Mutual Influence-aware Runtime Learning of Self-adaptation Behavior

STEFAN RUDOLPH, University of Augsburg, Germany

SVEN TOMFORDE, University of Kassel, Germany

JÖRG HÄHNER, University of Augsburg, Germany

Self-adaptation has been proposed as a mechanism to counter complexity in control problems of technical systems. A major driver behind self-adaptation is the idea to transfer traditional design-time decisions to runtime and into the responsibility of systems themselves. To deal with unforeseen events and conditions, systems need creativity—typically realized by means of machine learning capabilities. Such learning mechanisms are based on different sources of knowledge. Feedback from the environment used for reinforcement purposes is probably the most prominent one within the self-adapting and self-organizing (SASO) systems community. However, the impact of other (sub-)systems on the success of the individual system's learning performance has mostly been neglected in this context.

In this article, we propose a novel methodology to identify effects of actions performed by other systems in a shared environment on the utility achievement of an autonomous system. Consider smart cameras (SC) as illustrating example: For goals such as 3D reconstruction of objects, the most promising configuration of one SC in terms of pan/tilt/zoom parameters depends largely on the configuration of other SCs in the vicinity. Since such mutual influences cannot be pre-defined for dynamic systems, they have to be learned at runtime. Furthermore, they have to be taken into consideration when self-improving their own configuration decisions based on a feedback loop concept, e.g., known from the SASO domain or the Autonomic and Organic Computing initiatives.

We define a methodology to detect such influences at runtime, present an approach to consider this information in a reinforcement learning technique, and analyze the behavior in artificial as well as real-world SASO system settings.

CCS Concepts: • **Computer systems organization** → **Self-organizing autonomic computing**; • **Theory of computation** → *Multi-agent learning*;

Additional Key Words and Phrases: Mutual influences, self-organization, self-integration, self-adaption, machine learning, organic computing, interwoven systems

ACM Reference format:

Stefan Rudolph, Sven Tomforde, and Jörg Hähner. 2019. Mutual Influence-aware Runtime Learning of Self-adaptation Behavior. *ACM Trans. Auton. Adapt. Syst.* 14, 1, Article 4 (September 2019), 37 pages.

<https://doi.org/10.1145/3345319>

We gratefully acknowledge the financial support of the German Research Foundation (DFG) in the context of the project CYPHOC (HA 5480/3-1, SI 674/9-1).

Authors' addresses: S. Rudolph, University of Augsburg, Eichleitnerstr. 30, Augsburg, 86159, Germany; email: stefan.rudolph@informatik.uni-augsburg.de; S. Tomforde, University of Kassel, Kassel, Germany; email: stomforde@uni-kassel.de; J. Hähner, University of Augsburg, Augsburg, Germany; email: joerg.haehner@informatik.uni-augsburg.de.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Association for Computing Machinery.

1556-4665/2019/09-ART4 \$15.00

<https://doi.org/10.1145/3345319>

1 INTRODUCTION

The research field of self-adaptive and self-organizing (SASO) systems emerged as response to the dramatic change Information and Communication Technology (ICT) has undergone within the past two decades: Back then, we started with single isolated and fully comprehensible systems with well-defined system boundaries. Afterwards, ICT components have been continuously coupled with each other and single monolithic solutions have been replaced by collections of autonomous entities or *agents* that communicate with each other and cooperatively solve tasks. Driven by initiatives such as *Autonomic Computing* [1], *Organic Computing* [2], *Proactive Computing* [3], *Multi-Agent Systems* [4], *Interwoven Systems* [5], or *Collective Adaptive Systems* [6], a shift in responsibilities has been established: from design-time to runtime and from engineers to systems themselves. As a result, various architectures and techniques have been established that allow for runtime self-adaptation in technical systems [7].

These autonomous systems typically act in shared environments with the goal to achieve and maintain a given utility function by means of adapting their configuration or parameterization in response to internal and external changes. This is often combined with machine learning mechanisms to improve the adaptation behavior over time, which also allows for handling unexpected or disturbed conditions [8, 9]. However, these existing techniques mostly neglect that configuration and behavior of other systems in the environment and may have strong impact on the degree of goal or utility achievement. These impacts become even more challenging when considering system-of-system compositions [10], goal-oriented holonic system structures [11], or Interwoven System constellations [12], where dependencies between systems emerge that affect different abstraction layers or application domains. An illustrating example for the latter case is the coupling of ICT with the smart grid [13].

Consider a surveillance network consisting of a potentially large number of smart cameras (SCs) [14, 15] as illustrating example: Each SC acts autonomously and tries to optimize, e.g., the coverage of the observed area, the tracking success of conspicuous objects, or the quality of 3D reconstruction of objects. In all three goal functions, there is an impact of the current configuration of other SCs in the vicinity: The field of view should either have no overlap with the own (coverage), should have overlap at a specific position and time (tracking), or full overlap as result of a certain ideal configuration (3D reconstruction).

What we can see in this simple example is that the success of the self-adaptation strategy has not just a local origin that is under control of the system itself—it is highly influenced by the current configuration of other systems [16]. We call these effects “mutual influences” to highlight that these influences stem from other systems and not the environment or possibly other sources. The shift of decisions from design-time to runtime also implies that such mutual influences cannot be fully foreseen and planned anymore. Consequently, systems need a methodology to detect them in the first place. This also requires mechanisms to incorporate this information in the self-adaptation strategy and to consider this in the self-learning behavior [17].

To be applicable to SASO systems in general, such a mutual influence methodology has to fulfill several requirements that are outlined in the following:

- **Runtime capability:** Due to the complexity that is present in modern systems, it is an increasingly challenging task to foresee all situations a system will face during its lifetime. It becomes virtually impossible because of the appearance and vanishing of other systems that affect the outcome of certain processes. Therefore, an influence-detection algorithm has to be executable during the runtime of a system to allow for an adaption to an unanticipated interaction partner, such as new devices from other stakeholders, manufacturers, or owners.

- **Heterogeneity:** To reach the goal of self-integration [18], a system faces the challenge to interact with several other systems that can be very heterogeneous in several aspects, e.g., scale, virtual/physical components, or connectedness. A special focus of this article is on the heterogeneity induced by the ownership of a system that limits the access and control of the other systems.
- **Compatibility:** Today, we face several long-living systems that are in operation for over a decade. Replacing such systems would be costly and time-intensive. Therefore, to integrate with such systems, it is necessary to keep the influence detection compatible to a wide-range of systems. This especially includes systems that do not have a cooperation mechanism, are non-adaptive, or use different adaption techniques.
- **Autonomy:** Recently, we can observe a dramatic increase in interconnected ICT systems surrounding us. These structures of systems bear new challenges, since it is much more difficult to predict their behavior. To avoid time- and cost-intensive maintenance of these systems, we require an influence-detection mechanism that can become an essential part of autonomous systems, i.e., that can be integrated into a system and afterwards decide on its own which influences have to be addressed.

In this article, we present a novel methodology that fulfills these requirements. We refine our preliminary work on mutual influence detection [19, 20, 21] and present a novel influence-aware learning mechanism. In particular, this article goes far beyond preliminary work by providing the following contributions:

- A redefined method for the detection of influences among systems including multi-component influences that extends previous work presented in Reference [21] by introducing a general workflow and more complex use cases (see Section 4);
- A novel extended method of this approach that is runtime-capable and consequently distinguishes correlations in utility and performance with others' configurations from noisy or coincidental effects (see Section 5);
- A novel approach for incorporation of this mutual influence information within a reinforcement learning algorithm that is applied to the problem of self-adaptation of the system configuration (see Section 6).

2 SYSTEM MODEL

As basis for this article, we describe the system model and the requirements a system has to fulfill to allow for runtime influence detection. Furthermore, we introduce three example applications for the methodology that serve as illustration purposes throughout this article.

2.1 Target Systems

We assume an overall system that is a composition of subsystems $A_1, \dots, A_n$ in a virtual or physical environment. We refer to the term (*sub*)system using the terminology from the Organic Computing domain [2], and, for better readability, omit the *sub* if it is clear from the context that not the overall system is meant. However, one might prefer to use the terms *entity* or *agent* instead. Each subsystem in the overall system can assume different configurations. Such a configuration typically consists of different components. We define the entire configuration space of a subsystem A_i as Cartesian product $C_i = c_{i1} \times \dots \times c_{im}$, where c_{ij} are the components of the configuration.

Consider, e.g., a router A_1 in a computer network as an illustrating example following the ideas of Reference [22]. It can take varying configurations into account, such as the processed network protocol or parameter settings [23]; e.g., an interval $c_{11} = [0, 100]$ for the timeout parameter in

Fig. 1. The basic reinforcement learning model.

seconds and the set $c_{12} = \{1, 2, \dots, 16\}$ for the buffer size in kilobyte. The entire configuration space of system A_1 would then be $C_1 = [0, 100] \times \{1, 2, \dots, 16\}$.

A further assumption is that the particular configurations of individual systems are non-overlapping, meaning each subsystem has its own set of configurations and cannot control those of other subsystems. This does not mean that the configuration components have to be completely disjoint in structure and values of the contained variables. For instance, two subsystems might have the same capabilities, which would lead to the same set of possible configurations in these attributes—but for different subsystems. Such a relation is explicitly allowed within the model.

Besides the configuration space, we need to consider the local reward. To apply the proposed method, each subsystem has to estimate the success of its decisions at runtime—as a response to actions taken before. This is realized based on a feedback mechanism, with feedback possibly stemming from the environment of the subsystem (i.e., direct feedback) or from manual reward assignments (i.e., indirect feedback). This resembles the classic reinforcement model [24], where the existence of such a reward is one of the basic assumptions. If there is no obvious way to create such a signal, it can be useful to apply more structured approaches from the field of goal-oriented requirements engineering [25], such as the *knowledge acquisition in automated specification* (KAOS) goal model [26] that has been applied to similar tasks [27]. Looking at the router example, a useful local reward can have different forms depending on the application scenario. For instance, a useful measure of success could be the throughput (see Reference [28]).

