In International Conference on Safety, Security, and Rescue Robotics (SSRR). Würzburg, September 2019.
(preprint)

Team communication processing and process analytics for supporting robot-assisted emergency response

Christian Willms, Constantin Houy, Jana-Rebecca Rehse, Peter Fettke and Ivana Kruijff-Korbayová
{christian.willms, constantin.houy, jana-rebecca.rehse, peter.fettke, ivana.kruijff}@dfki.de

***Abstract*— Mobile robots can provide significant operational advantages in emergency response missions. With increasing autonomy robots need knowledge of the current mission in order to be able to properly contribute to it. We propose to acquire mission knowledge by interpreting the verbal communication among the human response-team members and to use process mining techniques to ground the interpretations in analyses of mission process data and corresponding reference models. We also present a novel concept of mission assistance that uses the acquired mission knowledge to support the first responders' work processes both during and after the mission. The assistance functions include process assistance for the coordination of human-robot team operations; automatic mission documentation generation; and process modeling for first responder training. We describe the architecture of our system and the design and current implementation state of its components: Speech Processing, Mission-Knowledge Management, Process Mining, and Process Assistance. We build on concepts that were evaluated and validated by first responders in a previous project; our extensions have been assessed qualitatively and will be further evaluated in the course of our current project.**

I. Introduction

Emergency response involves operation in high risk situations and making critical decisions despite partial and uncertain information, particularly in medium- to large-scale incidents, such as major fires, floods, landslides or collapsed buildings. The use of mobile robots to access dangerous or inaccessible areas can provide significant advantages for operational safety and capacity, and sometimes even decrease costs. First responders increasingly employ mobile robots, most often aerial vehicles, and sometimes also ground or aquatic robots. Robots are most frequently used for the reconnaissance of an incident site to increase situational awareness, but also for other tasks. For example, during the Notre Dame fire in Paris on April 15 2019, drones were deployed for aerial surveillance and a fire-fighting ground robot pulled a firehose inside the church and sprayed water on the burning ruble [24].

In real emergency response deployments to date, robots have been used as teleoperated tools. Since teleoperation can be very demanding, both cognitively for the operators and technically, autonomy and autonomous assistance functions are the subject of intensive research and development.

With increasing autonomy, robots should become agents capable to actively contribute to a mission. For this, the robotic system consisting of the robot(s) plus the encompassing software system, potentially including other (software) agents, needs *mission knowledge*, i.e., run-time awareness and understanding of the mission goals, the tasks of the human-robot team, how they are being carried out and the state of their execution. Correspondingly, the long-term goal of our research is *to provide emergency-response robotic systems with adequate mission knowledge* that will enable them to contribute to task execution according to first responders' expectations, take initiative in a meaningful way and, more broadly, support the first responders' work processes.

Where and how can the robotic system acquire knowledge of the current mission? The first responders have it; they explicitly communicate it to each other during the mission and they also rely on shared background knowledge. However, since first responders typically operate under high cognitive load and time pressure, it is not a viable option to put the additional burden on them to enter mission knowledge into the system. Our research addresses these questions:

(i) *how mission knowledge can be acquired by interpreting the verbal communication among the human response-team members and grounding it in analyses of mission process data and corresponding reference models*;
(ii) *how to use the acquired mission knowledge to assist the first responders during or after the mission*, for example, by supporting the real-time coordination of human and robot actions or by mission documentation generation.

What this paper is (not) about!

This paper presents initial outcomes of our research. We describe the design of our mission management-support system based on verbal team communication understanding, and its implementation to date. We do not describe the overall robotic system that this is part of nor how mission knowledge is obtained from other sources, e.g., sensors carried by the robots and/or manual input by the first responders.

Since the system is currently under development, we do not present formal evaluation results. However, the underlying concept and an initial design of the speech-processing components as well as an early idea concerning mission task management were previously implemented and evaluated with firefighters during the TRADR project (see II). Throughout the TRADR project, end users were continuously involved in development and evaluation. They also used the system during simulated operations to give us first-hand feedback on the system and its components. This means that a considerable part of the ideas and concepts, which this work develops further, have already been validated.

Further validation has been provided in a focus group held in the context of our current project A-DRZ [1] with (different) firefighters and aimed to complete the practical requirements and to harmonize our development goals with actual first responders' needs. They confirmed the practical usefulness of our current approach and system concept for supporting emergency response missions.

