Towards Safer Homes: Behavior-Adaptive Risk Assessment for Service Robots Using Semantic Context*

Sena Ishii[1], Ankit A. Ravankar[1], Jose Victorio Salazar Luces[1] and Yasuhisa Hirata[1]

[1]Tohoku University, Sendai, Miyagi, 9808579, Japan

Abstract

We present a novel framework for estimating accident-prone regions in everyday indoor scenes, aimed at improving real-time risk awareness in service robots operating in human-centric environments. As robots become integrated into daily life, particularly in homes, the ability to anticipate and respond to environmental hazards is crucial for ensuring user safety, trust, and effective human-robot interaction. Our approach models object-level risk and context through a semantic graph-based propagation algorithm. Each object is represented as a node with an associated risk score, and risk propagates asymmetrically from high-risk to low-risk objects based on spatial proximity and accident relationship. This enables the robot to infer potential hazards even when they are not explicitly visible or labeled. Designed for interpretability and lightweight onboard deployment, our method is validated on a dataset with human-annotated risk regions, achieving a binary risk detection accuracy of 75%. The system demonstrates strong alignment with human perception, particularly in scenes involving sharp or unstable objects. These results underline the potential of context-aware risk reasoning to enhance robotic scene understanding and proactive safety behaviors in shared human-robot spaces. This framework could serve as a foundation for future systems that make context-driven safety decisions, provide real-time alerts, or autonomously assist users in avoiding or mitigating hazards within home environments.

Keywords

Risk reasoning, Home accidents, Service Robots

1. Introduction

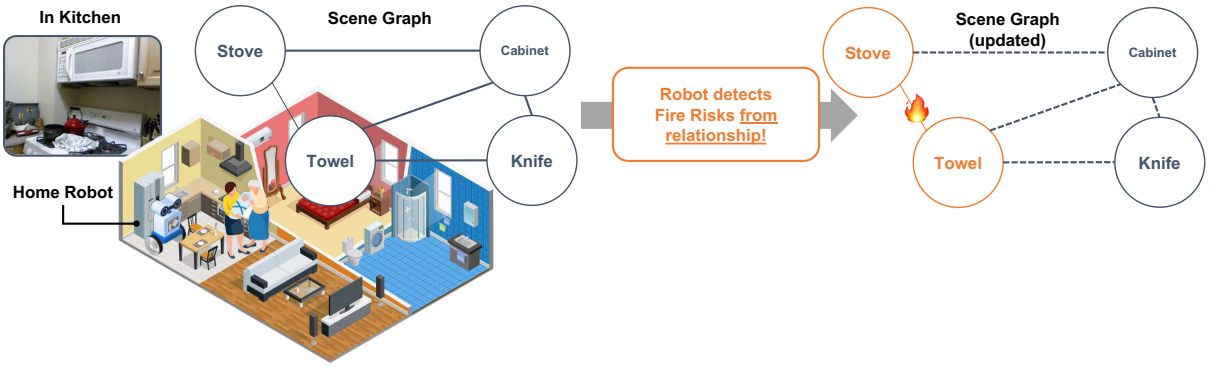

Figure 1: Overview of the proposed risk propagation framework in a household environment. The home robot detects potential fire risks by analyzing the spatial proximity and accident-related correlations among objects. In this scenario, the towel near the oven was initially not recognized as hazardous, but after risk propagation, it is identified as a high-risk object due to its proximity to the oven.

As service robots become increasingly integrated into daily life—supporting tasks such as cleaning, acting as communication companions, searching for objects or navigating shared spaces [1, 2, 3]—their

BEAR 2025 - Workshop on Benefits of pErsonalization and behAvioral adaptation in assistive Robots, within IEEE RO-MAN 2025, August 25-29 2025, Eindhoven, The Netherlands

*You can use this document as the template for preparing your publication. We recommend using the latest version of the ceurart style.

s.ishii@srd.mech.tohoku.ac.jp (S. Ishii); ankit@srd.mech.tohoku.ac.jp (A. A. Ravankar); j.salazar@srd.mech.tohoku.ac.jp (J. V. S. Luces); hirata@srd.mech.tohoku.ac.jp (Y. Hirata)

0009-0001-6978-4443 (S. Ishii); 0000-0003-0556-9194 (J. V. S. Luces)

roles are expected to expand beyond single-function behaviors. With recent advances in embodied intelligence and large language models, these robots are beginning to understand complex instructions and act autonomously in diverse home environments. While robot capabilities have advanced, developing systems that anticipate and mitigate household accidents remains critical, especially in aging societies where in-home safety is a pressing concern [4, 5]. Service robots, especially mobile home robots with their inherent mobility and environmental awareness, can play a vital role by predicting and responding to potential risks, both proactively preventing accidents and providing timely responses when incidents occur. By endowing these robots with contextual risk reasoning capabilities, their utility can be expanded beyond single-function tasks, leading to safer and more adaptive home environments.