Relations to Reinforcement Learning: We resemble the basic reinforcement learning [24] (RL) model and its extension to multi-agent reinforcement learning [29] (MARL). The following paragraphs discuss how this relates to the model used in this article:

Figure 1 illustrates this basic model for RL. It consists of an agent that interacts with its environment. It has the ability to sense the current state of the environment and manipulate it by applying an action to it. Furthermore, a reward is provided to the agent that reflects the usefulness of the applied action, which is possibly related to the current or resulting state and can have stochastic components. The goal of the agent is to maximize the expectation of this reward over the runtime of the system by choosing the best actions available in each state. To find such a strategy, several approaches have been presented [30]. The model is usually formalized as a Markov Decision Process (MDP), which is a tuple (S, A, T, r) defining the following components:

- a set of states S ,
- a set of actions A ,
- a transition function $T : S \times A \times S \rightarrow [0, 1]$, $T(s', a, s) \mapsto p(s'|s, a)$, i.e., a function that gives the probability that action a in the state s results in the state s' .
- a reward function $r : S \times A \times S \rightarrow \mathbb{R}$, which gives a reward for each state transition.

Its basic form assumes discrete time steps $t \in \mathbb{N}$. In each time step t , the agent senses the environment's state $s_t \in S$. It then selects and applies an action $a_t \in A$. In the next step, $t + 1$, the agent receives a reward r_{t+1} that reflects the quality of its action in the given state. It then senses the new state s_{t+1} and starts over again. Looking at the transition function T , we see that the probability of the appearance of a specific state can depend on the previous state and the selected action. The agent's goal is to find a *policy* $\pi(s_t)$ that maximizes the expected discounted reward

$$\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t, s_{t+1}), \quad (1)$$

where $0 < \gamma < 1$ is a discount factor for the future rewards.

A widely used and intensively studied RL algorithm is Q-learning [31]. Like most RL techniques, Q-learning tries to solve the general RL problem, i.e., to find an optimal policy for a given problem with respect to the long-term reward. The main idea is to find a Quality-function $Q : S \times A \rightarrow \mathbb{R}$ that approximates the reward for each state-action pair and takes into account the long-term reward. To reach this goal, the value for each state-action-pair is initialized according to some of the various proposed methods, e.g., they are all set to a fixed value or they are set to a random value and updated afterwards according to the rule

$$Q_{t+1}(s_t, a_t) = Q_t(s_t, a_t) + \alpha \left(r_{t+1} + \gamma \max_a Q_t(s_{t+1}, a) - Q_t(s_t, a_t) \right), \quad (2)$$

where $Q_t(s, a)$ denotes the old Q -Value and $Q_{t+1}(s, a)$ the new one, each for a given state-action pair (s, a) . Furthermore, r_{t+1} denotes the reward received in time step $t + 1$ and therefore is the immediate reward for the action a_t taken in time step t . The discount factor $\gamma \in [0, 1)$ determines the fraction of estimated future rewards that is taken into account in the present step. The learning rate $\alpha \in (0, 1]$ determines how much the current experience, i.e., the current reward, is taken into account for approximating the Q -value.

An extension to the single-agent RL is MARL [29], where several agents interact with a single environment. Each of these agents is considered to observe and manipulate only specific parts of the environment, resulting in the agents mutually influencing their rewards. The formal approach is a generalization of the MDP and is called a stochastic game (SG). It is defined by a tuple $(S, A_1, \dots, A_n, T, r_1, \dots, r_n)$ where n is the number of agents. The components are defined analogously to the single-agent case. Such SGs are classified depending on the reward structure. They are called fully cooperative if $r_1 \equiv \dots \equiv r_n$ and fully competitive if $r_1(s', a_1, \dots, a_n, s) + \dots + r_n(s', a_1, \dots, a_n, s) = 0$, where a_i is the action of the i th agent. Games between these border cases are called mixed games.

We see several similarities between the initial system model and that from the RL and MARL domains. The initially discussed configuration is similar to the action set, since both are encoding the *choice* of the agent/subsystem. However, while actions are often encoded in a relative manner, the configuration space has to be encoded absolute; e.g., a robot could have an action *turn left* that would result in a different orientation, depending on its previous orientation. For the configuration, an absolute encoding is necessary, e.g., the orientation as a cardinal point. This can result in an implicit integration of the state in the configuration, e.g., when the current orientation is part of the state. Furthermore, the local reward is a concept adopted from the reward in the RL domain, but the difference is the focus on the locality; e.g., in a case where the local reward of each agent is equal (as in a fully cooperative game), the influence detection can still be applied but might be of limited use, since it is only possible to measure which agent has influence on the overall result. This result can still be interesting but, in this article, the focus is on studying systems where the notion of a local reward is matched.

Fig. 2. An example smart camera as used in our lab.

2.2 Example Applications

In the following, we introduce several example applications for illustration purposes. We start with two elementary use cases that allow for a simple description and introduction in the appearing challenges. Afterwards, smart camera networks are introduced that define a real-world application often used to demonstrate SASO concepts.

2.2.1 Collaborative Box Manipulation. The collaborative box manipulation serves as a first, rather simple scenario where mutual influences occur. It includes two robots and a heavy box. The robots can be configured to push or pull the box, but none of them is able to move the box alone, since it is too heavy. Therefore, the robots have to cooperate, i.e., both push or pull the box. This means that the configuration space is $\{PUSH, PULL\}$. The robots receive a local reward that is 1 if the box moves forward or 0 otherwise. Obviously, the configuration of one robot has an influence on the success of the other.

2.2.2 Two-Man Saw. The two-man saw scenario is a second example that is quite similar to the previous one. Again, it includes two robots, but, in comparison to the box example, a two-man saw is operated instead of pushing a box. The robots can be configured to push or pull the two-man saw (which results in the same configuration space as before), but they cannot handle the saw alone. Therefore, the robots have to cooperate, i.e., one has to push and the other one has to pull the saw (and vice versa). The robots receive a local reward that is 1 if the saw moves or 0 otherwise. Again, a human can easily recognize that the robots influence each other.

2.2.3 Smart Camera Network. Smart cameras (SCs) possess a built-in computation unit that can be utilized for several tasks such as image processing, object localization, and object tracking. An example of such a SC can be seen in Figure 2. Most SCs have pan, tilt, and zoom capabilities, and the computation unit is used to determine beneficial alignments for the camera. Beyond, SCs are equipped with wired or wireless communication devices that allow for communication with neighboring cameras and, therefore, form an SC network. Such SCs can be utilized to achieve different goals, including the tracking of objects, the identification of new objects, or the 3D reconstruction of objects. To apply the methodology of detection of mutual influences to SCs, a local reward is required. Obviously, this reward function heavily depends on the chosen goal. Regarding the configuration, we consider the three adaptable parameters of the camera's alignment (pan, tilt, zoom). To cover different classes of possible mutual influences in SASO systems, several goals can be considered:

- the maximization of the observed space,
- the coverage of fast-changing environments [14],

- the detection of previously unknown objects [32], and
- the 3D-reconstruction of objects [21].

Concluding the former discussion and mapping this application to the model of the target system, it is a realistic assumption that the cameras' configuration space is composed of their pan, tilt, and zoom configuration. Allowing a pan between 0 and 360 degrees, a tilt between 0 and 90 degrees, and a zoom between 12 and 18, this would be $[0, 360] \times [0, 90] \times [12, 18]$ for each camera. As stated before, the configuration spaces of the cameras are identical in structure, but each camera can assume an individual configuration. For the local reward, we can choose one of the introduced functions.

3 RELATED WORK

In the following, the related work in the area of influence formalization and quantification is presented. Further related work can be found in Section 4.1, where several dependency measures are introduced.

In the literature, several approaches can be found that try to formalize the mutual influences. Most of these approaches focus on the influence through direct or indirect interactions. For instance, a model for interactions is proposed by Keil et al. [33], but a method to detect the implicit interactions is not provided. Another common approach is to use *stit* logic for modeling the interactions in multi-agent systems [34–36]. The focus of these works is on the system specification and verification and therefore differs fundamentally from the focus of our work where the goal is an adaption at runtime.

Multi-agent reinforcement learning (MARL) is an active field of research [29, 37]. An overview and a useful taxonomy has been introduced by Busoni et al. [29], for instance. Following their taxonomy, the approach presented here can be useful in fully competitive and mixed games (each static or dynamic). It can be used in fully cooperative games as well, especially if the global payoff is a function of the local payoffs of the agents (cf. the smart-camera application). As stated by the authors, the complexity resulting from coordination is a major issue in MARL systems. The influence-detection mechanism presented here can keep this complexity to a minimum, since it allows to only coordinate with relevant partners.

In the following, we outline the related work from the MARL domain. Kok et al. [38] presented work that is based on so-called coordination graphs and an approach to solve the global coordination problem on a local basis if it is possible to decompose the global payoff function into a sum of local payoff functions. A restriction of this method is that it relies on inference rules that are hand-crafted. Furthermore, they focus on discrete state variables, which is not the case in the approach presented in this article. But, since they assume a given graph, the influence detection could be used to infer such a graph based on the roles of the agents.

An extension to this approach has been presented in Reference [39]. There, similar to the influence-based approach, the coordination graph is inferred at runtime, creating a transition from independent learners to coordinated action selection. It is based on a t-test between the maximally possible expected reward when the agents act in common and the expected reward from independent decisions. The approach is for general sum games and is limited to discrete state and action spaces, which means that it cannot be simply adopted to a continuous case. Furthermore, the work only shows the applicability for a Q-learning-based approach, and there is no trivial way to make it applicable to a wide range of algorithms. Methodologically, a major difference here is that the method is focused on finding the states in which agents should collaborate. This requires that all agents are always willing and able to cooperate.

DeHauwere et al. [40] demonstrated a specialized solution for mazes with two robots. They used a generalized learning automaton that uses the distance to another robot to learn how to avoid a collision by identifying states in which they have to coordinate. While this approach might be adapted for other tasks, it would be necessary to hand-craft the states of the learning algorithm each time.

Later, the authors presented an article focusing on a method to generalize the learned behavior for a single state over several states. However, the presented approach needs to identify the states where coordination is necessary [41]. Their method is similar to Kok's. However, they assume that the agents have already learned an optimal policy if acting alone, since this is necessary to find states that need coordination that renders it unusable for learning at runtime. Furthermore, the work is focused on sparse interaction, i.e., it identifies the states for cooperation and not the systems. This means the approach is not appropriate to find useful collaborators from a set of agents. Furthermore, the approach is limited to discrete state and action sets. Another difference to the approach presented here is that it relies on the Kolmogorov-Smirnov test used as a goodness-of-fit test, i.e., it tests if the distribution of points fits a model. However, this approach depends on a model, which is not always present, e.g., if not Q-learning is used but a policy-based RL approach. This would be to the requirement of independence from the control mechanism of the system.

