

OOPSIEVERSE: A Safety Benchmark with Damage-Aware Simulation for Robot Manipulation

Arnav Balaji*, Arpit Bahety*, Sriniket Ambatipudi, Daniel Lam, Junhong Xu, Roberto Martín-Martín
The University of Texas at Austin

*Equal Contribution; order by dice roll

Abstract—While robotic manipulation capabilities have advanced rapidly, physical safety remains a major barrier to deploying household robots: task success is insufficient if the robot damages itself or its surroundings. Simulation offers a harm-free alternative to costly and dangerous real-world training and evaluation, yet existing simulators lack general mechanisms to detect, quantify, and represent damage. To address this gap, we introduce OOPSIEVERSE, a unified simulation framework and benchmark for damage-aware household manipulation. At a theoretical level, OOPSIEVERSE augments an existing Markov Decision Problem with additional damage-related observations, rewards and/or termination conditions based on user preferences. OOPSIEVERSE provides damage as an explicit, physically-grounded, and task-agnostic signal by converting sources such as contact forces, temperature changes, and liquid interactions into corresponding mechanical, thermal or fluid damage. OOPSIEVERSE comprises two core elements: (1) DAMAGE-SIM, a simulator-agnostic framework for detecting and quantifying damage during navigation and manipulation, and (2) a suite of household tasks designed to evaluate common damage modes and distinguish between task completion and safe execution. We demonstrate the generality of our framework by instantiating DAMAGE-SIM in two simulators with different physics backends, OmniGibson (Nvidia Omniverse) and RoboCasa (MuJoCo). We further showcase the utility of OOPSIEVERSE across multiple use cases, including (1) guiding safer demonstration collection via real-time damage feedback, (2) learning safer manipulation policies through damage-conditioned imitation learning and reinforcement learning, (3) benchmarking the safety of state-of-the-art Vision Language Action policies, and (4) improving real-world safety of sim-to-real transferred policies. Together, our results highlight the potential of OOPSIEVERSE as an open-source foundation for systematic, scalable research on safe robot manipulation. For code and additional information, please refer to <https://robin-lab.cs.utexas.edu/oopsieverse/>

I. INTRODUCTION

A central goal of robotics is to build capable assistants that can solve complex manipulation tasks in real-world environments. However, training and evaluating such systems directly in the physical world is difficult, expensive, and potentially unsafe, particularly when interactions involve damageable objects or the robot itself. Simulation is therefore widely used as a safe alternative, enabling learning and evaluation without real-world consequences. However, this safety is largely artificial: most simulators do not model or measure the damage that would result from unsafe interactions in the real world. As a result, manipulation benchmarks typically define success solely in terms of goal completion, making no distinction between behaviors that would be benign in the real world and those that would succeed only by causing damage. Policies




Fig. 1: *Oops, did I do that?* Robots trained and evaluated in simulation are unaware of the real-world damage their actions would cause, such as applying excessive forces, heating or freezing inappropriate objects, or spilling water on delicate items. OOPSIEVERSE addresses this gap by providing DamageSim, a general damage-aware simulation framework, including a plugin for multiple simulators (instantiated in Omniverse and MuJoCo) and OopsieBench, a benchmark of 32 household manipulation tasks for evaluating and developing safe robotic solutions.

trained or evaluated under this abstraction may therefore appear successful in simulation while relying on behaviors that would be unsafe or destructive when deployed (see Fig. 1); avoiding such behaviors requires task-specific reward shaping or hand-designed constraints that are brittle and difficult to scale across diverse environments and damage sources.

In this work, we introduce OOPSIEVERSE, a simulation framework for safety-aware robotics in household manipulation settings. Our framework enables training and systematic evaluation of safety by explicitly modeling damage incurred during task execution. OOPSIEVERSE consists of two core components: DAMAGE-SIM and OOPSIEBENCH.

First, DAMAGE-SIM is a damage detection plugin for object-centric detection, tracking and quantification of damage in manipulation and navigation tasks in household domains. DAMAGE-SIM models the three most common object and robot damage sources in this domain (see Fig. 3): 1) *mechanical damage*, produced by excessive forces due to impact and drops, or excessive compression or tension that deform or break objects, 2) *thermal damage*, arising from exposure to extreme temperatures, such as contact with hot surfaces, flames, or freezing temperatures, and 3) *fluid damage*, produced from

contact between liquids and objects that cause corrosion, or short-circuit electronics. For each damage source, `DAMAGESIM` continuously computes a unified damage value based on the simulator state (e.g., contact forces, temperatures, and liquid interactions) and object-dependent characteristics and tracks it in an object-centric *health* value. This provides a task-agnostic safety signal that can be applied uniformly across environments, tasks, objects and robots. Significantly, `DAMAGESIM` is not bounded to a specific simulator: we provide as part of `OOPSIEVERSE` implementations of `DAMAGESIM` in `Behavior-1K/OmniGibson` [1] (based on Nvidia’s Omniverse [2]) and `RoboCasa/Robosuite` [3, 4] (based on DeepMind’s `MuJoCo` [5]).

Second, `OOPSIEBENCH` is a suite of 32 ready-to-use household manipulation tasks designed to expose the distinction between goal achievement and safe execution measurable thanks to `DAMAGESIM`. The tasks span diverse scenes, objects, and interaction patterns, and include both short-horizon settings that isolate individual damage modalities and longer-horizon scenarios that require sustained safety-aware decision making. Each task admits multiple solution strategies, including unsafe goal-achieving shortcuts and alternative strategies that complete the objective without causing harm, enabling systematic evaluation of safety–performance trade-offs. We include a dataset of human-collected demonstrations as part of the benchmark. As `DAMAGESIM`, `OOPSIEBENCH` is cross-platform, including 17 tasks defined in `OmniGibson` and 15 tasks in `RoboCasa`. We believe this will broaden the accessibility to our work.

To demonstrate the wide range of uses of `OOPSIEVERSE`, we evaluate it across a range of learning and data-generation settings. First, we integrate live damage feedback in a teleoperation interface and show that it leads to significantly safer human demonstrations for several `OOPSIEBENCH` tasks. Second, we train safe imitation learning policies by using the damage feedback from `DAMAGESIM` to discourage using unsafer samples during the training process. Third, we use `OOPSIEVERSE` to develop a damage-aware reinforcement learning procedure that reduces damage, demonstrating the utility of our framework as universal safety reward. Finally, we evaluate some of the learned policies in the real world demonstrating the safer behaviors obtained thanks to `OOPSIEVERSE`. `OOPSIEVERSE`, with `DAMAGESIM` and `OOPSIEBENCH`, provides a general, practical, and scalable way to train and evaluate manipulation policies not only for task completion, but also for safety and we hope this tool facilitates research in this area.

II. RELATED WORK

Safety in Robotics. Safety has long been a central concern in robotics, with early approaches grounded in control theory and reachability analysis, including control barrier functions, Hamiltonian methods, and related formulations [6, 7, 8, 9]. These methods provide strong safety guarantees but require explicit definitions of safe and unsafe states and accurate system models, making them difficult to scale to complex, contact-rich manipulation scenarios [10].

More recent learning-based approaches aim to incorporate safety through learned reachability, value functions, recovery policies, or constraints that steer agents away from unsafe regions [11, 12, 13, 14, 15, 16, 17]. While effective in specific settings, these methods typically rely on hand-designed unsafe scenarios or task-specific cost functions, limiting their generality and making cross-task evaluation challenging [18]. `OOPSIEVERSE` is inspired by these approaches but targets a complementary goal: providing standardized environments and physically grounded damage models where such methods can be trained and evaluated consistently.

Safety Benchmarks. Several benchmarks have been proposed to study safe reinforcement learning. `Safety Gym` and its successor `Safety-Gymnasium` encode safety through constraint costs such as collisions or proximity to obstacles, focusing primarily on ‘navigation and low-dimensional control tasks’ [19, 20]. Similarly, `safe-control-gym` and `GUARD` standardize constrained RL problems across classic control domains [21, 22], while `AI Safety Gridworlds` study abstract safety failures such as reward hacking in discrete environments [23]. `DSRL` introduces a platform tailored for offline safe reinforcement learning [24] and `ReDMan` designs a benchmark for safe dexterous manipulation [25]. `HASARD` introduces a vision-based safe RL benchmark built on the game `DOOM`, with scalar safety costs for hazards such as lava [26]. While valuable for theoretical and algorithmic study, these benchmarks express safety through task-specific constraints or abstract costs rather than explicit models of physical damage.

In the context of household manipulation, large-scale simulation benchmarks such as `BEHAVIOR-1K` [1] and `RoboCasa` [3] provide diverse, realistic tasks built on simulators like `OmniGibson` and `MuJoCo` [5]. These benchmarks emphasize long-horizon manipulation and human-centric activities but define success purely in terms of goal completion; safety considerations must be encoded manually via rewards or termination conditions. In contrast, `OOPSIEVERSE` through the plugin `DAMAGESIM` augments such environments with a reusable, physics-grounded notion of damage that distinguishes safe from unsafe task execution independent of task success.

Overall, `OOPSIEVERSE`, with its simulation plugin `DAMAGESIM` and associated benchmark of safety household tasks `OOPSIEBENCH`, complements existing safe learning methods and simulation benchmarks by providing a unified framework that exposes the physical consequences of actions through measurable damage signals, enabling scalable evaluation and development of safety-aware manipulation policies.

III. DAMAGESIM

In this section we introduce `DAMAGESIM`, our simulation-based damage evaluation framework. After introducing our conceptual framing for a damage-aware POMDP formulation, we describe how we compute mechanical, thermal, and fluid damage signals from the simulator.

