

Dynamic and probabilistic safety zones for autonomous mobile robots operating near humans

Fabrice Saffre^{a,*}, Hanno Hildmann^{b,c}, Eetu Heikkilä^a, Timo Malm^a, Daniel Pakkala^a

^a VTT Technical Research Centre of Finland Ltd, Finland

^b Autonomous Systems & Robotics, TNO Netherlands Organization for Applied Scientific Research, the Netherlands

^c Smart Sensor Systems Lab, THUAS The Hague University of Applied Sciences, the Netherlands

ARTICLE INFO

Keywords:

Dynamic safety
Autonomous systems
Mobile robots
Simulation
Analysis

ABSTRACT

The inefficiency of maintaining static and long-lasting safety zones in environments where actual risks are limited is likely to increase in the coming decades, as autonomous systems become more common and human workers fewer in numbers. Nevertheless, an uncompromising approach to safety remains paramount, requiring the introduction of novel methods that are simultaneously more flexible and capable of delivering the same level of protection against potentially hazardous situations. We present such a method to create dynamic safety zones, the boundaries of which can be redrawn in real-time, taking into account explicit positioning data when available and using conservative extrapolation from last known location when information is missing or unreliable. Simulation and statistical methods were used to investigate performance gains compared to static safety zones. The use of a more advanced probabilistic framework to further improve flexibility is also discussed, although its implementation would not offer the same level of protection and is currently not recommended.

1. Introduction

For purpose of the present study, a safety zone is defined as an area in which special rules, such as a permanently reduced speed (in the static case) or mobile exclusion zones around pedestrians (in the dynamic case), apply to the movement of autonomous machinery (e.g., delivery robots or self-driving forklifts) to ensure the safety of the people present. Therefore, as such, a safety zone does not imply any restriction to human access.

1.1. Motivation

According to Ref. [1], more than half of the work-related fatalities of 2021 in the EU happened in industries where humans work alongside increasingly autonomous machinery. In this context, safety zones are an essential component of many industrial operations across a wide variety of activity sectors, from mining to shipping, from forestry to manufacturing. Yet by definition, a safety zone is detrimental to efficiency, performance and/or comfort, since it is necessarily an area in which some restrictions apply to ensure the safety of people working at the area. In some cases, it may be completely off-limits to some personnel or equipment (e.g. manufacturing robots isolated in physical

cages inside a factory), in others some actions that may be desirable are forbidden (e.g. light a fire), others that may be burdensome are mandatory (e.g. put on a helmet and a high-visibility jacket). Other aspects may include observing lower speed limits etc. In other words: safety zones have an operational cost, which individuals, companies and/or other organisations are willing to pay in order to avoid accidents and injuries that would come at an even heavier price.

Precisely because of the potentially heavy cost of accidents and injuries [2], safety zones tend to be managed conservatively. As a result, they often are larger, last longer and impose heavier restrictions than strictly necessary. In the “old” world, where real-time situational awareness was severely undermined by technology limitations, this was not only justified but unavoidable: if you cannot follow or predict an employee’s whereabouts precisely and with a high degree of certainty, you are forced to make sure that everywhere they could be is safe for them. This has led to current practice in different industries to set up static safety zones, which isolate people from automated production processes or operations. This practice has some drawbacks considering the current technological development. For example, efficient and dynamic co-operation between people and autonomous robots (augmented intelligence/capabilities of workers) is not possible if people and their intended collaborators are isolated physically from each other with

* Corresponding author.

E-mail address: fabrice.saffre@vtt.fi (F. Saffre).

<https://doi.org/10.1016/j.rineng.2024.102731>

Received 5 July 2023; Received in revised form 24 June 2024; Accepted 12 August 2024

Available online 14 August 2024

2590-1230/© 2024 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

static safety zones. We argue that considerable advances in sensors, communications and machine intelligence over the past few decades have rendered the static and deterministic safety zone paradigm obsolete and that the time has come to redesign the methods by which accidents and injuries are prevented at industrial production sites. Such a redesign can and should be evaluated in terms of the potential efficiency and performance benefits compared to the current static safety zone paradigm.

We are proposing to develop the concept of dynamic and probabilistic safety zones, informed, created, maintained and removed on-demand by leveraging current and soon-to-be available technology (e.g., 5G communications and advanced location tracking systems). With regard to these enabling technologies, we do not claim expertise in any specific approach or product line, the envisioned concepts are hardware agnostic. The primary motivation for this is to improve efficiency and performance of business operations by increasing flexibility and lifting unnecessary restrictions, for instance by “suspending” the effect of a safety zone when it is known with absolute certainty that the reasons for its existence have ceased to apply (e.g., there is no human being in the designated area). It may however be worth emphasising that the proposed approach could also result in extending the boundaries of a pre-defined safety zone if the situation requires it (e.g., an employee has been detected outside the restricted area in which special rules normally apply to ensure their safety). So, far from compromising safety, increased flexibility would be accompanied by further risk reduction.