The above approach has been extended to solve delayed coordination problems [42]. However, the limitation regarding the discrete state and action sets persists. Furthermore, the approach is still not independent from the control algorithm.

Lanctot et al. [43] proposed a measure called joint policy correlation. This has been applied prototypically to a two-agent laser-tag scenario. It is based on the repetition of the same scenario with several seeds that lead to different strategies of the agents. Afterwards, a matrix is formed that compares the average rewards of the agents against the agents from the other repetitions. The values against the initial opponent and the other opponents are then aggregated to create a measure that allows to see how much an agent has overfitted to its initial opponent. The goal of the work is very different from the one presented here, since it wants to create a measure on how much an agent overfits to the behavior of other agents. A measure of influence cannot be directly derived from this approach. Furthermore, in this article, we focus on a runtime learning approach that is not possible if the experiment has to be rerun several times.

The approach presented in this article has similarities with feature selection methods. These are categorized as *filter*, *wrapper*, or *embedded* methods. Wrapper methods are not applicable here, since they require multiple repetitions of the learning tasks, which is contrary to the goal of learning at runtime. In the field of filter methods, the minimum-redundancy-maximum-relevance [44] (MRMR) approach is frequently used. However, it relies on the elimination of redundancies between the features, which is not desired here, since it reduces the information one gets from the influence detection. Furthermore, it comes at an extensive computational cost compared to the methods proposed in this article.

Regarding the embedded methods, the most prominent instance is the regularized least-squares policy [45, 46, 47]. This and similar methods are tied to a specific learning algorithm and, therefore, breach the restriction of an applicability independent from the control algorithm.

Furthermore, there are attempts to combine self-organizing systems with learning algorithms, e.g., Boes et al. [48] introduced a MAS for the control of technical systems that has SASO characteristics. The method is called *Escher* and aims to deconstruct complex control problems by using an acyclic network of interacting agents that are able to learn from experiences. The algorithm has been applied to toy problems and an engine calibration. The approach differs from the one presented here, since they created a MAS to control a single technical system, while here we focus on several systems that are autonomous and have to adapt to the behavior of the other systems.

But, we believe that an application of the influence detection to the individual controller agents of Escher can be helpful if the different inputs of the controlled system have to be coordinated to reach optimal results.

Concluding the related work, there is no algorithm that allows to identify subsystems with continuous and discrete states and actions that influence other subsystems to find useful collaborators and is independent from the control algorithm.

4 MUTUAL INFLUENCES

In this section, we present our basic method for the detection of influences among autonomous SASO systems. This method is a refined version of preliminary work that can be found in References [19, 20, 21] and necessary to understand the subsequent variants presented in the following sections. Especially, the refinement brings the method into a wider scope by defining a general workflow and applies it to new scenarios. Besides extending the basic method in the following sections, we simultaneously incorporate more complex tasks, i.e., the detection at runtime and the possibilities of an intelligent SASO system to adapt its own behavior to detected influences.

4.1 Methodology for Detection

Based on the described target systems, we define the methodology for mutual influence measurement. The goal is to identify those components of the configuration of the other systems (i.e., value range and considered variables) that have influence on the system itself. After the identification of influencing configuration components, they can be addressed by a designer, e.g., by considering them in control algorithms, or by a self-adapting system itself, e.g., using a learning algorithm.

In general, we are interested in the question of whether a system as a whole is influencing other systems. However, to be more precise in the description of the influence, we want to detect those parameters where the optimal configuration values are somehow influenced by the current settings of the other systems. The basic idea of the following approach is to make use of stochastic dependency measures that estimate associations and relations between the configuration components of a system and the reward of a second system. The basic method assumes that the mutual influence between the systems is instantaneous, meaning that the reward of the system reacts to the configurations of the other systems in the same time step. However, the approach can be extended to detect delayed influences as well by measuring the dependency between the configuration components and a later reward [49].

In general, dependency measures are designed to find correlations between two random variables. In the following, we model the configurations of distributed systems as such random variables that reflect the system model where we face autonomous systems whose actions can be uncertain due to non-deterministic behavior or incomplete information, for instance.

Before introducing relevant candidates for the dependency measures, we outline the general workflow (cf. Figure 3) for an influence detection for each subsystem in the overall system.

- (1) *Observation*: Continuously observe the configuration and estimate the goal achievement of a productive system, i.e., the pan, tilt, and zoom of a camera and the corresponding reward in terms of the given goals. These observations will typically be done by the system itself, but they could as well come from an external entity.
- (2) *Distribution*: Gather configurations from other systems and provide your own one to them.
- (3) *Estimation*: Estimate the dependency value by relating the own local reward to the configurations of the other systems. The basic idea of the algorithm is to make use of stochastic dependency measures that estimate associations between the reward of a system A and the configuration components of a second system B . These dependency measures are

Fig. 3. The general workflow of influence detection.

designed to identify correlations between two random variables X and Y . The reward of A is identified with a random variable X and the configuration of entity B with a random variable Y . This mapping implies that if the association between X and Y is high, we also have a high influence of B on A , since it reflects that the configurations of B *matter* for the reward of A . Vice versa, if the association is low, we do not see an influence. There are several dependency measures available that might be suitable depending on the application. An overview for the most interesting measures is given in the following. Furthermore, it might be necessary to take other configuration parts into account. These are most likely the configuration parts of the influenced agent, i.e., the configuration of system A is relevant for the measurement of the correlation between the reward of system A the configuration of system B . In this case, we condition the calculation with the configuration of B . This means that we calculate the dependency of the reward and the configuration for each configuration of B separately in the discrete case. In the continuous case, we split the configuration in two (or possibly more) parts and calculate the dependency separately. For a more detailed explanation of the estimation please see References [19, 21].

- (4) *Evaluation*: This step compares the values calculated for the different (other) systems and their configuration components. There are several possibilities to do so. The first approach would be to compare the influence values to a fixed threshold, which is a valid option but can be difficult, since an appropriate threshold might not be available for each application. Second, the influence values can be compared on a relative basis between the systems. This leads to a ranking of systems according to their influence. A third possibility is to simulate an independent system with the same configuration space and compare the “artificial” value to the real values of the other systems. This allows to decide on a basis that only includes one other system. Additionally, it can be useful to calculate the p-value, i.e., a value that allows to estimate how likely the given outcome under the assumption of independence is.
- (5) *Adaptation*: This step addresses the influences in the control strategy, which can take various forms: Especially when applied at design-time, the designer can decide on a case-by-case basis. For a self-adaptive solution, we propose to use a learning algorithm that includes the configuration of the influencing systems in the situation description during runtime.

This process can be used either at design-time with a prototype or continuously at runtime. For the design-time variant, the system runs for a certain time while the configurations and rewards are logged. Afterwards, the calculation and decision process starts. During runtime, the steps have to be considered as a loop that runs continuously, but one can still decide on how long samples are gathered, since a distribution and recalculation may not be justified for each newly gathered sample.

Discussion of Dependency Measures. As mentioned before, our method relies on the utilization of basic dependency measures for the influence detection. As in Reference [19], we briefly introduce the most prominent measures and consider their advantages and drawbacks in the following list. The contained measures are then analyzed in the following evaluation to derive statements about their possible applicability for the ranking task in our method.

- *Pearson correlation:* Probably the most prominent instance is the Pearson correlation coefficient—sometimes just called “correlation coefficient” [50]. The main advantages are its simple implementation and its fast calculation. In this context, the major drawback is that only linear correlations can be detected, i.e., it can fail in case of more complex dependencies. Moreover, it is necessary to calculate the distance between realizations of the random variable that might make it not well suited for some problems. It assumes values between -1 and 1 , where -1 indicates a perfect negative linear correlation and 1 a perfect positive correlation. 0 means that there is no linear correlation. When comparing the influence of different systems, it should be considered to use the absolute values. The Pearson correlation coefficient r is defined as:

$$r = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^n (x_i - \bar{x})^2} \sqrt{\sum_{i=1}^n (y_i - \bar{y})^2}}, \quad (3)$$

where the x_i and y_i are the gathered samples, i.e., the configuration components and the reward. n is the number of samples and $\bar{x}, \bar{y}$ denote the mean values of the random variables.

- *Kendall rank correlation:* Another measure that is based on calculating the ranks of the gathered samples has been introduced by Kendall [51]. The measure can be computed rather fast. It can detect monotone dependencies, which is better than just linear dependencies but still can be insufficient for many applications. It assumes values between -1 and 1 , where -1 indicates a perfect monotone declining relationship and 1 a perfect monotone increasing relationship. For independent variables, a value around 0 should be considered. The Kendall rank correlation τ is calculated using:

$$\tau = \frac{(\#concordantPairs) - (\#discordantPairs)}{n(n-1)/2}. \quad (4)$$

A concordant pair are two samples $(x_i, y_i), (x_j, y_j)$ where the ranks of the elements agree, i.e., if $x_i > x_j$, then $y_i > y_j$; or if $x_i < x_j$, then $y_i < y_j$. The opposite of a concordant pair is called discordant. Cases with $x_i = x_j$ or $y_i = y_j$ are neither discordant nor concordant and are not handled in the basic variant.

- *Spearman rank correlation:* The rank correlation after Spearman is similar to Kendall’s, since it is based on ranks, too. Here, the ranks of the samples are calculated and instead of a comparison between the samples the Pearson correlation coefficient is calculated on the ranks. Therefore, the values range from -1 to 1 with 0 meaning that no dependency has

been detected. The Spearman correlation is calculated using:

$$\rho = \frac{\sum_{i=1}^n (rg(x_i) - \overline{rg_x})(rg(y_i) - \overline{rg_y})}{\sqrt{\sum_{i=1}^n (rg(x_i) - \overline{rg_x})^2} \sqrt{\sum_{i=1}^n (rg(y_i) - \overline{rg_y})^2}}, \quad (5)$$

where x_i and y_i denote the samples and n the number of samples. $rg(x)$ is short for the rank of x , i.e., the position of x if all samples are ordered by their value. $\overline{rg_x}$ is the average rank of the samples x_i and $\overline{rg_y}$ for y_i , respectively.

