OVER DATA.

Research Project

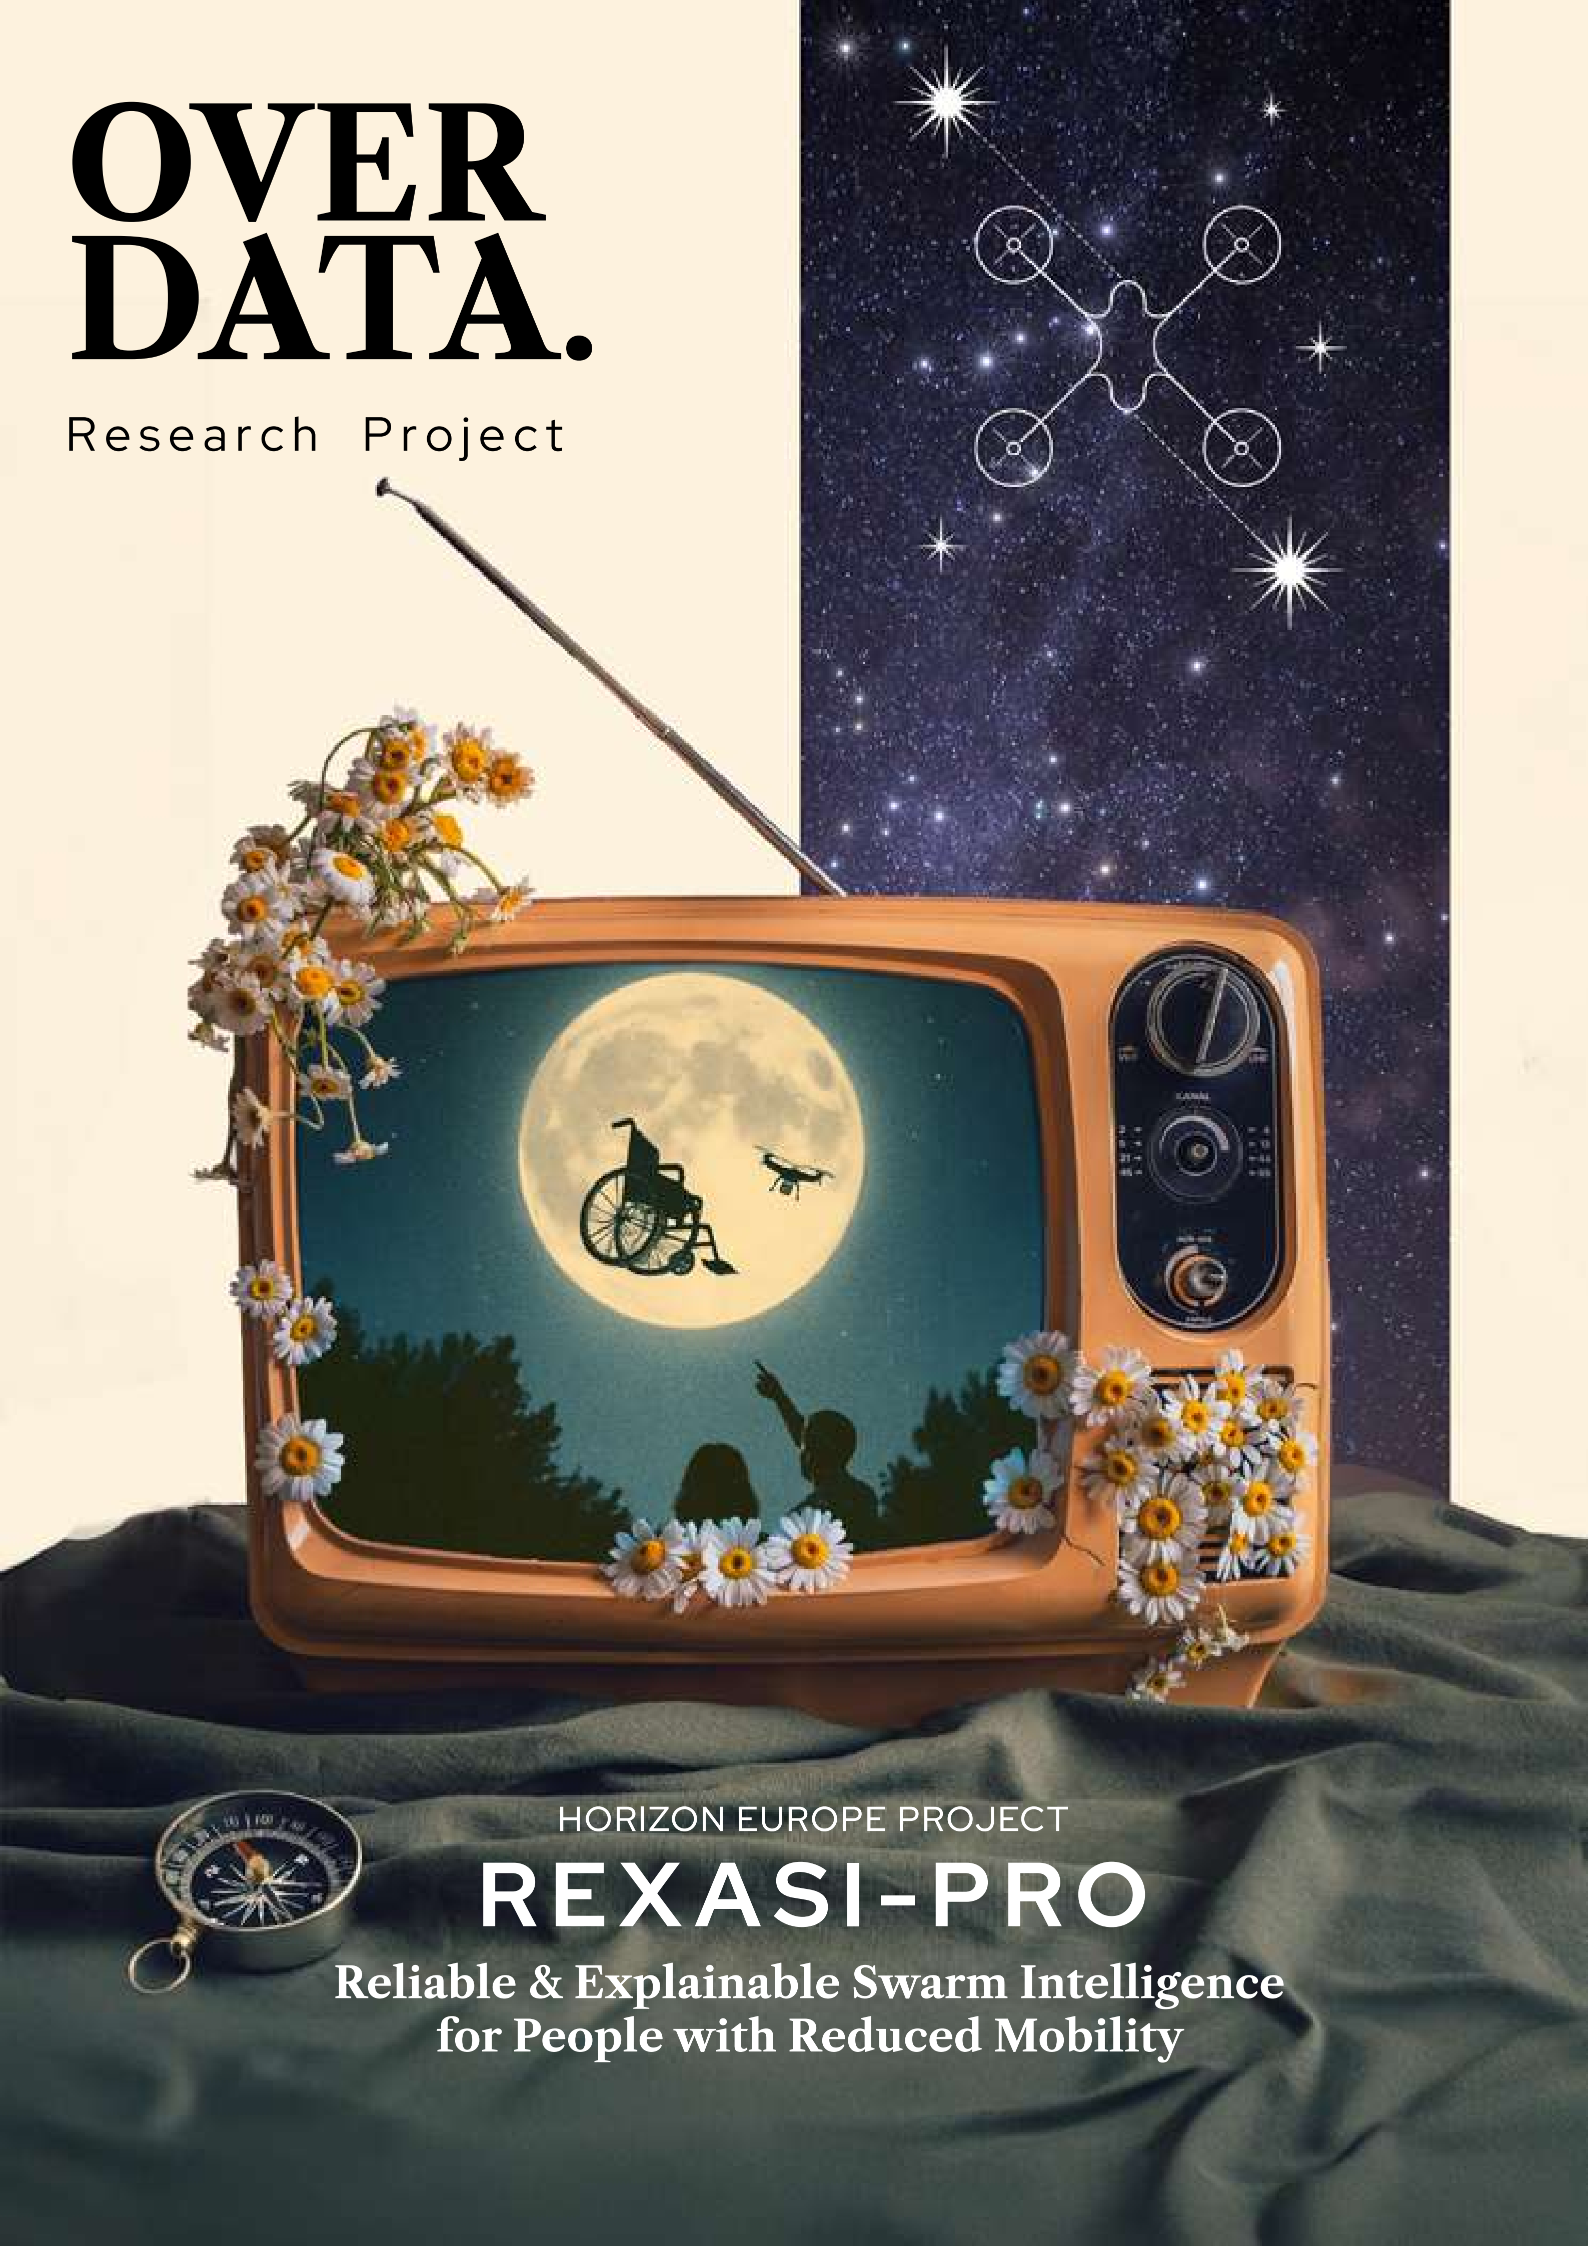

HORIZON EUROPE PROJECT

REXASI-PRO

Reliable & Explainable Swarm Intelligence for People with Reduced Mobility

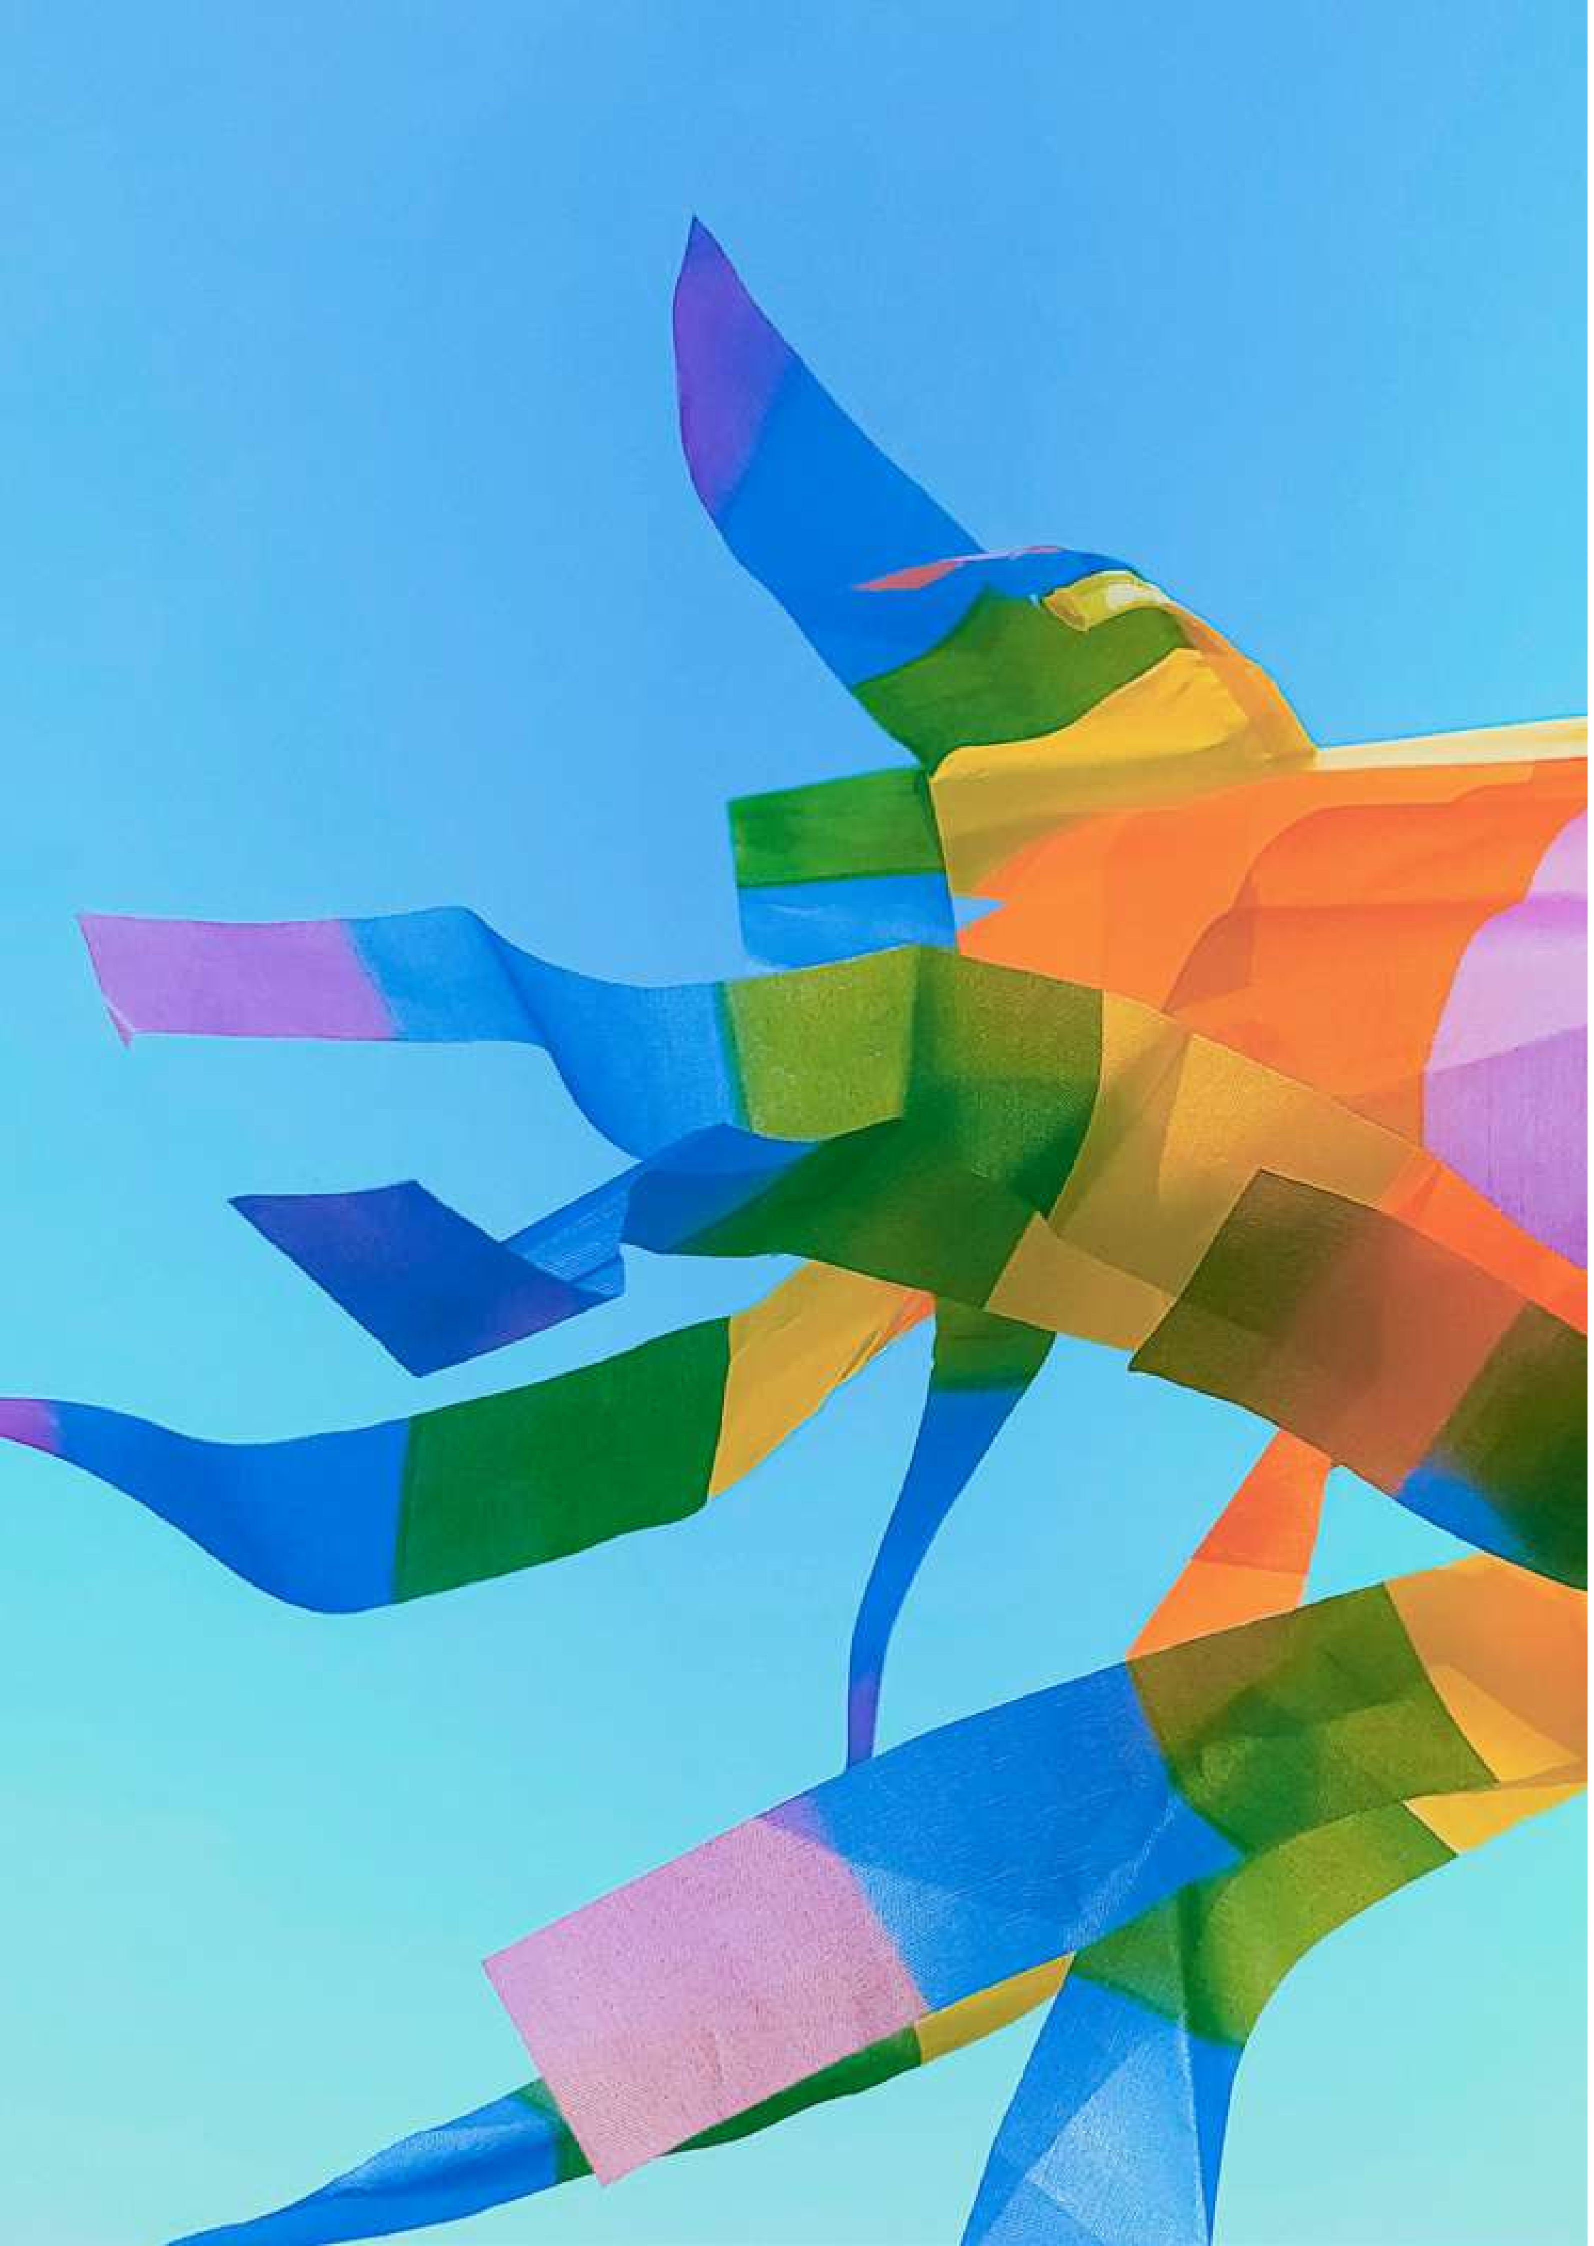

INDEX

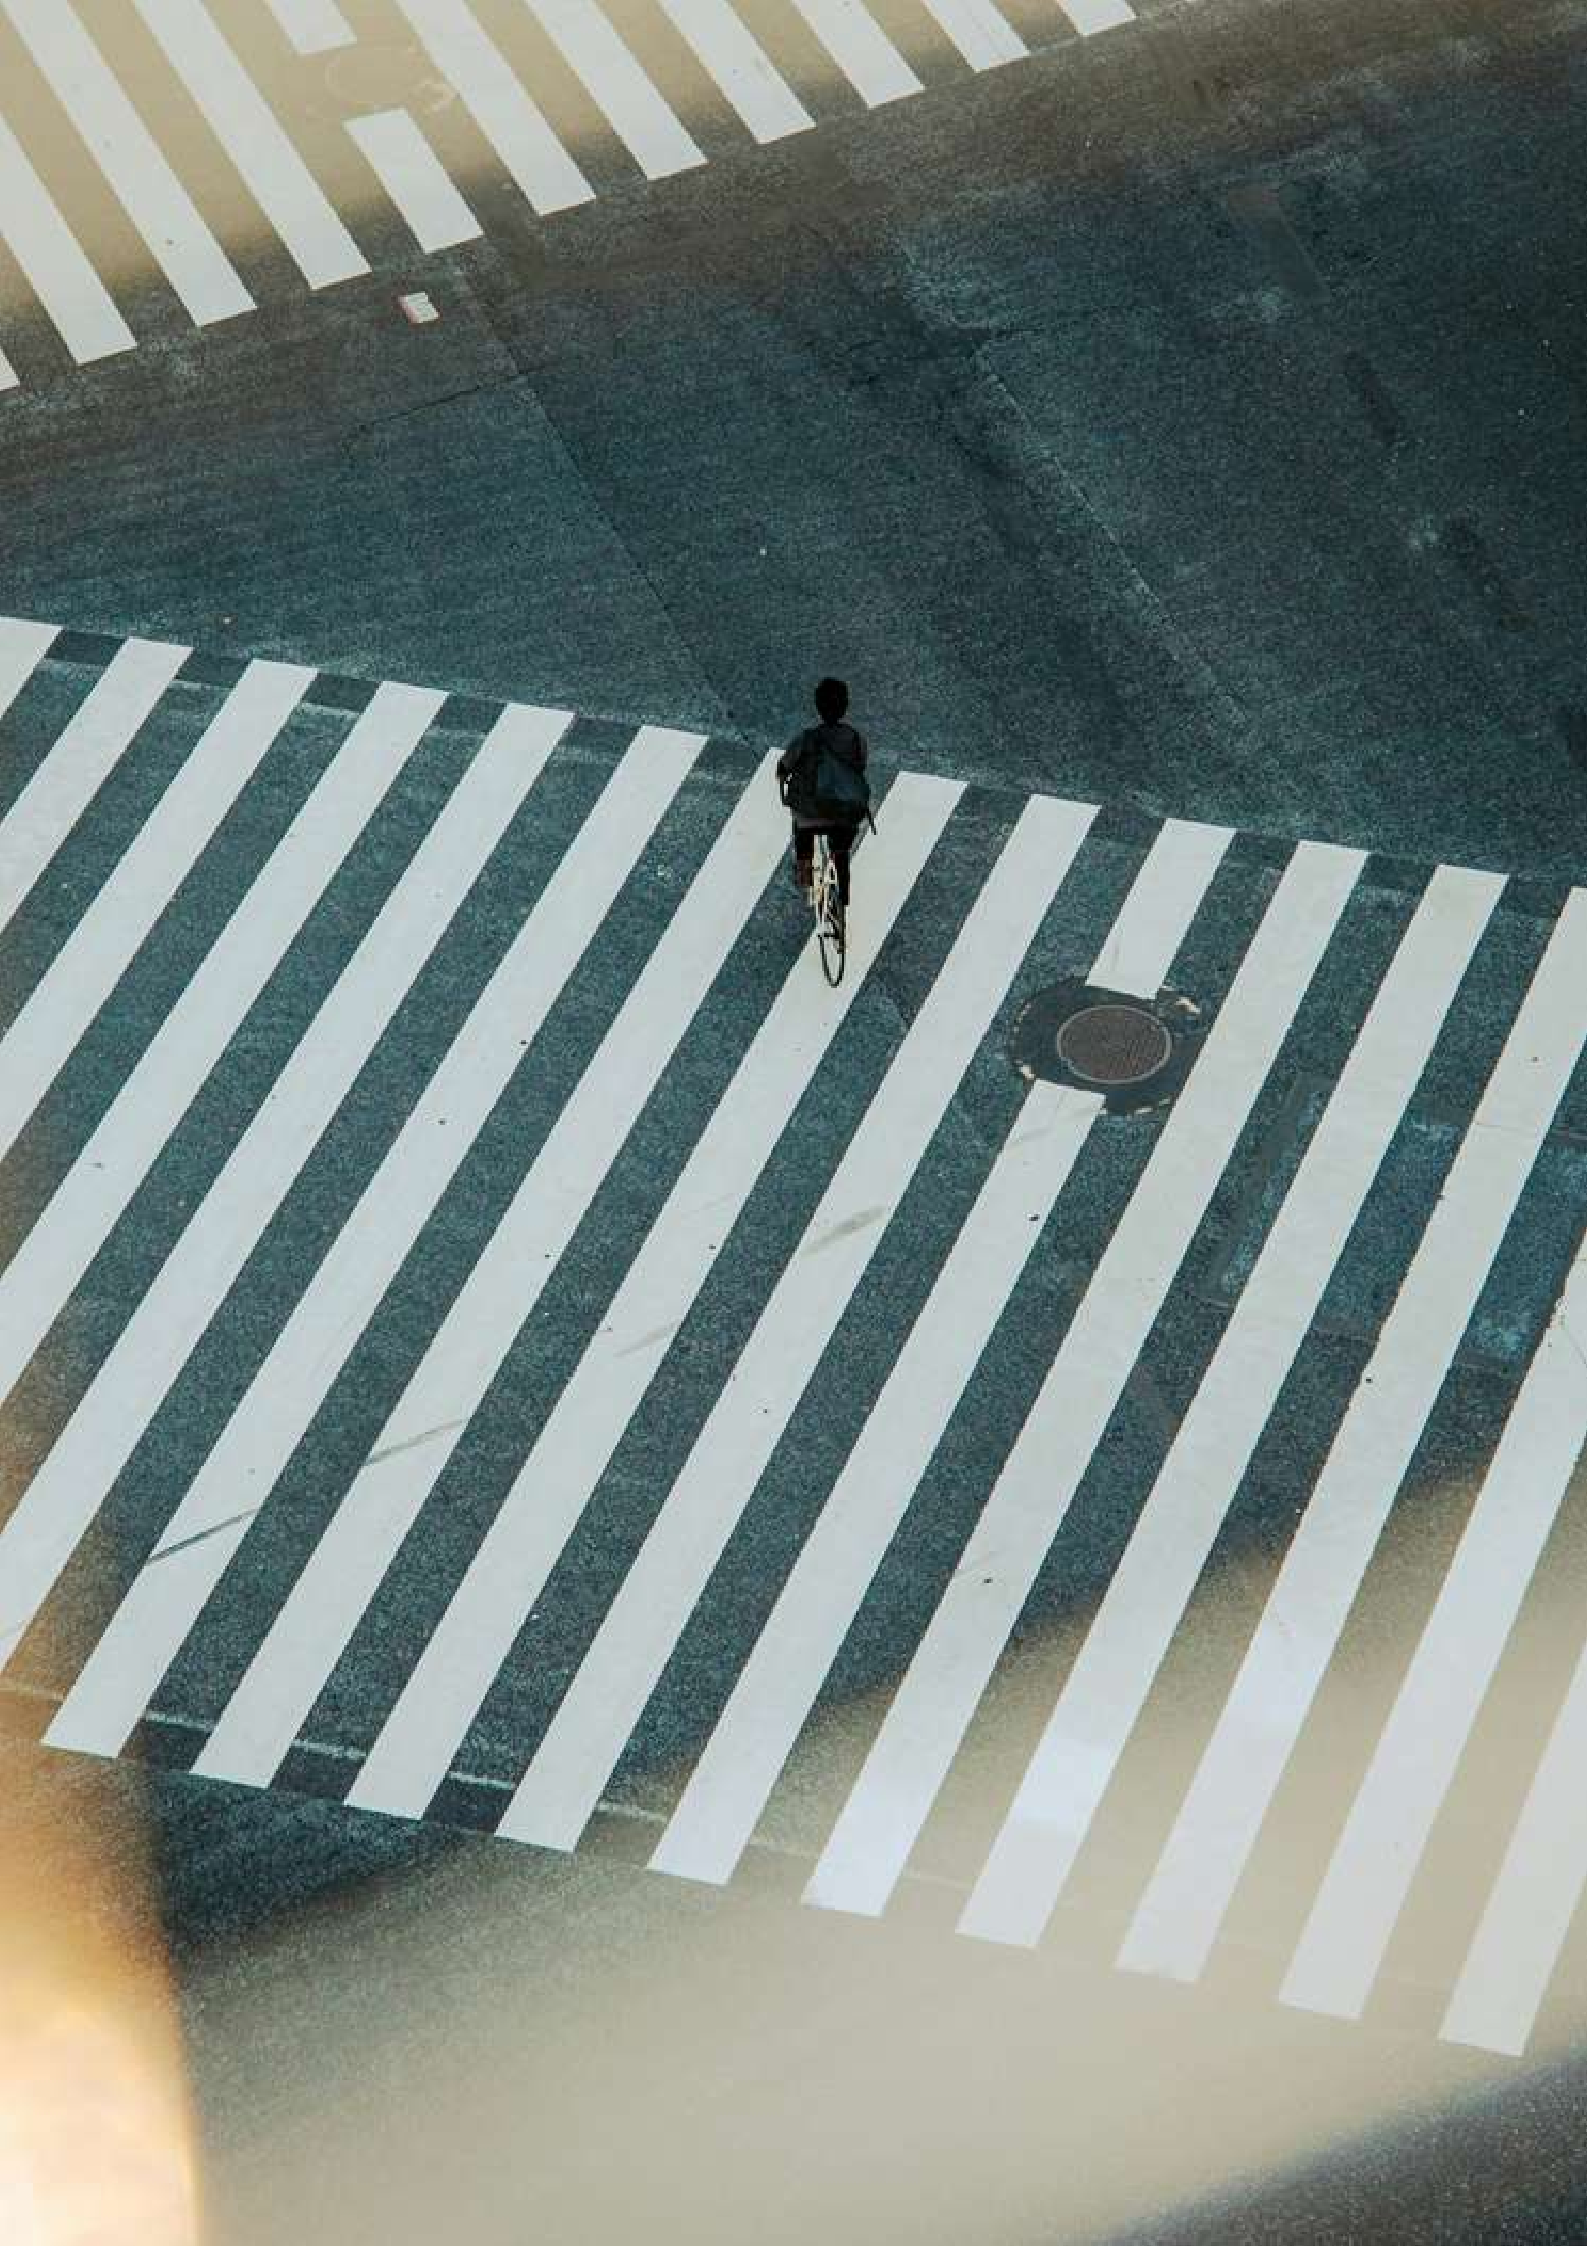

REXASI-PRO

Innovating Mobility Through Trustworthy AI

Reliable & Explainable Swarm Intelligence for People with Reduced Mobility

Introduction to the REXASI-PRO Project

The REXASI-PRO project is pushing the boundaries of **artificial intelligence** and **mobility by creating a more sustainable, trustworthy, and efficient system for people with reduced mobility.** Funded by the EU under Horizon Europe, the project aims to develop groundbreaking AI solutions that are "trustworthy-by-construction." This innovative approach will ensure that autonomous vehicles, like wheelchairs and flying robots, operate with the highest levels of safety and reliability. By combining social navigation algorithms and cutting-edge AI models, REXASI-PRO envisions a future where AI not only adapts to dynamic environments but also communicates effectively with its surroundings to ensure smooth, safe mobility for its users.

The project takes a multi-faceted approach to ensure the robustness of these AI systems. From **autonomous wheelchairs navigating through crowded spaces to flying robots assisting in emergency scenarios**, REXASI-PRO's framework integrates advanced machine learning techniques with a focus on **green AI solutions**. Through energy-efficient methods and a novel certification process, the framework will set new standards for AI-powered mobility solutions, helping improve the independence and safety of individuals with reduced mobility. With its commitment to innovation and ethical considerations, REXASI-PRO is transforming the way we think about mobility, autonomy, and the future of AI.

> **"The project aims to develop groundbreaking AI solutions that are 'trustworthy-by-construction.' This approach will ensure that autonomous vehicles, like wheelchairs and flying robots, operate with the highest levels of safety and reliability."**

This magazine provides an in-depth look at the **various contributions and milestones achieved by the partners involved** in REXASI-PRO. It includes detailed sections on the development of AI algorithms for both wheelchairs and drones, insights from leading institutions on safety, ethical considerations, and the environmental impact of these technologies. Featured articles dive into the technical, ethical, and societal challenges faced during the project, alongside practical demonstrations of how REXASI-PRO's innovations are shaping a more inclusive future. The structure of the magazine offers a comprehensive view of the project's scope, with contributions from project partners, highlighting the collaborative effort to create smarter, greener AI solutions for autonomous mobility.

REXASI
PRO

artificial intelligence by aitek

Aitek, with over 30 years of experience in intelligent systems for transport, security, and communication, will play a key role in the project by defining requirements for an innovative self-piloting wheelchair. Leveraging expertise in Deep Learning and Cyber-Physical Systems, Aitek will integrate sensors, drones, and cameras to realize this groundbreaking solution.

Aitek contribution

1.1 Aitek role in video analysis

Aitek in REXASI-PRO provides its central contribution in the video monitoring system installed at the CNR premises. Cameras are able to detect specific events and alert the orchestrator. Cameras were strategically positioned in certain points to comprehensively monitor wheelchair-accessible routes. The video monitoring system detects in real-time events through Deep Learning (DL) techniques, such as YOLO [1], carefully crafted and optimized to fulfill project's requirements and environment (i.e., fine-tuning with new collected data, adding pre-processing and post-processing layers, ensembling with other architectures). These types of Convolutional Neural Networks (CNNs) are able to classify different types of objects and detect their position in the image, determining the class and the bounding box of the objects of interest. The open-source models of YOLOv5 (https://docs.ultralytics.com/) range from nano to extra-large, with the nano model having the best performance in terms of frames per second and the extra-large obtaining the best accuracy value, measured by metrics such as mean average precision (mAP). In the task of object detection, the YOLO model is on topo of Single Shot Detectors (SSDs), those models capable of detecting both the class and position of a given object in a single shot. YOLO is not the only model belonging to the SSD class, but it is generally more efficient in terms of speed and accuracy [7]. In this specific video monitoring system, YOLO was trained both on open-source datasets available on the web and on custom data collected at the CNR premises.

The classes of interest are people and wheelchair, while the identify events are:

- Crowd in a particular area of interest. In particular, the level of crowd detected in a specific area that could change the path of the wheelchair.
- Stop of the wheelchair. If the wheelchair stays still for more than a particular time selected, for instance 10 seconds. This type of information could be useful to act quickly during a possible malfunction.
- Direction of a single person. For instance, if a person is going toward the wheelchair in opposite direction in a corner, we can notify the potential risk to the orchestrator and thus to the wheelchair.
- Running person. For instance, a person who is running in the corridor going in the wheelchair path and possibly met the wheelchair in a blind spot.

Identified the classes and the events, the video monitoring system notifies the orchestrator through a MQTT messages.

Figure 1 - Examples of framing and event detection through the smart monitoring system.

1.2 Verification and Validation framework

The objective of the REXASI-PRO V&V framework is to provide a structured and comprehensive approach to validate and verify the AI components in REXASI-PRO. By integrating simulation, explainability, reliability and cybersecurity, we aim to create a robust V&V process that guarantee the safety of the AI system. These aspects were introduced in Deliverable D2.3 in the requirements definition phase, and can be summarized in:

- Simulation: Simulation plays a crucial role in the validation process, allowing for testing AI systems in a controlled environment that replicates real-world conditions. Through extensive simulations, we can evaluate system's performance across a wide range of scenarios, including edge cases that might not be easily testable in the real world, such as the so called "black swans [8]".
- XAI: Explainability is an important feature of trustworthy AI. This framework emphasizes how the use of explainable AI models produces an AI system which is not only effective but also transparent.
- RAI: Reliability is essential in the definition of Trustworthy AI, especially in sa-

"The video monitoring system detects events in real time through Deep Learning, ensuring safer wheelchair navigation in dynamic environments."

fety-critical applications. This framework includes methodologies to ensure that AI models maintain consistent performance over time and across different operating conditions through the definition of probabilistic safety region.

By continuously looping between machine analysis and human insight through these statistical procedures, the W cycle in REXASI-PRO ensures transparency, adaptiveness, and a solution aligned with safety standards. This continuously process is applied in every step of the video monitoring system design and development:

- **Risk Analysis:** The process begins with a detailed analysis of the system's Operational Design Domain (ODD) and the identification of triggering conditions that could lead to unsafe behaviors. This phase results in the definition of safety goals for the video monitoring system, concrete countermeasures designed to mitigate potential failure modes reported in Deliverable D2.4.
- **Fair and ethical data collection:** The data collection phase includes a rigorous assessment of the dataset to ensure fairness, both in terms of class balance (e.g., person, wheelchair) and ethical representation (avoiding bias across demographics, gender or ethnic groups). This ensures that the training data reflects realistic and inclusive scenarios, forming a reliable foundation for the model.
- **Development and Training:** Multiple neural network models are trained and evaluated to determine the most robust architectures. During this phase, ensemble techniques or model combination strategies are applied to maximize accuracy and generalizability. This allows for redundancy and fault tolerance, aligning with the previously defined safety goals.
- **Deployment:** In the deployment phase, statistical methods such as conformal prediction, [6], are applied to monitor and validate the model's output in real time. These methods help assess whether predictions fall within the system's expected behavior boundaries, and trigger warnings when the model operates outside its validated ODD. This strengthens the video monitoring system's reliability in dynamic environments.

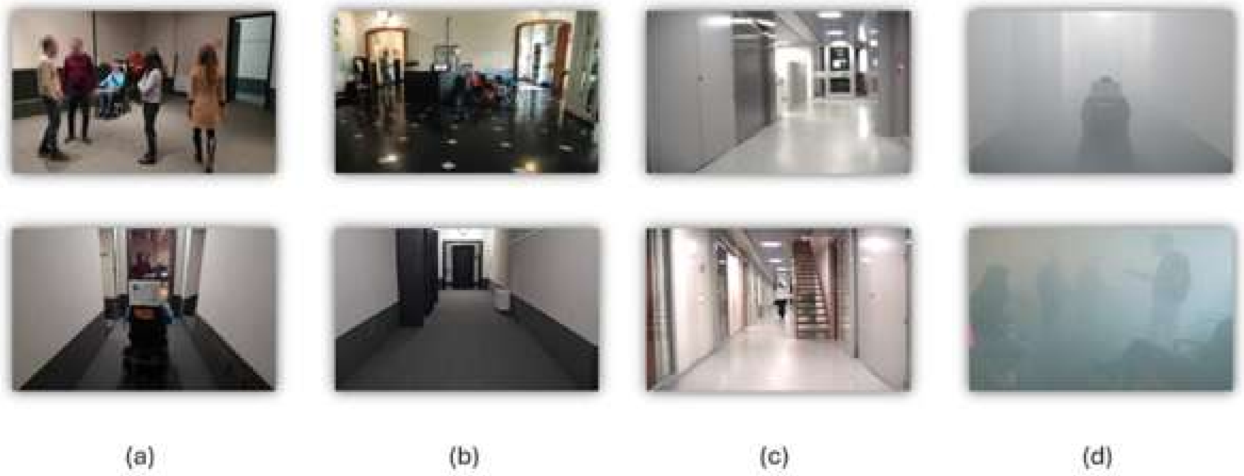

Figure 2 – example of dangerous scenarios considered: (a) risky scenes with obstacles or high crowd density; (b) safe scenes with clear forward paths; (c) safe scenes from the KTH-IDOL2 database; (d) perturbed scenes with fog or smoke, labeled as risky due to impaired visibility.

The W cycle thus reflects a complete loop, from proactive risk identification to ethical data handling, model refinement, and real-time statistical assurance. It ensures that the video monitoring component of REXASI-PRO is not only technically sound but also ethically aligned and safety-focused, contributing to the project's overarching goal of building Trustworthy AI solution.

1.2.1 Gray swan and scenario settings

This Section presents a benchmarking study of Vision-Language Models (VLMs) to support semantic scene understanding and safety assessment in autonomous wheelchair navigation. The evaluation is based on an indoor dataset recorded at the CNR Villa in Genoa, extended with synthetically perturbed images that simulate visual degradation (e.g., fog, smoke, reduced illumination). The safety classification task addresses four critical risk factors relevant to indoor assistive navigation: immediate obstacles, crowd density, visibility degradation, and fall risks due to stair-related hazards. The results highlight key trade-offs between reasoning power, inference speed, and robustness under varied visual conditions. To support live operation in robotic navigation, a hybrid pipeline is discussed that integrates VLM checkpoints with traditional perception modules, achieving both performance and explainability at the edge. Vision-Language Models (VLMs) offer a compelling new paradigm for semantic scene understanding in robotic navigation. Unlike traditional perception systems that rely solely on object detection and depth estimation, VLMs combine visual and textual modalities to reason about entire scenes. This is particularly advantageous in assistive robotics, where static images of crowded corridors or partially visible obstacles may require contextual interpretation that goes beyond geometric data. Models such as CLIP [2], MiniGPT-4 [3], and GPT-4 with Vision (GPT-4V) [4] enable scene-level interpretation through either similarity-based matching or natural-language-conditioned reasoning. Compared to traditional pipelines—which typically classify risks using a limited set of visual categories and predefined rules—VLMs offer more flexible, explainable, and abstract safety filtering. This section describes how each model was configured and evaluated to determine its effectiveness in classifying scenes as either "Safe to Proceed" or "Risky to Proceed". A key motivation for using VLMs in safety assessment is their ability to recognize unexpected or rare risk scenarios. This challenge is often conceptualized using the metaphor of "swans." Grey swans represent rare but partially understood risk conditions that can be uncovered through systematic performance evaluation. In contrast, black swans resist such treatment, as they stem from the fundamental incompleteness of human understanding—raising questions such as "when is enough, enough?" and pointing to the existence of risks that are "unknown, but unknowable" [5].