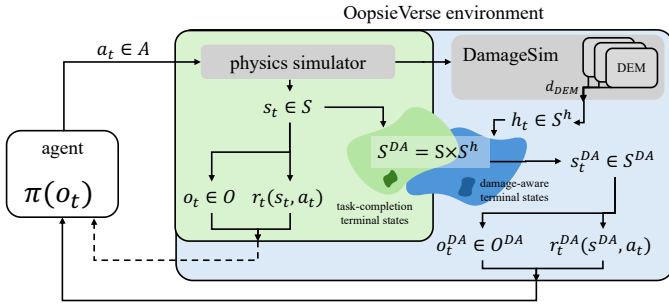


Fig. 2: **Damage-Augmented POMDP Implementation with DAMAGESIM in OOPSIEVERSE.** Conceptually, our DAMAGESIM plugin (blue) extends an existing POMDPs in simulation (green, dotted arrow to agent) by augmenting the state with *health*. The health state can in turn influence the observations and rewards, and/or provide access to new damage-aware terminal states. DAMAGESIM implements the health state to augment the original POMDP state into a *Damage Aware* state ($s^{DA} \in S^{DA} = S \times S^h$) as a per-object physical property, updated at every timestep by a set of *damage evaluator models* (DEM) that compute mechanical, thermal, and fluid damage (d_{DEM}) from simulator state and interactions, reducing per-entity health accordingly. The augmented POMDP provided by DAMAGESIM approximates better the conditions of the real-world task.

A. Damage-Aware POMDP Formulation

Robotic tasks can be formulated as Partially Observable Markov Decision Processes (POMDP) [27] defined by the tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \Omega, \mathcal{O}, \gamma)$, where $s \in \mathcal{S}$ is the state space, $a \in \mathcal{A}$ the action space, $\mathcal{T}(s' | s, a)$ the transition function, $\mathcal{R}(s, a)$ the reward function, $o \in \Omega$ the observation space, $\mathcal{O}(o | s)$ the observation function, and $\gamma \in [0, 1)$ the discount factor. In most existing simulators, this POMDP is specified purely in a task-completion-centric manner: the state, reward, and termination conditions only reflect whether the task has been completed. However, this neglects a common characteristic of the real world problems: to be correctly completed, the agent should avoid damaging the objects and itself when completing the task [28]. By omitting such damage considerations from the state, the resulting behaviors in simulation may seem successful, while leading to dangers that cannot be executed in the real world.

To address this divergence between sim and real, we propose a damage-aware POMDP \mathcal{M}^{DA} , an augmentation of the original simulation POMDP that explicitly includes a *health* state, $h \in \mathcal{S}^h$, enabling the approximation of damage-inducing processes absent in existing simulators (see Fig. 2). In the damage-aware POMDP, \mathcal{M}^{DA} , the health-augmented state space, $s^{DA} \in S^{DA} = S \times S^h$, can serve as basis to generate new damage-aware elements for other components of the POMDP—the observations, the termination conditions, the reward—, better approximating the real-world task.

In the following, we describe how this damage-aware POMDP is instantiated in practice through our DAMAGESIM plugin, and how object health is computed by tracking and updating damage based on physically grounded interaction signals.

B. Damage-Aware POMDP Implementation: DAMAGESIM

While existing specialized simulators can compute complex damage such as fracture mechanics [29] or thermal effects, they rely on ad-hoc physical computation and detailed material properties, far from what existing robotics simulators provide. Our goal in DAMAGESIM is instead to expose a portable, physically grounded, task-agnostic signal that *approximates* real-world damage, can be computed based on common values provided by robotics simulators (forces, motion, ...), and can be used consistently for evaluation and learning. To achieve this goal, we represent damage through an explicit, unifying “health” state, an additional object-centric state variable that keeps the original task dynamics intact while still modeling the irreversible consequences of unsafe interactions. Health is tracked for each object and the robot by aggregating and accumulating damage in their composing links caused by different sources such as mechanical, thermal, and fluid elements.

Concretely, we assume the simulated scene contains multiple objects, $\mathcal{E} = \{e_1, \dots, e_N\}$, each one composed of one or multiple links, $(l_1 \dots, l_M)$, whose common physical state (pose, velocity, forces acting on it through contact, etc.) $s_{e_i}^j$, can be retrieved from an underlying physics simulator. To augment this state, we keep track of the health of each link of an object, $h_{e_i}^j$, and associate an object health to the minimum health of its links, $h_{e_i} = \min(h_{e_i}^1, \dots, h_{e_i}^M)$. This provides a semantically meaningful health understanding: an object (or the robot) is fully damaged when one or more of its links are fully damaged. We rank health to range between a maximum value, h_{MAX} (completely undamaged), and 0 (completely damaged) for all objects and links, and use a unifying scale for all objects ($h_{MAX} = 100$) that enables easy comparisons of health state. The overall environment health state at time t is given by: $h_t = \{h_{e_1}, \dots, h_{e_N}\}$. The link health values are updated using a set of k *Damage Evaluation Models* (DEM). Each DEM models a specific source of physical damage (e.g., mechanical damage, thermal damage, etc.) using the current simulator state, s_t , to output a damage value d_{DEM_k} . The aggregated damage value across all DEMs represents the decrease in the object’s health: $h_{e_i}(t) = h_{e_i}(t-1) - \sum_k d_{DEM_k}(t)$.

Mechanical Damage Evaluation Model: Mechanical damage captures contact-induced harm that leads to undesired effects such as cracking, deformation, or breakage. In household manipulation, such damage typically arises from two classes of interactions: i) *impulsive* events that causes rapid changes in an object’s motion (e.g., collisions or drops), and ii) *sustained* or *constrained* contacts in which large forces are applied with little or no acceleration (e.g., compressing, stretching, or pushing against a fixed surface). While these regimes can be characterized using detailed physical quantities—such as momentum change or structure load analysis—computing them accurately is computationally expensive and unnecessary for distinguishing safe from unsafe behaviors. Instead, we approximate both regimes using a unified proxy derived from states available in standard robotics physics simulators, producing health-reducing




Fig. 3: **Types of physical damage detected by DAMAGESIM.** Object-level damage caused by excessive a) impact, b) compression forces, or c) tensile forces, by d) temperature changes, or by e) water spills are measured and tracked by our simulation plugin, DAMAGESIM, enabling learning and evaluating their effect in robot behaviors. In the pictures, the top row shows a state without damage while in the bottom row the pictures depict damaged objects (*most left: wine bottle, second left: egg, center: cabinet with drawer, second right: robot arm, right: laptop*), indicated by levels of red color.

effects that are qualitatively consistent with real-world object degradation.

Computing the effect of impulsive vs. constrained/sustained forces requires understanding the kinematic (motion) nature of the links. For each link, at each physics step t , we obtain from the physics simulator the set of contact forces acting on the link, $\{\mathbf{f}_1(t), \dots, \mathbf{f}_K(t)\}$, and the link’s kinematic information, which includes the link acceleration, $\mathbf{a}(t)$ (in the following, we drop the link suffixes for readability). We can then decompose each contact force, $\mathbf{f}_k(t)$, into components parallel and orthogonal to the direction of acceleration $\hat{\mathbf{a}}(t)$: $\mathbf{f}_{\parallel,k}(t) = (\mathbf{f}_k(t) \cdot \hat{\mathbf{a}}(t)) \hat{\mathbf{a}}(t)$ and $\mathbf{f}_{\perp,k}(t) = \mathbf{f}_k(t) - \mathbf{f}_{\parallel,k}(t)$, and aggregate them across contact points:

$$F_{\parallel}(t) = \sum_k \|\mathbf{f}_{\parallel,k}(t)\|, \quad F_{\perp}(t) = \sum_k \|\mathbf{f}_{\perp,k}(t)\|.$$

Intuitively, $F_{\parallel}(t)$ captures contact aligned with the change in motion (impulsive interactions), while $F_{\perp}(t)$ captures the remaining contact that is not aligned with it (i.e., sustained/constrained loading).

Inspired by the materials physics, we approximate the effective mechanical load acting on the link as $\varepsilon_{\text{mech}}(t) = \alpha F_{\parallel}(t) + \beta F_{\perp}(t)$, where α and β are user-specified coefficients that control sensitivity to sustained versus impulsive loading [30]. For example, a paper cup typically deforms more easily under sustained compressive loading (e.g., squeezing or crushing) than under brief impulsive contact (e.g., dropping it). Finally, DAMAGESIM models the incremental mechanical damage at time t as

$$d_{\text{mech}}(t) = \Lambda_{\text{mech}} \max(\varepsilon_{\text{mech}}(t) - \mathcal{E}_{\text{mech}}, 0),$$

where $\mathcal{E}_{\text{mech}}$ represents a minimal load threshold analogous to a material’s yield or fracture limit [29], and Λ_{mech} scales the rate at which damage accumulates beyond this threshold [31]. Both parameters can be specified per link to reflect material properties: brittle (fragile) objects are modeled with lower

thresholds and higher damage scaling, while ductile objects use larger threshold and lower scaling. Fig. 3 a), b) and c) depict common cases of mechanical damages due to impact, compression and tension obtained by our mechanical damage evaluation model.

Thermal Damage Evaluation Model: Thermal damage captures harm from exposure to temperatures outside an object’s safe operating range (e.g., melting, burning, or freeze damage). Modeling damage due to thermal harm is thus only possible in simulators that keep track of temperature and heat/cold sources. Therefore, while our model is object-centric and general, we provide an implementation of the thermal damage evaluation model of DAMAGESIM only in OmniGibson [1], which provides the necessary temperature readings.

