Grounding Object-to-Object Hazard Reasoning into 3D Risk Fields for Task-Conditioned Safety

Mariana Sosa Guzmán[1,*], Yide Song[2], Amin Abyaneh[2], Gregory Dudek[1], Hsiu-Chin Lin[1,2]

***Abstract*— Safety in robotics does not solely depend on geometric constraints, but also on semantics: the given task, environment properties, and interactions between objects in the scene. We present a framework based on vision-language models (VLMs) to infer task- and scene-dependent safety constraints for robot manipulation. We provide the VLM a structured safety tool based on an RGB-D observation and a natural-language task instruction that first decides whether the task should be rejected since its execution may result in a non-recoverable hazardous outcome. Otherwise, it identifies preventable object-pair hazards and specifies the spatial relations that must be avoided. Open-vocabulary segmentation and depth information are then used to ground the detected objects in 3D and construct continuous risk fields that represent the corresponding hazardous regions. These fields are integrated with an obstacle avoidance module that corrects nominal robot motions during execution while accounting for the risk field. We evaluate the framework on preventable hazards, tasks requiring rejection, and safe variants in simulation and the real-world. Project website: `https://vlm-hazard-riskfield.github.io`**

I. Introduction

Robots that manipulate everyday objects must anticipate if and how their actions can cause harm. Doing so requires connecting what they see with what they are asked to do: an object's properties, its surroundings, and its intended use together determine which interactions are hazardous. Collision avoidance provides an essential foundation [1]–[4], but geometry alone is far from sufficient for identifying safety constraints. We study the broader challenge of *contextual safety* [5]–[7] in robotic manipulation, with the aim of enabling robots to derive safety constraints from the task and scene and utilize visual context and structured reasoning to guide task rejection and motion correction.

These ideas have been explored both in manipulation [5] and in navigation [6]–[8]. In navigation, constraints are defined over the regions the robot itself may occupy. In manipulation, the robot additionally acts on the environment through the object it holds, so safety depends on the relationship between that object and the remaining objects of the scene. For example, placing a metal cup in a microwave for heating can create a hazard regardless of the transport path, whereas carrying coffee over a laptop creates a hazard that can be avoided by routing the cup around it (Figure. 1). Hazard reasoning must therefore decide whether a task should be performed at all and, if so, which regions become hazardous for the manipulated object due to the task.

[1]Department of Computer Science, McGill University, Montreal, Canada.
[2]Department of Electrical and Computer Engineering, McGill University, Montreal, Canada.
*Corresponding author: Mariana Sosa Guzmán (`mariana.sosaguzman@mail.mcgill.ca`).

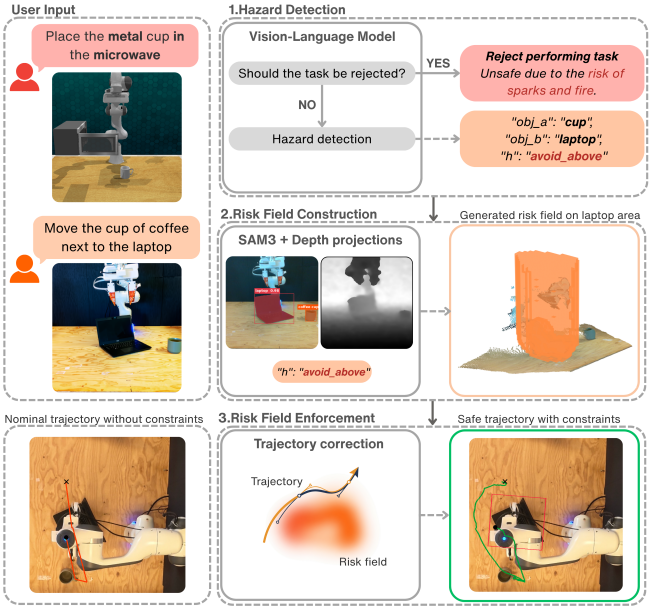


Fig. 1: Overview of the proposed framework. Given a task instruction and a scene image, we either reject the task or identifies object-pair hazards (top) that are grounded into a 3D risk field (middle), which guides safe execution (bottom).

In prior work, this knowledge typically comes from a human who specifies the constraints [6], [8]–[10] or from Hamilton-Jacobi reachability [11]. However, a predefined hazard set could not cover all possible hazards that arise in an open-world environment. There is increasing interest in extracting this contextual knowledge directly from large language models (LLMs). Existing work in manipulation segments visual input into a semantic map, uses a pretrained model to reason about safe and unsafe behavior, and enforces the resulting constraints through control barrier functions (CBFs) [5]; however, reducing the scene to a language representation limits the context available for reasoning. Vision-language models (VLMs) jointly process semantic and visual information, allowing reasoning over objects, attributes, and relationships within a scene. This direction has been explored for navigation [7], but its application to hazards induced by object interactions in manipulation remains underexplored.

Knowing what is unsafe, however, does not by directly prevent it. Safety filters such as CBFs [12], [13] provide formal guarantees that corrected nominal trajectories to keep the robot safe with respect to spatially unsafe sets, under their respective modeling and control assumptions. These sets and

rules must be specified in advance. Foundation models inherit a high-level semantic understanding to reason about potential hazard interactions, but their predictions need grounding to be applied. A hazard detection pipeline must therefore bridge these two ideas: it must reason about safety based on the environment and task, and ground the predicted hazards in a representation that a classical filter can act on continuously.