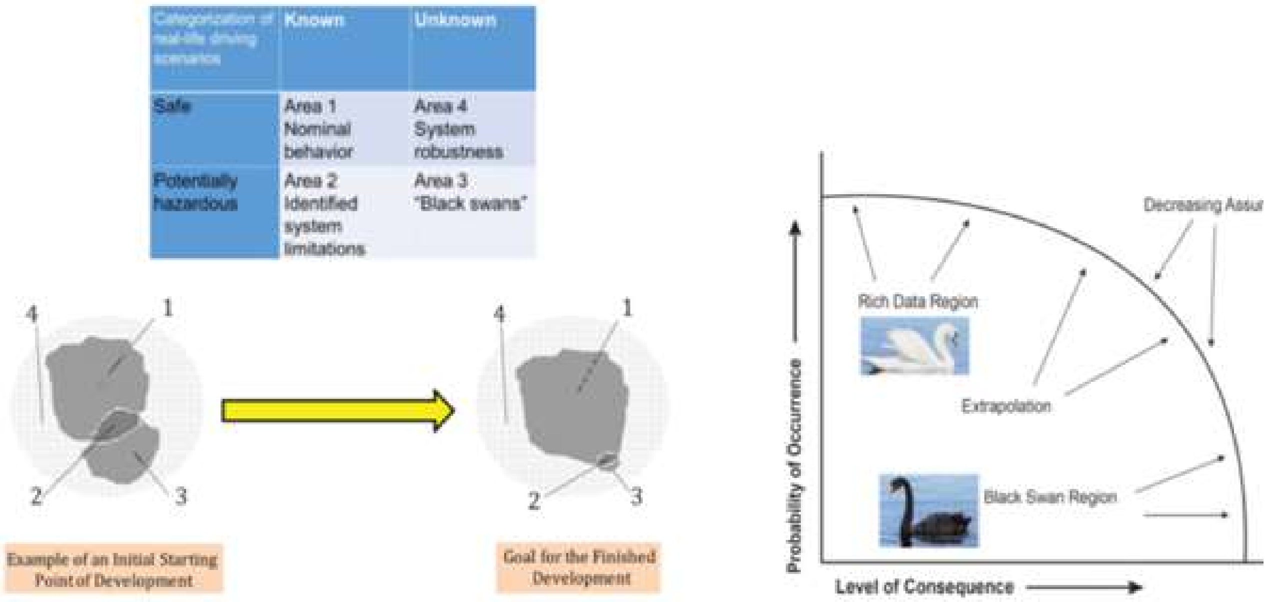

Figure 3: Swans in risk analysis. Grey swans refer to rare but partially understood risk conditions. Black swans represent fundamentally unknowable hazards that may remain invisible to structured evaluation frameworks.

Figure 3 illustrates this taxonomy. The VLM helps in the search of rare cases as it automatically analyses a large bunch of images, thus avoiding human brute-force checking. It comes out with unsafe conditions and explains inherent environmental conditions. In turn, this drives further testing of unforeseen conditions in the border between safety and unsafety. The final goal is to precisely specify requirements to avoid hazards, such as quantifying the danger of decrease of light at specific corners or increase of crowd density. The impact of a mix of those hazards on safety is difficult to understand. The VLM output for each image can be automatically parsed in this perspective so as to focus human attention to specific cases of interest with their explanations.

Figure 4 - Scene with no immediate obstacles or crowd, misclassified by all models.

Figure 4 shows a borderline case misclassified. It is a kind of "grey swan". The scene contains a pedestrian and a staircase at a short distance, but was annotated as "Safe to Proceed". The VLM misclassified it due to image resolution and ambiguous spatial cues, claiming the person to be too close to avoid, likely overestimating distance. How scenes like Figure 4 should be further treated is questionable. A specific training over an additional set of stairs around the environment would help increase driving awareness.

Figure 5 - Ambiguous navigation scenario: A person in a wheelchair near the forward path introduces

Conversely, Figure 5 illustrates a "black swan" case. Blackness is due to a wrong and counterintuitive conclusion of the analysis. A person in a wheelchair near the forward path creates ambiguity. While no direct obstacle exists, the path is not fully open. Annotators debated the correct label, and all three models flagged it as risky. In this case, the risky label maintains a prudent approach to the problem. However, too many false positives (no real risk exists in practice) would induce untuneful braking of the wheelchair, thus reducing the quality of the driving experience. As already specified by the grey case above, additional training with obstacles along the way and in the surrounding would increase the environment awareness as well. Overall, just two swans here suggest that there may be a significant lack of knowledge about the exact level of object proximity to consider a path safe without false positives. It is worth noting that a proper filtering of VLMs outputs, such as risky flags with objects explanation, gives an automatic way to facilitate human checking.

Overall, the study confirms that Vision-Language Models (VLMs), despite differences in architecture and deployment feasibility, can meaningfully contribute to safety assessment in assistive mobility systems. While current VLMs do not yet meet the latency and reliability requirements for real-time autonomous navigation, they offer strong potential for offline safety auditing, post-hoc validation, and debugging. Their ability to semantically interpret complex scenes makes them well-suited for identifying ambiguous or under-specified risk conditions that may be overlooked by traditional pipelines. These findings support the integration of semantic-level understanding into the broader explainability and safety assurance strategy of the REXASI-PRO framework, reinforcing the complementary role of language-guided reasoning alongside conventional perception systems.

References

[1] Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You Only Look Once: Unified, Real-Time Object Detection. In 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 779–788). https://doi.org/10.1109/CVPR.2016.91

[2] Alec Radford, Jong Wook Kim, Chris Hallacy, et al. Learning transferable visual models from natural language supervision. In Proceedings of the International Conference on Machine Learning (ICML), 2021.

[3] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigpt-4: Enhancing vision-language understanding with advanced large language models. arXiv preprint arXiv:2304.10592, 2023.

[4] OpenAI. Gpt-4v(ision) technical report, 2023. https://openai.com/research/ gpt-4v-system-card.

[5] M. Jones. Multiclass counterfactual explanations using support vector data description. Safety-Critical Systems eJournal, 2(2):69–84, 2023.

[6] Shafer, Glenn, and Vladimir Vovk. "A tutorial on conformal prediction." Journal of Machine Learning Research 9.3 (2008).

[7] Liu, L.; Ouyang, W.; Wang, X.; Fieguth, P.; Chen, J.; Liu, X.; Pietikäinen, M. Deep learning for generic object detection: A survey. Int. J. Comput. Vis. 2020, 128, 261–318.

[8] lisabeth Krausmann, Amos Necci, Thinking the unthinkable: A perspective on Natech risks and Black Swans, Safety Science, Volume 139, 2021, 105255, ISSN 0925-7535, https://doi.org/10.1016/j.ssci.2021.105255.

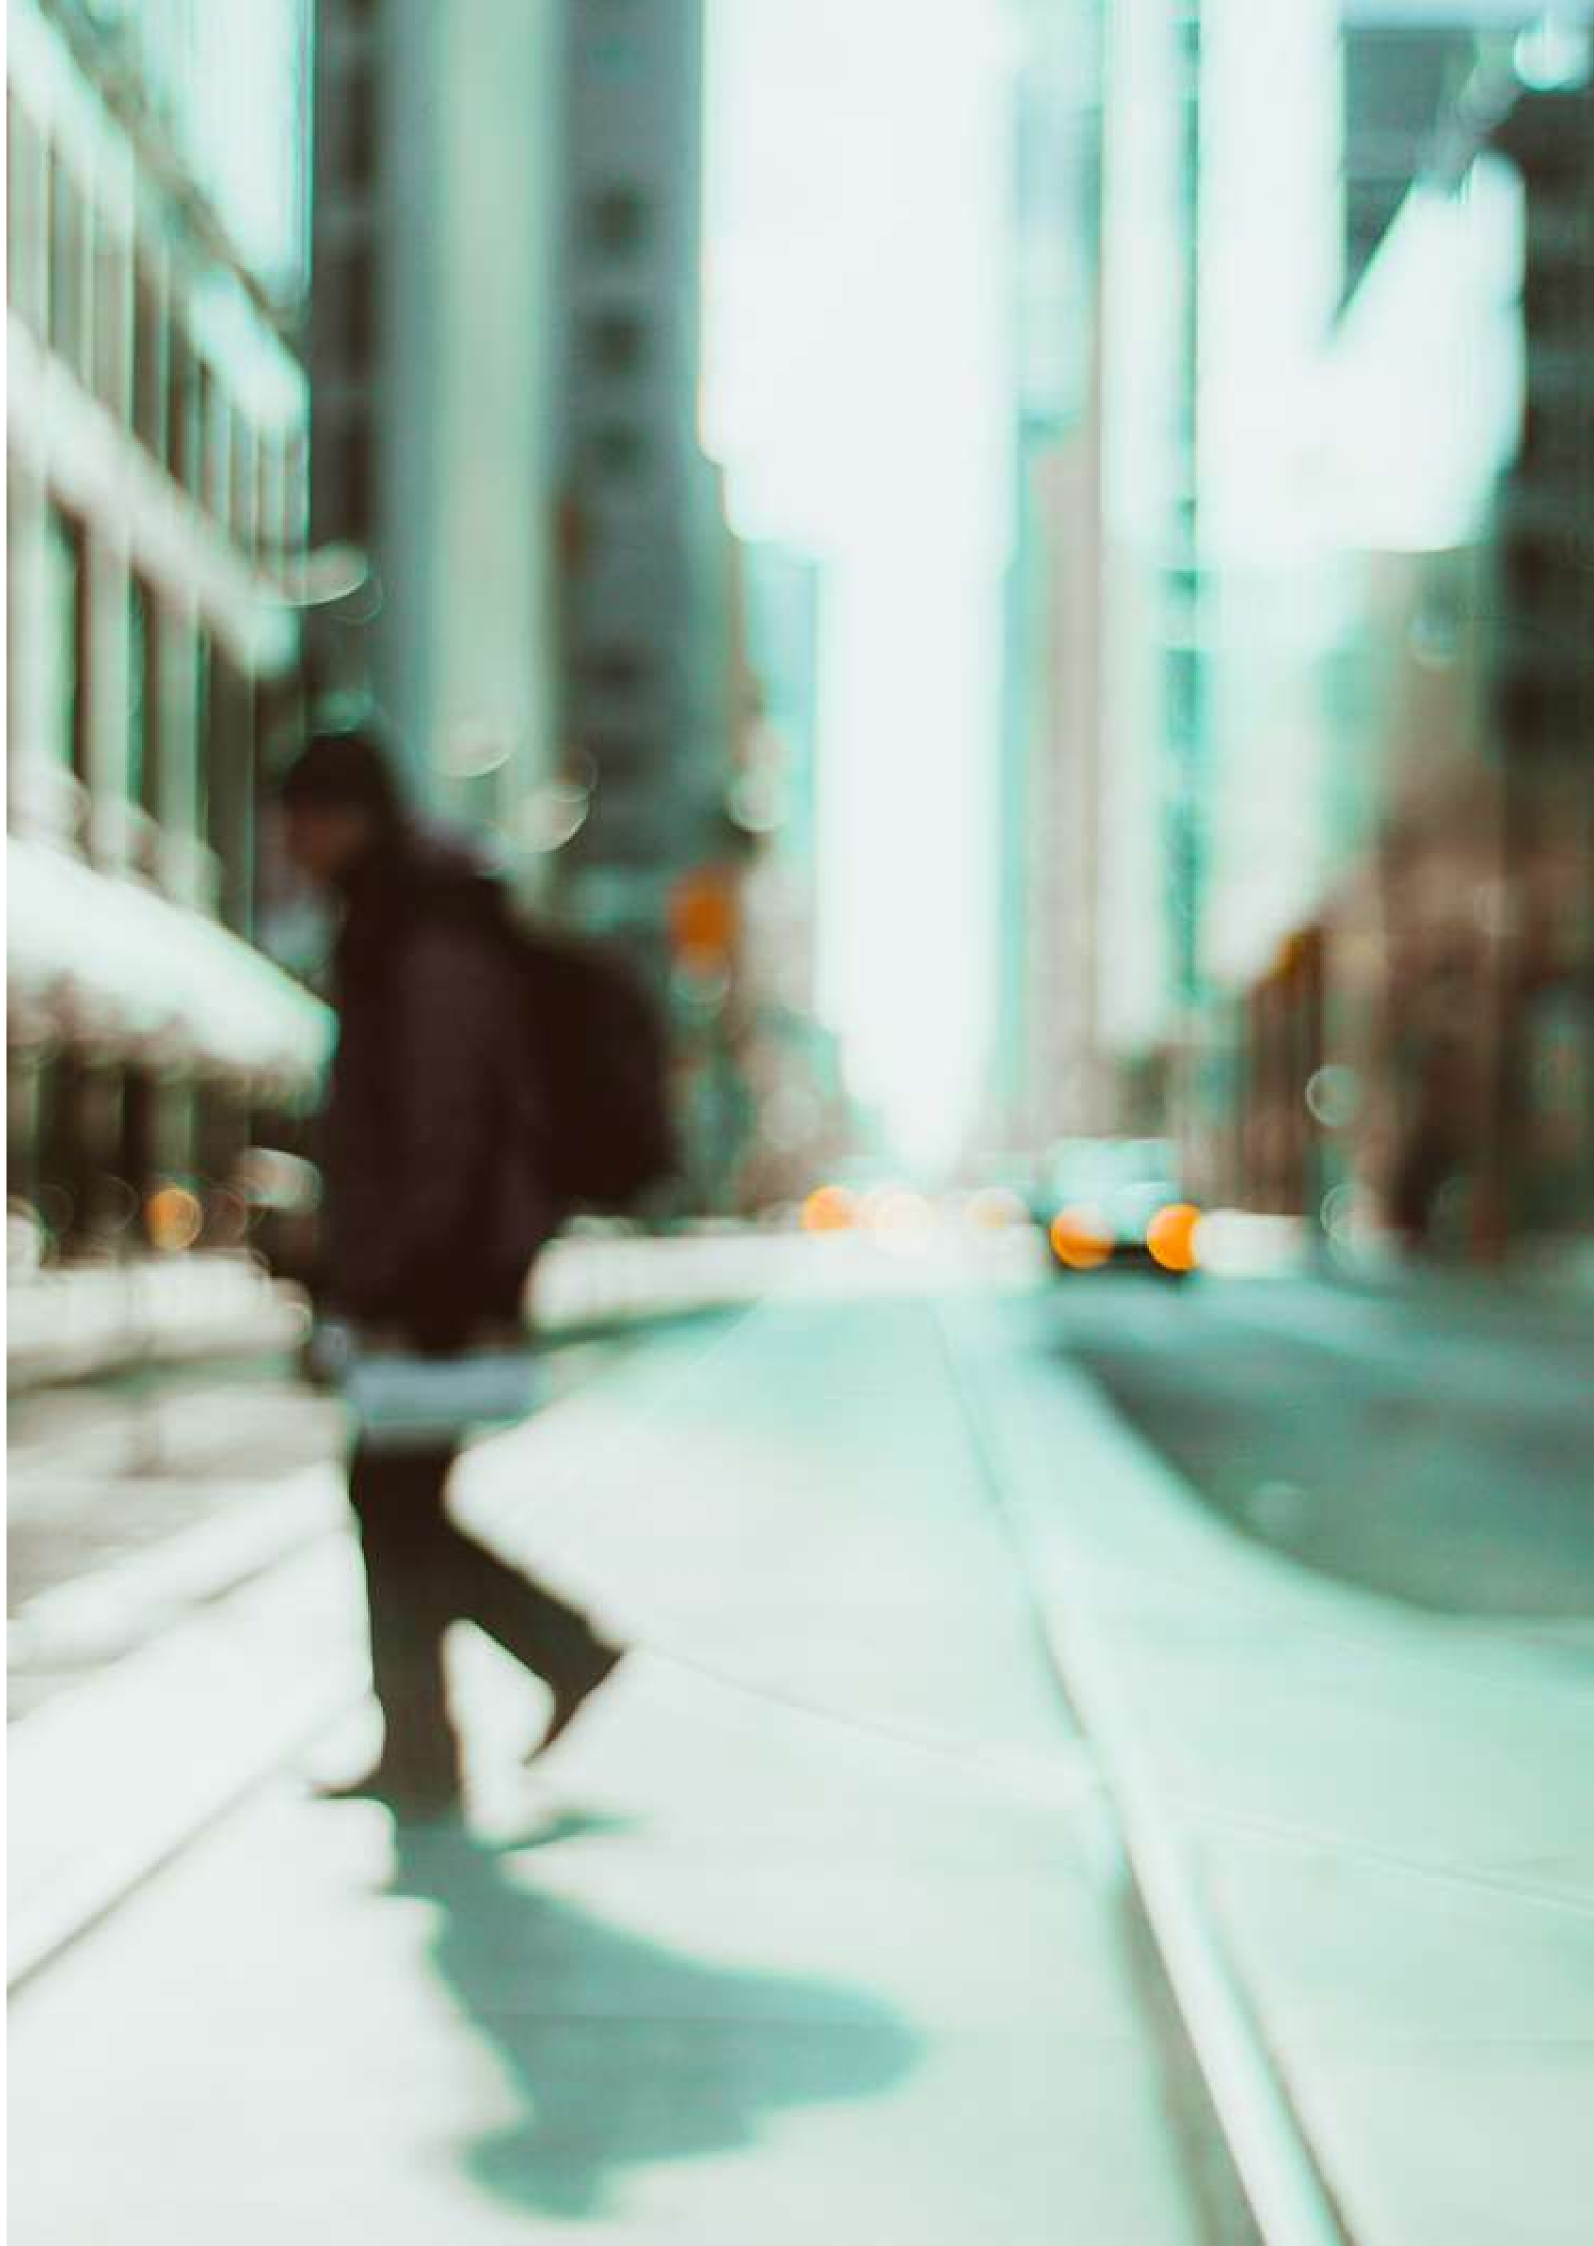

Centralized coordination of a fleet of drones while exploring an unknown area

Spindox Labs

Spindox Labs srl (SPXL, project coordinator) is the research and development center of Spindox SpA. Its mission is to support the process of digital transformation of their customers, by developing new technology and providing consultancy in five strategic domains: cloud computing, hybrid and native mobile applications, artificial intelligence, big data & analytics and internet of things.

Centralized coordination of a fleet of drones while exploring an unknown area

We all know the importance of developing efficient and effective solutions, that is solutions that require low energy consumption while ensuring a good quality of the proposed results. Designing green solutions is becoming increasingly essential for managing fleets of agents more effectively. By optimizing their tasks, the operational time required can significantly be reduced, leading to lower energy consumption and environmental impact.

In REXASI-PRO we try to achieve the reduction of power consumption of the agents, developing algorithms to manage them in an optimized way. These algorithms are used in a central orchestrator that communicates with the agents of the system and takes decisions based on the current state of the system. Coordinating agents allows them to be directed more efficiently towards their final objective.

Other aspects considered during the design and development of the algorithms implemented into the orchestrator are the use of open-source software and reuse of well-known existing algorithms. The development of simplified algorithms, which leverage efficient, pre-tested subprocesses or straightforward procedures with minimal parameter settings, also contributes to the goal of having green algorithms, intended as easy to implement, parameterize and use. Such approaches not only enhance code efficiency but also ensure a more sustainable use of resources, aligning with the principles of green computing. This dual focus on optimization and simplicity ultimately fosters a more eco-friendly technological environment.

The orchestrator is used to manage different types of agents in different use cases. The exploration of an unknown indoor area using a fleet of flying robots is one of the problems addressed in REXASI-PRO. A central orchestrator will manage a fleet of flying robots in order to perform a good exploration in less time, while reducing energy consumption. In the use cases the flying robots are provided by Hovering Solutions and therefore the developed algorithms are designed considering the technological characteristics of these flying robots. The orchestrator integrates the local information of each flying robot to have global information on which to make optimized decisions. The goal is to reduce the overall time of use of the fleet.

Autonomous exploration, using flying robots to map unknown environments, is a fundamental problem which finds various applications as inspection of large indoor buildings or pipe networks and search-and-rescue of people in dangerous zone without harming human operators. The result of the exploration is a 2D map of the environment showing free and occupied area, such as walls, furniture and other static features.

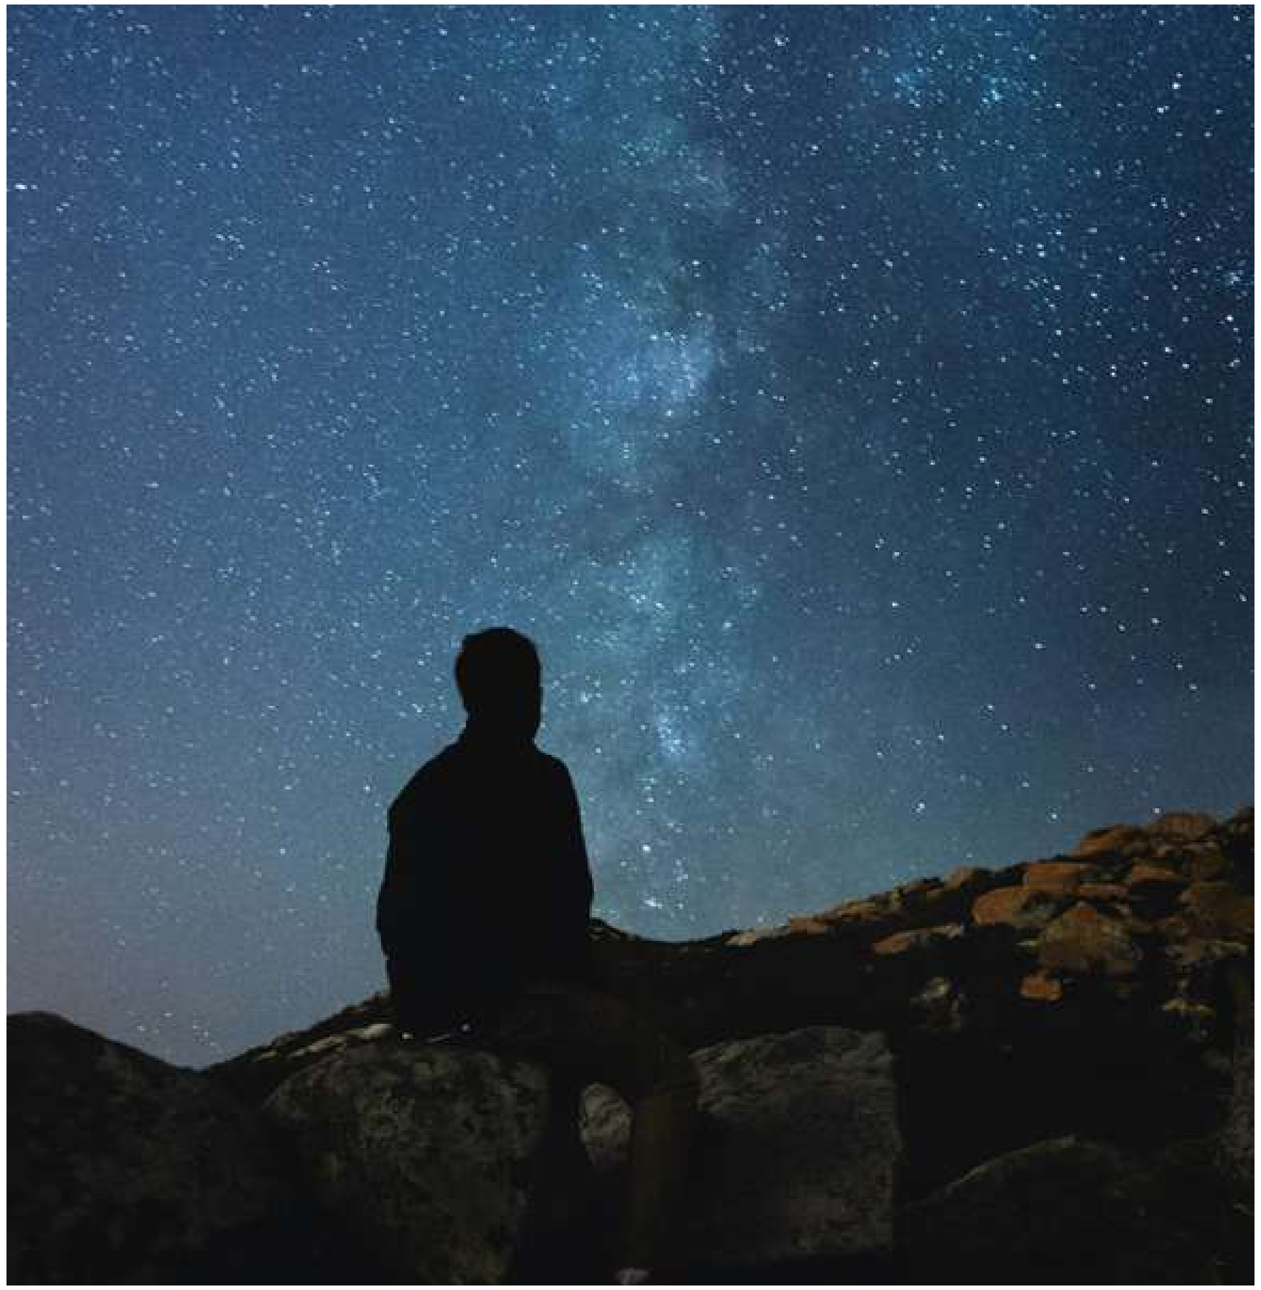

The flying robots can localize themselves in the environment and gather information from their surroundings (SLAM, Simultaneous Localization and Mapping) using on-board sensors. In this way, they can autonomously explore the environment, that is, they can move in the environment, find free and occupied areas and avoid collisions with mobile and fixed objects. Moreover, each flying robot can autonomously compute a path from its current position to a target point in a known area. Therefore, since the trajectory is computed by the flying robot itself, control algorithms for trajectory optimization are obviously computed by the flying robot itself.

