StressLess: Towards Generalizable Deformable and Fragile Object Grasping from Stress-Aware Demonstrations at Scale

Author Names Omitted for Anonymous Review

***Abstract*— Grasping fragile and deformable objects is challenging because contact stress beyond a material's yield stress can cause irreversible damage. Previous works typically learn a policy for such a task via human demonstrations and with additional hardware such as tactile sensors and/or compliant grippers. However, learning gentle grasping behavior from human demonstrations is costly since achieving coverage across object categories, poses, sizes, shapes, and materials requires extensive physical data collection. Towards addressing this challenge, we present STRESSLESS, a framework that learns manipulation policies from stress-aware demonstrations generated at scale in simulation. A finite-element surrogate predicts the internal stress induced by gripper closure, and a scripted planner jointly selects the grasp pose and closing width that minimize stress while maintaining sufficient grip to resist gravity. By extensive randomization, STRESSLESS generates over 6k diverse demonstrations across 34 object categories without any human tele-operation. We use these data to train a single point-cloud-conditioned Diffusion Policy and deploy it zero-shot in the real world. In challenging real-world evaluations across 11 object categories and broad variations in object pose and size, our method outperforms baselines trained on real-world demonstrations in both lifting and gentleness. These results indicate that physics-based simulated supervision can efficiently provide the broad coverage required for generalizable gentle manipulation.**

I. INTRODUCTION

Humans handle delicate items such as food almost effortlessly, yet the same task remains challenging for robots. Automating this capability is important in food processing, agriculture, where compression can tear, crack, or bruise produce and cause irreversible damage [1], [2], [3]. Gentle grasping requires the robot to jointly select the contact configuration, where and at what orientation the fingers close, and an appropriate closing width. These decisions are coupled: a closing width that is appropriate at one location may cause failure across a thin edge or when the gripper straddles a stem. Moreover, both decisions depend on the item's pose, size, and shape. These properties vary continuously within a food category and can differ entirely across categories, so success on one item does not necessarily transfer to another.

Existing approaches address this challenge primarily through hardware design and contact sensing. Compliant and pneumatic end-effectors mechanically limit contact pressure [4], while tactile or force-controlled grippers regulate contact at run time and can handle delicate items effectively [5]. These approaches improve how contact is regulated or observed, but they do not address the source of policy *supervision*: policies are still learned by imitation from human demonstrations recorded on real hardware, one item and one pose at a time. Achieving *broad coverage* of object categories and physical variations therefore remains a human-labour intensive problem.

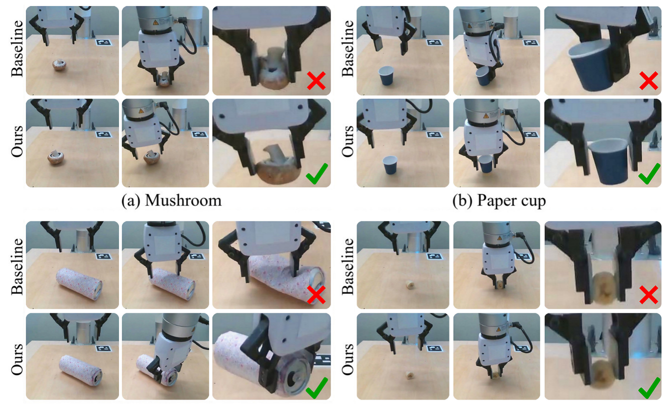

(a) Mushroom (b) Paper cup (c) Can (d) Banana

Fig. 1. **Cross-category gentle grasping.** To our knowledge, STRESSLESS is the first single policy trained entirely on simulation-generated, stress-aware demonstrations and deployed zero-shot across diverse deformable and fragile objects. Prior methods train per category [6], [5] or use contact feedback and real demonstrations [7], [8]. Here, STRESSLESS grasps safely, whereas a Diffusion Policy [9] trained on 120 real demonstrations fails.

The importance of such coverage is supported by recent scaling results. Lin *et al.* [10] show that policy performance follows a power law as the number of environments and objects increases, whereas the benefit of additional demonstrations per object saturates quickly. Thus, *diversity*, rather than training data size alone, drives generalization. The data requirement becomes even more severe for high-precision tasks: the number of demonstrations grows super-exponentially as the tolerance tightens, while the limiting precision itself depends on the quality of the expert [11]. Existing data-generation systems reduce collection costs by adapting a few hundred human demonstrations to new object placements, producing tens of thousands of episodes [12]. However, this adaptation (specifically on grasp width/strength) is primarily driven by the geometry of the object and does not necessarily adapt intelligently to the appropriate grasping force required for a novel item.

To address this limitation, we propose STRESSLESS, a framework that autonomously generates gentle demonstrations in simulation and uses them to train a single policy across multiple fragile food categories. The central insight is that internal *stress* of fragile objects, which determines whether an item is damaged but is not measured by any sensor on our robot, can be computed sufficiently accurate with soft body FEM solver. This allows gentleness to be formulated as a planning objective rather than as a signal that must be sensed at run time. A finite-element model predicts the stress induced by a candidate grasp pose and closing width, and a scripted demonstrator selects one that minimizes stress while still holding the item against gravity. We execute this demonstrator across 34 fragile food categories, randomizing object pose, size, shape, and material within each, to collect over 6k demonstrations without human teleoperation. Using the pooled data, we train a single policy and deploy the policy on a 7-DoF arm that observes the scene as a point cloud obtained from a single external depth camera, a sensing modality selected for its narrow sim-to-real appearance gap; no tactile or force sensing is used in the control loop. In real-world experiments, the policy grasps objects from 11 categories with a 67% success rate while handling in-category variation.

We make three contributions: (1) a stress-aware grasp planner based on a grasp-width-controlled, distributed-pad contact model that selects both the grasp location and closing width for a deformable item; (2) a single policy that handles *diverse* deformable and fragile object categories with only point-cloud observations; and (3) 275 real-world evaluations demonstrating strong performance across multiple categories and in-category variations in pose, size, and shape, over baselines. Together, these contributions represent a step towards broadly applicable gentle manipulation of everyday food.