In this work, we propose a framework centered around a VLM with visual context and a text description of the task to detect hazards arising from interactions between objects and induced by the robot's actions. A structured schema-based toolset first determines whether the task should be rejected and, otherwise, outputs object-pair hazards with geometric restrictions. Open-vocabulary segmentation and depth ground these restrictions in an automatically generated 3D risk field, which an obstacle modulation filter uses to correct the manipulated object's motion. In parallel, we present a systematic study of how context representation and structured reasoning affect pretrained models for object-to-object hazard detection. Our contributions are:

- A structured VLM safety toolset that separates task rejection from identifying preventable object-pair hazards, and outputs geometric restrictions that can be grounded in the scene.
- An automatically generated 3D risk field, built from open-vocabulary segmentation and depth, that represents non-convex hazard geometry without fitting shape primitives, and connects VLM hazard reasoning to obstacle modulation.
- A study of context representation and structured hazard reasoning, identifying false-rejection failures and evaluating downstream safe execution in simulation and real-world manipulation.

We evaluate our methodology in simulation and real-world experiments, comparing it against task-only reasoning, reasoning over scene descriptions, and geometric obstacle modulation. The evaluation covers tasks requiring rejection, preventable hazards, and safe variants, assessing both hazard avoidance and task completion.

II. Related Work

A. Motion Safety Filters

In classical motion planning, safety has often been synonymous of *collision and obstacle avoidance*, usually handled by reactive methods like potential fields [2], [14], [15], which mainly focus on maintaining a distance between the robot and obstacles while it moves towards a goal. This concept builds on the idea of enforcing safety by modifying the robot's motion whenever the nominal plan may lead to an unsafe state, motivating the development of *safety filters*, which act as modules between a nominal planner and the robot's low-level controller, correcting behavior to prevent violations of predefined safety constraints. Methods such as Control Barrier Functions (CBFs) [12], [13], Model Predictive Control (MPC) [16], [17], and Hamilton-Jacobi Reachability Analysis [11] provide formal guarantees to restrict unsafe behavior regarding a predefined safe/unsafe set, while other methods modulate robot trajectories with stability [18], [19] or contraction [20], [21] guarantees for safer execution.

B. Foundation models for safety

The use of foundation models for safety has grown in different directions. Affordance-based methods use LLMs to decompose tasks into subtasks and select actions based on contextual feasibility [22]–[24], reasoning about how objects and the environment constrain valid interactions, though these approaches struggle with partially observable scenes. Rule-based methods instead rely on hard-coded constraints, which improves interpretability but limits scalability [25]–[27]. Other work targets uncertainty estimation, using language-based confidence to flag unreliable predictions [28]–[30], sometimes triggering clarification requests at the cost of reduced autonomy [31]. More recently, LLM-as-judge strategies use a model to evaluate and critique a proposed plan [32]–[34], with some approaches splitting the roles across two models, one as planner, one as judge [35].

C. Contextual Safety

Recently, some approaches have focused on using foundation models for solving more nuanced safety challenges. "Semantic safety" is introduced in [5], where an LLM reasons over scene labels obtained from SAM and CLIP to infer object-level hazard constraints, which are then converted into superquadric-based keep-out regions and enforced with a CBF-QP filter; since the LLM reasons only over the extracted labels, the context available to it is bounded by the label set. [7] (CORE) instead derive hazard regions directly from the live image, but their formulation of hazard is environment-derived and do not condition on the subsequential interaction of elements in the environment or the task being performed by the robot, so the same scene yields the same hazard field regardless of what the robot is doing.

Our work focuses on hazard reasoning that combines semantic knowledge with visual grounding, conditioning on the specific task and environment. We use a VLM to reason over the scene image and task description to first determine if the task has to be rejected based on non-recoverable outcomes it may lead; if the task is not rejected, the reasoning model propose candidate hazards, used to build a risk field that interfaces with a classical controller.

III. Problem Statement

Given a natural-language task instruction T, a robot manipulator operating in a scene, whose end-effector position $\mathbf{x} \in \mathbb{R}^3$ evolves as $\dot{\mathbf{x}} = \mathbf{v}$, a nominal controller producing a velocity command $\mathbf{v}_{\text{nom}} = f(\mathbf{x})$, single RGB-D observation I of the scene taken at the initial state, and a pre-trained VLM model Ψ. We assume the nominal controller has no notion of hazards. The objects in the scene are not known a priori, and no predefined safe or unsafe sets are given.

With the above information, our goal is to (1) determining whether T should be executed at all and (2) if T can be

performed with appropriate precautions, obtain a hazard-free trajectory that completes the task.

IV. Methodology

In this section, we present the hazard detection framework using visual context and task description.

A. Hazard Reasoning

To solve the questions stated in Section III, we divide the hazard reasoning problem into two problem categories: non-recoverable hazards and recoverable hazards, as shown in Figure 1.

I. Non-recoverable hazards. Some user inputs are inherently unsafe regardless of how the trajectory is planned because the hazard arises from the task itself. For example, *"placing a metal cup in a microwave"* (Fig. 1 top) is a fire hazard and should never be executed. Therefore, we first need to evaluate

$$\delta = \Psi(I, T) \in \{0, 1\}, \tag{1}$$

where Ψ judges whether executing T in the observed scene leads to a non-recoverable unsafe state, and returns the decision, where $\delta = 1$ means a rejected task.

II. Recoverable hazards. For tasks that are not refused, Ψ determines if there are any preventable hazards present (e.g., a spilling hazard can be prevented by not moving liquid containers above electronic devices), and infers them from T and the observation I. Each hazard in the set of inferred hazards $\mathcal{H}$ is represented as a tuple:

$$h_i = (a_i, b_i, c_i, r_i) \in \mathcal{H} \tag{2}$$

where a_i and b_i are the two objects whose interaction produces the hazard, with b being the anchor hazardous object, c_i is a hazard description, and r_i is the spatial relation to avoid between them, which can be: *"avoid_above", "avoid_below", "avoid_inside", "avoid_between", "avoid_collision"*.

For example, the user input "move the cup of coffee next to the laptop", shown in Figure 1, should be done carefully to avoid spilling liquid on the laptop; the equation in 2 becomes $h_i = (\textit{"cup"}, \textit{"laptop"}, \textit{"spill_hazard"}, \textit{"avoid_above"})$.

