails rapid contact changes, and is subject to severe visual occlusion. We introduce *Right-Side-Out*, a zero-shot sim-to-real framework that effectively solves this challenge by exploiting task structures. We decompose the task into DRAG/FLING to create and stabilize an access opening, followed by INSERT&PULL to invert the garment. Each step uses a depth-inferred, keypoint-parameterized bimanual primitive that sharply reduces the action space while preserving robustness. Efficient data generation is enabled by our custom-built, high-fidelity, GPU-parallel Material Point Method (MPM) simulator that models thin-shell deformation and provides robust and efficient contact handling for batched rollouts. Built on the simulator, our fully automated pipeline scales data generation by randomizing garment geometry, material parameters, and viewpoints, producing depth, masks, and per-primitive keypoint labels without any human annotations. With a single depth camera, policies trained entirely in simulation deploy zero-shot on real hardware, achieving up to 81.3% success rate. By employing task decomposition and high fidelity simulation, our framework enables tackling highly dynamic, severely occluded tasks without laborious human demonstrations. More details and supplementary material are on the website: <https://right-side-out.github.io>.

I. INTRODUCTION

Garment manipulation is a core challenge in robotics, complicated by cloth’s high deformability, complex self-contact, and friction. While recent works have made significant progress in tasks such as folding [1, 2], smoothing [3, 4], and hanging [5, 6], almost all approaches share a critical assumption: the garment is already right-side out. Nevertheless, flipping a garment to its proper orientation, a fundamental step in real-world garment workflows (e.g. laundry) and a prerequisite to all aforementioned tasks, is long overlooked. This task (coined *Right-Side-Out*) restores the garment to its wearable state, exposes functional features, and reduces ambiguity for downstream tasks. Despite its ubiquity, this **dexterous, manually intensive** task remains largely unexplored in robotics.

Right-Side-Out presents three main challenges. (1) **Mixed dynamics**: the naturally hybrid dynamics of the task – alternating between quasi-static singulation and dynamic and momentum-assisted pull-throughs – complicate the learning of a unified policy. (2) **Complex topology and contact**: the task requires manipulating the garment’s topology by



Fig. 1: Our system flips garments right-side-out with three keypoint-conditioned bimanual primitives: (1) DRAG exposes a single-layer hem; (2) FLING brings the single layer to the top; (3) INSERT&PULL grasps the interior and pulls through; (4) final right-side-out state.

routing a gripper through a narrow, self-occluding boundary loop (e.g. the hem). It is also inherently bimanual: one arm anchors the clothes, while the other manipulates, imposing long-horizon gripper-cloth and arm-arm collision constraints. (3) **Garment variation**: high variance in garment type, size, and material results in a combinatorial explosion in configurations, complicating robust policy learning and making demonstration-based data collection prohibitively costly.

Recently, end-to-end learning-based paradigms have made rapid advances in garment manipulation. These approaches, from imitation learning [7, 8] to Vision-Language-Action (VLA) policies [9], have enabled end-to-end execution of complex tasks, such as shirt hanging in ALOHA Unleashed [10]. However, they rely on massive teleoperation datasets, which are labor-intensive and expensive to collect.

Our key insight is to decompose this complex task into a sequence of simpler, human-inspired sub-goals: (i) layer singulation to create an access opening, and (ii) in-garment insertion and pull-through to complete the flip, as illustrated in Fig. 5. Each step is implemented as a set of keypoint-conditioned primitives, predicted from depth and mask. This decomposition not only drastically reduces the action space compared to end-to-end methods, but also provides interpretable sub-goals and clear supervisions for robust learning.

In addition to the expensive labor cost of teleoperation,

* equal contribution.

¹gln0st@live.com, siiyuma@outlook.com, {wenxindu, zeshunzong, cffjiang}@ucla.edu, AIVC Lab, UCLA, USA.

²{xuchen.han, joe.masterjohn, alejandro.castro}@tri.global, Toyota Research Institute, USA.

³{xiaoxiaoxh, chenwendi-andy, lucewu}@sjtu.edu.cn, Shanghai Jiao Tong University, China.

⁴yin.yang@utah.edu, University of Utah, USA.

Right-Side-Out also necessitates extensive sample coverage of critical contact events during intense interactions as manifested in Fig. 5, which further complicates manual data collection. This motivates us to utilize faithful physical simulation for data-generation and policy learning. We developed a high-fidelity, Material Point Method (MPM)-based cloth simulator that accurately models frictional robot-cloth contact, self-collision, and codimensional deformations. Crucially, it supports batched multi-environment GPU simulation for scalable data generation. Our automated pipeline randomizes garment geometry, initial states, and material parameters to produce depth and mask labels without human annotation. Randomization in camera positions and parameters is also introduced for further data augmentation.

Thanks to the high fidelity of our physics-based cloth simulator, the comprehensive coverage of our data augmentation strategy, and the strong robustness of our primitive-based model, we are able to achieve zero-shot sim-to-real transfer with *only* one calibrated depth-camera: our policy is trained purely in simulation, and can be deployed on real robots without any additional fine-tuning.

To summarize, our contributions are as follows:

- We introduced the novel task of **robotic garment reversal**, referred to as *Right-Side-Out*.
- We proposed **parameterized bimanual action primitives** for layer singulation and in-garment pull-through, providing a compact yet robust action parameterization.
- We developed a **high-fidelity, GPU-parallel MPM simulator** with a fully automatic data generation and annotation pipeline for cloth manipulation.
- We demonstrated **zero-shot sim-to-real transfer** on diverse garments.

II. RELATED WORK

A. Learning for Garment Manipulation

Learning-based garment manipulation policies have attracted significant interest. Recent works span folding [1, 2, 11], unfolding [3, 4], hanging [5, 6, 10], and dressing [12, 13], often assuming canonical initial states [3, 14]. Orthogonally, cloth singulation has been pursued with task-specific hardware [15, 16]. Bimanual systems are increasingly employed to extend workspace and enable richer behavior, despite higher planning complexity and potential self-collisions. Beyond hardware considerations, action dynamics represent another key axis. Many pipelines remain quasi-static [1] for smoothing or folding, whereas unfolding from crumpled states has motivated dynamic actions (e.g., fling) to overcome occlusion and reach limits [3].