II. RELATED WORKS

A. Gentle manipulation of deformable and fragile objects

Damage-free handling of deformable and fragile objects is a long-standing goal in agricultural and food automation and remains open across the harvest pipeline [3]. Much of the existing work addresses this problem through end-effector mechanics: compliant, pneumatic, and enveloping grippers limit contact pressure by construction [4], [1]. Online feedback provides a complementary approach: tactile sensors enable fruit recognition and grasp-force prediction for damage-free handling [13]. These approaches mitigate damage through mechanical design or contact feedback. Our work instead focuses on the source of policy supervision, using a physics-aware objective to select both the grasp pose and closing width for a fixed parallel-jaw gripper.

A parallel line learns the behavior from demonstrations, often using contact feedback to modulate the grasp [6], [8], [5], [7]. Such feedback is valuable and orthogonal to our contribution; what it does not address is coverage. In these systems supervision comes from real interaction — teleoperated demonstrations or robot-collected trials — so broadening the range of objects a policy handles means collecting proportionally more real data. We target that bottleneck instead: because supervision is synthesized in simulation, a new object category costs a mesh and a material rather than a full-round real-world collection.

B. Grasp quality for deformable objects

Classical analytic grasp metrics assume rigid objects and quantify force closure or robustness to external wrenches [14]. Such criteria capture whether a grasp can resist disturbances, but not the deformation or damage caused by the grasp itself. For deformable objects, Pan *et al.* [15] address this requirement with a stress-minimizing metric that accounts for object material and identifies fragile regions. However, [15] assumes point contacts and general multi-contact grasps, while we target gripper finger pads. Finite-element simulation has also been used to characterize deformable grasp outcomes at scale. DefGraspSim [16] evaluates parallel-jaw grasps across a large object set and reports stress, deformation, and stability metrics with good sim-to-real correspondence. Le *et al.* [17] simulate deformation-dependent contacts to construct a grasp quality metric that predicts grasp success more reliably. Fragile objects impose the additional requirement of limiting internal stress. Both employ a simulation engine to perform the task while we use a computationally efficient linear FEM for grasp quality scoring.

C. Data diversity in robot learning

Data diversity is a key driver of generalization in robot learning. In imitation learning, Lin *et al.* [10] find that generalization follows a power law in the number of environments and objects, whereas the benefit of additional demonstrations per object saturates quickly; coverage therefore matters more than demonstration count. The need for coverage becomes even more pressing in high-precision tasks, for which [11] have observed a super-exponential growth in data requirements as the tolerance tightens. Play2Perfect [18] pretrains a task-agnostic dexterous policy through reinforcement learning on procedurally generated objects and random goal trajectories. It shows that increasing object diversity and sampling diverse trajectories during pretraining improve transfer to precise assembly. These results indicate diverse experience can be acquired through autonomous interaction.

Generating such diversity on physical systems, however, remains expensive. MimicGen [12] addresses this cost by transforming a few hundred human demonstrations into tens of thousands of episodes with new object placements. Although it increases geometric coverage, they retain the grasps selected by a human. Our method instead generates diverse, object-level demonstrations efficiently in simulation. By randomizing object pose, size, shape, and material and re-planning each grasp with a stress-aware physics objective, the demonstrator adapts both the grasp pose and closing width to every sampled configuration. The resulting supervision therefore captures physical as well as geometric variation without requiring human tele-operation. A similar line to our work is [19], [20] which trains an Reinforcement Learning expert and generate large set of demonstration for sim-to-real transfer. We adopt similar idea but with a scripted demonstrator.

III. METHOD

A. Problem Setting

We consider a robot arm with a parallel-jaw gripper that must lift a single fragile object from a table and hold it

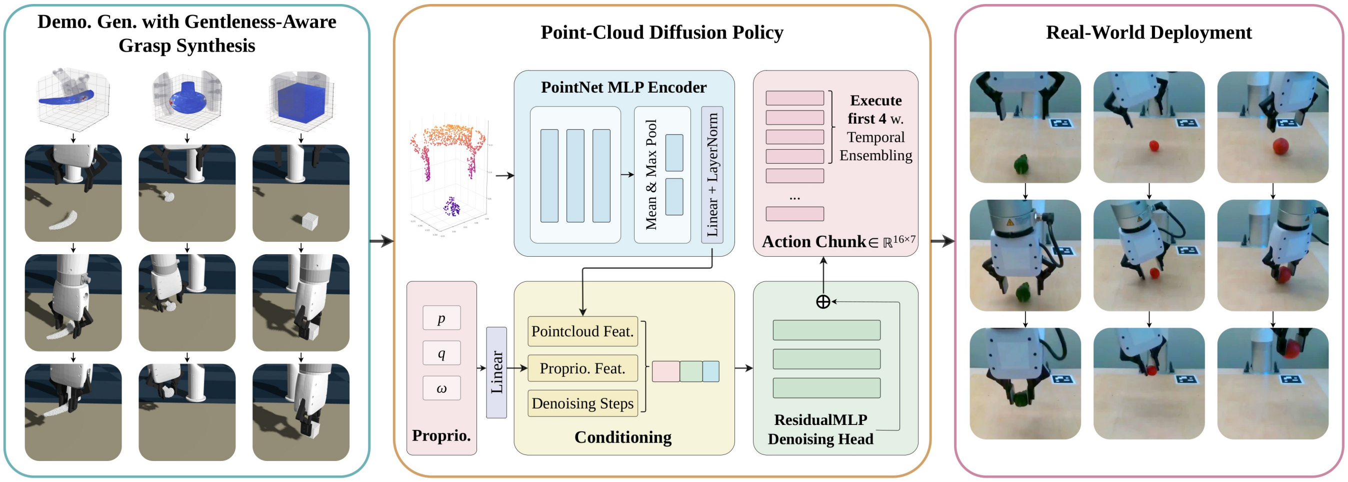

Fig. 2. **Overview** of StressLess. A finite-element contact model scores candidate grasps by the stress they impose, and the gentlest feasible grasp is executed in Material Point Method (MPM) simulation to generate demonstrations without tele-operation *(left)*. A single point-cloud diffusion policy is trained on these demonstrations across all object categories *(center)*, and transfers zero-shot to the real robot from one depth camera *(right)*

without damaging it. Object category, pose, size, shape and material vary across episodes. Perception is provided by a single fixed depth camera; there is no tactile or force sensing in the loop. We make this omission deliberately: tactile feedback has been shown to improve gentle manipulation [6], and we do not claim otherwise. Our question is instead one of coverage — how far a single policy trained on synthesized demonstrations extends across object categories — and tactile signals remain difficult to simulate faithfully, so including them would introduce a second, challenging sim-to-real transfer problem without bearing on that question.

