Safer Than Perception: Increasing Resilience of Automated Vehicles Against Misperception

Martin Fränzle[0000−0002−9138−8340]⋆1 and Andreas Hein[0000−0001−8846−2282]2

[1] Research Group Foundations and Applications of Systems of Cyber-Physical Systems
Carl von Ossietzky Universität Oldenburg, Germany
martin.fraenzle@uol.de
[2] Research Group Assistance Systems and Medical Technology
Carl von Ossietzky Universität Oldenburg, Germany
andreas.hein@uol.de

Abstract. Autonomous vehicles (AV) rely on environmental perception to take manoeuvre decisions. Safety assurance for AV thus hinges on achieving confidence in all percepts that are safe-guarding critical manoeuvres. As the safety targets for such critical manoeuvres are considerably higher than the statistical figures for the reliability of at least current learning-enabled classification algorithms within the environmental perception, we need means for assuring that the overall system is "'safer than perception"' in that the frequency of erratically adopting a critical manoeuvre is considerably lower than the frequency of individual misclassifications. We present a methodology for constructively generating reformulations of guard conditions that are more resilient to misperception than the original condition. The synthesized, rephrased guard conditions reconcile a given safety target, i.e. a given a societally accepted upper bound on erratically activating a critical manoeuvre due to a false positive in guard evaluation, with maximal availability, i.e. maximal true positive rate. By synthesizing a resilient rephrasing of the guard condition, the constructive setting presented herein complements the analytical setting from a previous companion paper [6], which merely analysed a given condition for its safety under uncertain perception.

Keywords: Safety-critical perception, Decision making, Robustification

1 Introduction

Decision-making based on — inherently uncertain to some extent— environmental perception is a key element of providing cyber-physical systems, like transportation systems in general and road vehicles in particular, with forms of autonomy, as in highly automated or autonomous driving. Such decision-making obviously is safety-critical, as the actions adopted in consequence of a decision

⋆ Supported by the Ministry of Science and Culture of the State of Lower Saxony under grant number ZN3493 as well as by Deutsche Forschungsgemeinschaft under grant no. DFG FR 2715/5-1.

have physical impact and can consequently incur risk to life, health, and property. Especially in the field of automated vehicles, societal expectations concerning the risk induced by automated driving functions, and thus ultimately for error rates decision-making, are very high: while manual driving already is amazingly safe at considerably more than a million kilometres driven on average between two accidents incurring some form of injury, the public debate as well as relevant authorities tend to require highly automated vehicles (HAV) to even further reduce the overall rate of injuries and fatalities compared to human-operated vehicles.

How rare accidents of HAV must be is a matter of ongoing societal debate, but the societal acceptance threshold will obviously be orders of magnitude below the misperception rates that can be realized by or guaranteed of [3] any perception system containing machine-learned components, which can only be trained and qualified on examples. Concerning the three types of uncertainties that these systems inevitably are prone to, namely

1. *existential uncertainty*, i.e. not knowing whether all or at least all relevant environmental objects have been detected,
2. *classification uncertainty*, i.e. uncertainty in exactly classifying the type, like "car", "adult pedestrian", "playing child", or "waste bin", of any detected object, and
3. *state uncertainty*, i.e. inaccuracies in determining relevant physical quanta, like speed or distance, of a classified object,

uncertainties especially concerning the first two remain relatively high. Even if those machine-perception systems could (and currently they cannot) guarantee significantly better performance w.r.t. these two criteria than human vision within complex street scenes and at any level of environmental disturbance, like rain, fog, or blinding sun, their error rates would still remain orders of magnitude higher than the inherently strict safety target expected of automated vehicles. This implies that a significant gap remains to be bridged here, namely the gap between actual performance of technical perception and expected societal acceptance thresholds for unjustified manoeuvres.

This paper sets out to narrow this gap by answering the following three questions affirmatively:

1. Can we provide a mathematical or logical formalization of relevance of a percept such that we understand when a misperception remains redundant to a decision, i.e. either does not propagate into a — then unjustified — decision or does not harm the safety of the decision?
2. Can we demonstrate the positive safety impact of such redundancies in that we rigorously show that actual guard conditions are "safer than perception" in that their evaluation exposes considerably lower error rates than the perception, which is their input?

[3] Note that for extremely low error rates, realizing them technically and providing evidence for their satisfaction are completely different, both very hard, issues.

3. Can we provide a mechanism automatically rewriting a safety-critical guard condition into a more robust variant that retains the logical content of the original condition, yet offers resilient evaluation under uncertainties in that it provides false positive rate below the societal acceptance threshold while maximizing true positive rates?

In answering the first question, we will follow the lines of the precursor paper [6], which also addressed the second question by means of safe, i.e. pessimistic approximation of the quantitative risk. We will herein complement its analysis by a precise closed-form analysis of a special case, which sheds more light on the actual safety level to be expected and confirms that safety gains overarching orders of magnitude are indeed plausible. The last question remained an open issue for future research in [6] and we are pleased of now being able to expose an algorithm that can constructively construct optimized guard formulations.

This paper is organised as follows: Section 2 sketches a reference architecture used throughout the subsequent discussions while 3 provides a simple example showcasing the effect of why and how the evaluation of a complex guard condition safeguarding a safety-critical manoeuvre can be "safer than perception". The subsequent Sections 4 and 5 develop the mathematical framework facilitating quantitative analysis of this effect and prove its existence. Section 6 then sketches an automatic rewriting technique maximizing the resilience of a critical guard condition while keeping detection performance at a requested level. Sections 7 and 8, finally, refer to related work and provide a summary and pointers to future work.

2 Preliminaries

For solidly basing our analysis, we postulate a certain reference architecture. As discussed in [6], which this section is based on, the exemplary reference architecture uses labelled occupancy grids for collecting the output of machine-learning based algorithms that classify objects in the environment of the ego-vehicle. Class labels are assigned according to a (generally partially ordered, e.g., collecting cars, trucks, motor-cycles, etc. into a super-class of vehicles) ontology. The occupancy grid partitions the geometric vicinity of the ego car into finitely many grid elements. Its grid elements are filled with the corresponding class labels from the ontology whenever they have been perceived as being occupied by an object. Postulating this particular reference architecture is a matter of convenience, as it provides the subsequent analysis with a defined basis, but does by no means imply that the analysis would fail for other architectures, like those representing the world model by an object list — in fact, it carries over, as the models are mostly isomorphic (we discuss this in some more depth in the conclusion). Our analysis ought consequently, cum grano salis, carry over to the highly proprietary implementations of original equipment manufacturers and their suppliers.