However, the flying robots do not communicate with each other, therefore a central orchestrator is necessary to manage the fleet of individual agents in a centralized way. Without centralized management, each one would explore the area independently. Although these multiple explorations could be more effective, in the sense that we will have multiple maps of the same area, reducing the margin of error and the data noise, this is obviously not efficient and can also introduce collision problems between agents. Therefore, a central orchestrator, knowing the position of all the flying robots, can efficiently coordinate their exploration suggesting different areas to explore.

Another reason for using a centralized orchestrator is that the flying robots locate themselves in the environment autonomously and each one is using its own reference system, since GPS is not available. So even if they could communicate, everyone would speak a different language, because the reference systems are different. Even the resulting map of the explored area is partial for each flying robot, because each one only knows a portion of space explored. Therefore, the orchestrator is also useful for merging all the maps coming from the individual flying robots.

Since the communication between the flying robots and the orchestrator is not so fast, as required in real-time exploration, the orchestrator will be used to manage the exploration of the fleet and not the exploration of each single agent. Therefore, the orchestrator will suggest relevant areas to be explored, as rooms and corridors, and each flying robot will autonomously explore the assigned area, that is the whole room, and it will call the orchestrator when the assigned area has been explored.

The orchestrator builds the global map iteratively using the maps provided by the drones. Initially, the global map is created using the map from the first drone that contacts the orchestrator, which then becomes the reference drone. The whole space exploration process starts when the orchestrator receives the first message. The pipeline operates as follows: when the orchestrator receives a message from a drone, it first integrates the drone's map with the existing global map through a merging process. After updating the global map, the orchestrator identifies unexplored areas by detecting frontiers—grid cells that separate explored regions from unexplored ones.

Once all frontiers are identified, the orchestrator assigns a frontier to the drone based on several criteria, including the drone's current position, the positions of other drones, and their respective assignments. Finally, the orchestrator computes a path to guide the drone to the assigned frontier, facilitating the continued exploration of the unknown environment. If no frontiers are detected, the exploration is deemed complete, and the orchestrator sends a message to the drone indicating that the environment has been fully explored.

Tests have been conducted using the flying robots provided by Hovering Solutions in their facility in Madrid, showing good performances in exploring a complex environment.

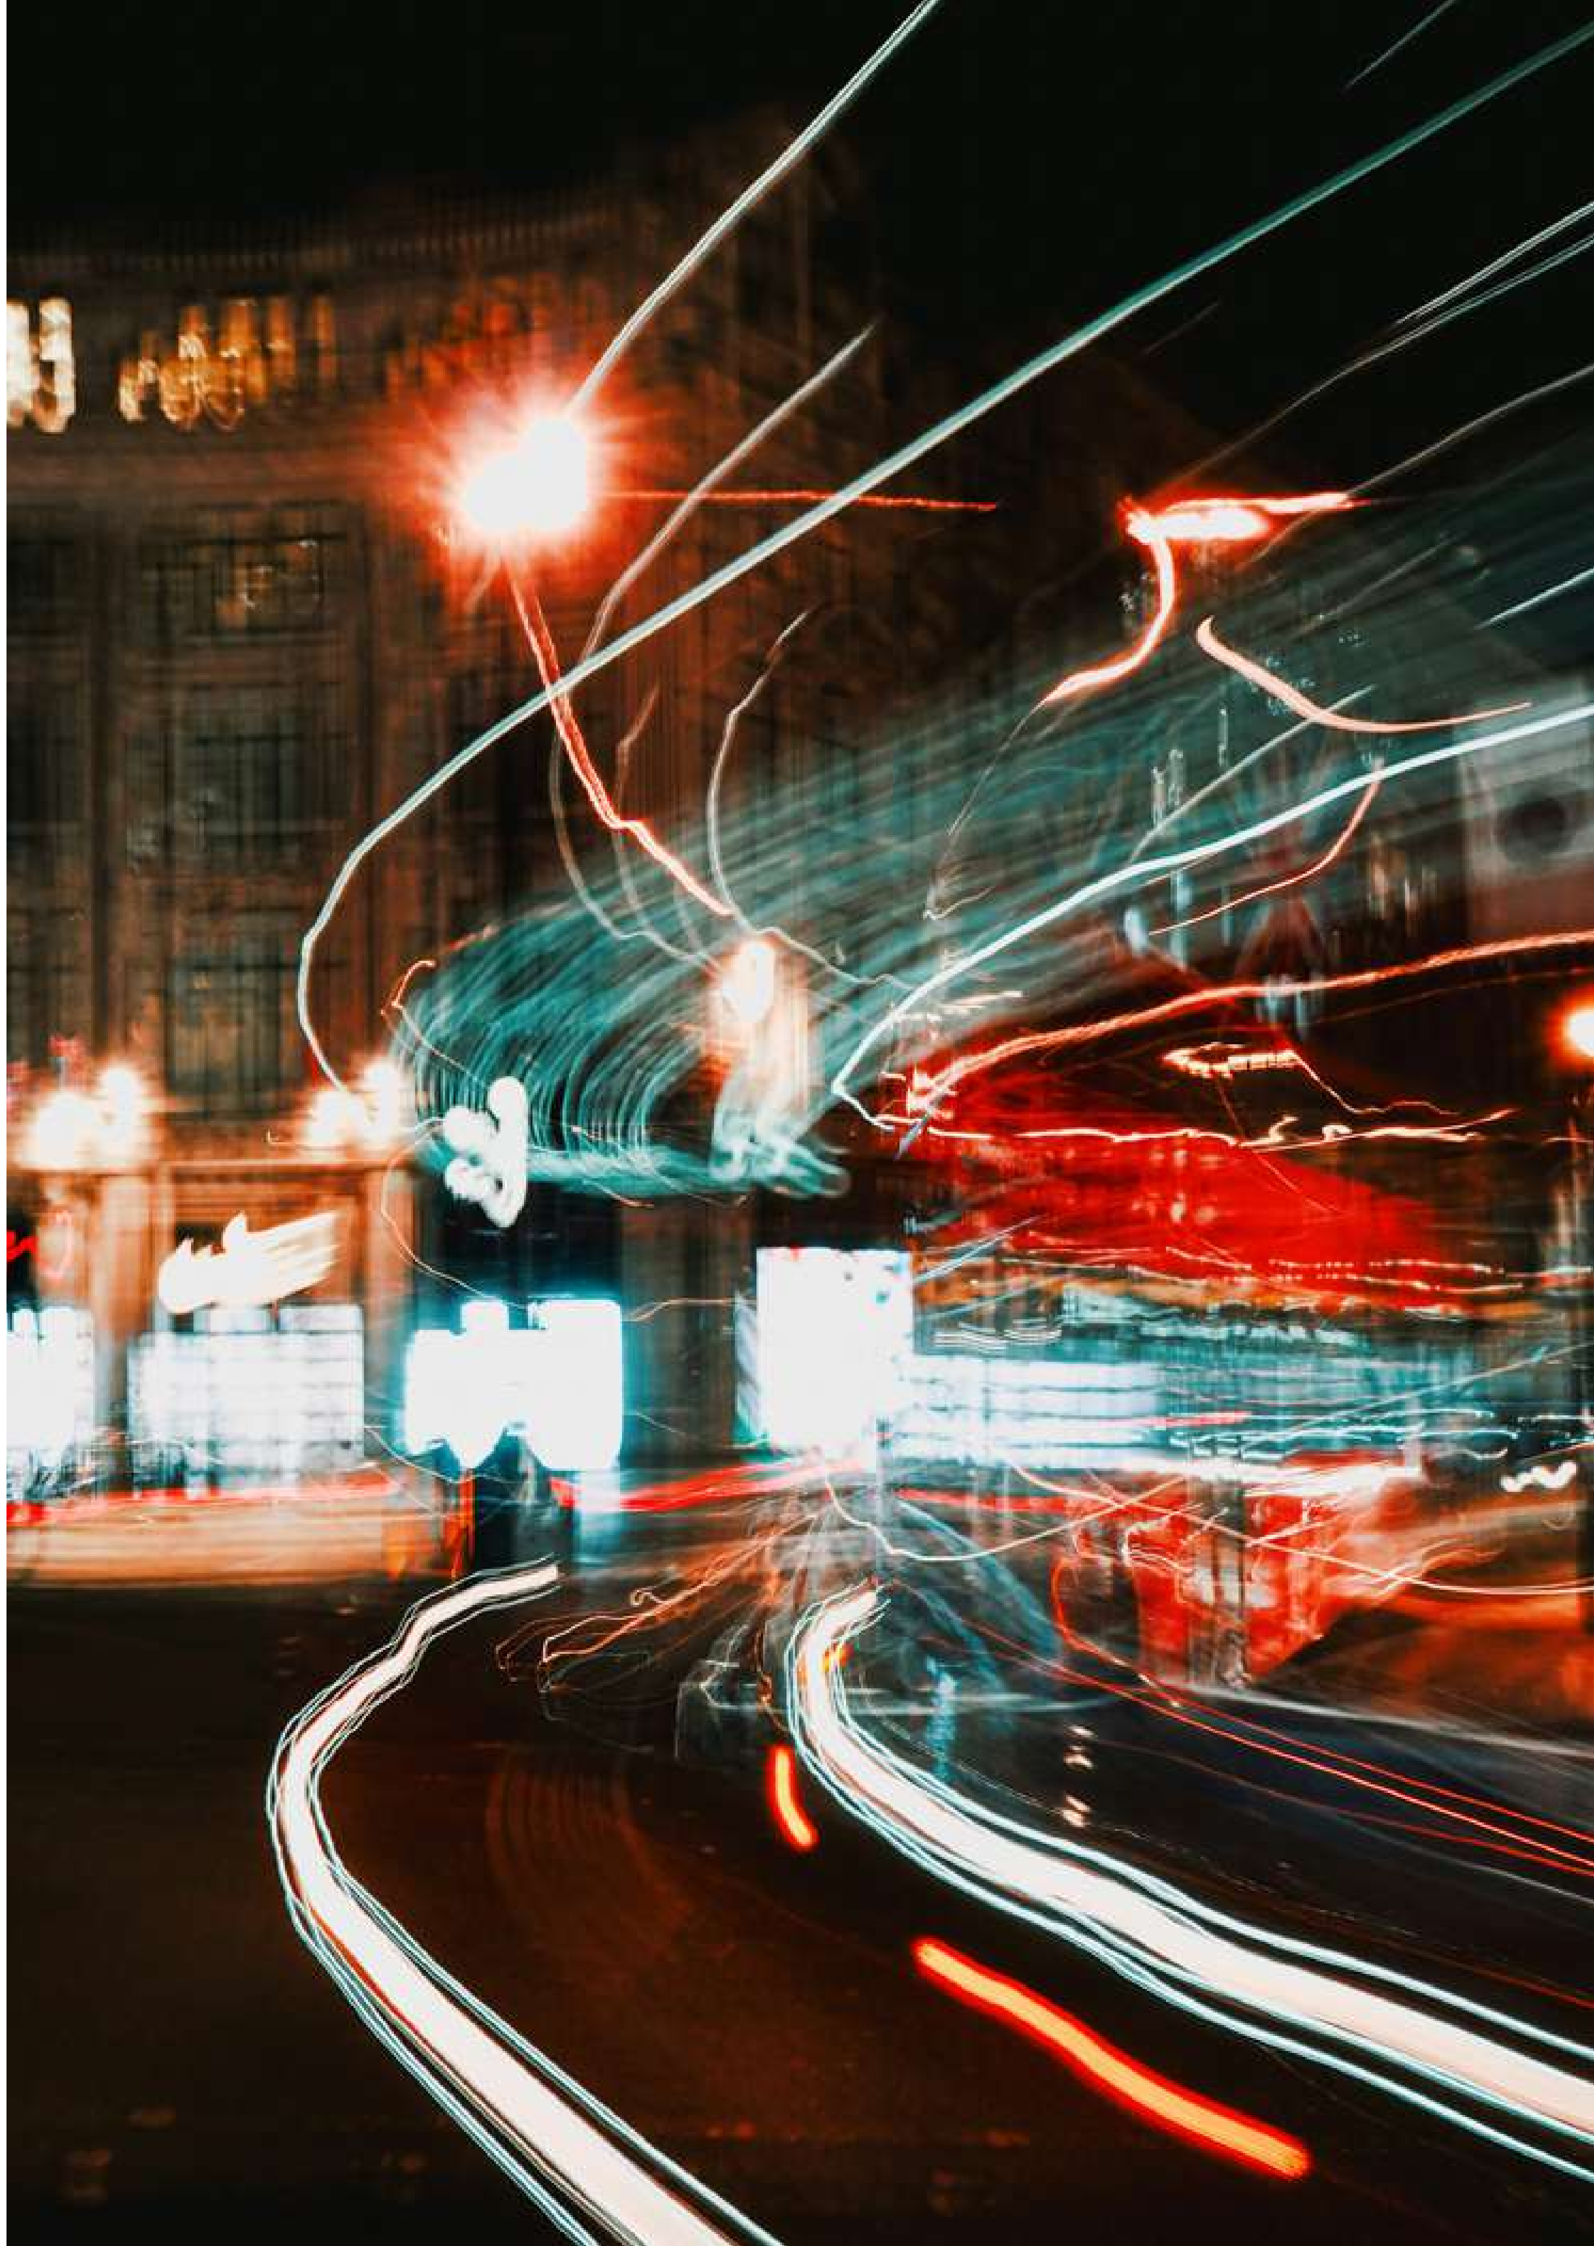

Robust AI Algorithms ensuring Safe and Reliable Autonomous Navigation

DFKI

The German Research Center for Artificial Intelligence (DFKI) is Germany's leading AI research institute. At its Bremen lab, the team develops smart wheelchair platforms, trains deep neural networks for safe and socially aware navigation, and applies formal methods to ensure compliance with the highest safety standards.

Robust AI Algorithms ensuring Safe and Reliable Autonomous Navigation

For the Rexasi-Pro partner German Research Center for Artificial Intelligence (DFKI) the central task was to develop robust algorithms for the autonomous and semi-autonomous navigation of an electrically powered wheelchair that had been equipped with sensors and a computing unit as part of the project (cf. Figure 1). The developed algorithms essentially consist of three central contributions: I) the implementation of the so-called 3D Safety Layer, a software layer for safeguarding driving commands that are either commanded by the human user or by an algorithm; II) the Deep Neural Network-based Local Navigation Approach (DNN-LNA), a neural network for obstacle-avoiding driving control; and III) a finely tuned NAV2 navigation stack that uses the SMAC Planner 2D for global navigation and the Model Predictive Path Integral Controller (MPPI) for local movement control. The NAV2-based navigation solution was developed in this context to serve both as a reference solution for DNN-LNA and as an alternative solution in the case where the DNN-LNA solution did not meet the desired performance criteria. The latter case could not be ruled out at the beginning of the project since the training of neural networks does not necessarily converge to a desired solution.

"A 3D safety layer and deep neural networks work together to detect obstacles and guide the wheelchair in real time."

3D Safety Layer

The basic idea behind the 3D safety layer is to represent the space surrounding the wheelchair using a stack of ground-parallel local obstacle maps. These are continuously updated by laser scanners and depth imaging cameras installed on the wheelchair. For each kinematically possible combination of translational and rotational speed, the so-called safety zones (cf. Figure 2) are calculated before the first use of the 3D safety layer. These zones describe the area the wheelchair will cover during the next time steps on each of the stacked local obstacle maps. During actual wheelchair use, the group of safety zones matching the pair of speed commands is selected from a pre-calculated table and projected into the current obstacle maps. Starting from the current pose of the wheelchair, the system then searches for the first overlap with an obstacle pixel, i.e., a potential collision. Finally, the position of the obstacle pixels relative to the center line of each wheelchair shape that makes up the safety zones determines whether an emergency stop is initiated, or an evasive maneuver is commanded.

Figure 1: The wheelchairs designed in the project and operated using the methods described here were equipped with depth imaging cameras, laser scanners (at the front and the back of wheelchair, not visible here), odometers, user interfaces, and a computing unit.

Deep Neural Network based Local Navigation Approach

Solutions to the problem of autonomous navigation in the robotics field usually consist of I) the topological planning layer, II) the planning layer based on global Simultaneous Localization and Mapping (SLAM) maps, and III) the obstacle-avoiding motion controllers used. Within the Rexasi-Pro project, levels II) and III) were implemented within the ROS2/NAV2 framework. The main focus was on developing a motion controller realizing level III), i.e., the deep neural network-based local navigation approach (DNN-LNA). Based on the concept of imitation learning, a real-world dataset of approximately 30 hours in length was first recorded, describing the navigation of experienced wheelchair users in partially crowded environments. This dataset was then used to train the actual network, which was based on architectural concepts such as the MobileNet backbone, attention blocks, positional encoding, and customized loss functions, to name just a few. Much more important for understanding the network architecture developed here is probably the fact that DNN-LNA is a regression network that basically maps the current pose and speed of the wheelchair, the current local obstacle map, and the current local target to a list of intermediate poses on the way to the target. In order to improve the model's performance, which was initially assessed as rather mediocre, additional simulation-based training data was added to the network in collaboration with project partners SPINDOX and SUPSI. Although it was shown that this improved model performance in terms of the loss function used, the quality achieved so far is only suitable for realizing short experimental driving distances on a real wheelchair.

NAV2-based Navigation Solution

The Model Predictive Path Integral Controller (MPPI) is a highly configurable motion controller that represents the de facto standard solution of the NAV2 robotics platform for various kinematic platforms. From the perspective of the three-layer navigation solutions presented above, MPPI represents the third layer. Based on the currently measured translational and rotational speed of the wheelchair, a global path to be executed, and the currently estimated pose with respect to this path, MPPI explores an adjustable set of potential trajectories up to an adjustable time horizon. Each of these trajectories obeys the configured kinematic boundary conditions of the wheelchair and allows for adjustable noise on these values. The final selection of the trajectory to be followed relies on a set of so-called critics. Each of these finely tuned modules evaluates a specific aspect of the proposed solutions, such as obstacle avoidance, progress along and orientation with respect to the global path, or compliance with kinematic constraints. This central component, MPPI, is only one component in the NAV2-based navigation solution we have assembled. Other noteworthy software modules we use in this context are the SLAM solution Real-Time Appearance-Based Mapping (RTAB-Map), the statistical self-localizer Adaptive Monte Carlo Localization (AMCL), and the global planner Smac2D. After intensive tuning of the components mentioned here, we were able to not only carry out test drives in busy indoor environments with a length in the single-digit km range as part of the Rexasi-Pro project, but also the final project evaluation in which, among other things, 16 test subjects were able to successfully use this navigation solution in a populated indoor navigation scenario.

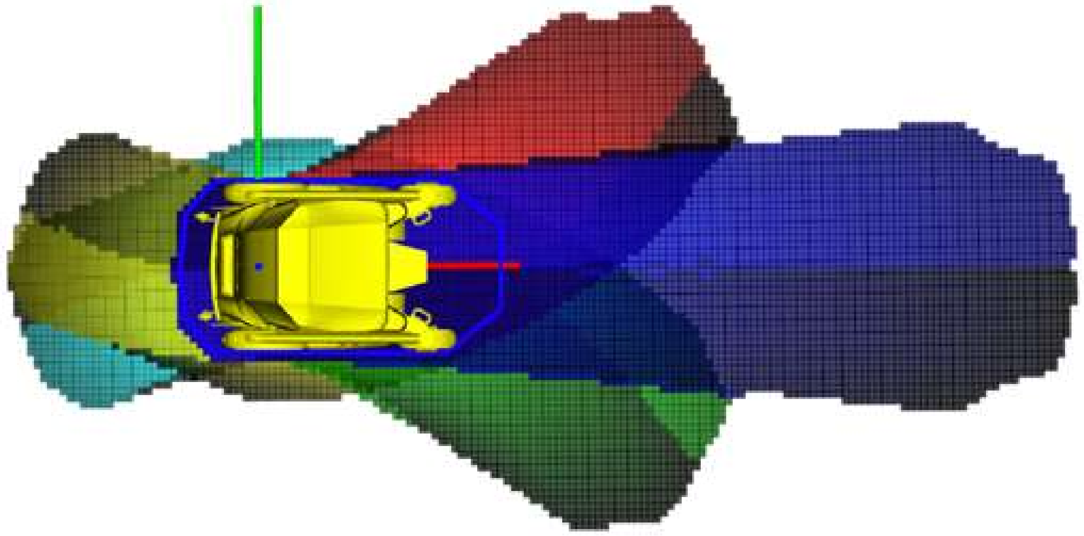

Figure 2: Precomputed safety zones for different combinations of translational and rotational velocities speed up the process of finding potential collisions and evasive movements.