At each control step, the policy network observes $o_t = (\mathbf{p}_t, \mathbf{q}_t, \omega_t, \mathcal{P}_t)$: end-effector position $\mathbf{p}_t$, orientation as a unit quaternion $\mathbf{q}_t$, measured gripper width ω_t, and a point cloud $\mathcal{P}_t$ in the robot base frame. It emits an *absolute* target pose and width $a_t = (\mathbf{p}_t^{\mathrm{cmd}}, \boldsymbol{\theta}_t^{\mathrm{cmd}}, \omega_t^{\mathrm{cmd}}) \in \mathbb{R}^7$, with orientation as Euler angles. An episode succeeds if the object centre enters a height band. We evaluate gentleness in two ways: (1) in simulation, whether the top-decile internal von Mises stress stays below the material's yield threshold; (2) in the real world, a 0–5 damage scale scored by more than one human evaluator.

B. Stress-Aware Grasp Synthesis

For each object instance the demonstrator synthesizes a 7-DoF grasp $g = (\mathbf{x}_g, \mathbf{r}_g, \omega)$: tool-center position $\mathbf{x}_g$, orientation $\mathbf{r}_g$, and commanded gripper width ω. The grasp pose is executed by a demonstration generator to generate policy training data in simulation.

The synthesizer scores candidates with a grasp-width-controlled FEM contact model, filters them through the feasibility gates of Table I, and searches for the gentlest stable grasp under a relaxation ladder. We detail each in turn below.

Grasp-width-controlled contact model. Common parallel jaw grippers are position-controlled: one commands a width and the grip force is an outcome. We model this directly. The object is tetrahedralised and a linear-elastic stiffness matrix

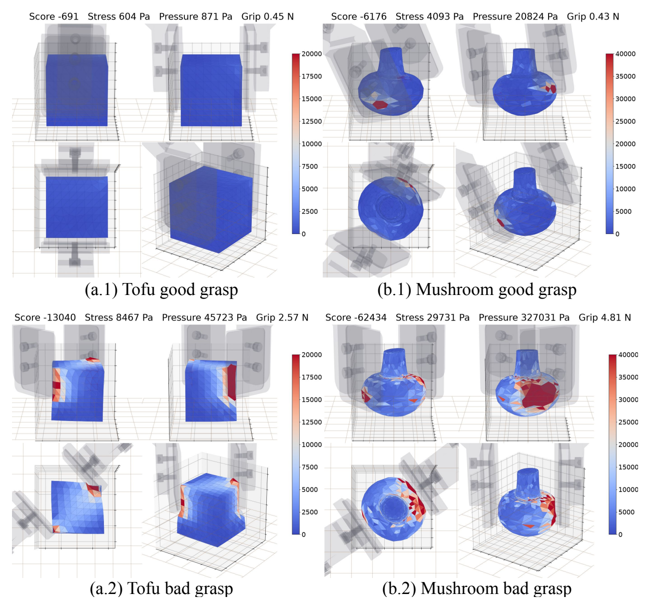

Fig. 3. Finite Element Method (FEM) simulation of good and sub-optimal grasps on two objects, tofu and mushroom.

assembled; since it is free-floating the system is singular in the six rigid-body modes, so we solve a bordered inertia-relief system, factored once per geometry. Two rigid pads, whose rectangular footprints are taken from the real gripper's finger meshes, are closed to width ω. Boundary nodes inside a footprint are pushed onto the pad's flat plane and constrained *only* along the closing axis, leaving tangential motion free so the material bulges; a thin smooth-step fillet at the pad edge avoids a flat-punch stress singularity. Contact is normal-only, and friction enters solely through the stable gate introduced below. In Figure 3 we show the modeling outcome of good and sub-optimal grasps on two objects, clear over deformation can be seen in (a.2) and (b.2).

Material scaling. Because the gripper finger displacement is prescribed rather than the force, the deformation field is independent of Young's modulus, instead is dictated by the

TABLE I

GRASP POSE FEASIBILITY GATES. SEE III-B FOR NOTATION.

Gate	Feasible when
Table	fingers stay above the table plane
Penetration	fingers have penetration into object (nominal) $\leq c$
Contact	each finger captures a non-empty contact patch
Indentation	$\delta \leq \delta_{\max}$ (linear-elastic regime)
Stable Grasp	$2\mu F \geq m\,(g+a)$
Torsion	$\lvert \boldsymbol{\tau} \cdot \hat{\mathbf{a}} \rvert \leq \frac{4}{3}\,\mu F \sqrt{A_{\min}/\pi}$
Yield	$\sigma_{\text{blk}} \leq \sigma_Y$

grasp width ω. Thus when below yield, both stress and grip force are linear in it,

$$\sigma(E) = E\,\sigma_1, \qquad F(E) = E\,F_1. \tag{1}$$

where σ_1 is the element stress field and F_1 the per-jaw normal reaction from a single solve at $E=1$. One solve at $E=1$ therefore yields every material by scalar multiplication, which makes randomizing stiffness across the demonstration set computationally efficient. We consider a grasp stable when Coulomb friction at the gripper's contact surfaces can withstand the object's mass m with an acceleration margin, $2\mu F(E) \geq m\,(g+a)$. We in particular use $a=g$.

Objective. A candidate is *feasible* when it passes every gate in Table I, where c bounds finger penetration into the nominal mesh, δ is the pad indentation, $\hat{\mathbf{a}}$ the closing axis, $\boldsymbol{\tau}$ the gravitational torque about the center of mass, and σ_Y the yield stress.

Among feasible candidates, we maximize a score:

$$S(x) \;=\; -\,\sigma_{\text{blk}} \;-\; w_{\text{pr}}\,\frac{F}{A_{\min}}, \tag{2}$$

where σ_{blk} be the mean of the top decile of element stresses with contact-adjacent elements masked out, so that it reports *bulk* stress rather than the mesh-dependent contact singularity. $A_{\min}$ is the smallest contact area among the left and right finger pads, thus the second term is the mean contact pressure. Bulk gentleness and local gentleness are thus both represented, and neither has a manually set floor.