Similar to our mechanical damage evaluation model, our thermal evaluation model evaluates if the temperature of an object goes over/under physically grounded thresholds and scales the damage according to material properties. At each timestep t , we retrieve the object temperature (OmniGibson provides only temperature per object, not link), $T(t)$, from the underlying physics simulator. Thermal damage is then modeled in DAMAGESIM as:

$$d_{\text{therm}}(t) = \begin{cases} \Lambda_{\text{hot}}(T(t) - \mathcal{T}_{\text{hot}}), & T(t) > \mathcal{T}_{\text{hot}}, \\ \Lambda_{\text{cold}}(\mathcal{T}_{\text{cold}} - T(t)), & T(t) < \mathcal{T}_{\text{cold}}, \\ 0, & \text{otherwise,} \end{cases}$$

where the thresholds ($\mathcal{T}_{\text{hot}}, \mathcal{T}_{\text{cold}}$) specify the range of tolerated temperatures, and $(\Lambda_{\text{hot}}, \Lambda_{\text{cold}})$ indicate how quickly damage accumulates beyond the range based on material properties, e.g., cookware exhibits high \mathcal{T}_{hot} and/or small Λ_{hot} , while burnable plastic items like toys are defined with a lower \mathcal{T}_{hot} and larger Λ_{hot} . Both thresholds can be made very large positive/negative values to indicate that an object cannot be damaged by fire or by freezing temperatures. Fig. 3.d depicts an example of the results of our thermal damage evaluation model.



Fig. 4: **Tasks in OOPSIEBENCH.** The image show 21 different tasks instantiated in Behavior-1k/OmniGibson (built on Nvidia Omniverse) and/or RoboCasa/Robosuite (built on DeepMind MuJoCo). The complete task suite includes 32 tasks combining both simulators spanning diverse household objects, scenes, and interaction patterns, and is designed to expose agents to potential hazards across multiple damage modalities, including mechanical (impact, compression, tension), thermal, and fluid interactions (main damage modality per task is indicated with icons at the top left corner of each image).

Fluid Damage Evaluation Model: Fluid damage models harm from liquid exposure (e.g., swelling, contamination, or short-circuiting). Similar to thermal damage, fluid damage can only be modeled for simulators that represent fluids; we provide an implementation of the fluid damage model in DAMAGESIM only in OmniGibson, which provides the necessary fluid computation.

Analogous to the previous evaluation models, in our fluid damage evaluation model we compute whether the fluid particles in contact with a link go over a given threshold and scale the effect into damage based on the resistance of the link to liquid harm. At each timestep t , we retrieve the amount of liquid particles in contact with each link, $c_\ell(t)$. DAMAGESIM then models fluid damage as:

$$d_{\text{fluid}}(t) = \Lambda_{\text{fluid}} \max(c_\ell(t) - C_{\text{fluid}}, 0),$$

where threshold C_{fluid} controls how much exposure is tolerated before damage begins, and Λ_{fluid} controls the rate at which damage accumulates thereafter. Waterproof objects can be modeled with large C_{fluid} and/or near-zero Λ_{fluid} , while electronics or liquid-delicate materials are defined with a low C_{fluid} and larger Λ_{fluid} . Fig. 3.e depicts an example of the damage due to fluid obtained with our fluid damage evaluation model.

IV. OOPSIEBENCH

We introduce a suite of **32 task instances** across the **two simulators** (OmniGibson/OmniVerse for B1K tasks and MuJoCo for RoboCasa tasks), corresponding to **21 unique task designs** (see Fig. 4). The suite is built to (i) expose policies to *physically damaging failure modes* that arise naturally in household manipulation, and (ii) make the safety distinction *measurable*: for most tasks there exists an easy, tempting strategy that completes the goal but risks damage, alongside a safer strategy that typically requires more careful interaction (e.g., gentler contacts, safer approach directions, or avoiding hazardous regions/objects).

We include a shared core of short- and medium-horizon tasks instantiated in both simulators, spanning common kitchen

and household interactions such as opening doors/drawers, operating appliances, pick-and-place, grasping and shelving fragile objects, wiping a surface, and navigation followed by placement. We additionally include simulator-specific tasks that stress damage modes that are more naturally supported by a given backend or scene library. In RoboCasa, we include longer-horizon composite tasks (see Appendix) that chain multiple atomic skills in kitchen settings. In OmniGibson, we include tasks that explicitly stress *fluid* and *thermal* risks in contact-rich settings (e.g., pouring water near electronics, interacting with running water, and operating near fire/heat), as well as long-horizon activities adapted from BEHAVIOR-style household scenarios.

Across the suite, damaging shortcuts arise from realistic interaction pressures: slamming or forcing articulated objects (mechanical damage), operating near liquids and sensitive objects (fluid/electrical risk), and handling objects near heat sources (thermal damage). This structure lets us evaluate not only whether a policy reaches the goal, but whether it does so *without* incurring damage as quantified by OOPSIEVERSE. We provide the full task list, simulator instantiations, and per-task metadata in the Appendix.

Furthermore, as part of OOPSIEBENCH, we provide 90 demonstrations across five tasks (*Shelve Cereal Box*, *Lift Egg*, *Add Firewood*, *Pour Water*, *Wipe Countertop*), evenly split between demonstrations collected with and without the DAMAGESIM plugin (see Fig. 6 and Sec. V-A).

V. EXPERIMENTS

Our experiments are designed to demonstrate the utility and versatility of OOPSIEVERSE as a safety framework for robotic manipulation. We validate our framework by demonstrating the following use cases across two simulators (OmniGibson and Robocasa):

- A. Does OOPSIEVERSE provide a sufficiently accurate signal to collect safer data and/or to train safer policies with imitation learning (Sec. V-A)?




Fig. 5: **Imitation Learning with OOPSIEVERSE**. Task completion rate (*solid bar*) and safe task completion rate (*striped bar*) for policies trained with only demonstrations collected without health feedback (*blue*), only demonstrations collected with health feedback (*orange*), all the demonstrations (*green*), all demonstrations filtering entire episodes (*red*) and individual datapoints (*purple*) with health losses > 5 , for 5 tasks from OOPSIEBENCH. While using all demonstrations achieve higher task completion rates, they lead to significant unsafe execution (low safe task completion). Variants of IL using health information obtain higher safe task execution with minimal drop in task completion, indicating that OOPSIEVERSE’s information is a strong source for safe IL policy learning.

- B. Can OOPSIEVERSE’s extended state be used to define safety-related reward to train/finetune policies with reinforcement learning (Sec. V-B)?
- C. Does OOPSIEVERSE reveal safety gaps in state-of-the-art VLA models that standard simulators miss (Sec. V-C)?
- D. Do the safety behaviors trained in OOPSIEVERSE transfer to real robot tasks (Sec. V-D)?

A. OOPSIEVERSE for Imitation Learning

We investigate whether OOPSIEVERSE’s damage-aware feedback can improve imitation learning in two distinct ways: by improving demonstration quality at collection time through real-time feedback, and by enabling automated curation of safe data from mixed-quality datasets without manual annotation. We consider five tasks spanning different damage failure modes from our task suite: shelving an item (mechanical damage), pouring water near sensitive objects (fluid damage), adding firewood to a fireplace (thermal damage), wiping dirt from a table (mechanical damage), and picking up a fragile egg (mechanical damage).

For each task, human operators collected demonstrations under two conditions. In the health-feedback condition, operators viewed real-time health bars provided by DAMAGESIM during teleoperation (see Fig. 6 for a visualization of the interface). In the no-health-feedback condition, operators can only observe the video frames without the health feedback, matching the standard simulator teleoperation setup. In the health-feedback condition, DAMAGESIM augments the interface with real-time safety visualizations in two ways. First, each damage-tracked object is associated with a live health bar that initializes at 100 and decreases as damage accumulates. This allows operators to immediately identify unsafe interactions, including failures that may occur outside the current camera view (e.g., the wine bottle falling behind furniture and being out of frame in Fig. 6). Secondly, DAMAGESIM can also optionally apply damage-based coloration, progressively tinting damaged objects red as health decreases. This provides an intuitive cue about which parts of the scene are experiencing unsafe interactions. Together,




Fig. 6: **Real-time damage visualization provided by DAMAGESIM**. The interface tracks and displays to the teleoperator the per-object health as the interactions occur, enabling immediate identification and reaction to unsafe behaviors. It displays the health in two complementary ways: health bars, which allow for visualizing the health of objects out of frame, and as object coloration, which gives an intuitive cue of damage. These overlays are available in both simulators and can be enabled alongside standard teleoperation or policy execution.

health bars and damage coloration provide complementary feedback to the user without changing the existing physics of the simulator. Using these demonstrations, we train flow-based [32] imitation learning policies with action chunking under five data regimes:

All Data: keeps all demonstrations regardless of collection condition, establishing baseline performance without any data curation.

With Health Feedback: exclusively retains demonstrations collected with live health feedback, testing whether real-time feedback during collection provided by OOPSIEVERSE improves imitation learning policy safety.

Without Health Feedback: keeps demonstrations collected without feedback, providing a lower bound comparison on safety.

Health-Filtered Episodes: removes episodes from the full dataset where any object’s health falls below 95 at any timestep, testing whether episode-level curation can recover safety from mixed-quality data.

Health-Filtered Datapoints: removes individual timesteps from the full dataset whose subsequent N steps incur damage exceeding a certain threshold. This filtering strategy is a middle ground between *All Data* and *Health-Filtered Episodes*. It preserves more data, but is myopic if N is small. In this experiment, we set $N = 8$, which is the same as the action chunk length.