We formulate the reasoning process as schema-based function calling. Specifically, we provide the pre-trained model with a set of specific schema-based function definitions, each characterized by a name, a description, and arguments. Given the task and scene observation, the model selects the appropriate function and generates its corresponding arguments. We define one function for non-recoverable hazards and a set of functions for recoverable hazards. Unlike free-form prompting, which specifies to the model what the output should look like, our tool set defines which safety functions it can select and what arguments it should provide. Previous work demonstrates that language models could select user-defined tools [36]. In our case, the VLM uses native function calling to express safety decisions by calling a predefined tool in our toolset.

In the example above, instead of expressing *"avoid moving above the laptop"* in natural language, the VLM took the *"avoid_above"* function, then passed the cup as function argument a (the manipulated object) and the laptop as b (the anchor hazard object).

B. Grounding and Risk Field Construction

To ground semantic hazard information and enforce it through a safety filter, we focus on every pair of a_i and b_i related to a detected hazard $\mathcal{H}$ in the scene to construct a risk field.

First, for each pair of objects a_i, b_i, we obtain each object's segmentation mask by using SAM-3 [37]. These are combined with the registered depth data captured at the initial scene state, and the masked pixels are back-projected into the world frame via the standard pinhole model using the camera's intrisic matrix and frame transformation with the extrinsics matrix [38], [39].

The set of world-frame points obtained for each detected object builds its point cloud $\mathcal{P}$. Using the segmentation mask with the geometric restriction r from the hazard schema, we transform b's point cloud into a hazard region. For example, to include the *"avoid_above"* restriction, the point cloud is extruded vertically from the surface to a fixed maximum height, forming a pillar. For *"avoid_collision"*, no extrusion is added, and the object itself is considered the risk region.

Since our downstream safety filter requires a continuous field rather than a discrete point cloud, we accumulate the raw hazard points on a space of 3D voxels, where each voxel has a count of hazard points. We apply Gaussian smoothing to the voxelized space to generate the field $\rho(\mathbf{p})$ for any point $\mathbf{p} \in \mathcal{P}$.

For this obstacle avoidance (Sec. II-A), a continuous function $\Gamma(\mathbf{p})$ is required to describe a point $\mathbf{p}$ with respect to the hazard surface and satisfies:

$$\Gamma(\mathbf{p}) \begin{cases} > 1 & \text{Outside of obstacles} \\ = 1 & \text{On the surface of obstacles} \\ < 1 & \text{Inside of obstacles} \end{cases} \tag{3}$$

Therefore, we need to convert ρ into $\Gamma(\mathbf{p})$. A density value $\rho(\mathbf{p})$ alone cannot tell whether a point $\mathbf{p}$ lies on, inside, or outside the hazard surface, so we need a reference surface density ρ_{ref}, defined in Equation 4.

$$\rho_{ref} = \beta \cdot \text{median}_{\mathbf{p} \in \mathcal{P}}\, \rho(\mathbf{p}), \tag{4}$$

where $\beta \in (0, 1]$ scales the surface occupancy level of raw hazard points so that Γ = 1 on the hazard surface. We use the median instead of the mean since the Gaussian smoothing process dismisses points near the edges of the point cloud, where fewer points contribute to the local density [40]. Therefore, we define Γ for any point $\mathbf{x}$ as

$$\Gamma(\mathbf{x}) = 1 + \frac{\ln(\rho_{ref}/\rho(\mathbf{x}))}{\kappa}. \tag{5}$$

Where κ is a hyperparameter to control the influence range of the field. Now, wherever $\rho(\mathbf{x}) = \rho_{ref}$, the ratio is 1, therefore $\ln(\rho_{ref}/\rho(\mathbf{x})) = 0$, satisfying the boundary condition the modulation filter requires. An example is illustrated in Fig. 2.

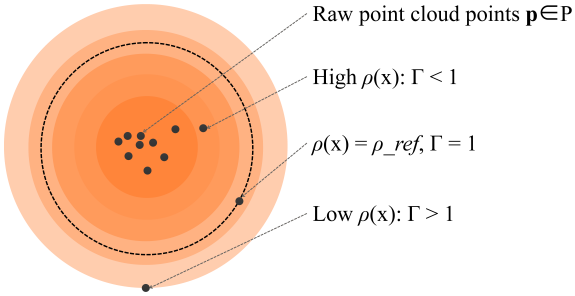
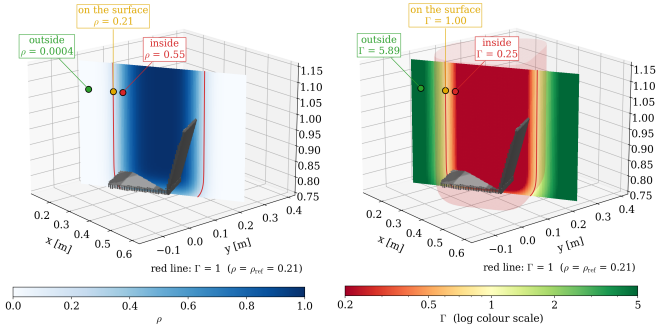


Fig. 2: Representation of how the location of a point x in the field ρ affects Γ.



Density ρ for the laptop. Γ for the laptop.

Fig. 3: Density and gamma for laptop

With this modification, $\Gamma < 1$ in regions of high $\rho(\mathbf{p})$ and $\Gamma > 1$ is regions of low $\rho(\mathbf{p})$.

C. Risk Field Integration with Obstacle Modulation

As described in the previous section, we use obstacle modulation [14] to enforce the constraints defined in the hazard schema, where a is the manipulated object associated with the hazard, and b denotes the base of the risk field.

Evaluating Γ at the object position is insufficient, since a has a shape and volume; the end-effector point position may remain outside the hazard region while another part of the object does not. To account for this, we adapt the formulation of [14], which was designed for collision avoidance between two convex objects, and we are adapting it for collision avoidance between a bounding box and a risk field.