Existing robotic technologies predict household accidents, with some, like AnomalyGen [6], leveraging large language models and 3D simulations. However, these approaches often determine risks based solely on observed objects, frequently misidentifying non-hazardous situations as risky. True risks, as humans intuitively recognize, emerge from objects' semantic roles, spatial relationships, and affordances, not just from the objects themselvess.

Prior works have addressed risk-aware navigation [7, 8, 9], fall-risk screening [10], and unsafe event detection [11]. Others use scene graphs [12], collaborative graph neural networks [13], or foundation models for task planning [14, 15, 16] to model object relationships and improve safety perception. Yet, these methods typically: (i) focus on traversability or physical terrain, (ii) rely on static rules or explicit task instructions, or (iii) lack real-time inference of latent risk in dynamic, unstructured home scenes.

Our work addresses this gap by focusing on *visual risk perception in everyday indoor environments*, where hazards depend on subtle object interactions and context. We propose a lightweight, onboard-friendly framework that integrates object-level risk estimation, semantic-spatial relationships, and asymmetric risk propagation through a scene graph. Unlike prior methods, our system does not rely on large symbolic models or handcrafted rules, inferring indirect risks (e.g., a bowl near a slippery surface) by holistically interpreting the scene.

To enhance robot-assisted safety monitoring, we propose a foundational algorithm for accident risk prediction in home environments. This system dynamically adjusts risk values based on environmental changes and personalizes assessments, reducing false positives and enabling robots to adapt to diverse contexts beyond immediate hazards.

We validate our method using a human-annotated dataset, achieving 75% binary risk detection accuracy. Our model's predictions align closely with human intuition, particularly in scenes with unstable placements, sharp tools, or cluttered environments. This work advances socially aware, perceptive robots by enabling them to understand risk as semantics-in-context, fostering safer and more intuitive human-robot coexistence.

2. Related Work

Recent advancements in robotics and assistive technologies have driven research toward improving safety and quality of life for elderly individuals and families with young children. While there has been a lot of studies involving robot safety and planning in factory environments, only a few studies have been done in home environments[17, 18]. With robots increasingly integrated into domestic spaces, safety-aware perception systems are becoming essential components of home robotics[19, 20].

Approaches like AnomalyGen [6] utilize large language models and 3D simulations for hazard detection. However, these systems often rely on static object-level detection, leading to high false positives when context is ignored, as human intuition about risk extends beyond isolated objects to semantic roles and spatial relationship

The SafetyDetect [21] employs scene graphs and LLMs to infer risk from textual relationships derived from image annotations. While semantically promising, its reliance on textual conversion rather than direct visual input can lose critical visual details for subtle risk interpretation.

Other efforts, such as Chen et al. [22]and EARBench [15], focus on visual grounding and multi-modal information fusion to improve obstacle detection and assess physical risk. These works highlight the

Table 1
Comparison of existing approaches in household accident risk assessment.

Method	Visual Input	Risk Adjustment	Context Awareness	Real-world Data
AnomalyGen [6]	✓	✗	✗	✓
SafetyDetect [21]	✗	✓	✗	✓
EARBench [15]	✓	✗	✓	✗
Ours	✓	✓	✓	✓

need for more grounded and interpretable solutions in real-world risk recognition and reaction.

In contrast, our approach directly utilizes visual (RGB-D) data for object detection and accident risk inference, bypassing intermediate text representations to capture nuanced spatial-semantic cues more robustly. Unlike prior methods reliant on language models or symbolic reasoning, our system integrates real-world accident statistics to derive data-driven risk scores. This allows for more intuitive and contextually grounded risk estimation, aligning closely with human risk perception, as summarized in Table 1

3. Methodlogy

3.1. Risk Estimation

We introduce the concept of a *Risk Score* to quantitatively evaluate the likelihood of household accidents. This score serves as an indicator of how prone a specific object is to causing a particular type of accident. It is computed using real-world statistical data, providing a robust and objective baseline for our model.