- *Distance covariance*: This measure is an extension to the Pearson correlation that takes the distance between the samples into account. Since these distances have to be calculated for the entire sample set, it is not suitable for an online calculation. The advantage of this method is that it is not limited to linear dependencies but can find all types of dependencies [52]. The distance covariance is calculated as follows: We first derive the Euclidean distances $a_{j,k} = \|x_j - x_k\|$ and $b_{j,k} = \|y_j - y_k\|$. Afterwards, we calculate $A_{j,k} := a_{j,k} - \overline{a_{j,\cdot}} - \overline{a_{\cdot,k}} + \overline{a_{\cdot,\cdot}}$ and $B_{j,k} := b_{j,k} - \overline{b_{j,\cdot}} - \overline{b_{\cdot,k}} + \overline{b_{\cdot,\cdot}}$ where $\overline{a_{j,\cdot}}$ is the mean of the j th row, $\overline{a_{\cdot,k}}$ is the mean of the k th column, and $\overline{a_{\cdot,\cdot}}$ is the mean of the whole matrix:

$$\frac{1}{n^2} \sum_{j=1}^n \sum_{k=1}^n A_{j,k} B_{j,k} \quad (6)$$

- *Mutual information*: A quite different approach has been taken by Shannon in the context of information theory [53]. The basic variant can be used for discrete random variables and can be calculated online. Furthermore, all types of dependencies can be found. A disadvantage is that the maximal possible value depends on the structure of the random variables, i.e., the values can be normalized between 0 and 1, but the comparability to other variables with different structures might be limited. The mutual information is defined as:

$$I(X; Y) := \sum_{x \in X} \sum_{y \in Y} p(x, y) \log \left(\frac{p(x, y)}{p(x)p(y)} \right), \quad (7)$$

where $p(x, y)$ is the joint probability of the events x and y and $p(x)$ and $p(y)$ are the marginal probabilities of x and y . The probabilities can be easily approximated using a frequency counting (i.e., a maximum-likelihood approach based on discrete finite random variables). There exists also a continuous variant of the mutual information measure. Here, the calculation is more complex, since the estimation from samples needs more advanced techniques. The most common approach is based on a k -nearest neighbor method [54]. Another approach is the manual binning of the values to use the discrete version, which can bear problems, because the results might depend highly on the chosen binning parameters. An automatic binning can be done with the maximal information coefficient (see below).

- *Maximal information coefficient*: This measure is an extension to the mutual information for real-valued data. It automatically calculates the binning that results in the maximal mutual information for a given data set. A drawback is that there is no method for the online calculation. The basic formula is:

$$MIC(X, Y) := \max_{n_x n_y < B} \frac{I(X; Y)}{\log(\min(n_x, n_y))}, \quad (8)$$

where n_x and n_y denote the number of bins for X and Y . This means that the number of bins is limited by a threshold B that is by default determined depending on the sample size.

The denominator serves for the normalization of the value. However, it is computationally expensive to determine the values for all possible binnings. Therefore, a heuristic is introduced. For details please see Reference [55].

It should be mentioned that the choice of a dependency measure does not have to be exclusive. It can be useful to calculate several measures in parallel or to perform an iterative process, which initially uses computationally light-weighted measures and adds more powerful and computationally expensive methods only if necessary.

4.2 Evaluation

Initially, we aim at investigating which measurement from the previous subsection is most suitable for the detection of mutual influences. Therefore, we consider the three initially introduced application scenarios that are characterized by increasing complexity.

4.2.1 Collaborative Box Manipulation. In this scenario, two robots have the task of pushing a box. Each of them can either *PUSH* or *PULL* the box, but the box only moves if both of them *PUSH*. This leads to a reward of 1 while all other combinations lead to a reward of 0. Since some of the dependency measures do not allow for categorical random variables—such as the *PUSH* and *PULL* used here—the configuration has been mapped to the numbers 0 and 1.

For the evaluation of this application, we conducted 100 experiments and measured the influence using each of the previously introduced dependency measures. For each run and each measure, two values have been calculated. The first one is the influence of robot *B* on robot *A*. The second one is the influence of a notional robot. This notional robot has the same capabilities as the real robot *B*, but his actions do not influence robot *A*. The value is calculated to assess how reliable the detection of the influencing robot is. This can be done by comparing in which fraction of the runs robot *B* has been found more influential than the notional robot.

The results are depicted in Figure 4. There, we see that most of the measures perform similarly in this scenario. Except for the continuous MI approximated with the Kraskov method, 40 steps are sufficient to distinguish the influencing robot from a non-influencing in each of the 100 runs. The vast majority is already correctly detected after 15 steps. The continuous MI eventually finds the influence in each run. However, the detection speed is rather slow. This is due to the fact that the continuous MI has problems with the discrete values that are assumed in the problem. Furthermore, the continuous MI has not been calculated for less than 10 steps, since a minimum number of samples is required. Therefore, the first three data points should be disregarded.

Concluding the results, the influence in this elementary use case can be detected quite easily with the previously introduced method. Even though the continuous mutual information shows a little slower detection, the selected measure for the task is not too important, since the detection is over all quite fast.

4.2.2 Two-man Saw. As a second example, we evaluate the conditioned measurement in an elementary use case, the two-man saw. The example is inspired by two robots that operate a saw that only moves if both of them move it in the right direction, i.e., one *PUSH*es and the other *PULL*s or vice versa, depending on the current position of the saw. If the saw moves, each robot gets a reward of 1, otherwise 0. Again, the categorical configurations *PUSH* and *PULL* have been mapped on numerical values to make all dependency measures applicable.

The results are depicted in Figure 5. As for the collaborative box manipulation, the actual influence has been calculated using the seven dependency measures and compared with a measurement of a notional robot that assumes uniformly distributed random configurations but has no influence on the actual outcome of the experiment. The figure shows in how many of the 100 independent

(a) An overview of the full 100 steps evaluated.

(b) A more detailed view on the first 30 steps.

Fig. 4. The results for the collaborative box manipulation. The graphs show the fraction of 100 runs in which the influence of robot B is detected higher as the influence of the notional robot.

runs the influencing robot has been found to be more influential than the notional (not-influencing) robot. The first graph shows the result for the method used in the previous evaluations, i.e., without a consideration of other configurations. The detection is between 30% and 50%, which is below the expected value of 50%. This is due to the fact that the run will only be counted as correctly detected if the value is higher but not if both values are equal, which is quite often the case in this scenario. The second graph shows the results with the consideration of the other configuration, i.e., the calculation of the dependencies is conditioned under the configuration of the first robot. We see that influences are almost perfectly detected after 30 steps, which is close to the result in the collaborative box manipulation. Also similar is the small weakness of the continuous MI measure, which needs about 90 steps to catch up.

(a) The detection if a single estimator is used and the own configuration is not considered.

(b) The detection if the own configuration is considered, i.e., there is one estimator for each of the configurations (Push and Pull).

Fig. 5. The results for the two-man saw use case.

Concluding the evaluation of the two-man saw application, we have seen that even in rather simple applications, the problem with other configurations can lead to influences that cannot be detected with the basic variant of the influence detection. However, the improvement that includes the other configurations by conditioning the measurement shows very good results for such cases as well.

4.2.3 Smart Camera Network. As a third application for the evaluation, we consider an example for the SC network domain. To recap briefly, SCs are surveillance cameras that are equipped with computational capabilities that can be used for several tasks, including image processing. Furthermore, these cameras are interconnected via a network that allows them to exchange data and to coordinate. For this article, we stick to so-called PTZ cameras that allow for an automatic adjustment of the pan angle, tilt angle, and zoom of the camera. As previously described, there are several

Fig. 6. *SCN 1*. A top-down view on a smart camera network. The black dots depict cameras surrounded by a circle that marks their potential observable area. The red shapes show the field of view for an exemplary PTZ configuration. The yellow arrow indicates from where and in which direction the objects of interest move.

reasonable goals for an SC network. For this evaluation, we stick to the goal of a 3D-reconstruction of the observed targets. In particular, this means to observe the targets from different perspectives and cameras at the same time. To achieve this, the cameras get a reward of 1 for each object that is observed by at least two cameras in a time step.

A top-down view on scenario *SCN 1* is depicted in Figure 6. There, we see three black dots that represent one camera each. Around the dots, there are colored circles. Each of them represents the area that is potentially observable by one of the cameras if it chooses an according PTZ configuration. Furthermore, we see example areas that are currently observed marked by red lines. Camera 1 and Camera 2 share a common area that can be observed by both at the same time. In contrast, Camera 3 is isolated and does not share a common area with one of the others. The yellow arrow marks the entry point and direction for objects that move through the scene in the area between Camera 1 and Camera 2.

Ten independent runs of this scenario have been conducted in a Mason¹ simulation, where each of the cameras assumes a uniformly sampled PTZ configuration in each time step. The pan angles are between 0 and 360 degrees, the tilt angle between 120 and 180 degrees, and the zoom between 12 and 18. In each of the runs, we compared the influence of Camera 2 on Camera 1 and the influence of Camera 3 on Camera 1. For each of the measures, there are three comparisons: one for the pan, one for the tilt, and one for the zoom of the cameras. It is expected that there is a clear trend towards a higher influence of Camera 2 over Camera 3 in general and especially for the pan and tilt, since these configurations determine if it is possible to gather a reward for Camera 1 or not.

The results are depicted in Figure 8. There, we see that after a few hundred steps most measures allow a correct detection in 100% of the cases. The continuous mutual information and maximal information coefficient, however, take a little longer to reach this level of certainty. For the pan, we see that the distance correlation shows the best result compared to the other measures. However, the other measures work as well with a higher sample size. Even though the MIC finds an influence quite reliably, the other measures do not show a definite result, which can be explained by the minor influence of this configuration's components, i.e., in most cases, the zoom does not determine if there is a positive reward at all but only the height of its positive reward.

Concluding the results, we have seen that it is possible to detect the influences in simple examples based on real-world applications, such as SC networks. We will see how to use this in scenarios with more complex requirements in the remainder of this article.

¹Mason is a multi-agent simulation framework for Java [56].

Fig. 7. *SCN 2*. A top-down view on a smart camera network. This scenario is only slightly changed from *SCN 1* by adjusting the flow of targets represented by the yellow arrows. The black dots depict cameras surrounded by a circle that marks their potential observable area. The red shapes show the field of view for an exemplary PTZ configuration.