1.2. Contributions of this article

The main contribution of the paper is not to demonstrate that dynamic and probabilistic safety zones yield efficiency gains, which is largely self-evident. Rather, the objective is to present a quantitative method to accurately measure the extent of the performance increase in different regions of the parameters space, thereby providing the decision-maker with the information needed to determine whether the cost/benefit ratio for their implementation in specific circumstances is sufficiently low. This is illustrated by considering a simple collision prevention scenario in which the relevant parameters are the density of the pedestrian population and the number of robots.

2. State of the art

2.1. Safety challenges of autonomous machinery

In recent years there has been a sharp increase in the industrial, commercial and professional use of mobile machines [3], which operate with increasing levels of autonomy. Such autonomous mobile robots (AMRs) [4], also referred to as “highly-automated heavy-duty mobile machinery” (HDMM) [5] offer a huge potential when it comes to improving efficiency and productivity, but they also pose a novel threat to the human personnel alongside which they are deployed: Villani et al. [6] identified safety issues as the primary main challenge for realizing human robot collaboration. And indeed, the New Machinery regulation proposal [7] (Proposal for a Regulation of European Parliament and of the Council on machinery product. Brussels, 25.1.2023) requires that autonomous mobile machinery products operate in an enclosed zone or object detecting devices are applied to prevent risks to health and safety, in the vicinity of the autonomous mobile machine. It is (will be) the first time autonomous mobile machines are mentioned in European machinery legislation.

From the presented two options the object detecting devices idea is more relevant to the dynamic safety zone concept. The technology that can be applied on dynamic safety zones depends on the required performance of the dynamic safety zone. The dynamic safety zone area, range and shape need to be changed rapidly according to known situation and its uncertainties. This means that fences [6] or light curtains are not useful, since they have to be physically installed and modified in

order to redefine a safety zone, thus lacking adequate dynamic properties.

Traditionally, detecting humans in a work area has been addressed centrally through CCTV systems [8], which can today identify specific humans based on their gait and can follow them even though they may be occluded by obstacles periodically [9]. With growing autonomy of the employed robotic systems (and due to their increasingly mobile nature) sensing capability is moving to the device level. In practice, the sensors that can detect objects beside the mobile machine tend to be mainly on-board, potentially augmented with infrastructure-mounted monitoring devices (e.g. fixed video cameras). The most applicable sensors, which can provide long detection range are LIDAR [10–12], RADAR [13,14], laser scanner and [15] camera technologies [11,12]. Several sensors are needed because overlapping of detection zones is required (for safety and reliability reasons), all possible blind spots must be covered and having a diversity of sensors helps compensate for the weaknesses of each sensor type. For example, dirt and water tend to weaken detection properties of light-based sensors. The number of object-detecting sensors depends on machine size but is usually a small number, partly for economic reasons; e.g., Garigipati et al. [12] report on performance comparison of SLAM algorithms using a combination of LIDAR, stereo camera, inertial measurement unit (IMU) and GNSS antennas. Small indoors mobile robots can cope with two or three sensors; e.g. Ref. [16], reports on an industrial mobile robot using two laser scanners. Some sensors can also be attached to the infrastructure, e.g., at crossings to track oncoming machines and persons. One additional way to avoid collisions is to tag all moving objects (persons and machines) and then keep separation distance adequate between all moving actors. In a centralised control scenario, this requires communication to fleet control, which keeps track of separation distances and commands the machines [17]. Fleet control cannot anticipate fortuitous hazards such as fallen rocks, which must be detected to avoid collisions. Once detected, a problematic rigid object can be marked as an obstacle on the digital terrain map, resulting in continuous updates.

One additional method for handling interaction between relevant agents would involve using state-of-the-art radio signal analysis techniques (which are becoming available, e.g., as a result of the deployment of 5G networks), e.g. Ref. [18], to build and maintain a real-time map of their whereabouts. It is unclear at this stage if the accuracy of such positioning methods would be sufficient to meet the stringent requirements of a safety zone management system but combining them with the probabilistic approach that we have developed may result in a viable solution.

There are some uncertainties when trying to detect objects using only on-board sensors. The causes for these uncertainties are, for example, uneven ground causing swaying of sensors and their detection zones, and adverse detection conditions. It is tricky to detect persons leaning against a wall or large object, or someone lying on the ground. It is also difficult to factor in additional dimensions of a machine, with certain elements (e.g., boom or load) potentially extending beyond the normal envelope. Environmental factors such as dust, water flows and darkness may also cause problems to some sensors [17]. Despite those limitations (which also apply to static safety zones), when cross-referencing enough reliable sources of sensor data, dynamic safety zones appear feasible. Obviously, safety aspects need to be evaluated on a case-by-case basis, with all possible sources of error (blind spots, capabilities of sensors, environmental conditions), and the characteristics of mobile machinery (momentum, braking distance etc.) taken into consideration.

