

A Doxastic Characterisation of Autonomous Decisive Systems*

Astrid Rakow

Carl von Ossietzky University, 26129 Oldenburg, Germany

a.rakow@uo1.de

A highly autonomous system (HAS) has to assess the situation it is in and derive beliefs, based on which, it decides what to do next. The beliefs are not solely based on the observations the HAS has made so far, but also on general insights about the world, in which the HAS operates. These insights have either been built in the HAS during design or are provided by trusted sources during its mission. Although its beliefs may be imprecise and might bear flaws, the HAS will have to extrapolate the possible futures in order to evaluate the consequences of its actions and then take its decisions autonomously. In this paper, we formalize an *autonomous decisive system* as a system that always chooses actions that it currently believes are the best. We show that it can be checked whether an autonomous decisive system can be built given an application domain, the dynamically changing knowledge base and a list of LTL mission goals. We moreover can synthesize a belief formation for an autonomous decisive system. For the formal characterization, we use a doxastic framework for safety-critical HASs where the belief formation supports the HAS's extrapolation.

1 Introduction

Nowadays highly autonomous systems (HASs) are already able to work within complex environments as e.g. the autonomous Tesla cars and the NASA's helicopter "Ingenuity" demonstrate. On the other hand, these systems have to satisfy high safety standards. This is especially so for human cyberphysical systems, where humans and autonomous systems symbiotically cooperate [23]. Accidents like [17, 20] show that there is still a long way to go before the high safety standards can be reliably implemented.

A HAS faces the great challenge of making good decisions despite an overwhelmingly complex and dynamic environment. In order to decide what to do next, a HAS builds a belief about its current situation, usually by combining incoming sensor data with prior data. As this data is often incomplete and data from different sources may be contradictory, the HAS can never be sure of its current state and thus deems several variants of the situation possible. The HAS extrapolates the possible futures that might develop from the current situation. Therefore, it simulates its action alternatives on a world model that describes the impact of its actions as well as how the environment might evolve. Hence, a HAS has to take decisions based on beliefs that concern its past, present and future and that can only approximate the real world.

In this paper, we focus on safety critical systems and formalize the notion of autonomous-decisive system. We call a system autonomous-decisive, if it can choose its actions based on the contents of its beliefs and then achieves its goals within the intended application domain. We formalize the notion within a game-theoretic doxastic framework, where we use the possible worlds semantics [11] to represent the beliefs of a system. A so-called *belief formation function* specifies the belief for a given history of perceptions and the current knowledge base. The latter holds knowledge and hard beliefs. It captures

*This work is partly funded by the German Research Council (DFG) as part of the project "Integrated Socio-technical Models for Conflict Resolution and Causal Reasoning", grant no. FR 2715/5-1, DA 206/12-1

constraints that an engineer chooses to enforce on the system's beliefs (e.g. insights about the application domain) as well as statements that the HAS gets from trusted sources during its mission. As we focus on safety-critical systems, we use a two-player zero-sum reactive game. That is, we consider the environment as adversarial and the system has to operate successfully irrespective of the environment's choices. By considering the belief formation itself as a special kind of observation-based strategy, we can synthesize a witness belief formation that enables a HAS to autonomously achieve its goals.

Conceptually, the presented formalization will be useful in the early design, where simple and abstract models are considered and decisions about the resources of a system like its sensors and capabilities are made. The formalism is applicable to any system that has to take autonomous decisions. We imagine that the application domain and the mission goals are formally described. When during the design the possible beliefs of the system have been fixed, our framework can be used to examine whether the system can successfully take autonomous decisions within the targeted application domain. When the setting does not allow to decide autonomously, the design of its possible beliefs, the system's sensing capabilities or actions might have to be adjusted. In certain cases, even the model of the application domain has to be revised [8].

Outline Next we introduce the ingredients of our formal framework. In Sect. 3 we formalize autonomous-decisive systems as systems that are able to take decisions based on their internal assessment of the situation. We prove that it is decidable whether such a system exists for a given application domain, available knowledge/hard beliefs, possible beliefs and mission goals. In Sect. 4 we present related and future work and conclude in Sect. 5.

2 Ingredients of our Doxastic Framework

In this section, we will introduce the ingredients our framework alongside a running example. As a toy example, we consider the set-up as sketched in Fig. 1(a) where two cars are on separate lanes heading towards each other. The left car is called *ego* and is our HAS. It has to avoid collisions and it has to take the left turn. The right car, *other*, is uncontrolled (from *ego*'s perspective).


2.1 A World

We assume that the HAS's intended application domain is captured via a Kripke structure, modelling static and dynamic objects and their interactions. This so called *design-time world* W_D captures the information about the application domain as well as test criteria that the HAS has to master as part of the development and certification process.

Definition 1 (*world*) Formally, a world W is a labelled Kripke structure $W = (S, Ed, L, I)$, where S is the set of states, $Ed \subseteq S \times S$ are edges between states, $L = L_S \cup L_{Ed}$ where $L_{Ed} = Ed \rightarrow 2^{Act}$ labels edges with a subset of the finite set of actions Act and $L_S : S \rightarrow 2^{AP}$ labels states with valuations of a finite set of Boolean propositions AP , and $I \subseteq S$ is the set of initial states. $Act = Act_{ego} \times Act_{env}$ is a set of tuples defining the simultaneous actions of our autonomous agent *ego* and its environment *env*, which may include other agents. We assume that the transition relation is defined for all states and actions, i.e., $\forall s \in S, \forall a \in Act, \exists e \in Ed : a \in L_{Ed}(e)$.

In order to express that an action is not enabled at s , W can transition into a dedicated state s_{undef} that is accordingly labelled.¹ A sequence of states $\pi = s_0 s_1 \dots s_n \in S^* \cup S^\omega$ is a *path* in W iff $\forall i, 0 \leq i < |\pi| :$

¹In this paper any strategy has hence to avoid s_{undef} .

Figure 1: A simple world W_D .

$(s_i, s_{i+1}) \in Ed$. $\pi(i)$ is the state s_i , $\pi_{\leq m}$ denotes the prefix $s_0 \dots s_m$ and $last(\pi)$ is the last state of a finite path π . π is initial iff $s_0 \in I$.

Example 1 (A world) *In our running example, the actions of ego are f , “moving one step forward if possible”, t , “turn and move one step forward”. Other is either a slow car and moves one tile forward, if possible, or it is a fast car. In that case, it moves two positions forward, if it is at its initial position, else it moves one tile forward, if possible. Other’s actions are f and F , “move two positions forward”. The actions of ego and other are annotated by pale blue and dark blue arrows in Fig. 1(a).*

The propositions $AP_{pos} = AP_{x_e} \cup AP_{y_e} \cup AP_{x_o}$ with $AP_{x_e} = \{x_e = i | 1 \leq i \leq 4\}$, $AP_{y_e} = \{y_e = i | 1 \leq i \leq 3\}$, $AP_{x_o} = \{x_o = i | 1 \leq i \leq 4\}$ encode the positions of the two cars, where $x_e = i$ and $y_e = i$ represent the horizontal and vertical position of ego, and $x_o = i$ represents the horizontal position of other. Its vertical position is always two. Other’s car type is encoded via the propositions slow (s) and fast (f). We assume that ego cannot observe other’s car type directly, but it has sensors perceiving other’s colour, which is either red or blue. The proposition blue (red, fast, slow) is true, iff other is a blue (red, fast, slow) car. The propositions $blue_p$ and red_p encode what ego perceives as other’s colour. They are used to modeling ego’s imperfect colour recognition, while the propositions blue and red encode the true colour of other. We assume that ego’s colour perception works correctly, when ego and other are less than two tiles apart, otherwise the sensor switches colours ($blue_p = \neg blue$, $red_p = \neg red$). Let $AP_{cartype}$ be the set $\{fast, slow, blue, red, blue_p, red_p\}$. The propositions in our example are hence $AP = AP_{pos} \cup AP_{cartype} \cup \{undef\}$, where $undef$ labels the sink state.