After the case that has been solved without the consideration of other configuration components, we focus on a scenario with few changes from scenario *SCN 1*. A top-down view on scenario *SCN 2* is depicted in Figure 7. In the figure, we see three black dots that mark the position of the cameras. Camera 1 and Camera 2 are quite close and potentially share a common field of view that is marked by the colored areas around the cameras. Camera 3, however, is separated and cannot observe the same areas as the other two cameras. The only difference to the previously evaluated scenario is where the objects of interest appear and move. In the last scenario, they were limited to the area between Camera 2 and 3 appearing from the south. Here, we see that they appear from the north and the south on the entire area around Camera 1 and 2, i.e., the two cameras are surrounded by objects. The influences from Camera 2 and Camera 3 on Camera 1 have been measured for the pan, the tilt, and the zoom. We expect that the pan and the tilt of Camera 2 have the most influence. Since a reward greater than 0 is much more likely in this scenario than in the previous, we also expect the zoom to play a more important role.

The results are depicted in Figures 9 and 10. As before, 10 independent runs have been conducted, and we see the ratio of runs in which the configuration components of Camera 2 have been identified as more influencing than those of Camera 3. In Figure 9, we see the results for the previous method. Even though the camera placement is identical to the previous scenario, we see that the pan is not detected within the first 2K steps. Furthermore, the tilt cannot be detected by the Pearson correlation. This is because it is necessary to consider the other configuration components in this scenario. Therefore, we adopted the conditioned calculation, i.e., we split each of the other configurations into two parts and sort the points in different buckets, depending on which configuration has been assumed. Since we face five other configurations, this leads to $2^5 = 32$ buckets. The values calculated for the buckets are then added up to get an aggregated result. The results achieved by this method are depicted in Figure 10. We see that, using this method, the issues experienced in the first experiment do not appear, and a flawless detection is ensured.

Concluding the results, we have seen that small changes in the setting can make it necessary to consider the other configuration components in the influence detection. A fast and reliable detection can be achieved by conditioning of the dependency measures.

(a) The results for the pan.

(b) The results for the tilt.

(c) The results for the zoom.

Fig. 8. The results for scenario *SCN 1* using the general method. Each graph shows the fraction of runs in which the influence of the configuration component (pan, tilt, or zoom) of Camera 1 on Camera 0 is detected to be higher than the influence of Camera 2.

5 INFLUENCE DETECTION AT RUNTIME

Modern ICT systems often cannot be tested entirely at design-time, since not all circumstances the system faces can be foreseen due to a lack of awareness or time constraints. This applies especially to the type of systems that this article focuses on: self-adaptive and self-organizing

(a) The results for the pan.

(b) The results for the tilt.

(c) The results for the zoom.

Fig. 9. The results for the scenario SC2 using no conditioning. Each graph shows the fraction of runs in which the influence of the configuration component (pan, tilt, or zoom) of Camera 1 on Camera 0 is detected to be higher than the influence of Camera 2.

systems acting in shared environments. Obviously, it is hardly possible to foresee all situations a system will face if the considered system has to act in parallel to other systems (that may be under the control of another stakeholder) and both should be able to adapt to new goals or environmental changes. Therefore, this section shifts the focus towards the detection of influences at runtime. A

(a) The results for the pan.

(b) The results for the tilt.

(c) The results for the zoom.

Fig. 10. The results for the scenario SC2 using two parts for conditioning. Each graph shows the fraction of runs in which the influence of the configuration component (pan, tilt, or zoom) of Camera 1 on Camera 0 is detected to be higher than the influence of Camera 2.

special examination of this is satisfied, because of the problem of static behavior of the systems that introduces disturbances in the measurement by causing correlations that are not based on causality but coincidence.

5.1 Methodology for Detection

Regarding the detection at runtime, we have to consider that the proposed influence-detection algorithm relies on dependency measures that estimate the correlation between two random variables. If the method is applied at runtime, this can lead to wrongly detected influences, because such correlations can appear without an underlying causality. The previous experiments demonstrated that it is possible to infer a causality from the correlation between the actions of two sub-systems, since we enforced the configurations to be randomly selected from independent uniform distributions. However, for the detection at runtime, it has to be factored in that we can face autocorrelations and other disturbances within the configurations of the systems.

The approach to avoid this is to rely mostly on randomized configurations that appear naturally in applications that use reinforcement learning during runtime. This is due to the need of exploration in these tasks to find optimal strategies and to avoid to “getting stuck” in a local optimum. This happens often with greedy approaches, because the algorithm only tries one behavior, and if it is “good,” it sticks to it and misses out on other strategies that may lead to better results. There are different strategies to avoid this behavior. An easy and reliable approach is to use an ϵ -greedy action selection, i.e., the algorithm will stick to the action that it has evaluated as the best one so far most of the time. However, with a probability of ϵ , it will use a random action regardless of the evaluated usefulness of it so far. This means a natural approach to the issue of correlation without causality is to only use the samples that are formed from such exploration steps. But this would lead to a significant lower amount of samples. To avoid this, we will also examine in detail how much “randomness” is necessary to allow the measurement to function properly by conducting experiments with different levels of ϵ . Identifying the lowest level of ϵ will ensure the fastest detection of influences without running the risk of falsifying the measurement.

One might expect that a very high level is crucial for a successful detection. However, the correlation mainly appears due to the repetition of specific patterns in both systems in the same frequency. If such patterns are brought out of sync, a correct detection is possible despite the repetition, i.e., a high autocorrelation does not necessarily mean that the measurement is flawed. For example, two systems *A* and *B* each switch back and forth between their two configurations 1 and 2. The system *A* gets a reward of 0.5 if it applies 1 and 0.8 if it applies 2. If both systems switch their configurations in the same time step, the measurement will be falsified, because system *B*’s configuration will correlate with the reward of system *A*. If one of the systems randomly does not switch the configuration, the behavior becomes “asynchronous” and the measurement works well.

5.2 Evaluation

In the following, the evaluation of the runtime detection and the method to adapt to the influences are evaluated. For the remainder, the evaluation parts of the article focus on the smart-camera scenarios only. At this point, the consideration of the elementary-use cases would not add to the discussion of the approach and are therefore omitted.

We start on SCN 2, which was introduced in Figure 7. Briefly recapping, we have three cameras, where Camera 1 and Camera 2 influence each other due to an overlap in their potential observable space. Camera 3 does not overlap with one of the other cameras and, therefore, it does not influence the other two. Previously, we saw how the detection works when the configuration is chosen randomly from the full configuration space.

In the following, the experiments are based on control mechanisms that use Q-learning, using discretized states and relative actions. Q-learning has been chosen since it is widely used and well understood. For the concrete implementation, we limit the alignment of the camera to 12 pan angles, i.e., it is from $s_p = \{0, 30, 60, \dots, 300, 330\}$, to 3 tilt angles, i.e., it is from $s_t = \{120, 150, 180\}$, and to 2 zoom levels, i.e., it is from $c_z = \{12, 18\}$. This results in $12 \cdot 3 \cdot 2 = 72$ states for each camera. The camera can increase, decrease, or leave the pan, tilt, and zoom in each time step resulting in $3 \cdot 3 \cdot 3 = 27$ actions that can be applied. This means that the configuration space is $C = s_p \times s_t \times s_z$ for each camera.

The first evaluation step is to find out how the limitation of the configuration space and the order of states that are visited affect the influence detection. The results can be found in Figure 11 and 12. The system has run for 10K steps by applying random actions, and the figure shows the percentage of runs in which Camera 1 detects Camera 2 as more influencing than Camera 3. The results for the pan and zoom are depicted in Figure 11. There, we see that each measure can find the influences very reliably within the 10K steps with the distance correlation being much more unreliable using less samples. In Figure 12, the results for the tilt are depicted. We can see that only few steps are needed to find the influence, i.e., in only 600 steps each of the measures can detect the influence very reliably.

In the next step, we analyze how the detection rate changes if we switch from a purely random action selection to ϵ -greedy. The ϵ is varied between 1 and 0.05 according to Table 1. It reflects that ϵ is set to 1 at start and is then decreased by 10% each 10K steps. The results can be found in Figure 13. There, we see the detection rate if single-agent Q-learning with an ϵ -greedy action selection is used and the last 10K samples are used for the detection. The Q-learning algorithm uses a low $\alpha = 0.1$ to handle fluctuations in the reward signal and a high $\gamma = 0.9$ to allow the learning of a sequence of actions. Figure 13(a) illustrates that for the pan the detection works well for the first 5K steps, i.e., for $\epsilon > 0.66$. Afterwards, the detection works worse for some of the measures with an average of about 80%. The continuous mutual information sticks out here with a perfect detection even with low values of ϵ . For the tilt, in Figure 13(b), we see a similar effect, but it is way less distinct and occurs only for lower ϵ . Figure 13(c) shows the graph for the zoom, which turns out to be way more sensitive regarding ϵ and makes the measurement unreliable.

Concluding the results, we have seen how the application of a control algorithm affects the influence detection in contrast to a random selection of configurations. Due to the disturbances and autocorrelations, it is possible that influences are not detected correctly. However, if it is ensured that an adequate amount of “randomness” is used to select the actions, it results in proper results. For the example examined here, we have seen that the usage of 66% randomly selected actions allows for a precise identification of influencing systems. If the usage of less randomly selected actions is necessary, an amount of greedy actions should be removed from the set for the calculation of influences until the required amount of “randomness” can be ensured.

6 SELF-ADAPTING TO INFLUENCES

In the previous sections, we focused on the detection of the influences, which is a crucial step towards the goal of optimal behavior. In this section, we introduce an approach to self-adaption to these identified influences, i.e., the exploitation of the influences if they have been discovered. For this adaption, the focus is not on the presentation of a single algorithm but on giving a general methodology based on the RL model and realizing it with Q-learning.

6.1 Methodology for Self-adaption

In general, there is a variety of options available to address the influences, since the detection is designed to work independently from the control algorithm of the system. Possible reactions to

(a) The results for the pan.

(b) The results for the zoom.

Fig. 11. The results for the detection of influences in SCN 2. The influence-detection rate for the **pan** and **zoom** of Camera 1 versus Camera 2 based on 10 independent runs. The measurement is similar to the graphs before but with the discretization of the state space and a Q-learning that applies the actions randomly. The graphs show the number of runs, in which the influence of Camera 1 is measured higher than the influence of Camera 2, for different numbers of samples.

influences range from a hand-crafted solution during design-time to a system that adapts during runtime in static patterns and finally to a full self-learning behavior. For the latter, we get an impression on how the subsystem's learning algorithm and influence-detection workflow (cf. Figure 3) interact by studying Listing 1. Potential candidates for such a learning algorithm can be found in the area of MARL. However, we rely on a basic principle that unifies three properties:

- (1) It is possible to start out with single learners and later on learn the cooperation,
- (2) it can be applied to many reinforcement learning algorithms,

Fig. 12. The results for the detection of influences in SCN 2. The influence-detection rate for the **tilt** of Camera 1 versus Camera 2 based on 10 independent runs. The measurement is similar to the graphs before but with the discretization of the state space and a Q-learning that applies the actions randomly. The graphs show the number of runs, in which the influence of Camera 1 is measured higher than the influence of Camera 2, for different numbers of samples.