The remainder of this paper is organized as follows: Sec. II describes relevant aspects of the previous TRADR project and other related work. Sec. III presents our current efforts in the A-DRZ project: it describes architectural framework, the design and current implementation status of the processing components of our system, and the process assistance functions. Sec. IV concludes and points out the potential benefits and limitations of our approach.

II. Related Work

Team communication processing to support the strategic management of large-scale emergency response operations, in particular resource management of personnel and vehicles was introduced in the SHARE project [15]. The system captured radio and text messages, distributed them in communication groups, indexed them by keywords and stored the resulting facts in a semantic repository for future retrieval; mission knowledge was represented in an ontology [14].

Projects addressing robot-assisted emergency response mostly strive to improve autonomous robot operation. The NIFTi project [3] was the first to focus on human-robot teaming and interaction for collaborative exploration and assessment of an incident site, including the use of speech to direct robot operation [17]. Also [33] present a framework for interpreting spoken commands combined with pointing gestures to assign tasks to robots in a simulated search&rescue environment. The TRADR project [5] was the first to investigate the spoken communication among the human members of the robot-assisted emergency response team and process it with the aim to understand task assignments and status updates during a mission [10]. A processing pipeline for this purpose and an initial version of the speech processing component was developed in TRADR (cf. Section III-B), as well as a simple rule-based task-management module [31] working on the mission status stored in an ontology, and exemplary working agreements for task allocation within the team [23]. A logging tool was developed for mission documentation purposes [12]. The mission management support based on spoken team communication understanding received very positive feedback in the final evaluation of the TRADR integrated system by first responders.

Our current work builds on the approach introduced in TRADR and aims to overcome its limitations regarding scalability, flexibility and robustness. We use essentially the same processing pipeline, but substantially improve the speech processing and mission-knowledge management components. Mission knowledge is represented in an ontology managed by a semantic repository used as a long-term memory for the mission status and the assistance functions. We replace the TRADR task management and working agreements with methods from business process management.

There have been several attempts to use business process models to support the planning of emergency response missions. However, each mission is unique, which makes it impossible to strictly follow pre-defined processes, like it is done in manufacturing scenarios. This is why [26] speak of the "Myth of Business Process Modelling for Emergency Management Planning". Their literature overview includes numerous reports on (more or less successful) attempts of using process modeling in the emergency domain, e.g., [13], [25], [30]. Our system does not primarily address the planning of emergency response missions, but uses a data-driven approach to mission process management instead. We aim at overcoming the challenges of process modeling for emergency management planning by gathering real-time mission process logs and discovering process models from these logs, which can be used to support the documentation, control, and analysis of emergency response mission processes. This means that compared to previous works, we are using a bottom-up process mining approach instead of a purely top-down process modeling approach. In order to replace and improve upon the working agreements used in TRADR, in a second step (not described in this paper), this control-focused process view will be supplemented with techniques for organization mining, supporting the modeling and management of human-robot teams and their capabilities, the assignment of specific tasks to individual team members, and their subsequent monitoring.

III. System Description

The aim of our system is to acquire mission knowledge from the verbal communication among the human response-team members and use it to provide process assistance to them in any phase of a deployment, i.e., when approaching the incident site, during the operation and for post-processing. In this section, we first describe the overall architecture design of the system, the main processing components, and finally the process assistance functions.

A. Architecture

The system has a modular architecture consisting of four major components: (1) Speech Processing, (2) Mission-Knowledge Management, (3) Process Mining, and (4) Process Assistance. Fig. 1 depicts the work and data flow.

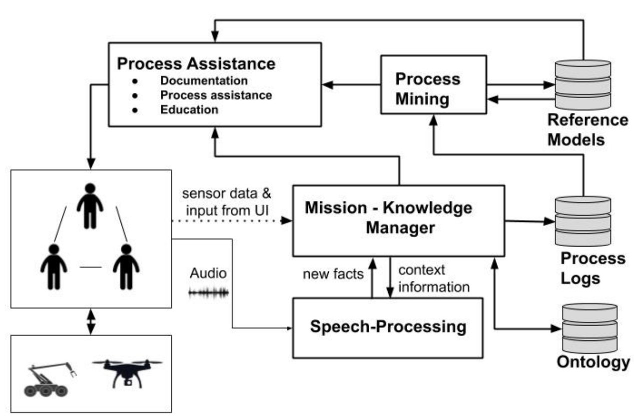

Fig. 1. System architecture overview

The team communication between the human team members is continuously captured and processed, i.e., interpreted, by the Speech Processing component. The interpretations are then used by the Mission-Knowledge Manager to update the ontology, which functions as a long-term memory for the mission knowledge/status. The mission knowledge is also continuously updated with facts derived from other sources, e.g., sensor data and user input. All changes to the ontology are stored in a shared log file, which later serves as input for Process Mining. Finally, the Process Assistance component uses the gathered information to provide domain-specific support to the human team members. The following sections describe the four components in more detail.