Garments are particularly difficult to manipulate due to their large configuration spaces, severe self-occlusions, and complex dynamics. Prior work has explored various learning paradigms to address these challenges. These include heuristic or keypoint-centric policies that detect features like wrinkles, corners, and edges [17, 18]; imitation learning from expert teleoperation [9, 10] and human preference [19, 20]; reinforcement learning, often bootstrapped by demonstrations or simplified goals [21, 22]; and goal-conditioned or




Fig. 2: **Real-world robot setup.** Two 6-DoF arms and an overhead RealSense D415; see Section V-A.

self-supervised methods [23, 24]. In this work, we adopt depth-driven, keypoint-conditioned bimanual primitives that reduce the action space for the long-horizon *Right-Side-Out* task, unifying quasi-static and dynamic operations within a single sim-to-real framework.

B. Physically-based Garment Dynamics

Classical cloth solvers discretize garments as thin shells and must resolve *codimensional* contact robustly. FEM pipelines typically rely on explicit collision handling with contact energies [25, 26]. Incremental Potential Contact (IPC) [27–29] introduces logarithmic contact barriers and continuous collision detection for intersection-free guarantees, at the cost of higher Newton iterations. Position-based Dynamics (PBD) [30, 31] trades accuracy for speed via constraint-based formulation and iterative solver, but slow convergence can accumulate unrecoverable penetrations in multi-layer contact. In contrast, the Material Point Method (MPM) [32, 33] uses a background Eulerian grid for particle-grid transfers, easing frictional contact and self-collision handling, and scales well on GPUs. Popular robotic simulation platforms reflect these trade-offs: PyFleX [34, 35] and SoftGym [36] adopt PBD for real-time control benchmarks; IPC-based systems [37–39] target intersection-free, two-way soft-rigid coupling; and MPM has been two-way coupled with rigid and integrated into Drake recently [40, 41].

Our work opts for a high-fidelity MPM simulator to provide scalable, simulation-only supervision – without human demonstrations or annotation – with physically faithful garment dynamics. To the best of our knowledge, we are the first to leverage an MPM-based solver to achieve zero-shot sim-to-real transfer in garment manipulation.

III. PROBLEM STATEMENT

Settings. We consider a dual-arm robot with parallel grippers operating over a flat tabletop and observed by an overhead depth camera (Fig. 2). The garment is distinguishable from the workspace, and the *face polarity* (outside vs. inside) is visually identifiable from texture/appearance cues in the observation. We study sleeveless upper garments (e.g.,




Fig. 3: **Success metric.** From the top-down frame, we segment the garment with SAM [42] and classify face polarity (outside, blue; inside, pink). *Right-side-out coverage* is the fraction of pixels within the garment mask labeled as outside; an episode is a success if coverage ≥ 0.80 . Case I succeeds (100%); Cases II and III fail (71.09%, 26.42%).

tank tops) of varying size and shape; topologically, each is an orientable thin-shell surface whose boundary consists of four disjoint loops, $\partial\Omega = \{\mathcal{C}_{\text{hem}}, \mathcal{C}_{\text{collar}}, \mathcal{C}_{\text{arm}}^L, \mathcal{C}_{\text{arm}}^R\}$, i.e., armholes are present but no attached sleeves. The garment is placed on the table *inside-out* at the start of an episode.

Task and success metric. The objective is to control the bimanual robot to produce a final configuration in which the garment is right-side out. We evaluate performance using a binary success metric computed from the final top-down frame. The metric is based on *right-side-out coverage*, defined as the fraction of pixels within the garment’s mask whose visible face is classified as the outside surface. An episode is considered a success if the coverage rate is $\geq 80\%$; otherwise, it is a failure, as illustrated in Fig. 3.

To ensure our metric focuses on the flipping task alone, we initialize garments in a canonicalized pose (following [4]) but with an inverted (inside-out) face polarity. After the policy runs, we allow an optional canonicalization step that preserves face polarity (e.g., light shaking or flinging [3]) before measurement. This ensures the garment is in a clean state for downstream tasks.

IV. METHOD

We outline our framework as follows. Section IV-A describes our multi-environment GPU-parallel MPM simulator with efficient collision handling and physically faithful garment dynamics. Section IV-B defines the policy decomposition, where we decompose *Right-Side-Out* into DRAG, FLING and INSERT&PULL. Section IV-C presents the action-primitive networks that take depth and mask as input and predict per-primitive value maps. Finally, Section IV-D details the sim-to-real transfer for zero-shot deployment. The overall pipeline is illustrated in Fig. 4.

A. MPM Garment Simulation Environment

We introduce a GPU-parallel Material Point Method (MPM) environment tailored for thin-shell garments and robotic manipulations. Our codimensional MPM solver offers robust collision handling and excellent parallel efficiency: its regular grid and localized memory access avoid irregular contact traversals, yielding good load balance and

scaling to multi-environments. We use the anisotropic elastoplastic constitutive model for codimensional materials (cloth) so that frictional contact of garment is expressed directly in the continuum; in practice, contact processing happens on the grid and requires no explicit collision queries.

a) *State and discretization:* Let \mathcal{S} be a triangulated garment surface with N_v vertices and N_f faces. The stacked vertices position vector is $\mathbf{x}_p = [\mathbf{x}_{p,1}^\top, \dots, \mathbf{x}_{p,N_v}^\top]^\top \in \mathbb{R}^{3N_v}$. For each of the N_f triangular faces, one quadrature point is chosen at its centroid, yielding N_f quadrature particles $\mathbf{x}_c = [\mathbf{x}_{c,1}^\top, \dots, \mathbf{x}_{c,N_f}^\top]^\top \in \mathbb{R}^{3N_f}$. They are evolved with a background Eulerian grid of spacing $h > 0$ via APIC transfers [44]. Particles carry mass, elastic deformation \mathbf{F} , and material parameters; the grid nodes $\{i\}$ hold the DOFs velocities \mathbf{v}_i on which boundary conditions are imposed. The in-plane energy for the garment is modeled as hyperelastic. We resolve frictional contact for self-collision and inter-garment interactions using the anisotropic elastoplastic cloth model of [32]. We refer readers to [44] for MPM details.