- (3) it can be used with systems that have heterogeneous configuration spaces and control algorithms, e.g., some of the systems have a static behavior.

The basic adaption principle is depicted in Figure 14 and shows how the *adapt learning algorithm* from Listing 1 has been realized. There, we adopt the basic RL model introduced in Section 2.1. We assume that previous measurements have resulted in a detection of an influence from system A on system B, meaning B should react properly regarding the configuration of A. Therefore, we integrate the value of the configuration c_t at time t into the state s_t resulting in a state s'_t . For example,

Table 1. The Values for ϵ in the ϵ -greedy Action Selection

Step	Start	10k	20k	30k	40k	50k	70k	80k
ϵ	1.0	0.9	0.81	0.73	0.66	0.59	0.53	0.48
Step	90k	100k	110k	120k	130k	140k	150k	160k
ϵ	0.43	0.39	0.35	0.31	0.28	0.25	0.23	0.20
Step	170k	180k	190k	200k	210k	220k	230k	240k
ϵ	0.19	0.17	0.15	0.14	0.12	0.11	0.1	0.09
Step	250k	260k	270k	280k	290k	300k	310k	320k
ϵ	0.08	0.07	0.06	0.06	0.05	0.05	0.05	0.05

It starts at 1.0 and decreases by 10% every 10k steps until it reaches 5%.

ALGORITHM 1: The main loop for each subsystem using the influence detection and adaption at runtime.

initialization;

for each step do

 observe own configuration and reward;

 distribute observations to other subsystems;

if enough observations then

 estimate influences;

 evaluate estimated influences;

if influence found then

 adapt learning algorithm;

end

end

 learning algorithm decides next action;

 update learning algorithm with new observation;

end

if B 's states are given through a position represented as (x, y) -coordinate between 0 and 1 each, its state space is $S = [0, 1] \times [0, 1]$. Assuming the influence detection has found A 's configuration component c_B , which represents the speed of B and is between -1 and 1 , as influencing, the state space S will be extended to a new state space $S' = S \times [-1, 1] = [0, 1] \times [0, 1] \times [-1, 1]$.

There are two points to discuss in this approach. The first one is how the influenced system B will know the current configuration of A . This is not an issue with delayed influences but only with immediate. It can be resolved by letting the influenced system wait until it can observe the decision either by a message of the influencing system or through sensors. If several influences in the system are detected, this can lead to chains of systems that wait on each other and it is necessary to check for each of the integrations in the state space that the graph they form is acyclic.

As mentioned before, this method is independent from the learning algorithm as long as it fits the RL model. But the extension of the state space can be seen as a special instance of transfer learning in the RL domain [57]. This makes algorithms that have a strong transfer capability more suitable for this task. In this article, we rely on Q-learning, since it gives an easy and natural way to realize the transfer during runtime that will be explained in the following: Assuming that c_B can take k different values from each state s , there will be k new states $s'_i, i \in \{1, \dots, k\}$. Before the expansion of the state space, the system holds a current Q-value for $Q(s, a_j)$, where $a_j \in A$ are the possible actions. To transfer the already-learned knowledge, the new Q-values are set to

(a) The results for the pan.

(b) The results for the tilt.

(c) The results for the zoom.

Fig. 13. The results for the detection of influences in SCN 2. The influence-detection rate for the pan, tilt, and zoom of Camera 1 versus Camera 2 based on 10 independent runs. The measurement is similar to the graphs before, but here the camera is controlled by a Q-learning algorithm with ϵ -greedy strategy, where the ϵ is falling from 1 to 0.05 in 10%-declines each 10K steps. The graphs show the number of runs, in which the influence of Camera 1 is measured higher than the influence of Camera 2, using the last 10K steps as samples.

Fig. 14. The proposed method to adapt to other influencing agents in terms of the reinforcement learning model.

$Q(s'_i, a_j) := Q(s, a_j)$, for each $i \in \{1, \dots, k\}$ and $a_j \in A$. The advantage here is that in states that are not affected by the other system, the Q -value will remain at the correct value. In turn, in states where the system is influenced, the values will be updated to their true value using the upcoming experiences.

If the state space is continuous, the standard Q -learning cannot be used. An alternative would be for example the extended classifier system (XCS-R) with continuous states [58, 59] that has good online learning capabilities and is transferable [60].

A further issue is that in the previous experiments about runtime detection, we have seen that a small amount of runs might detect the independent systems as more influencing than they actually are. This behavior is temporary and can be corrected by using a bigger sample size. However, during the runtime of the system, it is not possible to determine if it is a malicious detection. Therefore, we introduce the factor that the real measurement has to be higher than the calculated notional independent counterpart that is used as baseline. Furthermore, the different configuration parts are ranked and after each influence calculation only the highest-ranked configuration is integrated in the state space of each camera. Even though this mechanism is likely to catch most erroneous integration, they cannot be avoided entirely. But, if a wrong integration is detected later, it can be corrected easily. To revert the false state space extension, the newly created states are merged back together by averaging the Q -values of the states, i.e., $Q(s, a_j) = \sum_{i=1}^k Q(s'_i, a_j)$ for every action $a_j \in A$.

6.2 Evaluation

In the next part, we examine how the influence measurement reacts to the adaption that has been introduced previously. To do so, we integrated the three configuration components of Camera 2 in the state of Camera 1 at the start of the simulations. In Figure 15, we see similar results for the independent learning with especially bad results as for the maximal information coefficient and the distance correlation.

As a last comparison in this scenario, we have a look at the rewards received with the different methods: the independent learners in Figure 16(a), the integration of the configuration of Camera 2 in the state of Camera 1 in Figure 16(b), and the dynamic integration of the configurations at runtime in Figure 16(c). We can see that the independent learners are not able to reach an appropriate result, since they lack a coordination mechanism. The integration at start can reach an optimal result of about 25 after 250K steps. The dynamic expand can reach a similar result but learns slower. This is because the transfer mechanism cannot play out its advantages here, since the Q -values differ for all states in this example.

In Figure 17, we see scenario SCN 3, another scenario from the SC domain. The figure shows a top-down view on an SC network with six cameras depicted as black dots. Each of the cameras has a corresponding partial circle around it that visualizes the space potentially observable space

(a) The results for the pan.

(b) The results for the tilt.

(c) The results for the zoom.

Fig. 15. The results for the detection of influences in SCN 2. The influence-detection rate for the pan, tilt, and zoom of Camera 1 versus Camera 2 based on 10 independent runs. The cameras are controlled with the same algorithms as before but the pan, tilt, and zoom configurations of Camera 2 are integrated in the state space of Camera 1 at the start of the runs. The graphs show the number of runs, in which the influence of Camera 1 is measured higher than the influence of Camera 2, using the last 10K steps as samples.

(a) The results with single learners.

(b) The results if the pan, tilt, and zoom configurations of Camera 2 are integrated in the state space of Camera 1 at the start of the simulation.

(c) The results if the configurations are integrated in the state space of Camera 1 dynamically during runtime.

Fig. 16. The figure shows the rewards received by the system if the Q-learning algorithm is applied to SCN 2 with single learners, with an integration at the start, and dynamic integration.

of the cameras. The circle is only partial, since the vision is blocked by pillars that are shown as gray squares. There are several streams of targets entering and crossing the scene from all sites marked by yellow arrows. We see that not each pair of cameras shares a common area both can observe. Therefore, only the pairs (1, 2), (2, 4), (3, 4), (3, 5), (5, 6) influence each other and have to consider the configuration of the other camera.

Again, we apply Q-learning to this scenario but with a small adaptation: We do not use relative actions, e.g., increase or decrease, but absolute ones, i.e., each of the defined states can be accessed from every other state, leading to 72 actions. Since this does not require to learn sequences of action to reach one alignment from the previous, the discount factor is set to $\gamma = 0$.

Fig. 17. *SCN 3*. A top-down view on a larger smart camera network with six cameras. The position of the cameras are depicted by the black dots. Each of them has a colored circle around it marking the area that can be potentially observed by it and a red shape that shows an exemplary observed area. The gray squares show the position of pillars limiting the potentially observable areas. The yellow arrows mark entry points of observation targets and their movement direction.

The results are shown in Figure 18. The single-agent learning case in Figure 18(a) is clearly worse than the other variants only reaching a reward of 2 on average, since the cameras do not coordinate well. We see that the algorithm gets stuck in local optima when ϵ drops low at around 250K steps, since single configurations are valued the highest due to the lack of knowledge of the influencing systems. The other two variants show the dynamic integration in the state space. In Figure 18(b), we see how the system performs when 3K samples are used for the detection; in Figure 18(c), 10K samples are used. The graphs are nearly identical, and both reach an optimal behavior that gives an average reward of above 3.

Finally, in Figure 19, we see the cooperation network that evolved after 200K steps in one example run. Each arrow means that the configuration parts of the camera that are pointed at are integrated in the state space of the first camera. We see that the structure that formed reflects the structure of influences described before.

Concluding the results, we have seen how single learners compete with *a priori* adapted systems and a stagewise integration of influencing configurations in the state space of an RL algorithm. We have seen that the single learner cannot compete on the same level as influence-respecting variants. The influence-detecting approaches have shown similar results, which makes them useful options for the design of systems.

(a) The cameras are controlled by single agent Q-learner.

(b) The results if 3k samples are used to determine the candidates for an integration.

(c) The results if 10k samples are used to determine the candidates for an integration.

Fig. 18. The figure shows the average reward over 10 runs in the scenario SCN 3. The cameras are each controlled by a Q-learning algorithm with an ϵ -greedy action selection, where the ϵ decreases. It uses single learners or the integration of configurations at runtime.

Fig. 19. A snapshot of the structure of the learning network in an exemplary run on SCN 3 after 200K steps.

7 CONCLUSION

Self-adaptive and self-organizing systems typically act in a shared environment with other systems of the same kind. This results in the need to consider the current strategies followed by other systems within their own decision and learning process. We call such effects between actions and their effect on the achieved utility “mutual influences.”