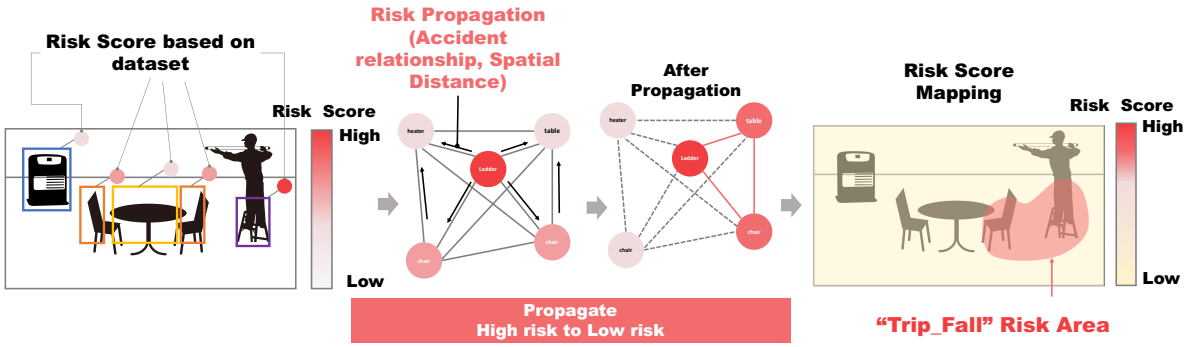

Figure 2: Illustration of risk propagation in a household environment. This figure shows how initial object risk scores (left) propagate through spatial and semantic relationships (center) to generate a final risk map, highlighting amplified risk areas (right).

To derive these scores, we use the Accident Information Database System[23] provided by the Consumer Affairs Agency and the National Consumer Affairs Center of Japan. This web-based database contains approximately 400,000 reports of actual accidents, detailing object involvement, accident causes, and damage outcomes. We query the database using both single-object terms (e.g., *Heater*) and object-pair combinations (e.g., *Towel near Heater*). The single-object queries inform individual risk scores, while pairwise queries capture contextual accident correlations between co-occurring objects.

The risk score $R(o, a)$ for object o and accident type a is initially a raw ratio based on accident counts. To ensure stability for objects with few reports, we apply Laplace-style Bayesian smoothing using Equation1:

$$R(o, a) = \frac{\text{count}(o, a) + k}{\text{total}(o) + k \cdot N} \tag{1}$$

Here, $\text{count}(o, a)$ is the number of accidents of types a involving objects o, $\text{total}(o)$ is the total reports involving o, N is the total accident types, and k is a smoothing parameter (e.g., $k = 1$). This method

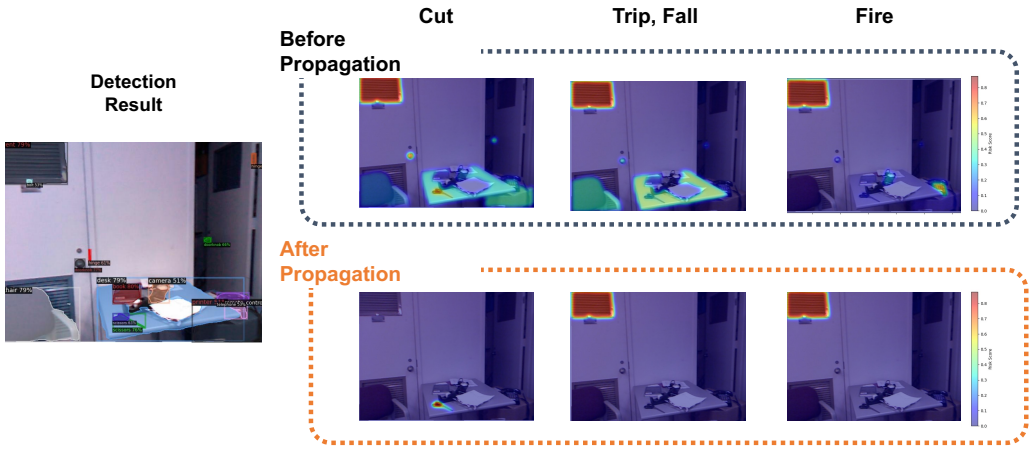

Figure 3: Comparison of risk score maps before (top row) and after (bottom row) risk propagation for Cut, Trip/Fall, and Fire accident types. The scene depicts a room with a table, where scissors are placed. In the Cut category, the initial assessment (top row) shows high risk for the scissors but also erroneously for a broader table area. After propagation (bottom row), the high risk is appropriately localized to the scissors only. Conversely, in Trip/Fall and Fire scenarios, risk scores remain largely unchanged.