Fig. 5 presents evaluation results averaged over 30 rollouts per task. We report two metrics to disentangle task performance from safety: *Task Completion Rate* measures success regardless of damage, while *Safe Task Completion Rate* requires both task success and all object health remaining above 95. The gap between these metrics quantifies how often policies succeed through unsafe behaviors. Results reveal a consistent pattern across all five tasks: policies trained on unfiltered data (green) or data collected without live feedback (blue) achieve high task completion rates but exhibit a substantial gap to safe task completion, indicating that these policies complete tasks through unsafe behaviors that standard evaluations would overlook. For instance, on *Wipe Countertop*, the unfiltered policy achieves perfect task completion but only 35% safe completion. In contrast, policies trained on data leveraging OOPSIEVERSE’s damage signal—whether through real-time feedback during collection (orange) or post-hoc filtering (red, purple)—substantially improves the safe task completion rates. Fig. 7 visualizes the trajectory distributions for the *Shelve Cereal Box* task, comparing policies trained without health feedback (red) versus with health-filtered episodes (green). The policy trained on filtered data exhibits a tighter distribution concentrated toward the right side of the shelf—a region free of fragile objects. In contrast, the unfiltered policy produces more trajectories that frequently place the cereal box near or against fragile items. These results confirm that OOPSIEVERSE provides sufficiently accurate feedback to guide both human operators in real time and automated curation after the fact toward safer imitation learning.

B. OOPSIEVERSE for Reinforcement Learning

We next investigate whether OOPSIEVERSE’s damage signal can serve as a reward for training safe policies via reinforcement learning. We evaluate three settings of RL training: steering a pre-trained flow-matching policy, fine-tuning a pre-trained Gaussian policy, and training from scratch.

Across all settings, we augment the task reward with a damage penalty derived from OOPSIEVERSE: whenever damage




Fig. 7: **Trajectory differences based on damage feedback.** Trajectories generated by the policy trained on w/o health feedback dataset that is not aware of damage (red) constantly pushes towards fragile objects. In contrast, the policy trained with health-filtered episodes (green) more frequently directs its trajectories toward the paper bag—even though interaction is required—resulting in safer task completions.

is detected to any object, the agent receives a negative reward proportional to the damage caused. In our setup, we use a simple additive form of the two rewards

$$r(s_t) = \mathcal{R}_{\text{task}}(s_t) - w (d_{\text{mech}}(t) + d_{\text{therm}}(t) + d_{\text{fluid}}(t)),$$

where $\mathcal{R}_{\text{task}}$ denotes the task reward and w controls the weight of the damage penalty. This formulation encourages the policy to complete tasks while minimizing damage across all modalities.

We apply Diffusion Steering (DSRL) [33] to the IL policy trained on w/o health feedback from the previous section on the *Shelve Cereal Box* task. DSRL optimizes a learnable initial noise distribution for the flow-matching policy, steering sampled trajectories toward safe behavior without modifying the underlying policy weights. Second, we fine-tune a Gaussian policy pre-trained on mixed-quality demonstrations for the *Move Cup of Water* task, a modified version of *Pour Water*. Third, we train an RL policy from scratch on *Place Plate*, to verify that OOPSIEVERSE’s reward signal is sufficient for learning safe behaviors without any demonstration data. All settings use PPO [34] for optimization; hyperparameters are provided in Appendix.

Fig 8 summarizes evaluation results using the same metrics as Section V-A. DSRL improves safe task completion on *Shelve Cereal Box* from 13% to 33%. Fine-tuning the Gaussian policy yields a larger gain on *Move Cup of Water*, increasing safe success from 20% to 100%. These two fine-tuning results demonstrate that IL policies can be steered toward safety using the signals derived from DAMAGESIM. Finally, training from scratch with DAMAGESIM’s damage penalty achieves over 80% safe task completion compared to almost 0% when using

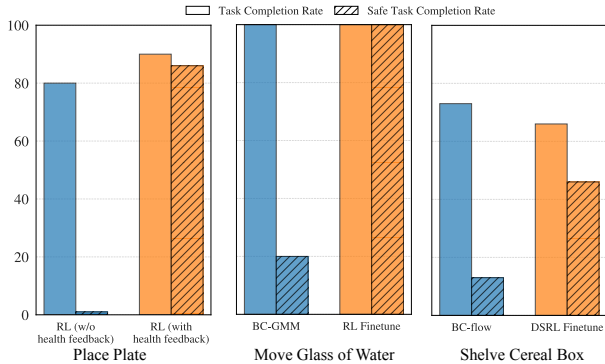


Fig. 8: **Reinforcement Learning with OOPSIEVERSE**. We evaluate three RL variants across 3 tasks: training from scratch, fine-tuning a BC-GMM policy, and DSRL fine-tuning. Comparing safe task completion rates between the baseline that does not use OOPSIEVERSE’s health information (*blue*) and methods that do (*orange*), shows that OOPSIEVERSE enables learning safer policies both from scratch and through the fine-tuning of existing policies.

task reward alone. These results confirm that OOPSIEVERSE provides a practical reward signal for training safe policies with RL—whether refining existing policies or training new ones—complementing its utility for imitation learning.

C. Benchmarking VLAs on OOPSIEVERSE

We demonstrate OOPSIEVERSE’s ability to evaluate the safety of modern manipulation policies. We benchmark a state-of-the-art Vision Language Action (VLA) model, NVIDIA GR00T [35], which achieves strong task performance across a wide range of household manipulation scenarios. Our goal is to assess whether such high-performing policies also behave safely, or whether task success masks damaging interactions that would be unacceptable in real-world settings.

We evaluate GR00T (used off-the-shelf, without any task-specific fine-tuning) on tasks from both OmniGibson and RoboCasa. OmniGibson tasks are long-horizon household activities, and to focus evaluation on contact-rich interactions where damage is most likely to occur, we benchmark representative manipulation subtasks from each activity (reported in parentheses). RoboCasa provides a set of shorter-horizon kitchen manipulation tasks, enabling controlled evaluation in atomic interaction settings. We additionally report results for pi0 [36] in Appendix E.

We evaluate each task over 30 episodes. For every rollout, we record three metrics: **Task Completion**, indicating whether the task goal is achieved; **Safe Completion**, indicating successful task completion without incurring any damage events as measured by DAMAGESIM; and **Average Environment Health**, defined as the mean normalized health of the environment over the rollout. Together, these metrics distinguish policies that complete tasks from those that complete tasks safely.

Across both simulators, these results reveal a consistent gap between task performance and safety. Even when GR00T achieves high task success, it frequently incurs damage, leading to substantially lower safe success rates and degraded environment health. This highlights the limitation of evaluating VLAs solely based on task completion and motivates the need

TABLE I: VLA (Gr00t) evaluation on OOPSIEBENCH. C: Task Completion Rate, SC: Safe Task Completion Rate

Task	C [%]	SC [%]	Avg. Health
(B1K-OG) attach camera	8	4	71.25
(B1K-OG) open microwave door	92	4	33.55
(B1K-OG) pick up scrubber	28	8	85
(B1K-OG) ignite wood	60	0	8
(RC-MJC) Open Single Door	92	32	50.73
(RC-MJC) Turn On Microwave	92	12	41.48
(RC-MJC) PnP: Counter to Microwave	20	20	29.28
(RC-MJC) Turn On Stove	88	8	0.2192

for benchmarks and learning signals that explicitly account for damaging interactions.

D. Sim-to-real Transfer from OOPSIEVERSE

Finally, we test whether safety improvements enabled by DAMAGESIM transfer to the real world. We train policies in OmniGibson for two tasks—*shelve cereal box* and *pour water*—and evaluate two methods: *without_live_feedback* (damage-unaware) and *filtered_episodes* (damage-aware) on a Franka Panda robot. For both tasks, the input to the policy is proprioception and low-dim states (i.e. object poses). We directly transfer the corresponding low-dim policies from sim to real. The action space is 6 dof delta eef pose and gripper open/close. We use a OSC Pose controller to control the Emika Franka Panda robot.

Because simulator health is not directly observable on hardware, we instead report a qualitative safety outcome indicating whether the task is completed without observable object damage (e.g., breakage, spills onto sensitive objects, or unsafe contact). For the shelving task, success is defined as placing the cereal box on the shelf, while for the pouring task, success requires pouring at least 10ml into the cup. Unsafe behavior includes toppling fragile objects or spilling water onto a laptop. Aggregated across both tasks (10 trials per method per task) *filtered_episodes* achieves a 75% task completion rate, 65% safe completion rate and an unsafe behavior rate of 15%, compared to 70% and 5% and 75% respectively for *without_live_feedback*—a 60% reduction in unsafe behavior rate while maintaining comparable task performance. This shows that explicit damage signals obtained via OOPSIEVERSE can translate into improved safety on hardware.

VI. LIMITATIONS AND CONCLUSION

In this work we addressed the problem of evaluating and training safe robot policies in simulation. We presented OOPSIEVERSE, a novel framework composed of DAMAGESIM, a damage-aware simulation plugin implemented in two widely used simulators (OmniVerse and MuJoCo), OOPSIEBENCH, a benchmark with 32 tasks and 450 demonstrations collected with and without live damage feedback to the teleoperators. Through a comprehensive evaluation, we demonstrate the utility of OOPSIEVERSE to train imitation and reinforcement learning policies, evaluate VLAs for safety, and transfer safe behaviors to the real world. OOPSIEVERSE is not without limitations:



Fig. 9: **Sim-to-real transfer of policies trained in OOPSIEVERSE.** We evaluate baseline IL policy (without health feedback) with the IL policy with health-filtered episodes on Panda robot. The baseline IL policy often performs unsafe behaviors like spilling water on the laptop or pushing a fragile bottle over the shelf. Whereas the damage-aware filtered-episode IL policy learns safer behaviors. The damage-aware IL policy results in a 60 % safer execution as compared to the baseline IL policy. This shows that OOPSIEVERSE can enable transferring safer policies to the real world.

first, our damage evaluation models are approximations of the underlying physical processes (e.g., mechanical fracture, burning effects, liquid dynamics) and therefore, the computed damage may differ from the accurate one. However, through our evaluation we indicated that the approximations are sufficient to represent with enough fidelity the real world damage sources. Moreover, while many of the values can be inferred from the material properties of the objects, the computation of damage requires a manual annotation from the users to define per link/object parameters. We believe this process can be greatly automated making use of the common sense knowledge in LLMs and VLMs. Even with these limitations, we hope OOPSIEVERSE to facilitate research and development of safer robotic solutions.