Algorithm 1 Gentleness-aware grasp synthesis

Require: mesh, pose, material (E, ρ, μ), table height
Ensure: grasp $g = (\mathbf{x}_g, \mathbf{r}_g, \omega)$

```
assemble FEM, factor the bordered system once
G0 ← antipodal grasps ∪ medial-axis grasps
for tier t = 0, 1, 2 do
    bounds ← RELAX(t)
    G1 ← {g ∈ G0 : table, orientation, penetration}
    score G1; G2 ← top K among G1
    G3 ← CMA-ES(G2) over all 7 DoF
    G4 ← spatially distinct poses of G3
    g* ← arg max_g S(g) over a 1-DoF width sweep at each pose of G4
    if g* is a stable grasp then
        return g*
    end if
end for
return most-centered feasible grasp seed    ▷ fallback when no stable grasp solution exists
```

Grasp search. The grasp search is a funnel (Algorithm 1). Grasp candidates are seeded by two geometric generators — antipodal surface pairs within the friction cone, and medial-axis points closing across the local tangent — yielding 6000 7-DoF poses with ω initialized to the object local cross-section. Seeds are filtered on the two efficiently batched gates (table and penetration); survivors are then scored by Equation 2 and filtered on feasibility (Tab. I) using FEM solvers. The top K candidates serve as seeds for subsequent optimization with CMA-ES. An orientation box bounds the optimization. Poses outside the bound are never evaluated. It holds roll and pitch away from tilts at which the wrist would strike the table, and limits yaw so the fingers do not fully occlude the object in the camera view the policy must act on. The CMA-ES optimization searches all seven degrees of freedom of a grasp jointly. At last, a final pass then fixes each of the best spatially distinct poses and sweeps width ω alone on a grid finer than the optimizer's step, retaining the highest-scoring stable grasp.

Relaxation ladder. After a single round of search, edge cases on some objects give no feasible solution, typically because the stable grasp gate conflicts with the penetration and/or indentation gates. The funnel is therefore run inside a short ladder. The first pass uses nominal bounds. The second widens the orientation box and the penetration tolerance while *doubling* the pressure weight of Eq. (2), so that the constraints loosen and the objective tightens together. The third additionally releases the yield gate, accepting a grasp that to some extent damages the object rather than one that fails to lift it. This third pass is used rarely — it was reached on 2.9% of planned grasps during data collection. The ladder stops at the first pass that yields a stable grasp thus costs nothing on objects that do not need it; if none of the three succeeds, the synthesizer returns its most-centered (relative to center of mass) feasible grasp candidate.

C. Demonstration Generation

Planned grasps are executed in a soft-body simulator, producing demonstration trajectories. The executor runs a phase machine — approach, settle, close, lift, hold.

Randomization. Each episode samples the object's pose, uniform scale, procedural shape and material parameters. Shape is perturbed by Barr bend [21], twist and taper along the object's longest axis, and each category carries several meshes representing distinct variants within the type (Fig. 4).

Drawing insights from OmniReset [20], we further randomize the robot arm initial condition and inject a mid-episode disturbance, so that the demonstrations encourage recovery rather than only clean executions. With probability $\mathbf{P}_{\text{home}} = 0.6$ the arm begins at the nominal pose; otherwise the gripper starts at a random workspace pose ($\mathbf{P} = 0.15$), or 4–12 cm above the grasp point ($\mathbf{P} = 0.15$), or part-way along the path to it ($\mathbf{P} = 0.1$), with the jaws at a random width so that the policy must also re-open before grasping. With probability $\mathbf{P}_{\text{disturb}} = 0.1$ the object is briefly given a lateral velocity during the approach, sliding it by a few centimeters; once it comes to rest, the grasp pose is translated by the measured displacement and the approach restarts from wherever the gripper then is. The demonstrations therefore

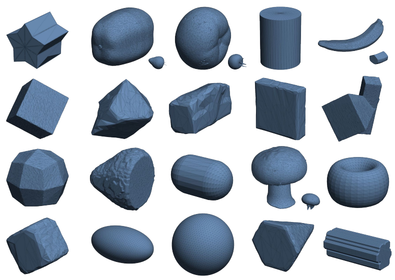

Fig. 4. The representative objects we use in simulation. Several categories have more than one mesh that we randomize over (e.g. mushroom), an example is shown at the bottom-right corner for each such a category. Meshes of organic objects are generated by using TripoSG [25] with manually taken images.

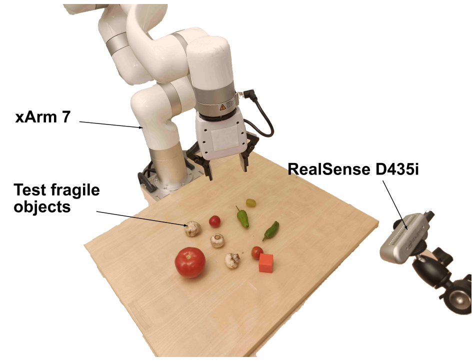

Fig. 5. Real-world experiment setup.

contain targets that move away mid-reach, which a policy trained only on clean trajectories is difficult to handle.

Approach trajectory. Rather than interpolating linearly from the initial arm pose to the grasp pose, the approach runs in two phases. Let a be the grasp's approach axis (the gripper's tool $+z$ axis at the grasp pose). The first phase interpolates position to a stand-by point $\mathbf{x}_{\text{stand}} = \mathbf{x}_g - d\,a$ and slerps orientation to the grasp orientation, where d is the start pose's own distance along a, clamped to $[4, 10]\,\text{cm}$; the gripper opens over this phase. The second phase holds the orientation fixed and translates straight along a into the grasp, so the open fingers clear the object. Both phases are commanded at a constant speed.

D. Policy Learning

As shown in Figure 2, center sub-figure, we train a single point-cloud diffusion policy on all synthesized demonstrations, following 3D Diffusion Policy [22]: a PointNet encoder maps the observed cloud to a global feature, which is concatenated with a short proprioceptive history and conditions a denoising diffusion model over action chunks. Because the demonstrations are generated in simulation, we augment the point cloud during training to make the policy robust to the sim-to-real sensor gap [23]: each cloud is perturbed — distance-dependent axial noise, lateral jitter, random point dropout, and occasional spurious point clusters — together with a fixed extrinsic offset of a few millimeters that stands in for residual calibration error. A chunk of 16 actions is predicted and the first 4 executed before re-planning. At deployment we use temporal ensembling: the overlapping predictions covering each executed step are averaged with exponentially decaying weights [24], which smooths the commanded trajectory across chunk boundaries.