2.2. Dynamic safety zone concept in research literature

In the research literature, the concept of dynamic safety zones has been studied primarily in the manufacturing industry. Mostly the research is related to industrial and collaborative robots. Especially in collaborative robotics, various dynamic approaches based on different sensors and algorithms have been developed and tested. For example,

Malm et al. [19,20] have developed a dynamic safety system based on monitoring the position and speed of the robot and human operator. Similarly, Byner et al. [21], Vicentini et al. [22] as well as Scalera et al. [23,24] have proposed approaches where the size of the safety zone around an industrial robot is determined based on various combinations of the robot and human positions, velocities, and trajectories. Kanazawa et al. [25] have developed a probabilistic approach for collaborative robots, focusing on probabilistic prediction of the human's motion for collision avoidance. Dynamic safety systems for industrial robots have also been developed in the automotive manufacturing, for example by the Volkswagen group [26].

In these use cases related to industrial and collaborative robotics, however, the robots are stationary, and their actuators can only move within the range of the physical dimensions of the robot. In mobile robotics, the complexity further increases as not only humans, but also the robots themselves, are moving simultaneously. In these types of more dynamic environments and involving mobile robots, dynamic safety zones have been considerably less studied. One example is by Adam et al. [27] who have proposed a dynamic safety system for mobile robots. However, this system is rule-based as opposed to the dynamic and probabilistic approach proposed in this article.

In mobile robotics, a lot of research has been directed towards route planning and route optimization of robots or robot fleets [3,4], also in environments with humans. However, these studies are typically not related to heavy industrial robotics nor the fulfilment of machinery safety standards and requirements. Instead, the mobile robots rely on on-board collision avoidance systems for safety.

3. Material and methods

Although we would argue that dynamic safety zones are applicable to various domains and therein to a range of specific scenarios in which the location and whereabouts of human personnel may be subject to uncertainty (due to e.g., signal blocking obstacles as found in underground mining operations or for deployment in contested or signal denied environments in the defense domain) [3], we chose to illustrate the concept and conduct our performance analysis in the more generic and relatable case of autonomous mobile robots sharing an urban environment with humans. A suitable example is found in the box-sized ground delivery robots that have become a common sight in districts of Helsinki, Finland, where the service is already available.¹

The results presented in this paper have been obtained primarily through stochastic numerical experiments (Monte Carlo simulation). When other methods are used, as with the deterministic calculation of probabilistic weights, it is clearly stated in the text.

3.1. Dynamic safety zones

In the first simulation, we look exclusively at the dynamic element: when dynamic safety zones are in effect, their boundaries are redrawn in real-time based on information the availability, completeness, timeliness, and reliability of which is assumed perfect. The proof-of-concept scenario is as follows. A fleet of autonomous mobile vehicles (e.g., robotic forklifts) must repeatedly cross the modelled environment from East to West (and back). This environment is discretised into 1024 cells, arranged in a square lattice (32×32). Whenever a robot reaches its destination, a new destination is chosen at random, on the opposite side, resulting in a "pendulum" movement. There can only ever be one robot (or human) inside a cell: in the rare event that a planned movement would break this rule, it is simply cancelled in the simulation. By default, on every time-step, every autonomous vehicle is attempting to move closer to its destination if possible. Depending on whether machine and destination are already on the same line (identical vertical coordinate)

or not, there can be one or two such preferred options. If no such move is allowed (the corresponding cell contains another robot or a person), the robot will attempt a sidestep or backtrack manoeuvre (i.e., moving to a cell that is one step further from the destination than its current location). If this option is also unavailable (e.g., because it is outside the system's limits), the unit's location remains unchanged (stop).

The static safety zone reference scenario consists in designating a vertical (North-South) corridor 16 cells wide in the middle of the environment as a restricted area. When in this zone, autonomous vehicle movement is limited to half-speed, which is simulated by allowing robots to move only every other time-step (see Fig. 1(A)). The safety zone is meant to represent a pedestrian crossing: a predetermined number of humans are using it to traverse the environment from North to South (and vice-versa).