Figure 1(b) shows the Kripke structure of our design time world. States are labelled with the propositions that hold in the respective state. The label $abcdef \in \mathbb{N}^3 \times \{f, s\} \times \{b, r\} \times \{b_p, r_p\}$ encodes that $x_e = a, y_e = b, x_o = c$ are true and other’s car type is d , its colour is e whereas the perceived colour is f . The valuations of all other propositions are false. Likewise, the label $undef$ encodes that the only valid proposition is $undef$. Edges are labelled with sets of actions. We omit the sets’ brackets. For example, the label ff denotes the set $\{ff\}$, that contains the one action ff , where ego and other simultaneously move one step forward, if possible. Actions that are not enabled at a state lead to the sink state s_{undef} . We omit the sink state and sink transitions in the sections that follow. ■

2.2 Goal List

Our HAS, ego , has to achieve a prioritized list of goals during a manoeuvre. A goal φ is a linear-time temporal logic (LTL) formula [3]. We denote the temporal operators “globally” by \square , “eventually” by \diamond , “next” by X and “until” by U . We interpret the LTL formulae over (infinite) traces, which are infinite

sequences $t = t_0 t_1 \dots \in (2^{AP})^\omega$ of valuations of AP . Satisfaction of an LTL formula φ by a trace t is denoted as $t \models \varphi$. A goal list $\psi = (\Phi, \text{prio})$ consists of a set Φ of LTL formulae and a priority function $\text{prio} : \Phi \rightarrow \{1, \dots, |\Phi|\}$ where $\varphi \in \Phi$ is more important than $\varphi' \in \Phi$ iff $\text{prio}(\varphi) < \text{prio}(\varphi')$. We say that a trace t satisfies ψ with priority up to n if n is the highest priority up to which t satisfies all goals of ψ , i.e., $t \models \varphi$ for all $\varphi \in \Phi$ with $\text{prio}(\varphi) \leq n$. A set of traces T satisfies ψ up to n , if all $t \in T$ satisfy ψ up to a priority m with $n \leq m$ and n is the greatest such priority. For convenience we extend goal lists to include the trivial goal true, i.e., $\text{true} \in \Phi$ and $\text{prio}(\text{true}) = 0$.

Example 2 (Prioritized Goals) We formalize collision freedom as $\varphi_c = \Box(x_e = 2 \wedge y_e = 2 \Rightarrow x_o \neq 2)$, and $\varphi_t = \Diamond(y_e = 3)$ expresses that ego eventually does the turn. In order to rule out that disabled transitions are taken, we use the goal $\varphi_u = \Box \neg \text{undef}$. The priorities are given by $\text{prio}(\varphi_u) = 1, \text{prio}(\varphi_c) = 2, \text{prio}(\varphi_t) = 3$.

Let us now take a closer look at what ego should do in order to accomplish its goals. If other is slow, then ego should not take the turn, but instead it should drive straight on, in order to avoid the collision. If other is fast, then ego can take the turn and accomplish all its goals. ■

2.3 Observations

Ego, being highly autonomous, will take decisions based on the beliefs that it has constructed about the world in which it operates. *Ego* derives its beliefs from the observations made so far and the knowledge/hard beliefs it has about the world.

A world is usually only partially perceivable by *ego* via observations. Observations \mathcal{O} are propositions of W_D whose valuations *ego* can assess and that represent e.g. sensor readings or received messages from other agents. Observations shed light on W_D , but they do not have to be truthful, as illustrated in Exp. 1, where initially $\text{blue}_p, \text{red}_p$ were switched and hence *ego* does not perceive the correct colour all the time. Nevertheless, a HAS has to reason about the current state of the world based on its past observations. Let P be a set of propositions, $P \subseteq AP$. We call h a (P -)history of a state s , if there is an initial path π of W leading to s and $h = h_0 h_1 \dots h_n$ is the sequence of state labels along π , $h_i = L(\pi(i)) \cap P, \forall i, 0 \leq i \leq n$. We denote the set of P -histories as \mathbb{H}_P . We say h is observable iff $P \subseteq \mathcal{O}$.

Example 3 (Observable History) In our example *ego* cannot observe other's position due to a broken distance sensor, but it can observe its own position and the colour of other, so that $\mathcal{O} := AP_{x_e} \cup AP_{y_e} \cup \{\text{undef}, \text{blue}_p, \text{red}_p\}$. Given the world of Fig. 1(b), $h = 114sbr_p, 213sbb_p, 312sbb_p, 411sbb_p$ is the history along the path $s1, s2, s3, s5$ wrt AP , whereas $11r_p, 21b_p, 31b_p, 41b_p$ is the observable history wrt \mathcal{O} . ■

2.4 Beliefs

A belief describes what *ego* currently thinks. For instance, *ego* may think that it saw an approaching vehicle and that this vehicle is a slow car (or a fast car). Depending on the type of car, *ego* will imagine different possible future evolutions of the current situation.

Definition 2 (belief, reality) A belief B is the set of alternative realities that *ego* currently deems possible, $B = \{r_0, \dots, r_n\}$. A reality is a pair $r = (W, S_c)$ of a (possible) world W and a set of believed current state $S_c \subseteq S$, where any current state is reachable from one initial state but no initial state is reachable from a current state.

A reality describes possible pasts, current states and futures: pasts are captured by the set of paths between initial states I and current states S_c , the possible futures are paths from the current states.

Example 4 (*Alternative Realities and Beliefs*) To illustrate the notion of belief as introduced above, let us consider two alternative realities of ego in Fig. 2(a)+(b). The believed past is marked by framing state labels.




Figure 2: Two alternative realities of ego

The alternative reality of (a) describes that other is slow and that ego itself is at the initial state b_0 of that world. In (b) other is fast and ego is at c_1 , the “second” state of that world. The believed history is 114fr,212fr. Since a belief is a set of alternative realities, singletons of either (a) or (b) form a belief. Also, the set of (a) and (b) forms a belief, where ego thinks both alternatives are possible.

Note, *ego* may believe in possible worlds that are substantially different from W_D , since it usually captures the real world by simplified concepts and rules, that may reflect the inner workings of the application domain only coarsely but sufficiently. Since a HAS has only finite resources, we assume that the set of possible beliefs \mathbb{B} and the set of possible worlds $\mathbb{W} := \bigcup_{(s_c, W) \in \mathbb{B}} W$ are finite. The choice of \mathbb{B} constitutes a design decision within the development process of the HAS.