Hovering Solutions: Enabling Indoor Autonomy through Flying Robots

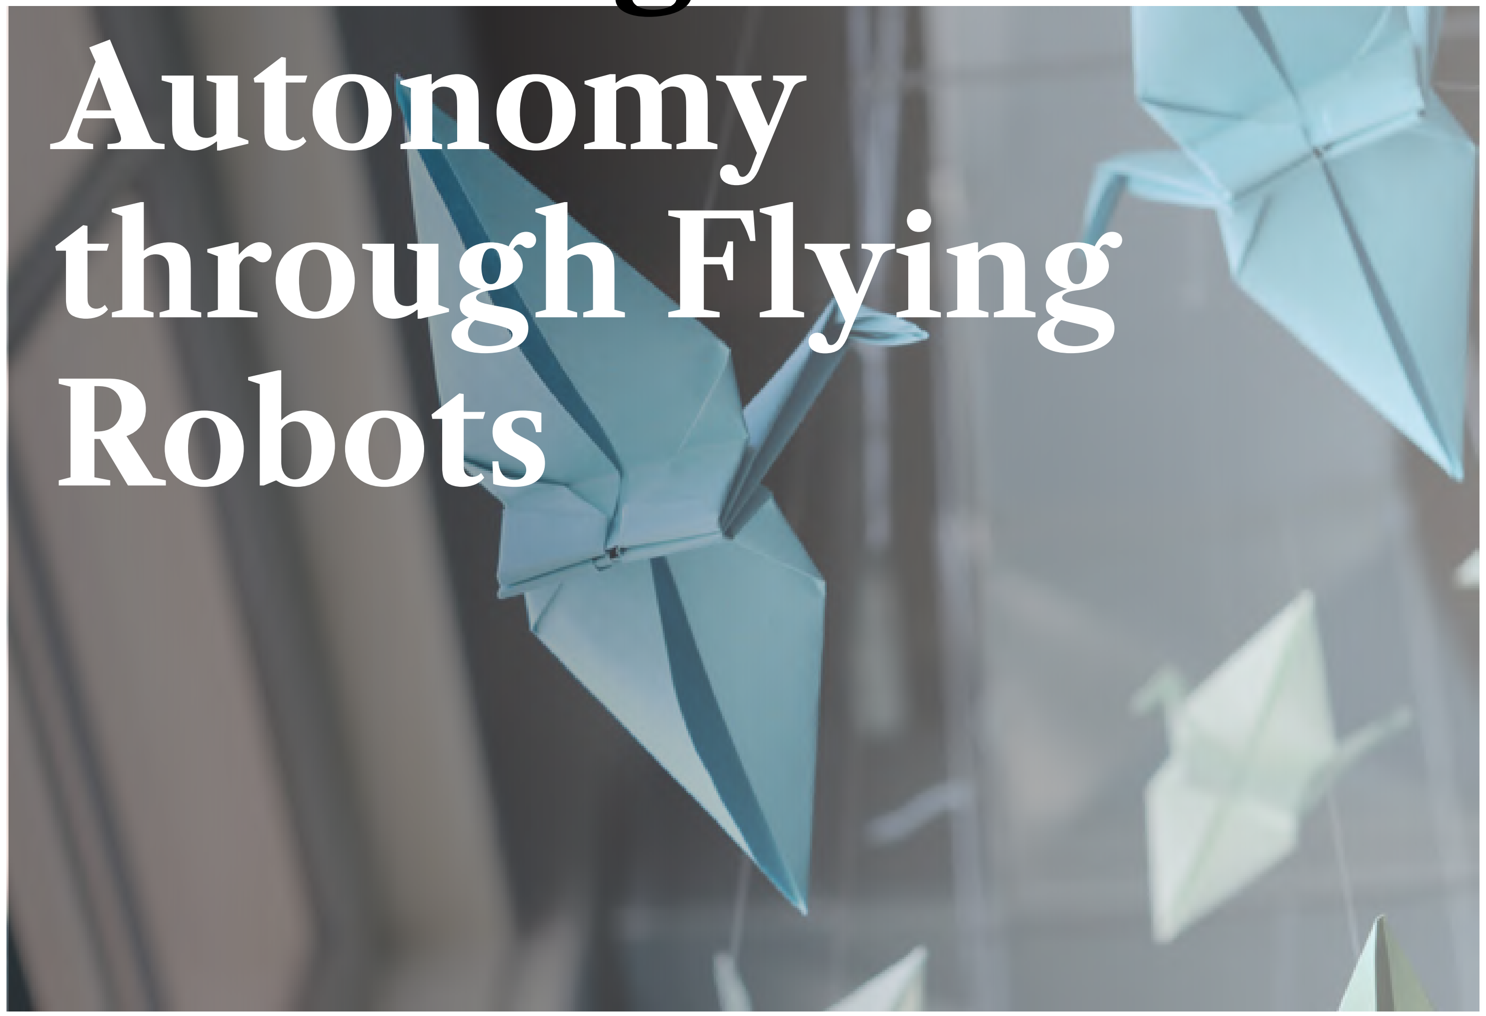

Hovering Solutions

Hovering Solutions, founded in Madrid in 2016, develops autonomous aerial robots for GNSS-denied environments, delivering high-resolution 3D mapping and imaging to help industries like mining, hydropower, water utilities, and construction digitalize critical, hard-to-access infrastructure safely and efficiently

Autonomous drones by HSOL are redefining indoor navigation for assistive robotics – enabling safe, self-guided flight in human environments as part of the REXASI-PRO project.

In the landscape of trustworthy AI and mobility robotics, UAVs from Hovering Solutions (HSOL) are contributing a crucial layer to REXASI-PRO: autonomous flying robots designed for indoor navigation in support of people with reduced mobility.

Where ground-based platforms reach their limits due to terrain, crowding, or accessibility constraints, indoor UAVs offer a complementary perspective, enabling quick situational awareness, cooperative navigation, and autonomous inspection in complex environments.

HSOL's mission within the project is clear: integrate flying platforms that are capable of indoor operation, spatial mapping, and real-time decision-making in close proximity to humans.

From Concept to Corridor: Real-World Validations

As part of the project's validation phase, HSOL successfully completed a series of flight tests in an indoor environment in Italy. The team demonstrated robust autonomous navigation under real-world constraints.

The drones flew through narrow corridors and rooms without external infrastructure, relying entirely on their onboard systems, including SLAM, inertial sensing, and embedded computing. These tests not only validated the system's robustness, but also its safe operation in human-inhabited areas.

Figure 1 HSOL's drone enters a meeting room during live autonomous flight – highlighting precision in confined, active spaces.

As shown in Figure 1, one of the aerial platforms successfully entered a populated workspace, demonstrating its ability to operate precisely and unobtrusively in confined environments.

HSOL's Core Technical Contributions

HSOL's role within REXASI-PRO extends across hardware development, autonomous navigation, and real-time spatial awareness. The aerial platforms are built with integrated sensing and onboard computing, enabling them to operate safely and efficiently in indoor spaces that lack GPS access and are often cluttered or occupied by people.

The system achieves full indoor autonomy, managing flight control and navigation without relying on any external infrastructure. It performs simultaneous localization and mapping (SLAM) in real time, allowing the drones to explore and build accurate representations of unknown environments as they fly. Navigation is enhanced by dynamic path planning capabilities, which help the drones adapt to changing layouts and avoid obstacles, even in narrow corridors or through tight doorways.

Safety is built into the hardware itself. Each unit is designed specifically for indoor environments, with features such as propeller shielding and system redundancy to reduce risk when operating in close proximity to people and structures. Because the platform requires no external markers or prior mapping, it can be rapidly deployed in new and unmapped buildings — making it ideal for dynamic missions where speed and adaptability are critical.

Data-Driven Autonomy and Future Integration

HSOL drones don't just fly, they generate rich, actionable data essential for au-

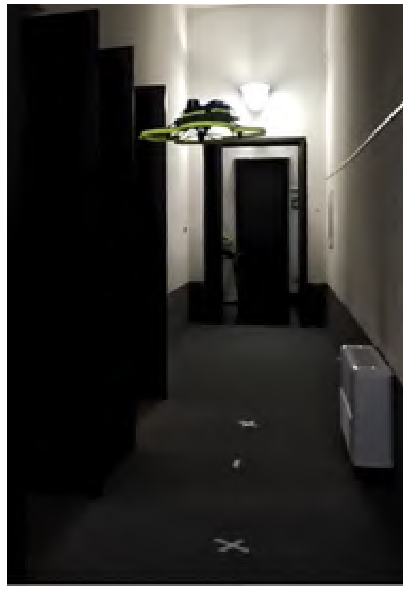

Figure 2 Autonomous indoor flight through a corridor during testing. The drone adapts to tight spaces using onboard SLAM and visual cues.

tonomy and integration. During each mission, the system produces localized 2D and 3D maps that allow the drone to understand and reconstruct its surrounding environment in real time. In addition, it logs flight trajectories and paths taken throughout the building, which are crucial for post-mission analysis and optimization. Sensor-based readings further enhance environmental awareness, enabling the platform to

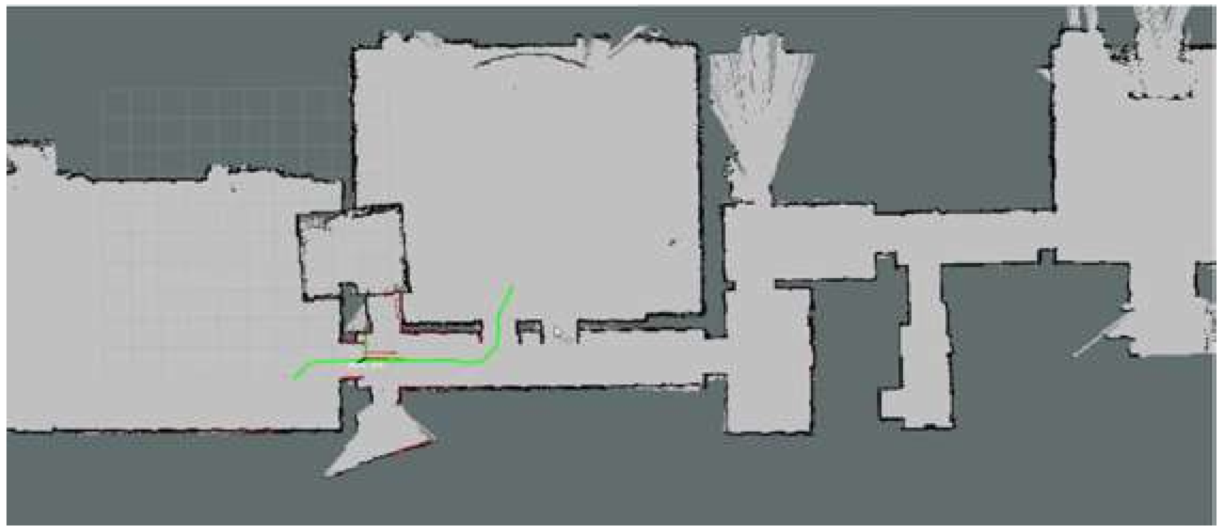

Figure 3 Path planning using the map generated by the drone in the CNR. In green, the safe route to reach a commanded point.

make intelligent decisions during dynamic navigation tasks.

These continuous data streams support the UAV's onboard autonomy while also contributing to the broader multi-agent orchestration layer developed within REXASI-PRO. Although the coordination and planning logic for "greener AI" is led by other partners, HSOL's high-quality mapping and localization data form a foundational element for those collaborative decisions.

Further data processing is currently underway. This includes visualizations of the generated maps and the extraction of key flight behaviour patterns as shown in figure 4, which will be featured in upcoming demonstrations and publications as part of the project's dissemination strategy.

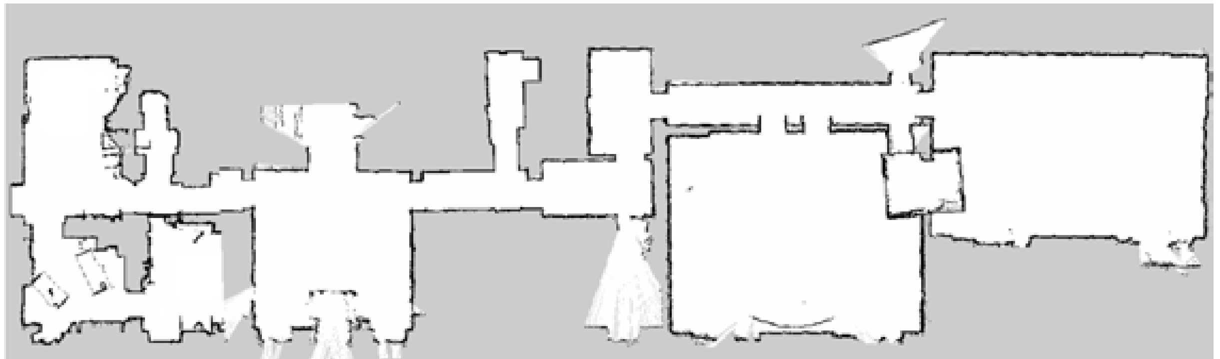

Figure 4 Map of the CNR generated by fusing the wheelchair's and the drone's maps

Impact: From Prototypes to Applications

Autonomous indoor UAVs have significant potential for practical applications, especially in environments that are dynamic, human-centered, or difficult to access by conventional robots. Their ability to operate safely and intelligently in indoor spaces opens doors to use cases such as emergency response, hospital logistics, and elderly care. In these contexts, timely information, fast mobility, and adaptable autonomy can make a measurable difference in safety and service quality.

With the initial validation flights completed and data analysis underway, HSOL is now focusing on refining its platform to address the next phases of development. These include improving interoperability with other robotic agents, scaling the solution to larger environments, and integrating seamlessly into coordinated multi-agent missions — all key objectives as the project moves forward.

As REXASI-PRO continues to evolve, HSOL remains committed to advancing trustworthy AI through aerial autonomy — not just as a technical achievement, but as a meaningful contributor to smarter, more inclusive environments.

A participatory approach to defining the ethics requirements of REXASI-PRO European project.

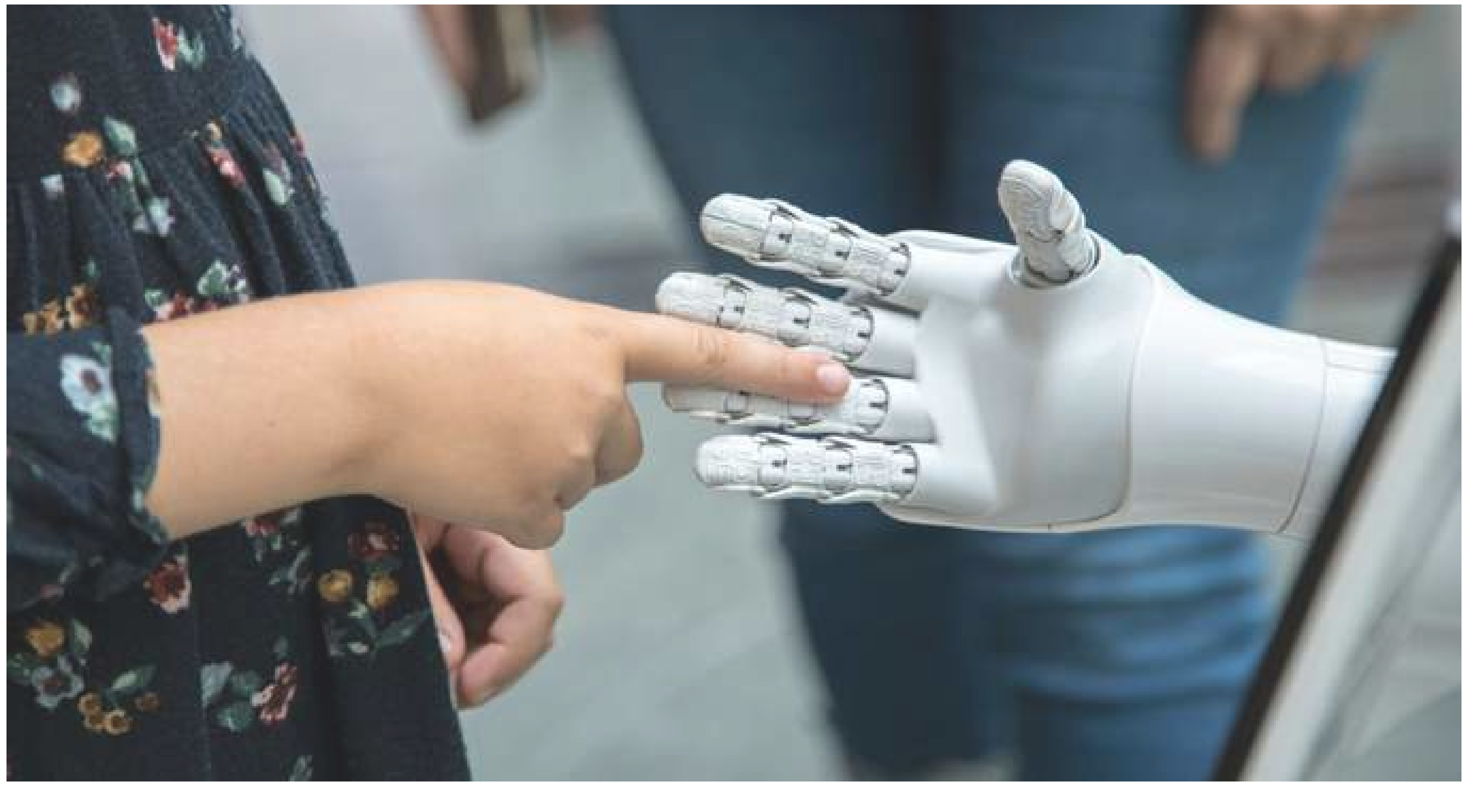

Scuola di Robotica

Scuola di Robotica (SDR) is a non-profit association founded in 2000 by a group of robotics and human science scholars. The main objective of Scuola di Robotica is the promotion of culture through education, training, education and dissemination of the arts and sciences involved in the process of development of robotics and new technologies.

REXASI-PRO has adopted a participatory and inclusive approach to define the ethics requirements, engaging partners and stakeholders in a co-creation process to ensure alignment with ALTAI recommendations and societal values.

Ethics-by-Design solutions

The Horizon Europe REXASI-PRO project has based its design and development on the methodology called Ethics by Design that refers to the systematic integration of ethical principles from the very early stages of technological design and development (European Commission document, 2021). It implies embodying ethical requirements into each phase of the development process allowing that ethical issues arising in itinere are addressed as early as possible and followed up closely during research activities Operto, 2011). The aim of REXASI-PRO Ethics by Design methodology is to make all partners aware of the potential ethical concerns in every phase of their development of the system. This approach, tailored to the objectives of the research proposed, takes into consideration that, for instance, ethics risks can be different during the research phase from that of the deployment or implementation phase (Veruggio et al, 2016).

Following an initial analysis of ethical challenges - which could have been necessarily partial and in progress – and an in-depth assessment by School of Robotics (Report on Ethical, Legal and Societal Requirements for the Design and Development of the REXASI System), at the end of the first year of the project, partners were invited to:

- participate in a Survey to assess the consortium level of knowledge, awareness, and attention with respect to ESL (Ethical, Legal and Social) Issues and ALTAI recommendations;
- to adapt the ethical requirements identified in the project's DoA to the developments achieved during the first year.
- identify the main challenges the project is going to encounter applying the EC

guidelines for Ethical requirements to the system design.

ALTAI recommendations and REXASI-PRO Ethical Self-Assessment

This process of aligning the project phases with the final ethical requirements is suggested by ALTAI, (Assessment List for Trustworthy Artificial Intelligence, 2020), a tool developed by EU expert group's AI HLEG set up by the European Commission to help organizations and developers implement and assess Trustworthy AI, in line with the seven key requirements identified by in the EU expert group report, that are:

- Human agency and oversight – meaningful human involvement and oversight,
- Technical robustness and safety – resilience, reliability, and safety,
- Privacy and data governance – respect for privacy and sound data management,
- Transparency – traceability, explainability, and clear communication,
- Diversity, non-discrimination and fairness – inclusiveness and fairness,
- Societal and environmental well-being – positive impact on society and the environment.
- Accountability – responsibility and auditability.

Applied in REXASI-PRO, these principles have been developed in accordance with the project's objectives:

- respect for human agency (autonomy, dignity and freedom);
- privacy, personal data protection and data governance;
- fairness (Avoidance of algorithmic bias:
- Universal accessibility; Fair impacts)
- individual, social, and environmental well-being;
- transparency;
- accountability and oversight.

The analysis of the partners' comments and responses to the Ethical Survey were incorporated into an Ethical Self-Assessment approved by the consortium, in which the main ethical issues and the adopted solutions were identified.

Ethical solutions

Requirements of safety, security, and explainability are entangled with ethical requirements because the fundamental ethical principle is to ensure the safety, security, well-being, and comfort of the users. The project is developing a novel trustworthy-by-construction solutions for social navigations and a methodology to certify the robustness of AI-based autonomous vehicles for people with reduced mobility (Narteni, et al, 2022). The social navigation algorithms are exploiting mathematical models of social robots (in the project, the robotics wheelchair) that are trained by using both implicit and explicit communication. REXASI-PRO methodology augments existing system-level and item-level engineering frameworks by leveraging novel

eXplainability methods to improve the entire system's robustness.

In operating a multi robot system for Assistive robotics like REXASI-PRO, the ethical requirements are related to the need of the assisted a person with reduced mobility and operating in an unstructured environment inhabited by humans who are not necessarily aware of the action of the robotics system.

The methodology will be used to certify the robustness of both autonomous wheelchairs and flying robots. The flying robots are equipped with unbiased machine learning solutions for people detection that will be reliable also in an emergency. Thus, REXASI-PRO is making the AI solutions greener. The REXASI-PRO framework will be demonstrated by enabling the collaboration among autonomous wheelchairs and flying robots to help people with reduced mobility.

To adopt this technology in real-life scenarios, new trustable social navigation approaches are required. Among the challenges, the wheelchair adapts to the dynamics of the environment (avoid collisions, avoid unexpected displacements), maximizes driver well-being, guarantees indoor localization and target identification.

We have considered some challenging environmental characteristics for such a system. If we consider, for example, that the system is to operate in a hospital and is to autonomously transport the user through several floors and various corridors, a control center (orchestrator) supervises the path and can alert some control points located along the way. If we design that the system must operate outdoors where video cameras cannot be placed, or where their reliability is uncertain, the use of the drone could be necessary.

Having analyzed how ethical principles should be integrated into the project, it is important that the design and development applied the best engineering solutions to ensure the most effective functioning of the system In compliance with the ethical requirements.

References

Ethics By Design and Ethics of Use Approaches for Artificial Intelligence, 2021, document drafted by a panel of experts at the request of the European Commission (DG Research and Innovation) a https://ec.europa.eu/info/funding-tenders/opportunities/docs/2021-2027/horizon/guidance/ethics-by-design-and-ethics-of-use-approaches-for-artificial-intelligence_he_en.pdf

Operto F. (2011). "Ethics in Advanced Robotics". in IEEE Robotics & Automation Magazine – Special Issues on Robot Ethics, Robotics and Automation Society", n. 18, v. 1, anno 2011, pp. 72-78.

Veruggio G., Bekey G., Operto, F. (2016). "Roboethics: Social and Ethical Implications", in B. Siciliano e O. Khatib (Eds), The Springer Handbook of Robotics, pp. 2135-2160.