The movement of pedestrians is handled in a manner similar but not identical to that of autonomous vehicles: on every time-step, every simulated person has a probability P of trying to move into their main direction of travel (North or South, depending on where they entered the environment). By default, $P = 0.75$. In the $1-P$ fraction of cases when they don't try to move in a straight line, or if the path is blocked, a sidestep is attempted. Unlike for robots, backtracking is never an option. Pedestrians also never leave the safety zone (i.e., a side-step is only allowed if it is still inside the designated area). Upon reaching the edge of the environment, a simulated human is removed and another one "enters" at a random location selected among the north-most or south-most edge cells of the safety zone. This results in a constant pedestrian density (one of the key parameters of the simulation).

In the dynamic safety zone alternative, individual safety zones are substituted to the static pedestrian crossing corridor. At the beginning of every robot movement phase, an exclusion zone is drawn around every pedestrian (see Fig. 1(B)). No autonomous vehicle is allowed to enter that zone (i.e., the corresponding cells are removed from available options). If a robot finds itself inside such a "personal" safety zone (which may only happen if one of the pedestrians moves towards it), options are limited to the cells outside the restricted area (i.e., it can leave the safety zone but not travel within it).

The main parameters to be explored are the number of autonomous vehicles, the number or density of pedestrians and the size of individual safety zones (by default, 3×3 cells). Performance is measured as the number of trips completed by the fleet of robots over a fixed period.

3.2. Dynamic and probabilistic safety zones

In the real world, perfect availability, completeness, timeliness, and reliability of information can never be fully achieved. Consequently, redrawing the boundaries of dynamic safety zones in real-time with impeccable accuracy is likely to remain impractical. It is therefore necessary to incorporate a probabilistic element, to account for errors or gaps in data that may adversely affect situational awareness. In order to develop a feasible conceptual framework to handle such uncertainty in a principled and scientifically rigorous manner, we focused on the specific case in which errors take the shape of missing information (e.g., a malfunctioning beacon that only transmits updates at irregular intervals). The reader will understand that this methodology can be adapted to other sources of inaccuracy such as erroneous localisation, time-lag, radio shadows etc.

Concretely, in the numerical experiments, whenever a pedestrian moves from one cell to the next, a random test is performed against an error probability P^* . If the test is successful, i.e., when the randomly generated number x in the $[0,1]$ interval is such that $x \geq P^*$, the updated location is successfully registered (and used to update the corresponding safety zone). When the test fails however, the pedestrian goes temporarily "missing" and the personal safety zone is treated differently (NB: this is assumed to be a "known unknown", i.e., we do not consider the scenario in which an update is wrongly thought to have been received). Basically, the system then extrapolates the probable location of the

¹ <https://www.s-kaupat.fi/sivu/robokuljetus>.




Fig. 1. Static (A) and dynamic (B) safety zone examples for a generic urban environment (as opposed to an industrial scenario). A robot within the static safety zone is restricted to half-speed. No robot is allowed to enter a dynamic safety zone. For easier visualisation, the environment size is only a quarter of the default value used in simulation (16×16 instead of 32×32).

person from the last known location and the expected direction of movement (reminder: by default, there is a 75 % chance that pedestrians move in the primary direction of travel, i.e., northward or southward, 25 % that they make a side-step). To account for the possibility that the path may have been blocked, these values are multiplied by a factor 0.99, resulting in 1 % chance of no movement. This leads to an asymmetrical expansion of the safety zone around all possible locations of the missing pedestrian (variable statistical weights).

In case of multiple consecutive failures, the extrapolation continues, leading to an increasingly large and “diluted” safety zone (see Fig. 2). When an update is finally received and the precise location of the pedestrian is known again, this “fuzzy” safety zone is collapsed around the correct position (NB: this makes the implicit assumption that individuals are identifiable, as it is obviously necessary to determine which “reappearing” pedestrian is which previously missing one to be able to update the corresponding safety zone correctly).

In order to quantify the performance of a dynamic and probabilistic safety zones scenario, an additional metric is necessary. It is not sufficient to count the number of trips to determine if there is an efficiency gain, we must also consider how frequently the probabilistic element creates hazardous situations. This is done by counting how many times a robot enters a safety zone by mistake due to the uncertainty surrounding the actual whereabouts of pedestrians. It should be noted that, if the rule is that the probability of a person being in a given location must be

strictly zero for it not to be added to the list of cells around which the safety zone is drawn, then such an event can never occur. However, such a rigid rule could also quickly lead to system “paralysis” (i.e., the entirety of the environment being declared unsafe for robots to enter) if the value of P^* is non-negligible.

A most interesting parameter to explore is therefore the value of this threshold, e.g., should a cell in which the probability that a pedestrian is present is < 0.01 be considered safe enough to enter, or should this be < 0.0001 ? It should be noted that this value is not the probability that an accident will occur, only that of a person and a robot finding themselves in closer proximity than intended. In addition to this and the other parameters already relevant in the deterministic case, the value of P^* is obviously also a critical one to investigate.