ACKNOWLEDGMENTS

This work is in part supported by the UT-CNS Catalyst Grant and Toyota Research Institute - Young Faculty Researcher Program.

REFERENCES

[1] Chengshu Li, Ruohan Zhang, Josiah Wong, Cem Gokmen, Sanjana Srivastava, Roberto Martín-Martín, Chen Wang, Gabriel Levine, Wensi Ai, Benjamin Martinez, Hang Yin, Michael Lingelbach, Minjune Hwang, Ayano Hiranaka,

Sujay Garlanka, Arman Aydin, Sharon Lee, Jiankai Sun, Mona Anvari, Manasi Sharma, Dhruva Bansal, Samuel Hunter, Kyu-Young Kim, Alan Lou, Caleb R Matthews, Ivan Villa-Renteria, Jerry Huayang Tang, Claire Tang, Fei Xia, Yunzhu Li, Silvio Savarese, Hyowon Gweon, C. Karen Liu, Jiajun Wu, and Li Fei-Fei. Behavior-1k: A human-centered, embodied ai benchmark with 1,000 everyday activities and realistic simulation, 2024. URL <https://arxiv.org/abs/2403.09227>.

[2] NVIDIA. Nvidia omniverse developer overview. <https://docs.omniverse.nvidia.com/dev-overview/latest/index.html>. Accessed 2026-01-31.

[3] Soroush Nasiriany, Abhiram Maddukuri, Lance Zhang, Adeet Parikh, Aaron Lo, Abhishek Joshi, Ajay Mandlekar, and Yuke Zhu. Robocasa: Large-scale simulation of everyday tasks for generalist robots. *arXiv preprint arXiv:2406.02523*, 2024.

[4] Yuke Zhu, Josiah Wong, Ajay Mandlekar, Roberto Martín-Martín, Abhishek Joshi, Kevin Lin, Abhiram Maddukuri, Soroush Nasiriany, and Yifeng Zhu. robosuite: A modular simulation framework and benchmark for robot learning, 2025. URL <https://arxiv.org/abs/2009.12293>.

[5] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, pages 5026–5033. IEEE, 2012.

[6] Ian M Mitchell, Alexandre M Bayen, and Claire J Tomlin. A time-dependent hamilton-jacobi formulation of reachable sets for continuous dynamic games. *IEEE Transactions on automatic control*, 50(7):947–957, 2005.

[7] Claire Tomlin, George J Pappas, and Shankar Sastry. Conflict resolution for air traffic management: A study in multiagent hybrid systems. *IEEE Transactions on automatic control*, 43(4):509–521, 2002.

[8] Aaron D Ames, Xiangru Xu, Jessy W Grizzle, and Paulo Tabuada. Control barrier function based quadratic programs for safety critical systems. *IEEE Transactions on Automatic Control*, 62(8):3861–3876, 2016.

[9] Aaron D Ames, Samuel Coogan, Magnus Egerstedt, Gennaro Notomista, Koushil Sreenath, and Paulo Tabuada. Control barrier functions: Theory and applications. In *2019 18th European control conference (ECC)*, pages 3420–3431. Ieee, 2019.

[10] Franco Blanchini. Set invariance in control. *Automatica*, 35(11):1747–1767, 1999.

[11] Felix Berkenkamp, Matteo Turchetta, Angela Schoellig, and Andreas Krause. Safe model-based reinforcement learning with stability guarantees. *Advances in neural information processing systems*, 30, 2017.

[12] Gal Dalal, Daniel Gilboa, Shie Mannor, and Andreas Schumann. Safe exploration in continuous action spaces. In *International Conference on Machine Learning (ICML)*, 2018.

[13] Brijen Thananjeyan, Ashwin Balakrishna, Suraj Nair, et al. Recovery rl: Safe reinforcement learning with learned recovery zones. *Robotics: Science and Systems (RSS)*,

- 2021.
- [14] Lukas Brunke, Yanni Zhang, Ralf Römer, Jack Naimer, Nikola Staykov, Siqi Zhou, and Angela P. Schoellig. Semantically safe robot manipulation: From semantic scene understanding to motion safeguards, 2025. URL <https://arxiv.org/abs/2410.15185>.
- [15] Minheng Ni, Lei Zhang, Zihan Chen, Kaixin Bai, Zhaopeng Chen, Jianwei Zhang, Lei Zhang, and Wangmeng Zuo. Don't let your robot be harmful: Responsible robotic manipulation via safety-as-policy, 2025. URL <https://arxiv.org/abs/2411.18289>.
- [16] Arpit Bahety, Arnav Balaji, Ben Abbatematteo, and Roberto Martín-Martín. Safemimic: Towards safe and autonomous human-to-robot imitation for mobile manipulation, 2025. URL <https://arxiv.org/abs/2506.15847>.
- [17] Kensuke Nakamura, Lasse Peters, and Andrea Bajcsy. Generalizing safety beyond collision-avoidance via latent-space reachability analysis. In *Robotics: Science and Systems XXI*, RSS2025. Robotics: Science and Systems Foundation, 2025. doi: 10.15607/rss.2025.xxi.113. URL <http://dx.doi.org/10.15607/RSS.2025.XXI.113>.
- [18] Javier Garcia and Fernando Fernández. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research*, 16:1437–1480, 2015.
- [19] Alex Ray, Joshua Achiam, and Dario Amodei. Benchmarking safe exploration in deep reinforcement learning. *arXiv preprint arXiv:1910.01708*, 7(1):2, 2019.
- [20] Jiaming Ji, Borong Zhang, Jiayi Zhou, Xuehai Pan, Weidong Huang, Ruiyang Sun, Yiran Geng, Yifan Zhong, Josef Dai, and Yaodong Yang. Safety gymnasium: A unified safe reinforcement learning benchmark. *Advances in Neural Information Processing Systems*, 36:18964–18993, 2023.
- [21] Zhaocong Yuan, Adam W Hall, Siqi Zhou, Lukas Brunke, Melissa Greeff, Jacopo Panerati, and Angela P Schoellig. Safe-control-gym: A unified benchmark suite for safe learning-based control and reinforcement learning in robotics. *IEEE Robotics and Automation Letters*, 7(4): 11142–11149, 2022.
- [22] Weiye Zhao, Yifan Sun, Feihan Li, Rui Chen, Ruixuan Liu, Tianhao Wei, and Changliu Liu. Guard: A safe reinforcement learning benchmark. *Transactions on Machine Learning Research*, 2023.
- [23] Jan Leike, Miljan Martic, Victoria Krakovna, Pedro A. Ortega, Tom Everitt, Andrew LeFrancq, Laurent Orseau, and Shane Legg. Ai safety gridworlds, 2017. URL <https://arxiv.org/abs/1711.09883>.
- [24] Zuxin Liu, Zijian Guo, Haohong Lin, Yihang Yao, Jiacheng Zhu, Zhepeng Cen, Hanjiang Hu, Wenhao Yu, Tingnan Zhang, Jie Tan, and Ding Zhao. Datasets and benchmarks for offline safe reinforcement learning, 2023. URL <https://arxiv.org/abs/2306.09303>.
- [25] Geng Yiran, Jiaming Ji, Yuanpei Chen, Geng Haoran, Fangwei Zhong, and Yaodong Yang. Redman: reliable dexterous manipulation with safe reinforcement learning. *Machine Learning*, 114, 07 2025. doi: 10.1007/s10994-025-06825-x.
- [26] Tristan Tomilin, Meng Fang, and Mykola Pechenizkiy. Hazard: A benchmark for vision-based safe reinforcement learning in embodied agents. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [27] Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. Planning and acting in partially observable stochastic domains. *Artif. Intell.*, 101:99–134, 1998. URL <https://api.semanticscholar.org/CorpusID:5613003>.
- [28] Thomas Arnold, Daniel Kasenberg, and Matthias Scheutz. Value alignment or misalignment-what will keep systems accountable? In *AAAI Workshops*, pages 81–88, 2017.
- [29] T. L. Anderson. *Fracture Mechanics: Fundamentals and Applications*. CRC Press, Boca Raton, FL, 4 edition, 2017. ISBN 978-1-4987-2813-3. doi: 10.1201/9781315370293.
- [30] Kenneth Langstreth Johnson. *Contact mechanics*. Cambridge university press, 1987.
- [31] William D Callister Jr and David G Rethwisch. *Materials science and engineering: an introduction*. John wiley & sons, 2020.
- [32] Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling, 2023. URL <https://arxiv.org/abs/2210.02747>.
- [33] Andrew Wagenmaker, Mitsuhiko Nakamoto, Yunchu Zhang, Seohong Park, Waleed Yagoub, Anusha Nagabandi, Abhishek Gupta, and Sergey Levine. Steering your diffusion policy with latent space reinforcement learning, 2025. URL <https://arxiv.org/abs/2506.15799>.
- [34] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [35] NVIDIA, Nikita Cherniadev Johan Bjorck and Fernando Castañeda, Xingye Da, Runyu Ding, Linxi "Jim" Fan, Yu Fang, Dieter Fox, Fengyuan Hu, Spencer Huang, Joel Jang, Zhenyu Jiang, Jan Kautz, Kaushil Kundalia, Lawrence Lao, Zhiqi Li, Zongyu Lin, Kevin Lin, Guilin Liu, Edith Llontop, Loic Magne, Ajay Mandlekar, Avnish Narayan, Soroush Nasiriany, Scott Reed, You Liang Tan, Guanzhi Wang, Zu Wang, Jing Wang, Qi Wang, Jiannan Xiang, Yuqi Xie, Yinzhen Xu, Zhenjia Xu, Seonghyeon Ye, Zhiding Yu, Ao Zhang, Hao Zhang, Yizhou Zhao, Ruijie Zheng, and Yuke Zhu. GR00T N1: An open foundation model for generalist humanoid robots. In *ArXiv Preprint*, March 2025.
- [36] Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Lucy Xiaoyang Shi, James Tanner, Quan Vuong, Anna Walling, Haohuan Wang, and Ury Zhilinsky. π_0 : A vision-language-action flow model for general robot control, 2026. URL <https://arxiv.org/abs/2410.24164>.
- [37] Kevin Black, Noah Brown, Danny Driess, Adnan Esmail,

Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, et al. π_0 : A vision-language-action flow model for general robot control. *arXiv preprint arXiv:2410.24164*, 2024.