b) *Garment-garment contact modeling:* Garment friction is formulated as a Coulomb-type yield condition on the Cauchy stress $\boldsymbol{\sigma}$ in the continuum [32]. Let \mathbf{n} denote the unit normal vector to the local garment manifold. We define a codimensional Coulomb yield function

$$\mathcal{Y}(\boldsymbol{\sigma}; \mathbf{n}) := \|(\mathbf{I} - \mathbf{n}\mathbf{n}^\top)\boldsymbol{\sigma}\mathbf{n}\| - \mu_{\text{cloth}} \mathbf{n}^\top \boldsymbol{\sigma} \mathbf{n} \leq 0, \quad (1)$$

which states that the magnitude of tangential traction is bounded by the friction coefficient μ_{cloth} times the compressive normal traction. At each step, we compute a trial elastic state and apply a return-mapping projection with no change to in-manifold strain, but rescales the shear components orthogonal to the surface so that Eq. (1) holds. This yields sticking, sliding, and separation behavior, consistent with Coulomb friction for codimensional garment.

c) *Garment-robot contact modeling:* We treat each robot link ℓ as well as the workzone table as an immovable collider and enforce contact by a projected Coulomb update on grid velocities by one-way coupling. At a contacted grid node i with $\phi_\ell(\mathbf{x}_i) \leq 0$ and normal \mathbf{n} , let $\mathbf{u} = \mathbf{v}_i - \mathbf{v}_\ell(\mathbf{x}_i)$, $u_n = \mathbf{u}\mathbf{n}$, and $\mathbf{u}_t = \mathbf{u} - u_n\mathbf{n}$. The post-contact grid velocity \mathbf{v}_i^+ is then

$$\mathbf{v}_i^+ = \begin{cases} \mathbf{v}_i - u_n \mathbf{n} - \min\left(\|\mathbf{u}_t\|, \mu_{\text{robot}}(-u_n)\right) \frac{\mathbf{u}_t}{\|\mathbf{u}_t\| + \varepsilon}, & \text{if } \phi_\ell(\mathbf{x}_i) \leq 0 \text{ and } u_n < 0, \\ \mathbf{v}_i, & \text{otherwise,} \end{cases}$$

where μ_{robot} is the friction coefficient and regularization $\varepsilon > 0$. The above design yields a contact-robust, GPU-efficient simulator for large-scale rollout of long-horizon garment interactions with realistic thin-shell deformations. We evaluate speed and fidelity in Section V-B.

B. Policy Parameterization

We reduce the long-horizon, contact-rich problem to parameterized bimanual primitives instantiated by a few 2D image-plane keypoints from a calibrated top-down depth

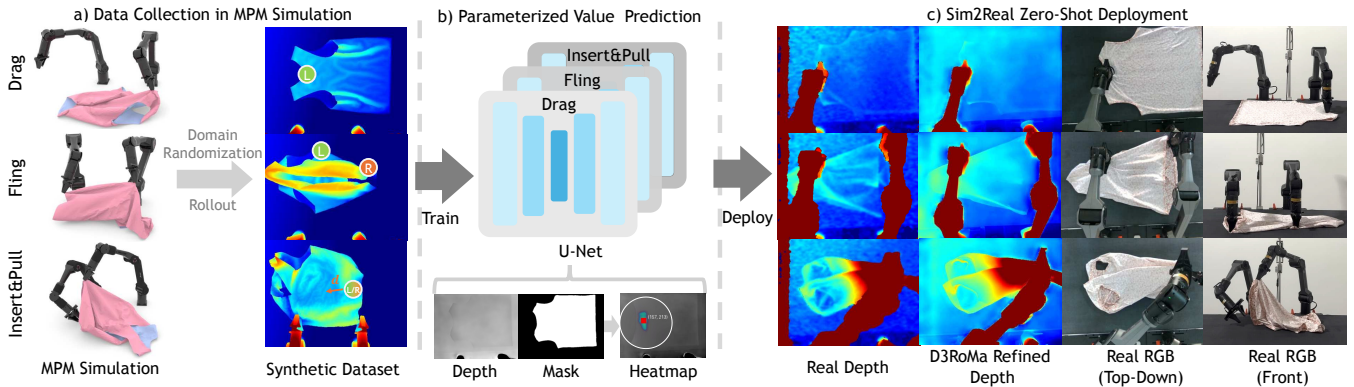


Fig. 4: **System overview:** (a) MPM simulation executes bimanual primitives on domain-randomized garments to generate synthetic depth, masks, and supervision; (b) per-primitive U-Nets take depth and mask as input and predict dense value maps for keypoint selection; (c) the same networks deploy zero-shot on real data using D3RoMa [43] and SAM masks [42].

view captured by an overhead camera. This simple yet effective approach drastically reduces the action space, while also facilitates sim-to-real transfer. Each primitive $m \in \{\text{DRAG}, \text{FLING}, \text{INSERT\&PULL}\}$ is described by a few 2D keypoints, which are predicted from the depth images and masks. Those 2D keypoints are then mapped back to world space, to trigger end-effector motions with fixed templates. We detail each primitive action below. See also Fig. 5 for a step-by-step illustration.