4. Results

4.1. Dynamic safety zones

As expected, for the reference or benchmark scenario (static safety zone), performance is strongly dependent on the number of autonomous devices and only weakly affected by the density of pedestrians, since the latter only influences a robot’s movement insofar as it cannot enter a cell when it is occupied by a person (which can be interpreted as the effect of a secondary collision detection mechanism). Apart from in this




Fig. 2. Probabilistic safety zones influenced by the absence of information about one pedestrian (6 consecutive failed updates, last known location marked by the icon with a question mark). The statistical weights (based on likely direction of movement) are indicated. (A) Threshold is zero (i.e., safety zone extends around any possible location, however unlikely). (B) Threshold is 0.001, resulting in a somewhat smaller exclusion zone (but proportionally increased risk).

particular event, the actual presence of pedestrians has no effect, since the rules only stipulate that autonomous vehicles are limited to half-speed whenever inside the safety zone (whether it is crowded or empty is not taken into consideration).

The simplest case has a single robot crossing the empty environment back and forth. The average trip length in a 32×32 discrete environment is ~ 41.5 steps: minimum 31 steps (straight line), maximum 62 steps (diagonal). Half of these are inside the safety zone, in which speed is divided by two, meaning that the average duration of one trip is ~ 62.5 time-steps. The simulation run is halted after 1000 time-steps. The autonomous vehicle is therefore expected to complete about 16 trips. The average over 10000 repetitions was 15.77, due to discretisation effects (the final, uncompleted trip is not counted), in good agreement with this simple theoretical projection (benchmark).

As the number of pedestrians present at any one time in the safety zone increases, performance slowly degrades due to collision avoidance effects (i.e., the robot stopping or taking a detour to avoid a person) but the effect is weak. Although congestion manifest itself in the fact that the slope of the curve plotting the number of trips completed as a function of the number of robots is lower than 15.77 in the absence of pedestrians, the relationship remains linear in the region of the parameters space that was considered. In summary: as anticipated, static safety zones are almost insensitive to population density fluctuations in the chosen scenario.

In the dynamic case, efficiency gains are substantial when the number of pedestrians is low. This is an intuitive result: in the absence of a static safety zone, robots may travel at full speed across the entire environment unless blocked by the smaller “interdiction zones” surrounding individuals. If this is a rare event, there is obviously a net benefit. Simulation results show the average number of trips completed to be an exponential decay function of the number of pedestrians in both the static and dynamic scenarios ($r^2 > 0.99$, Fig. 3).

However, the slope is much steeper in the latter, leading to the number of trips completed falling below that observed in the static scenario for ≈ 25 pedestrians concurrently present, for the default parameter values (see Fig. 3). Additional results (not shown) also indicate that this is almost independent of the number of robots with only a very minor increase of this threshold value: ≈ 26 individuals for 20 robots instead of for the 10 robots, for which the results are shown in Fig. 3, below (intersection values calculated from exponential fitting functions).

If increasing the size of personal safety zones to 5×5 , the threshold value falls to ≈ 9 people, so almost the same ratio as between the surface




Fig. 3. Sample results illustrating the effect of enabling dynamic safety zones, as a function of pedestrian density. Every data point is the average number of trips completed after 1,000 time-steps, summed over the entire robot population, for 10,000 independent realisations. Exclusion zone is 3×3 cells around every person, the environment is 32×32 cells, with the central 16×32 area being designated as the static safety zone (see Fig. 1).

areas ($\frac{25}{9}$). The small but noticeable deviation from the exponential decay fitting (see Fig. 4, especially the “10 robots (dynamic)” scenario) is likely attributable to nonlinearities such as the increased probability of significant overlap as dynamic safety zones grow larger.

These results point to the need to limit the use of dynamic safety zones to cases in which certain conditions are met (here, sufficiently low pedestrian traffic). Obviously, this numerical experiment and its specific threshold values are meant as a proof-of-concept. The operational conditions in which the introduction of dynamic safety zones can lead to significant efficiency gains will vary from one application to the other, and the corresponding boundaries must be calculated on a case-by-case basis. Our findings merely demonstrate that this is feasible, with a high degree of confidence, provided that a sufficiently accurate model exists.

4.2. Dynamic and probabilistic safety zones

For a P^* value of 0.5, indicating that updates are received, statistically, every other step, most results are unsurprising. As expected, for the same number of pedestrians, the frequency of safety incidents decreases progressively as the tolerance threshold (probability value above which a cell is included in the safety zone) is lowered from 0.1 to 0.01 and 0.001. Similarly, the number of trips completed also decreases, albeit more slowly, as the “paralysing” effect of faster-expanding safety zones becomes stronger for lower threshold values. The number of robots (5 or 10) appears to affect both variables (frequency of safety incidents and number of trips) near linearly, as in the “perfect positioning accuracy” scenario, which is also an intuitive result. The number of pedestrians (population density) appears to have the strongest effect over the number of trips completed, emphasising the main drawback of the probabilistic approach: for a large number of potential hazards, even limited uncertainty ($P^* = 0.5$) is highly detrimental to performance, compounding the effect already observed for the dynamic scenario (cf. Section 4.1). These results are summarised in Fig. 5.