Narteni S., Orani V., Vaccari I., Cambiaso E., and Mongelli M. (2022). "Sensitivity of Logic Learning Machine for Reliability in Safety-Critical Systems". IEEE Intell. Syst. 37(5): 66-74 (2022)

A Modular Framework for Robot Navigation and Simulation

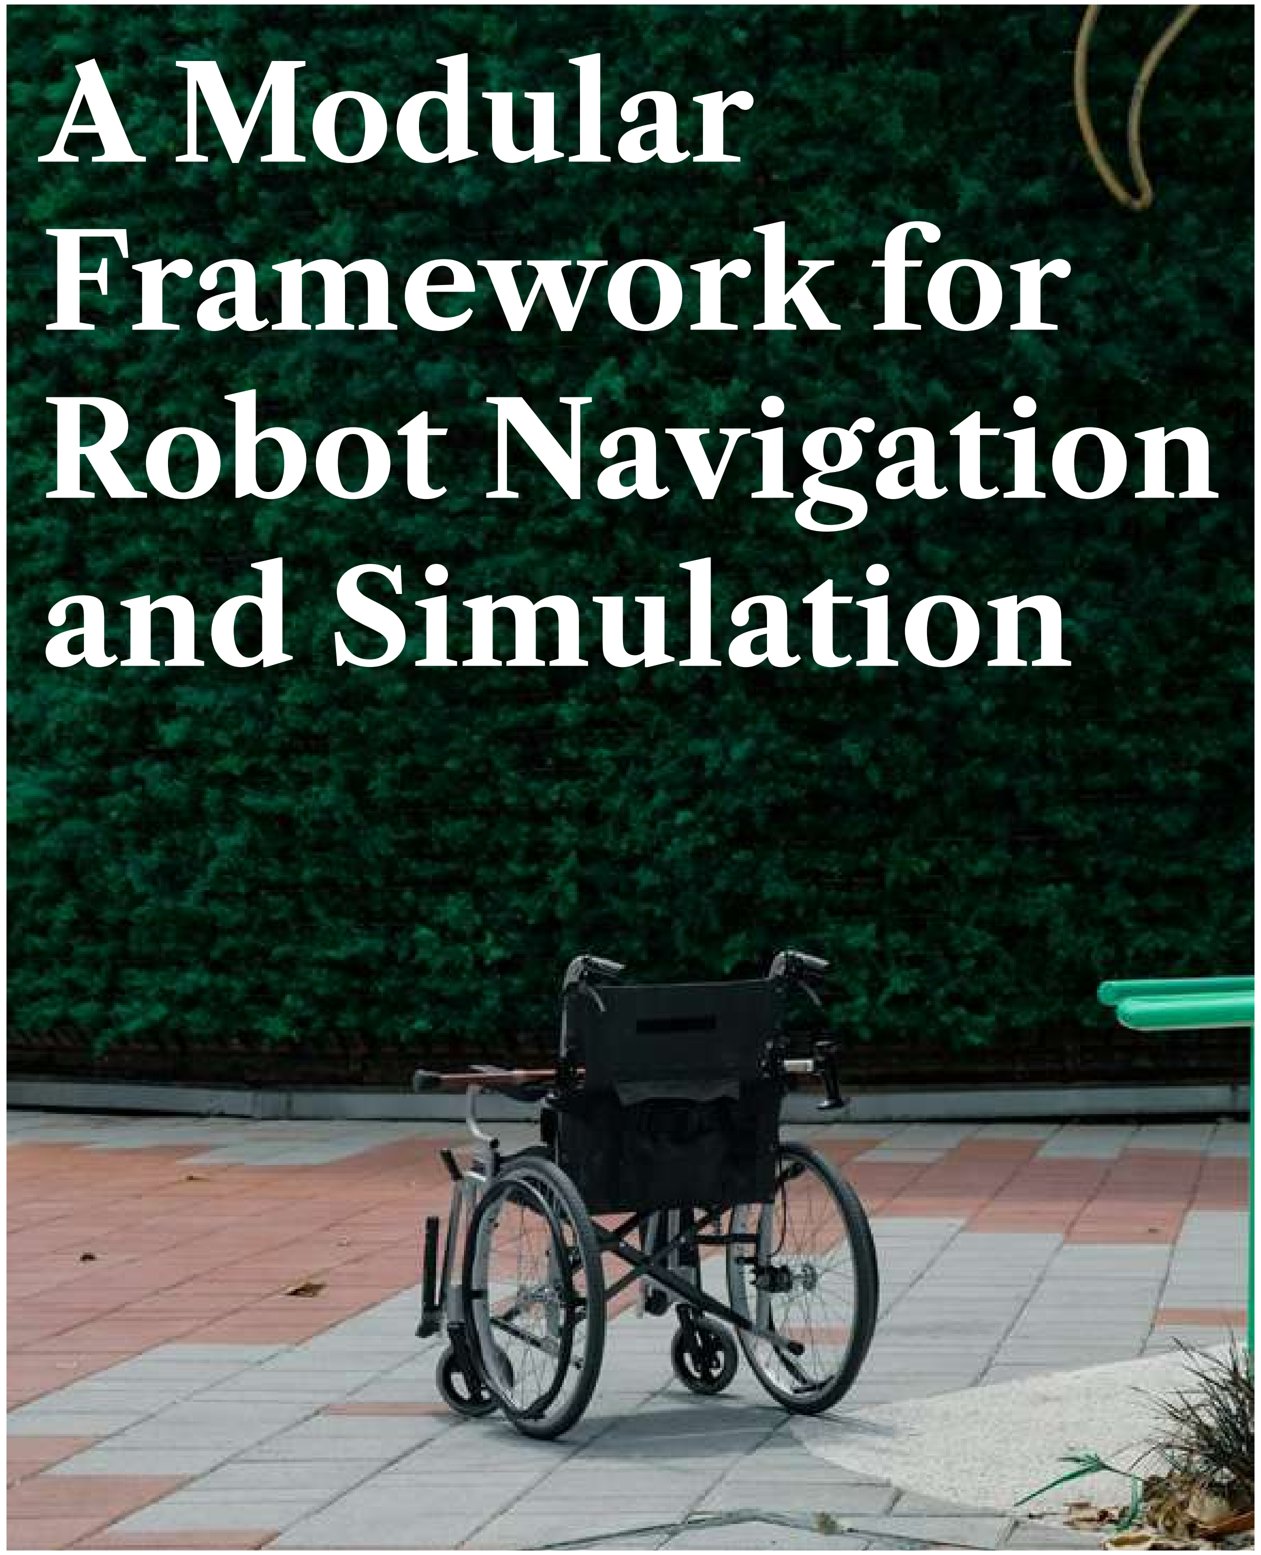

SUPSI

University of Applied Sciences and Arts
of Southern Switzerland
SUPSI

The University of Applied Sciences and Arts of Southern Switzerland (SUPSI) contributes expertise in Artificial Intelligence through research performed at the Dalle Molle Institute for Artificial Intelligence (IDSIA)

Navigation

Navigation is a fundamental skill for humans, animals, and mobile robots. The word originates from the Latin name for ship, meaning "the ability to drive a ship"; however, it refers to the more generic ability of moving towards a target. A look at ship navigation helps us to understand the concept better. Ship navigation is a collaborative effort: the navigator plans a path to the destination, the helmsman steers to stay on track, avoiding hitting any object, while sailors manage how the ship is propelled. Similarly, a robot uses different controllers to plan a route, follow it while avoiding obstacles, and control the speed of motors. Moreover, like the helmsman takes account of the currents and the winds, a robot needs to consider the environmental conditions. In summary, robotics navigation refers to the ability of following a direction while avoiding obstacles; which obstacles, and how to avoid them, depends on the context.

Smart wheelchairs

In the context of REXASI-PRO, the focus shifts from ships to smart motorized wheelchairs that navigate autonomously in spaces shared with people. How should smart wheelchairs navigate? For starters, they should be safe: no collisions or capsizing should occur, and movements should be gentle to minimize discomfort. At the same time, they should reach their target without excessive maneuvering, which costs time, energy, and patience. Additionally, moving among people presents challenges such as respecting personal spaces and other unwritten social norms: for example, passing too close to people may inadvertently create embarrassment.

Playground

Our goal is to equip smart wheelchairs with the best navigation algorithms. As a first step, this requires a precise definition of a navigation algorithm and a way to compare different navigation algorithms. For this, we developed Navground, which stands for navigation playground. Navground defines the minimal software interface that navigation algorithms should satisfy and enables running experiments to assess their performance. Defining these interfaces required an act of balance, because simpler, more abstract interfaces describe a broader set of problems but diminish their practical utility. We resolved this dilemma pragmatically by distilling a minimal interface from a collection of concrete navigation algorithms, and subsequently verifying its suitability for implementing novel algorithms.

Modular and extensible

Like kids bringing their own toys and tools when playing in a sandbox, the navigation playground provides some navigation algorithms while simultaneously encouraging users to contribute their own algorithms. The playground is indeed modular and extensible: it defines the shape of playing blocks (what navigation behaviors are, what kinematics are, and so on), it implements some blocks (e.g., several model-based navigation behaviors from the robotics literature), and lets users add more blocks (e.g., novel navigation behaviors) using C++ or Python as they prefer. Python allows rapidly prototyping algorithms that can then be ported to C++ to increase computational efficiency.

Simulation

Simulation plays a pivotal role in testing robotics algorithms. It enables us to validate that an algorithm is safe by simulating extensive hours of agents navigating under the modeled environmental conditions, ensuring that no collisions occur. Furthermore, it facilitates the selection of the most effective algorithms by running the same scenarios for each algorithm and comparing performance metrics. Additionally, it simplifies the tuning of algorithm parameters by simulating any combination of parameters. Our navigation playground fulfills the essential requirements for such assessments. Computational costs are low: for instance, the simulation of model-based navigation behaviors implemented in C++ requires only a few micro-seconds per agent per update cycle, which results in simulating one hour of navigation within a few seconds. Experiments are reproducible: executing multiple instances of an experiment with the same configuration yields identical outcomes. Lastly, the configuration of an experiment is user-friendly: whether via Python, C++, or YAML, users can set up complex experiments, incorporating random variability if desired, in a few lines of code.

Simulation at work

Building upon Navground, we explored different directions to enhance the navigation of smart wheelchairs in simulation. For instance, we devised strategies to mitigate traffic congestion caused by narrow passages, where only a single individual or wheelchair can pass at a time, like when passing through doors. We also integrated a wheelchair-specific path planner that computes trajectories with reduced curvature to maximize user comfort. Furthermore, we evaluated the performance of bio-inspired model-based navigation algorithms to simulate people but also to control wheelchairs moving among people. Additionally, to replicate the wheelchair developed by DFKI, we generated realistic perception.

Machine-learning for navigation

Machine-learning (ML) has three primary applications in robotics navigation: perception models (models that extract an environment description, like the position of obstacles, from raw sensor readings), policies (navigation algorithms that generate actions from observations), and assessment models (models that extract a robust and/or interpretable assessment from a large array of simulations). In REXASI-PRO, all three applications play an important role: we built ML-based pipelines for people detection, we trained several navigation policies, and we developed and applied XAI methods for navigation. Navigation policies are particularly noteworthy in REXASI-PRO, as they offer an alternative to model-based navigation algorithms that may increase performance but reduce transparency. To study them in simulation, we developed a machine-learning extension of Navground where users can train and deploy navigation policies. This allows for the creation of mixed groups where some agents adhere to the ML policy while others employ model-based behaviors. For instance, we conducted experiments where smart wheelchairs learn to move among pedestrians whose navigation was governed by rules derived from the social science literature.

Implicit and explicit communication

Proxemics, a form of implicit communication, governs human perception and use of space. For instance, standing in front of a doorstep implicitly communicates our intention to enter. Similarly, advancing towards someone indicates our intention to interact. This form of communication is crucial for coordinating among individuals, in particular during navigation. Therefore, it should

be considered when developing navigation algorithms for smart wheelchairs that promote local coordination, such as the formation of lanes of flow. In situations with limited free space and increased probability of conflicts, people resort to more explicit rules and forms of communication. For instance, they may nod or wave to signal precedence, similarly to how drivers use turning lights. In REXASI-PRO, we addressed this challenge in two ways. We trained machine-learning models that simultaneously learn to communicate and to communicate with neighbors. Additionally, we designed model-based communication strategies that make the smart wheelchair's intents more explicit, such as projecting the desired direction or trajectory on the floor.

Testing with humans-in-the-loop

Following the evaluation of navigation algorithms in simulation, it is crucial to validate their performance in a real-world setting. In fact, simulating human behavior may introduce inconsistencies as behavioral models are inherently incomplete, failing to encompass all potential navigation or interaction scenarios. To address this, we developed tools that connect a realistic robotics simulation, backed by Navground, to Virtual Reality visors, which enable users to immerse themselves within a virtual environment populated with navigating agents. This approach allows us to equip virtual smart wheelchairs with diverse navigation algorithms and assess their behavior in the presence of real individuals who are free to move independently. Additionally, the setup enables us to assess how people perceive the robots' behavior and the effectiveness of explicit communication and coordination strategies.

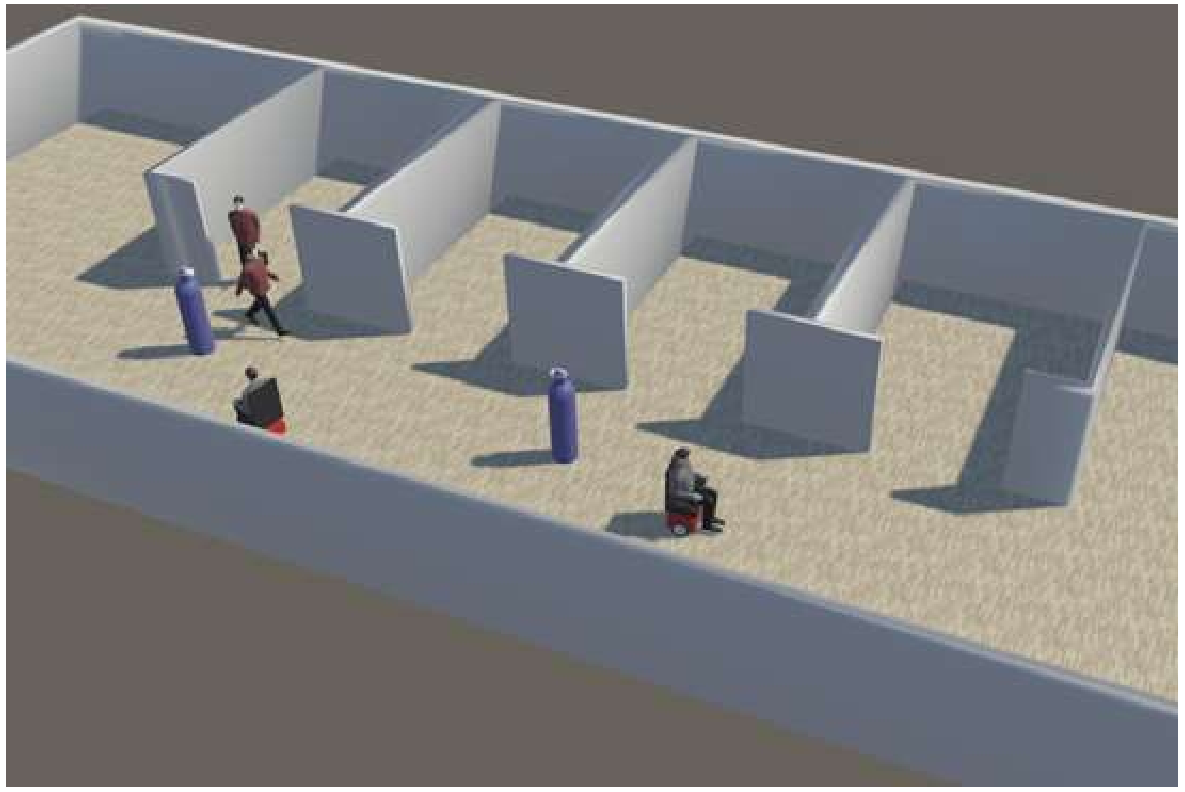

Simulation of a navigation scenario with two virtual smart wheelchairs, two virtual pedestrians, and the blue avatars of two users wearing a Virtual Reality visor.

Open source

We released Navground and its extensions as open source contributions (https://github.com/idsia-robotics/navground). They are distributed on PyPI and supported by a comprehensive documentation, including numerous reproducible tutorials. While they were developed to meet REXASI-PRO requirements, we think that they will facilitate the development and benchmarking of generic robotics navigation.

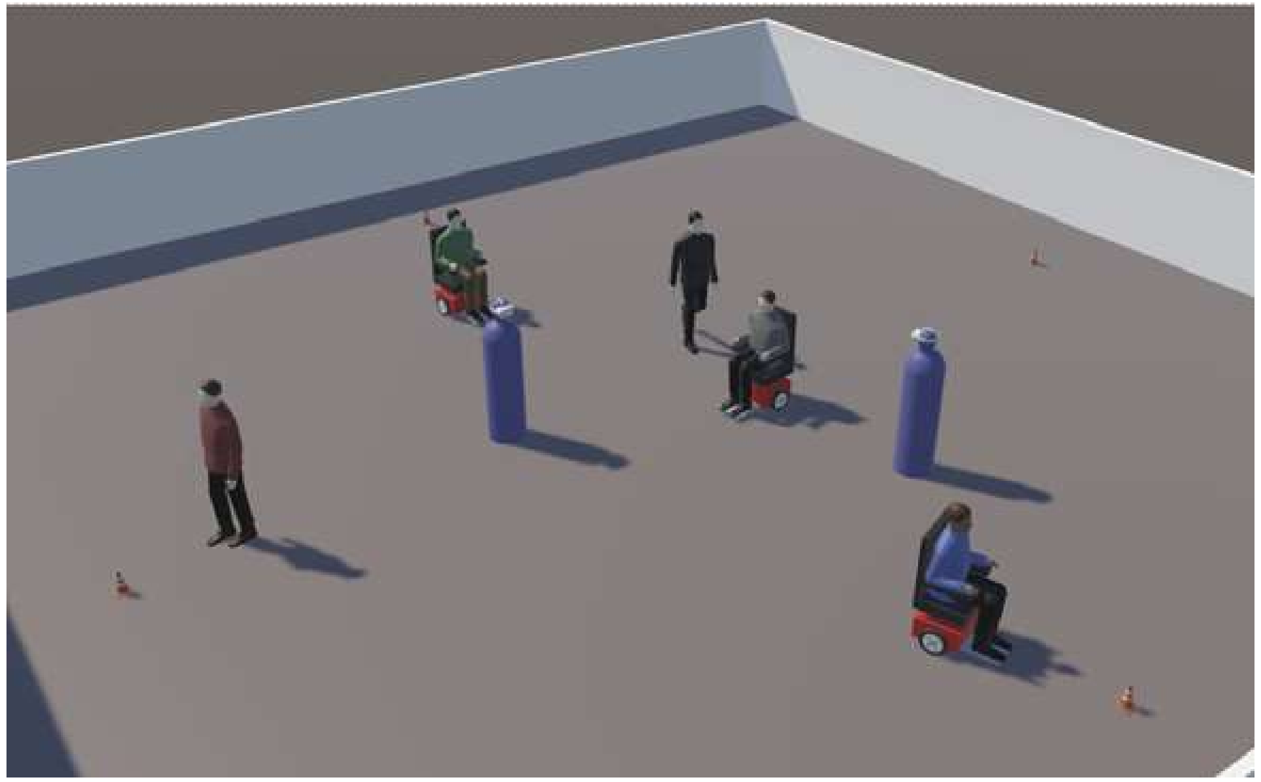

Simulation of a navigation scenario that test the crossings of simulated smart wheelchairs with simulated and real (blue avatars with VR visor) pedestrians.