Typical conditions enabling or blocking — and thereby meant to safeguard — critical manoeuvres then take the form of Boolean combinations of statements

concerning the occupancy of certain elements of the occupancy grid, with these elements together forming grid areas which correspond to subspaces of the surrounding traffic space. The atoms of such statements query occupancy of a particular grid element by certain object types named in the ontology, plus maybe additional unlocalized environmental conditions, like general visibility conditions. As an example take an evasive manoeuvre of a car across the curb to the footpath in order to make room for an emergency vehicle: Initiation of such a manoeuvre by the ego car would naturally be safeguarded by a Boolean condition requiring (1) presence of an emergency vehicle somewhere in the occupancy grid elements belonging to the traffic space reasonably close behind the ego car, (2) absence of vulnerable road users within some sufficiently large and connected group of occupancy grid elements belonging to the bike lane and footpath just ahead of the ego vehicle, (3) absence of any obstacles, including parked or stopped cars, on the line between the current ego position and the space for evasion identified via the previous condition, and finally (4) general (like absence of dense fog, presence of illumination) and geometric (like absence of occlusions) visibility conditions pertaining to the critical objects mentioned throughout the previous conditions.

While all the sub-conditions of the above guard condition intuitively make sense as being necessary conditions for safe execution of the safety-critical evasive manoeuvre, a safety risk due to misperception of some of the atomic statements occurring in the guard condition prevails, as no technical (nor a biological) perception system is perfect. In complex road scenes, we can neither expect to detect all potentially relevant objects nor are safe from misclassification of harmful objects as harmless and irrelevant. With absolute object detection rates often dropping below $\frac{2}{3}$ and classification accuracy easily falling below 90% in non-ideal visibility conditions [8], would our reliance into the evaluation of the guard condition drop into similar ranges due to the weakest link principle?

As the condition guarding the manoeuvre decision is a massive Boolean combination of atomic percepts, individual misperceptions might mask each other: not every single pedestrian needs to be detected, as safely crossing a pedestrian lane does not depend on the particular number of pedestrians being present. Likewise, slight misplacements of perceived objects is irrelevant, as e.g. a slight offset in locating a cyclist will not change drivability of the manoeuvre. When the guard condition reflects these properties, this can induce a considerably lower misevaluation rate for the overall condition than for its constituents, i.e. than for the atomic percepts dealing with detecting, locating, and classifying objects. Within this note, we are rendering this intuitive argument rigorous and formal, thus lifting reliability levels of combinatorial critical environmental perception well beyond the figures for atomic percepts achieved by state-of-the-art perception [8] paired with fusion techniques [14].

The main result of this paper is a methodology for, first, formally establishing and, second, constructively optimizing refined bounds on the risks of misperception for guard conditions concerning safety-critical manoeuvres based on the rates of misperception of atomic environmental artefacts. It complements formal synthesis-based approaches towards achieving safe controllers as well as

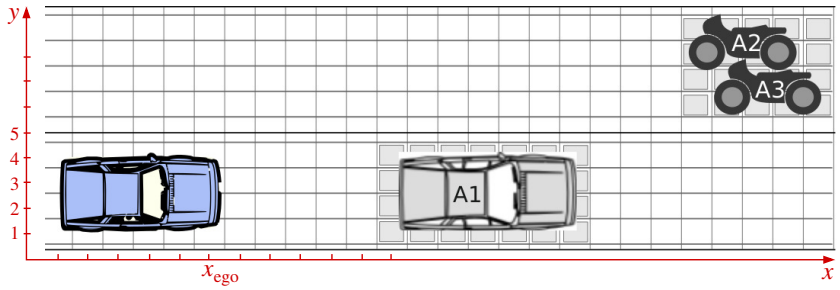

Fig. 1. A traffic situation requiring a safety-critical manoeuvre decision

engineered control architectures, as our reference architecture does not restrict the typically highly proprietary planning and manoeuvre control of HAV, and instead provides a generic interface between any such proprietary solutions and the perception chain. It does so by allowing to tune guard conditions pertaining to critical manoeuvre decisions based on the confidence level of the underlying individual percepts, allowing to optimise the trade-off between availability (induced from enabling the guarded action) and safety (bounding the error of erratically enabling an action due to misconception to societally accepted risk).

3 A simple example

The following example is taken from [6]. Consider the blue ego car being in the situation depicted in Fig. 1. When the ego car detects the obstacle A1, it will consider an evasive manoeuvre across the dividing line into the oncoming lane. This manoeuvre would, however, only be adopted if (1) it is necessary to avoid collision on the originally planned track and (2) it is considered safe w.r.t. the available information about the environmental state.

The manoeuvre would thus be (safe-)guarded by a guard condition g defined as $g \equiv \mathit{necessary} \wedge \mathit{safe}$, where

$$\mathit{necessary} \equiv \bigvee_{x=6}^{7} \bigvee_{y=1}^{5} \mathit{obstacle}@(x_{\text{ego}} + x, y)$$

denotes that some type of obstacle is detected as being present on the own lane, i.e. between 1 and 5 in y position, within relevant x distance (here, for the sake of being able to depict the example, shown as just 6 to 7 grid elements ahead in the x direction; the real figure would be considerably larger). Within the ontology, *obstacle* denotes an arbitrary type of road-blocking object and is defined as a disjunction about different basic object classification labels, like *trash container*, *tire*, *debris*, or *parked/slow car* (cf. Fig. 1). Note that this very definition already induces some fault-tolerance w.r.t. to misperception of atomic percepts: identifying the necessity for circumvention neither requires identifying the full back frontier of the obstacle, as the disjunction across y positions would

evaluate to true already if only a fraction of the frontier is detected, nor identifying correctly the exact type of obstacle, as *obstacle* is a disjunction across numerous obstacle types. Even identification of the x position of the obstacle would permit for tolerances if circumvention manoeuvres are dimensioned with a safety margin: locating A1 further left than it actually is would not cause risk (yet extend the circumvention), while locating it too far right stays collision-free if the misplacement remains within the safety margin. Note the combinatorially vast number of distorted perceptions of A1 that would thus still lead to the same truth value as the ground truth does. The likelihood of failing to detect the necessity of a circumvention consequently remains considerably lower than the unreliability of atomic percepts. This implies that the rate of false negative verdicts in the evaluation of *necessary* remains comparatively low. We will later see that, by just some rewriting to the way *necessary* is expressed, we will also be able to reduce the false-positive rate of the evaluation of *necessary* further to a frequency well below the false-positive rate of the atomic percepts.