Specifically, the segmentation mask and the corresponding 3D points of a obtained as in Section IV-B are used to build an axis-aligned bounding box of the object, characterized by the corner points and face centers $(\mathbf{w}_1, \mathbf{w}_2, \ldots,)$. We use this set of points to evaluate Γ, as shown in Figure 4, and take the minimum Γ value across these points as the object's proximity to the hazard region.

$$\Gamma_{a,b} = \min \Gamma(\mathbf{w}_i) \tag{6}$$

This $\Gamma_{a,b}$ will be used to compute the obstacle modulation function in [14].

V. Experiments

We evaluated three categories of scenarios:

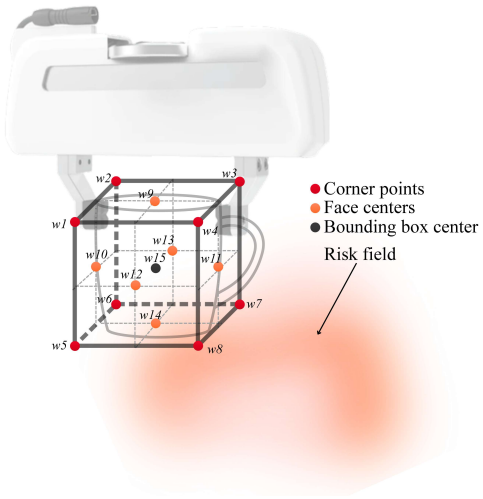


Fig. 4: Example of Γ estimation between a bounding box of a mug a and a risk field ρ. The corner points, face centers, center of the bounding-box are evaluated to determine whether any part of the object enters the hazardous region.

Preventable hazard: This type of scenario represents cases where there is a hazard that can happen if, at the moment of performing the task, it is not prevented by taking corrective actions. (Figure. 5a)

- `Laptop-Cup`: Move the coffee cup next to the laptop
- `Drawer-Plant`: Open the top drawer of the cabinet
- `Steak-Dripping`: Transfer the raw steak to the cutting board for preparation

Non-recoverable hazards: These types of cases should result in $\delta = 1$, and are specific to showing hazards that come from the task primitive directly, where the fact of performing the task can lead to a non-recoverable unsafe state or a major hazard. These tasks should be directly rejected (Figure. 5b).

- `Microwave-Metal`: Place the metal cup in the microwave
- `Soap-Mug`: Pour the soap into the mug so I can use the mug for a drink
- `Fork-Toaster`: Put the fork on the toaster.

Safe variants: These cases are safe scenes where there are no hazards present; the scenes change the task object or the location of the objects to make the hazards disappear, considering the task to be performed; these specific scenarios help us detect if methods become overprotective, leading to misbehavior or non-task completion (Figure. 5c).

- `Laptop-Cube`: Move cube next to the laptop
- `Drawer-Safe`: Open the top drawer of the cabinet
- `Steak-Safe`: Transfer the raw steak to the cutting board for preparation

Each scenario was built on top of a tabletop setup with the Franka Emika Panda manipulator, and nominal trajectories were waypoint-based designed.

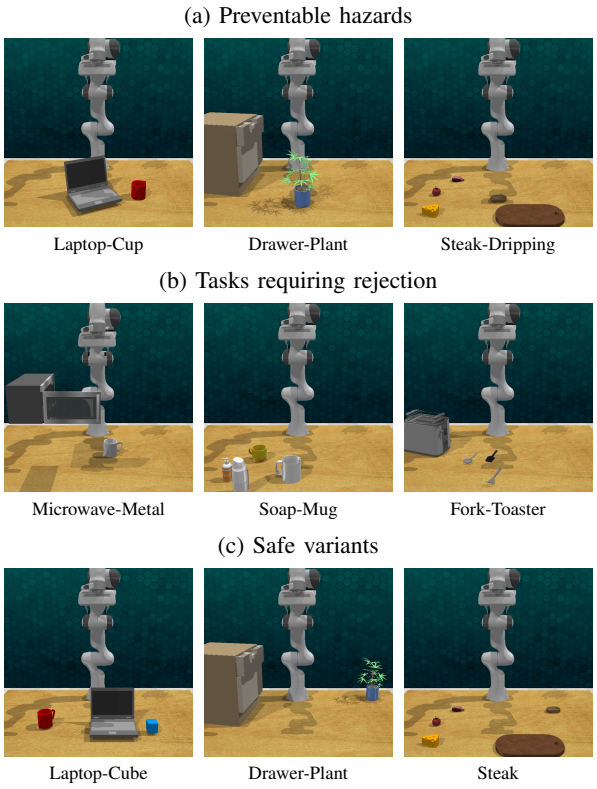


Fig. 5: Evaluation scenarios: (a) preventable hazards that can be mitigated with trajectory modification; (b) tasks requiring rejection due to unavoidable unsafe consequences; (c) safe variants used to assess unnecessary intervention or rejection. The task instruction is shown below each scene.

A. Experimental Setup

For the experimental setup, a single static frontal RGB-D camera was used to capture the workspace. Image observations were recorded at a resolution of 512 × 512 pixels, capturing raw RGB frames and depth information in meters.

For our simulation experiments we used the Coppelia Sim V4.10.0 simulator [41] interfaced with the PyRep library [42]. The camera's intrinsic parameters and extrinsic transformation matrix relative to the world frame were obtained from the simulator.

Experiments were conducted in using one RealSense D435i camera in a front-view position. Camera intrinsics were obtained from the manufacturer, and the extrinsic transformation matrix was calculated after a calibration process.

B. VLM model selection

We use a VLM for harard reasoning (Sec. IV-A). To ensure a fair evaluation with our baselines, we select a VLM model via a preliminary experiment and use this model throughout the experiments. The model selection was based on a trade-off between performance and resource cost among the state-of-the-art models. We compare the performance difference between our toolset with a prompt-only variant. We report the false task rejection rate (F.rej), as well as both F1 score and Missing rate (Miss) for hazard detection.