In this article, we presented a novel methodology to identify and consider these mutual influences by extending and refining prior work and investigating novel solutions to consider them in runtime self-adaptation. The detection is realized with dependency measures that have different characteristics and are consequently compared regarding their applicability in various settings. The adaption to the influences is realized using the example of Q-learning. The techniques are evaluated on two elementary use cases, the collaborative box manipulation and the two-man saw, and a smart camera application with multiple scenarios.

In the remaining paragraphs, we discuss the scalability of the approach and give an outlook on future work.

7.1 Considerations about the Scalability of the Approach

While the scalability is not the main focus of this work, we want to give an impression of this by considering and discussing two potential bottlenecks of the approach and afterwards give practical advice on how to handle it in situations where they might become an issue.

The first one is the computational overhead: Besides the calculations for the learning algorithm, this leads to $n(n-1)k$ calculations of the dependency measures, where n is the number of sub-systems and k is the number of configuration components on each agent, meaning the overhead will grow quadratic with the number of agents. However, this number will reduce vastly after the first components have been found influencing and are integrated into the learning algorithm, since many possible paths are not considered any longer. For an experiment in the SC3 scenario with $n = 6$ and $k = 3$, this leads to a maximum of 90 calculations for each estimation of the influence.

The experiments have been run sequentially on an Intel i5-2500K@3.3GHz and the approximate runtime for 150K steps (which is approximately the time until an optimal solution has been found) is about 1.9h, including the time for the simulation itself and the learning algorithm. Given the remark before, it is clear that the average runtime of the influence estimation is much higher initially than in the end.

The second potential bottleneck is the distribution of the realizations of the configuration components. Depending on the system architecture, it is possible that each subsystem has to send its status to every other subsystem. Denoting the number of agents as n and the number of components as k , this leads to $(n - 1)k$ values sent by each agent or $n(n - 1)k$ values sent in the whole system in each step. These values can have different size, depending on the nature of the configuration space. Common value types are Boolean, integer, and floating point numbers. Using the values of scenario SC3 with $n = 6$ and $k = 3$, this would result in 15 values sent in each step by each agent.

While we believe that these requirements regarding the computational power and network capacity can be satisfied for a variety of systems that can benefit from our approach, it is clear that some systems have limitations that might prevent the application of the influence detection. In this case, it is obviously possible to tweak the communication complexity by using algorithms that are more efficient given the system's architecture or to rely on less computationally intensive dependency measures. Besides such implementation details, we want to discuss two basic ideas that can lead to a better scalability in general. These are the sequential analysis of systems and the reduction of network load.

Sequential Analysis of Systems. In the conducted experiments, each system measures the influence of each other system on itself. This can lead to a high system load and possibly to a backlog in the detection of influences, especially in large-scale systems.

To avoid such situations, it is possible to analyze each other system in a sequential order instead of simultaneously (possibly with multiple measures), i.e., a system creates an order and starts analyzing the influence from the first system in the sequence. As soon as it is done analyzing the first system, it can switch to the next and so on. This order can be random or based on the physical or virtual distance, for instance. This reduces the computational peak load heavily.

Reduction of Network Load. In the conducted experiments, we assumed that the influence originates from an arbitrary node in the overall system, i.e., each of the systems distributes its configuration via the communication network to the other systems. Depending on the connectivity of the systems, this can lead to a high network load affecting the communication necessary to operate the systems, e.g., the results in a wireless sensor network cannot be forwarded to the sink.

Similar to the reduction of computational resources, this can be avoided by analyzing the systems sequentially. Another effective way that can be used for some system types is to send the reward to the other systems instead of the configuration. This switches the places where the influences are calculated, i.e., instead of system A receiving the configurations of system B and calculating the influence of system B on itself, it sends the rewards of itself to system B , which in turn calculates the influence. Since the configuration will most often create more traffic than just the reward, this can reduce the computational overhead, but this approach is limited to systems that are trusted and are willing to share their resources.

7.2 Future Work

In the following, two directions for future work are outlined. They aim mainly at further automating the detection process by an automatic combination of the dependency measures and extending the applicability to systems that are not covered by the current system model.

Automated Weighting of Dependency Measures. In this work, we used multiple dependency measures for the detection. This is due to the different detection rates for the various classes of systems. An interesting approach to extend the approach shown here is to combine the dependency measures to get a more reliable detection. This can be achieved by relying on simple techniques such as majority voting, but a useful approach could be to consider ensemble learning techniques [61]. Especially promising are methods where the combination of the learners is done by determining the weights for the individual learners on a case-by-case basis regarding the structure of the system. The main challenge here lies in the engineering of a feedback signal or a training set that allows the combinator to find optimal weights for each situation.

Detecting Influences in Systems without a Local Reward Signal. In this article, the system model relies on a local reward that is a feedback signal to the subsystems regarding their decisions. If there is no obvious choice for such a reward signal, it can be useful to have an influence-detection mechanism that is able to work without such. Therefore, an interesting approach for future work is to change the system model to a type where only the configuration and a current state of the subsystems is necessary. The dependency can then be measured between one system's configuration and another system's situation. However, this introduces the problem that the situation is very likely not a one-dimensional variable but has several dimensions, e.g., the position is composed of an x - and y -coordinate. Therefore, it will be necessary to use dimension-reduction techniques on the situation of a system or to measure the dependency between two groups of random variables, e.g., with a canonical correlation analysis [62].

Besides these two concrete directions, it is promising to apply the influence detection to further real-world systems, especially with different learning algorithms and influence structures.

ACKNOWLEDGMENTS

We would like to thank our project partners Bernhard Sick, Christian Gruhl, Henner Heck (Uni Kassel, IES group) and Arno Wacker (Universität der Bundeswehr München) for the fruitful discussions in the project.

REFERENCES

- [1] J. O. Kephart and D. M. Chess. 2003. The vision of autonomic computing. *Computer* 36, 1 (Jan. 2003), 41–50. DOI: <http://dx.doi.org/10.1109/MC.2003.1160055>
- [2] Christian Müller-Schloer and Sven Tomforde. 2018. *Organic Computing—Technical Systems for Survival in the Real World*. Springer International Publishing. DOI: http://dx.doi.org/10.1007/978-3-319-68477-2_7
- [3] David Tennenhouse. 2000. Proactive computing. *CACM* 43, 5 (May 2000), 43–50. DOI: <http://dx.doi.org/10.1145/332833.332837>
- [4] Michael Wooldridge. *An Introduction to Multiagent Systems*. John Wiley & Sons.
- [5] S. Tomforde, J. Hähner, and B. Sick. 2014. Interwoven systems. *Inform.-Spekt.* 37, 5 (2014), 483–487. DOI: <http://dx.doi.org/10.1007/s00287-014-0827-z>
- [6] Serge Kernbach, Thomas Schmickl, and Jonathan Timmis. Collective adaptive systems: Challenges beyond evolvability. *Arxiv abs/1108.5643* ([n.d.]).
- [7] Danny Weyns, Bradley Schmerl, Vincenzo Grassi, Sam Malek, Raffaella Mirandola, Christian Prehofer, Jochen Wuttke, Jesper Andersson, Holger Giese, and Karl M. Göschka. 2013. On patterns for decentralized control in self-adaptive systems. In *Software Engineering for Self-Adaptive Systems II*. Springer, 76–107.
- [8] Sven Tomforde and Christian Müller-Schloer. 2014. Incremental design of adaptive systems. *Journal of Ambient Intelligence and Smart Environments* 6, 2 (2014), 179–198. DOI: <http://dx.doi.org/10.3233/AIS-140252>
- [9] Christian Krupitzer, Felix Maximilian Roth, Sebastian VanSyckel, Gregor Schiele, and Christian Becker. 2015. A survey on engineering approaches for self-adaptive systems. *Pervas. Mob. Comput.* 17 (2015), 184–206.
- [10] M. W. Maier. 1998. Architecting principles for systems-of-systems. *Syst. Eng.* 1, 4 (1998), 267–284.
- [11] Ada Diaconescu, Sylvain Frey, Christian Müller-Schloer, Jeremy Pitt, and Sven Tomforde. 2016. Goal-oriented holonics for complex system (self-)integration: Concepts and case studies. In *Proceedings of the 10th IEEE International Conference on Self-Adaptive and Self-Organising Systems*. IEEE, 100–109.

