Jane Cleland-Huang Ankit Agrawal janehuang@nd.edu aagrawa2@nd.edu Computer Science and Engineering University of Notre Dame Notre Dame, IN, USA Michael Vierhauser michael.vierhauser@jku.at LIT Secure and Correct Systems Lab Johannes Kepler University Linz Linz, Austria

Michael Murphy Mike Prieto murphym18@gmail.com mprieto2@nd.edu Computer Science and Engineering University of Notre Dame Notre Dame, IN, USA

ABSTRACT

The MAPE-K feedback loop has been established as the primary reference model for self-adaptive and autonomous systems in domains such as autonomous driving, robotics, and Cyber-Physical Systems. At the same time, the Human Machine Teaming (HMT) paradigm is designed to promote partnerships between humans and autonomous machines. It goes far beyond the degree of collaboration expected in human-on-the-loop and human-in-the-loop systems and emphasizes interactions, partnership, and teamwork between humans and machines. However, while MAPE-K enables fully autonomous behavior, it does not explicitly address the interactions between humans and machines as intended by HMT. In this paper, we present the MAPE-K_{HMT} framework which augments the traditional MAPE-K loop with support for HMT. We identify critical human-machine teaming factors and describe the infrastructure needed across the various phases of the MAPE-K loop in order to effectively support HMT. This includes runtime models that are constructed and populated dynamically across monitoring, analysis, planning, and execution phases to support human-machine partnerships. We illustrate MAPE-K_{HMT} using examples from an autonomous multi-UAV emergency response system, and present guidelines for integrating HMT into MAPE-K.

CCS CONCEPTS

• Human-centered computing → Collaborative interaction; HCI theory, concepts and models.

KEYWORDS

Self-Adaptive Systems, Human-Machine Teaming, Autonomous Systems, MAPE-K

ACM Reference Format:

Jane Cleland-Huang, Ankit Agrawal, Michael Vierhauser, Michael Murphy, and Mike Prieto. 2022. Extending MAPE-K to support Human-Machine Teaming. In 17th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS '22), May 18–23, 2022, PITTSBURGH, PA, USA. ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/ 3524844.3528054

SEAMS '22, May 18-23, 2022, PITTSBURGH, PA, USA

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 978-1-4503-9305-8/22/05...\$15.00 https://doi.org/10.1145/3524844.3528054

1 INTRODUCTION

The MAPE-K feedback loop [7, 43], is a well-adopted reference model for managing and controlling autonomous and self-adaptive systems, and its use has enabled significant advances in autonomous systems over the past decades, for example, in areas such as autonomous driving and traffic management [34], Unmanned Aerial Vehicles [55, 59], Smart Home and IoT applications [6, 36], and assistive robots [39]. Furthermore, rapid advancements in Artificial Intelligence (AI), supported by frameworks such as MAPE-K, have shifted the focus from traditional human-directed robots to fully autonomous ones that do not require explicit human control. These systems, which are commonly developed as "Human-onthe-Loop" (HotL) [30] systems, differ from "Human-in-the-Loop" (HitL) systems in several important ways. In HitL systems, humans make decisions at key points of the system's execution; while HotL systems take full advantage of machine autonomy to perform tasks independently, efficiently, and quickly.

However, given technological advances in autonomic computing, a more advanced form of collaboration, referred to as Human Machine Teaming (HMT) has emerged [47]. HMT emphasizes interactions, partnership, and teamwork between humans and machines. It capitalizes upon the respective strengths of both the human and the machine, whilst compensating for each of their potential limitations [47, 54]. According to McDermott et al., effective HMT requires transparency of the machine's progress and plans, as well as augmented cognition to empower the machine to adapt as needed, keep the human partner aware of critical problems, and allow both human and machine to explore the shared solution space. Coordination between humans and machines establishes shared knowledge and trust between both human and machine partners and empowers a human partner to direct the machine's behavior when desired [56, 57]. Finally, the machine's self-adaptation capabilities are extended to allow it to dynamically configure and reconfigure its interactions with human partners throughout the mission.

HMT systems incorporate aspects of both Cyber-Physical Systems (CPS) [74], and Socio-Technical Systems (STS) [25, 81]. In the context of HMT, systems are still expected to operate fully autonomously, with all the capabilities that MAPE-K is designed to support. In fact, not only are the machines capable of performing their tasks autonomously, but they are perceived as true *partners* and not just "tools" in achieving mission goals. To make this transition from the HotL paradigm to HMT, humans and machines must interact more closely – not in a way that reduces or curtails the autonomous behavior of the machine, but in one that leverages that behavior to create meaningful partnerships.

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 the author(s) 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.

SEAMS '22, May 18-23, 2022, PITTSBURGH, PA, USA



Figure 1: An overview of MAPE-K_{HMT} showing machine activities (outer circle), and human activities (inner circle). Phases are mapped to Situational Awareness levels (L1-L3). Examples of runtime models are shown for each phase.

The primary goal of any feedback control system is to remove humans from the loop, and therefore MAPE-K focuses upon autonomous decision-making and self-adaptation without emphasizing the human aspects of a CPS. This was reflected in the results of a recent systematic literature review [10] which reported that runtime models associated with self-adaptation primarily target the architecture, structure of the system, and/or its goals, but hardly incorporate any human-related factors or activities, such as user interaction or situational awareness [26]. This creates a gap between the existing MAPE-K framework and the capabilities needed by an autonomous system to fully interact with human partners in an HTM environment. Kephart proposed bridging this gap through creating highly interactive "dialogs" between humans and machines; however, his examples are all drawn from information systems and not real-time robotic environments [42].

We propose a solution for bridging this gap in more diverse system environments through MAPE- K_{HMT} , which enhances the fundamental MAPE-K loop with runtime support for HMT, and aligns *teaming factors*, identified from the HMT literature [57] with the different phases of the MAPE-K loop. We then present a set of runtime models for use in MAPE- K_{HMT} and describe how they enable support for bidirectional human-machine interactivity and decision-making. Our approach is illustrated using a set of worked examples taken from a multi-agent system of autonomous Unmanned Aerial Vehicles (UAV) for supporting emergency response missions [21, 24]. Based on these examples, we have derived a lightweight process and recommendations for integrating HMT into MAPE-K.

The remainder of this paper is structured as follows. Section 2 introduces $MAPE-K_{HMT}$ and its extensions to the MAPE-K loop.

Section 3 introduces our case study system, while Section 4 presents six examples of HMT-related runtime models and their integration into $MAPE-K_{HMT}$. Section 5 then describes a process for utilizing $MAPE-K_{HMT}$ in a MAPE-K system. Finally, Sections 6 to 8 present threats to validity, related work, and conclusions.

2 THE $MAPE-K_{HMT}$ **LOOP**

With *MAPE-K_{HMT}* we aim to leverage the benefits of active human engagement while preserving the autonomous behavior of the selfadaptive system. *MAPE-K_{HMT}* follows the same general structure of MAPE-K, however, as illustrated in Fig. 1, each of the phases is augmented with additional capabilities targeted to support HMT. The original MAPE-K loop [43] consists of four pivotal phases: *Monitoring* in which information is collected from the environment, *Analysis* where data is analyzed to determine if adaptations need to be performed, *Planning* where corresponding actions and adaptations are planned, and finally *Execution* in which the proposed plans are enacted. Additionally, the "K" stands for an underlying knowledge base, accessible to all other parts of the MAPE loop, and often supported by runtime models [35].