We now turn to the safety condition, yet do in this note simplify its exposition slightly by omitting some additional conditions that are structurally perfectly similar to the ones shown. These omissions deal with occluded areas and are perfectly symmetric to the conditions on oncoming traffic explicated in the sequel. With these simplifications, the safety condition reads

$$safe \equiv \neg \bigvee_{y=6}^{10} \begin{pmatrix} \bigvee_{x=1}^{20} pedestrian@(x_{\text{ego}} + x, y) & \vee \\ \bigvee_{x=-4}^{40} car@(x_{\text{ego}} + x, y) & \vee \\ \bigvee_{x=-1}^{60} motorcycle@(x_{\text{ego}} + x, y) & \end{pmatrix} .$$

Its evaluation determines the presence of critical objects in the oncoming lane within the ego car's vicinity, constituting the safety condition that may block the circumvention manoeuvre when its execution may become hazardous. As this condition *safe* structurally resembles *necessary* with an outermost negation added, its fault-tolerance properties are in principle dual: where *necessary* is massively disjunctive and therefore tolerant against some or even numerous lacking or inaccurate percepts, *safe* as a negation over a disjunction essentially is conjunctive and consequently seems to require completeness of all percepts across the large set of atomic observations it mentions. This would imply that the very safety condition *safe* were not only as, but even orders of magnitude more fragile against misperception than any of the atomic percepts involved! Sufficiently reliable evaluation of the safety condition would consequently seem elusive, given that reliability of atomic percepts already falls considerably short of our actual safety targets. Fortunately, we will see that also here, a rewriting of the condition *safe* can help. Due to the duality, this rewriting now would have to reduce the false-negative rate in the evaluation of *safe*, thus being dual to the one applied to *necessary*.

The argument that we want to either minimize false-positive rate or false-negative rate of a disjunctive state condition, depending on the polarity of its occurrence, does however only apply when formula like *necessary* or *safe* stand in isolation. Boolean combinations of such disjunctive state conditions, where the

satisfying violations of subformulae occurring in opposite polarity (like *necessary* and *safe* in the example) can overlap, may require compromises. Furthermore, an — in principle desirable — minimization of false detections (i.e. of false-positive rate or false-negative rate, resp.) may not always be appropriate, as it unfortunately also tends to minimize the true detection (i,e, true-positive rate or true-negative rate, resp.), thus maximizing safety at the price of minimizing availability of a — presumably useful — action.

Luckily, this problem can be alleviated by careful analysis (and modification whenever beneficial) of the Boolean problem structure of the conditions *safe*, *necessary*, and the guard condition g within the general mathematical framework provided in the next section. Subsequent sections will then exploit the framework to rigorously quantify the reliability gain that the Boolean structure of the guard condition provides over the atomic percepts, as well as show how to constructively rewrite the guard condition by phrasing its true-positive and false-positive rates as a constrained optimization problem.

4 Boolean guard formulae as classifiers under uncertainty

Let Φ be a formula that guards a safety-critical manoeuvre in the sense that the driving function will only adopt the manoeuvre when it has positive evidence of the validity of Φ in the current situation, implying that the manoeuvre would be avoided (and a safer substitute adopted) whenever Φ is violated *or* evaluation of Φ remains inconclusive. The formula $g \equiv \mathit{necessary} \wedge \mathit{safe}$ from the previous section is an example of such a guard condition Φ.

Generally, such formulae Φ comprise massive Boolean combinations of conditions on individual cells of the occupancy grid, where both the particular cells referenced and the individual conditions vary situationally. E.g. Φ may safeguard a transit through a shared traffic space by ensuring that there are no vulnerable road users in the street, where the geometric position of the referenced areas of the occupancy grid depend on the own car's position as well as the particular geometry of the shared space and the planned trajectory through that space. In this particular setting, $\Phi = s_1 \wedge s_2 \wedge \cdots \wedge s_n$ is a conjunction of statements $s_i = \neg o_i$, where $o_i = \mathit{pedestrian}_i \vee \mathit{cyclist}_i$ is a disjunction of atomic percepts $a_{i,j}$ expressing the property "there is a vulnerable road user of type j at the cell c_i of the occupancy grid". The truth value of each atom $a_{i,j}$ therein directly depends on a classifier output, which is a classifier for the object classes "pedestrian" or "cyclist", resp., in this particular example.

The central problem we are facing obviously is that the percepts $a_{i,j}$ are not reliable wrt. ground truth $A_{i,j}$, as there is a non-trivial risk for lacking perception of an object or for misclassification of a perceived object. We consequently have to distinguish between the ground truth $A_{i.j}$ underlying such an atomic percept and the possibly distorted percept $a_{i,j}$. E.g., it may be true that there is a pedestrian at cell i (i.e. $A_{i,pedestrian}$ holds), but we misperceive her as a waste bin (i.e. $\neg a_{i,pedestrian}$ as well as $a_{i,waste\ bin}$ hold).

The first — rather trivial, yet crucial — observation is that there is no direct need to align $a_{i,j}$ with $A_{i,j}$, i.e. to minimize the misperception rates of individual classifiers, but only a need to sufficiently reduce the misevaluation rate of the compound condition Φ. Let us denote by $\mathcal{GT}(\Phi)$ the formula $\Phi[a_{1,1},\ldots,a_{n,m}/A_{1,1},\ldots,A_{n,m}]$ where all percepts $a_{i,j}$ have been replaced by their (factually unknown) ground truth $A_{i,j}$. Then, in any situation σ assigning truth values to all ground-truth atoms $A_{i,j}$ as well as to all percepts $a_{i,j}$v, the truth value $\mathcal{GT}(\Phi)(\sigma)$ represents the (desired, yet unknown in practice) actual value of the guard condition Φ, while $\Phi(\sigma)$ is the result of evaluating Φ on the potentially distorted percepts. We follow the tradition to write $\sigma \models \psi$ if $\psi(\sigma)$ evaluates to true and $\sigma \not\models \psi$ if $\psi(\sigma)$ evaluates to false, for any formula ϕ. Thus, we call σ a *false positive* for Φ iff $\sigma \models \mathcal{GT}(\Phi)$ while $\sigma \not\models \Phi$. We call σ a *true positive* for Φ iff $\sigma \models \mathcal{GT}(\Phi)$ and $\sigma \models \Phi$.