IV. EXPERIMENT

A. Experiment Setup

We share the same experimental setup in simulation and real-world. As shown in Figure 5 We use a 7-DoF Ufactory xArm 7 equipped with its standard two-finger parallel gripper. An external Intel RealSense D435-i RGB-D camera is positioned in front of the arm, observing the workspace diagonally. Simulation is conducted using Genesis [28].

B. Simulation Experiment

Grasp synthesizer comparison. We compare various grasp synthesizer in simulation against our proposed stress-aware grasp planner, for deformable and fragile object manipulation. Each baseline replaces only the grasp-selection step: the same controller executes the returned pose and width, point-cloud method (GPD [27]) sees a simulated depth observation and mesh-based methods the same tet-mesh surface. We discard only method-native candidates whose fingers would pass through the table or whose approach exceeds the arm's orientation limits, executing each method's own top-ranked survivor unmodified. Here, we compare the following methods:

- *Naive Top-Down*. Gripper is at the object center, with uniform random yaw. The gripper width is set to the object's measured cross-section along that axis. It encodes no contact reasoning.
- *Antipodal*. Samples surface point pairs whose connecting line lies inside both friction cones [26] and ranks them by cone margin, commanding the pair's separation as the width.
- *GPD* [27]. Samples 6-DoF candidates in the point cloud and scores them with its released CNN. GPD returns, for each candidate, the extent of the observed points enclosed between the fingers—a contact width—which we command directly.
- *Geometry-only (Ours w/o FEM)*. Retains our candidate seeding, CMA-ES search and 1-D width refinement, but substitutes a purely geometric SDF objective [29] and omits the FEM-derived feasibility gates. It constitutes an ablation of our method without its FEM component.

Metrics. Each cell reports a result over 60 episodes. *Succ* is the success rate. Gentleness is measured from the simulated stress field: at every simulation step we take the mean of the top decile of per-particle von Mises stress, keep the maximum over the episode, and normalize by the object's yield stress σ_y, giving a peak-stress ratio per episode. *Sub-y* is the percentage of episodes whose ratio stays below 1, and $\tilde{\sigma}/\sigma_y$ is the median ratio. Both are computed over *successful* episodes only. This is because stress recorded on lift-failed episodes (e.g. the jaws close in air) may have low stress and dilute the aggregated metrics. Lastly, we also report the final mean success rate conditioned on object sub-yield, denote

TABLE II

GRASP SYNTHESIS COMPARISON ON FIVE DEFORMABLE OBJECTS IN MPM SIMULATION ENGINE. METRICS ARE EXPLAINED IN SECTION IV-B. **BOLD** = BEST, UNDERLINE = SECOND BEST; PER-OBJECT STRESS IS CONDITIONED ON SUCCESSFUL GRASPS. *: GPD FAILS COMPLETELY ON STRAWBERRY, SO THE AVERAGE VALUE HERE IS OVER FOUR OBJECTS, NOT FIVE.

Method	Tofu			Mushroom			Strawberry			Pepper			Cherry tomato			Mean Succ.↑	Mean Sub-y↑	Mean Cond.-Succ.↑
	Succ.↑	Sub-y↑	$\tilde{\sigma}/\sigma_y \downarrow$	Succ.↑	Sub-y↑	$\tilde{\sigma}/\sigma_y \downarrow$	Succ.↑	Sub-y↑	$\tilde{\sigma}/\sigma_y \downarrow$	Succ.↑	Sub-y↑	$\tilde{\sigma}/\sigma_y \downarrow$	Succ.↑	Sub-y↑	$\tilde{\sigma}/\sigma_y \downarrow$			
Naive top-down	13.3	100	0.23	71.7	**100**	**0.20**	25.0	**100**	**0.27**	5.0	**100**	**0.23**	86.7	**100**	0.62	40.3	**100.0**	40.3
Antipodal [26]	88.3	100	0.28	90.0	30	1.28	76.7	**100**	0.38	83.3	75	0.84	**100.0**	73	0.79	87.7	75.5	65.4
GPD [27]	1.7	100	**0.22**	25.0	**100**	0.23	0.0	N/A	N/A	20.0	**100**	0.33	73.3	91	**0.58**	24.0	97.7*	22.7
Ours w/o FEM	85.0	100	0.27	**100.0**	30	1.08	**100.0**	45	1.03	**98.3**	19	1.11	81.7	2	1.20	**93.0**	39.2	36.1
STRESSLESS (Ours)	**100.0**	100	0.29	**100.0**	**100**	0.32	96.7	**100**	0.41	73.3	98	0.67	83.3	51	0.98	90.7	89.7	**82.1**

"Cond.-Succ." in Table II, last column.

Results. As shown in Table II, success rate alone does not separate the methods completely: Ours 90.7% and its FEM-free ablation 93.0% are indistinguishable. The separation is in stress. Removing the FEM term drops sub-yield from 89.7% to 39.2%, and the ablation's median peak stress reaches or exceeds yield on four of five objects (1.03–1.20), meaning its typical successful grasp damages the object it lifts. Antipodal sampling behaves similarly at a smaller scale: competitive success rate 87.7% with 75.5% sub-yield and a median of 1.28 on mushroom. The naive baseline inverts the trade-off, achieving 100% sub-yield at 40.3% success — it is gentle only on the subset of configurations it manages to lift. GPD transfers poorly at this scale, succeeding in 24.0% of attempts and not at all on strawberry, so its gentleness figures rest on very few grasps. Our synthesizer is the only method that is both reliable and gentle. However, its failure mode is also visible from the table: sub-yield falls to 51% on cherry tomato, the smallest object with low yield stress in the set, where the feasible margin is narrow.

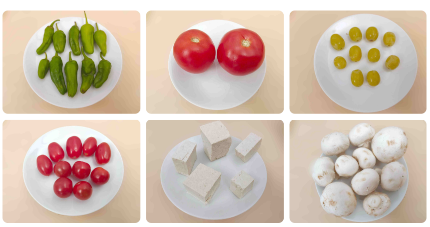

Fig. 6. Real demonstration training set, comprising Padrón peppers, tomatoes, grapes, cherry tomatoes, tofu, and mushrooms, with variation in size and shape within each category. We collect 20 demonstrations per category, summing up to 120 demonstrations for baseline training.