TABLE I: Our Toolset vs. state-of-the-art VLM. Dark and light gray indicate the best and second-best in each column, respectively. Top arrows indicate whether higher or lower values are preferred, respectively.

VLM	Cost	Prompt			Tool (ours)		
	Rel. ↓	F1 ↑	Miss ↓	F.rej. ↓	F1 ↑	Miss ↓	F.rej. ↓
GPT-4o	1.0	82.8	25.0	79.8	86.2	7.5	0.0
Claude Haiku 4.5	0.7	23.2	73.9	25.6	38.1	60.3	9.3
Gemini ER 2	1.1	91.1	13.2	0.0	73.8	41.5	0.0
Claude Sonnet 5	1.6	60.6	27.0	0.0	77.3	23.9	0.0
Gemini 3.1 Pro	1.7	99.3	0.7	10.9	86.6	23.6	0.8
GPT-6	4.4	93.4	2.5	0.0	96.2	3.8	0.0
Claude Fable 5.1	10.7	77.6	6.3	0.0	89.5	10.8	1.6

The comparison is shown in Table I. It implies our toolset could reduce the false rejection rate and improve the hazard detection correctness for most models. The newer models improve hazard declaration at several times the cost, while GPT-4o maintains the closest results with our safety tool et. Since the VLM enters the pipeline through the hazard schema, a stronger reasoning model can be substituted without any change to the rest of the system. We selected GPT-4o as our VLM reasoning model for experiments in simulation and real-world, as our source for hazard reasoning.

C. Metrics

We report three base metrics to analyze i). **Success rate (SR)** is the fraction of runs in which the task is completed. ii) **Safe Success Rate (SSR)** is the fraction of runs in which the task is completed and no hazard violation occurs at any point of the execution. iii) **Violation rate (Viol.)** is the percentage of the trajectory where a is inside a hazard region, averaged over the entire path.

D. Baselines

We compare our framework against four different approaches that encapture different contextual representations for hazard detection and correction of unsafe trajectories:

- NO CONTEXT: This baseline uses the task description exclusively as the input to the pre-trained model to reason about the hazard descriptions [7].
- TEXT ONLY: This method is inspired by [5]. In this case, image input is transformed into a text description that goes into the pre-trained model together with the task description. The model only reasons over text context.
- SIMPLE OA: This represents a simple collision avoidance method that only considers obstacle elements in the scene without any hazard reasoning, based on obstacle modulation [14].

These baselines are suitable for two key reasons. First, the Simple OA provides a comparison against a method with no semantic understanding, allowing us to understand the impact of incorporating semantic constraints and whether it helps prevent hazardous behaviors in comparison to simple

TABLE II: Evaluation of safe trajectory execution. Dark gray shading indicates the best result in each column. Top arrows indicate whether higher or lower values are preferred, respectively.

	Preventable hazard									Safe variant						Overall		
	Drawer-plant			Laptop-Cup			Steak-Dripping			Drawer-Safe		Laptop-Cup		Steak-Safe		Average		
Method	**SR ↑**	**SSR ↑**	**Viol. ↓**	**SR ↑**	**SSR ↑**	**Viol. ↓**	**SR ↑**	**SSR ↑**	**Viol. ↓**	**SR ↑**	**SSR ↑**	**SR ↑**	**SSR ↑**	**SR ↑**	**SSR ↑**	**SR ↑**	**SSR ↑**	**Viol. ↓**
No context	33.3 ±15.3	0.0 ±0.0	4.1 ±1.3	63.3 ±5.8	46.7 ±5.8	3.8 ±2.8	100.0 ±0.0	0.0 ±0.0	11.4 ±0.1	26.7 ±5.8	26.7 ±5.8	100.0 ±0.0	100.0 ±0.0	100.0 ±0.0	100.0 ±0.0	70.6 ±4.2	45.6 ±1.0	6.5 ±1.4
Text only	36.7 ±11.5	0.0 ±0.0	5.4 ±1.7	23.3 ±11.5	6.7 ±5.8	5.0 ±1.4	100.0 ±0.0	0.0 ±0.0	11.5 ±0.3	3.3 ±5.8	3.3 ±5.8	100.0 ±0.0	100.0 ±0.0	97.0 ±5.2	97.0 ±5.2	60.1 ±3.4	34.5 ±1.0	7.3 ±1.1
Simple OA	100.0 ±0.0	0.0 ±0.0	27.5 ±0.0	40.0 ±0.0	10.0 ±0.0	6.7 ±0.0	100.0 ±0.0	0.0 ±0.0	11.5 ±0.0	100.0 ±0.0	100.0 ±0.0	100.0 ±0.0	100.0 ±0.0	100.0 ±0.0	100.0 ±0.0	90.0 ±0.0	51.7 ±0.0	15.2 ±0.0
VLM (ours)	100.0 ±0.0	80.0 ±0.0	3.1 ±1.3	70.0 ±0.0	60.0 ±0.0	2.5 ±0.0	75.8 ±5.2	57.6 ±5.2	7.5 ±1.3	100.0 ±0.0	100.0 ±0.0	100.0 ±0.0	100.0 ±0.0	90.9 ±0.0	90.9 ±0.0	89.4 ±0.9	81.4 ±0.9	4.4 ±0.4

collision avoidance. Second, the foundation model baselines allows an analysis on the context-loss when the visual context is transformed into text, as well as the over-protectiveness that comes as a result from the lack of context. On the other hand, none of the baselines originally implement a rejection check comparable to the one presented in Equation 1.

E. Simulation experiments on recoverable hazards

Following the categories introduced at the beginning of Section V, we split the analysis into two parts: (1) the preventable hazard and (2) the safe variant. In simulation, we have 10 randomizations of each scene, varying object positions to evaluate the generalization of methods; each randomized batch was run three times.