The first step, denoted the singulation stage, is to form a drag point near the hem where the front and back layers are separated. This allows one robot arm to lift the front layer, creating an access opening for the insertion of the second arm. Most prior singulation methods rely on specialized hardware. [45] recognizes single layers with parallel grippers without extra sensors, but its bag-specific perception and grasp-height sensitivity hinder transfer to diverse garments. We observe that, in a canonical pose, the front neckline typically lies below the back neckline because of the front neck drop for anatomical fit. Although the collar loop $\mathcal{C}_{\text{collar}}$ is typically too small for insertion, we can “transport” this single-layer patch towards the hem to create a larger opening, using a drag followed by a brief fling.

For each pixel u in the image plane, let $L(u)$ denote the number of garment layers at pixel u . For instance, around the collar where the back neckline is exposed, $L(u) = 1$, whereas regions around the chest typically have $L(u) = 2$.

DRAG. From a pick point p_{drag} in the single layer region of the collar $\{u \in \text{collar} : L(u) = 1\}$, the robot drags the garment a fixed distance upward and then outward, using workspace friction to separate the two layers and advect the single-layer patch towards the hem region. That is, after this step, $\exists u \in \text{hem}$ s.t. $L(u) = 1$.

FLING. To successfully create an opening with one arm, the hem must have a single-layer on the upper side, i.e. there exists a neighborhood N of pixels s.t. $\forall u \in N, L(u) = 1$ and this layer is the upper layer (i.e. the front layer). We achieve this by grasping a hem-spanning pair (p_L, p_R) and executing a fixed stretch-and-fling: lift to a set height, stretch to a force threshold, then apply a brief forward-backward

pulse. Following prior dynamic-cloth work [3], we fix the fling kinematics and only learn the grasp pair. This action is parameterized by (p_L, p_R) .

INSERT&PULL. This action is parameterized by $(p_{\text{lift}}, p_{\text{pull}})$. The left arm grasps a point $p_{\text{lift}} \in N$ and lifts a fixed distance upwards to p'_{lift} , to create an access opening for subsequent insertion. The right arm positions its gripper directly beneath the lifted p'_{lift} (where the opening is largest), keeping its jaws parallel to the workspace, slides along $d = p_{\text{pull}} - p_{\text{lift}}$ to p_{pull} , raises for a fixed distance and closes the jaw. Gravity drapes the garment into the jaw, yielding a stable grasp; a fixed kinematic pull-out along d completes extraction.

C. Action Primitive Networks

We learn dense per-pixel action-value maps with a U-Net backbone [46]. Given an overhead observation $O \in \mathbb{R}^{C \times H \times W}$ (depth + binary masks) at the start of each primitive, the network outputs an $H \times W$ map whose argmax yields a pixel-wise target. Pixels are back-projected to 3D; controllers fix the z height and kinematic profile, so networks predict only planar (x, y) .

Unimanual primitive (DRAG). Using collar annotations and a ray intersector, we deterministically form a dense supervision map $Y[u] = \mathbb{1}\{u \in \text{collar} : L(u) = 1\}$. We train a single network Q_{drag} with pixel-wise binary cross-entropy loss (BCE): $\mathcal{L}_{\text{drag}} = \text{BCE}(Q_{\text{drag}}[\cdot; O], Y)$.

Bimanual primitives (FLING, INSERT&PULL). For bimanual primitives specified by a pair (p_a, p_b) , we use two networks sequentially: $Q_a(O)$ to select p_a , and $Q_b(O, p_a)$ to select p_b conditioned on p_a , via concatenating a one-hot map of p_a to O . We optimize

$$\mathcal{L}_b = \text{BCE}(Q_b[p_b; O, p_a], S), \mathcal{L}_a = \text{BCE}(Q_a[p_a; O], \hat{S})$$

with $\hat{S} = \mathbb{1}\{\max_y Q_b[y; O, p_a] > \tau\}$ and $\tau=0.5$. We heuristically sample candidates within annotated garment regions. Stage success S is evaluated in simulation for each rollout. We instantiate:

- **FLING:** $(p_a, p_b) = (p_L, p_R)$, $(Q_a, Q_b) = (Q_{\text{fling}}^L, Q_{\text{fling}}^R)$; $S=1$ iff the area of the hem single-layer set $\{u : L(u) = 1\}$ exceeds a threshold τ_{area} .




Fig. 5: **Primitives parameterization.** (a) DRAG, (b) FLING, (c) INSERT&PULL; see Section IV-B.

- INSERT&PULL: $(p_a, p_b) = (p_{\text{lift}}, p_{\text{pull}})$, $(Q_a, Q_b) = (Q_{\text{lift}}, Q_{\text{pull}})$; $S=1$ iff *right-side-out coverage* ≥ 0.80 on mask (illustrated in Fig. 3).

D. Sim-to-Real Transfer and Deployment

To mitigate the physics gap, we apply domain randomization to material and contact parameters: Young’s modulus $E \in [300, 3000]$ kPa, density $\rho \in [1000, 3000]$ kg/m³, friction coefficients $\{\mu_{\text{cloth}}, \mu_{\text{ground}}, \mu_{\text{robot}}\} \in [0.3, 0.95]$, and RPIC damping $c \in [0.0, 0.5]$; shearing stiffness γ is sampled as a fixed fraction of E . For geometric variation, we procedurally generate garments with diverse C_{hem} and C_{collar} shapes, assign rule-based component annotations, and randomize overall size. To reduce the visual gap, we randomize camera intrinsics/extrinsics and add noise to depth and segmentation masks during training. At deployment, we use a well-calibrated overhead camera and D3RoMa-refined depth maps [43] to suppress depth noise. Without any real-world fine-tuning, the simulation-trained policy can be deployed on a dual-arm robot, achieving zero-shot sim-to-real transfer.