C. Real-World Experiment

We compare against four Diffusion Policy (DP) baselines [9] trained on the real-world demonstrations of object categories in Fig. 6.

- The three RGB baselines use a ResNet-18 encoder: *RGB-frozen* fixes the ImageNet-pretrained encoder and trains only the policy head; *RGB-finetune* jointly trains the pretrained encoder and policy head; and *RGB-scratch* trains both from random initialization.
- *PCD* follows DP3 [22], using point-cloud observations, trained on the same 120 real demonstrations.

For all baselines, we train them with real-world human-teleop demonstration (while our policy is purely trained in sim). This is to test whether a typical budget of human demonstration serves a single policy to handle different deformable and fragile items well. We use six objects (pepper, tomato, grape, cherry tomato, mushroom, and tofu) shown in Figure 6. We collect 20 episodes each, summing to 120 episodes of training data. The evaluation set includes the same six categories represented in the real demonstration set and five absent from it (banana, plum, bread, beverage can, and paper cup). Note that the in-domain evaluation objects do not necessarily perfectly match the profile of that in training set (as they are purchases on different days), we treat this as a real challenge that a policy would face for food manipulation.

Fig. 7. Evaluation sets covering 11 categories: tofu, banana, pepper, cherry tomato, mushroom, bread, tomato, beverage can, paper cup, grape, and plum. Each method uses the same number of instances per category.

Evaluation protocol. We follow recommended experiment protocol proposed by [30]: Two evaluators, blinded to the policy identity, assess each roll-out under similar lighting conditions. Both receive the same category-specific scoring rubrics before evaluation and record quantitative scores and qualitative observations. All methods are run in-turn in the same session to maximize fairness.

Each evaluator assigns two scores in increments of 0.5. The lift score $s_{\text{lift}} \in [0,5]$ is zero if the object is not lifted and five if it reaches at least 20 cm above the tabletop; intermediate scores reflect the estimated lift height. The damage score $s_{\text{damage}} \in [0,5]$ assesses damage throughout approach, grasp, lift, and hold: zero denotes an intact object, and five denotes severe damage, such as crushing or irreversible deformation. Intermediate damage scores follow empirically defined category-specific rubrics that are made clear to each evaluator.

We derive three metrics from each evaluator's scores. Each lies in $[0,1]$, with higher values indicating better performance:

- LIFT $= s_{\text{lift}}/5$ measures normalized lift height.

TABLE III

REAL-WORLD GRASPING RESULTS. COLUMN HEADERS REPORT CATEGORY | NUMBER OF TRIALS. EACH CELL REPORTS (LIFT↑ |SAFE↑ |S-LIFT↑) AS DESCRIBED IN SECTION IV-C. **BOLD** = BEST AND UNDERLINE = SECOND BEST.

Method | # Trials	Mean | 55	Banana | 4	Tofu | 7	Pepper | 2	Plum | 2	Tomato | 2
RGB-frozen	0.18 | 0.53 | 0.36	0.00 | 0.50 | 0.25	0.00 | 0.50 | 0.25	0.00 | 0.50 | 0.25	0.50 | 0.65 | 0.58	0.00 | 0.50 | 0.25
RGB-finetune	0.48 | 0.67 | 0.58	0.00 | 0.50 | 0.25	0.00 | 0.50 | 0.25	1.00 | 0.95 | 0.98	1.00 | 0.80 | 0.90	1.00 | 0.96 | **0.98**
RGB-scratch	0.57 | 0.67 | 0.62	0.50 | 0.61 | **0.56**	0.18 | 0.61 | 0.39	1.00 | 0.95 | 0.98	0.50 | 0.71 | 0.61	0.50 | 0.68 | 0.59
PCD	0.43 | 0.67 | 0.55	0.50 | 0.63 | **0.56**	0.29 | 0.63 | **0.46**	0.50 | 0.75 | 0.63	0.50 | 0.71 | 0.61	0.50 | 0.70 | 0.60
STRESSLESS (Ours)	**0.67** | **0.78** | **0.73**	0.25 | 0.60 | 0.43	0.19 | 0.60 | 0.39	1.00 | 1.00 | **1.00**	1.00 | 0.90 | **0.95**	1.00 | 0.93 | 0.96

Method | # Trials	Mushroom | 6	Bread | 4	Cherry tomato | 10	Grape | 5	Can | 5	Paper cup | 8
RGB-frozen	0.33 | 0.53 | 0.43	0.00 | 0.50 | 0.25	0.01 | 0.52 | 0.26	0.02 | 0.45 | 0.24	0.60 | 0.58 | 0.59	0.50 | 0.62 | 0.56
RGB-finetune	0.33 | 0.62 | 0.48	0.25 | 0.59 | 0.42	0.30 | 0.63 | 0.47	0.00 | 0.50 | 0.25	0.80 | 0.68 | 0.74	0.60 | 0.65 | 0.63
RGB-scratch	0.83 | 0.65 | 0.74	0.25 | 0.61 | 0.43	0.55 | 0.72 | 0.64	0.60 | 0.71 | **0.65**	0.78 | 0.58 | 0.68	0.58 | 0.54 | 0.56
PCD	0.67 | 0.75 | 0.71	0.53 | 0.73 | **0.63**	0.20 | 0.58 | 0.39	0.40 | 0.70 | 0.55	0.00 | 0.50 | 0.25	0.63 | 0.64 | 0.63
STRESSLESS (Ours)	0.83 | 0.86 | **0.85**	0.25 | 0.63 | 0.44	0.60 | 0.77 | **0.69**	0.20 | 0.60 | 0.40	1.00 | 0.88 | **0.94**	1.00 | 0.86 | **0.93**