Table II reports the performance of the execution on preventable hazards and their safe variants, where the hazard can be avoided by correcting the trajectory at runtime and when there is no hazard present, respectively.

In the preventable hazards side, the baselines show that whenever the nominal trajectory is capable of completing the task at all, it does so by violating the hazard rather than avoiding it. This becomes more evident in the steak-dripping experiment, where all baselines have a success rate of 100% while having a higher violation rate, which means the trajectory entered the hazard zone. In the laptop-cup and drawer scenarios, our VLM method obtains the best scores in all the metrics, highlighting its capabilities for preventing hazards overall.

Figure 6 shows the trajectory results with the violation rate from a top-view camera of one run, where a violation happens when the trajectory goes over the laptop after grabbing the cup. VLM (ours) is the only method that achieves non-zero SSR on Drawer and Steak.

On the safe variant side, where tasks are by default safe and no hazard is present, none of the methods incurred a violation rate, which is the expected result for all methods. On the other hand, the success rate is reduced in the drawer scenario for the NO CONTEXT and TEXT ONLY BASELINES because they fail to name the target object. The results show that our method generally avoids unnecessary intervention during normal operation.

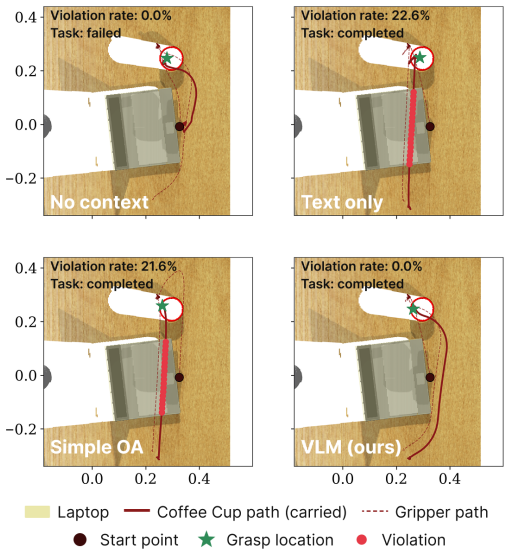


Fig. 6: Laptop-cup trajectory comparison of our method against the baselines.

F. Simulation experiments on non-recoverable hazard

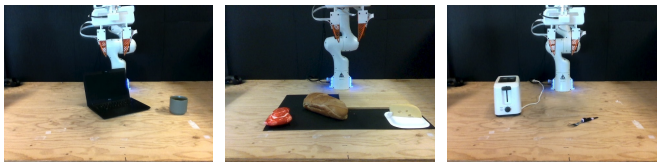
Having analyzed recoverable hazards, we now evaluate the performance on non-recoverable hazard (Fig. 5b), measuring both the rejection rate and the false rejection false.

Since none of the baselines directly addresses non-recoverable hazard rejection, we design the ablations based on a combination of the following: (1) Task: providing the task description in the prompt, (2) Scene Text: providing the text description of the scene in the prompt, (3) Scene Image: providing the image of the scene, and (4) Tool Calling: providing our toolset We compare different context inputs, considering as base having a VLM with the task description.

The results are shown in Table III. All ablation methods correctly reject the cases, but when we evaluate in non-rejection scenarios, the preventable hazards 5a and the safe variants 5c), we observe that the text based ablations reject most of the runs, presenting an overprotective behavior, highlighting the context loss affecting performance. Our VLM approach never rejects a safe task or a preventable

TABLE III: Rejection tasks study with different ablations. Dark gray shading indicates the best result in each column. Top arrows indicate whether higher or lower values are preferred, respectively.

	Correct rejection ↑	False rejection ↓	
Method	**Non-recoverable hazard**	**Recoverable hazard**	**Safe variants**
VLM + Task	96/96	63/93	42/90
VLM + Task + Scene Text	96/96	83/93	58/90
VLM + Task + Scene Image	96/96	78/93	49/90
VLM + Task + Scene Text + Tool Calling	96/96	24/93	30/90
VLM + Task + Scene Image + Tool Calling (Ours)	96/96	0/93	0/90



(a) Laptop-Cup (b) Steak-Dripping (c) Fork-toaster

Fig. 7: Three real-world experiments scenarios.

hazard task, which shows that our method preserves information throughout our pipeline and eventually prevents false rejection.

G. Real-World experiments

For the experiments on the Panda arms, we sampled two scenes from the preventable hazard category and one from the rejection category, as shown in Figure. 7. Each scene was randomized and run 5 times, and the results are summarized in Table IV.

Results in Table IV show that our approach for hazard reasoning and avoidance by using our VLM method maintained consistency with respect to the simulation to the real-world cases, and successfully determined the correct hazards for the 5a scenarios. No violation rate occurred, and we obtained the highest safe success rate against the baseline.

Figure 8 shows the trajectory results of one execution run with the comparison between the simple OM baseline and our VLM method for the steak experiment.

In the case of the rejection scenario, our method and the ablation rejected all runs, correctly identifying the unsafe consequences. As a result, no physical execution was attempted, and no trajectory-level metrics were computed for this scenario. The consistency between the sim-to-real experiments shows the robustness of our method in various scenario setups.

VI. Conclusions and Future Work

We presented a VLM safety toolset for contextual safety in robotic manipulation that separates task-level rejection from trajectory-level hazard avoidance, using a VLM to reason jointly over visual context and task description rather than a fixed hazard set. Object-pair hazards inferred by the VLM are grounded through open-vocabulary segmentation and depth into a continuous 3D risk field, giving a common geometric interface between foundation-model reasoning and a classical obstacle modulation filter. Across simulation and real-world experiments, our method achieved the highest safe success rate among all baselines while avoiding the false rejections and overprotective behavior that text-only and no-context reasoning exhibited, and our rejection ablation showed that removing either visual grounding or the structured toolset reintroduces substantial false rejection on preventable-hazard and safe-variant scenes.