As false positives induce risk (e.g. due to suggesting overtaking when it actually is unsafe) while true positives are constitutional for system performance (e.g. enabling overtaking whenever safely possible), our obligation then is to ensure that the false-positive rate remains below a defined threshold Θ pertaining to societally acceptable risk while maximizing true-positive rate. As Φ is a complex Boolean combination of atomic statements $A_{i,j}$, this is not identical to the problem of maximizing the true-positive rates and minimizing the false-positive rates for the classifiers generating the percepts $a_{i,j}$. Simple as the above observation is, it already has profound consequences for the pragmatics of developing safety-critical autonomous systems, as it implies that the currently prevailing isolated optimization of computer vision for high detection rates may not be the most effective approach towards overall system safety and performance. We will not elaborate on this issue within this note, yet leave it to future exploration.

5 Equivalence of logically distinct guard conditions with respect to ground truth

The second — and more profound — observation is that the above problem statement provides us with liberty in phrasing the condition Φ. Especially if the real world satisfies some relevant invariants ι — which it inevitably does — then we can rephrase Φ into Φ' such that in all worlds satisfying the invariant ι, the two formulae Φ and Φ' evaluate identically, i.e.

$$\iota \models \mathcal{GT}(\Phi) \iff \mathcal{GT}(\Phi') \ , \tag{1}$$

where, as usual, $\phi \models \psi$ denotes that any model of ϕ (i.e. every world satisfying ϕ) also is a model of the formula ψ (i.e ψ holds too in those worlds). In the sequel, we call formulae satisfying Equation (1) world equivalent (W-equivalent for short):

Definition 1. *Given a propositional invariant ι, two propositional formulae Φ and Φ' over atoms $a_{i,j}$ referring to percepts are called* W-equivalent modulo ι *iff property* (1) *holds.*

The interesting property is that W-equivalent formulae, despite always agreeing over the ground truth, may well feature substantially different positive rates, both for true and for false positives.

For an extremely simplified example, consider that flat obstacles are only relevant if they have a size of at least 6 grid elements — all smaller ones we circumvent while staying in lane or drive over them, letting them pass between our wheels. The invariant ι for relevant flat obstacles thus is that they cover six grid elements at least, such that in the ground truth, either none or at least six grid elements feature a relevant flat obstacle. Then the formula

$$\phi \;\equiv\; \bigvee_{x=1}^{3} \bigvee_{y=1}^{5} \mathit{flatobstacle}@(x_{\text{ego}} + x, y) \tag{2}$$

expressing presence of a flat obstacle directly in front of the ego car is, for any $k \in \{1, \ldots, 6\}$, W-equivalent wrt. ι to

$$\phi_k \equiv \sum_{x=1}^{3} \sum_{y=1}^{5} \mathit{flatobstacle}@(x_{\text{ego}} + x, y) \geq k \;\;, \tag{3}$$

where we adopt the standard convention to identify false with 0 and true with 1 when taking the sum. Nevertheless, their positive rates vary obviously, as satisfying formula ϕ_k gets harder for larger k, with the easiest instance given by $k = 1$ being logically equivalent to ϕ. Being harder to satisfy means that the positive rates get smaller. This applies both to the false positive rates — which is beneficial — and — detrimentally — to the true positive rates.

To provide an analytically solvable example of this effect, consider the stronger invariant ι' that any relevant flat obstacle covers *exactly* six grid elements. Then all instance ϕ_k of (3) are still W-equivalent modulo ι' and their false-positive rates and true-positive rates can, by straightforward reduction to binomial distributions, be analytically described by the formulae in Table 1, where *tp*, *fp*, *tn* and *fn* denote the true-positive rates, false-positive rates, true-negative rates, and false-negative rates, resp., of the atomic classifiers. For simplicity of the example, these are assumed to be uniform and stochastically independent

obstacles present	$fp(\phi_k)$	$tp(\phi_k)$
0	$\sum_{i=k}^{15} \binom{15}{i} fp^i\, tn^{15-i}$	—
1	—	$\sum_{i=k}^{15} \sum_{j=0}^{i} \binom{6}{j} tp^{\,j}\, fn^{6-j} \binom{15-6}{i-j} fp^{i-j}\, tn^{15-6-(i-j)}$
2	—	$\sum_{i=k}^{15} \sum_{j=0}^{i} \binom{12}{j} tp^{\,j}\, fn^{12-j} \binom{15-12}{i-j} fp^{i-j}\, tn^{15-12-(i-j)}$

Table 1. False-positive rates and true-positive rates for the guard formula ϕ_k from (3) in dependence of $k \in \{1, \ldots, 6\}$ for the cases of none to two obstacles of size 6 grid cells being present in the critical region.

	$k=1$	$k=2$	$k=3$	$k=4$	$k=5$	$k=6$	
0 o.	**0.965**	0.833	0.602	0.352	0.164	**0.061**	$fp(\phi_k)$
1 o.	$1-2\cdot10^{-6}$	$1-6\cdot10^{-4}$	0.999	0.992	0.954	0.834	$tp(\phi_k)$
2 o.	$1-6\cdot10^{-11}$	$1-5\cdot10^{-9}$	$1-1\cdot10^{-7}$	$1-3\cdot10^{-6}$	$1-4\cdot10^{-5}$	$1-4\cdot10^{-4}$	$tp(\phi_k)$

Table 2. False- and true-positive rates for the guard formula ϕ_k from (3) in dependence of $k \in \{1,\ldots,6\}$ when the detection rates for the atomic classifiers are $tp = 0.85$, $fp = 0.2$, $tn = 0.8$, and $fn = 0.15$.

across the occupancy grid. In practice, stochastic dependencies are obviously to be expected between adjacent grid elements. Extending the analysis to such spatial and furthermore to spatio-temporal dependencies is subject to future work. Qualitatively, the expected effects are, however, similar, as spatial dependencies affect both the ground truth (where geometric connectedness of objects permits additional inference) and the perception (where mispercepts may be correlated). These true-positive rates etc. are in practice determined empirically by the usual statistical testing methods and thus known (up to a confidence) for the operational design domain.

When these classification rates are $tp = 0.85$, $fp = 0.2$, $tn = 0.8$, $fn = 0.15$ then this results in the false and true positive rates for ϕ_k reported in Table 2. As expected, the positive rates decrease rapidly when k grows. For the safety-critical false-positive rate, we observe a $\approx$ 16-fold reduction when going from the original formula ϕ or its logical equivalent ϕ_1 to ϕ_6: In detail, $\phi = \phi_1$ has $\approx$ 5-fold fp of the atomic sensor, while ϕ_6 has just $\approx \frac{3}{10} fp$ of the atomic sensor (cf. entries in Table 2 marked in bold-print). The true-positive rate (i.e., performance) degrades also significantly to 83.4% in worst case (cf. second line of Table 2), but still remains at the level of an individual classifier, which features $tp = 85\%$.