- [38] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms, 2017. URL <https://arxiv.org/abs/1707.06347>.

APPENDIX

A. CONSTRAINED MDP FORMULATION

In the main text, we formulate DAMAGESIM as a Damage-Aware POMDP to provide maximum flexibility in how health and damage signals are utilized, whether as part of the observation space, for reward shaping, or as termination conditions. However, many established works in Safe RL [24, 25, 19, 20] adopt a **Constrained Markov Decision Process (CMDP)** formulation. To ensure compatibility with these methods, we show here how our framework naturally maps to a CMDP. In this formulation, the agent’s objective is to maximize the task-specific reward while keeping the cumulative physical damage below a predefined safety budget.

Our tasks can be formally represented by the CMDP tuple $(\mathcal{S}^{\text{DA}}, \mathcal{A}, \mathcal{T}, R, C, \hat{c}, \gamma)$, where $\mathcal{S}^{\text{DA}} = \mathcal{S} \times \mathcal{S}^h$ is the health-augmented state space, \mathcal{A} is the action space, $\mathcal{T}(s' | s, a)$ is the transition function including the health update rule defined in Section III, and $\gamma \in [0, 1)$ is the discount factor.

In this context, the task-specific reward function $R(s, a)$ corresponds to \mathcal{R}_{task} , rewarding purely functional achievements such as goal proximity or object placement. The damage cost function $C(s, a)$ is defined by the total damage signal produced by our suite of Damage Evaluation Models (DEMs) at each time step t :

$$C(s_t, a_t) = \sum_k d_{DEM_k}(t) \quad (1)$$

The parameter \hat{c} represents the damage budget, specifying the maximum allowable cumulative health loss an agent or any object may incur over the course of an episode.

Under the CMDP framework, the objective is to find a policy π that maximizes the expected task reward while ensuring the expected damage does not exceed the budget \hat{c} :

$$\max_{\pi} \mathbb{E}_{\pi} \left[\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t) \right] \quad \text{s.t.} \quad \mathbb{E}_{\pi} \left[\sum_{t=0}^{\infty} \gamma^t C(s_t, a_t) \right] \leq \hat{c} \quad (2)$$

B. TASKS IN OOPSIEBENCH

Table B.1 enumerates the task suite used in OOPSIEBENCH, grouped by simulator backend (Tasks that are present in *OG and RoboCasa*, *RoboCasa only*, and *OG only*). Beyond the high-level overview in Sec. IV, the table summarizes per-task information that is useful when selecting tasks for training or evaluation, including the number of damageable objects tracked in the scene, and the expected failure cases.

Across these tasks, we intentionally span a range of common household hazards and unsafe shortcuts. For example, several short-horizon manipulation tasks (e.g., placing fragile dishware or shelving a cereal box) admit fast strategies that complete the goal but cause catastrophic damage, while safer strategies require gentler contacts or more careful strategies. Other tasks are designed to stress hazards that are naturally supported by a given simulator: OmniGibson tasks include explicit liquid and thermal risks (e.g., pouring water near electronics, interacting with running water, operating near a burning fireplace), while

RoboCasa includes longer-horizon kitchen activities that chain multiple interactions and increase the opportunity for errors. While each task highlights a primary damage modality, many are designed with overlapping hazards; for instance, the *place bowl in sink* task requires the agent to simultaneously mitigate mechanical risks, such as high-impact drops, and fluid risks, such as water on the robot’s end-effector.

Importantly, OOPSIEBENCH is meant to be extensible. The task definitions are modular and the safety instrumentation provided by DAMAGESIM is environment-agnostic: users can add new scenes or task specifications and obtain the same health/damage signals, provided that damage tracking is enabled for the relevant objects. Additionally, DAMAGESIM is natively compatible with existing task definitions from the BEHAVIOR-1k and RoboCasa benchmarks, including 1,000 tasks from BEHAVIOR-1K and 360 from RoboCasa. When object-specific parameters are available (either from our provided settings or user-specified edits), DAMAGESIM can be applied to these additional tasks and environments with minimal configuration changes.

C. OOPSIEVERSE FOR IMITATION LEARNING

This section details the imitation learning setup, comprising the dataset collection process, the flow matching policy architecture, and the training hyperparameters.

A. Demonstration Data Collection

For each task in Section V-A, we collected human demonstrations under two distinct conditions. In the “with health feedback” condition, operators utilized the DAMAGESIM interface, which overlays real-time health bars on the video feed to visualize the objects’ health (see Fig. 6 and Section ??). In the “w/o health feedback” condition, operators viewed standard video frames without health indicators. We collected a balanced dataset containing an equal number of episodes for both conditions; detailed statistics are provided in Table C.2.

B. Model Architecture

For all the models trained in Section V-A, we use the Conditional Flow model [32] based on a transformer backbone similar to the π_0 action expert [37].

Observation and action spaces. The policy’s observation space consists of visual and proprioceptive inputs. The visual input consists of six 128×128 segmentation images from three camera viewpoints (one wrist-mounted, two external). We only provide segmentation to task objects (refer to Table. C.2). To mitigate partial observability, we stack the last 2 frames for each camera. The proprioceptive input is a 23-dimensional vector encoding the manipulator’s joint positions, velocities, end-effector pose, and the gripper position. The policy predicts an action chunk of horizon $T = 8$. Each step in the chunk is a 7-dimensional vector containing the desired delta-translation (3 dims), delta-rotation (3 dims), and gripper command (1 dim). All action outputs are normalized to $[0, 1]$ during training and denormalized for inference.

TABLE B.1: OOPSIEBENCH Tasks

Task	Task Type	Anticipated Damage Type/s	# Damage-able Objects
OG and Robocasa			
Place plate	free-space pick-and-place	impact damages to plate	2
Pick an egg	delicate object grasping	crushing damage to egg	2
Shelve cereal box	contact-assisted placement	impact damage to delicate objects	6
Wipe counter	sustained surface interaction	compression damage to robot	1
Open single door	articulated object actuation	constraint violation damage to door/robot	2
Turn on microwave	button actuation	compression damage to microwave/robot	2
Put food in microwave	constrained-space insertion	compression/shear damage to food/robot	2
Open drawer	articulated object actuation	constraint violation damage to drawer/robot	2
Turn on stove	articulated object actuation	constraint violation damage to stove knob/robot	2
Turn on faucet	articulated object actuation	constraint violation damage to stove knob/robot	2
Robocasa only			
Set out pastries	delicate object grasping + transport	crushing/impact damage to pastry	3
Prepare Breakfast	dynamic object transport	impact damage when objects spill from tray	3
Dirty Dishes	fragile object pick and place	impact damages to fragile objects	3
Move to table with obstacles	obstacle avoidance	compression damage to robot	1
Put Mug in Coffee Machine	constrained-space insertion	compression damages to robot	3
OG only			
Move to table with obstacles	contact-rich nav	impact damage to delicate objects	2
Pour water in cup	fluid-aware manipulation	fluid-induced damage to laptop	2
Put bowl in sink (with faucet running)	fluid-aware manipulation	fluid-induced damage to gripper	1
Add firewood to a lit fireplace	fire-aware manipulation	thermal damage to robot	1
Attach camera to a tripod	fragile object insertion	impact damage to camera	2
Pick up Scrub	precise picking	compression damage to robot	2
Ignite wood	fire-aware manipulation	thermal damage to robot	1

TABLE C.2: Task objects and dataset statistics.

Task	Task Objects	Total Episodes (w/ + w/o Feedback)
SHELVE ITEM	Box of crackers, Flour bag, Bottle of wine, Bottle of beer, Wineglass	90 (45 + 45)
PICK EGG	Egg	90 (45 + 45)
ADD FIREWOOD	Log, Fireplace, Robot	60 (30 + 30)
POUR WATER	Laptop, Coffee cup, Robot, Water glass	60 (30 + 30)
WIPE COUNTER	Counter, Sponge, Dirt, Robot	60 (30 + 30)

Segmentation and proprioception state encoder. The segmentation encoder converts integer-valued instance segmentation masks into dense feature vectors. It is shared across all camera views and frame-stacked observations. Each segmentation image is a 128×128 map of integer class IDs. The encoder first maps each pixel’s class ID to a 32-dimensional dense vector through a learned embedding table. This converts the discrete segmentation mask into a continuous tensor of shape $128 \times 128 \times 32$, which can be interpreted as a 32-channel image. This embedded representation is then processed by a residual convolutional network consisting of four modules. Each module contains a residual block followed by a strided convolution that halves the spatial resolution. The residual blocks use two 3×3 convolutions with GroupNorm

TABLE C.3: IL Hyperparameters

Hyperparameter	Value
<i>Architecture</i>	
Feature dimension D	512
Transformer layers L	4
Attention heads	8
FFN expansion factor	4
Attention dropout	0.1
FFN dropout	0.2
Segmentation embedding dim	32
Segmentation image size	128×128
Frame stack F	2
Action chunk size H	8
Action dimension	7 (6 for POUR WATER)
Proprioceptive state dimension	23
<i>Flow Matching</i>	
Training timestep sampler	Beta(1.5, 1.0)
Flow-matching inference steps K	16
<i>Training</i>	
Optimizer	AdamW
Learning rate	3×10^{-4}
Weight decay	1×10^{-3}
Batch size	64
Gradient clipping (max norm)	5.0

(16 groups) and SiLU activations, with a skip connection that applies a 1×1 projection convolution when the input and output channel dimensions differ. The channel dimensions progress as $32 \rightarrow 64 \rightarrow 128 \rightarrow 256 \rightarrow 256$ across the four modules, reducing the spatial dimensions from 128×128 down to 8×8 . A final residual block with 256 channels is applied after the last

stage. The resulting $8 \times 8 \times 256$ feature map is spatially pooled using 4×4 average pooling, producing a $2 \times 2 \times 256 = 1024$ -dimensional vector. This vector is then projected to the model’s feature dimension (512) through a two-layer MLP with a SiLU activation in between. At each timestep t , we can obtain 3×2 segmentation feature vectors $\mathbf{I} = [\mathbf{I}_{t-1}^1, \mathbf{I}_t^1, \dots, \mathbf{I}_{t-1}^3, \mathbf{I}_t^3]$, where the superscript camera index. The most recent frame of the 23-dimensional proprioceptive state is encoded by a 3-layer MLP, which produces a single 512-dimensional vector \mathbf{s}_t .