In the context of human-machine teaming, a number of capabilities are required to ensure both successful human-machine interaction and cooperation [57]. In Table 1, we summarize some of the primary dependencies. *Transparency* is supported by "Observability" (TF1) of the autonomous partner's task progress and "Predictability" (TF2) of its future plans. *Cognition* is augmented through "Directing Attention" (TF3) to critical problems, for example by raising meaningful alerts to increase situational awareness. "Solution Exploration" (TF4) and "Adaptability" (TF5) imbue both

| | Teaming Factor | | Definition | Key example of HMT realization in MAPE-K _{HMT} |
|---------------------|----------------|----------------|-------------------------------------|---|
| ransparency | TF1 | Observability | Visibility of the task progress of | [M ⁺] Data related to status & task progress is collected from sensors, |
| | | | the automated partner and human | UI inputs, and software probes. |
| | | | actions. | [A ⁺] Progress is dynamically visualize for humans. |
| | TF2 | Predictability | Transparency of future intentions, | [A ⁺] Diverse views generated to provide situational awareness of the |
| | | | states, and activities. | machine's activities & intent. [P ⁺] Analysis & planning results |
| | | | | communicated to the human partner. |
| Augmented Cognition | TF3 | Directing | Keeping human partners aware of | [M ⁺] Relevant data is collected & analyzed. |
| | | Attention | critical problems through warnings, | [A ⁺] Based on factors such as state or user role, runtime user-alerts |
| | | | recommendations, and indicators. | are raised, prioritized, and displayed. |
| | TF4 | Solution | Ability for both partners to | Human partners leverage diverse interactive views and simulations |
| | | Exploration | leverage multiple views, | to explore the solution space $[A^+]$ and to make plans $[P^+]$. |
| | | | knowledge, & candidate solutions. | |
| | TF5 | Adaptability | Ability to address potentially | [P ⁺] Human engages in adaptation planning alongside machine. |
| | | | unexpected evolving, dynamic | [E ⁺] Machine self-adapts its interactions with humans according to |
| | | | situations through adaptation. | human behavior and context. |
| Coordination | TF6 | Directabilty | Humans ability to direct/redirect | [A ⁺] Humans draw on their observations & analysis to |
| | | | an automated partner's resources, | $[\mathbf{P}^+]$ intervene in the machine's plans. $[\mathbf{E}^+]$ The machine requests |
| | | | activities, and priorities. | support from the human when its confidence is low. |
| | TF7 | Calibrated | Trustworthiness indicators of | $[M^+]$ Based on collected data $[M^+]$ the machine $[P^+]$ computes |
| | | Trust | machine's ability to make correct | reliability of its own autonomous decisions and actions, and displays |
| | | | decisions in current context. | appropriate trust-related indicators at runtime [E ⁺]. |
| | TF8 | Common | Shared beliefs, assumptions, and | [M ⁺] The Human understands data collection, analysis & use, [A ⁺], |
| | | Ground | intents across human and | & the machine's capabilities $[P^+/E^+]$. |
| | | | automated partners. | $[E^+]$ The machine responds to human directions as expected. |

Table 1: Key impacts of HMT factors upon phases of the MAPE-K loop. The factors (TF1-TF8) are proposed by McDermott *et al.* [56] and augmented to reflect the bidirectional partnerships proposed in this paper.

the machine and humans with the knowledge and capabilities they need to make and enact decisions. Finally, *Coordination* is supported through "Directability" (TF6), "Calibrated Trust" (TF7), and establishing "Common Ground" (TF8) to enable informed, trustworthy, and trusted partnerships. Achieving each of these capabilities requires consideration for how humans and machines can work in teams to accomplish their goals, then establishing runtime models to collect, aggregate, and visualize information that supports HMT. As a result, human and machine collaborators can engage in meaningful interactions. These eight capabilities represent Teaming Factors that are applied across the MAPE-K loop to augment each phase as M⁺, A⁺, P⁺, E⁺, and K⁺ respectively. In the following sections, we explain how each phase is augmented to support HMT.

2.1 The Monitoring Phase (M)

In the MAPE-K loop, monitoring is primarily concerned with *collecting data* from the self-adaptive system and the environment in which it operates. Hardware sensors provide raw data such as temperature, distance to potential obstacles, video streams, or GPS locations, while software probes provide data from the running system such as its resource usage, response times, and currently executing tasks [37]. This data is collected, persisted, and used in runtime models to guide subsequent analysis and self-adaptation decisions [40, 79].

M⁺: To forge effective human-machine partnerships, HMT environments must provide bidirectional situational awareness; therefore, the machine not only collects data about its own state but

also collects human-initiated data reflecting human goals, directives, workload, and response times [26]. This data is collected via Graphical UIs (GUIs) and via hardware interfaces such as radio controllers (RC), audio devices, pointing devices, and even eye-trackers [17, 32, 62, 63] or brain interfaces [33, 61]. Furthermore, given the dissonance between the way humans and machines perceive the world [85], supplementary data is needed to support human awareness and interactions and must be collected from all three sources (i.e., machine, environment, and human inputs). The data collection process is therefore expanded accordingly and used across subsequent analysis, planning, and execution phases to support the HMT goals of transparency, cognition, and coordination, with an emphasis on the teaming factors of *observability, directing attention, calibrated trust*, and *common ground*.

2.2 The Analysis Phase (A)

The MAPE-K analysis phase is concerned with determining whether adaptation actions are required, based on the current and predicted state of the system, the environment it is operating in, and its defined goals, safety constraints, and quality of service specifications. Automated analysis enables timely and fast reactions to changes in the environment and emergent situations.

A⁺: HMT recognizes the value of augmenting machine analysis capabilities with human perspectives and inputs [85]. This requires information to be exchanged between machines and humans through human-facing interfaces that provide insights into the autonomous machine's intent, performance, plans, and reasoning processes' [38, 71]. Human analysis often relies upon a combination of real world observations as well as software supported interactions. In direct real-world analysis, for example, the human uses their physical senses to observe and analyze the machine's behavior, and to intervene quickly and directly through either a hardware or software interface, potentially causing a temporary interrupt to the MAPE-K loop. In contrast, in software-supported analysis, the human analyzes a representation of the machine's behavior and/or state through a GUI. As humans need time to understand and analyze the information and formulate decisions, static snapshots become stale within seconds - even milliseconds. This leads to the conclusion that GUI support for HMT must include dynamic runtime views that reflect current system states and historical information about past actions, rather than just static snapshots of the system. They are therefore often built on top of existing runtime models with continually updated views. Historical views that enable the human to explore why a machine made a past decision (e.g., [67, 70]), or what it plans to do next, are also needed. The MAPE-K_{HMT} analysis has a broad impact across almost all of the teaming factors as analysis is a precursor to human engagement.

2.3 The Planning Phase (P)

In the planning phase the machine plans self-adaptation actions such as switching states to perform different tasks, reconfiguring existing features, activating or deactivating sensors, or modifying polling frequencies to preserve power or to collect additional information about the system or its environment.

P⁺: HMT introduces two additional considerations to the planning phase. First, the human leverages their observations of the machine, its operating environment, interactions with other team members, and profound experience as a "human knowledge base" to engage directly in the planning process. Humans might reconfigure the mission's goals or plans, or temporarily intervene in the operation of the machine, for example by assuming manual control of a task that the machine is not able to perform autonomously or when the machine malfunctions. However, this introduces a potential tug-ofwar that can occur when humans and machines create competing plans [22]. This was catastrophically illustrated in the crash of Lion Air Flight 610 and Ethiopian Airlines Flight 302 in which the MCAS (Maneuvering Characteristics Augmentation System) incorrectly perceived the angle of attack to exceed predefined limits and therefore pushed the nose of the plane down, whilst pilots struggled to push it back up [31]. The system was not designed to detect and mitigate this type of tug-of-war, and ultimately the machine "won", causing the planes to crash. Achieving effective coordination between humans and machines is a challenging problem which we discuss further in Section 4.3.

HMT has a second major impact on the planning phase, as the machine may make additional self-adaptation plans targeted at enhancing the human's interactive experience. For example, the machine might self-adapt its internal alert system to adjust the type and frequency of human alerts if the system perceives that human response time is starting to lag [3].

2.4 The Execution Phase (E)

During the execution phase, the previously generated plan or adaptation strategy is executed on the physical machine or device.

 E^+ : In an HMT setting, both the machine and the human partners enact plans – sometimes closely coordinating their work whilst at other times working more independently on tasks that each partner is best suited to perform.

2.5 The Knowledge Base (K) + Runtime Models

MAPE-K employs diverse runtime models that represent the structure, behavior, and/or goals of a system at runtime. These models are pivotal for guiding autonomous adaptation decisions [8]. In a recent systematic literature review Bencomo *et al.*, [10] reported that *Models@run.time* have been used in numerous ways – for example to depict the current state of the system [28] and its behavioral dynamics specifying exactly what the system is able to do from its current state [29], model system goals [18, 69, 77] and functional and non-functional requirements [20, 41], and to depict product variability [82]. They provide a bidirectional reflection layer, such that changes in the runtime model trigger changes in the goals, structure, and/or behavior of the underlying system, whilst changes in the system are reflected in the model. As such, runtime models support system autonomy, imbuing the system with the ability to sense, analyze, predict, and to make independent decisions.

K⁺: In *MAPE-K_{HMT}* runtime models must also support the HMT goals of transparency, augmented cognition, and coordination between human and machine partners. Depending on the type of system, application domain, or the type of missions that are executed, this requires different types of runtime models that are explicitly designed to provide information to the user or relay critical information from the user to the system so that informed adaptation decisions can be made.

3 CASE PROJECT: A MULTI-UAV SYSTEM

Throughout this paper, we illustrate MAPE-K_{HMT} with examples drawn from our *Drone Response* system – a flexible and configurable framework for multi-UAV missions [21, 24].