Example 5 (*Possible Beliefs \mathbb{B}*) Our ego has been designed to represent a certain set of scenarios, for which it can evaluate what to do by extrapolating the future. Figures 3(a)-(d) sketch a set of the alternative realities via their initial set-up. The other car may be fast or slow, the road may be up to 6 tiles long, the intersection may be at $x = 2$ or $x = 3$, and the start position of other varies from $x = 4$ to $x = 6$. Note, that W_D of Fig. 1 is described by Fig. 3 (a). Let ego’s possible beliefs \mathbb{B} be the beliefs that




Figure 3: Sketch of *ego*’s possible beliefs. If *other* is fast, it uses the dashed arrow at first and then the solid arrows. If it is slow, it uses the solid arrow.

canonically evolve from these initial scenarios.

2.5 Knowledge/Hard Beliefs

Knowledge/Hard Beliefs are statements with which an engineer equips a system during the system design or that are provided by certain trusted sources (e.g. a traffic control system) during the HAS's mission. We think of statements like (φ_a) “In settled areas the speed limit is 50 km/h”, (φ_b) “I will be on the highway for the next 20 mins.” or rules like (φ_c) “If A promises to give way, I can rely on it.”. We specify these statements via an LTL variant, which we call Belief-LTL (BLTL). A BLTL formula φ is an LTL formula that is preceded by either the operator K or B. $K\varphi$ reads as “ego believes that it knows φ ”. A belief B satisfies $K\varphi$, $B \models K\varphi$, if in all alternative realities of B , φ holds on all traces of all initial paths. Note, $K\varphi$ may hold, although φ may not hold on W_D . $B\varphi$ is defined as $\neg K\neg\varphi$ and reads as “ego believes that φ is possible”. Likewise, we define K^c or B^c on traces from current states. A finite set of BLTL formulae constitutes a *knowledge base* $\mathcal{K} \in \mathbb{K}$. A belief B satisfies \mathcal{K} , $B \models \mathcal{K}$, if $B \models \varphi$ holds for all $\varphi \in \mathcal{K}$. We assume that *ego*'s knowledge base varies over time. Therefore, we extend the labelling function L by $L_{\mathbb{K}} : S \rightarrow \mathbb{K}$ specifying the available knowledge base at a state. Given a history h , \mathcal{K}_h denotes the knowledge base of $last(h)$.

Example 6 (A Knowledge Base) *Let ego have the following knowledge base \mathcal{K} at all states.*

1. $\varphi_z = K\Box((\neg x_o = 5 \wedge \neg x_o = 6) \vee undef)$ (other is at most at $x = 4$),
2. $\varphi_t = K\Box(\neg x_e = 2 \Rightarrow (y_e = 1 \vee undef))$ (a turn may only be possible at $x = 2$),
3. $\varphi_i = Kx_e = 1$ (ego starts at $x = 1$) and
4. $\varphi_{ct} = K\Box \bigwedge_{t \in \{slow, fast\}} (t \Rightarrow X(t \vee undef)) \wedge (\neg t \Rightarrow X(\neg t \vee undef))$ (the initial car type does not change). ■

2.6 Belief Formation

Ego updates its beliefs e.g. when it gets new information from its sensors, a clock tick or a message from another agent. The belief formation function \mathcal{B} captures formally how *ego* builds its belief.

Definition 3 (belief formation function, knowledge-consistent) *The belief formation function \mathcal{B} , $\mathcal{B} : \mathbb{H}_{\mathcal{O}} \times \mathbb{K} \rightarrow \mathbb{B}$, specifies the belief B that ego derives after perceiving history h of observations $\mathcal{O} \subseteq \mathbb{O}$ and while trusting in \mathcal{K}_h .*

We call a belief formation function \mathcal{B} knowledge-consistent, if its beliefs satisfy the current knowledge base, i.e., $\mathcal{B}(h, \mathcal{K}_h) \models \mathcal{K}_h$ for all h .

Note, that $L_{\mathbb{K}}$ is a mean to define a relation between *ego*'s beliefs and the ground truth, when the belief formation is knowledge-consistent. In that case we can for instance specify that the formed beliefs have to be consistent with the last two observations, or with messages of the traffic control system.

For a history of observations $h = h_0h_1 \dots h_n$ and a history of knowledge bases $\mathcal{K} = \mathcal{K}_0\mathcal{K}_1 \dots \mathcal{K}_n$, we denote by $\tilde{\mathcal{B}}(h, \mathcal{K})$ the resulting history of beliefs, $\mathcal{B}(h_0, \mathcal{K}_0)\mathcal{B}(h_0h_1, \mathcal{K}_1) \dots \mathcal{B}(h_0h_1 \dots h_n, \mathcal{K}_n)$. We usually write $\mathcal{B}(h)$ instead of $\mathcal{B}(h, \mathcal{K}_h)$. Likewise, we write $\tilde{\mathcal{B}}(h)$ for $\tilde{\mathcal{B}}(h, \mathcal{K})$.

Example 7 (Knowledge-Consistent Belief Formation) *Let us now see how ego's belief formation could be like. Let ego have the knowledge base \mathcal{K} of Exp. 6 and let ego also be convinced that a red car is fast, while a blue car is slow, i.e., $\varphi_b = \Box(\text{fast} \Leftrightarrow \text{red}_p) \wedge \Box(\text{slow} \Leftrightarrow \text{blue}_p)$, so all states of W_D are labelled with $\mathcal{K}' := \mathcal{K} \cup \varphi_b$ by $L_{\mathbb{K}}$. In order to satisfy \mathcal{K} , only beliefs sketched in (a) of Fig. 3 remain possible. The initial belief of a knowledge-consistent belief formation has to consist of the alternative realities depicted in Fig. 4(a)+(c). When ego in the real world scenario of Figure 4(b) incorrectly perceives that other is blue, it will form B_b , the belief of Figure 4(a). B_b expresses that ego thinks it knows that other is blue and hence a slow car, although it is red in the real world. It moreover captures ego's expectation of*




Figure 4: The two alternative realities of *ego* at its initial state (a)+(c); the real world scenarios (c)+(d) with observable values in bold face type.

how the future will develop. Similarly, Fig. 4(c) shows the belief B_r , which *ego* forms, when it is in the real world scenario of Fig. 4(d) and incorrectly perceives that other is red.

So, since *ego*'s perception is mistaken, *ego* is initially convinced that there is a fast car, when there is a slow car and vice versa. *Ego* hence thinks that it should do the turn, when it should not. But *ego* does not have to take the decision whether to turn at that time. It first moves one tile forward, while other simultaneously moves either one or two tiles forward. In both cases, *ego* then perceives other's colour correctly.

Let us say, in retrospective *ego* is aware of the faulty perception and corrects its belief on other's car type and colour. It updates the belief to Fig. 5(a) when it is in the scenario Fig. 4(b), and to Fig. 5(b), when in Fig. 4(d). Figure 6 sketches the belief formation so far. The observed history $11r_p$ is mapped to belief $B_{0,1}$ and $11b_p \mapsto B_{0,2}$, $11r_p, 21b_p \mapsto B_{1,1}$ and $11b_p, 21r_p \mapsto B_{1,2}$. The sketched belief formation is knowledge-consistent. Although *ego* changed its belief about other's type, φ_{ct} holds for each formed belief.




Figure 5: Possible worlds

Since *ego*'s beliefs match now the reality, *ego* is able to assess the best strategy. *Ego* can develop its strategy based on its sequence of beliefs, arguing along the lines "Initially I thought the car is red and fast and that it is a good idea to do the turn. Now I think the car is blue and slow and then the turn is not good idea, since I would collide with other. Since I believe, that my current belief matches the reality, I choose to drive straight on." ■

In the above example, the belief is made up of singletons of alternative realities only. The next example illustrates the use of several alternative realities.