B. Speech Processing

The speech processing pipeline uses state of the art natural language processing components, notably Automatic Speech Recognition (ASR) and Natural Language Understanding (NLU), as shown in Fig. 2.

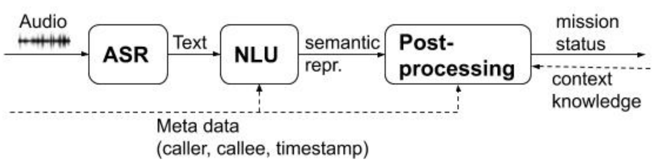

Fig. 2. The speech processing pipeline

First, the raw team communication audio is sent to the ASR module, which outputs text. We have integrated two ASR implementations, namely Nuance Mix ASR [4] and Kaldi [28], [22]. Nuance Mix ASR is a cloud-based solution using pre-trained commercial quality language models provided by Nuance. Domain adaptation is also possible by uploading specific vocabularies. Kaldi is a local solution where arbitrary models can be used. We use existing open source models and plan to adapt them to our domain by transfer learning. We decided to use two ASR modules for redundancy reasons. Nuance ASR performs better off-the-shelf, but needs a network connection. Since internet access is sometimes restricted or unavailable in emergencies, Kaldi serves as a locally running fallback solution.

Second, the text output of the ASR is interpreted by the NLU component. This consists, for each utterance, of dialogue act (DA) recognition and task interpretation. In TRADR we used Nuance Mix intent recognition which intertwined both of these. This had many disadvantages. We have therefore implemented a separate DA recognition module [7]. It uses the Keras [9] implementation of FastText, a library for sentence classification created by Facebook's AI Research lab [21], trained on data acquired during the TRADR project [6] and manually annotated with DAs (e.g., request, offer, inform, answer), task types (e.g., explore_area, take_photo) and task-related information (e.g., actor, location, point of interest). The DA recognition module considers not only the text input from the ASR module but also meta information, e.g., the speaker. Since we use a push-to-talk voip system (mumble [2]), the speakers are automatically identified by their role, e.g., Team leader (TL) or UAV operator (UAV1). The DA recognition module outputs a list of the most likely DAs, compare Fig. 3. It has accuracy of approx. 78%. Re-implementation of the task interpretation is work in progress.


```
 "Speaker": "TL",
 "Text": "Explore area A and report.",
 "DA": { "Request": "0.90350974",
         "Inform": "0.019799994",
         ...}
"Task": "Exploration_Task"
"Task_Parameter": {"Area A"}
```


Fig. 3. Example for the combined output (DA and Task interpretation) from the NLU module.

The post-processing module fills in missing information based on the context and re-ranks the NLU hypotheses if needed. The most likely semantic representation is transformed into a corresponding RDF representation and sent to the Mission-Knowledge Manager. For example, assuming that the TL placed a ContactRequest on UAV1 before assigning the task in Fig. 3, we can conclude that this task was assigned to UAV1, even if this was not mentioned in the utterance itself. Re-ranking the DA results is not necessary here, as they make sense w.r.t. the dialogue structure.

C. Mission-Knowledge Manager

The Mission-Knowledge Manager is a data hub fulfilling the tasks of managing, storing, and distributing the data within our framework. It consists of a semantic repository and reasoning component and an ontology representing the mission knowledge itself. The semantic repository manages the facts in the ontology providing functions for manipulating but also querying the ontology. The reasoner is used to infer new information based on the ontology.

We employ HFC [16] for the semantic repository and reasoner. HFC supports RDF n-tuples, and thus allows us to encode each fact in the ontology as an RDF-triple (according to the w3c standard) and annotate it with additional arguments, such as transaction times. The annotation with time stamps is particularly important for the creation of process logs and thus for the process assistance functions.

The ontology we have developed is modular, consisting of several sub-ontologies as shown in Fig. 4. We have partially adapted the ontologies used in TRADR, and have also incorporated other research results, e.g., Robin Murphy's ontology for the integration of robot and agent systems [19].