3.1 The Drone Response Ecosystem

Drone Response is fully deployable with both physical and simulated UAVs [46, 65]. The *Drone Response* architecture includes diverse user interfaces, a Ground Control Station (GCS), and autonomy capabilities located onboard each UAV. The MAPE-K infrastructure is distributed across the *Drone Response* ecosystem as follows.

• UAV Onboard Pilot: The Onboard Pilot module acts as an application layer for the autopilot stack which in turn includes flight control software and hardware for executing plans. Its internal *State Machine* receives mission specifications and instantiates itself accordingly. It then uses its onboard state machine to progress through a series of tasks, with transitions triggered by events it detects by analyzing its own sensor data and/or messages received from the GCS or other UAVs. MAPE-K's monitoring, analysis, planning, and execution occur onboard each UAV with an emphasis on managing the tasks and adaptations of each individual UAV.

SEAMS '22, May 18-23, 2022, PITTSBURGH, PA, USA



Figure 2: Human-Machine Teaming is supported by GUIs, hardware UIs, and a set of closely integrated runtime models.

• *Ground Control Station:* The Control Station adheres to the microservices and publish-subscribe architectural style [60], with each microservice providing a specific capability such as airspace leasing, multi-UAV coordination, or safety checking mission specifications. During a mission, microservices receive messages from both the human operators and UAVs, construct runtime models, and leverage them for analysis and planning purposes. The knowledge base is thereby distributed across runtime models managed by various microservices and supported by a shared in-memory database. Status data (e.g., GPS location, battery, health) and task progress updates (e.g., current task, potential adaptations) sent to the GCS by both humans and UAVs support MAPE-K's mission-level monitoring, analysis, and planning.

• *Graphical and Hardware Interfaces: Drone Response* uses graphical and hardware user interfaces to enable human-machine interactions. Most of the interactions use GUIs; however, in case of emergency, or to temporarily assume control for tasks that the machine has not yet been trained to perform, humans can directly issue commands to UAVs via hand-held radio controllers. The GUI components are built on a centrally hosted web application and asynchronously send and receive status data and video streams over a mesh-radio via GCS's message broker and the onboard pilot modules [3]. Many of the *MAPE-K_{HMT}* runtime models are coupled with one or more GUIs that provide situational awareness to the human for planning and analysis purposes while also potentially monitoring aspects of human interaction behavior.

Within *Drone Response*, HMT is supported by diverse runtime models integrated across the system (cf. Fig. 2). For example, in order to generate meaningful explanations where multiple UAVs are involved, information about their current state is collected from each of the UAVs. Humans interact directly with the UAVs using physical devices (1) and indirectly via a GUI (2). HMT Interaction Models (4) that process data and control human-machine interactions in the GUI (3), receive input from underlying runtime-models (5), as well as from other sources such as UAVs (7) and humans via the GUIs. Communication between these major components is achieved using the MQTT message broker with service level agreements guaranteeing fast response times when needed (6).

3.2 A Motivating Scenario

Drone Response's UAVs are designed to leverage computer vision and work alongside humans in emergency response missions such as search-and-detect, surveillance, and rescue scenarios. As depicted in Fig. 3, the UAVs first use their onboard computer vision to search for, and potentially detect the victim. They then notify the human responders about the victim's location. Drone Response then has the option of delivering a flotation device, as shown in Fig. 4; however, this requires very close human-machine teaming. For example, if a physical rescue from a boat is imminent, then simultaneously dropping a flotation device in the location of the victim could hinder the rescue operation and introduce a potential safety concern. This example provides a clear differentiation between HiTL, HoTL, and HMT paradigms. In a HiTL environment, the UAV would wait to be dispatched by a human and subsequently request input regarding specific tasks to be performed. In a HoTL environment, the UAV might decide to initiate the delivery of the flotation device independently; however, the human could decide to intervene and cancel the operation. In contrast, while human intervention is definitely part of HMT, the HMT scenario calls for joint decision-making, whereby the human and machine leverage a shared set of beliefs and understanding of the mission to evaluate their capabilities and opportunities, weigh the options with respect



Figure 3: Drone Response streams video of victims once detected by the onboard Computer Vision.



Figure 4: Proof-of-concept Drone Response delivery of a defibrillator. (See https://youtu.be/AleVXc3QWIk)

SEAMS '22, May 18-23, 2022, PITTSBURGH, PA, USA

Cleland-Huang, Agrawal, Vierhauser, Murphy, Prieto

to global mission goals, and ultimately make a joint decision for the good of the mission.

All eight HMT teaming factors (cf. Table 1) play a vital role in fostering collaborations between emergency responders and the UAVs. Because humans need to develop trust in the autonomy capabilities of the UAV during the mission, they require observability of their tasks, and expect predictability of their decisions and actions. The system must raise appropriate alerts in order to help humans to maintain situational awareness. In the spirit of the HMT partnership, the human can directly contribute to the search - for example, providing information for use by the computer vision algorithm or hints to the route planner, that the victim is wearing a red sweater and that sightings have been reported in a certain part of the search area. To increase the likelihood of a successful search. the UAVs and humans must establish a shared conceptual model of the mission and humans must be able to explore the solution space through user-facing perspectives of the current mission, including the location, status, and progress of each UAV.

4 APPLYING MAPE-K_{HMT}

In this section, we describe six different runtime models from our *Drone Response* project that support HTM within the MAPE-K loop. Each of the runtime models contributes towards one or more HMT goals of transparency, augmented cognition, or human-machine coordination. Furthermore, each model not only reflects the runtime behavior or structure of the machine, but also provides direct or indirect support for a human-facing UI component in order to support human-machine collaboration. In the following sections we describe models that are particularly pertinent to transparency, cognition, and coordination respectively.

4.1 Runtime Models for Transparency

HMT's transparency goal focuses on observability and predictability and therefore aligns closely with Endsley's Situational Awareness goals of observing and understanding [26, 27]. Observability means that humans are aware of what their autonomous partner is doing, including its goals, current tasks, future intentions, status, progress, ability to adapt to changing contexts, rationales for adaptations, and challenges or constraints that impact its ability to solve the current problem. Additionally, predictability helps to remove potential surprises introduced by the machine's decisions, provides insights into uncertainties [9] and ways in which reliability of the autonomous partner changes over time, under what circumstances it changes, and how its decisions are made. Given the importance of the transparency goal, and the many facets of observability and predictability, systems will likely have several associated runtime models. These models tend to use runtime data about the machine's progress and its environment collected using runtime monitors. Based on this, visualization and interactivity support is provided for the analysis phase. Drone Response primarily supports transparency through the use of several map-based views that depict the current location and status of each UAV, as well as task-based views such as the one presented in the following example.

- Task-Centric Model: In a multi-agent system, the human needs to understand exactly what each individual agent is currently doing



Figure 5: This 'Multi-Agent Tracking' view displays the task progress of all active UAVs. It is managed by a dedicated microservice, which continually aggregates states and transition paths for all active UAVs and uses each UAV's uniquely colored token to mark their current state.

and its progress towards the overall mission plan. The *Drone Response* task-centric view generates a global state transition model combining the states and transitions from all active UAVs' onboard state machines, and then tracks and displays the current task performed by each UAV in the merged state transition model. Progress is then visualized in a UI showing the current task of each UAV as a colored token assigned to a specific node. This is illustrated in Fig. 5, which shows that the Green (G) delivery UAV is *on standby*, Red (R) and Orange (O) UAVs are *searching*, the Purple (P) UAV is performing *surveillance*, and the Blue (B) UAV has *detected a victim*.

4.2 Runtime Models that Augment Cognition

The cognition goal extends far beyond basic transparency and is intended to help humans understand emergent problems and their



Figure 6: The 'Alert Triage' model stores user-created and/or machine-learned prioritization rules, and triages the display of alerts for each view (e.g., Map view on the right).

causes, provide insights into the decisions and actions taken by the autonomous partner, and allow the user to explore different perspectives and solutions [50]. $MAPE-K_{HMT}$ runtime models support cognition goals through enabling the right information to be presented or available to the human at the right time, without overloading their cognitive abilities [3, 45, 83]. The two examples we present here focus on *triaging runtime alerts* and *explaining autonomous actions* of the machine by generating human-readable explanations of autonomous behavior; however, our previously presented task-centric model (cf. Fig. 5) provides an additional example, as its interactive GUI also provides interactive and more detailed views of individual UAV's tasks and progress.