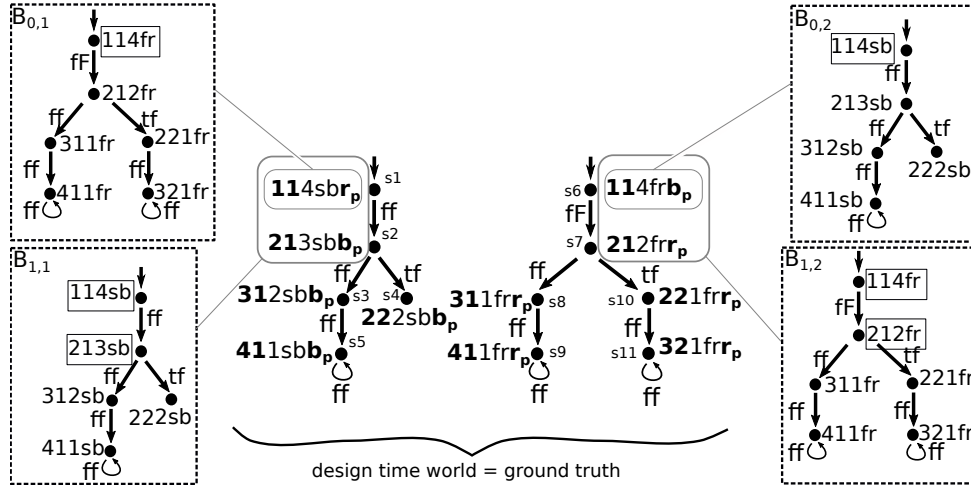


Figure 6: Sketch of a belief formation function

Example 8 (Alternative Realities) Let us assume that *ego* is unsure of its own initial position, thinking that it may initially be at $x = 1$ or $x = 2$, as sketched in Fig. 3(a)+(c). So *ego* deems two realities possible, when at state s_1 . In both realities other is red, but in one reality, r_1 , *ego* is at $x = 1$, in the other, r_2 , at $x = 2$. Similarly, *ego* deems two realities possible when at s_6 . ■




Figure 7: Sketch of a belief formation function when *ego* is unsure of its initial position

2.7 Doxastic Model

To summarize, a doxastic model D is given by a tuple $(W_D, \psi, \mathbb{K}, \mathcal{O}, \mathbb{B}, \mathcal{B})$ of the design-time world W_D , the prioritized list of goals ψ , the knowledge bases \mathbb{K} , a set of observations \mathcal{O} , the set of possible beliefs \mathbb{B} of a HAS and a belief formation \mathcal{B} . The belief formation describes how a system links observations made within the world W_D to the HASs inner representation of the world, i.e. the beliefs \mathbb{B} that a HAS can possibly build. The world W_D is considered as ground truth during the design. Later design steps have to take care of the gap between W_D and the real world.

3 Autonomous Decisions

In this section, we formalize what it means that *ego* decides autonomously. *Ego* decides autonomously, if its decision is the result of *ego*'s simulative assessment of its actions on content of its beliefs and the decision is sensible wrt *ego*'s observations.

Our formalization hence distinguishes between autonomous systems –systems that take decisions based on the belief content– and automatic systems –systems that play out rule-based decisions that have been taken at design time. In order to formalize *autonomous-decisive* we contrast (i) strategies that have access to ground-truth with (ii) strategies that can only observe the formed beliefs, and (iii) strategies that run as simulation within the beliefs. The different strategy notions are introduced step by step in the sequel and an overview of the notions is given in Table 1.

3.1 Truth-Observing and Doxastic Strategies

In our framework, we use truth-observing strategies as reference of what would be achievable, if *ego* could directly access the ground truth W_D via a set of propositions P . To this end, we say *ego* implements a *P-truth-observing strategy* $s_t : (2^P)^+ \rightarrow Act_{ego}$, if *ego* chooses its actions based on the history of values of $P \subseteq AP_d$ as observed in W_D . Implementing s_t means that *ego* decides what to do at a state $\pi(i)$, reached via the path π in W_D , based on $L(\pi_{\leq i})|_P$, i.e. based on the true past values of P observed thus far. A truth-observing strategy s_t together with a sequence of environment actions $e \in Act_{env}^\omega$ determines a set of traces, $T(e, s_t)$. Formally, $T(e, s_t) = \{t_0 t_1 \dots \in (2^{AP_d})^\omega \mid \exists \text{ path } \pi \text{ from } I_D, \forall i \geq 0 : t_i = L_D(\pi(i)) \wedge a_i := s_t(L_D(\pi_{\leq i})|_P) \wedge (a_i, e(i)) \in L_D(\pi_i, \pi_{i+1})\}$.

Since *ego* has no direct access to the ground truth, we also formalize that *ego* decides based on the history of beliefs. At a state $\pi(i)$ in W_D *ego* takes a decision based on the history of its beliefs $b_0 \dots b_i$ that *ego* has built along $\pi_{\leq i}$. *Ego* implements a doxastic strategy $s_d : \mathbb{B}^+ \rightarrow Act_{ego}$ on W_D , if *ego* chooses its actions based on the history of its beliefs. A strategy s_d together with a sequence of environment actions $e \in Act_{env}^\omega$ determines a set of traces in W_D , just like for truth-observing strategies. The set of traces is $T(e, s_d) = \{t_0 t_1 \dots \in (2^{AP_d})^\omega \mid \exists \text{ path } \pi \text{ from } I_D, \forall i \geq 0 : t_i = L_D(\pi(i)) \wedge a_i := s_d(\mathcal{B}(\pi_{\leq i})) \wedge (a_i, e(i)) \in L_D(\pi_i, \pi_{i+1})\}$. Note that a doxastic strategy also depends on observations of the doxastic model D due to the belief formation \mathcal{B} (cf. Def. 3).

Since truth-observing and doxastic strategies both determine traces for a given sequence of environment actions, we can compare them straight forwardly: *Ego* achieves ψ up to n on W_D by strategy s , if for all $e \in Act_{env}^\omega$ all $t \in T(e, s)$ satisfy ψ up to n . A strategy s ψ -dominates s' on W , $s' \leq_W s$, iff s' achieves ψ up to m and s up to $k \geq m$.

Example 9 (Truth-Observing and Doxastic Strategies) *As an example of a dominant P-truth-observing strategy, let us consider the goal list of Exp. 2, the world model in Fig. 1(b), $P := \{x_e, \text{slow}, \text{fast}\}$ and the strategy s_t that chooses to drive straight on, if other is fast, and it chooses to turn, if other is slow (it maps $1s \mapsto f$, $1s, 2s \mapsto f$, $1s, 2s, 3s \mapsto f$, $1s, 2s, 3s, 4s \mapsto f$, and $1f \mapsto f$, $1f, 2f \mapsto t, \dots$). Strategy s_t achieves ψ up to 2 and is a dominant (P-truth-observing) strategy, since in all cases collision-freedom is guaranteed and in case the car is slow, no other strategy is able to realize both, the turn and collision freedom.*

Let us now define a doxastic strategy s_d for a belief formation, as sketched in Fig. 6. Let s_d be a doxastic strategy with $B_{0,1} \mapsto f$, $B_{0,1} B_{1,1} \mapsto f$, \dots and $B_{0,2} \mapsto f$, $B_{0,2} B_{1,2} \mapsto t$, \dots . Just like s_t , s_d chooses to turn when other is fast ($B_{0,2} B_{1,2} \mapsto t$), and it chooses to drive straight on, if other is slow ($B_{0,1} B_{1,1} \mapsto f$). As there is “no better” strategy, s_d is dominant. ■

Next, we want to capture that *ego* chooses its actions based on the *content* of its beliefs. In order to motivate our formalization, let us consider the following example.