V. EXPERIMENTS

A. Experimental Setup

a) *Real-world deployment:* Our real-world experiment setup consists of two 6-DoF AgileX PiPER arms with parallel grippers. The top-down RGB-D image is captured with an Intel RealSense D415, generating observations $O \in \mathbb{R}^{640 \times 480}$. A 3cm foam mat is placed on the workspace to allow the grippers to reach below the garment without

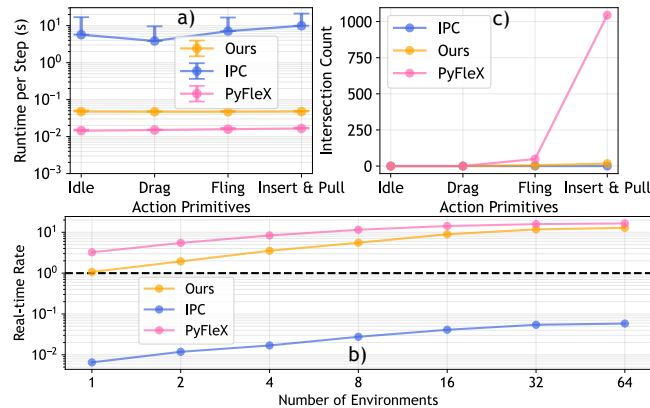


Fig. 6: **Profiling.** (a) per-step runtime by primitive; (b) real-time rate vs. parallel environments (dashed = 1×); (c) max triangle–triangle intersections per step.

colliding, as shown in Fig. 2. Action primitives are executed at a control frequency of 20 Hz.

b) *Simulation data collection:* All simulation and training are run on a single NVIDIA RTX 4090 (24 GB). The control frequency is aligned with the real-world setup and the simulation takes a timestep $\Delta t = 7 \times 10^{-4}$ s with explicit integration under the CFL restriction. Using 64 parallel environments, a heuristic policy collects 20,000 trajectories with 128 synthetic garment meshes; end-to-end data generation takes approximately 25 GPU-hours. We render aligned depth images and garment masks with the SAPIEN renderer [47]. For each primitive, we train a standard U-Net with 5 levels of downsampling/upsampling for 400 epochs; a full training run for each network takes about 1 hour.

B. Simulator Performance and Fidelity Evaluation

We evaluate the throughput and physical fidelity of our MPM simulator against two commonly used baselines for deformable manipulation: (i) a multi–environment IPC solver [38, 39] and (ii) PyFleX [34, 35], adopted in prior Sim-to-Real garment pipelines [3, 5]. For a fair comparison, we only change the physics backend while keeping the data-generation policy, controllers, and control rate fixed. We run IPC at a timestep equal to the control step, owing to its Newton-based solver and continuous collision handling. PyFleX uses its default time step size $\Delta t = 0.01$ s.

We benchmark three aspects: (1) single-environment runtime per physics step at a 20 Hz control loop; (2) real-time rate as the number of parallel environments increases; and (3) garment self-penetration, measured as the maximum number of triangle–triangle intersections per step during the rollout. Results are summarized in Fig. 6.

Our MPM simulator achieves per-step runtime comparable to PyFleX across primitives, below 50 ms per step. IPC is about $10^2 \times$ slower, making large-scale data generation impractical for long-horizon interaction. Real-time rate grows linearly with the number of environments for our simulator, reaching roughly $10 \times$ real-time rate at 64 environments. IPC remains sub-real-time rate even with many environments, reflecting its heavier contact solve. Fidelity difference is most pronounced in the contact-rich INSERT&PULL stage. IPC, by construction, enforces non-penetration, though at a substantial computational cost. PyFleX exhibits severe interpenetration ($> 10^3$), whereas our MPM simulator has much fewer penetrations (< 20) in INSERT&PULL, sufficient to execute the primitive robustly.

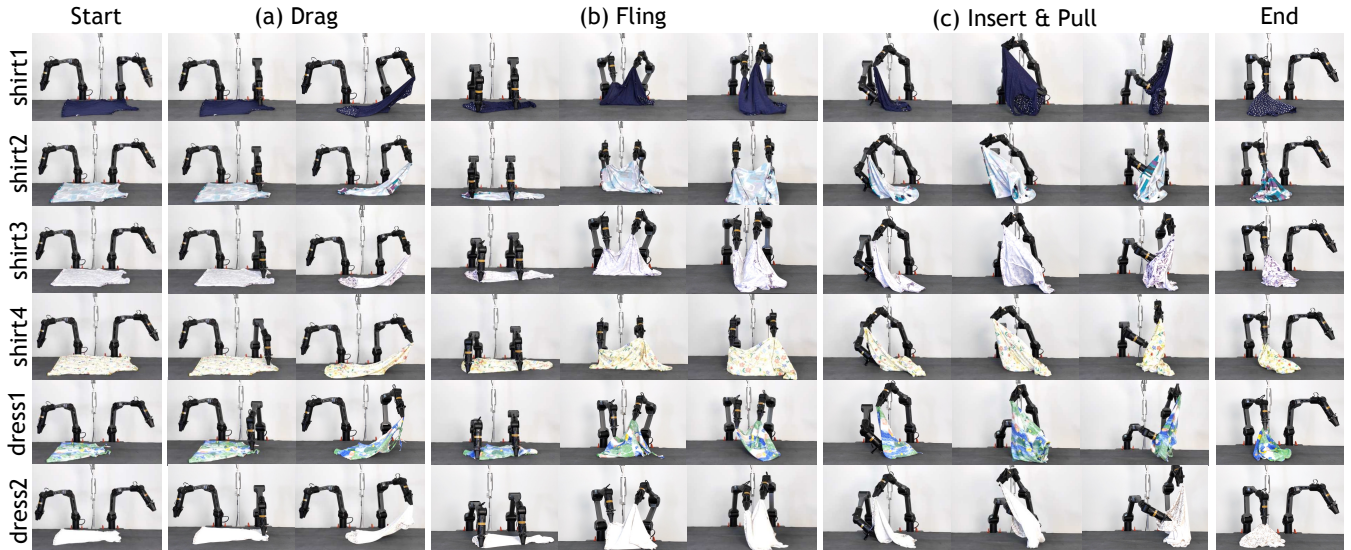


Fig. 7: **Results.** Real-world deployment of our system across diverse garments.



Fig. 8: **Real-world garment assets.** The garments in experiment vary in type, size, color, fabric, and collar/hem shapes.