- Alert Prioritization Model: The Drone Response alert prioritization model is designed to avoid the situational awareness "design demon" of information overload [26], and is built upon a formal meta-model for human-UAV collaborations [4]. It focuses particularly on the HMT goal of augmenting cognition with an emphasis on directing human attention to important messages. In Fig. 6 alert rules and priorities are shown on the left. They are initially provided as default values by human stakeholders but can be dynamically adapted at runtime, by both the human and the machine. They specify essential alerts which must always be displayed (\checkmark), and prioritized alerts (1-5) which will only be displayed if they don't cause the maximum threshold to be exceeded. The triage part of the model (e.g., Map View), is dynamically maintained by the system at runtime and is responsible for managing alerts in each active GUI view. It is notified whenever an alert is generated by a runtime model hosted on the UAV or on a GCS microservice. It is also notified by the GUI server whenever a new GUI is activated or deactivated. The alert prioritization model thus builds upon services already available in the basic MAPE-K loop, by collecting, aggregating, and processing the data they produce, in order to support the HMT-focused capability of triaging alerts.

Type H/M Explanation Template Ext М UAV-{id/color} identified {Event} in the environment. Therefore, adapting {Action - internal changes} to {Rationale} Ext UAV-{id/color} identified {Event} in the environment. Η Therefore, need {Desired Changes} to {Rationale} UAV-{id/color} observed {Event}. Therefore, {Action Int Μ internal changes} to {Rationale} UAV-{id/color} observed {Event} due to {cause}. Η Int Therefore, need {Desired Changes} to {Rationale}

(a) Explanation templates for internally and externally triggered adaptation events initiated by either the human (H) or machine (M).



(b) An example explanation generated by the runtime model.

Figure 7: The Autonomy Explanation Model generates an explanation for all major adaptation decisions.

- Adaptation Explanation Model: Finally, the explanation model generates explanations for UAV autonomous decisions so that humans can gain insights into the machine's reasoning and assess the appropriateness of individual adaptations [3, 48]. These insights potentially strengthen the human's trust and confidence in its machine partner. The predefined explanation templates shown in Table 7a are used to dynamically generate human-readable textual explanations for all adaptations performed by a UAV. Whenever a UAV self-adapts, it collects three types of information. These are explanation snippets describing relevant external events (e.g., "misty weather conditions" or "victim detected"), the UAV's response to the events (e.g., "reduced altitude by 8 m" or "switched to tracking mode"), and finally, the rationale behind those actions (e.g., "limited visibility", or "high confidence in victim sighting"). An example of a weather-related adaptation explanation is shown in Figure 7b. Upon receipt of the adaptation message, the runtime model selects the appropriate template and generates the explanation by filling in the missing parts with the data provided by the UAV [3]. This model, therefore, utilizes the outputs of existing runtime adaptation models to provide explanations that are critical for HMT.

4.3 Runtime Models for Coordination

Coordinating the actions of both humans and machines represents a challenging problem [51] that is exacerbated by the differing cadences of human and machine response times. Successful coordination depends on many factors, including a shared conceptual

SEAMS '22, May 18-23, 2022, PITTSBURGH, PA, USA



Figure 8: The UAV's Computer Vision component considers confidence, reliability, and calibrated trust levels, to determine when to request help from the human partner [1].

model of the operating environment, respect for each partner's capabilities, and well-calibrated bidirectional trust. Humans trust the machine to operate autonomously when capable and to request help when needed. Conversely, machines accept interventions from humans and expect feedback when requested, if the human is available. In this context, we present three different runtime models covering the cases of machine-initiated and human-initiated coordination, as well as the particularly challenging case that occurs when humans and machines generate conflicting plans of action.

- Machine Initiated Coordination: In Drone Response much of the autonomous behavior of the UAV is supported by its onboard CV capabilities (cf. Fig. 8) [1]. For example, when searching for a drowning victim, the UAV uses CV to continuously analyze the video stream and detect objects classified as "person". For each detected object, the CV module generates two scores. The confidence score represents the probability that the object is correctly classified, while the reliability score accounts for any uncertainty arising from image noise or mismatch between the current context and the training data [64]. Based on learned threshold values, the UAV uses these scores to decide whether it can autonomously decide on its actions (e.g., track the object vs. continue to search) or should request help from the human. Both the CV model and the related humanmachine coordination mechanism are supported by runtime models which take a coordination specification as input (cf. Algorithm 1), instantiate a simple state machine (managed by a microservice on the GCS), and use the existing messaging system to choreograph human and machine tasks. In closely related work Li et al., also proposed a form of choreography that provides users with advanced knowledge of tasks they would be requested to participate in [51]. - Human Initiated Coordination Models: In order for the human to engage in meaningful interactions with the machine, they must have high degrees of situational awareness and a clear understanding of what they can and cannot do at any point during a mission. Furthermore, the different cadences of human and machine response times within MAPE-K_{HMT} create a significant coordination challenge resulting in problems such as the human making plans based on stale data or intervening after the machine has already completed an action. An example of a relevant runtime model determines currently available human interaction options and dynamically adapts each active GUI to activate and deactivate widgets (e.g., icons, buttons, menu options) according to the ways that humans can feasibly

| Algorithm 1: Human-Machine coordinated decision- making addressing Computer Vision reliability problems | | | | | |
|---|--|--|--|--|--|
| if Object detected at low reliability then | | | | | |
| UAV raises alert and requests help from human; | | | | | |
| if Human is available and responsive then Human evaluates video stream, makes decision, and selects CONFIRM or REJECT option; | | | | | |
| if Human confirms victim sighting then UI Server sends CONFIRMATION message to UAV; | | | | | |
| else | | | | | |
| Human refutes sighting; | | | | | |
| UI Server sends REFUTATION message to UAV; | | | | | |
| end | | | | | |
| if no response from human within waiting_period thenNO RESPONSE message sent to UAV; | | | | | |
| 'human failure to respond' event is logged; | | | | | |
| Responsibility reverts to UAV; | | | | | |
| end | | | | | |
| end | | | | | |

interact with the machine given the current state of the mission. For example, if the UAV is currently in active search mode, it is reasonable for the human to request a view of the UAV's annotated video stream; however, this option should not be available if the UAV is in RTL mode with cameras turned off to preserve limited power. A suitable affordance (e.g., a button) should be activated and deactivated accordingly). Tracking these currently available human actions is handled by a dedicated runtime model.

- Mitigating Tug-of-War Scenarios: Finally, a significant challenge in HMT is reconciling potentially conflicting actions taken by the human and machine, a problem that is again exacerbated by the different operating speeds of the machine and human. As the machine has the advantage of faster cadence, it will often win any disagreements with potentially devastating results, as illustrated by the case of the Lion Air Crash (cf., Section 2.3). Similar scenarios play out across other domains. For example, a UAV may place itself close to the river to collect better imagery of the riverbank. On a sunny day, reflections from the light on the water could impact the UAV's sensors causing sudden altitude fluctuations and triggering a land-in-place failsafe mechanism to activate. An alerted human might quickly intervene to prevent the UAV from landing in the river directing it to ascend to a safe altitude; however, once the UAV's autonomy kicks back into gear, the whole cycle could repeat itself - causing a tug-of-war with respect to the ideal altitude placement. HMT systems must allow humans and/or machines to detect and break interwoven cycles of human and machine actions that are indicative of a tug-of-war. As a simple solution, detected cycles could be broken by curtailing the UAV's autonomy until reestablished by the human.

4.4 Integrating the RunTime Data Probes

While we have presented these six example models as independent entities, many of them share common data inputs (e.g., data read from sensors and software probes), or rely on data/outputs from

another runtime model. For example, in order to generate meaningful explanations where multiple UAVs are involved, information is collected from the individual UAV onboard runtime models. The Autonomy Explanation Model relies upon "adaptation notification events" provided by each individual UAV, which in turn employs their own state machine. This information is then aggregated and post-processed, so that the Alert Prioritization Model (cf. Fig. 6) can select and display critical alerts from individual UAVs or global alerts from the system. While both the knowledge base, and (parts) of the runtime models are distributed across the respective agents they share data as needed. This means that the UAV retains full autonomy when a person is detected, and adaptation decisions are made (onboard), but at the same time raises alert events which support decision-making of the human partner. This dependence on accurate and appropriate information between the different models calls for thorough design and planning to ensure that the right information is available at the right time in a potentially resourceconstrained operating environment.

5 ANALYSIS AND TAKE-AWAYS

In this paper, we have proposed augmenting the MAPE-K loop with support for HMT. *MAPE-K_{HMT}* assumes that both humans and machines are capable of autonomous behavior and decision-making, and that mission goals are achieved jointly through an interactive partnership. Based on our own experiences in applying *MAPE-K_{HMT}* to the *Drone Response* system, we lay out an initial process for designing and deploying *MAPE-K_{HMT}*.