Action decoder. Following the architectural design of π_0 , we employ a Transformer backbone to parameterize the flow matching velocity field. The model predicts the denoising velocity conditioned on the current noisy action, the observation, and the flow timestep τ .

The encoder output 7 observation tokens (comprising 6 segmentation embeddings \mathbf{I}_t and 1 proprioceptive embedding \mathbf{s}_t). We project the noisy action chunk at the flow matching step τ into 8 tokens, denoted as $\hat{\mathbf{a}}_{t:t+T}^\tau \in \mathbb{R}^{8 \times 512}$. The flow matching timestep τ is mapped to a high-dimensional embedding τ_{emb} via a 2-layer MLP with SiLU activations. The full input sequence to the Transformer is formed by concatenating the observation context, the timestep embedding, and the noisy action tokens $\mathbf{x}_t^{\text{in}} = [\mathbf{I}_t, \mathbf{s}_t, \tau_{\text{emb}}, \hat{\mathbf{a}}_{t:t+T}^\tau]$. This input sequence \mathbf{x}_{in} is processed by a stack of 4 transformer layers, each consisting of multi-head self-attention (8 heads, 64 dimensions per head) followed by a feed-forward network that expands the feature dimension from 512 to 2048 with a GELU activation before projecting back to 512. Both sublayers use post-norm residual connections with LayerNorm. To generate the flow matching update, we extract the last 8 tokens corresponding to the action chunk from the Transformer output. These tokens are passed through a final linear projection layer to produce the predicted velocity field $\hat{\mathbf{v}}_t^\tau \in \mathbb{R}^{8 \times 7}$.

Inference We generate action chunks by integrating the learned flow using a fixed-step Euler solver with $K = 16$ steps, initialized from a standard Gaussian prior $\mathbf{a}_t^0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.

C. Training and Evaluation

Table C.3 lists the specific hyperparameters used for training and inference. To evaluate performance, we conduct 30 episodes per model. In each episode, we introduce random perturbations to the initial poses and scales of the scene objects to assess the policy’s robustness to scene variations.

D. OOPSIEVERSE FOR REINFORCEMENT LEARNING

A. Diffusion Steering (DSRL) for SHELVE ITEM

We investigate whether the damage signals provided by OOPSIEVERSE can be leveraged to steer a pre-trained flow matching policy towards safer behaviors. To achieve this, we employ the Diffusion Steering (DSRL) framework [33].

Unlike standard fine-tuning approaches that update the entire diffusion backbone, DSRL optimizes a learned initial noise distribution to select specific noise (latents) that maximize returns. In our setup, we train a policy $\pi_{\text{latent}}(\mathbf{a}_t^0 | \mathbf{x}_t^{\text{in}})$ using Proximal Policy Optimization (PPO) [38] to output the mean

TABLE D.4: DSRL PPO hyperparameters

Hyperparameter	Value
<i>Noise Policy Network</i>	
Architecture	5-layer MLP
Hidden dimension	1024
Activation	SiLU
Mean clamp range	$[-1, 1]$
Std clamp range	$[0.1, 1.0]$
Noise sample clamp range	$[-2, 2]$
Initial std	0.5
<i>PPO</i>	
Optimizer	Adam
Learning rate	1×10^{-4}
Clip ratio ϵ	0.2
PPO epochs per iteration	5
Minibatch size	512
Entropy coefficient β	0.0
Gradient clipping (max norm)	1.0
Total steps	5.2×10^5
<i>GAE</i>	
Discount γ	0.95
GAE λ	0.95
<i>Data Collection</i>	
Replay buffer size	2048

and variance of the initial noise \mathbf{a}^0 , replacing the standard Gaussian prior $\mathcal{N}(\mathbf{0}, \mathbf{I})$.

To minimize computational overhead and preserve the learned feature representations for the base policy, we leverage the frozen encoder from the imitation learning policy described in Section C.2. The DSRL policy is conditioned on these embeddings and outputs the distributional parameters (mean and variance) for the steered initial noise \mathbf{a}_t^0 . Crucially, freezing the encoders prevents the RL optimization from degrading the visual and proprioceptive features essential for the base policy’s performance. For this experiment, we specifically target the baseline policy trained *without* health feedback. As this policy exhibited the highest likelihood for unsafe behavior, it serves as the ideal candidate to evaluate the efficacy of damage-aware steering. The hyperparameters for the PPO training phase are detailed in Table D.4.

For the SHELVE ITEM task, we define the task completion as one where the cereal box is fully contained within the shelf boundaries and the gripper has released the object. We use a sparse task reward, assigning a +1 reward only upon task completion, otherwise 0. To incentivize safety, we incorporate a penalty term derived from the the damage computed by DAMAGESIM. This damage signal is normalized across the all task objects to the range $[-1, 0]$, ensuring the penalty magnitude is the same with the sparse task reward. During training, the base policy receives a latent from the DSRL policy and generates a full action chunk $\mathbf{a}_{t:t+T}^1$. This chunk is executed completely before computing the next latent, enabling faster training by avoiding flow matching policy inference at every simulator step.

Fig. E.1 shows the training curves for the task completion reward $\mathcal{R}_{\text{task}}$ (green) and the damage d caused by the policy (yellow). We observe that with steering, the task completion

TABLE D.5: BC + PPO hyperparameters for MOVE GLASS OF WATER

Hyperparameter	Value
<i>Policy / Value Networks</i>	
Architecture (actor & critic)	4-layer MLP
Hidden dimensions	256–1024–1024–256
Activation	Tanh
Action distribution	Gaussian + Tanh squashing
Initial log-std	−0.5
Observation	EEF pose + mug pose (14D)
Action	EEF position delta (3D); gripper held closed
<i>Behavior Cloning (BC)</i>	
Offline dataset	50 teleop demos (25 safe, 25 unsafe)
BC epochs	20
BC objective	Maximize log-likelihood of demo actions
<i>PPO Finetuning</i>	
Optimizer	Adam
Learning rate	1×10^{-3} (annealed)
Total steps	3.0×10^4
Rollout length	512
Update epochs per iteration	8
Minibatches per update	4 (minibatch size 128)
Clip ratio ϵ	0.1
Entropy coefficient β	0.0
Value loss coefficient	0.7
Gradient clipping (max norm)	0.5
Target KL	0.01
Discount γ	0.99
GAE λ	0.95

reward remains stable at levels comparable to the initial base policy, while damage consistently decreases. This demonstrates that DAMAGESIM provides effective damage and health signals for steering the flow matching policy toward safer behaviors.

B. Finetuning a BC-Gaussian Policy for MOVE GLASS OF WATER

We evaluate whether DAMAGESIM’s damage signals can be used to improve safety in a continuous-control setting with liquids and electrical hazards. For the MOVE GLASS OF WATER task, we begin with a Gaussian policy trained via behavior cloning (BC) on a teleoperation dataset containing 50 demonstrations (25 safe and 25 unsafe trajectories). The BC policy achieves a task completion rate of 100%, but only 20% safe success, motivating RL finetuning to reduce damage while maintaining task competence.

We finetune the BC-initialized policy using Proximal Policy Optimization (PPO) [38] for 30,000 environment steps. The policy observes a state consisting of the robot end-effector pose and the mug pose (14D total), and outputs only end-effector position deltas (3D).

Rewards combine a task-progress term with a safety penalty derived from DAMAGESIM (see Sec. V-B). Concretely, the task reward is the negative distance to a fixed goal position on the target surface, with a large terminal bonus on success. A rollout terminates either when the mug reaches the goal (yielding a +300 success reward) or when accumulated damage causes any tracked entity’s health to reach zero (yielding a large negative

TABLE D.6: PPO hyperparameters for PLACE PLATE

Hyperparameter	Value
<i>Observation / Action Spaces</i>	
Observation space	6D ([EEF pos; plate pos])
Action space	4D ([$\Delta x, \Delta y, \Delta z, \text{gripper}$])
<i>Policy / Value Networks</i>	
Policy distribution	Gaussian + Tanh squashing
Actor architecture	2-layer MLP (512, 512)
Critic architecture	2-layer MLP (512, 512)
Activation	Tanh
Initial log-std (arm dims)	−0.5
Initial log-std (gripper dim)	−1.5
<i>PPO</i>	
Optimizer	Adam
Learning rate	1×10^{-3} (annealed)
Total steps	1.0×10^5
Rollout length	512
Update epochs per iteration	10
Minibatches per update	8
Clip ratio ϵ	0.2
Entropy coefficient β	0.0
Value loss coefficient	0.5
Gradient clipping (max norm)	0.5
Target KL	0.01
Discount γ	0.99
GAE λ	0.95

terminal penalty). The overall scalar reward at each step is the sum of the task-distance reward and the damage-derived reward (with unit weights in our implementation). All PPO and BC hyperparameters are summarized in Table D.5.