Overall, our GPU-parallel MPM simulator provides a balanced operating point, combining fast rollouts with reliable handling of frictional and self-collision-rich interactions, essential for *Right-Side-Out*.

C. Policy Comparison in Simulation

For each garment we run 16 trials and report the final success rate S together with three stage-conditional rates S_1 , S_2 , S_3 . S counts episodes whose final state achieves right-side-out coverage (Fig. 3). S_1 is the success rate of DRAG, producing a hem-side single-layer region; S_2 is the success rate of FLING, exposing an upper-side hem aperture, conditioned on S_1 success; and S_3 is the final success rate conditioned on S_1 and S_2 . Note that $S = S_1 \times S_2 \times S_3$.

We evaluate the policy on unseen garments in simulation; the results are reported in Table I. The *Heuristic* baseline is a rule-based controller that utilizes the garment mesh state; these manually specified rules underperform our learned primitive policies, which were trained on a large amount of data with supervised labels, collected and evaluated using our simulator. We also train two end-to-end imitation learn-

TABLE I: Method comparison in simulation.

Methods	Output	$S_1(\%) \uparrow$	$S_2(\%) \uparrow$	$S_3(\%) \uparrow$	$S(\%) \uparrow$
Heuristic	Prim	100.0	81.3	61.5	50.0
	Prim	100.0	93.8	86.7	81.3
DP [7]	Joint	87.5	50.0	42.9	18.8
ACT [8]	Joint	93.8	66.7	55.6	31.3

ing methods, Action-chunked Transformer (ACT) [8] and Diffusion Policy (DP) [7], using the successful portion of the same data with the same perception. These benchmarks take the per-frame observation as input and output 14-DoF bimanual joint actions at 20Hz, producing around 800 action steps per episode. In contrast, our method predicts only five image-plane keypoints to parameterize the primitives. While ACT and DP learn DRAG and FLING reasonably well, they struggle in the long-horizon, contact-rich INSERT&PULL stage under severe occlusions and diverse observation states.

D. Policy Comparison in Real-world

We evaluate the zero-shot transferability of the methods in Section V-C by training only in simulation and deploying on hardware without any fine-tuning. We test four garments (#1-#4 in Fig. 8) under identical initializations and action budgets for a fair comparison. Results are in Table II.

Imitation learning baselines such as ACT exhibit a severe sim-to-real gap and rarely succeed in the real world, even in DRAG and FLING. In contrast, our keypoint-conditioned primitive policy transfers zero-shot robustly. We further ablate the D3RoMA [43] depth refinement, success rates drop due to noisy single-view depth and D3RoMA [43] markedly alleviates these artifacts and restores performance by stabilizing keypoint predictions and grasp selection.

E. Generalization of Garments in Simulation

We assess material-level generalization entirely in simulation using an unseen garment mesh. The density is fixed at 1000 kg/m^3 and the Young's modulus ranges from 300 to 3000 kPa, which lies within the training randomization range. We vary E as a representative stiffness parameter; other

TABLE II: Method comparison in real-world.

Methods	$S_1(\%) \uparrow$	$S_2(\%) \uparrow$	$S_3(\%) \uparrow$	$S(\%) \uparrow$
ACT [8]	12.5	0.0	0.0	0.0
Ours (w/ [43])	100.0	93.8	86.7	81.3
Ours (w/o [43])	62.5	60.0	50.0	18.8

TABLE III: Simulation performance across materials.

E (kPa)	$S_1(\%) \uparrow$	$S_2(\%) \uparrow$	$S_3(\%) \uparrow$	$S(\%) \uparrow$
300	93.8	93.3	85.7	75.0
1000	100.0	93.8	93.3	87.5
3000	100.0	100.0	87.5	87.5

randomized material/contact parameters are fixed here for clarity. As reported in Table III, the trained policy maintains a consistently high success rate across this stiffness range, indicating that (i) our domain-randomized data generation adequately covers the task’s material variability and (ii) the simulator serves as a reliable testbed for collecting and evaluating policies under diverse physical parameters.

F. Generalization of Garments in Real-World

We evaluate zero-shot transfer on ten upper garments spanning variations in color, size, fabric, and collar/hem geometry (Fig. 7). The set includes eight sleeveless tops (#1-#8) and two dresses (#9 and #10) to stress test category shift (Fig. 8). Across items, mass lies in [0.055, 0.105] kg, body length in [57, 70] cm, hem width in [47, 56] cm, collar-part length in [20, 25] cm, and collar-part width in [4, 10] cm.

We summarize results in Table IV. Overall, the policy attains a high success rate on sleeveless tops, proving that the keypoint-conditioned primitives generalize across textures and geometry changes. Success tends to increase with garment size: larger hems provide a wider access opening, which eases gripper insertion and stabilizes INSERT&PULL. The fabric of #1 and #2 is Modal[®], where pronounced wrinkling during manipulation reveals a small sim-to-real gap. For stress tests, #9 exhibits the lowest performance because its opening is comparable to the manipulator’s effective cross-section, making insertion kinematically challenging; in contrast, #10 demonstrates that the method reasonably generalizes to novel categories as well when geometry permits.

Failure modes. Rare failures arise when rollouts drift outside the primitives’ expected state distribution. The most common failure is an undersized hem opening: minor pose errors and friction during DRAG/FLING or suboptimal INSERT&PULL keypoint selection yield an aperture that exists but is too narrow or unstable for insertion. A distinct failure is that DRAG/FLING do not expose an upper-side single-layer hem region, so no opening can be formed and INSERT&PULL cannot be initiated. We also observe occasional perception U-Net misprediction under severe self-occlusion. A lightweight state-monitoring and recovery module, e.g., an auxiliary U-Net that classifies the current state and triggers autonomously re-canonization could improve robustness.