A more surprising result was that the number of safety incidents as a function of pedestrian population density goes through a maximum in certain regions of the parameters space. This pattern, clearly visible in Fig. 6, is most likely attributable to a “herding” effect: when uncertainty is high ($P^* = 0.9$), the “opportunities” for incidents (when the statistical weight of a person’s actual position falls below the interdiction threshold) increases with the number of pedestrians. However, when the crowd becomes even denser, the expanding probabilistic safety zones increasingly overlap, meaning that individuals can “protect” each other. In short: a robot may unknowingly keep its distance to one “missing” person by purposefully avoiding another. Since this “herding” effect is also directly proportional to population density, it counteracts the “law




Fig. 4. Identical representation as in Fig. 3 but with 5×5 (instead of 3×3) personal safety zone around individual pedestrians (all other parameter values are unchanged). The “static” data is the same as in Fig. 3 and is provided for comparison purposes.




Fig. 5. Average number of safety incidents and of trips completed as a function of the tolerance threshold for a constant pedestrian density of 20 (A and B) and as a function of the number of n pedestrians for a constant tolerance threshold of 0.001 (C and D). NB: the vertical scale on A is logarithmic.




Fig. 6. Average number of safety incidents as a function of the number of pedestrians n for a constant tolerance threshold of 0.001 and error probability. P^* of 0.9.

of large numbers” and brings down the probability of safety breaches.

5. Conclusion

The main conclusion of this work is that dynamic safety zones could provide a feasible solution to introduce efficient automation (self-driving vehicles and other robots) without compromising safety standards for human workers, given adequate levels of situational awareness (communication, localisation etc.) Our results clearly demonstrate that a transition from static to dynamic safety zones would yield a significant

increase in productivity in certain circumstances, e.g. when the density of relevant actors is low, and the precautionary exclusion of a large area designated as a safety zone severely complicates the navigation of autonomous vehicles (see Section 3).

However, as stated in the introduction, our intention is not to emphasise that dynamic and probabilistic safety zones have advantages, which is obvious. Rather, we wanted to demonstrate the use of a quantitative method for evaluating performance gains in a simple collision prevention scenario, emphasising the presence of a “law of diminishing return” in certain circumstances. Our results show that above a critical pedestrian population density, the use of dynamic safety zones can have an adverse effect, reducing performance (in terms of trips completed over a fixed period) instead of increasing it (see, e.g., Figs. 3 and 4). Other non-trivial findings include the presence of a maximum in the number of safety incidents recorded in the probabilistic case when plotted as a function of pedestrian density, which we interpret as a “herding” effect (see Fig. 6).

One of the most important findings is that there are reliable algorithmic methods to deal with specific types of uncertainty, which was illustrated by the use of expanding safety zones around the last known location of a human worker. It is therefore essential to categorise uncertainty and especially to understand how it can be affected very differently by lack of either confidence or precision: a high degree of confidence in some information indicates that there is a high probability of it being correct, notwithstanding its specificity. By contrast, a high degree of precision indicates that some information has a sufficient level of detail, but it does not imply that it is accurate.

Our results indicate that it is possible to maintain safety when facing uncertainty resulting from lack of precision, but only if confidence remains high. For instance, a robot may be able to guarantee that it does

not compromise the safety of Bob or Alice as long as it is 99.9999 % confident that Alice cannot be less than 8 m away and Bob less than 12 m away, even if it doesn't know their actual location. On the contrary, being 90 % confident that Bob is 15 m to the Northeast and 85 % sure that Alice is 10 m to the South does not allow to make such a determination, even though the localisation data is more precise. Accordingly, we conclude that technologies capable of increasing information timeliness and reliability (e.g., excellent synchronisation, ultra-low latency, accurate quality of service assessment), both of which can increase confidence (or flag the lack thereof when correctly identified as compromised), will be more critical enablers of dynamic safety zones than those emphasising high-precision measurements (e.g., 3D mapping of the autonomous robot's environment).

CRediT authorship contribution statement

Fabrice Saffre: Conceptualization, Data curation, Formal analysis, Funding acquisition, Investigation, Methodology, Software, Supervision, Validation, Visualization, Writing – original draft, Writing – review & editing. **Hanno Hildmann:** Validation, Writing – review & editing. **Eetu Heikkilä:** Writing – original draft, Investigation. **Timo Malm:** Investigation, Writing – original draft. **Daniel Pakkala:** Funding acquisition, Investigation, Project administration, Resources, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgements

This research was supported by the Next Generation Mining (NGMining) project funded by Business Finland (Grant no 43670/31/2020).