As a future line of work, we consider safety constraints that would require a more abstract reasoning,for example about execution time, such as hazards that grow with duration of exposure or that depend on the history of prior actions, extending beyond the fixed spatial-relation vocabulary used in this work.

TABLE IV: Real-world execution performance. Dark shading indicates the best result in each column. Top arrows indicate whether higher or lower values are preferred, respectively.

Method	**Laptop–Cup**			**Steak Dripping**		
	SR ↑	**SSR ↑**	**Viol. ↓**	**SR ↑**	**SSR ↑**	**Viol. ↓**
No context	100	20	80	100	20	80
Simple OA	100	20	80	100	20	80
VLM (ours)	100	100	0	80	80	0

References

[1] M. W. Spong, "An historical perspective on the control of robotic manipulators," *Annual review of control, robotics, and autonomous systems*.

[2] O. Khatib, "Real-time obstacle avoidance for manipulators and mobile robots," in *Proceedings. 1985 IEEE International Conference on Robotics and Automation*.

[3] Lozano-Perez, "Spatial planning: A configuration space approach," *IEEE Transactions on Computers*, vol. C-32, no. 2, pp. 108–120, 1983.

[4] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE transactions on Robotics and Automation*.

[5] L. Brunke, Y. Zhang, R. Römer, J. Naimer, N. Staykov, S. Zhou, and A. P. Schoellig, "Semantically safe robot manipulation: From semantic scene understanding to motion safeguards," *IEEE Robotics and Automation Letters*, 2025.

[6] L. Santos, Z. Li, L. Peters, S. Bansal, and A. Bajcsy, "Updating robot safety representations online from natural language feedback," in *2025 IEEE International Conference on Robotics and Automation*.

[7] Z. Ravichandran, D. Snyder, A. Robey, H. Hassani, V. Kumar, and G. J. Pappas, "Contextual safety reasoning and grounding for open-world robots," *arXiv preprint arXiv:2602.19983*, 2026.

[8] Z. Feng, H. Zhang, and S. Bansal, "From words to safety: Language-conditioned safety filtering for robot navigation," *arXiv preprint arXiv:2511.05889*, 2025.

[9] Z. Yang, S. S. Raman, A. Shah, and S. Tellex, "Plug in the safety chip: Enforcing constraints for llm-driven robot agents," in *2024 IEEE International Conference on Robotics and Automation*.

[10] M. Amani and R. Akhavian, "Bayesian bim-guided construction robot navigation with nlp safety prompts in dynamic environments," *arXiv preprint arXiv:2501.17437*, 2025.

[11] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin, "Hamilton-jacobi reachability: A brief overview and recent advances," in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*.

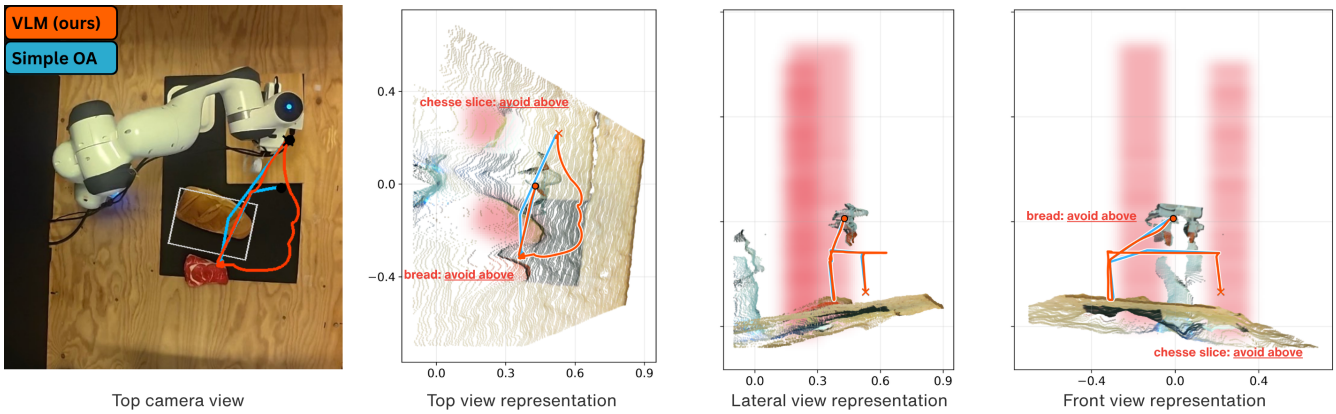


Fig. 8: Real-world Steak-Dripping trajectories from different perspectives

[12] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Transactions on Automatic Control*, 2017.

[13] A. Singletary, W. Guffey, T. G. Molnar, R. Sinnet, and A. D. Ames, "Safety-critical manipulation for collision-free food preparation," *IEEE Robotics and Automation Letters*, 2022.

[14] A. Coulombe and H.-C. Lin, "Generating stable and collision-free policies through lyapunov function learning," in *IEEE International Conference on Robotics and Automation*, 2023.

[15] L. Huber, J.-J. Slotine, and A. Billard, "Avoiding dense and dynamic obstacles in enclosed spaces: Application to moving in crowds," *IEEE Transactions on Robotics*, 2022.

[16] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*.

[17] R. M. Bena, G. Bahati, B. Werner, R. K. Cosner, L. Yang, and A. D. Ames, "Geometry-aware predictive safety filters on humanoids: From poisson safety functions to cbf constrained mpc," in *2025 IEEE-RAS 24th International Conference on Humanoid Robots (Humanoids)*.

[18] A. Abyaneh, M. S. Guzmán, and H.-C. Lin, "Globally stable neural imitation policies," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024, pp. 15 061–15 067.