allows systematic evaluation of object risk for specific accident categories, focusing on cut injuries, fire-related accidents, and falls in this study.

3.2. Risk Propagation Algorithm

In household scenarios, an object's danger often arises from its interaction with nearby items (e.g., a towel near a stove increases fire risk). To emulate this human intuition for context-sensitive hazards, we propose an asymmetric Risk Propagation Algorithm based on an object-centric scene graph. In this graph, each object is a node with an initial risk score, iteratively updated by considering semantic relationships and spatial proximity to neighboring objects.

The propagation weight w_{ij} between object i and its neighbor j is defined as:

$$w_{ij} = \phi_{\text{accrel}}[i,j] \cdot (1 - \phi_{\text{distance}}[i,j]) \cdot \max(0, r_j - r_i) \tag{2}$$

where r_i and r_j are current risk values. Risk propagates unidirectionally from higher-risk objects to lower-risk ones($r_j > r_i$), reflecting the intuition that dangerous objects influence their surroundings. ϕ_{accrel}, the Accident-Related Correlation, quantifies how strongly two objects interact to cause accidents, capturing context-dependent interactions (e.g., a stuffed toy on a stove significantly increases fire risk). ϕ_{accrel} is calculated using:

$$\phi_{\text{accrel}}(o_1, o_2, a) = \frac{\text{count}(o_1, o_2, a) + k}{\text{total}(o_1, o_2) + k \cdot N} \tag{3}$$

This value indicated accident correlation, enhancing risk propagation.

As shown in Fig. 2, the process begins with analyzing RGB-D images to detect objects and obtain their three-dimensional positions. Object detection is performed using Detic [24], and the centroid coordinates of the bounding box are used as the object's location. Each object is initially assigned a risk score $R(o, a)$, and a scene graph is constructed using spatial and semantic edges.

To spatially visualize risk scores, heatmaps are generated over RGB images using object detection results and segmentation masks. The process involves:

- Assigning risk scores to pixels within the detected object masks.
- Applying Gaussian filters ($\sigma = 3$) to smooth abrupt transitions between high-risk and low-risk areas.

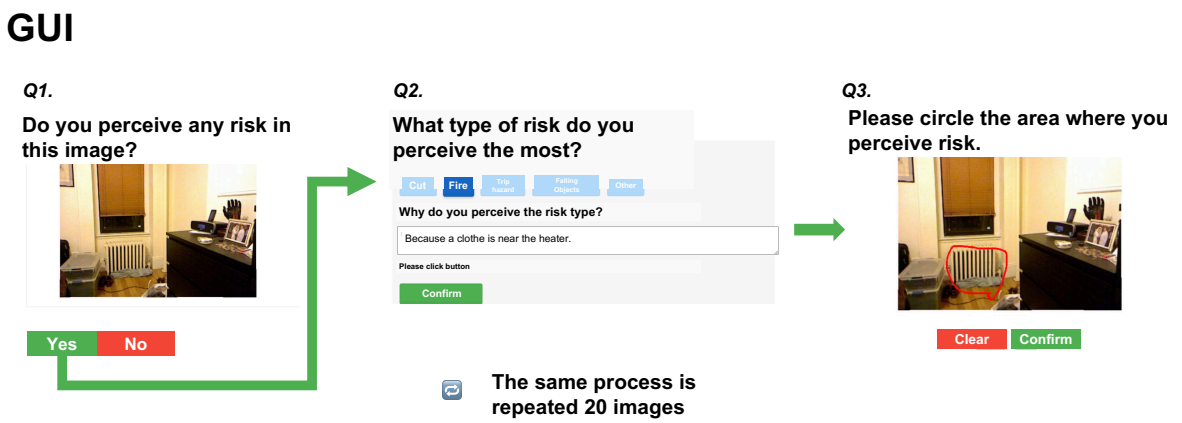

Figure 4: Evaluation GUI Flow

- In the heatmap visualization, red indicates high-risk areas, while blue denotes low-risk areas.
- Normalizing risk scores to the range $[0, 1]$ for interpretability.