VI. CONCLUSION

We presented *Right-Side-Out*, a zero-shot sim-to-real framework for garment reversal that factors the task into DRAG, FLING, and INSERT&PULL. The policy operates with depth-driven observations and compact primitives parameterization, and is supervised entirely in a high-fidelity, GPU-parallel MPM simulator that models thin-shell deformation, self-collision, and frictional robot–cloth contact. On a dual-arm platform, the approach achieves robust right-side-out

TABLE IV: Real-world performance across garments.

Metrics	Garment ID									
	#1	#2	#3	#4	#5	#6	#7	#8	#9	#10
$S_1(\%) \uparrow$	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	75.0	87.5
$S_2(\%) \uparrow$	93.8	81.3	93.8	100.0	93.8	87.5	81.3	87.5	58.3	71.4
$S_3(\%) \uparrow$	80.0	84.6	86.7	100.0	86.7	85.7	84.6	71.4	57.1	70.0
$S(\%) \uparrow$	75.0	68.8	81.3	100.0	81.3	75.0	68.8	62.5	25.0	43.8

performance across diverse garments, and ablations confirm the importance of task decomposition and simulator fidelity.

While we have focused on sleeveless upper garments, extending the framework to broader categories (e.g., long-sleeve shirts and pants) that introduce new topology and self-occlusion patterns is a promising next step.

REFERENCES

- [1] A. Ganapathi, P. Sundaresan, B. Thananjeyan, A. Balakrishna, D. Seita, J. Grannen, M. Hwang, R. Hoque, J. E. Gonzalez, N. Jamali *et al.*, “Learning dense visual correspondences in simulation to smooth and fold real fabrics,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 11 515–11 522.
- [2] J. Hietala, D. Blanco-Mulero, G. Alcan, and V. Kyrki, “Learning visual feedback control for dynamic cloth folding,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 1455–1462.
- [3] H. Ha and S. Song, “Flingbot: The unreasonable effectiveness of dynamic manipulation for cloth unfolding,” in *Conference on Robot Learning*. PMLR, 2022, pp. 24–33.
- [4] A. Canberk, C. Chi, H. Ha, B. Burchfiel, E. Cousineau, S. Feng, and S. Song, “Cloth funnels: Canonized-alignment for multi-purpose garment manipulation,” *arXiv preprint arXiv:2210.09347*, 2022.
- [5] Y. Chen, S. Wei, B. Xiao, J. Lyu, J. Chen, F. Zhu, and H. Wang, “Robohanger: Learning generalizable robotic hanger insertion for diverse garments,” *IEEE Robotics and Automation Letters*, 2025.
- [6] W. Chen, K. Li, D. Lee, X. Chen, R. Zong, and P. Kormushev, “Graphgarment: Learning garment dynamics for bimanual cloth manipulation tasks,” *arXiv preprint arXiv:2503.05817*, 2025.
- [7] C. Chi, Z. Xu, S. Feng, E. Cousineau, Y. Du, B. Burchfiel, R. Tedrake, and S. Song, “Diffusion policy: Visuomotor policy learning via action diffusion,” *The International Journal of Robotics Research*, p. 02783649241273668, 2023.
- [8] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, “Learning fine-grained bimanual manipulation with low-cost hardware,” *arXiv preprint arXiv:2304.13705*, 2023.
- [9] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter *et al.*, “ $\pi 0$: A vision-language-action flow model for general robot control. corr, abs/2410.24164, 2024. doi: 10.48550,” *arXiv preprint ARXIV.2410.24164*, 2024.
- [10] T. Z. Zhao, J. Tompson, D. Driess, P. Florence, K. Ghasemipour, C. Finn, and A. Wahid, “Aloha unleashed: A simple recipe for robot dexterity,” *arXiv preprint arXiv:2410.13126*, 2024.
- [11] Y. Avigal, L. Berscheid, T. Asfour, T. Kröger, and K. Goldberg, “Speedfolding: Learning efficient bimanual folding of garments,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 1–8.
- [12] F. Zhang and Y. Demiris, “Learning garment manipulation policies toward robot-assisted dressing,” *Science robotics*, vol. 7, no. 65, p. eabm6010, 2022.
- [13] S. Kotsovolis and Y. Demiris, “Model predictive control with graph dynamics for garment opening insertion during robot-assisted dressing,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 883–890.