Each sub-ontology covers a distinct domain and is thus easily exchangeable. The Actors ontology features the capabilities and affiliations of human and robots. The Communication and Dialogue ontologies represent communication events during a mission. This includes roles such as caller and callee, the communicated information as well as dialogue acts following the ISO 24617-2 Standard [8].

The main ontology combines these sub-ontologies by using equivalence relations for common concepts and by

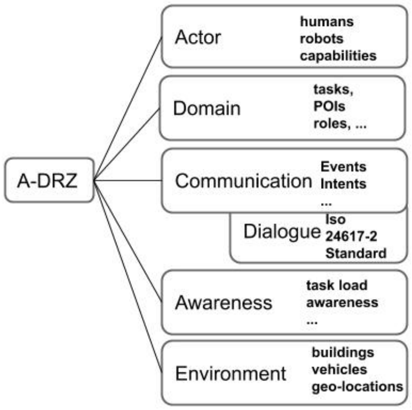

Fig. 4. The high-level structure of the used ontology.

introducing new relations to connect, for example, particular communication events to tasks and Points of Interest (POIs).

The modular structure of the ontology allows us to easily adapt it towards specific use-cases by exchanging sub-ontologies or adding new domain-specific sub-ontologies.

D. Process Mining

The Mission-Knowledge Manager generates real-time process logs, which document actions and communication during a mission. These logs are combined with existing knowledge about relevant emergency processes, established practices, and regulations in disaster scenarios, codified in an extensive collection of *reference models* [11], which we developed for the A-DRZ project. They were derived from first responders' handbooks and training material. The reference models provide a generic blueprint for how a mission could or should be conducted, e.g., in case of a burning chemical plant. They document existing expertise on efficient action, such as handling injuries, and communication, such as radio protocols, but are easier for humans to understand and reason about than the existing text-heavy handbooks.

The Process Mining module is responsible for processing the logs to be used for the process-model-based assistance functions, e.g., activity recommendation, in the next step. Its main purpose is to provide a bridge between actions and communication during a mission, captured in the ontology, and the mission-independent knowledge represented in the reference models. The module contains generic methods that support the three main functions of process mining [32]:

- Process Discovery: The sequential log data is transformed into a two-dimensional graphical process model, easier to read and understand for humans. We use state-of-the-art approaches such as the Inductive Miner [18].
- Conformance Checking: Next, we compare the as-is mission process with the to-be reference process to check its conformance. The goal is to find and explain deviations between mission process and reference process, which could present risks to the mission success.
- Model Enhancement: The mission logs can also be used to extend or improve the reference models. For example, if they describe a scenario, which is not yet documented in the reference model, or if compliance checking reveals that the prescribed process is impractical to follow during a mission, the reference models can be updated accordingly, using for example inductive methods for reference model development [29].

We realize these non-domain-specific functions by using existing implementations from the ProM framework for process mining[1] and integrating them with process analytics components from our RefMod-Miner[2] research prototype [20]. The results are used in the next component to provide domain-specific assistance functions for emergency response.

E. Assistance Functions

The components described above establish the foundation for a set of assistance functions making the collected information accessible to the team members. In discussions with first responders in TRADR and A-DRZ we identified the following assistance functions as most useful. We have developed a domain-specific visualization for each of them, which helps the responders to understand and benefit from the data produced by the process mining module. The underlying functionalities will be implemented using the existing implementations from the process mining module.

Mission documentation: First responders write a report after each mission, explaining why they made certain decisions, e.g., regarding deployment. We support them by automatically creating mission documentation from the collected data and process models, cf. Fig. 5. The process model shows how the mission progressed and can be enriched with the corresponding events, the elapsed times, a summary of the team communication and audio recordings. The system can also provide the corresponding reference models for comparison. The team can then discuss, explain, and document the potential deviations from the reference model detected during conformance checking in a debriefing session.