- [12] S. Tomforde, J. Hähner, H. Seebach, W. Reif, B. Sick, A. Wacker, and I. Scholtes. Engineering and mastering interwoven systems. In *Proceedings of the 27th International Conference on Architecture of Computing Systems, Workshop (ARCS'14)*. 1–8. Retrieved from: <http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=6775093>.
- [13] Sven Tomforde, Stefan Rudolph, Kirstie L. Bellman, and Rolf P. Würtz. An organic computing perspective on self-improving system interweaving at runtime. In *Proceedings of the IEEE International Conference on Autonomic Computing (ICAC'16)*. 276–284.
- [14] Stefan Rudolph, Sarah Edenhofer, Sven Tomforde, and Jörg Hähner. Reinforcement learning for coverage optimization through PTZ camera alignment in highly dynamic environments. In *Proceedings of the International Conference on Distributed Smart Cameras (ICDSC'14)*. 19:1–19:6.
- [15] Claudio Piciarelli, Lukas Esterle, Asif Khan, Bernhard Rinner, and Gian Luca Foresti. 2016. Dynamic reconfiguration in camera networks: A short survey. *IEEE Trans. Circ. Syst. Vid. Technol.* 26, 5 (2016), 965–977.
- [16] Lukas Esterle. 2017. Centralised, decentralised, and self-organised coverage maximisation in smart camera networks. In *Proceedings of the IEEE 11th International Conference on Self-Adaptive and Self-Organizing Systems (SASO'17)*. IEEE, 1–10.
- [17] Kirstie L. Bellman, Jean Botev, Ada Diaconescu, Lukas Esterle, Christian Gruhl, Christopher Landauer, Peter R. Lewis, Anthony Stein, Sven Tomforde, and Rolf P. Würtz. 2018. Self-improving system integration—Status and challenges after five years of SISSY. In *Proceedings of the IEEE 3rd International Workshops on Foundations and Applications of Self* Systems (FAS*W'18)*. 160–167. DOI: <http://dx.doi.org/10.1109/FAS-W.2018.00042>
- [18] Kirstie L. Bellman, Christian Gruhl, Christopher Landauer, and Sven Tomforde. 2019. Self-improving system integration—On a definition and characteristics of the challenges. In *Proceedings of the IEEE 4th International Workshops on Foundations and Applications of Self* Systems (FAS*W'19)*. 1–6.
- [19] Stefan Rudolph, Rainer Hihn, Sven Tomforde, and Jörg Hähner. Comparison of dependency measures for the detection of mutual influences in organic computing systems. In *Proceedings of the 29th International Conference on Architecture of Computing Systems (ARCS'16)*. 334–347. DOI: http://dx.doi.org/10.1007/978-3-319-30695-7_25
- [20] Stefan Rudolph, Sven Tomforde, and Jörg Hähner. A mutual influence-based learning algorithm. In *Proceedings of the 8th International Conference on Agents and Artificial Intelligence (ICAART'16), Volume 1*. 181–189.
- [21] Stefan Rudolph, Sven Tomforde, Bernhard Sick, and Jörg Hähner. A mutual influence detection algorithm for systems with local performance measurement. In *Proceedings of the IEEE 9th International Conference on Self-Adaptive and Self-Organizing Systems*. 144–149.
- [22] S. Tomforde, B. Hurling, and J. Hähner. 2010. Dynamic control of mobile ad-hoc networks—Network protocol parameter adaptation using organic network control. In *Proceedings of the 7th International Conference on Informatics in Control, Automation, and Robotics (ICINCO'10)*. INSTICC, Setubal, 28–35.
- [23] S. Tomforde, M. Steffen, J. Hähner, and C. Müller-Schloer. 2009. Towards an organic network control system. In *Proceedings of the 6th International Conference on Autonomic and Trusted Computing (ATC'09)*. Springer Verlag, 2–16.
- [24] Richard S. Sutton and Andrew G. Barto. *Introduction to Reinforcement Learning* (1st ed.). The MIT Press, Cambridge, MA.
- [25] A. van Lamsweerde. Goal-oriented requirements engineering: A guided tour. In *Proceedings of the 5th IEEE International Symposium on Requirements Engineering*. 249–262. DOI: <http://dx.doi.org/10.1109/ISRE.2001.948567>
- [26] Axel van Lamsweerde. *Requirements Engineering: From System Goals to UML Models to Software Specifications* (1st ed.). Wiley Publishing.
- [27] Erik M. Fredericks, Byron DeVries, and Betty H. C. Cheng. Towards run-time adaptation of test cases for self-adaptive systems in the face of uncertainty. In *Proceedings of the 9th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS'14)*. ACM, New York, NY, 17–26. DOI: <http://dx.doi.org/10.1145/2593929.2593937>
- [28] S. Tomforde and J. Hähner. 2011. *Biologically Inspired Networking and Sensing: Algorithms and Architectures*. IGI, Chapter: Organic Network Control—Turning Standard Protocols into Evolving Systems, 11–35.
- [29] L. Busoniu, R. Babuska, and B. De Schutter. 2008. A comprehensive survey of multiagent reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 38, 2 (2008), 156–172. DOI: <http://dx.doi.org/10.1109/TSMCC.2007.913919>
- [30] Marco Wiering and Martijn van Otterlo. *Reinforcement Learning: State-of-the-Art*. Springer Publishing Company, Incorporated.
- [31] Christopher J. C. H. Watkins and Peter Dayan. Technical note Q-learning. 8 ([n.d.]), 279–292.
- [32] S. Rudolph, S. Tomforde, B. Sick, H. Heck, A. Wacker, and J. Hähner. An online influence detection algorithm for organic computing systems. In *Proceedings of the 28th International Conference on Architecture of Computing Systems (ARCS'15)*. 1–8.
- [33] David Keil and Dina Q. Goldin. Modeling indirect interaction in open computational systems. In *Proceedings of the 12th IEEE International Workshops on Enabling Technologies (WETICE'03)*. 371–376.

- [34] Robert Logie, Jon G. Hall, and Kevin G. Waugh. Towards mining for influence in a multi agent environment. In *Proceedings of the IADIS European Conference on Data Mining*, Ajith Abraham (Ed.). IADIS, 97–101.
- [35] Robert Logie, Jon G. Hall, and Kevin G. Waugh. 2010. Investigating agent influence and nested other-agent behaviour. *Int. J. Adv. Intell. Syst.* 2, 4 (2010).
- [36] Jan M. Broersen. CTL.STIT: Enhancing ATL to express important multi-agent system verification properties. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'10)*, Volume 1–3. 683–690.
- [37] Peter Stone and Manuela Veloso. Multiagent systems: A survey from a machine learning perspective. 8, 3 ([n.d.]), 345–383.
- [38] Jelle R. Kok, Matthijs T. J. Spaan, and Nikos Vlassis. 2003. Multi-robot decision making using coordination graphs. In *Proceedings of the International Conference on Advanced Robotics (ICAR'03)*, A. T. de Almeida and U. Nunes (Eds.). 1124–1129.
- [39] Jelle R. Kok, Pieter Jan 't Hoen, Bram Bakker, and Nikos Vlassis. 2005. Utile coordination: Learning interdependencies among cooperative agents. In *Proceedings of the IEEE Symposium on Computational Intelligence and Games (CIG'05)*. 29–36.
- [40] Yann-Michaël De Hauwere, Peter Vrancx, and Ann Nowé. 2009. Learning what to observe in multi-agent systems. In *Proceedings of the 20th Belgian-Netherlands Conference on Artificial Intelligence*. 83–90.
- [41] Yann-Michael De Hauwere, Peter Vrancx, and Ann Nowé. 2010. Learning multi-agent state space representations. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 715–722.
- [42] Yann-Michaël De Hauwere, Peter Vrancx, and Ann Nowé. 2011. Solving delayed coordination problems in MAS. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'11)*, Volume 3. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1115–1116. Retrieved from: <http://dl.acm.org/citation.cfm?id=2034396.2034445>.
- [43] Marc Lanctot, Vinicius Zambaldi, Audrunas Gruslys, Angeliki Lazaridou, Karl Tuyls, Julien Pérolat, David Silver, and Thore Graepel. 2017. A unified game-theoretic approach to multiagent reinforcement learning. In *Proceedings of the Conference on Advances in Neural Information Processing Systems*. 4190–4203.
- [44] Hanchuan Peng, Fuhui Long, and Chris Ding. Feature selection based on mutual information: Criteria of max-dependency, max-relevance, and min-redundancy. 27, 8 ([n.d.]), 1226–1238.
- [45] Amir Massoud Farahmand, Mohammad Ghavamzadeh, Csaba Szepesvári, and Shie Mannor. Regularized policy iteration. In *Proceedings of the 22nd Conference on Advances in Neural Information Processing Systems*. 441–448.
- [46] J. Zico Kolter and Andrew Y. Ng. 2009. Regularization and feature selection in least-squares temporal difference learning. In *Proceedings of the 26th International Conference on Machine Learning (ICML'09)*. ACM, New York, NY, 521–528. DOI : <http://dx.doi.org/10.1145/1553374.1553442>
- [47] De-Rong Liu, Hong-Liang Li, and Ding Wang. Feature selection and feature learning for high-dimensional batch reinforcement learning: A survey. 12, 3 ([n.d.]), 229–242.
- [48] Jérémy Boes and Frédéric Migeon. 2017. Self-organizing multi-agent systems for the control of complex systems. *J. Syst. Softw.* 134 (2017), 12–28.
- [49] Stefan Rudolph, Rainer Hihn, Sven Tomforde, and Jörg Hähner. Towards discovering delayed mutual influences in organic computing systems. In *Proceedings of the 30th GI/ITG International Conference on Architecture of Computing Systems (ARCS'17)*. 39–46.
- [50] Karl Pearson. Note on regression and inheritance in the case of two parents. 58, 347–352 ([n.d.]), 240–242.
- [51] Maurice G. Kendall. 1938. A new measure of rank correlation. *Biometrika* 30, 1/2 (1938), 81–93.
- [52] Gábor J. Székely, Maria L. Rizzo, and Nail K. Bakirov. 2007. Measuring and testing dependence by correlation of distances. *The Annals of Statistics*. 35, 6 (2007), 2769–2794. DOI : <http://dx.doi.org/10.1214/009053607000000505>
- [53] Claude Shannon and Warren Weaver. *The Mathematical Theory of Communication*. University of Illinois Press.
- [54] Alexander Kraskov, Harald Stögbauer, and Peter Grassberger. 2004. Estimating mutual information. *Phys. Rev. E* 69 (2004), 066138. Issue 6. DOI : <http://dx.doi.org/10.1103/PhysRevE.69.066138>
- [55] David N. Reshef, Yakir A. Reshef, Hilary K. Finucane, Sharon R. Grossman, Gilean McVean, Peter J. Turnbaugh, Eric S. Lander, Michael Mitzenmacher, and Pardis C. Sabeti. 2011. Detecting novel associations in large data sets. *Science* 334, 6062 (2011), 1518–1524. DOI : <http://dx.doi.org/10.1126/science.1205438>
- [56] Sean Luke, Claudio Cioffi-Revilla, Liviu Panait, Keith Sullivan, and Gabriel Balan. 2005. Mason: A multiagent simulation environment. *Simul.: Trans. Soc. Model. Simul. Int.* 82, 7 (2005), 517–527.
- [57] Matthew E. Taylor and Peter Stone. 2009. Transfer learning for reinforcement learning domains: A survey. *J. Machine Learn. Res.* 10 (July 2009), 1633–1685.
- [58] S. W. Wilson. Classifier fitness based on accuracy. 3, 2 ([n.d.]), 149–175.

- [59] Stewart W. Wilson. 2000. Get real! XCS with continuous-valued inputs. In *Learning Classifier Systems*, Pier Luca Lanzi, Wolfgang Stolzmann, and Stewart W. Wilson (Eds.). Springer Berlin, 209–219.
- [60] Xianneng Li and Guangfei Yang. 2016. Transferable XCS. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO'16)*. ACM, New York, NY, 453–460. DOI: <http://dx.doi.org/10.1145/2908812.2908859>
- [61] Bartosz Krawczyk, Leandro L. Minku, João Gama, Jerzy Stefanowski, and Michał Woźniak. 2017. Ensemble learning for data stream analysis: A survey. *Inform. Fus.* 37 (2017), 132–156.
- [62] Wolfgang Härdle and Léopold Simar. 2007. *Applied Multivariate Statistical Analysis*. Vol. 22007. Springer.

Received September 2018; revised May 2019; accepted July 2019