Example 10 (*Decisions Not Based on the Belief Content*) We modify our running example slightly: Let us assume the colour perception is severely broken and permanently switches red to blue and vice versa. In Fig. 8 the changed world model is given, and a belief formation is sketched, that derives a belief that trusts in the colour perception, i.e., if the sensors say the other car is red, then ego believes the other car is red.




Figure 8: Sketch of a belief formation function for severely broken colour perception

Let s_d' be a doxastic strategy with $B_{0,1} \mapsto f$, $B_{0,1}, B_{1,2} \mapsto f$, ... and $B_{0,2} \mapsto f$, $B_{0,2}, B_{1,1} \mapsto t$, ... Just as s_t and s_d , the strategy s_d' realizes a turn on W_D , if other is fast ($B_{0,2}, B_{1,1} \mapsto t$), and ego drives straight on, if other is slow ($B_{0,1}, B_{1,2} \mapsto f$). So s_d' is also dominant, but s_d' makes no sense from ego's perspective. In case other is fast, ego believes that other is slow, since it trusts its sensors. In this case, ego extrapolates that doing the turn would cause a collision, but according to s_d' , ego chooses to take the turn. Vice versa, s_d' chooses to drive straight on, when ego believes other is fast and although ego extrapolates that taking the turn is alright. ■

3.2 Autonomous-decisive System

We now formalize what it means that *ego* decides based on the content of its belief. *Ego* always chooses an action that belongs to a strategy that is dominant in *ego*'s current belief B . To capture this formally, we introduce the notion of *possible-worlds strategy*. A *possible-worlds strategy* is a function $s_p : (2^{AP_B})^+ \rightarrow Act_{ego}$ and it is simulated on the alternative realities of *ego*'s current belief B . This results in "believed traces". We define this set of traces in an alternative reality $r = (W, S_c) \in B$ for a (believed) sequence of environment actions $e \in Act_{env}^\omega(W)$ as $T(e, s_p, r) = \{t_0 t_1 \dots \in (2^{AP})^\omega \mid \exists \text{ path } \pi \text{ in } W \text{ from } I : \forall i \geq 0 : t_i = L(\pi(i)) \wedge a_i := s_p(L(\pi_{\leq i})) \wedge (a_i, e(i)) \in L(\pi_i, \pi_{i+1})\}$. We generalize the notion of ψ -dominance to possible-worlds strategies. A possible-worlds strategy s_p ψ -dominates a possible-worlds strategy s_p' in B , if s_p ψ -dominates s_p' in all realities $r \in B$.

Example 11 (*Possible-Worlds Strategy*) Consider the possible-worlds strategy that chooses to turn, if other is fast, and to drive straight on, if other is slow, i.e., consider s_p with $114fr \mapsto f$, $114fr, 212fr \mapsto t$, $114fr, 212fr, 221fr \mapsto f$, ..., $114sb \mapsto f$, $114sb, 213sbb_p \mapsto f$, ... In the belief $B_{0,1}$ of Fig. 8 ego thinks that it is at the initial state $114fr$. Ego follows s_p by choosing $s_p(114fr) = f$. In contrast, the belief $B_{1,2}$ captures

Strategy types:

- *truth-observing strategy* $s_t : (2^P)^+ \rightarrow Act_{ego}$
observes the ground truth world W_D via $P \subseteq AP_{W_D}$ and takes decisions based on their history; serves as comparative reference of what is achievable given P could be observed directly
- *doxastic strategy* $s_d : \mathbb{B}^+ \rightarrow Act_{ego}$
observes the beliefs to take decisions and takes decisions based on their history; represents the decision making of autonomous and automatic systems
- *possible-worlds strategy* $s_p : (2^{AP_B})^+ \rightarrow Act_{ego}$
captures how a HAS “simulates” its strategies within the alternative realities; decisions are taken based on the believed history within the respective alternative reality

A strategy s ψ -dominates s' , $s' \leq s$, iff s' achieves ψ up to m' but s achieves ψ up to m' with $m' \leq m$.

Table 1: Strategy types & dominance in a nutshell

that ego thinks to have already made one move and is now at state 212fr. Ego follows s_p by choosing $s_p(114fr, 212fr) = t$. ■

A peculiarity of possible-worlds strategies is, that they can be indecisive for a belief: A belief may contain alternative realities encoding even contradictory information. E.g. let us assume that *ego* is currently not sure where its goal position is. In the alternative reality r_1 its goal position is to its left while in r_2 the goal position is right of *ego*. *Ego*'s strategy is to move to the left, when in r_1 and to move to the right when in r_2 . Since *ego* deems both realities possible, it cannot decide whether to turn right or left.

Definition 4 (*current-state decisive possible-worlds strategy, current-state choice*) A possible-worlds strategy s_p is current-state decisive in a belief B , if s_p specifies one and the same action for all the believed current states of all alternative realities of this belief, i.e.

$\exists a_B \in Act_{ego}$ such that $\forall r = (W, S_c) \in B : \forall$ paths π in $W : \pi(0) \in I \wedge last(\pi) \in S_c \Rightarrow s_p(\pi) = a_B$. Let $Act(s_p, B)$ denote the action a_B that s_p chooses at the current states $\bigcup_{(., S_c) \in B} S_c$ of B . We call $Act(s_p, B)$ be the current-state choice of s_p .

We can now define our notion of autonomous-decisive system. We consider a system to be autonomous-decisive, if (i) it takes its decisions based on the content of its beliefs and (ii) the resulting behaviour is sensible. We capture (i) by requiring that there is a doxastic strategy s_d , i.e., a strategy that can only observe the beliefs. Moreover, s_d takes only actions that seem the best choice according to the simulation within its beliefs. With other words, s_d only takes actions that are the current-state choice of a dominant, current-state decisive possible-worlds strategy. We capture (ii) by requiring that s_d does not perform worse than a strategy that decides based on the made observations, i.e. s_d is not ψ -dominated by a \mathcal{O} -truth-observing strategy. We hence require that the belief formation captures the gist of observations.

Definition 5 (*doxastic system, autonomous-decisive system*) Let a doxastic model $D = (W_D, \psi, \mathbb{K}, \mathcal{O}, \mathbb{B}, \mathcal{B})$ be given.

A doxastic strategy s_d is autonomous-decisively achieving ψ up to n , if s_d achieves ψ up to n and it always follows dominant, current-state decisive possible-worlds strategies, i.e., for all observed histories h , it holds that there is a dominant possible-worlds strategy s_p with $s_d(\tilde{\mathcal{B}}(h)) = Act(s_p, \mathcal{B}(h))$ that is current-state decisive in $\mathcal{B}(h)$.