Note that these reductions in positive rates come with no change in the ground-truth semantics, as W-equivalent formulae evaluate identically wrt. the ground truth according to Definition 1 and property (1). It also is worth noting that the above reductions do not degrade when going from this toy example to realistically fine and large grids — to the contrary, due to the binomials involved, these results obviously become exponentially better over a finer grid!

6 Synthesizing optimal representations of guard conditions

Having seen that the replacement of a safety-critical guard condition by a W-equivalent variant can significantly alter true-positive and false-positive rates, an obvious follow-up question is whether we can effectively rewrite such a guard condition into a more appropriate form. The exact problem we would like to solve is the following:

Problem 1. Given a guard condition Φ and a societally accepted maximum risk $\theta \in]0,1[$, automatically synthesize a formula Φ' that is W-equivalent to Φ and satisfies the two requirements

Safety: the false positive rate of Φ' remains below θ, i.e.

$$fp(\Phi') \leq \theta \ , \tag{4}$$

Performance: the true positive rate of Φ' be maximal among the safe W-equivalent rewritings of Φ, i.e.

$$tp(\Phi') = \max\{tp(\Phi'') \mid \phi'' \text{ W-equivalent to } \Phi, tp(\Phi'') \leq \theta\} \ . \tag{5}$$

To solve Problem 1, we first observe that it is a non-standard instance of don't-care optimization. For the sake of rendering the solution representable in a conference paper, we adopt an even smaller example than in Section 5. Assume three observational atoms A_1, A_2, A_3 and a guard condition $g(A_1, A_2, A_3)$ as given in Table 3 as a truth table. The ground-truth invariant is that never a single atom A_i can be true. Entries in the truth table satisfying $\sum_{i_1}^{3} A_i = 1$ are consequently "don't-cares", as they cannot arise in reality.

The don't-care entries in the truth table allow for setting them arbitrarily to 0 or 1. It would, however, in general be a bad idea to set all of them to 0, as this nicely minimizes $fp(g')$, but also minimizes $tp(g')$ unfortunately, thus optimizing safety at the price of minimizing performance. Vice versa, setting all don't-cares to 1 would maximize both $fp(g')$ and $tp(g')$, thus optimizing performance while minimizing safety. In general, we need clever compromises, to be achieved by a differentiated setting of the individual don't-cares to 0 or 1.

To achieve such, we observe that both the true-positive rate and the false-positive rate of g' can be represented as affine terms over the don't-care assignments as follows: Denote by p_{abc}^{xyz} the likelihood of perceiving the ground truth

A_1	A_2	A_3	$g(A_1, A_2, A_3)$
0	0	0	0
0	0	1	$*$
0	1	0	$*$
0	1	1	1
1	0	0	$*$
1	0	1	1
1	1	0	1
1	1	1	0

Table 3. Truth table of a guard condition. Entries satisfying $\sum_{i_1}^{3} A_i = 1$ have an arbitrary "don't care" truth value (marked with $*$), as they do not arise in ground truth due to a ground-truth invariant that no single atom A_i can be true in isolation.

$(A_1, A_2, A_3) = (a, b, c)$ as $(a_1, a_2, a_3) = (x, y, z)$. Then

$$\begin{aligned} fp(g') = {} & p_{000}^{001}\, x_1 + p_{000}^{010}\, x_2 + p_{000}^{011}\ + p_{000}^{100}\, x_3 + p_{000}^{101}\ + p_{000}^{110}\ + \\ & p_{111}^{001}\, x_1\ + p_{111}^{010}\, x_2\ + p_{111}^{011}\ \ + p_{111}^{100}\, x_3\ + p_{111}^{101}\ \ + p_{111}^{110}\ , \end{aligned} \tag{6}$$

$$\begin{aligned} tp(g') = {} & p_{011}^{001}\, x_1 + p_{011}^{010}\, x_2 + p_{011}^{011}\ + p_{011}^{100}\, x_3 + p_{011}^{101}\ + p_{011}^{110}\ + \\ & p_{101}^{001}\, x_1 + p_{101}^{010}\, x_2 + p_{101}^{011}\ + p_{101}^{100}\, x_3 + p_{101}^{101}\ + p_{101}^{110}\ + \\ & p_{110}^{001}\, x_1 + p_{110}^{010}\, x_2 + p_{110}^{011}\ + p_{110}^{100}\, x_3 + p_{110}^{101}\ + p_{110}^{110} \end{aligned} \tag{7}$$

holds, where x_1, x_2, and x_3 are the truth values assigned to the three don't-cares $(A_1, A_2, A_3) = (0, 0, 1)$, $(A_1, A_2, A_3) = (0, 1, 0)$, and $(A_1, A_2, A_3) = (1, 0, 0)$, respectively.

An assignment $x_1 \in \{0, 1\}$, $x_2 \in \{0, 1\}$, and $x_3 \in \{0, 1\}$ to the don't-cares satisfying the above two requirements Safety and Performance can now mechanically be found by solving the following 0-1 integer-linear program:

$$\begin{aligned} &\text{Maximize } tp(g') \\ &\text{subject to } fp(g') \leq \theta \text{ and } x_1 \in \{0, 1\},\ x_2 \in \{0, 1\},\ x_3 \in \{0, 1\}, \end{aligned}$$

where $fp(g')$ and $tp(g')$ are the affine expressions from the right-hand sides of Equations (6) and (7). Note that the objective function $tp(g')$ and the domain constraint $fp(g') \leq \theta$ do directly encode the two requirements (5) and (4) from Problem 1. The above 0-1 integer-linear program can routinely be solved by any integer-linear programming (ILP) solver. The values reported for x_1 to x_3 in the optimal solution do then directly fill the don't-care entries in Table 3 if a solution exists. If no solution exists, then it is impossible to satisfy the societal acceptance threshold Θ on false positives by an W-equivalent rewriting.