References

- [1] Marea de Koning, Tyrone Machado, Andrei Ahonen, Nataliya Strokina, Morteza Dianatfar, Francesco De Rosa, Tatiana Minav, Reza Ghabcheloo, A comprehensive approach to safety for highly automated off-road machinery under regulation 2023/1230, *Saf. Sci.* 175 (2024) 106517, 10.1016/j.ssci.2024.106517, <https://www.sciencedirect.com/science/article/pii/S0925753524001073>.
- [2] Okonkwo Chinedu, Okpala Ifeanyi, Awolusi Ibukun, Nnaji Chukwuma, Overcoming barriers to smart safety management system implementation in the construction industry, *Results in Engineering* 20 (2023) 101503, 10.1016/j.rineng.2023.101503, <https://www.sciencedirect.com/science/article/pii/S2590123023006308>.
- [3] Safa Jameel Al-Kamil, Robert Szabolcsi, Optimizing path planning in mobile robot systems using motion capture technology, *Results in Engineering* 22 (2024) 102043, 10.1016/j.rineng.2024.102043, <https://www.sciencedirect.com/science/article/pii/S2590123024002962>.
- [4] Nattapong Promkaew, Sippawit Thammawiset, Phiranat Srisan, Phurichayada Sanitchon, Thananop Tummawai, Somboon Sukpancharoen, Development of metaheuristic algorithms for efficient path planning of autonomous mobile robots in indoor environments, *Results in Engineering* 22 (2024) 102280, 10.1016/j.rineng.2024.102280, <https://www.sciencedirect.com/science/article/pii/S2590123024005358>.
- [5] de Koning, Marea and Andrei Ahonen and Nataliya Strokina and Reza Ghabcheloo, Safety architectures for cyber-physical systems: review of state-of-the-art and outlook for heavy duty mobile machinery, in: T. Minav, J. Uusi-Heikkilä (Eds.), *The Eighteenth Scandinavian International Conference on Fluid Power, SICFP'23*, 2023, pp. 224–238, *Scandinavian International Conference on Fluid Power*; Conference date: 30-05-2023 through 01-06-2023.
- [6] Valeria Villani, Fabio Pini, Francesco Leali, Cristian Secchi, Survey on human-robot collaboration in industrial settings: safety, intuitive interfaces and applications, *Mechatronics* 55 (2018) 248–266, 10.1016/j.mechatronics.2018.02.009, <https://www.sciencedirect.com/science/article/pii/S0957415818300321>.
- [7] Proposal for a regulation of european parliament and of the council on machinery product, Online: <https://data.consilium.europa.eu/doc/document/ST-5617-2023-INIT/EN/pdf> (01 2023).
- [8] Mohammad Naim Uddin, Hussain Nyeem, Engineering a multi-sensor surveillance system with secure alerting for next-generation threat detection and response, *Results in Engineering* 22 (2024) 101984, 10.1016/j.rineng.2024.101984, <https://www.sciencedirect.com/science/article/pii/S2590123024002378>.
- [9] D. Sangeetha, P. Deepa, A low-cost and high-performance architecture for robust human detection using histogram of edge oriented gradients, *Microprocess. Microsyst.* 53 (2017) 106–119, 10.1016/j.micpro.2017.07.009, <https://www.sciencedirect.com/science/article/pii/S0141933116302216>.
- [10] Xiao Teng, Zhenjiang Shen, Lin Huang, Hui Li, Wankai Li, Multi-sensor fusion based wheeled robot research on indoor positioning method, *Results in Engineering* 22 (2024) 102268, 10.1016/j.rineng.2024.102268, <https://www.sciencedirect.com/science/article/pii/S2590123024005231>.
- [11] Pengcheng He, Haopeng Duan, Jinyu Luo, Xinyan Du, Congyan Jia, Jin Cui, An integration tool of safety and security requirements for autonomous vehicles, in: 2023 International Conference on Artificial Intelligence of Things and Systems (AloTSys), 2023, pp. 118–124, <https://doi.org/10.1109/AloTSys58602.2023.00040>.
- [12] Bharath Garigipati, Nataliya Strokina, Reza Ghabcheloo, Evaluation and comparison of eight popular lidar and visual slam algorithms, in: 2022 25th International Conference on Information Fusion (FUSION), 2022, pp. 1–8, <https://doi.org/10.23919/FUSION49751.2022.9841323>.