In the Cut category, the scenario presented in Figure 3 highlights a critical refinement by our propagation model. Initially, the scissors on the table correctly receive a high risk score, as depicted in the top row. However, a broader area of the table is also erroneously attributed a high risk. Ideally, the high risk should be localized solely to the hazardous object itself. Through risk propagation, shown in the bottom row, the risk score for the surrounding table area is effectively reduced, and the high risk is appropriately highlighted for the scissors only. This demonstrates our model's selectivity and effectiveness in correcting initial overestimations. Conversely, in Trip/Fall and Fire scenarios, where strong contextual correlations for significant changes are absent in this instance, risk scores remain largely unchanged, further emphasizing the model's selective propagation.

4. Experimental Results

To validate our proposed risk propagation algorithm, we conducted a two-stage user study assessing its alignment with human perception in detecting and classifying household accident risks. The evaluation utilized the NYU Depth V2 dataset [25], a publicly available resource providing diverse RGB-D imagery of domestic scenes. This allowed us to benchmark the system's core reasoning capabilities in a controlled setting, minimizing variability from robot-collected data. The study aimed to address two key research questions:

- Can the system accurately identify whether a household scene contains potential accident risks?
- Can it correctly classify the identified risk into one of the following categories: cut, fire, or tirp/fall?

To collect reference data, we developed a web-based GUI (Fig. 4) to gather human annotations from 14 participants on 20 NYU RGB-D images. Participants were asked to: (1) indicate perceived accident risk, (2) select the risk type (cut, fire, or fall), and (3) annotate hazardous regions. This yielded three types of ground-truth data: binary risk (present/absent via majority vote), risk type labels (cut, fire, fall, or none), and averaged heatmaps of perceived risk areas. These annotations served as a human-grounded reference for evaluating the algorithm's semantic accuracy (correct risk type identification) and semantic alignment (correspondence between generated heatmaps and human-marked hazardous areas).

For both the algorithm and the human annotations, the risk maps were normalized to the range $[0, 1]$ for consistency.

Although this evaluation was conducted using a benchmark dataset, it represents a critical first step in validating our system's real-world reasoning capabilities. The use of a well-curated dataset enabled

controlled comparisons and reproducibility while removing confounding factors such as sensor noise or robot navigation limitations.

A follow-up study is planned to deploy the system on a real robot equipped with an RGB-D camera to assess its effectiveness in dynamic, cluttered, and potentially unstructured household environments. In such future settings, the proposed risk propagation mechanism can serve not only to detect risky configurations in real time but also to proactively guide the robot in mitigating those risks—laying the groundwork for safe and intelligent robotic assistance in domestic contexts.

4.1. Evaluation of Risk Detection Accuracy

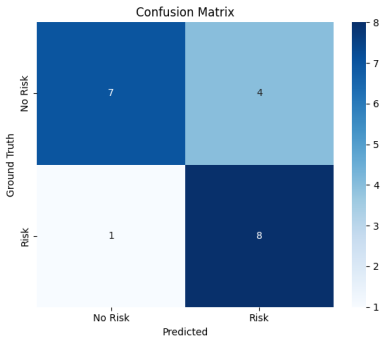

Figure 5: Risk Detection Confusion Matrix

We assessed the system's ability to identify accident risk using a binary classification evaluation against human annotations. Each test image was labeled as 'Risk' or 'No Risk' based on participant consensus. System predictions were determined by whether algorithm-generated heatmaps (for cut, fire, or trip/fall categories) exceeded a predefined threshold.

The confusion matrix (Figure 5) for this binary classification task shows the proposed system achieved an overall accuracy of 75% in identifying the presence or absence of risk. A notably low false negative rate is critical for safety-critical applications. For this evaluation, predictions were based solely on heatmaps corresponding to the annotated risk type, ensuring consistent assessment aligned with human perception.

It is important to note that for this evaluation, predictions were made using only the heatmap corresponding to the risk type annotated for each image. This risk-type-specific evaluation avoids artificial performance inflation that could arise from combining unrelated heatmaps and allows a more consistent assessment of the system's alignment with human risk perception.