The above construction, however, does not scale. Being based on enumerating the entries of the truth table, its size is strictly exponential in the number of atoms A_i involved in the guard condition. The construction consequently becomes impractical when considerably more than 20 atoms are involved, which still constitutes a clearly minuscule occupancy grid. But luckily there is a lot of symmetry in formulae (6) and (7): one would for example expect that $p_{000}^{001} = p_{000}^{010} = p_{000}^{100}$, as all of them involve flipping one bit from 0 to 1 in $(A_1, A_2, A_3) = (0, 0, 0)$. Likewise, $p_{111}^{001} = p_{111}^{010} = p_{111}^{100}$ and $p_{100}^{011} = p_{010}^{101} = p_{001}^{110}$ and $p_{101}^{001} = p_{101}^{100} = p_{110}^{010} = p_{110}^{100} = p_{011}^{010} = p_{011}^{001}$ etc. Grouping together equal factors and exploiting the symmetry in the solution space stemming from the fact that for subexpressions of the form $ax_1 + ax_2 + ax_3$ only the sum $x_1 + x_2 + x_3$ is decisive while it is irrelevant which of x_1, x_2, and x_3 is set to 1, we can replace above ILP by the ILP

$$\begin{aligned} &\text{Maximize } tp(g'') \\ &\text{subject to } fp(g'') \leq \theta \text{ and } x \in \{0, \ldots, 3\}, \end{aligned}$$

where $fp(g')$ and $tp(g')$ are defined by the affine integer expressions

$$\begin{aligned} fp(g'') &= 3p_{000}^{001}\,x + 3p_{111}^{001}\,x + 3p_{000}^{011} + 3p_{111}^{110} \\ &= \tbinom{3}{0}\tbinom{3}{1}\tbinom{0}{0}p_{000}^{001}\,x + \tbinom{3}{3}\tbinom{0}{0}\tbinom{3}{2}p_{111}^{001}\,x + \tbinom{3}{0}\tbinom{3}{2}\tbinom{0}{0}p_{000}^{011} + \tbinom{3}{3}\tbinom{0}{0}\tbinom{3}{1}p_{111}^{110} \quad , \end{aligned} \tag{8}$$

$$\begin{aligned} tp(g'') &= 6p_{011}^{001}\,x + 3p_{011}^{011} + 3p_{011}^{100}\,x + 6p_{011}^{101} \\ &= \tbinom{3}{2}\tbinom{1}{0}\tbinom{2}{1}p_{011}^{001}\,x + \tbinom{3}{2}\tbinom{1}{0}\tbinom{2}{0}p_{011}^{011} + \tbinom{3}{2}\tbinom{1}{1}\tbinom{2}{2}p_{011}^{100}\,x + \tbinom{3}{2}\tbinom{1}{1}\tbinom{2}{1}p_{011}^{101} \quad . \end{aligned} \tag{9}$$

Note that the binomial factors in front of the probabilities p_{abc}^{xyz} directly reflect the numbers of bits set in the ground truth (a, b, c) in the first binomial factor, the number of bits flipped from 0 to 1 among the 0 bits in the ground truth (a, b, c) to obtain the perception (x, y, z) in the second binomial factor, and the number of bits flipped from 1 to 0 in the third binomial factor. Therefore, the rather compact — and therefore as computationally inexpensive to formulate and solve — formulae (8) and (9) can be constructed directly by combinatorial reasoning without enumerating truth table entries. This process would in practice start from a Don't-Care-BDD representation of the guard condition g rather than a truth table like that from Table 3 used here for illustration. Implementation of this procedure is underway such that experimental results cannot yet be reported.

7 Related Work

Partially or fully autonomous cyber-physical systems, like highly automated vehicles, operate in an uncertain dynamic environment, which they have to perceive and understand in order to draw often safety-critical decisions. Such systems consequently tend to be learning-enabled — not necessarily end-to-end, but at least in central components relevant to perception and situation assessment. Their perception of the environment, i.e. the detection of properties about the dynamic environment, is enabled through inherently noisy sensors and subsequent machine-learnt classifiers. Especially in environmental perception based on computer vision, the uncertainties and the misperception rates induced by such machine-learnt classification algorithms remain substantial when visibility conditions are non-ideal [10]. The resultant misperception rates are orders of magnitude higher than the safety targets of, e.g., HAV [15].

Characterisation and control of perceptive uncertainty can be achieved at multiple stages of an architecture for automated cyber-physical systems, starting from the individual sensor level over fusion of multiple sensors and sensing modalities to control of uncertainty propagation through the decision and planning layers of a robotics architecture. The measures taken at these different stages complement each other, with our contribution being located at the last of the aforementioned three stages.

Representations of uncertainty impacting the inferences underlying planning decisions have been investigated within the paradigm of probabilistic robotics [23], among others, particularly as applied to vehicle localization in urban environments [13, 19, 14], with localization being a special and historically more well-understood

instance of the general problem of safe-guarding critical manoeuvre decisions. In these and related works such as [1, 17], the environment uncertainty is usually represented as probabilistic beliefs. Our constructive approach in this article, as well as its analytical counterpart in [6], complements such approaches by analysing and optimising the uncertainty propagation through the complex Boolean guard conditions usually employed for enabling and safeguarding safety-critical manoeuvres, be it as enabling preconditions of such manoeuvres or as side-conditions in safety shields [2] for AI components [12] or as fallback conditions in SIMPLEX-type fault-tolerant control architectures [22, 16].

It ought to be noted that approaches confining and controlling error propagation in the decision layer complement optimizations on the previous layers of sensing and sensor fusion, directly benefiting from, but also enhancing the impact of, enhancements at these layers. Various approaches to combining multiple classifiers can be found in the literature, e.g., see [7, 21] for an overview. The goal of such a combination is often to compensate for individual shortcomings in the performance by a better performance of the multitude of classifiers [21]. While in the pre-classification level the combination happens at the sensor or raw data level, the focus of this note is on fusion of classifiers at the post-classification level, as on-the-fly combinations of the decision of multiple atomic classifiers are considered. A major challenge for fusion on the decision level arises from the fact that the least genuine information about the object of observation is available at this level [11].

Our approach inherits the traditional setting of balancing between true positive rates, i.e. performance, and false positive rates, i.e. quantitative risk induced by misclassification, of adjustable classifiers, which often is pursued by analysing the empirical ROC (receiver operating characteristics) curve [3, 20]. Such analysis facilitates the optimisation of individual classifiers as well as their combination to obtain a better performance by a multitude of classifiers [21].