- (1) Stakeholder Identification: HMT systems inherently involve humans, and therefore, it is important to identify CRACK (Collaborative, Representative, Authorized, Committed, Knowledgeable) [12] stakeholders serving as direct users and domain experts. Engaging with stakeholders, who will become the human partners, helps to uncover interactions and expectations users have on the system, for example, by carefully exploring the humanmachine interactions related to mission-related scenarios [76].
- (2) Elicitation of HMT requirements: Once stakeholders are identified, requirements must be elicited. As a starting point, Mc-Dermott et al. [57] has described a detailed elicitation process, supported by a list of key questions associated with each of the HMT factors [56]. For example, to understand "predictability" requirements, analysts must discover (a) automation goals, abilities, and limitations, (b) how the human partner's goals and priorities are tracked, (c) reliability of different automated tasks within different contexts, and (d) the types of changes that are expected to occur and trigger subsequent adaptations.
- (3) Requirements Analysis and Specification: The elicited requirements are subsequently analyzed to negotiate and reconcile trade-offs, and to identify and specify requirements that support human-machine interactions [68]. While our process is agnostic to specific techniques, the domains in which MAPE-K operates generally dictate that the requirements process includes a rigorous safety analysis (e.g., [23, 49, 80]), and determines that requirements should be specified sufficiently formally to capture timing and other performance constraints and/or to establish

formal goal models. In many cases, the requirements specification (e.g., Goal Models or state transition diagrams) provides the foundation for the respective runtime models [5, 10, 15].

- (4) Design and Integration of HMT runtime models: MAPE- K_{HMT} emphasizes the importance of HMT-related runtime models. We, therefore, start the design process by taking an inventory of existing runtime models, identifying gaps where HMT requirements are not adequately supported by existing models, designing new runtime models as needed, and then finally composing models into workflows to service each HMT requirement. This involves assessing who (machine or human-role), when, and where each model will be used and updated, and what data sources are required as inputs (e.g., probes, message subscriptions, or human-initiated data and events). Furthermore, as actions are performed at different speeds, the required refresh frequencies must be determined for each constituent element of each runtime model, and a system-wide plan established so that each collected data attribute satisfies refresh frequencies of all relevant models.
- (5) User Interface Design: User interfaces need to be designed to provide interactive support for humans. Depending on the data that is displayed, or the input that is expected from the human, this may include GUIs, or hardware interfaces such as Joysticks or radio controllers. Mission- and safety-critical information needs to be provided to the human in a timely manner without creating cognitive overload, and must not be "hidden" in submenus or views that require multiple steps to access [26].
- (6) Implementation, Testing, and Deployment: The final steps involve implementing the system, verifying that required runtime models, UIs, and supporting features are implemented as intended, and finally validating that the deployed system satisfies its stated requirements and supports the desired HMT. Further discussion of these steps is outside the scope of this paper.

6 THREATS TO VALIDITY

Our work is subject to three primary threats to validity. First, all of the runtime models described in this paper are designed to support HMT in our own Drone Response system, representing a multi-agent, multi-human system operating in a mid-level safety-critical domain [53, 78]. As types of interactions are influenced by the operating domain, the human-machine partnerships in Drone Response may differ significantly from other application domains - for example, those with a single operator or a single machine, or operating in a highly safety-critical domain. On the other hand, instead of proposing a specific set of HMT-related runtime models, MAPE-KHMT provides a simple process for identifying, developing, and integrating context-specific models in support of transparency, augmented cognition, and coordination goals, which are known to be applicable across diverse operating environments [47, 56, 57]. Further, our process could be easily extended - for example, by integrating a more formal safety analysis for more critical domains. Our future work will apply this process to more varied systems.

Second, while our runtime models have been prototyped in various forms, only some have been integrated into the full version of *Drone Response*. As such, this paper represents a vision, supported by concrete examples. Future work will focus on concrete implementation details. For example, while we have discussed the creation of shared data probes that service the refresh frequencies of all runtime models, and have developed monitoring capabilities to achieve this, this aspect of the work needs further investigation.

Finally, while *Drone Response* has been deployed in the physical world, formal evaluations of HMT user interfaces have relied upon user studies conducted in *Drone Response*'s simulator (e.g., [2, 3, 79]). However, the *Drone Response* GUI is identical for both physical and simulated UAVs, and our past experiences have shown that findings from these user studies have been effective when deployed in physical field tests. Nevertheless, further experiments targeted specifically at HMT need to be conducted in real-world settings where humans collaborate with UAVs under far more noisy, volatile, and potentially stressful conditions.

7 RELATED WORK

A significant body of work, primarily in the HCI community has focused on designing UIs to support situational awareness of Cyber-Physical Systems (e.g., [2, 4, 11, 66, 72, 75]). The focus has primarily been on enabling users to perceive, understand, and make effective decisions [26, 27]. While this related work supports HMT goals, it puts little emphasis on integration with underlying runtime models which are needed in order to deliver accurate, timely, and often aggregated data for use in the UIs. Kephart advocated for increased interactivity between humans and users in adaptive decision-making [42]; however, their examples were all taken from domains in which humans had plenty of time to consider their decisions. Integrating HMT into the MAPE-K loop for real-time robotics systems introduces additional, and very challenging, timing constraints.

In the HMT domain, researchers have explored many facets of human-machine teaming. For example, Klein *et al.* [44] identified challenges associated with achieving shared goals, preventing breakdowns in team coordination, and fostering communication and collaboration. Furthermore, Schmid *et al.* [73] studied ways to adjust system automation in complex, safety-critical environments in order to better support human operators. While their work has significant relevance to *MAPE-K*_{HMT}, it perceives humans as "operators" rather than true partners.

Calhoun *et al.* [16] proposed a flexible architecture that allows the degree of automation to vary according to the human's current engagement and workload. These goals are reflected in our discussion on adaptation in *MAPE-K_{HMT}*, which embraces the notion of adapting for improved human collaboration and performance. To increase context awareness in order to better engage humans in the decision-making process, Li *et al.* [51] proposed a formal framework based on probabilistic reasoning to determine when advanced notifications are useful for humans interacting with selfadaptive systems. Their work specifically addresses the cadence problem in which humans may be required to respond quickly but require "thinking time". However, their evaluation was conducted in a robotics goods delivery domain, which has a lower decision cadence than a multi-UAV system.

Finally, several papers have explored self-adaptation and/or human interactions in the UAV domain. Vierhauser *et al.* [80] identified human-UAV interaction hazards and potential mitigations, whilst Miller *et al.* [58] explored different techniques by which users could interact with multiple UAVs. However, these papers focus on Human-on-the-loop rather than HMT. Other researchers have proposed UAV-related self-adaptation frameworks. For example, Braberman *et al.* [13, 14] presented a MAPE-K reference architecture for unmanned aerial vehicles, while Yu *et al.* [84] proposed a self-adaptive framework for UAV forensics. Neither of these explored the HMT aspects of adaptation. Finally, in more closely related work Lim *et al.* [52] explored ways to adapt human-robot interactions in a multi-UAV system, with a focus on modulating automation support according to the cognitive states of the human operator. This creates a subtle, but important difference from our *MAPE-K_{HMT}* goal, as it centers primarily around the needs of the operator rather than optimizing teamwork goals.

8 CONCLUSION

This paper has described the techniques and process we have used to augment MAPE-K in order to address the three HMT goals of transparency, augmented cognition, and coordination. We have described how MAPE-K provides a meaningful self-adaptation framework in which humans and machines collaborate together, and have mapped HMT factors to the MAPE-K phases and then used a series of examples to illustrate how carefully designed runtime models enable meaningful support for human-machine partnerships. Rather than diminishing the autonomy of machines, HMT draws upon the autonomous abilities of both humans and machines to deliver an even stronger solution realized through developing meaningful teamwork.

In conducting this work we have identified two key challenges. The first stems from the very different cadences at which humans and machines operate. This creates the potential for humans to make decisions based upon stale data, and for directives interjected into the machine's plans to inadequately reflect the current state of the system. Existing HMT solutions, such as turn-taking [19], are only effective in scenarios which can tolerate slower response times. In Drone Response we partially address this challenge by dynamically adapting the UIs to only reflect currently available human interaction options, and by implementing a reliable messaging system which ensures that messages are only sent to the UAV when it is in a state to handle them; however, the problem is a complex one which warrants further exploration. The second challenge relates to designing and supporting the integrated MAPE-K_{HMT} environment. As different runtime models are required for effective human-machine teaming, and these need to be identified, designed, integrated, deployed, and effectively maintained at runtime.