A (doxastic) system is a pair $\mathcal{S} = (D, s_d)$ of a doxastic model D and a doxastic strategy s_d . \mathcal{S} is an autonomous-decisive system, if there is an n such that s_d autonomously achieves ψ up to n and s_d is not ψ -dominated by a \mathcal{O} -truth-observing strategy.

Example 12 (Autonomous, Non-Autonomous, Automatic)

Let us consider an example of an autonomous ego in the setting of Fig. 6, where the sensor only initially switches colours, and consider the possible-worlds strategy s_p of Exp. 11 (turn, if other is fast, and drive on, if other is slow). When ego initially evaluates its situation in s_1 of W_D , it believes that the situation is as described by $B_{0,1}$, i.e. other is a fast, red car. Ego can decide to follow s_p in $B_{0,1}$, as it seems a good choice – s_p is dominant and current-state decisive in $B_{0,1}$. According to its extrapolation, it would move one step forward, and then it would successfully take the turn. After actually moving forward, ego evaluates the situation in s_2 . In s_2 ego believes in $B_{1,1}$ reflecting that ego now truthfully perceives other's colour as blue (cf. Exp. 7). Again, s_p is a dominant current-state decisive possible-worlds strategy and determines f as the next move. Along this line, it is easy to see that ego can implement a doxastic strategy s_d that chooses the action that s_p determines for the respective $\mathcal{B}(h)$.

For an example of a non-autonomous ego, we modify our running example slightly. Let us assume that ego is unsure of its own initial position, thinking that it may initially be at $x = 1$ or $x = 2$, as sketched in Fig. 7. Let a belief formation be given that evolves the initial beliefs $B_{0,3}$ and $B_{0,4}$ analogously to Exp. 7, that is, ego perceives the correct colour after moving one step forward. The possible-worlds strategy s_p is still a dominant strategy but not current-state decisive, since for example in $B_{0,3}$ ego would do the turn at s_1 due to reality r_2 , and it would also drive straight on due to r_1 . Hence, there is no dominant possible-worlds strategy in $B_{0,3}$ that is able to determine one action. Ego cannot decide autonomously.

Note that we nevertheless can specify a dominant doxastic strategy s_d for this case, but its actions are not chosen based on the belief content: Ego chooses to turn after one step when its initial belief was $B_{0,4}$ ($B_{0,4}, B_{1,4} \mapsto t$), otherwise it drives straight on. This strategy is dominant and could be used to build an automatic system, where ego just plays out s_d . Such a strategy might be useful when an engineer knows that ego will start from $x = 1$ but did not equip ego with this information. ■

A system that is not autonomous-decisive cannot determine by itself which action is currently appropriate. A goal for the design of a HAS is to ensure that a system is autonomous-decisive.

Theorem 1 (Autonomous Decisiveness) Let $D = (W_D, \psi, \mathcal{K}, \mathcal{O}, \mathbb{B})$, a doxastic model without belief formation function, be given.

We can decide whether there is a knowledge-consistent belief formation \mathcal{B} and a doxastic strategy s_d such that $\mathcal{S} = (D, \mathcal{B}, s_d)$ is an autonomous-decisive system that achieves ψ up to n . If such \mathcal{B} and s_d exist, we can synthesize them.

Proof Sketch 1 We build a Kripke structure W'_D such any truth-observing strategy s_t in W'_D encodes a belief formation \mathcal{B} and a doxastic strategy s_d , such that \mathcal{B} is knowledge-consistent and s_d always follows dominant, current-state decisive possible-worlds strategies, and if s_t achieves ψ up to n , also s_d does.

We first determine the set of actions $\text{Act}(B)$ that an autonomous-decisive system can choose when currently having the belief $B = \{r_1, \dots, r_n\} \in \mathbb{B}$. An action a can only be chosen by a doxastic strategy of an autonomous-decisive system, if it is chosen by a dominant, current-state decisive possible-worlds strategy s_p for the current states of B (cf. Def. 5). Note that there may be several dominant possible-worlds strategies for B . We build a single reality r_B by the disjoint union of all alternative realities $r_B := ((\dot{\cup}_{r_i} \{S_i\}, \dot{\cup}_{r_i} E d_i, \dot{\cup}_{r_i} L_i, \dot{\cup}_{r_i} I_i), \dot{\cup}_{r_i} \{S_{ci}\})$. If necessary, we can make the realities disjoint by renaming their states but keeping their structure. In order to judge how well a current-state decisive, possible-worlds strategy has to perform in order to be dominant, we determine the maximal n_m up to which ψ can

be achieved by any possible-worlds strategy. Therefore, we check whether we can synthesize a strategy in r_B that achieves ψ up to $n_m = n$ down to $n_m = 0$, if necessary [21]. Note that in the worst case only the trivial goal true, i.e. $n_m = 0$, can be achieved. In order to check whether a is chosen by a dominant, current-state decisive strategy, we first check whether a is enabled at all current states of r_B . If a is enabled, we make a the only choice at the current states. Therefore, we build the modified Kripke structure $r_{B,a}$. We replace all arcs from the current states of r_B that are not labelled with a by arcs leading to s_{undef} . If we can synthesize a strategy $r_{B,a}$ that achieves ψ up to n_m , a is a current-state choice of that current-state decisive, possible-worlds strategy. By iterating over all actions, we collect the set of the actions $\text{Act}(B)$.

In order to judge how well the doxastic strategy s_d has to perform, we determine the maximal n_m up to which ψ can be achieved by any \mathcal{O} -truth-observing strategy in W_D . If $n_m < n$, then there is no (\mathcal{B}, s_d) such that \mathcal{S} achieves ψ up to n . Hence let $n \leq n_m$. We build the modified Kripke structure W'_D as follows. If W_D has an a -labelled arc from a state s labelled, we keep it only if there is a belief $B \in \mathbb{B}$ with $a \in \text{Act}(B)$ and $B \models L_{\mathbb{K}}(s)$ [7]. Otherwise, we redirect the arc to s_{undef} . We synthesize a dominant strategy s_t on W'_D [21]. We derive the belief formation \mathcal{B} from s_t by $\mathcal{B}(h) := B$ for a belief B with $s_t(h) \in \text{Act}(B)$. We define s_d to be the doxastic strategy defined by $s_d(\mathcal{B}(h)) := s_t(h)$. In case s_t achieves n , \mathcal{S} is an autonomous-decisive system. Otherwise the truth-observing strategies perform better. \square

4 Related and Future Work

Epistemology is the theory of knowledge and concerned with information-processing and cognitive success [10, 24]. Doxastic means “relating to belief” [9]. By using the term “doxastic”, we want to stress that our formalism focuses on beliefs. Note in particular that an alternative reality can vary arbitrarily from the ground truth W_D , and a HAS does not necessarily believe that the ground truth is even possible. In the epistemic logic literature, the semantics of doxastic languages are often given via *doxastic models*, that are special Kripke structures [12]. A doxastic model (S, v, \rightarrow_i) consists of a set of nodes S representing possible worlds w , a valuation function $v : S \rightarrow 2^{AP}$ for the set of atomic facts AP and a belief relation \rightarrow_i for each player i , that specifies “ i deems w' possible in w ” if “ $w \rightarrow_i w'$ ”. With other words, the belief of i at w is defined as the worlds accessible via the agent i ’s belief relation, \rightarrow_i [12, 19]. In this paper, we use complex possible worlds instead of the plain nodes of a Kripke structure. Our alternative realities encode the believed histories, the current states and possible futures. A HAS uses alternative realities to simulate its strategy in order to decide on its current action. In our framework, a reality constitutes an extensive form two-player zero-sum game, where the winning condition is defined by the list of LTL goals. While the possible-world strategies have perfect information on the beliefs, the belief formation, which we also regard as a strategy, is based on partial observations and the currently available knowledge/hard beliefs.