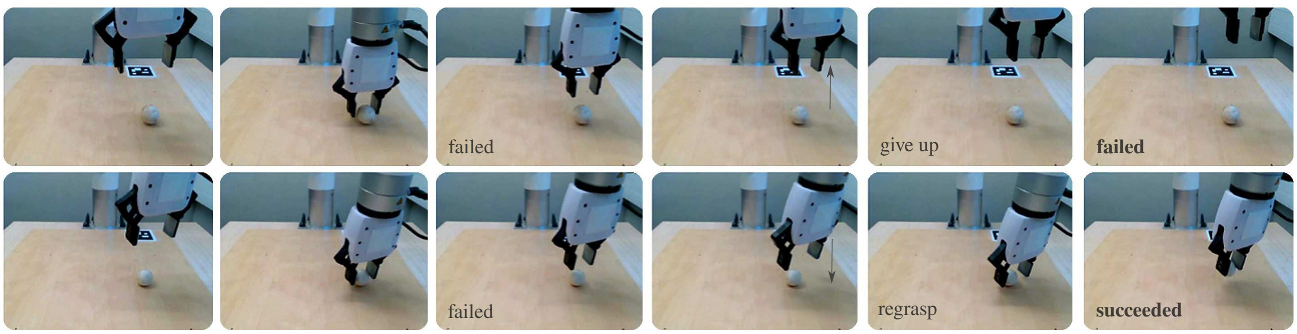

Fig. 8. **Recovery after failure:** The same initial condition where baseline fails a grasp and never retry, while our method can recover from initial failure.

- SAFE $= 0.5 + 0.1r$ is a lift-conditioned safety score, where

$$r = \begin{cases} -s_{\text{damage}}, & s_{\text{lift}} = 0, \\ 5 - s_{\text{damage}}, & s_{\text{lift}} > 0. \end{cases}$$

  Thus, an intact but unlifted object scores 0.5, a severely damaged unlifted object scores 0, and an intact lifted object scores 1.
- S-LIFT $= (\text{LIFT} + \text{SAFE})/2$ combines lifting and safety with equal weight.

We compute these metrics per evaluator and trial, average them within each category, and then average equally across categories.

Trials use a $26 \times 30\,\text{cm}$ workspace, with 55 trials per method distributed across 11 categories as shown in Figure 7 and Table III. Instances vary in size, shape, and fragility within each category. And each object instance is allowed to be tested only once to avoid any chance of accumulated damage that is falsely attributed to a single method. Initial poses are randomized within the workspace, with the same pose used across methods for each corresponding trial index within a category.

Experiment results. Table III reports the results on 11 object categories, averaged over two evaluators. The Mean column reports the three scores averaged across all 11 categories. Our method achieves the highest category-averaged LIFT, SAFE, and S-LIFT scores (0.67, 0.78, and 0.73), compared with the best baseline scores of 0.57, 0.67, and 0.62, respectively. The gains are not uniform across categories: baselines obtain higher S-LIFT scores on banana, tofu, bread, and grape.

TABLE IV

POLICY RETRY STATISTICS. LIFT SPLITS INTO FIRST-ATTEMPT LIFT SCORE AND RECOVERED CONTRIBUTIONS. THE BASELINES NEVER EXHIBIT RETRY BEHAVIORS.

	LIFT			Failed on first attempt		
Method	first-try	recovered	total	#	retried	recovered
RGB-frozen	0.18	0.00	0.18	45	0	0
RGB-finetune	0.48	0.00	0.48	34	0	0
RGB-scratch	0.57	0.00	0.57	25	0	0
PCD	0.43	0.00	0.43	33	0	0
Ours	0.49	**0.18**	**0.67**	31	**25**	**10**

As shown in Figure 8, our policy exhibits retry behavior, whereas the baselines execute a single descend–grasp–lift sequence. Table IV decomposes LIFT by whether a second attempt was needed. Of the 31 trials whose first attempt did not lift the object, our policy reopened and re-approached in 25 and went on to lift in 10; only 6 ended without a further attempt. This recovery contributes 0.18 of our 0.67 LIFT. Most first-attempt failures occur because the commanded width is slightly loose, and retrying gives the policy opportunity to close further on coming attempts, a behavior we attribute to the diverse reset and disturbance conditions in the simulated demonstrations.

V. CONCLUSION AND DISCUSSION

This work presents a stress-aware framework for learning gentle visuomotor grasping policies. Its core component is a grasp synthesis algorithm that uses simulated internal stress to generate demonstrations for fragile objects at scale. Policies trained on these demonstrations transfer zero-shot to the real world without tactile sensing or real-world policy demonstrations. Simulation experiments demonstrate improved grasp synthesis performance, while real-world evaluation yields higher lifting and safety scores than policies trained on real demonstrations.

Although our demonstration dataset captures diverse categories, extending it to more object categories is a promising next step. As the dataset grows, deliberate sampling or weighting will be needed to prevent underrepresented categories from limiting policy performance [31]. Simulation–real co-training could further improve real-world robustness [32]. Finally, our category-specific, multi-evaluator protocol provides a foundation for assessment of object damage in fragile object manipulation tasks. Vision-language-model evaluators could reduce the associated human effort, although damage outside the available camera views may remain difficult to detect.

REFERENCES

[1] Y. Zhu, D. Yang, and Y. Lee, "Deformable and fragile object manipulation: A review and prospects," *Sensors*, vol. 25, p. 5430, 2025.

[2] D. Blanco-Mulero, Y. Dong, J. Borràs, F. T. Pokorny, and C. Torras, "T-DOM: A taxonomy for robotic manipulation of deformable objects," *arXiv preprint arXiv:2412.20998*, 2024.

[3] Q. Wang, Y. Tu, W. Xu, J. Zhang, A. Knoll, M. Zhou, and Y. Ying, "Towards damage-less robotic fragile fruit grasping: A systematic review on system design, end effector, and visual and tactile feedback," *Journal of Field Robotics*, vol. 42, no. 8, pp. 4521–4543, 2025.

[4] Q. Wang, Y. Zhang, W. Liu, Q. Li, J. Zhang, A. Knoll, M. Zhou, H. Jiang, and Y. Ying, "Toward damage-less robotic fragile fruit grasping: A closed-loop force control method for pneumatic-driven soft gripper," *Soft Robotics*, 2024.

[5] X. Kang, T. Tian, S.-W. Lee, B. Huang, Y. Li, and Y.-L. Kuo, "Learning force-regulated manipulation with a low-cost tactile-force-controlled gripper," *arXiv preprint arXiv:2602.10013*, 2026.

[6] A. Swann, A. Qiu, M. Strong, A. Zhang, S. Morstein, K. Rayle, and M. Kennedy, "Dexfruit: Dexterous manipulation and gaussian splatting inspection of fruit," *IEEE Robotics and Automation Letters*, vol. 11, no. 2, pp. 2034–2041, 2025.