Even beyond these challenges, much additional work is needed in order to realize the vision of trusted, coordinated, and accountable teams of humans and machines. Our ongoing work will therefore investigate solutions for addressing these challenges, create a fully integrated MAPE- K_{HMT} environment for evaluation purposes, conduct field-tests with physical UAVs, and explore the deployment of MAPE- K_{HMT} across a much broader set of different domains.

ACKNOWLEDGMENT

The work in this paper has been funded by the US National Science Foundation (NSF) under Grant# CNS-1931962 and also by the Linz Institute of Technology (LIT-2019-7-INC-316).

SEAMS '22, May 18-23, 2022, PITTSBURGH, PA, USA

REFERENCES

- [1] Sophia J. Abraham, Zachariah Carmichael, Sreya Banerjee, Rosaura G. VidalMata, Ankit Agrawal, Md Nafee Al Islam, Walter J. Scheirer, and Jane Cleland-Huang. 2021. Adaptive Autonomy in Human-on-the-Loop Vision-Based Robotics Systems. In Proc. of 1st Workshop on AI Engineering – Software Engineering for AI. IEEE.
- [2] Ankit Agrawal, Sophia J. Abraham, Benjamin Burger, Chichi Christine, Luke Fraser, John M. Hoeksema, Sarah Hwang, Elizabeth Travnik, Shreya Kumar, Walter J. Scheirer, Jane Cleland-Huang, Michael Vierhauser, Ryan Bauer, and Steve Cox. 2020. The Next Generation of Human-Drone Partnerships: Co-Designing an Emergency Response System. In Proc. of CHI Conference on Human Factors in Computing Systems. ACM, New York, 1–13. https://doi.org/10.1145/3313831. 3376825
- [3] Ankit Agrawal and Jane Cleland-Huang. 2021. Explaining Autonomous Decisions in Swarms of Human-on-the-Loop Small Unmanned Aerial Systems. In Proc. of Ninth AAAI Conference on Human Computation and Crowdsourcing, Vol. 9. 15–26.
- [4] Ankit Agrawal, Jan-Philipp Steghöfer, and Jane Cleland-Huang. 2020. Model-Driven Requirements for Humans-on-the-Loop Multi-UAV Missions. In Proc. of Workshop on Model-Driven Requirements Engineering. IEEE, 1–10.
- [5] Amal Ahmed Anda and Daniel Amyot. 2019. Arithmetic semantics of feature and goal models for adaptive cyber-physical systems. In Proc. of 27th International Requirements Engineering Conference. IEEE, 245-256.
- [6] Paolo Arcaini, Raffaela Mirandola, Elvinia Riccobene, Patrizia Scandurra, Alberto Arrigoni, Daniele Bosc, Federico Modica, and Rita Pedercini. 2020. Smart home platform supporting decentralized adaptive automation control. In Proc. of 35th Annual ACM Symposium on Applied Computing. ACM, New York, 1893–1900.
- [7] Paolo Arcaini, Elvinia Riccobene, and Patrizia Scandurra. 2015. Modeling and Analyzing MAPE-K Feedback Loops for Self-Adaptation. In Proc. of 10th International Symposium on Software Engineering for Adaptive and Self-Managing Systems. IEEE, 13–23. https://doi.org/10.1109/SEAMS.2015.10
- [8] Uwe Aßmann, Sebastian Götz, Jean-Marc Jézéquel, Brice Morin, and Mario Trapp. 2014. A Reference Architecture and Roadmap for Models@run.time Systems. Springer International Publishing, Cham, 1–18. https://doi.org/10.1007/978-3-319-08915-7_1
- [9] Nelly Bencomo and Amel Belaggoun. 2014. A world full of surprises: bayesian theory of surprise to quantify degrees of uncertainty. In Proc. of the 36th International Conference on Software Engineering, Companion Proceedings. ACM, 460–463. https://doi.org/10.1145/2591062.2591118
- [10] Nelly Bencomo, Sebastian Götz, and Hui Song. 2019. Models@run.time: a guided tour of the state of the art and research challenges. Software and Systems Model 18, 5 (2019), 3049–3082. https://doi.org/10.1007/s10270-018-00712-x
- [11] Francesco N Biondi, Monika Lohani, Rachel Hopman, Sydney Mills, Joel M Cooper, and David L Strayer. 2018. 80 MPH and out-of-the-loop: Effects of real-world semi-automated driving on driver workload and arousal. In Proc. of Human Factors and Ergonomics Society Annual Meeting, Vol. 62. SAGE Publications, 1878–1882.
- [12] Barry W. Boehm and Richard Turner. 2004. Balancing Agility and Discipline: Evaluating and Integrating Agile and Plan-Driven Methods. In Proc. of the 26th International Conference on Software Engineering. IEEE Computer Society, 718– 719. https://doi.org/10.1109/ICSE.2004.1317503
- [13] Victor Braberman, Nicolas D'Ippolito, Jeff Kramer, Daniel Sykes, and Sebastian Uchitel. 2015. Morph: A reference architecture for configuration and behaviour self-adaptation. In Proc. of 1st International Workshop on Control Theory for Software Engineering. ACM, New York, 9–16.
- [14] Victor Braberman, Nicolas D'Ippolito, Jeff Kramer, Daniel Sykes, and Sebastian Uchitel. 2017. An Extended Description of MORPH: A Reference Architecture for Configuration and Behaviour Self-Adaptation. In Software Engineering for Self-Adaptive Systems III. Assurances. Springer International Publishing, Cham, 377–408.
- [15] Jennifer Brings, Marian Daun, Thorsten Weyer, and Klaus Pohl. 2020. Goal-based configuration analysis for networks of collaborative cyber-physical systems. In *Proc. of 35th Annual ACM Symposium on Applied Computing*. ACM, New York, 1387–1396.
- [16] Gloria L. Calhoun, Heath A. Ruff, Kyle J. Behymer, and Elizabeth Marie Mersch. 2017. Operator-Autonomy Teaming Interfaces to Support Multi-Unmanned Vehicle Missions. In Advances in Human Factors in Robots and Unmanned Systems. Springer, 113–126.
- [17] Radu Calinescu, Naif Alasmari, and Mario Gleirscher. 2021. Maintaining driver attentiveness in shared-control autonomous driving. In Proc. of 2021 International Symposium on Software Engineering for Adaptive and Self-Managing Systems. IEEE, 90–96. https://doi.org/10.1109/seams51251.2021.00021
- [18] Lorena Castañeda, Norha M. Villegas, and Hausi A. Müller. 2014. Self-Adaptive Applications: On the Development of Personalized Web-Tasking Systems. In Proc. of 9th International Symposium on Software Engineering for Adaptive and Self-Managing Systems. ACM, New York, 49–54.
- [19] Crystal Chao and Andrea Thomaz. 2016. Timed Petri nets for fluent turn-taking over multimodal interaction resources in human-robot collaboration. *The International Journal of Robotics Research* 35, 11 (2016), 1330–1353.