A couple of epistemic temporal logics have been suggested for specifying aspects of knowledge throughout time for multi-agent systems. These logics combine temporal logics with knowledge operators, like KCTL [4], KCTL* or HyperCTL*_{lp} [6]. They are interpreted over Kripke structures. But since an agent i has its local view, only certain propositions are assumed to be observable, so that an observational equivalence relation \sim_i on the traces arises. “Agent i knows φ ” then means that φ holds on all i -equivalent initial traces. The alternating time temporal logics (ATL) [1] has been developed for reasoning about what agents can achieve by themselves or in groups throughout time. In ATL, the path quantifiers of CTL are replaced by modalities that allow to quantify paths in the control of groups of agents. ATL is interpreted over concurrent game structures (CGS), which are labelled state transition

systems. By adding a knowledge operator, ATL has been extended to an epistemic variant, ATEL [15]. To this end the concurrent game is extended by an observational equivalence relation per agent modelling the agent’s limited view. In contrast to the above logics, a BLTL formula is interpreted on a belief B , i.e. a set of alternative realities. Since the set of possible beliefs \mathbb{B} is finite, a formula $K\phi$ means the finite conjunction $\bigwedge_{r \in B} r \models \phi$. But just like the logics above, we assume that a HAS can only partially observe the ground truth – our doxastic strategies cannot distinguish between histories that are observational equivalent. Our beliefs, however, cannot straightforwardly be expressed in terms of an equivalence on the ground truth, since an alternative reality may be a distinct Kripke structure and a belief does not have to include the ground truth.

The field of epistemic planning is concerned with computing plans (“a finite succession of events” [18]) that achieve the desirable state of knowledge from a given current state of knowledge [5]. DEL, dynamic epistemic logic, is a formalism to describe planning tasks succinctly by a semantic and action model based approach. Epistemic models capture the knowledge state of the agents, and epistemic action models describe how these are transformed. An evolution results from a stepwise application of the available actions. In [18] distributed synthesis of observational-strategies for multiplayer games are considered. While ATEL and DEL allow for reasoning about a combination of knowledge and strategies, we are interested in the belief *formation*. We ask whether there exists a belief formation that justifies a strategy that successfully achieves temporal goals within a given ground truth world.

Properties of belief formation are studied in the field of belief revision and update. Belief revision is done when a new piece of information contradicts the current information, and it aims to determine a consistent belief set. Belief updates may be necessary when the world is dynamic [13]. The works in this field are concerned with rational belief formation, following e.g. some guiding principle like making minimal changes [13]. In our work, we consider very general belief formation functions, since we focus on HASs. The belief of a HAS may be determined by a composition of different components, and there may not necessarily be an entity that ensures that the resulting belief is rational. We imagine that during the design, requirements regarding the rationality might be specified. So, approaches examining whether an appropriate belief formation exists might be valuable tools for the development of HASs.

BDI agents are rational agents with the mental attitudes of belief (B), desire (D) and intention (I) [22]. Beliefs describe what information the agent has, desires represent the agent’s motivational state and specify what the agent would like to achieve, while intentions represent the currently chosen course of action. These attitudes allow an agent balancing between deliberation about its course of action and its commitment to the chosen course of action. In our framework, an agent deliberates about its course of action at each state. We do not enforce commitment to a certain course of action, as we are interested in whether some belief formation exists. Nevertheless, the framework conceptually allows capturing notions of commitment, and we plan to examine these in future work. Basically, a chosen action represents a set of believed best possible world strategies. These can be considered as the current intent. So, a notion of commitment could require that (some) strategies of the previous belief are still best strategies in the current belief. An engineer may then specify when a HAS should be committed.

In this paper, we call a system autonomous-decisive, if it can choose its actions based on its belief content, i.e., when it always chooses actions that it believes can be extended to a best strategy. The concept that an agent adjusts its actions following the believed most promising strategy is certainly not new. For instance in [25], Wang et al. present an approach to motion planning for autonomous cars in competitive scenarios. *Ego* solves iteratively several rounds of trajectory optimization, whose fixed point represents a Nash equilibrium in the trajectory space of all cars. In [2], a two player game with communication is considered, and a generic framework is defined that uses belief revision techniques to take the communication into account. Player B makes announcements and then player A revises its

preferences by analysing what its best response would be.

Although we motivated our term of *autonomous-decisive* by referring mostly to HASs, we think the term is not limited to systems of a high level of autonomy, like level 3 to 5 of the norm SAE J3016. Our notion characterizes how a system takes decisions, while the SAE levels classify the interaction between driver and system. A lot of works characterize notions of autonomous systems. For instance, [14] focuses on the challenges of a foundation for autonomous systems. According to the authors, autonomous behaviour embodies five autonomous behavioural functions: perception, model update, goal management and planning and self-adaptation. They further state that decision-making is a central aspect and at the very heart of autonomous system engineering. In this paper, we clearly focus on decision-making. The paper [26] explores the intelligence and system foundations of autonomous systems. According to the authors, autonomy aggregates hierarchically from reflexive, imperative and adaptive intelligence. They consider autonomy as a property that enables a system to change their behaviour in response to unanticipated events during operation (thereby following [27]) and “without human intervention” (following [16]). Our notion of autonomous-decisiveness ensures that no human intervention is necessary for the decision-making. As our *ego* can only partially observe the world, one might argue that behaviour is unexpected, when observations of the ground truth do not match with the extrapolation of prior beliefs. When *ego* adjusts its beliefs and adopts its behaviour, it hence autonomously adjusts to unexpected events. We rather envision a different research emphasis to approach his topic. In order to characterize *ego*’s ability to adjust to unexpected events, we envision to develop notions of strategic/belief formation robustness that specify when an application domain differs too much from the envisioned design time domain to allow *ego* taking autonomous decisions.

Moreover, we plan to use the framework to tackle the question “What is relevant for the HAS?”, capturing the trade-off between observations and knowledge when the belief formation is given.

5 Conclusion

In this paper, we are concerned with safety critical systems that autonomously assess their situation and decide what to do based on this assessment. Such HASs operate in overwhelmingly complex environments and have to deal with partial observability and distorted perceptions. They construct an inner representation of the world that is based on insights of the world and develops over time, when new observations are made. We develop a doxastic framework in order to formalize, what it means that such a system autonomously takes decisions. In contrast to automatic systems, that decide based on their beliefs how to play out pre-determined strategies, a system decides autonomously, if it chooses an action that is justified by extrapolation of the system’s beliefs. Given a prioritized list of LTL goals, a given design context modelled as Kripke structure and given it has been decided on the build-in knowledge and the system’s possible beliefs, it is decidable whether a belief can be formed that enables a system to autonomously achieve its goals and we can synthesize a witness. We consider this work as a start to examining belief formation with respect to HAS’s strategic success.

Acknowledgements I like to thank the anonymous reviewers for their valuable comments.