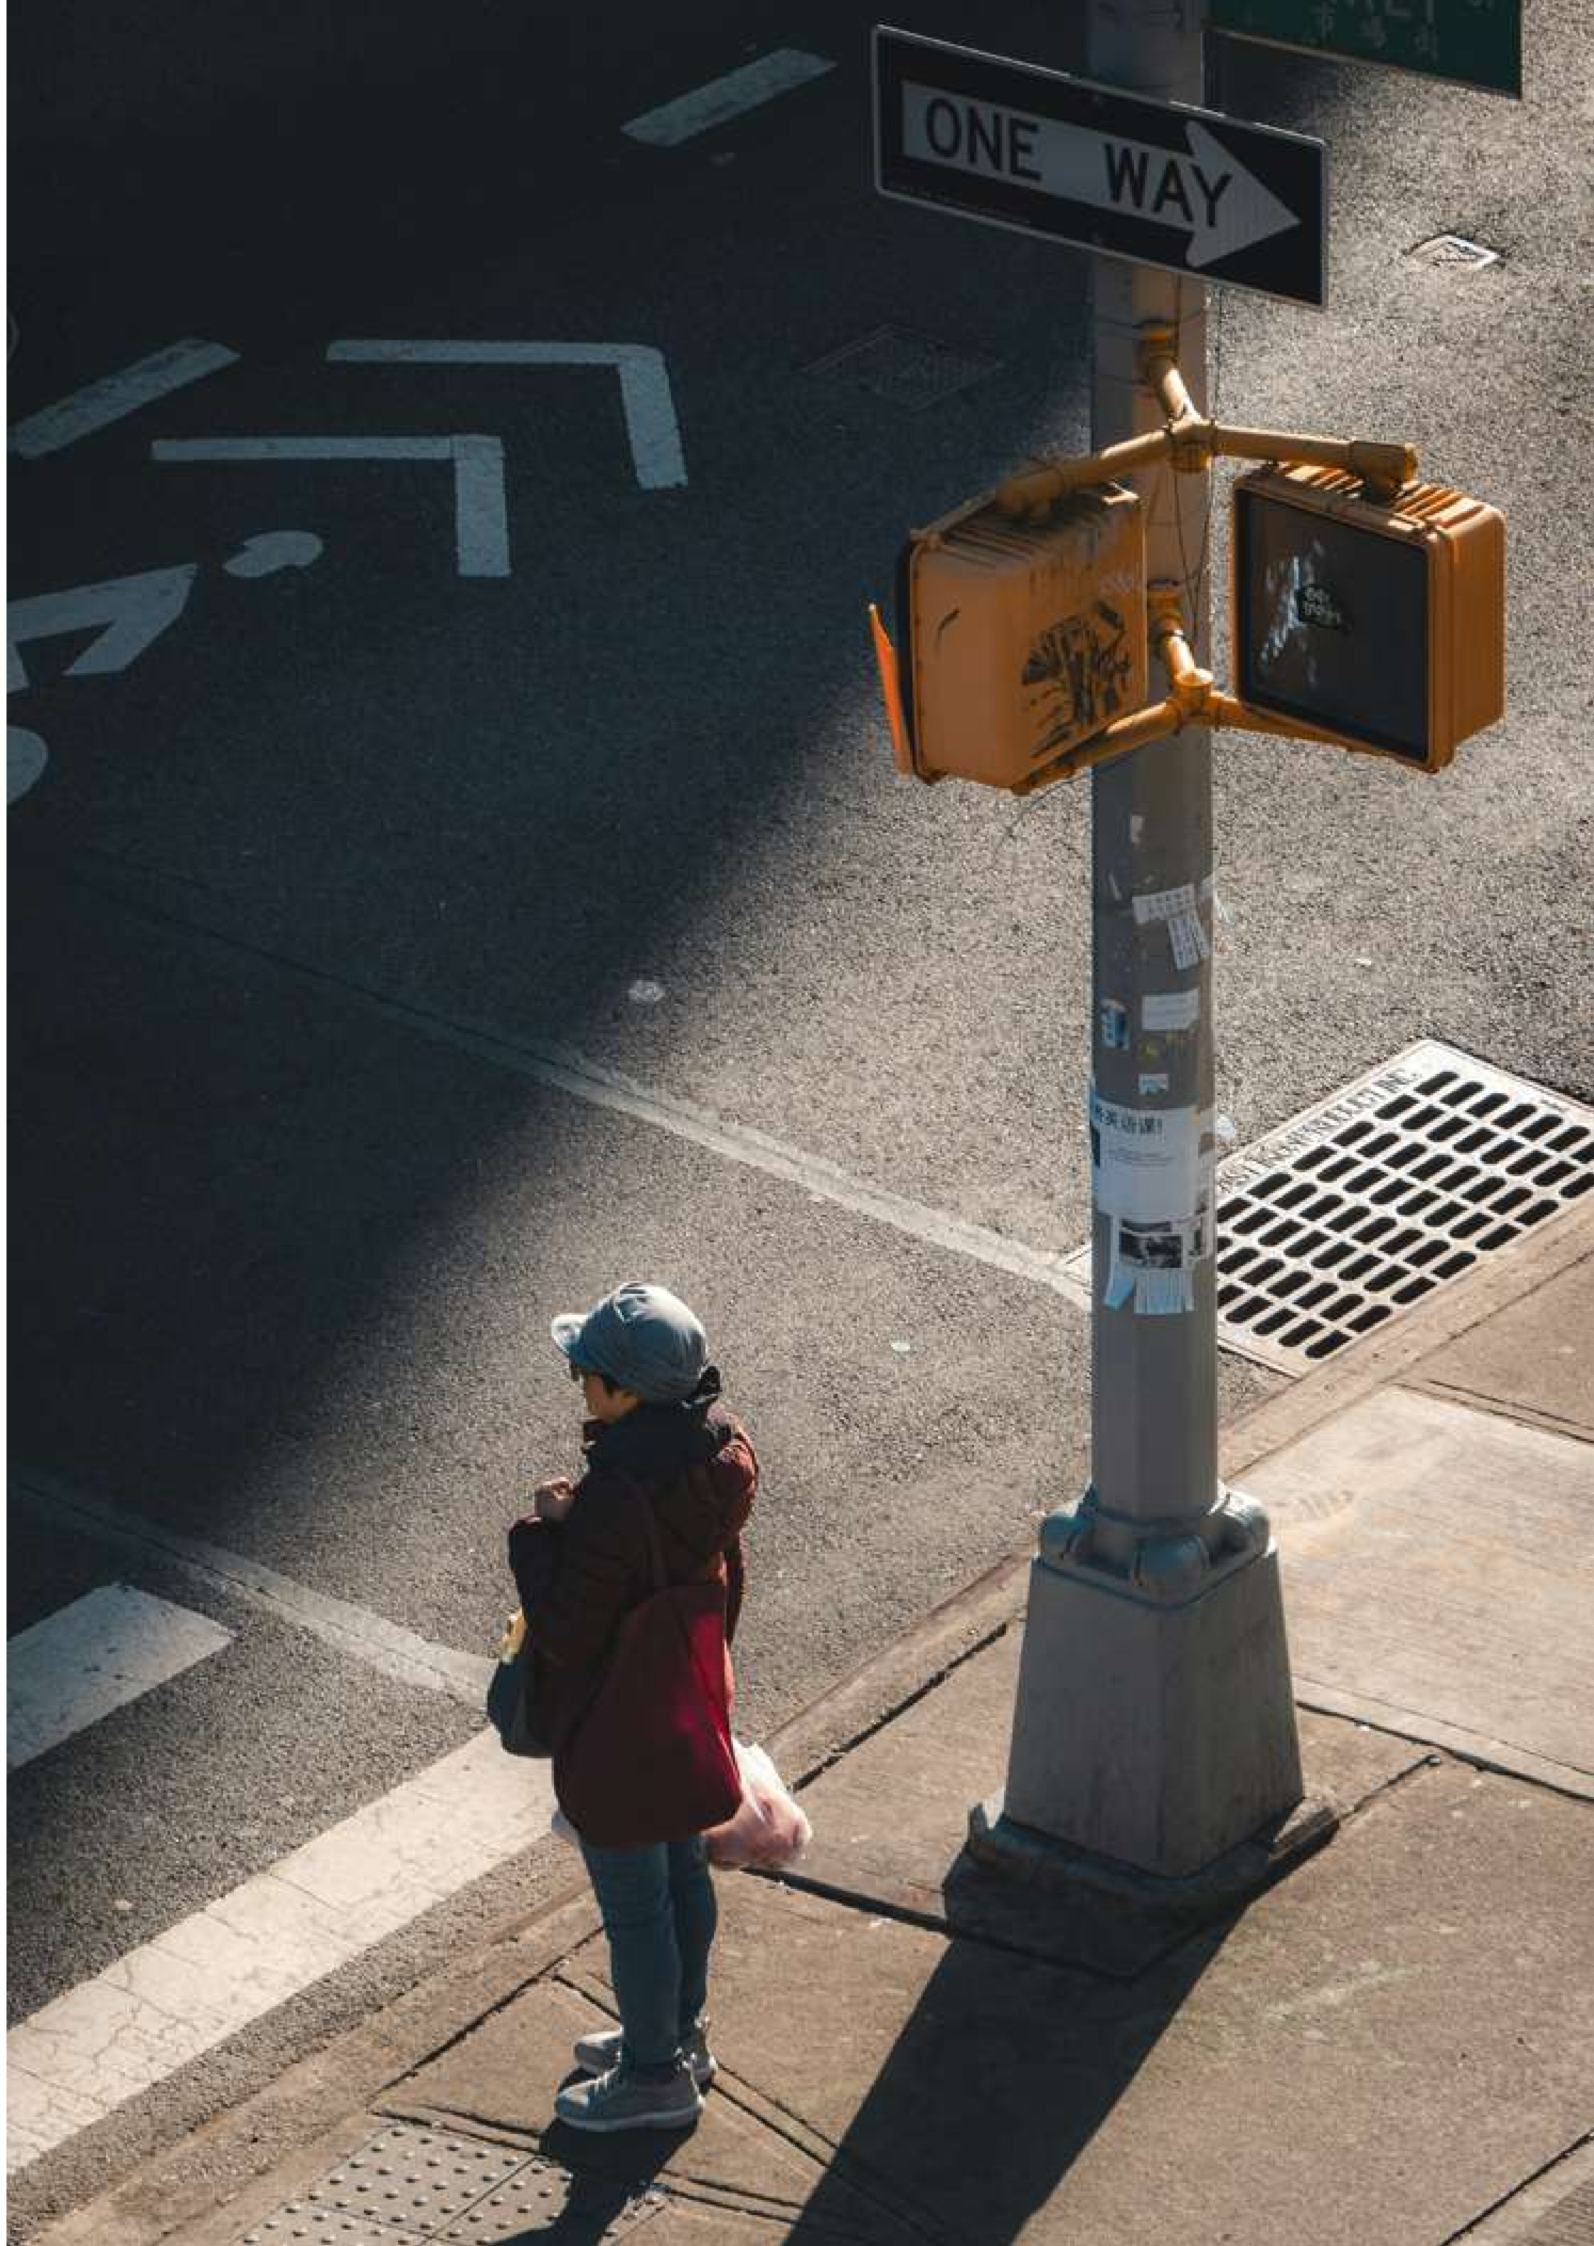

Model compression for Green AI

SIE

SIEMENS

Siemens (SIE) is a global leader in healthcare, electrical, and electronic industries, providing advanced infrastructure and sustainable green technologies. Siemens Advanta develops hardware and software systems, as well as ICT solutions and services, driving innovation within Siemens. With broad expertise in software development, it operates across multiple sectors, including healthcare, energy, mobility, and industrial manufacturing.

Green AI: a new frontier

In an era where artificial intelligence is rapidly reshaping industries and societies, a critical new frontier is emerging: Green AI.

This movement addresses the growing concern over the substantial environmental footprint of developing and deploying AI models. As AI systems become increasingly complex and powerful, their demand for computational resources—and consequently, energy—has surged, leading to significant carbon emissions. Green AI, therefore, represents a fundamental shift towards developing more sustainable and energy-efficient artificial intelligence. It focuses on creating models and systems that are not only powerful and accurate but also mindful of their ecological cost, championing efficiency in computation, data usage, and hardware to ensure that the progress driven by AI does not come at the expense of our planet.

In the World of Artificial Intelligence, a network can become more efficient by compressing it. The compression process involves making a trade-off between performance and efficiency.

The 3 main reasons of why efficiency is important for every type of artificial intelligence model are:

- **Costs:** The economic benefits of AI efficiency are a powerful driver of technological adoption worldwide. For businesses, running compressed and efficient models significantly cuts costs by reducing expenses for cloud computing and enabling the use of cheaper, less powerful hardware in manufactured goods. This dual reduction in operational and production costs has a direct, positive impact on the end-user. These savings translate into more affordable retail prices for smart devices and lower subscription fees for AI-powered services. Ultimately, AI efficiency accelerates the democratization of technology, making advanced capabilities more financially accessible to a global consumer base.
- **Environment and Pollution:** As artificial intelligence becomes globally essential in 2025, its environmental impact is a critical concern. Inefficient AI models drive this problem by consuming vast amounts of electricity for processing, much of it from a global grid still reliant on fossil fuels. However, the environmental costs

are also indirect, including the immense energy required to cool the data centers that house the hardware, and the "embodied carbon" from the resource-intensive manufacturing of powerful computer chips. AI efficiency directly tackles these issues. By requiring less computational power, efficient models reduce direct energy use, lessen the need for cooling, and can run on simpler, less environmentally costly hardware. Therefore, developing efficient AI is a core part of the Green AI movement and a global responsibility for sustainable technological progress.

- **User Experience:** The impact of a more efficient network on the user experience is profound, extending far beyond a simple increase in speed to fundamentally shape how a user perceives and interacts with a product. In the world of human-computer interaction, responsiveness is paramount, and the primary enemy is latency—the perceptible delay between a user's action and the system's response. An inefficient, computationally heavy AI model introduces significant latency, which manifests as frustrating lag, making an application feel clunky, unintuitive, and unreliable. For instance, consider the real-time portrait mode effect in a smartphone camera; an efficient network allows the background blur to be applied instantaneously and seamlessly in the viewfinder, creating a magical "what you see is what you get" experience. A less efficient model would result in a choppy, delayed preview that breaks the user's creative flow.

Knowledge Distillation: A Core Strategy

A pivotal strategy within the Green AI movement is the application of knowledge distillation, a technique that directly addresses the environmental impact across the entire lifecycle of an AI model. This process involves training a large, computationally expensive "teacher" model to achieve high accuracy, and then transferring its learned intelligence to a much smaller, more efficient "student" model. Comparing it with other compression techniques knowledge distillation is not subtractive, it does not involve taking parts away from a network like pruning and quantization. It is also important to note that this technique can transfer, so called: "dark knowledge", meaning that it replicates the entire thought process of the teacher, in a richer, more detailed manner. Also, the student is not architecturally constrained by the teacher.

The environmental benefits are comprehensive: while the initial training of the teacher model is energy-intensive, the resulting lightweight student model drastically cuts down on the energy consumed during the inference phase, which is where a deployed model spends most of its operational life. This efficiency cascade means lower electricity bills for data centers, a reduced need for powerful, resource-intensive hardware (thereby lessening the environmental cost of manufacturing and e-waste), and a significantly smaller carbon footprint for large-scale because of energy consumption, real-world applications that may handle billions of daily requests.

For the REXASI-PRO project, 3 different students were created, more of which are improving the operational cost of the model quite significantly. The teacher model is generically named "Teacher", and the students are named "Student", "Tiny-Student", "Micro-Student".

The architecture of all students and teacher was selected as the DR-SPAAM network, where for the teacher model we use the identical network and for the "Student" we used half of the original architecture features, for the "Tiny-Student" we used a quarter of the original architecture features and for the "Micro-Student" we used one eight of the features.

The models were created using knowledge distillation techniques. We have experimented with 5 different methods, out of which one method was proven to perform better. The method is named "Knowledge Review" and implements the idea of training a given level of student network by using information from multiple layers of the teacher network. Knowledge review uses residual learning to distill the knowledge. This approach is in contrast with traditional knowledge distillation

methods where the information is distilled on corresponding levels. In specialized literature, an analogy to the human learning was done, where reviewing previously learned material is crucial for deeper understanding. Furthermore, attention-based fusion is used for the student network to focus on the most relevant information from the teacher and it is used with a hierarchical context loss designed to capture and transfer contextual information at different levels.

Results and Impact

Additional changes were conducted to the method where for the distillation a dynamically weighted mean squared error loss was used. This loss penalizes the student even more on cases that are not considered learned.

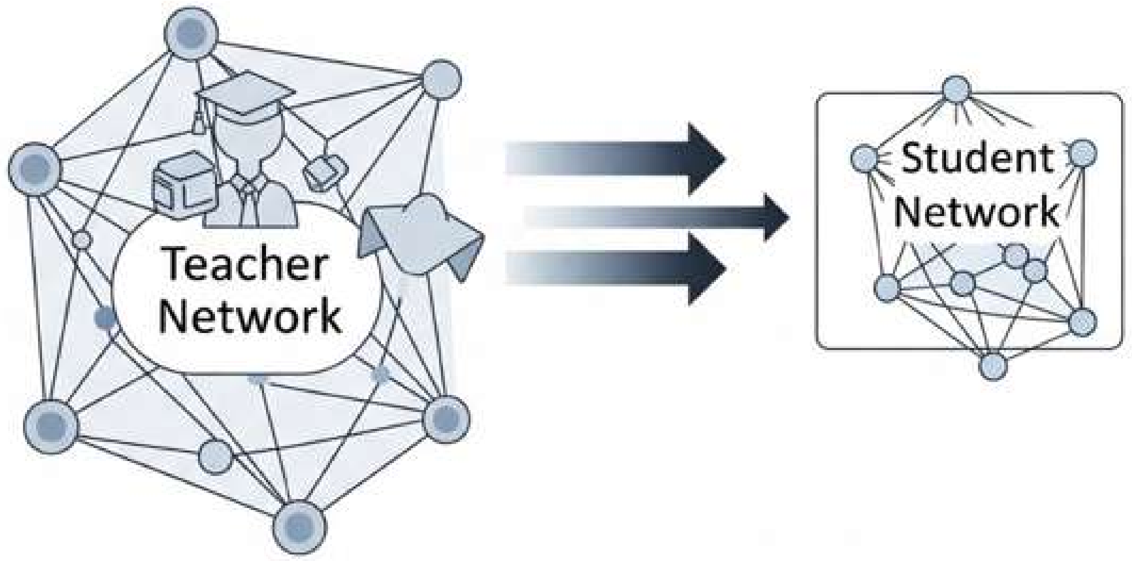

To evaluate the results, the same evaluation was used as the original DR-SPAAM network. The metrics used were AP (average precision), F1-score and EER (equal error rate). The results were highly satisfactory because it is quite rare that a student network improves the performance of its teacher. This was achieved by "Student" network which has double the efficiency of the Teacher's operational cost and had improved the performance. The "Tiny-Student" has about the same performance and four times better efficiency, whereas the "Micro Student" sacrifices some of the performance but has 8 times better efficiency than the "Teacher". The results are a big improvement in speed and computational cost. This aligns with one of aspects of the Green AI movement that implies AI to be more sustainable.