In a sense, our approach can be interpreted as a combination of classifiers too, namely one per grid element, albeit with given Boolean combination logic, which distinguishes it from the aforementioned fusion approaches, where the combination is to be designed based on the mutual performance figures of the multiple classifiers. The only degree of freedom we have thus is to modify the combinatory logic such that it maintains the same logical function on all ground-truth instances (see Def. 1, yet still enhances resilience to misclassification. The key to modification of the combinatory logic is the identification of don't cares in its truth tables due to ground-truth invariants. The exploitation of such ground-truth invariants itself is not original; it has already been investigated in visual classification tasks to enhance the accuracy of scene segmentation tasks, e.g. in [5], where the natural vertical layering (e.g. that vehicles stand on the tarmac and not vice versa) of visual scenes is used as invariant.

It is also interesting to note that the paradigm of occupancy grids and hence the approach suggested carries well beyond automotive manoeuvre planning and similar real-time path planning problems requiring distance between objects. Similar approaches have been used to cover safety problems of contact robotics,

like the safe use of robot manipulators in collaborative scenarios (cobots) as well as physical human-robot interaction. To generate collision-free trajectories, models of human motion ought to be integrated for better estimation, and path planning needs to be optimised for execution speed and safety. Both the environment and the human operator are represented via occupancy grids in [24] and exploited for planning that adapts to different human operators or their hand positions. If direct contact between the manipulator and the human is relevant to the task, the contact forces must also be limited. [18] have integrated biomechanical injury information into the robot controller for this purpose, with different force thresholds applying to different body parts, thus requiring occupancy maps for representing the positions of human body parts. Symmetrically, occupancy maps of the full robot arm, rather than just the tool centre point, are employed in [9] to limit joint torques in whole arm manipulations with their multiple contact points, inducing spatially distributed safety constraints.

8 Conclusion

Reliable guarantees for the safety of autonomous systems are a prerequisite for their societal acceptance. The quest for such guarantees cannot easily be served, at least not at the appropriate quantitative safety levels for safety-critical systems like autonomous vehicles at usual speed of traffic, due to the relatively high misperception rates of technical perception chains. When mapping an autonomous vehicle's vicinity, their error rates tend to — currently as well as for the foreseeable future — be orders of magnitude higher than the pertinent safety targets for autonomous operation. Any compositional analysis inducing a weakest-link principle, i.e. suggesting that the overall system's reliability would be bounded by perception performance, is thus bound to generate grossly insufficient evidence of system safety.

We consequently are in need of analytical methods or even constructive means to ensure that our systems actually are "safer than perception". More precisely, this requires a rigorous assessment of the likelihood that a safety-critical manoeuvre is erroneously adopted, and this assessment has to provide much tighter bounds for such erratic manoeuvre adoption than for any misperception. In a companion report [6], we have addressed the analytical problem of rigorously proving a quantitative reliability figure for the evaluation of a complex Boolean guard condition that is safeguarding a safety-critical manoeuvre, in the sense that its evaluation to true is a necessary prerequisite for adoption of the manoeuvre. We have been able to show that for complex guard conditions, the rates of critical misevaluations can be proven to be significantly lower than misperception rates concerning atomic percepts.

In this article, we drove this analysis further and gave it a constructive tweak in that we asked for solving an optimization problem that deals with finding that rewriting of a given guard condition that is most resilient to misperceptions while retaining the semantics of the original guard condition. Concretely, we asked for finding a formula rewriting that

1. is a equivalent to the original guard over all ground-truth instances,
2. reaches the societally given safety target in that its rate of false positives, i.e. of suggesting the critical manoeuvre when it is undue, remains below the societal acceptance threshold, and
3. optimizes performance in that it yields maximal true positive rate, i.e. actually enables the critical manoeuvre when it is due.

We showed that this problem has a constructive solution by reducing it to integer-linear programming, thereby automatically synthesizing a formula rewriting satisfying the aforementioned three requirements.

Practical implementation of the algorithm and of modifications leveraging symmetries as well as symbolic reasoning for enhancing scalability are underway as a student project. Future work will deal with a spatio-temporal rather than just propositional analysis, refining our analysis by topological and geometrical properties induced by the grid structure and temporal correlations induced by the dynamics. Obviously, a slight misplacement of a detected object both is more likely to happen and more unlikely to change a guard's perceived truth value than a large displacement. Similar arguments apply in the temporal domain, where true positives and true negatives tend to have a higher temporal persistence than false positives or false negatives, respectively. While these properties have extensively been studied for effectively filtering atomic mispercepts, their impact on the evaluation of complex spatio-temporal conditions serving as guards remains a subject of ongoing research [4].

Another interesting question concerns transfer of the results to other reference architectures than occupancy grids, especially to object list representations of the environment. These do classify objects, locate them at an environmental position, and add a bounding box characterizing their geometric extent. The latter is often inexact, giving rise to quality measures like the (relative) area of intersection over union for the perception. Related perception problems are detecting an object part (e.g., a car backlight) instead of the full object (the car) due to visibility problems like occlusion, then attributing a subcomponent label and an accordingly smaller bounding box. Both phenomena fit our analysis in principle, as again a reasonable guard condition will talk about a non-trivial Boolean combination of more than percept (now in terms of class, relative position and size of the bounding box), and rephrasing it to increase resilience is as relevant. Object lists add convexity properties and shape constraints due to the bounding boxes, but can otherwise be seen as mostly isomorphic to 2.5-dimensional occupancy grids, locating the bounding boxes in the perspective plane rather than the street plane. It will again be an issue of future research to fill in the details.

Acknowledgements: The research reported herein has been supported by the State of Lower Saxony within the Zukunftslabor Mobilität as well as by Deutsche Forschungsgemeinschaft under grant no. DFG FR 2715/5-1 "Konfliktresolution und kausale Inferenz mittels integrierter sozio-technischer Modellbildung". It furthermore benefitted from technical discussions with Krzysztof R. Apt, Werner Damm, Willem Hagemann, Hardi Hungar, and Paul Kröger, as well as from the comments of the anonymous reviewers. Their support is gratefully acknowledged.