- [20] Tao Chen and Rami Bahsoon. 2017. Self-Adaptive and Online QoS Modeling for Cloud-Based Software Services. *IEEE Transactions on Software Engineering* 43, 5 (2017), 453–475. https://doi.org/10.1109/TSE.2016.2608826
- [21] Jane Cleland-Huang, Ankit Agrawal, Md Nafee Al Islam, Eric Tsai, Maxime Van Speybroeck, and Michael Vierhauser. 2020. Requirements-driven configuration of emergency response missions with small aerial vehicles. In Proc. of 24th ACM International Systems and Software Product Line Conference. ACM, New York, 1–12. https://doi.org/10.1145/3382025.3414950
- [22] Rogério de Lemos. 2020. Human in the Loop: What is the Point of No Return?. In Proc. of IEEE/ACM 15th International Symposium on Software Engineering for Adaptive and Self-Managing Systems. ACM, New York, 165–166. https://doi.org/ 10.1145/3387939.3391597
- [23] Ewen Denney, Ganesh Pai, and Josef Pohl. 2012. Heterogeneous aviation safety cases: Integrating the formal and the non-formal. In Proc. of the 17th International Conference on Engineering of Complex Computer Systems. IEEE, 199–208.
- [24] Dronology. 2020. Research Incubator and Dataset. https://dronology.info. [Last accessed 01-01-2022].
- [25] F.E. Emery and E.L. Trist. 1960. Socio-technical systems. In In: Churchman, C.W., Verhulst, M. (Eds.), Management Science Models and Techniques, Vol. 9. Pergamon, 83–97.
- [26] Mica R. Endsley. 2011. Designing for Situation Awareness: An Approach to User-Centered Design, Second Edition (2nd ed.). CRC Press, Inc., Boca Raton, FL, USA.
- [27] Mica R Endsley. 2017. Autonomous driving systems: A preliminary naturalistic study of the Tesla Model S. *Journal of Cognitive Engineering and Decision Making* 11, 3 (2017), 225–238.
- [28] Naeem Esfahani, Eric Yuan, Kyle R. Canavera, and Sam Malek. 2016. Inferring Software Component Interaction Dependencies for Adaptation Support. ACM Transactions on Autonomous and Adaptive Systems 10, 4 (2016), 26:1–26:32. https: //doi.org/10.1145/2856035
- [29] Antonio Filieri, Giordano Tamburrelli, and Carlo Ghezzi. 2016. Supporting Self-Adaptation via Quantitative Verification and Sensitivity Analysis at Run Time. *IEEE Transactions on Software Engineering* 42, 1 (2016), 75–99. https: //doi.org/10.1109/TSE.2015.2421318
- [30] Joel E Fischer, Chris Greenhalgh, Wenchao Jiang, Sarvapali D Ramchurn, Feng Wu, and Tom Rodden. 2021. In-the-loop or on-the-loop? Interactional arrangements to support team coordination with a planning agent. *Concurrency and Computation: Practice and Experience* 33, 8 (2021), e4082.
- [31] Flight Safety Foundation. 2019. Preliminary Report B737-800MAX. https:// flightsafety.org/preliminary-report-b737-800max-et-avj. [Last accessed 01-01-2022].
- [32] Lex Fridman, Bryan Reimer, Bruce Mehler, and William T Freeman. 2018. Cognitive load estimation in the wild. In Proc. of 2018 CHI Conference on Human Factors in Computing Systems. ACM, New York, 1–9.
- [33] Antoine Gaume, Gérard Dreyfus, and François-Benoît Vialatte. 2019. A cognitive brain-computer interface monitoring sustained attentional variations during a continuous task. *Cognitive Neurodynamics* 13, 3 (2019), 257–269.
- [34] Ilias Gerostathopoulos and Evangelos Pournaras. 2019. TRAPPed in traffic? A self-adaptive framework for decentralized traffic optimization. In Proc. of 14th International Symposium on Software Engineering for Adaptive and Self-Managing Systems. IEEE, 32–38.
- [35] Holger Giese, Nelly Bencomo, Liliana Pasquale, Andres Ramirez, Paola Inverardi, Sebastian Wätzoldt, and Siobhán Clarke. 2014. Living with Uncertainty in the Age of Runtime Models. 47–100. https://doi.org/10.1007/978-3-319-08915-7_3
- [36] Muhammad Usman Iftikhar, Gowri Sankar Ramachandran, Pablo Bollansée, Danny Weyns, and Danny Hughes. 2017. Deltaiot: A self-adaptive internet of things exemplar. In Proc. of 12th International Symposium on Software Engineering for Adaptive and Self-Managing Systems. IEEE, 76–82.
- [37] Didac Gil De La Iglesia and Danny Weyns. 2015. MAPE-K formal templates to rigorously design behaviors for self-adaptive systems. ACM Transactions on Autonomous and Adaptive Systems (TAAS) 10, 3 (2015), 1-31.
- [38] Chen J. Y., Procci K., Boyce M., Wright J.and Garcia A., and Barnes M. 2014. Situation Awareness–Based Agent Transparency. Report No. ARL-TR-6905. Aberdeen Proving Ground, MD: U.S. Army Research Laboratory. (2014).
- [39] Pooyan Jamshidi, Javier Cámara, Bradley Schmerl, Christian Käestner, and David Garlan. 2019. Machine learning meets quantitative planning: Enabling selfadaptation in autonomous robots. In Proc. of 14th International Symposium on Software Engineering for Adaptive and Self-Managing Systems. IEEE, 39–50.
- [40] Yuchen Jiang, Shen Yin, and Okyay Kaynak. 2018. Data-driven monitoring and safety control of industrial cyber-physical systems: Basics and beyond. *IEEE Access* 6 (2018), 47374–47384.
- [41] Donia Kateb, Nicola Zannone, Assaad Moawad, Patrice Caire, Grégory Nain, Tejeddine Mouelhi, and Yves Le Traon. 2015. Conviviality-Driven Access Control Policy. *RE Journal* 20 (11 2015). https://doi.org/10.1007/s00766-014-0204-0
- [42] Jeffrey Kephart. 2020. Keynote: Viewing Autonomic Computing through the Lens of Embodied Artificial Intelligence: A Self-Debate. In Proc. of the 15th International Symposium on Software Engineering for Adaptive and Self-Managing Systems. IEEE.
- [43] J.O. Kephart and D.M. Chess. 2003. The vision of autonomic computing. Computer 36, 1 (2003), 41–50. https://doi.org/10.1109/MC.2003.1160055

SEAMS '22, May 18-23, 2022, PITTSBURGH, PA, USA

- [44] Glen Klien, David D Woods, Jeffrey M Bradshaw, Robert R Hoffman, and Paul J Feltovich. 2004. Ten challenges for making automation a" team player" in joint human-agent activity. *IEEE Intelligent Systems* 19, 6 (2004), 91–95.
- [45] Torkel Klingberg. 2009. The overflowing brain: Information overload and the limits of working memory. Oxford University Press.
- [46] Nathan Koenig and Andrew Howard. 2004. Design and Use Paradigms for Gazebo, An Open-Source Multi-Robot Simulator. In Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems. IEEE, Sendai, Japan, 2149–2154.
- [47] Margarita Konaev and Husanjot Chahal. 2021. Building trust in humanmachine teams. https://www.brookings.edu/techstream/building-trust-inhuman-machine-teams. [Last accessed 01-01-2022].
- [48] Jeamin Koo, Jungsuk Kwac, Wendy Ju, Martin Steinert, Larry Leifer, and Clifford Nass. 2015. Why did my car just do that? Explaining semi-autonomous driving actions to improve driver understanding, trust, and performance. *International Journal on Interactive Design and Manufacturing (IJIDeM)* 9, 4 (2015), 269–275.
- [49] Nancy G. Leveson. 2011. The Use of Safety Cases in Certification and Regulation. Technical Report. MIT.
- [50] Nianyu Li, Sridhar Adepu, Eunsuk Kang, and David Garlan. 2020. Explanations for human-on-the-loop: A probabilistic model checking approach. In Proc. of IEEE/ACM 15th International Symposium on Software Engineering for Adaptive and Self-Managing Systems. ACM, New York, 181–187.
- [51] Nianyu Li, Javier Cámara, David Garlan, Bradley Schmerl, and Zhi Jin. 2021. Hey! Preparing Humans to do Tasks in Self-adaptive Systems. In Proc. of 2021 International Symposium on Software Engineering for Adaptive and Self-Managing Systems. IEEE, 48–58. https://doi.org/10.1109/SEAMS51251.2021.00017
- [52] Yixiang Lim, Nichakorn Pongsakornsathien, Alessandro Gardi, Roberto Sabatini, Trevor Kistan, Neta Ezer, and Daniel J. Bursch. 2021. Adaptive Human-Robot Interactions for Multiple Unmanned Aerial Vehicles. *Robotics* 10, 1 (2021). https: //doi.org/10.3390/robotics10010012
- [53] Robert Loh, Y. Bian, and Tim Roe. 2009. UAVs in civil airspace: Safety requirements. *IEEE Aerospace and Electronic Systems Magazine* 24, 1 (2009), 5–17. https://doi.org/10.1109/MAES.2009.4772749
- [54] Azad M. Madni and Carla C. Madni. 2018. Architectural Framework for Exploring Adaptive Human-Machine Teaming Options in Simulated Dynamic Environments. Systems 6, 4 (2018). https://doi.org/10.3390/systems6040044
- [55] Paulo Henrique Maia, Lucas Vieira, Matheus Chagas, Yijun Yu, Andrea Zisman, and Bashar Nuseibeh. 2019. Dragonfly: a tool for simulating self-adaptive drone behaviours. In Proc. of 14th International Symposium on Software Engineering for Adaptive and Self-Managing Systems. IEEE, 107–113.
- [56] Patricia McDermott, Cindy Dominguez, Nicholas Kasdaglis, Matthew Ryan, and Isabel Trahan. 2018. Human-Machine Teaming Systems Engineering Guide, Technical Report # MP180941. (2018).
- [57] Patricia L McDermott, Katherine E. Walker, Cynthia O Dominguez, Alex Nelson, and Nicholas Kasdaglis. 2018. Quenching the Thirst for Human-Machine Teaming Guidance: Helping Military Systems Acquisition Leverage Cognitive Engineering Research, The MITRE Corporation. https://www.mitre.org/publications/technical-papers/quenching-the-thirstfor-human-machine-teaming-guidance-helping. [Last accessed 01-01-2022].
- [58] Christopher A. Miller, Joseph Mueller, Christopher Geib, David LaVergne, Phillip Walker, and Joshua Hamell. 2019. User Interaction Approaches For Managing Multiple UAS In The National Airspace. In 2019 Integrated Communications, Navigation and Surveillance Conference (ICNS). IEEE, 1–16. https://doi.org/10. 1109/ICNSURV.2019.8735205
- [59] Gabriel Moreno, Cody Kinneer, Ashutosh Pandey, and David Garlan. 2019. DART-Sim: An exemplar for evaluation and comparison of self-adaptation approaches for smart cyber-physical systems. In Proc. of 14th International Symposium on Software Engineering for Adaptive and Self-Managing Systems. IEEE, 181–187.
- [60] Henry Muccini and Mahyar Tourchi Moghaddam. 2018. IOT Architectural Styles. In Proc. of 2018 European Conference on Software Architecture. Springer, 68–85.
- [61] Amin Nourmohammadi, Mohammad Jafari, and Thorsten O Zander. 2018. A survey on unmanned aerial vehicle remote control using brain-computer interface. IEEE Transactions on Human-Machine Systems 48, 4 (2018), 337–348.
- [62] Oskar Palinko, Andrew L Kun, Alexander Shyrokov, and Peter Heeman. 2010. Estimating cognitive load using remote eye tracking in a driving simulator. In Proc. of 2010 Symposium on Eye-Tracking Research & Applications. ACM, New York, 141–144.
- [63] Colin Paterson, Radu Calinescu, Suresh Manandhar, and Di Wang. 2019. Using Unstructured Data to Improve the Continuous Planning of Critical Processes Involving Humans. In Proc. of 14th International Symposium on Software Engineering for Adaptive and Self-Managing Systems. IEEE, 25–31. https: //doi.org/10.1109/SEAMS.2019.00013
- [64] Alexandre Pouget, Jan Drugowitsch, and Adam Kepecs. 2016. Confidence and certainty: distinct probabilistic quantities for different goals. *Nature neuroscience* 19, 3 (2016), 366.
- [65] PX4. 2022. jMAVSim. https://docs.px4.io/master/en/simulation/jmavsim.html. [Last accessed 01-01-2022].