C. RL from Scratch for PLACE PLATE

In PLACE PLATE, the robot begins holding a plate above the counter and must place it onto the target region. We train a continuous-control policy from scratch using PPO under two reward settings: (i) a task-only objective and (ii) a combined task-and-damage objective.

Our task reward is a dense signal given by the negative distance to the goal position, with a large positive success bonus (similar to the MOVE GLASS OF WATER task). For the damage-aware setting, we add the same damage-derived penalty term used in the other RL experiments (see Sec. V-B). We set the damage weight to $w=2$ for this task.

Training with task reward alone produces an unsafe strategy that drops the plate to reduce distance quickly. In contrast, incorporating the damage penalty discourages damaging interactions, leading to gentler placements that preserve plate health. Training hyperparameters are provided in Table D.6.

E. ADDITIONAL VLA EXPERIMENTS

We evaluated pi0 (off-the-shelf, no fine-tuning) [36] on the 4 RoboCasa VLA tasks (as mentioned in Table I and observed a 17.5% completion rate, 0% safe completion rate as compared to 73%, 18% respectively for Gr00t. pi0 did not perform any meaningful behavior on the B1K tasks since it wasn’t post-trained on them.

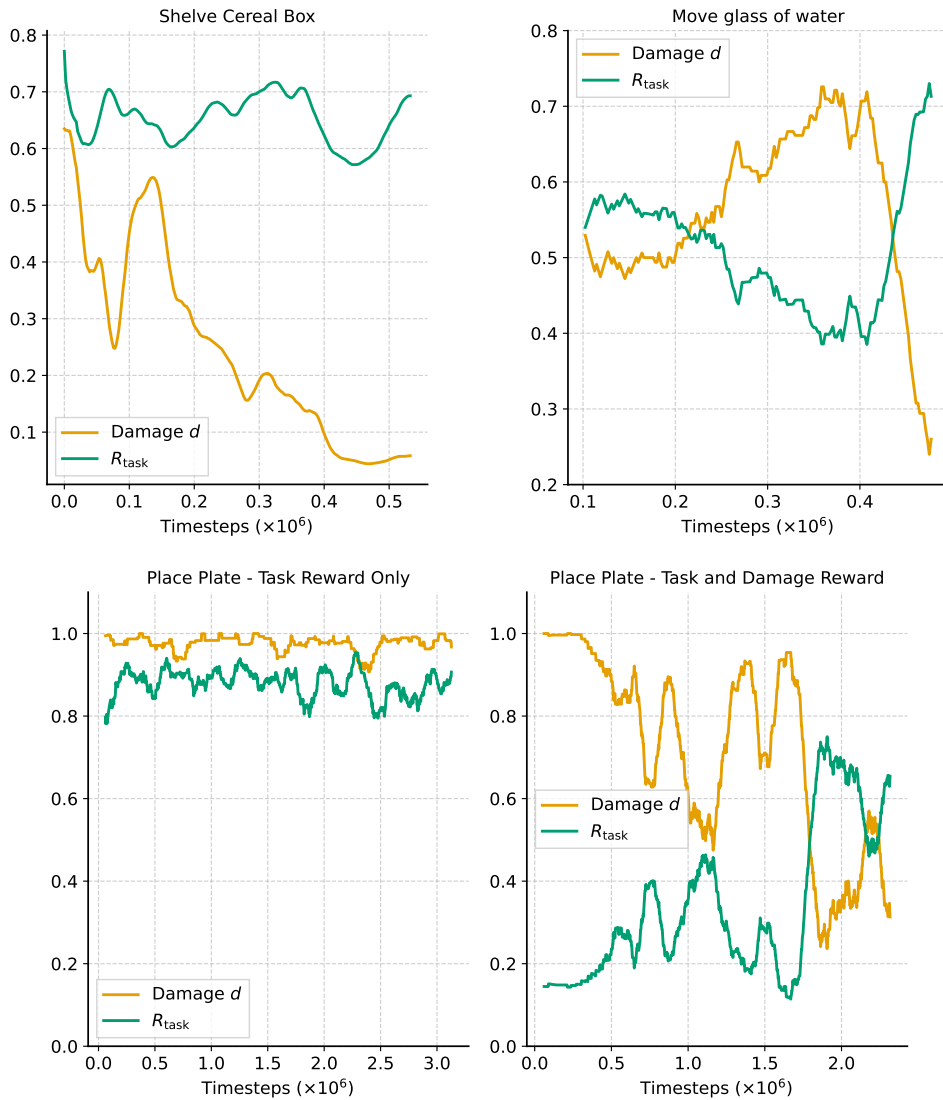


Fig. E.1: **Reinforcement Learning with OOPSIEVERSE**. We use DAMAGE SIM’s damage and health signal for RL training for SHELVE ITEM with DSRL, and MOVE GLASS OF WATER and PLACE PLATE with PPO. We plot task return (green) and damage (yellow) to decouple task performance from safety metrics, demonstrating that policies trained with DAMAGE SIM can maintain high task success while progressively reducing the damage. (*top row:*) SHELVE ITEM with DSRL (*left*) and MOVE GLASS OF WATER with PPO (*right*). (*bottom row:*) PLACE PLATE trained with PPO using task reward only (*left*) and task reward with the damage penalty enabled (*right*), illustrating that without damage feedback the policy attains the goal but incurs substantially higher damage. For visualization, the curves are smoothed using a moving average with a window of the last 10 episode returns for SHELVE ITEM (top left) and the last 60 episode returns for the other three plots.

F. SETTING OBJECT PARAMETERS

DAMAGE SIM models mechanical, thermal, and fluid damage using the object-specific parameters introduced in Section III. Because real objects exhibit diverse failure behaviors that depend on properties such as material, geometry, etc., there is no single set of parameters that accurately captures all objects. Instead, DAMAGE SIM exposes these parameters to the user and treats them as part of the environment specification.

DAMAGE SIM parameters are set through the same configuration file (or dictionary) used to construct the environment. Each object’s damage parameters are specified alongside standard per-object attributes already present in simulator configs (e.g.,

pose, scale). We provide parameter settings for all objects used in our tasks and experiments, but the configuration is intentionally editable: users can easily make objects more or less fragile based on the behaviors they wish to model.

Damage tracking can be enabled for any subset of objects in the scene. For articulated objects, users may optionally restrict tracking to a subset of rigid parts (“links”) and can further provide link-specific overrides. This level of specification is *not required* for typical use (the default is to track all links with shared parameters), but it is useful when different parts of an object exhibit different behaviors (e.g., the different parts of a Tiago robot), and a user wishes to model them.

Crucially, specifying parameters for new objects does not

require strong domain expertise. The objective is to establish a physically grounded proxy for unsafe interactions rather than perfectly modeling every material property. As such, users can utilize our provided benchmark objects as starting points for objects with similar properties and perform minimal iterative tuning to achieve effective results.

Tables F.7, F.8, and F.9 list the per-object parameter values used in our experiments for the mechanical, thermal, and fluid evaluators, respectively.

Object	α	β	$\mathcal{E}_{\text{mech}}$	Λ_{mech}
default	1.0	1.0	30.0	0.1
agent (robot)	0.01	1.0	70.0	0.2
microwave	1.0	1.0	100.0	1.0
camera_tripod	0.1	1.0	150.0	1.0
digital_camera	1.0	0.5	60.0	100.0
scrub_brush	0.01	0.01	300.0	100.0
bottle_of_wine	1.0	0.5	50.0	100.0
wineglass	1.0	0.5	15.0	100.0
bottle_of_beer	1.0	0.5	15.0	100.0
bag_of_flour	0.1	0.1	150.0	100.0
box_of_crackers	0.1	0.8	200.0	1.0
stand	0.001	0.001	500.0	1.0
laptop	1.0	0.5	80.0	100.0
water_glass	1.0	0.5	50.0	100.0
coffee_cup	1.0	0.5	150.0	1.0
plate	1.0	0.5	50.0	10.0
vase	1.0	0.5	30.0	10.0
pedestal_table	0.1	0.1	150.0	1.0
swivel_chair	0.1	0.1	500.0	1.0

TABLE F.7: Mechanical damage parameters per object. Symbols match Sec. III: $\varepsilon_{\text{mech}} = \alpha F_{\parallel} + \beta F_{\perp}$ and $d_{\text{mech}} = \Lambda_{\text{mech}} \max(\varepsilon_{\text{mech}} - \mathcal{E}_{\text{mech}}, 0)$.

Object	\mathcal{T}_{hot}	$\mathcal{T}_{\text{cold}}$	Λ_{therm}
agent (robot)	40.0	-20.0	1.0
camera_tripod	50.0	-30.0	0.1
digital_camera	50.0	-30.0	0.1
laptop	50.0	-20.0	0.1
vase	60.0	-40.0	0.01
pedestal_table	60.0	-40.0	0.01
swivel_chair	60.0	-60.0	0.01

TABLE F.8: Thermal damage parameters per object. Symbols match Sec. III: $d_{\text{therm}}(t) = \Lambda_{\text{therm}} \max(T - \mathcal{T}_{\text{hot}}, 0) + \Lambda_{\text{therm}} \max(\mathcal{T}_{\text{cold}} - T, 0)$ (our code uses a single scale Λ_{therm} applied symmetrically).

Object	C_{fluid}	Λ_{fluid}
agent (robot)	10.0	10.0
digital_camera	20.0	5.0
laptop	20.0	5.0

TABLE F.9: Liquid-exposure parameters per object. Symbols match Sec. III: damage begins once exposure exceeds C_{fluid} and then accumulates as $d_{\text{fluid}}(t) = \Lambda_{\text{fluid}} \max(C(t) - C_{\text{fluid}}, 0)$.