[7] W. Xie, S. Caldararu, and N. Correll, "Just add force for contact-rich robot policies," *arXiv preprint arXiv:2410.13124*, 2024.

[8] Y. Han, K. Yu, R. Batra, N. Boyd, C. Mehta, T. Zhao, Y. She, S. Hutchinson, and Y. Zhao, "Learning generalizable vision-tactile robotic grasping strategy for deformable objects via transformer," *IEEE/ASME Transactions on Mechatronics*, vol. 30, no. 1, pp. 554–566, 2024.

[9] 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*, vol. 44, no. 10–11, pp. 1684–1704, 2025.

[10] F. Lin, Y. Hu, P. Sheng, C. Wen, J. You, and Y. Gao, "Data scaling laws in imitation learning for robotic manipulation," in *International Conference on Learning Representations (ICLR)*, 2025, oral. arXiv:2410.18647.

[11] C. Xu, Y. Xu, M. Xue, J. Lin, J. Wang, X. Zhang, Y. Wang, and J. Yu, "The curse of precision: A data scaling law for high-precision robotic manipulation," *arXiv preprint arXiv:2607.23108*, 2026.

[12] A. Mandlekar, S. Nasiriany, B. Wen, I. Akinola, Y. Narang, L. Fan, Y. Zhu, and D. Fox, "MimicGen: A data generation system for scalable robot learning using human demonstrations," in *Conference on Robot Learning (CoRL)*, 2023, arXiv:2310.17596.

[13] R. Huang, W. Zheng, B. Zhang, J. Zhou, Z. Cui, and Z. Zhang, "Deep learning with tactile sequences enables fruit recognition and force prediction for damage-free grasping," *Computers and Electronics in Agriculture*, vol. 211, p. 107985, 2023.

[14] C. Ferrari and J. Canny, "Planning optimal grasps," in *IEEE International Conference on Robotics and Automation (ICRA)*, 1992, pp. 2290–2295.

[15] Z. Pan, X. Gao, and D. Manocha, "Grasping fragile objects using a stress-minimization metric," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 517–523.

[16] I. Huang, Y. Narang, C. Eppner, B. Sundaralingam, M. Macklin, R. Bajcsy, T. Hermans, and D. Fox, "DefGraspSim: Physics-based simulation of grasp outcomes for 3d deformable objects," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 6274–6281, 2022.

[17] T. N. Le, J. Lundell, F. J. Abu-Dakka, and V. Kyrki, "A novel simulation-based quality metric for evaluating grasps on 3d deformable objects," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 3123–3129.

[18] T. G. W. Lum, K. Kedia, C. K. Liu, and J. Bohg, "Play2Perfect: What matters in dexterous play pretraining for precise assembly?" *arXiv preprint arXiv:2606.26428*, 2026.

[19] R. Singh, A. Allshire, A. Handa, N. Ratliff, and K. Van Wyk, "Dextrah-rgb: Visuomotor policies to grasp anything with dexterous hands," *arXiv preprint arXiv:2412.01791*, 2024.

[20] P. Yin, T. Westenbroek, Z. Zhang, I. Dagnino, E. Shilamkar, N. Mbiziwo-Tiapo, S. Bagaria, X. Liu, G. Mullins, A. Kolobov, *et al.*, "Emergent dexterity via diverse resets and large-scale reinforcement learning," in *International Conference on Learning Representations*, vol. 2026, 2026, pp. 88 343–88 356.

[21] A. H. Barr, "Global and local deformations of solid primitives," *ACM Siggraph Computer Graphics*, vol. 18, no. 3, pp. 21–30, 1984.

[22] Y. Ze, G. Zhang, K. Zhang, C. Hu, M. Wang, and H. Xu, "3D diffusion policy: Generalizable visuomotor policy learning via simple 3D representations," *Robotics: Science and Systems*, 2024.

[23] Y. Dong, Z. Sun, L. Yang, M. Baum, K. Ikemura, S. Song, F. T. Pokorny, and X. Cheng, "Robustness of robotic manipulation: Foundations and frontiers," *arXiv preprint arXiv:2606.31494*, 2026.

[24] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, "Learning fine-grained bimanual manipulation with low-cost hardware," *Robotics: Science and Systems*, 2023.

[25] Y. Li, Z.-X. Zou, Z. Liu, D. Wang, Y. Liang, Z. Yu, X. Liu, Y.-C. Guo, D. Liang, W. Ouyang, *et al.*, "Triposg: High-fidelity 3d shape synthesis using large-scale rectified flow models," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.

[26] V.-D. Nguyen, "Constructing force-closure grasps," *The International Journal of Robotics Research*, vol. 7, no. 3, pp. 3–16, 1988.

[27] A. Ten Pas, M. Gualtieri, K. Saenko, and R. Platt, "Grasp pose detection in point clouds," *The International Journal of Robotics Research*, vol. 36, no. 13-14, pp. 1455–1473, 2017.

[28] G. A. Team, "The role of simulation in scalable robotics, genesis world 1.0, and the path forward," *Genesis AI Blog*, May 2026.

[29] S. Yi, X. Bai, A. Singh, J. Ye, M. T. Tolley, and X. Wang, "Co-design of soft gripper with neural physics," *arXiv preprint arXiv:2505.20404*, 2025.

[30] H. Kress-Gazit, K. Hashimoto, N. Kuppuswamy, P. Shah, P. Horgan, G. Richardson, S. Feng, and B. Burchfiel, "Robot learning as an empirical science: Best practices for policy evaluation," *arXiv preprint arXiv:2409.09491*, 2024.

[31] J. Hejna, C. Bhateja, Y. Jiang, K. Pertsch, and D. Sadigh, "ReMix: Optimizing data mixtures for large scale imitation learning," in *Proceedings of the 8th Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, vol. 270. PMLR, 2025, pp. 145–164.

[32] A. Maddukuri, Z. Jiang, L. Y. Chen, S. Nasiriany, Y. Xie, Y. Fang, W. Huang, Z. Wang, Z. Xu, N. Chernyadev, S. Reed, K. Goldberg, A. Mandlekar, L. Fan, and Y. Zhu, "Sim-and-real co-training: A simple recipe for vision-based robotic manipulation," in *Proceedings of Robotics: Science and Systems*, 2025.