- [66] Nicolas Regis, Frédéric Dehais, Emmanuel Rachelson, Charles Thooris, Sergio Pizziol, Mickaël Causse, and Catherine Tessier. 2014. Formal Detection of Attentional Tunneling in Human Operator-Automation Interactions. *IEEE Transactions* on Human-Machine Systems 44, 3 (2014), 326–336. https://doi.org/10.1109/THMS. 2014.2307258
- [67] Owen Reynolds, Antonio García-Domínguez, and Nelly Bencomo. 2020. Automated provenance graphs for models@run.time. In Proc. of the ACM/IEEE 23rd International Conference on Model Driven Engineering Languages and Systems. ACM, New York, 1–10. https://doi.org/10.1145/3417990.3419503
- [68] Günther Ruhe, Armin Eberlein, and Dietmar Pfahl. 2003. Trade-off analysis for requirements selection. International Journal of Software Engineering and Knowledge Engineering 13, 04 (2003), 345–366.
- [69] Luca Sabatucci and Massimo Cossentino. 2015. From Means-End Analysis to Proactive Means-End Reasoning. In Proc. 10th International Symposium on Software Engineering for Adaptive and Self-Managing Systems. IEEE, 2–12.
- [70] Lucas Sakizloglou, Sona Ghahremani, Matthias Barkowsky, and Holger Giese. 2021. Incremental execution of temporal graph queries over runtime models with history and its applications. *Software and Systems Modeling* (2021).
- [71] Lindsay Sanneman and Julie A. Shah. 2020. A Situation Awareness-Based Framework for Design and Evaluation of Explainable AI. In Explainable, Transparent Autonomous Agents and Multi-Agent Systems. Springer International Publishing, Cham, 94–110.
- [72] Filippo Santoni de Sio and Jeroen van den Hoven. 2018. Meaningful Human Control over Autonomous Systems: A Philosophical Account. Frontiers in Robotics and AI 5 (2018), 15. https://doi.org/10.3389/frobt.2018.00015
- [73] Daniela Schmid, Bernd Korn, and Neville Stanton. 2020. Evaluating the Reduced Flight Deck Crew Concept Using Cognitive Work Analysis and Social Network Analysis: Comparing Normal and Data-link Outage Scenarios. Cognition, Technology & Work 22 (02 2020). https://doi.org/10.1007/s10111-019-00548-5
- [74] J. Shi, J. Wan, H. Yan, and H. Suo. 2011. A survey of Cyber-Physical Systems. In Proc. of 2011 International Conference on Wireless Communications and Signal Processing. 1–6. https://doi.org/10.1109/WCSP.2011.6096958
- [75] Tim Claudius Stratmann and Susanne Boll. 2016. Demon Hunt The Role of Endsley's Demons of Situation Awareness in Maritime Accidents. In *Human-Centered* and Error-Resilient Systems Development. Springer International Publishing, Cham, 203–212.
- [76] Alistair G Sutcliffe, Neil AM Maiden, Shailey Minocha, and Darrel Manuel. 1998. Supporting scenario-based requirements engineering. *IEEE Transactions on Software Engineering* 24, 12 (1998), 1072–1088.
- [77] Aurélien Vialon, Kenji Tei, and Samir Aknine. 2017. Soft-Goal Approximation Context Awareness of Goal-Driven Self-Adaptive Systems. In Proc. of 2017 IEEE International Conference on Autonomic Computing. IEEE, 233–238. https://doi. org/10.1109/ICAC.2017.25
- [78] Michael Vierhauser, Sean Bayley, Jane Wyngaard, Wandi Xiong, Jinghui Cheng, Joshua Huseman, Robyn R. Lutz, and Jane Cleland-Huang. 2021. Interlocking Safety Cases for Unmanned Autonomous Systems in Shared Airspaces. *IEEE Transactions on Software Engineering* 47, 5 (2021), 899–918. https://doi.org/10. 1109/TSE.2019.2907595
- [79] Michael Vierhauser, Jane Cleland-Huang, Sean Bayley, Thomas Krismayer, Rick Rabiser, and Pau Grünbacher. 2018. Monitoring CPS at runtime-A case study in the UAV domain. In Proc. of 44th Euromicro Conference on Software Engineering and Advanced Applications. IEEE, 73–80.
- [80] Michael Vierhauser, Md Nafee Al Islam, Ankit Agrawal, Jane Cleland-Huang, and James Mason. 2021. Hazard analysis for human-on-the-loop interactions in sUAS systems. In Proc. of 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering. ACM, New York, 8–19.
- [81] Guy H. Walker, Neville A. Stanton, Paul M. Salmon, and Daniel P. Jenkins. 2008. A review of sociotechnical systems theory: a classic concept for new command and control paradigms. *Theoretical Issues in Ergonomics Science* 9, 6 (2008), 479–499. https://doi.org/10.1080/14639220701635470 arXiv:https://doi.org/10.1080/14639220701635470
- [82] Martin Weißbach, Philipp Chrszon, Thomas Springer, and Alexander Schill. 2017. Decentrally Coordinated Execution of Adaptations in Distributed Self-Adaptive Software Systems. In Proc. of 11th International Conference on Self-Adaptive and Self-Organizing Systems. IEEE, 111–120.
- [83] Robert H Wortham, Andreas Theodorou, and Joanna J Bryson. 2017. Robot transparency: Improving understanding of intelligent behaviour for designers and users. In Annual Conference Towards Autonomous Robotic Systems. Springer, 274–289.
- [84] Yijun Yu, Danny Barthaud, Blaine A. Price, Arosha K. Bandara, Andrea Zisman, and Bashar Nuseibeh. 2019. LiveBox: A Self-Adaptive Forensic-Ready Service for Drones. *IEEE Access* 7 (2019), 148401–148412. https://doi.org/10.1109/ACCESS. 2019.2942033
- [85] Zijian Zhang, Jaspreet Singh, Ujwal Gadiraju, and Avishek Anand. 2019. Dissonance Between Human and Machine Understanding. In Proc. of ACM on Human-Computer Interaction, Vol. 3. ACM, New York, 1–23. https://doi.org/10.1145/ 3359158