- [14] X. Lin, Y. Wang, Z. Huang, and D. Held, "Learning visible connectivity dynamics for cloth smoothing," in *Conference on Robot Learning*. PMLR, 2022, pp. 256–266.
- [15] C. Zhao, C. Jiang, J. Cai, M. Y. Wang, H. Yu, and Q. Chen, "Flipbot: learning continuous paper flipping via coarse-to-fine exteroceptive-proprioceptive exploration," *arXiv preprint arXiv:2304.02253*, 2023.
- [16] C. Zhao, C. Jiang, L. Luo, S. Yuan, Q. Chen, and H. Yu, "Learning thin deformable object manipulation with a multi-sensory integrated soft hand," *IEEE Transactions on Robotics*, 2025.
- [17] L. Sun, G. Aragon-Camarasa, P. Cockshott, S. Rogers, and J. P. Siebert, "A heuristic-based approach for flattening wrinkled clothes," in *Conference Towards Autonomous Robotic Systems*. Springer, 2013, pp. 148–160.
- [18] B. Willimon, S. Birchfield, and I. Walker, "Model for unfolding laundry using interactive perception," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2011, pp. 4871–4876.
- [19] H. Xue, Y. Li, W. Xu, H. Li, D. Zheng, and C. Lu, "Unifolding: Towards sample-efficient, scalable, and generalizable robotic garment folding," *arXiv preprint arXiv:2311.01267*, 2023.
- [20] W. Chen, H. Xue, F. Zhou, Y. Fang, and C. Lu, "Deformpam: Data-efficient learning for long-horizon deformable object manipulation via preference-based action alignment," *arXiv preprint arXiv:2410.11584*, 2024.
- [21] J. Matas, S. James, and A. J. Davison, "Sim-to-real reinforcement learning for deformable object manipulation," in *Conference on Robot Learning*. PMLR, 2018, pp. 734–743.
- [22] Y. Wu, W. Yan, T. Kurutach, L. Pinto, and P. Abbeel, "Learning to manipulate deformable objects without demonstrations," *arXiv preprint arXiv:1910.13439*, 2019.
- [23] R. Lee, D. Ward, V. Dasagi, A. Cosgun, J. Leitner, and P. Corke, "Learning arbitrary-goal fabric folding with one hour of real robot experience," in *Conference on Robot Learning*. PMLR, 2021, pp. 2317–2327.
- [24] 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 *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 4568–4575.
- [25] R. Bridson, R. Fedkiw, and J. Anderson, "Robust treatment of collisions, contact and friction for cloth animation," in *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, 2002, pp. 594–603.
- [26] D. Harmon, E. Vouga, R. Tamstorf, and E. Grinspun, "Robust treatment of simultaneous collisions," in *ACM SIGGRAPH 2008 papers*, 2008, pp. 1–4.
- [27] M. Li, Z. Ferguson, T. Schneider, T. R. Langlois, D. Zorin, D. Panozzo, C. Jiang, and D. M. Kaufman, "Incremental potential contact: intersection-and inversion-free, large-deformation dynamics," *ACM Trans. Graph.*, vol. 39, no. 4, p. 49, 2020.
- [28] M. Li, D. M. Kaufman, and C. Jiang, "Codimensional incremental potential contact," *ACM Trans. Graph. (SIGGRAPH)*, vol. 40, no. 4, 2021.
- [29] W. Du, C. Yu, S. Ma, Y. Jiang, Z. Zong, Y. Yang, J. Masterjohn, A. Castro, X. Han, and C. Jiang, "Embedded ipc: Fast and intersection-free simulation in reduced subspace for robot manipulation," *arXiv preprint arXiv:2409.16385*, 2024.
- [30] M. Müller, B. Heidelberger, M. Hennix, and J. Ratcliff, "Position based dynamics," *Journal of Visual Communication and Image Representation*, vol. 18, no. 2, pp. 109–118, 2007.
- [31] M. Macklin, M. Müller, and N. Chentanez, "Xpbd: position-based simulation of compliant constrained dynamics," in *Proceedings of the 9th International Conference on Motion in Games*, 2016, pp. 49–54.
- [32] C. Jiang, T. Gast, and J. Teran, "Anisotropic elastoplasticity for cloth, knit and hair frictional contact," *ACM Transactions on Graphics (TOG)*, vol. 36, no. 4, pp. 1–14, 2017.
- [33] X. Han, T. F. Gast, Q. Guo, S. Wang, C. Jiang, and J. Teran, "A hybrid material point method for frictional contact with diverse materials," *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, vol. 2, no. 2, pp. 1–24, 2019.
- [34] M. Macklin, M. Müller, N. Chentanez, and T.-Y. Kim, "Unified particle physics for real-time applications," *ACM Transactions on Graphics (TOG)*, vol. 33, no. 4, pp. 1–12, 2014.
- [35] Y. Li, J. Wu, R. Tedrake, J. B. Tenenbaum, and A. Torralba, "Learning particle dynamics for manipulating rigid bodies, deformable objects, and fluids," *arXiv preprint arXiv:1810.01566*, 2018.
- [36] X. Lin, Y. Wang, J. Olkin, and D. Held, "Softgym: Benchmarking deep reinforcement learning for deformable object manipulation," in *Conference on Robot Learning*. PMLR, 2021, pp. 432–448.
- [37] W. Du, S. Yao, X. Wang, Y. Xu, W. Xu, and C. Lu, "Intersection-free robot manipulation with soft-rigid coupled incremental potential contact," *IEEE Robotics and Automation Letters*, vol. 9, no. 5, pp. 4487–4494, 2024.
- [38] S. Ma, W. Du, C. Yu, Y. Jiang, Z. Zong, T. Xie, Y. Chen, Y. Yang, X. Han, and C. Jiang, "Grip: A general robotic incremental potential contact simulation dataset for unified deformable-rigid coupled grasping," *arXiv preprint arXiv:2503.05020*, 2025.
- [39] Y. Li, W. Du, C. Yu, P. Li, Z. Zhao, T. Liu, C. Jiang, Y. Zhu, and S. Huang, "Taccel: Scaling up vision-based tactile robotics via high-performance gpu simulation," *arXiv preprint arXiv:2504.12908*, 2025.
- [40] Z. Zong, C. Jiang, and X. Han, "A convex formulation of frictional contact for the material point method and rigid bodies," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 1831–1838.
- [41] C. Yu, W. Du, Z. Zong, A. Castro, C. Jiang, and X. Han, "A convex formulation of material points and rigid bodies with gpu-accelerated async-coupling for interactive simulation," *arXiv preprint arXiv:2503.05046*, 2025.
- [42] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo *et al.*, "Segment anything," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 4015–4026.
- [43] S. Wei, H. Geng, J. Chen, C. Deng, C. Wenbo, C. Zhao, X. Fang, L. Guibas, and H. Wang, "D³roma: Disparity diffusion-based depth sensing for material-agnostic robotic manipulation," in *ECCV 2024 Workshop on Wild 3D: 3D Modeling, Reconstruction, and Generation in the Wild*, 2024.
- [44] C. Jiang, C. Schroeder, A. Selle, J. Teran, and A. Stomakhin, "The affine particle-in-cell method," *ACM Transactions on Graphics (TOG)*, vol. 34, no. 4, pp. 1–10, 2015.
- [45] L. Y. Chen, B. Shi, R. Lin, D. Seita, A. Ahmad, R. Cheng, T. Kollar, D. Held, and K. Goldberg, "Bagging by learning to singulate layers using interactive perception," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 3176–3183.
- [46] 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.
- [47] F. Xiang, Y. Qin, K. Mo, Y. Xia, H. Zhu, F. Liu, M. Liu, H. Jiang, Y. Yuan, H. Wang *et al.*, "Sapien: A simulated part-based interactive environment," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 097–11 107.