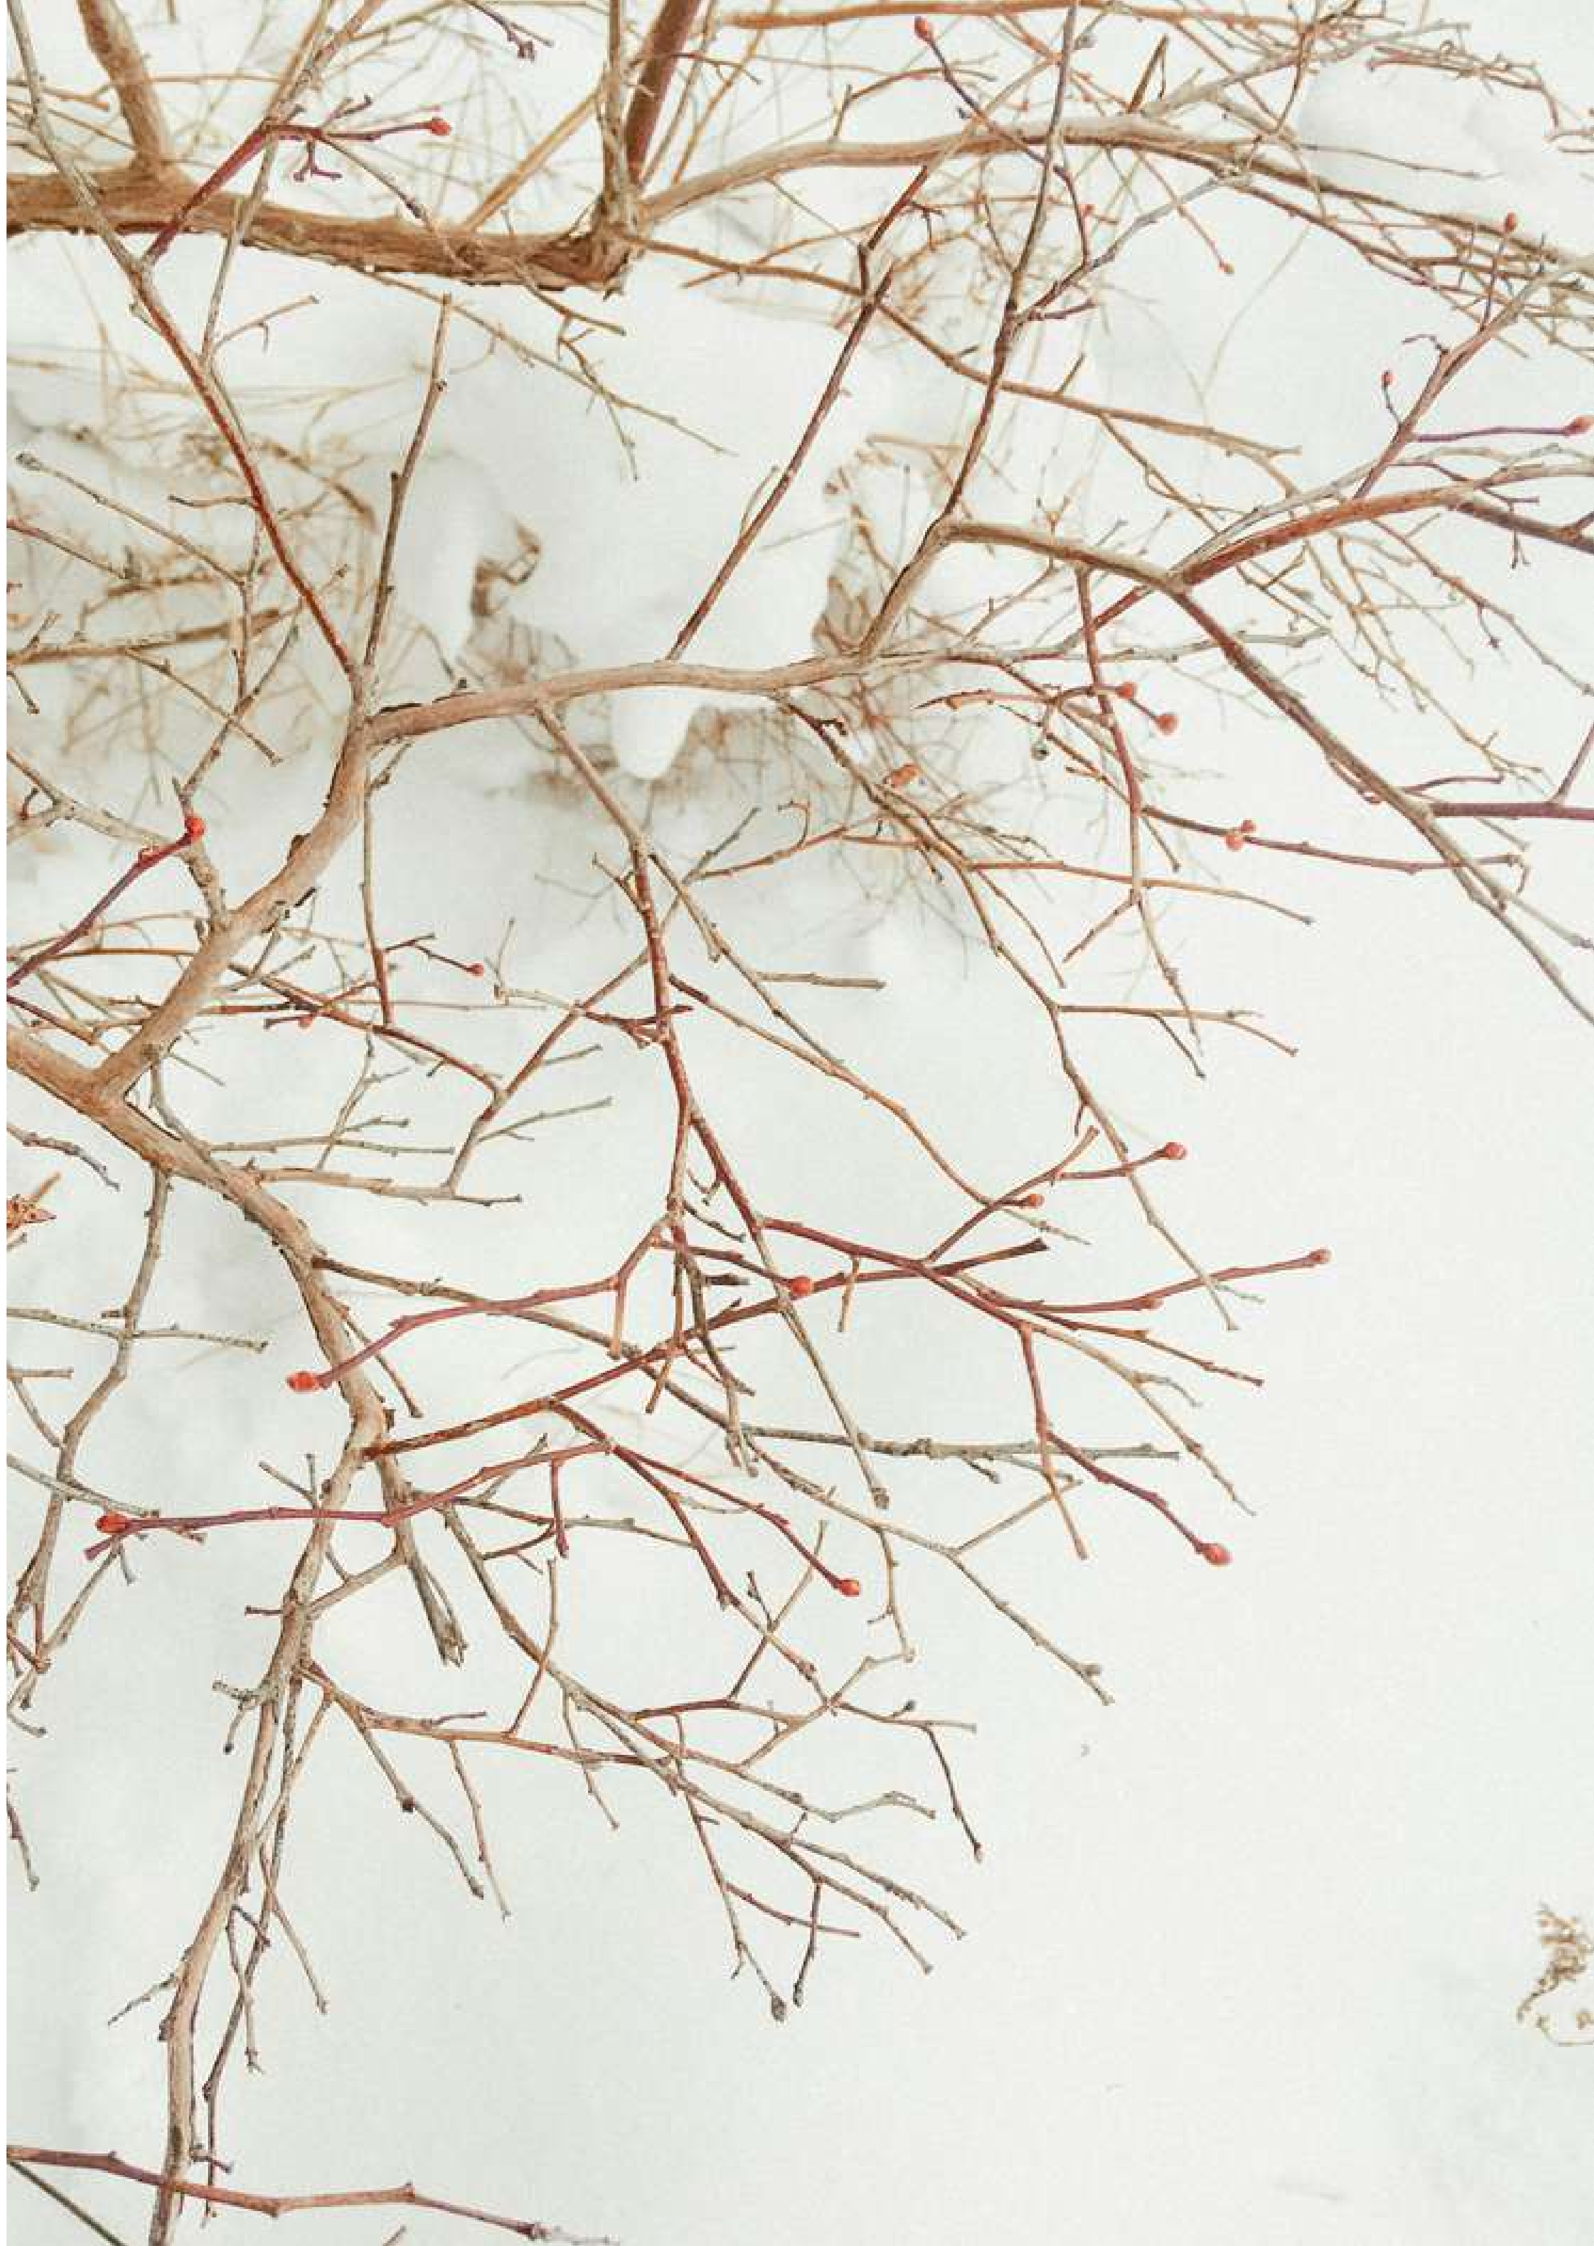

How topology aids in understanding social navigation and optimizing data selection for AI model training.

University of Seville

The University of Seville (USE) is a public institution whose mission is to foster learning and to conduct, promote, disseminate, transfer, and enhance research activities across the main fields of knowledge and their applications for scientific, technological, economic, and social development.

How Topology Aids in Understanding Social Navigation and Optimizing Data Selection for AI Model Training

Safe navigation for smart wheelchairs is a critical challenge in an era where mobility assistance aims to be increasingly autonomous and reliable. In this context, topology [2,3], a fascinating branch of mathematics, emerges as a powerful tool to address two fundamental aspects of this goal: 1) optimizing data selection for artificial intelligence (AI) model training and 2) analyzing the navigation behaviors of agent fleets. Using topological data analysis tools like persistent homology [3] and persistent entropy in the context of fleet behavior modeling is

to achieve a more nuanced understanding of system dynamics. These tools enable the characterization of distinct patterns and structures in the movement behaviors of different entities within the environment (e.g., pedestrians, wheelchairs).

Topology in the Service of Green AI and Data Selection

In the field of Artificial Intelligence (AI), there's a growing concern about the environmental impact of Deep Learning (DL) models, which often demand vast volumes of data and computational resources for their development. This trend, known as Red AI, contrasts with the Green AI [9] approach, which seeks to balance model performance with energy efficiency and sustainability. The work of USE aligns with this latter objective by exploring how topology can support data reduction without compromising model quality.

Historically, AI research has focused on improving performance at any cost, leading to substantial resource consumption. However, increasing awareness of AI's carbon footprint has led to a search for more efficient alternatives. One way to make AI more sustainable is to train models using only a fraction of the

data, commonly known as data reduction methods [7], while still maintaining high accuracy. If the selected data are representative, the training can be significantly faster, cheaper, and greener, yet equally effective.

In this context, USE has proposed a topological metric called epsilon-representativeness [8] to compare different training datasets. This measure provides a quantifiable way to assess the "similarity" between an original dataset and a reduced one. Intuitively, if we think of a dataset as a cloud of points in space, topology allows us to analyze its underlying "shape" or structure. Epsilon-representativeness indicates whether, after removing some points, the overall topological structure remains intact. The lower this value, the more representative the reduced dataset is of the original, implying that training a model on such a reduced dataset could yield similar performance to one trained on the full dataset, and better performance than one trained on a less representative reduction.

Moreover, as part of the REXASI-PRO project, the University of Seville has reviewed several data reduction techniques, comparing their effectiveness across different tasks, including classification on tabular datasets and image detection tasks aimed at identifying people and individuals in wheelchairs. These efforts are closely aligned with the project's broader goal of promoting sustainability in the use of Artificial Intelligence. The work is explored in greater detail in the publication "An in-depth analysis of data reduction methods for sustainable deep learning" [1].

Distinguishing Navigation Behaviors and Detecting Safe Simulations with Topological Methods

Beyond data selection, topology also played a key role in analyzing and enhancing the behavior of agent fleets, particularly in the context of safe navigation for smart wheelchairs. USE's objective was to make these models more reliable, ensuring simulations that were both safe (i.e., free from collisions) and efficient (i.e., without deadlocks or congestion).

To this end, USE proposed the use of topological data analysis (TDA) techniques such as persistent homology [3] and persistent entropy [4]. These methods leveraged the geometric and topological structures inherent in the data, allowing the capture of high-level spatial and relational patterns in agent behavior and configurations. Unlike traditional approaches, which often rely on handcrafted features or strong statistical assumptions, TDA offered an interpretable and robust framework capable of revealing qualitative differences in agent dynamics. The primary goal of applying persistent homology and persistent entropy in this context was to better understand the dynamics of fleet behavior. These tools helped to distinguish and quantify patterns in how entities moved and interacted in shared environments. By identifying such patterns, it became possible to predict potential conflicts, optimize wheelchair routing, and ultimately improve both safety and efficiency in navigation strategies. In parallel, USE also explored graph-based modulation techniques to identify narrow areas where congestion was likely to occur, helping to prevent crowding, collisions, or blockages among agents [5]. Additionally, topological descriptors such as persistent entropy were used to define interpretable "safety regions" within the environment, regions that could classify simulations as safe or unsafe, efficient or inefficient [6]. Matching diagrams and persistent homology were also employed to track the evolution of agent behaviors over time, enabling the identification of diverse navigation strategies across simulations.

All these tools contributed to the development of autonomous systems that were safer, more efficient, and more explainable, highlighting the added value that topology brought to the analysis of collective behavior in complex environments. This part of the work was explored in greater detail in the publication "Topology-based analysis and optimization for agent fleet behavior" [10].

References

[1] Perera-Lago J, Toscano-Duran V, Paluzo-Hidalgo E et al. An in-depth analysis of data reduction methods for sustainable deep learning. Open Res Europe 2024, 4:101 (https://doi.org/10.12688/openreseurope.17554.2)

[2] Gunnar Carlsson. Topology and data. Bulletin of the American Mathematical Society, 46(2):255–308, 2009

[3] Herbert Edelsbrunner and John L Harer. Computational topology: an introduction. American Mathematical Society, 2022.

[4] Harish Chintakunta, Thanos Gentimis, Rocio Gonzalez-Diaz, Maria-Jose Jimenez, and Hamid Krim. An entropy-based persistence barcode. Pattern Recognition, 48(2):391–401, 2015.

[5] Javier Perera-Lago, Álvaro Torras-Casas, Jérôme Guzzi, and Rocio GonzalezDiaz. The induced matching distance: A novel topological metric with applications in robotics. arXiv preprint arXiv:2502.02112, 2025.

[6] Victor Toscano-Duran, Sara Narteni, Alberto Carlevaro, Rocio Gonzalez-Diaz, Maurizio Mongelli, and Jerome Guzzi. Safe and efficient social navigation through explainable safety regions based on topological features. arXiv preprint arXiv:2006.16824, 2025.

[7] Cody Coleman, Christopher Yeh, Stephen Mussmann, Baharan Mirzasoleiman, Peter Bailis, Percy Liang, Jure Leskovec, and Matei Zaharia. Selection via proxy: Efficient data selection for deep learning. arXiv, 2019. doi:10.48550/arXiv.1906.11829.

[8] Rocio Gonzalez-Diaz, Miguel A. Gutiérrez-Naranjo, and Eduardo PaluzoHidalgo. Topology-based representative datasets to reduce neural network training resources. Neural Computing and Applications, 34(17):14397–14413, September 2022. doi:10.1007/s00521-022-07252-y

[9] Roy Schwartz, Jesse Dodge, Noah A. Smith, and Oren Etzioni. Green AI. Communications of the ACM, 63(12):54–63, 2020. doi:10.1145/3381831.

[10] Perera-Lago, J., Toscano-Duran, V., Torras-Casas, A., and Gonzalez-Diaz, R. "Topology-based analysis and optimization for agent fleet behavior". Open Res Europe 5:200 (2025), https://doi.org/10.12688/openreseurope.20760.1

Cybersecurity and Compliance: A Smarter Way Forward

V-Research

Founded in Italy in 2020, V-Research is a private R&D lab focused on cybersecurity for complex IT/OT systems. It provides risk-driven analysis, gap identification, and remediation plans. Within REXASI-PRO, V-Research contributes to advancing its cybersecurity risk assessment tool-chain from TRL 3 to TRL 5–7.

V-Research's contribution to the REXASI-PRO project

In an age where cybersecurity threats evolve faster than regulations can respond, businesses find themselves under increasing pressure to demonstrate compliance with a growing number of frameworks. This scenario presents not only a technical challenge but also a managerial one, involving high volumes of documentation, fragmented processes, and complex interdepartmental coordination.

Within the REXASI-PRO project, V-Research has addressed this issue head-on by developing Alpha, an innovative SaaS platform designed to support organizations in managing cybersecurity posture and regulatory compliance. Born from the intersection of risk assessment, automation, and governance technologies, Alpha simplifies the journey toward compliance, while also enhancing an organization's ability to detect, assess, and mitigate cybersecurity threats.

Why a Smart Compliance Framework Matters

Regulatory compliance has traditionally been associated with laborious documentation tasks, audits, and checklists. For cybersecurity in particular, it means aligning with standards such as ISO/IEC 27001, ISA/IEC 62443 NIST Cybersecurity Framework, or the more recent EU NIS2 Directive and AI Act. But organizations today face three key challenges:

- **Overload of documentation** – critical compliance data is dispersed and unstructured.
- **Fragmented processes** – security and compliance processes often lack integration across departments and systems.
- **High operational costs** – maintaining compliance the traditional way drains time and resources.

Alpha was conceived to tackle these pain points in a unified and scalable way, becoming the digital backbone for compliance management in the context of REXASI-PRO's cybersecurity-focused mission.

From Framework to Function: What Alpha Does

Alpha is a centralized, cloud-based platform that allows organizations to orchestrate all compliance and cybersecurity processes in one place. Its architecture is designed for flexibility and scalability – key for supporting both SMEs and large enterprises. Some of its core functionalities include:

- **Multi-client management:** Ideal for consultancies or large institutions, Alpha allows clustering and segmentation of clients or departments, with tailored configurations.
- **Integrated standards support:** Alpha enables users to carry out assessments based on major cybersecurity standards like ISO/IEC 27001, ISA/IEC 624434, NIST, NIS2, and AI Act. These are broken down into structured requirements, ensuring precise and traceable evaluations.
- **Posture visualization:** Dashboards provide a real-time overview of cybersecurity posture, mapped against control requirements and associated documentation.
- **Risk and gap assessments:** For each control, Alpha supports both qualitative and quantitative risk analysis – estimating impact and likelihood while identifying critical vulnerabilities.
- **BPMN-based process modeling:** Business processes and mitigation strategies can be visualized, tracked, and refined using BPMN (Business Process Model and Notation), a widely accepted standard.
- **Role-based responsibility mapping:** Through integrated RACI matrices, responsibilities across teams and departments are clearly assigned and tracked.

Usability Meets Rigor

One of the main strengths of Alpha lies in its structured but user-friendly approach. While cybersecurity and compliance are domains known for technical complexity, Alpha provides a streamlined user experience that favors collaboration between security professionals, auditors, managers, and compliance officers. Key features such as:

- Mitigation task management
- Document-to-control linking
- Comment and approval workflows

ensure that communication is traceable and decisions are documented in a transparent manner. This is especially valuable in audit contexts or when preparing for certifications. Furthermore, Alpha's support for iterative risk evaluation (including residual risk tracking) ensures that cybersecurity becomes an ongoing and dynamic process – not a one-off compliance event.

A Strategic Fit for REXASI-PRO

REXASI-PRO's mission to build ethical, explainable, and predictive tools for managing prediabetes involves handling sensitive health data and advanced machine learning models. In this complex ecosystem, the **cybersecurity and compliance layer is fundamental**. Alpha contributes by:

- Providing a structured environment for documenting compliance with NIS 2, AI Act, and health-related cybersecurity standards.
- Offering a scalable tool that supports project partners (from clinics to tech developers) in managing their internal and shared risk assessments.
- Facilitating secure and auditable documentation workflows required by ethics committees, regulatory bodies, and medical data governance authorities.

V-Research's role was not limited to engineering a SaaS tool: it was about **designing a framework for digital trust**, tailored for multidisciplinary and multi-partner settings like REXASI-PRO.

Looking Ahead: Toward Standardization and Certification

As the European Union moves forward with tighter digital regulations (including NIS2 and the EU AI Act), solutions like Alpha are no longer a luxury – they are becoming a requirement. Alpha's structured approach positions it as a candidate to support future certification processes, enabling organizations to not only comply but **demonstrate** compliance in a measurable, auditable, and cost-efficient way.

Its modularity also paves the way for integrating further verticals; including **AI-specific compliance modules, medical device cybersecurity frameworks, or sector-specific checklists.**

Conclusion

In the ever-evolving landscape of cybersecurity and digital compliance, Alpha represents a concrete step forward. Developed by V-Research within REXASI-PRO, it brings together automation, structure, and strategic visibility; transforming compliance from a burden into a business enabler and building a resilient digital infrastructure, aligned with the highest standards of security, governance, and trust.

Exploiting the Results of REXASI-PRO Project: From Research to Market Impact

ALPHA Consult

Alpha Consult

ALPHA Consult is a European consultancy delivering tailor-made solutions supported by senior staff with significant strategic development, market assessment and business modelling experience.

The project

REXASI-PRO is an European Horizon initiative focused on developing a novel engineering framework for safe, green, reliable, and trustworthy artificial intelligence (AI) solutions, with a specific focus on the use of autonomous vehicle to assists people with reduced mobility. The project aims to create an AI-based system that enables the collaboration between autonomous wheelchairs and flying robots (drones) in order to provide seamless assistance to people with disabilities. From a technical point of view, the project will combine AI, video analytics, and real-time 3D maps that will support the drones to better understand the surrounding environment and detect potential unsafe scenarios. The combination of these technologies in a single architecture, will allow the system to increase its capability to decide in unknown situations with a fast reaction time. The use of an integrated approach involving different agents (e.g., cameras, flying robots, autonomous wheelchairs) will allow drones to explore the environment autonomously, thanks to the ability of the system to define "safety trajectory" in complex environments. In this sense, the characteristic of the system to continuous monitor the environment and update the trajectories with real-time feedback, will ensure a constant safety mobility of the vehicles in the environment.

The Key Exploitable Results

During the project, our main objectives has been the identification of potential Key Exploitable Strategies for the knowledge developed during the initiative and the assessment of their impact, both on the market and in term of financial performance.

Cleary the main use of this specific technology is in the autonomous wheelchairs industry. This sector is particularly attractive indeed, according to our analysis, the global wheelchair market size is valued at USD 5.91 billion in 2024 and is expected to reach USD 9.93 billion by 2030 with a projected CAGR of 9.06% from 2024-2030. On the same time, the market growth is expected to be driven by an increasing need for mobility assistance, particularly among the elderly population. For this reason, in the next years, the electric wheelchair market is projected to grow for different reasons: government initiatives, increment in the demand, healthcare improvements and rising expenditures.

Nevertheless, our analysis went beyond the application of the REXASI's technology for the development of autonomous wheelchairs; in fact, in addition to this, we have identified other 4 potential applications for REXASI. The first one is related to the use of the Ublique Decision Intelligence platform in the field of logistic and mobility. In particular, the new modules of this platform that has been conceived, in the context of REXSI PRO, for the coordination of agents (in this case the wheelchairs) can be adopted also in other field and one of this is in the sector of Supply Chain Managements Systems. Indeed, thought this platform and its technology, it would be possible use this technology for the management of complex supply chains. This is a very interesting application, in fact one of the main challenges for companies in the next years, will be the management of supply chains, where the difficult in their management is expected to growth in the next years for both internal factors (increase of volumes and supply chain globalization) and external factors (pandemics, wars and geopolitical challenges).

The second potential application is the use of this technology for the coordination of swarm drones used in the indoor environmental mapping industry. Indeed, the technology developed in the project enables the coordination of autonomous wheelchairs in specific indoor environments, nevertheless the same technology can be applied also on agents that are not constituted by wheelchairs but by drones. In this sense, thanks to the application of this specific technology it would be possible to coordinate flying drones with the purpose to navigate and generate 3D models of specific indoor environments (e.g. buildings). This technology performs well in extremely dangerous scenarios where there is no ambient light, no satellite signal, no radio communication, and no line of sight to the pilot. For this reason, these drones can be used in different and complex scenarios such as mining exploration, building constructions, search and rescue operations and dangerous plant monitoring.

In addition to this, the application of REXASI PRO's technology can be applied also in the industry of cybersecurity systems. In this sense, the increment of cyberattacks and regulations have caused the rise in the demand of this kind of system. For this reason, according to the trends of the market, during the project has been developed a complex Cyber-Physical Systems (CPS) that integrates not only Operational Technologies (OT) and hardware components but also leverages AI and SaaS (Software as a Service) to enhance the system's capabilities. These cybersecurity systems will be based on a yearly subscription, and it will address the Critical Infrastructure Operators impacted by the NIS2 regulations.

Finally, during the project it has been developed a video analytics platform whose objective was to monitor and support the coordination of wheelchairs in indoor environments through the use of intelligent cameras and a specific platform based on AI. This technology can be easily applied also in other contexts, such as the video surveillance industry. Indeed, through the use of this platform and intelligent cameras, it would be possible to identify and monitor specific environments leveraging on real-time data. In this sense the possibility to leverage on real-time data and on intelligent cameras based on Artificial Intelligence, will enable users to alert public authorities in a faster way and to guarantee an effective and reliable area monitoring.

About Alpha Consult

ALPHA Consult is a European consultancy delivering tailor-made solutions supported by senior staff with significant strategic development, market assessment and business modelling experience. Alpha's key competences lie in: 1) Exploitation: business plans, impact assessments, market / competition assessments, go-to-market strategy, IPR; 2) Communication & Dissemination: on-site promotion (events, workshops), on-line promotion (websites, social media); 3) Support to Project Management: quality assurance plans, risk management, ethics.

Alpha Consult plays a pivotal role in the project, defining the business plan and exploitation strategy, driving the REXASI PRO's impact, and ensuring its innovations transition from research to market.

For more information on Alpha Consult, please visit https://alphacons.eu/.

Funded by
the European Union

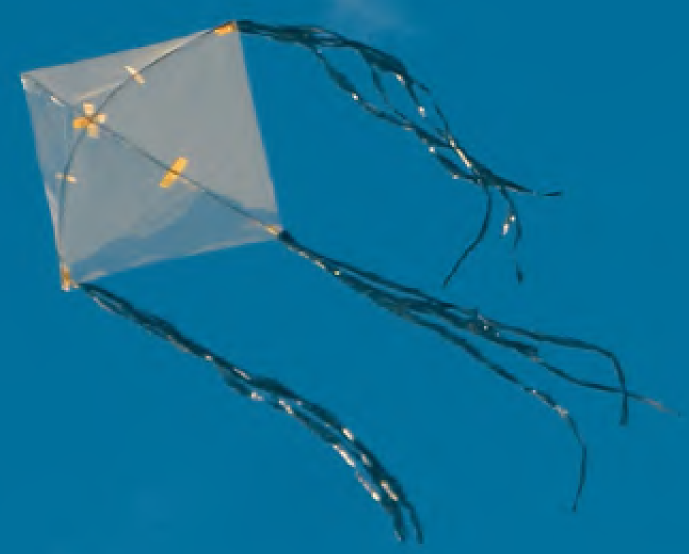

This project has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement No 101070028-2.

OVER DATA.

Research Project

A magazine on research and development topics of interest to the Spindox Group.

More information about the research project

info-rexasipro@spindox.it
rexasi-pro.spindoxlabs.com

Contact us

info@spindox.it
www.spindox.it

SPINDOX
MAKE A MARK