4.2. Evaluation of Risk Type Classification Accuracy

Beyond risk detection, we evaluated the system's ability to classify perceived risk types (cut, fire, or trip/fall hazard) and their spatial alignment. For spatial assessment, we employed a centroid-based evaluation metric, calculating the Euclidean distance d between predicted and ground truth heatmap centroids:

$$d = \|\mathbf{c}_{\text{pred}} - \mathbf{c}_{\text{gt}}\|_2, \tag{4}$$

where c_{pred} and c_{gt} are the centroid coordinates. This distance is interpreted as an accuracy-like measure$(1/d)$, useful for evaluating heatmaps that may vary in shape but indicate similar risk zones.

Figure 6 summarizes the risk classification results, demonstrating clear improvement after risk propagation, particularly in cut-related scenes where contextual reasoning enhanced detection of localized hazards. While also evaluating semantic classification accuracy using Intersection over Union (IoU), the centroid-based metric provided a complementary perspective for intuitive spatial reasoning.

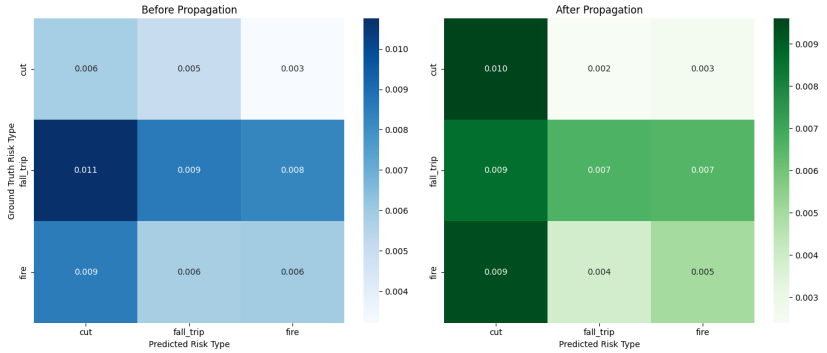

Figure 6: Confusion matrices showing risk type prediction results before (left) and after (right) propagation. The improvement is especially evident in cut-related scenarios.

Together, these evaluations offer a holistic view of the algorithm's ability to detect, classify, and spatially infer risks aligned with human judgment

5. Discussions

This study evaluated the feasibility of risk reasoning in domestic environments using statistical data and object relationships. While our approach demonstrated capability in identifying risky situations, several limitations impact its overall performance and robustness.

A primary cause of false positives is object over-detection; assigning risk scores to every detected item, even benign ones (e.g., a blanket detected independently from a bed), can lead to risk overestimation. This highlights a current limitation of object detectors in safety-critical applications.

Another challenge is the quality and interpretability of accident data from governmental databanks. Despite being statistically correct, some entries are counterintuitive (e.g., cuts involving 'blankets' may be indirectly caused by needles inside fabric), affecting the reliability of initial object risk scores

To mitigate these issues, a filtering mechanism based on accident rate thresholds could be beneficial. Objects accounting for less than a specified percentage (e.g., 0.5%) of reported accidents could be excluded from propagation, reducing noise and improving accuracy.

While our algorithm effectively identifies cut-related risks, its performance on fall-trip and fire scenarios is less consistent. This is likely because these risk types depend on object states (e.g., pose, occlusion for fall-trip; physical states like visible flames for fire) not explicitly represented in current features. Incorporating object-level state recognition (e.g., 'lying on the floor,' 'burning') could significantly improve detection accuracy for these categories.

The current framework, in its nascent stage, lacks comprehensive situational context and human activity. Future work will integrate detailed contextual information via VLMs and Semantic Maps to enable richer situational understanding and personalized risk assessment[26]. Multi-robot task planning and allocation with foundation modeling and probabilistic inferencing can also significantly improve the efficiency of search in complex environments[27, 28]. We plan to implement these frameworks in real robotic systems for quantitative evaluation of risk reduction, enhancing safety and reliability in human-robot coexistence.

Future work will prioritize improving initial object risk estimation by integrating higher-quality accident data and contextual and developing adaptive thresholding methods to reduce false alarms effectively.