- [13] Ankith Manjunath, Ying Liu, Bernardo Henriques, Armin Engstle, Radar based object detection and tracking for autonomous driving, in: 2018 IEEE MTT-S International Conference on Microwaves for Intelligent Mobility (ICMIM), 2018, pp. 1–4.
- [14] Onur Tokur, Suleiman Alsweiss, mmwave radar based approach for pedestrian identification in autonomous vehicles, in: 2020 SoutheastCon, 2020, pp. 1–2, <https://doi.org/10.1109/SoutheastCon44009.2020.9249704>.
- [15] Oscar Real-Moreno, Julio C. Rodriguez-Quinonez, Oleg Sergiyenko, Luis C. Basaca-Preciado, Hernandez-Balbuena, ~ Daniel, Moises Rivas-Lopez, Wendy Flores-Fuentes, Accuracy improvement in 3d laser scanner based on dynamic triangulation for autonomous navigation system, in: 2017 IEEE 26th International Symposium on Industrial Electronics (ISIE), 2017, pp. 1602–1608, <https://doi.org/10.1109/ISIE.2017.8001486>.
- [16] Golnaz Raja, Teemu Makkonen, Reza Ghabcheloo, Safe Control Using Occupancy Grid Map-Based Control Barrier Function (ogmcbf), 2024 10703 arXiv:2405.
- [17] T. Malm, D. Pakkala, E. Heikkilä, Autonomous Mobile Machines in Mines Using 5G Enabled Operational Safety Principles, No. 412 in VTT Technology, VTT Technical Research Centre of Finland, Finland, 2023, <https://doi.org/10.32040/2242-122X.2023.7412>.
- [18] Reza Ghabcheloo, A. Pedro Aguiar, Antonio Pascoal, Carlos Silvestre, Synchronization in multi-agent systems with switching topologies and non-homogeneous communication delays, in: 2007 46th IEEE Conference on Decision and Control, 2007, pp. 2327–2332, <https://doi.org/10.1109/CDC.2007.4434831>.
- [19] T. Malm, R. Tiisanen, E. Heikkilä, T. Ahonen, J. Sarsama, New safety concepts for autonomous mobile machines: EasyChair preprint no. 5239, in: Automaatiopäivä 24, 2021, automaatiopäivä 24 : Automaatio, kestävää kehitystä ja tulevaisuutta; Conference date: 13-04-2021 Through 14-04-2021. URL <https://www.automaatioseura.fi/automaatiopäivä24/>.
- [20] T. Malm, T. Salmi, I. Marstio, J. Montonen, Dynamic safety system for collaboration of operators and industrial robots, *Open Eng.* 9 (1) (2019) 61–71, <https://doi.org/10.1515/eng-2019-0011> [cited 2024-06-23].
- [21] C. Byner, B. Matthias, H. Ding, Dynamic speed and separation monitoring for collaborative robot applications – concepts and performance, *Robot. Comput.-Integr. Manuf.* 58 (C) (2019) 239–252, <https://doi.org/10.1016/j.rcim.2018.11.002>.
- [22] F. Vicentini, N. Pedrocchi, M. Giussani, L. Molinari Tosatti, *Dynamic Safety in Collaborative Robot Workspaces through a Network of Devices Fulfilling Functional Safety Requirements*, 2014, pp. 1–7.
- [23] L. Scalera, R. Vidoni, A. Giusti, Optimal scaling of dynamic safety zones for collaborative robotics, in: 2021 IEEE International Conference on Robotics and Automation (ICRA), IEEE Press, 2021, pp. 3822–3828, <https://doi.org/10.1109/ICRA48506.2021.9561611>, 10.1109/ICRA48506.2021.9561611.
- [24] L. Scalera, A. Giusti, R. Vidoni, A. Gasparetto, Enhancing fluency and productivity in human-robot collaboration through online scaling of dynamic safety zones, *Int. J. Adv. Des. Manuf. Technol.* 121 (9) (2022) 6783–6798, <https://doi.org/10.1007/s00170-022-09781-1>, 10.1007/s00170-022-09781-1.
- [25] A. Kanazawa, J. Kinugawa, K. Kosuge, Adaptive motion planning for a collaborative robot based on prediction uncertainty to enhance human safety and work efficiency, *IEEE Trans. Robot.* 35 (4) (2019) 817–832, <https://doi.org/10.1109/TRO.2019.2911800>.
- [26] R. Evans, Dynamic safety zones enable safe cooperation between people and robots. *Automotive Testing Technology International*, 2018. <https://www.automotivetestingtechnologyinternational.com/news/rd/dynamic-safety-zones-people-robots.html>.
- [27] S. Adam, M. Larsen, K. Jensen, U. Schultz, Rule-based dynamic safety monitoring for mobile robots, *Journal of Software Engineering for Robotics* 7 (2016) 120–141.