References

- [1] Rajeev Alur, Thomas A. Henzinger & Orna Kupferman (2002): *Alternating-Time Temporal Logic*. *J. ACM* 49(5), p. 672–713, doi:10.1145/585265.585270.

- [2] Guillaume Aucher, Bastien Maubert, Sophie Pinchinat & François Schwarzentruber (2015): *Games with Communication: From Belief to Preference Change*. In: *PRIMA 2015: Principles and Practice of Multi-Agent Systems*, Springer, pp. 670–677, doi:10.1007/978-3-319-25524-8_50.
- [3] Christel Baier & Joost-Pieter Katoen (2008): *Principles of Model Checking*, chapter 5, pp. 229–239. The MIT Press.
- [4] Mario Benevides, Carla Delgado, Carlos Pombo, Luis Lopes & Ricardo Ribeiro (2008): *A Compositional Automata-based Approach for Model Checking Multi-Agent Systems*. *Electronic Notes in Theoretical Computer Science* 195, pp. 133–149, doi:10.1016/j.entcs.2007.08.030. Proc. of the Brazilian Symposium on Formal Methods (SBMF 2006).
- [5] Thomas Bolander (2017): *A Gentle Introduction to Epistemic Planning: The DEL Approach*. In: *Proc. of the 9th Workshop on Methods for Modalities, M4M@ICLA 2017, EPTCS 243*, pp. 1–22, doi:10.4204/EPTCS.243.1.
- [6] Laura Bozzelli, Bastien Maubert & Sophie Pinchinat (2015): *Unifying Hyper and Epistemic Temporal Logics*. In: *Foundations of Software Science and Computation Structures*, Springer, pp. 167–182, doi:10.1007/978-3-662-46678-0_11.
- [7] Edmund M Clarke, Orna Grumberg & Doron A. Peled (1999): *Model checking*. MIT Press, doi:10.1002/9780470050118.ecse247.
- [8] Werner Damm & Bernd Finkbeiner (2011): *Does It Pay to Extend the Perimeter of a World Model?* In: *FM 2011: Formal Methods*, Springer, pp. 12–26, doi:10.1007/978-3-642-21437-0_4.
- [9] Collins English Dictionary (2022): *doxastic*. Available at www.collinsdictionary.com/de/worterbuch/englisch/doxastic.
- [10] Collins English Dictionary (2022): *epistemics*. Available at www.collinsdictionary.com/de/worterbuch/englisch/epistemics.
- [11] Ronald Fagin, Joseph Y. Halpern, Yoram Moses & Moshe Y. Vardi (2003): *Reasoning About Knowledge*. MIT Press, doi:10.7551/mitpress/5803.001.0001.
- [12] Paolo Galeazzi & Emiliano Lorini (2016): *Epistemic logic meets epistemic game theory: a comparison between multi-agent Kripke models and type spaces*. *Synthese* 193(7), pp. 2097–2127, doi:10.1007/s11229-015-0834-x.
- [13] Peter Gärdenfors (1992): *Belief Revision*, chapter Belief Revision: An Introduction, pp. 1–28. Cambridge University Press, doi:10.1017/CBO9780511526664.001.
- [14] David Harel, Assaf Marron & Joseph Sifakis (2020): *Autonomics: In search of a foundation for next-generation autonomous systems*. *Proceedings of the National Academy of Sciences* 117(30), pp. 17491–17498, doi:10.1073/pnas.2003162117.
- [15] Wiebe van der Hoek & Michael Wooldridge (2002): *Tractable Multiagent Planning for Epistemic Goals*. In: *Proc. of the 1st Int. Joint Conference on Autonomous Agents and Multiagent Systems: Part 3, AAMAS '02*, Association for Computing Machinery, p. 1167–1174, doi:10.1145/545056.545095.
- [16] Ming Hou, Simon Banbury & Catherine Burns (2014): *Intelligent Adaptive Systems: An Interaction-Centered Design Perspective*. CRC Press, doi:10.1201/b17742.
- [17] Timothy B. Lee: *Self-driving cars – How terrible software design decisions led to Uber’s deadly 2018 crash*. Available at arstechnica.com/cars/2019/11/how-terrible-software-design-decisions-led-to-ubers-deadly-2018-crash.
- [18] Bastien Maubert, Aniello Murano, Sophie Pinchinat, François Schwarzentruber & Silvia Stranieri (2020): *Dynamic Epistemic Logic Games with Epistemic Temporal Goals*. *CoRR* abs/2001.07141, doi:10.3233/FAIA200088.
- [19] John-Jules Ch. Meyer (2003): *Modal Epistemic and Doxastic Logic*. In: *Handbook of Philosophical Logic*, Springer, pp. 1–38, doi:10.1007/978-94-017-4524-6_1.

- [20] Dorde Petrovic, Radomir Mijailovic & Dalibor Pesic (2020): *Traffic Accidents with Autonomous Vehicles: Type of Collisions, Manoeuvres and Errors of Conventional Vehicles' Drivers*. *Transport Infrastructure and systems in a changing world. Towards a more sustainable, reliable and smarter mobility. TIS Roma 2019 Conference Proceedings* 45, pp. 161–168, doi:10.1016/j.trpro.2020.03.003.
- [21] A. Pnueli & Roni Rosner (1989): *On the synthesis of a reactive module*. *Automata Languages and Programming* 372, pp. 179–190, doi:10.1145/75277.75293.
- [22] Anand S. Rao & Michael P. Georgeff (1998): *Decision Procedures for BDI Logics*. *Journal of Logic and Computation* 8(3), pp. 293–343, doi:10.1093/logcom/8.3.293.
- [23] David Romero, Peter Bernus, Ovidiu Noran, Johan Stahre & Åsa Fast-Berglund (2016): *The Operator 4.0: Human Cyber-Physical Systems & Adaptive Automation Towards Human-Automation Symbiosis Work Systems*. In: *Advances in Production Management Systems. Initiatives for a Sustainable World*, Springer, pp. 677–686, doi:10.1007/978-3-319-51133-7_80.
- [24] The University of Sheffield (2021): *Epistemology*. Available at www.sheffield.ac.uk/philosophy/research/themes/epistemology.
- [25] Mingyu Wang, Zijian Wang, John Talbot, J. Christian Gerdes & Mac Schwager (2021): *Game-Theoretic Planning for Self-Driving Cars in Multivehicle Competitive Scenarios*. *IEEE Transactions on Robotics* 37(4), pp. 1313–1325, doi:10.1109/TRO.2020.3047521.
- [26] Yingxu Wang, Konstantinos N. Plataniotis, Sam Kwong, Henry Leung, Svetlana Yanushkevich, Fakhri Karay, Ming Hou, Newton Howard, Rodolfo A. Fiorini, Paolo Soda, Edward Tunstel, Jianmin Wang & Shushma Patel (2019): *On Autonomous Systems: From Reflexive, Imperative and Adaptive Intelligence to Autonomous and Cognitive Intelligence*. In: *2019 IEEE 18th International Conference on Cognitive Informatics & Cognitive Computing*, pp. 7–12, doi:10.1109/ICCICC46617.2019.9146038.
- [27] David Watson & David Scheidt (2005): *Autonomous systems*. *Johns Hopkins APL Technical Digest (Applied Physics Laboratory)* 26, pp. 368–376. Available at www.jhuapl.edu/Content/techdigest/pdf/V26-N04/26-04-Watson.pdf