6. Conclusions

In this paper, we introduced a novel approach to risk reasoning by modeling accident-prone scenes as relational graphs and propagating risk asymmetrically based on contextual cues. Our method restricts

propagation to flow from higher-risk to lower-risk objects, guided by accident relevance and distance, ensuring stable convergence and interpretable risk dynamics.

Our model achieved 75% binary classification accuracy against user-annotated risk maps, correctly distinguishing risk-present from non-risk scenes. Risk propagation improved spatial accuracy in cut-related scenes, though gains were limited for fire and fall/trip categories. These findings suggest context-aware propagation meaningfully enhances risk estimation where individual object risk is insufficient.

This work lays the groundwork for real-world deployment, designing the framework for integration with robotic systems equipped with RGB-D sensors, despite initial evaluation on static benchmark images. Future work will focus on real-time deployment, enabling robots to proactively respond to risks in dynamic home environments.

Future extensions aim to improve robustness by integrating temporal reasoning, multimodal perception, and richer object state representations. A key development area is integrating human presence and behavior into risk reasoning, as many hazards are only truly risky when people are nearby. Distinguishing between static and person-dependent risks will be essential for reducing false positives and increasing real-world usability. The framework is generalizable to broader household risks beyond the three common types studied.

Ultimately, we envision extending this system into a personalized, household-specific risk mapping framework that adapts to user lifestyles and evolving home conditions. Such maps can guide robot behavior, prioritizing high-risk areas, enabling autonomous patrols, and providing timely, personalized support to address risks in daily life.

Acknowledgments

This work was partially supported by JST Moonshot R&D [Grant Number JPMJMS2034] and JSPS Kakenhi [Grant Number JP24K07399].

References

[1] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter, et al., pi_0: A vision-language-action flow model for general robot control, arXiv preprint arXiv:2410.24164 (2024).

[2] A. A. Ravankar, S. A. Tafrishi, J. V. Salazar Luces, F. Seto, Y. Hirata, Care: Cooperation of ai robot enablers to create a vibrant society, IEEE Robotics Automation Magazine 30 (2023) 8–23. doi:10.1109/MRA.2022.3223256.

[3] A. Chikhalikar, A. A. Ravankar, J. V. S. Luces, Y. Hirata, Enhancing object search in indoor spaces via personalized object-factored ontologies, in: 2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), IEEE, 2025, pp. 19210–19217.

[4] Ministry of Health, Labour and Welfare of Japan, Report on Accident Risk Assessment in Household Environments, Technical Report 0000047617, Ministry of Health, Labour and Welfare of Japan, 2013. URL: https://www.mhlw.go.jp/file/05-Shingikai-12201000-Shakaiengokyokushougaihokenfukushibu-Kikakuka/0000047617.pdf, [Online; accessed March 22, 2025].

[5] Ministry of Health, Labour and Welfare of Japan, Vital Statistics of Japan, 2023, Technical Report 15_all, Ministry of Health, Labour and Welfare of Japan, 2023. URL: https://www.mhlw.go.jp/toukei/saikin/hw/jinkou/kakutei23/dl/15_all.pdf, [Online; accessed March 22, 2025].

[6] Z. Song, G. Ouyang, M. Fang, H. Na, Z. Shi, Z. Chen, Y. Fu, Z. Zhang, S. Jiang, M. Fang, et al., Hazards in daily life? enabling robots to proactively detect and resolve anomalies, arXiv preprint arXiv:2411.00781 (2024).

[7] B. Li, J. Wang, X. Chen, Trg-planner: Traversal risk graph-based path planning in unstructured environments, arXiv preprint arXiv:2501.01806 (2024).

[8] S. Wang, Y. Kobayashi, A. A. Ravankar, A. Ravankar, T. Emaru, A novel approach for lidar-based robot localization in a scale-drifted map constructed using monocular slam, Sensors 19 (2019) 2230.

[9] A. Ravankar, A. Ravankar, Y. Hoshino, Y. Kobayashi, Virtual obstacles for safe mobile robot navigation, in: 2019 8th international congress on advanced applied informatics (IIAI-AAI), IEEE, 2019, pp. 552–555.