References

1. Baig, Q., Perrollaz, M., Laugier, C.: A robust motion detection technique for dynamic environment monitoring: A framework for grid-based monitoring of the dynamic environment. IEEE Robot. Automat. Mag. **21**(1), 40–48 (2014)
2. Bloem, R., Könighofer, B., Könighofer, R., Wang, C.: Shield synthesis: - runtime enforcement for reactive systems. In: Baier, C., Tinelli, C. (eds.) Tools and Algorithms for the Construction and Analysis of Systems - 21st International Conference, TACAS 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015. Proceedings. Lecture Notes in Computer Science, vol. 9035, pp. 533–548. Springer (2015). https://doi.org/10.1007/978-3-662-46681-0_51, https://doi.org/10.1007/978-3-662-46681-0_51
3. Fawcett, T.: An introduction to ROC analysis. Pattern recognition letters **27**(8), 861–874 (2006)
4. Finkbeiner, B., Fränzle, M., Kohn, F., Kröger, P.: A truly robust signal temporal logic: Monitoring safety properties of interacting cyber-physical systems under uncertain observation. Algorithms **15**(4), 126 (2022). https://doi.org/10.3390/A15040126, https://doi.org/10.3390/a15040126
5. Fouopi, P.P.: Holistische Modellierung und Interpretation von Szenen und Situationen basierend auf symbolischen, probabilistischen und subsymbolischen Modellen. Ph.D. thesis, University of Oldenburg, Germany (2019), http://oops.uni-oldenburg.de/4601
6. Fränzle, M., Hagemann, W., Damm, W., Rakow, A., Swaminathan, M.: Safer than perception: Assuring confidence in safety-critical decisions of automated vehicles. In: Haxthausen, A.E., Huang, W., Roggenbach, M. (eds.) Applicable Formal Methods for Safe Industrial Products - Essays Dedicated to Jan Peleska on the Occasion of His 65th Birthday. Lecture Notes in Computer Science, vol. 14165, pp. 180–201. Springer (2023). https://doi.org/10.1007/978-3-031-40132-9_12, https://doi.org/10.1007/978-3-031-40132-9_12
7. Galar, M., Fernandez, A., Barrenechea, E., Bustince, H., Herrera, F.: A review on ensembles for the class imbalance problem: bagging-, boosting-, and hybrid-based approaches. IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews) **42**(4), 463–484 (2011)
8. Geirhos, R., Janssen, D.H.J., Schütt, H.H., Rauber, J., Bethge, M., Wichmann, F.A.: Comparing deep neural networks against humans: object recognition when the signal gets weaker. CoRR **abs/1706.06969** (2017), http://arxiv.org/abs/1706.06969
9. Gliesche, P., Kowalski, C., Pfingsthorn, M., Hein, A.: Geometry-based two-contact inverse kinematic solution for whole arm manipulation. In: IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2021, Prague, Czech Republic, September 27 - Oct. 1, 2021. pp. 8269–8274. IEEE (2021). https://doi.org/10.1109/IROS51168.2021.9636664, https://doi.org/10.1109/IROS51168.2021.9636664
10. Huber, L.S., Geirhos, R., Wichmann, F.A.: The developmental trajectory of object recognition robustness: children are like small adults but unlike big deep neural networks. CoRR **abs/2205.10144** (2022). https://doi.org/10.48550/ARXIV.2205.10144, https://doi.org/10.48550/arXiv.2205.10144
11. Khreich, W., Granger, E., Miri, A., Sabourin, R.: Iterative Boolean combination of classifiers in the ROC space: An application to anomaly

detection with HMMs. Pattern Recognition **43**(8), 2732–2752 (2010). https://doi.org/10.1016/j.patcog.2010.03.006
12. Könighofer, B., Rudolf, J., Palmisano, A., Tappler, M., Bloem, R.: Online shielding for reinforcement learning. Innov. Syst. Softw. Eng. **19**(4), 379–394 (2023). https://doi.org/10.1007/S11334-022-00480-4, https://doi.org/10.1007/s11334-022-00480-4
13. Levinson, J., Montemerlo, M., Thrun, S.: Map-based precision vehicle localization in urban environments. In: Proceedings of Robotics: Science and Systems. Atlanta, GA, USA (June 2007). https://doi.org/10.15607/RSS.2007.III.016
14. Levinson, J., Thrun, S.: Robust vehicle localization in urban environments using probabilistic maps. In: IEEE International Conference on Robotics and Automation. pp. 4372–4378 (2010)
15. Maurer, M., Gerdes, J.C., Lenz, B., Winner, H.: Autonomous Driving: Technical, Legal and Social Aspects. Springer Publishing Company, Incorporated, 1st edn. (2016)
16. Mitsch, S., Platzer, A.: Modelplex: verified runtime validation of verified cyber-physical system models. Formal Methods Syst. Des. **49**(1-2), 33–74 (2016). https://doi.org/10.1007/S10703-016-0241-Z, https://doi.org/10.1007/s10703-016-0241-z
17. Moras, J., Cherfaoui, V., Bonnifait, P.: Moving Objects Detection by Conflict Analysis in Evidential Grids. In: IEEE Intelligent Vehicles Symposium (IV 2011). pp. 1120–1125 (2011)
18. Palleschi, A., Hamad, M., Abdolshah, S., Garabini, M., Haddadin, S., Pallottino, L.: Fast and safe trajectory planning: Solving the cobot performance/safety trade-off in human-robot shared environments. IEEE Robotics Autom. Lett. **6**(3), 5445–5452 (2021). https://doi.org/10.1109/LRA.2021.3076968, https://doi.org/10.1109/LRA.2021.3076968
19. Petrovskaya, A., Thrun, S.: Model based vehicle detection and tracking for autonomous urban driving. Auton. Robots **26**(2-3), 123–139 (2009)
20. Powers, D.: Evaluation: From precision, recall and f-measure to ROC, informedness, markedness & correlation. Journal of Machine Learning Technologies **2**(1), 37–63 (2011)
21. Sagi, O., Rokach, L.: Ensemble learning: A survey. WIREs Data Mining and Knowledge Discovery **8**(4), e1249 (2018). https://doi.org/10.1002/widm.1249
22. Seto, D., Krogh, B., Sha, L., Chutinan, A.: The simplex architecture for safe online control system upgrades. In: Proceedings of the 1998 American Control Conference. ACC (IEEE Cat. No.98CH36207). vol. 6, pp. 3504–3508 vol.6 (1998). https://doi.org/10.1109/ACC.1998.703255
23. Thrun, S., Burgard, W., Fox, D.: Probabilistic Robotics (Intelligent Robotics and Autonomous Agents). The MIT Press (2005)
24. Zanchettin, A.M., Messeri, C., Cristantielli, D., Rocco, P.: Trajectory optimisation in collaborative robotics based on simulations and genetic algorithms. Int. J. Intell. Robotics Appl. **6**(4), 707–723 (2022). https://doi.org/10.1007/S41315-022-00240-4, https://doi.org/10.1007/s41315-022-00240-4