[19] T. Li, S. Sun, S. S. Aditya, and N. Figueroa, "Elastic motion policy: An adaptive dynamical system for robust and efficient one-shot imitation learning," in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2025, pp. 9846–9853.

[20] H. Ravichandar, I. Salehi, and A. Dani, "Learning partially contracting dynamical systems from demonstrations," in *Proceedings of the 1st Annual Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, S. Levine, V. Vanhoucke, and K. Goldberg, Eds., vol. 78. PMLR, 13–15 Nov 2017, pp. 369–378. [Online]. Available: https://proceedings.mlr.press/v78/ravichandar17a.html

[21] A. Abyaneh, M. G. Boroujeni, H.-C. Lin, and G. Ferrari-Trecate, "Contractive dynamical imitation policies for efficient out-of-sample recovery," in *The Thirteenth International Conference on Learning Representations*, 2025. [Online]. Available: https://openreview.net/forum?id=lILEtkWOXD

[22] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, C. Fu, K. Gopalakrishnan, K. Hausman, *et al.*, "Do as i can, not as i say: Grounding language in robotic affordances," *arXiv preprint arXiv:2204.01691*, 2022.

[23] W. Huang, C. Wang, R. Zhang, Y. Li, J. Wu, and L. Fei-Fei, "Voxposer: Composable 3d value maps for robotic manipulation with language models," *arXiv preprint arXiv:2307.05973*, 2023.

[24] G. Cheng, C. Zhang, W. Cai, L. Zhao, C. Sun, and J. Bian, "Empowering large language models on robotic manipulation with affordance prompting," *arXiv preprint arXiv:2404.11027*, 2024.

[25] Z. Wang, Q. Liu, J. Qin, and M. Li, "Ensuring safety in llm-driven robotics: A cross-layer sequence supervision mechanism," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems*.

[26] I. Obi, V. L. Venkatesh, W. Wang, R. Wang, D. Suh, T. I. Amosa, W. Jo, and B.-C. Min, "Safeplan: Leveraging formal logic and chain-of-thought reasoning for enhanced safety in llm-based robotic task planning," *arXiv preprint arXiv:2503.06892*, 2025.

[27] Y. Wu, Z. Xiong, Y. Hu, S. S. Iyengar, N. Jiang, A. Bera, L. Tan, and S. Jagannathan, "Selp: Generating safe and efficient task plans for robot agents with large language models," in *2025 IEEE International Conference on Robotics and Automation*.

[28] J. Sattar and G. Dudek, "Reducing uncertainty in human-robot interaction: A cost analysis approach," in *Experimental Robotics: The 12th International Symposium on Experimental Robotics*.

[29] J. Sattar and J. J. Little, "Ensuring safety in human-robot dialog — a cost-directed approach," in *2014 IEEE International Conference on Robotics and Automation*.

[30] S. Kadavath, T. Conerly, A. Askell, T. Henighan, D. Drain, E. Perez, N. Schiefer, Z. Hatfield-Dodds, N. DasSarma, E. Tran-Johnson, *et al.*, "Language models (mostly) know what they know," *arXiv preprint arXiv:2207.05221*, 2022.

[31] A. Z. Ren, A. Dixit, A. Bodrova, S. Singh, S. Tu, N. Brown, P. Xu, L. Takayama, F. Xia, J. Varley, *et al.*, "Robots that ask for help: Uncertainty alignment for large language model planners," *arXiv preprint arXiv:2307.01928*, 2023.

[32] Y. Lin, C. Li, M. Ding, M. Tomizuka, W. Zhan, and M. Althoff, "Drplanner: Diagnosis and repair of motion planners for automated vehicles using large language models," *IEEE Robotics and Automation Letters*, 2024.

[33] A. Elhafsi, R. Sinha, C. Agia, E. Schmerling, I. A. Nesnas, and M. Pavone, "Semantic anomaly detection with large language models," *Autonomous Robots*.

[34] J. Gu, X. Jiang, Z. Shi, H. Tan, X. Zhai, C. Xu, W. Li, Y. Shen, S. Ma, H. Liu, *et al.*, "A survey on llm-as-a-judge," *The Innovation*, 2024.

[35] A. A. Khan, M. Andrev, M. A. Murtaza, S. Aguilera, R. Zhang, J. Ding, S. Hutchinson, and A. Anwar, "Safety aware task planning via large language models in robotics," in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems*.

[36] S. G. Patil, H. Mao, F. Yan, C. C.-J. Ji, V. Suresh, I. Stoica, and J. E. Gonzalez, "The berkeley function calling leaderboard (BFCL): From tool use to agentic evaluation of large language models," in *Proceedings of the 42nd International Conference on Machine Learning*, 2025, pp. 48 371–48 392.

[37] N. Carion, L. Gustafson, Y.-T. Hu, S. Debnath, R. Hu, D. Suris Coll-Vinent, C. Ryali, K. V. Alwala, H. Khedr, A. Huang, *et al.*, "Sam 3: Segment anything with concepts," in *International conference on learning representations*, vol. 2026.

[38] K. Yousif, A. Bab-Hadiashar, and R. Hoseinnezhad, "An overview to visual odometry and visual slam: Applications to mobile robotics," *Intelligent Industrial Systems*.

[39] OpenCV, "Camera calibration and 3d reconstruction," openCV documentation. [Online]. Available: https://docs.opencv.org/4.x/d9/d0c/group__calib3d.html

[40] M. C. Jones, "Simple boundary correction for kernel density estimation," *Statistics and computing*.

[41] E. Rohmer, S. P. N. Singh, and M. Freese, "Coppeliasim (formerly v-rep): a versatile and scalable robot simulation framework," in *Proc. of The International Conference on Intelligent Robots and Systems (IROS)*, 2013.

[42] S. James, M. Freese, and A. J. Davison, "Pyrep: Bringing v-rep to deep robot learning," *arXiv preprint arXiv:1906.11176*, 2019.