[10] A. Georgiev, A. Mihailidis, Robotics-enabled in-home environment screening for fall risks, ACM Transactions on Accessible Computing 14 (2021) 1–25.
[11] H. Alemzadeh, V. Raman, C. Li, K. Glisson, Real-time context-aware detection of unsafe events in robot-assisted surgery, in: Proceedings of the 2020 IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), IEEE, 2020, pp. 25–37.
[12] F. Dellaert, M. Kaess, Factor graphs for robot perception, Foundations and Trends® in Robotics 6 (2017) 1–139.
[13] Y. Chen, X. Yan, J. Liu, F.-Y. Wang, Multi-robot collaborative perception with graph neural networks, arXiv preprint arXiv:2201.01760 (2022).
[14] A. Chikhalikar, A. A. Ravankar, J. Victorio Salazar Luces, Y. Hirata, Semantic-based multi-object search optimization in service robots using probabilistic and contextual priors, IEEE Access 12 (2024) 113151–113164. doi:10.1109/ACCESS.2024.3444478.
[15] Z. Zhu, B. Wu, Z. Zhang, L. Han, Q. Liu, B. Wu, Earbench: Towards evaluating physical risk awareness for task planning of foundation model-based embodied ai agents, arXiv preprint arXiv:2408.04449 (2024).
[16] A. Chikhalikar, A. A. Ravankar, J. V. S. Luces, Y. Hirata, Integrating semantic awareness and probabilistic priors for object search in indoor environments, in: The Proceedings of JSME annual Conference on Robotics and Mechatronics (Robomec) 2023, The Japan Society of Mechanical Engineers, 2023, pp. 1P1–C25.
[17] Y. Tsushima, S. Yamamoto, A. A. Ravankar, J. V. S. Luces, Y. Hirata, Task planning for a factory robot using large language model, IEEE Robotics and Automation Letters 10 (2025) 2383–2390.
[18] A. Ravankar, A. A. Ravankar, Y. Kobayashi, T. Emaru, Path smoothing extension for various robot path planners, in: 2016 16th international conference on control, automation and systems (ICCAS), IEEE, 2016, pp. 263–268.
[19] Y. Hirata, J. V. S. Luces, A. A. Ravankar, S. A. Tafrishi, Cooperation of assistive robots to improve productivity in the nursing care field, in: The International Symposium of Robotics Research, Springer, 2022, pp. 287–294.
[20] S. Coşar, M. Fernandez-Carmona, R. Agrigoroaie, J. Pages, F. Ferland, F. Zhao, S. Yue, N. Bellotto, A. Tapus, Enrichme: Perception and interaction of an assistive robot for the elderly at home, International Journal of Social Robotics 12 (2020) 779–805.
[21] J. F. Mullen, P. Goyal, R. Piramuthu, M. Johnston, D. Manocha, R. Ghanadan, "don't forget to put the milk back!" dataset for enabling embodied agents to detect anomalous situations, IEEE Robotics and Automation Letters 9 (2024) 9087–9094. doi:10.1109/LRA.2024.3430129.
[22] R. Chen, Y. Li, J. Wang, H. Liu, M. Li, Deep-learning-based context-aware multi-level information fusion framework for indoor mobile robots, Sensors 23 (2023) 996.
[23] Consumer Affairs Agency, Accident Information Data Bank System, https://www.jikojoho.caa.go.jp/ai-national/, 2024. Accessed: Feb. 11, 2025.
[24] X. Zhou, R. Girdhar, A. Joulin, P. Krähenbühl, I. Misra, Detecting twenty-thousand classes using image-level supervision, in: European conference on computer vision, Springer, 2022, pp. 350–368.
[25] P. K. Nathan Silberman, Derek Hoiem, R. Fergus, Indoor segmentation and support inference from rgbd images, in: ECCV, 2012.
[26] A. Chikhalikar, A. A. Ravankar, J. V. S. Luces, Y. Hirata, Open vocabulary object search utilizing large language models and fuzzy inferencing, in: 2025 IEEE/SICE International Symposium on System Integration (SII), 2025, pp. 345–351. doi:10.1109/SII59315.2025.10870891.
[27] A. Ravankar, A. A. Ravankar, Y. Kobayashi, Y. Hoshino, C.-C. Peng, M. Watanabe, Hitchhiking based symbiotic multi-robot navigation in sensor networks, Robotics 7 (2018) 37.
[28] A. Ravankar, A. A. Ravankar, Y. Kobayashi, T. Emaru, Avoiding blind leading the blind: Uncertainty integration in virtual pheromone deposition by robots, International journal of advanced robotic systems 13 (2016) 1729881416666088.