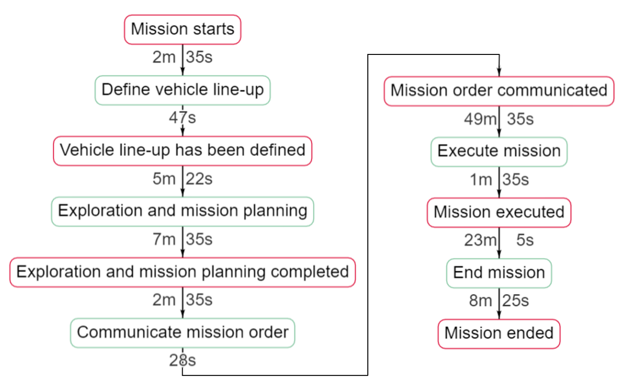

Fig. 5. Mined process model for specific deployment documentation

Process assistance: During a deployment, the officer in charge of team coordination oversees many simultaneous activities, from both humans and robots. To reduce cognitive

[1]http://www.processmining.org/prom/start
[2]https://refmod-miner.dfki.de

load, the system provides a graphical process assistance. The process model presents an overview of completed activities and the steps that should be taken next. If the current activity has multiple potential outcomes, which could initiate different next steps, these options are also displayed. For example, if a robot inspects an unknown object for potentially hazardous chemicals, the officer will see which steps should be taken, if such substances are found. This assistance function is realized by combining real-time process logs with process knowledge from the reference model. Fig. 6 shows an example where the officer oversees two robots (blue and yellow), which are currently completing two different tasks at different locations. Process model abstraction methods [27] are used for displaying the models with an appropriate level of detail, differentiating between different command levels.

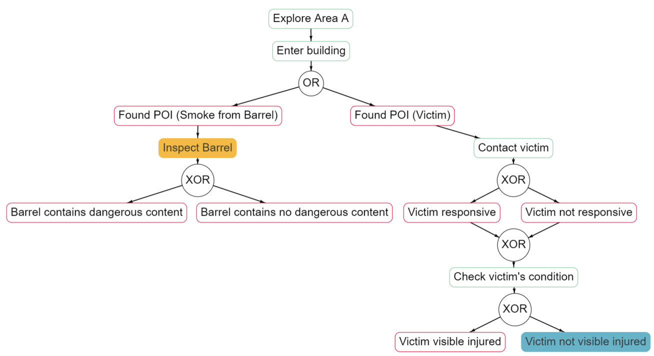

Fig. 6. Model-based process assistance in real time

Education: When training new first responders, it is important that they not only study guidelines and prescribed behavior from handbooks and reference models, but also experience how these guidelines are enacted in a real mission. A discovered process model can be enriched with additional information such as regulations or audio recordings, to help the trainees understand the involved team members' actions, decisions, and communication. This assistance function can also be used to (re-)train personnel for specific emergency situations, e.g., nuclear incidents. Conformance checking plays a vital role here, as it helps trainees to understand the role and relevance of the reference models. Fig. 7 shows a comparison between a deployment model and a reference model, indicating documented deviations.

In addition to the above process assistance functions, a situation awareness system is being developed as part of the robotic system in the A-DRZ project. It has 2D and 3D maps of the incident site showing, e.g., the positions of the robots and points of interest. There will also be an option to annotate these maps with thermal or CBRN-related information. Our assistance functions will be tightly integrated with this system.

IV. Conclusions

We presented a system that uses team communication processing and process analytics techniques for supporting robot-assisted emergency response. The main novel idea is

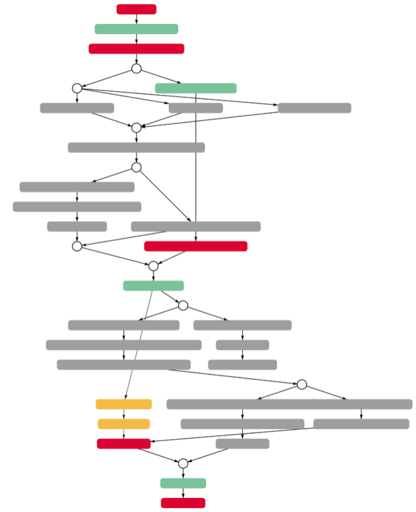

Fig. 7. Structural deviations between mission and reference model (missing activities in grey, additional ones in yellow)

to acquire operational mission knowledge by understanding the verbal communication among the human team members.

The system captures the verbal team-communication and uses state-of-the-art natural language processing tools to interpret it. A mission knowledge module with an ontology stores the information extracted from the team communication and to link it to knowledge from additional sources. All changes to this mission knowledge are reflected in a shared process log, which is continuously analyzed by the process mining module to combine the verbal input of mission process data with the corresponding reference models. Both the mission knowledge management and the process mining module support the outlined process assistance functions, i.e., deployment documentation, process assistance, and education. These functions will be connected to the situation awareness system and thus accessible to the first responders during (online) and after (offline) the mission.

We build on concepts previously validated in the TRADR project. Our extensions have been assessed qualitatively so far; further qualitative and quantitative evaluation will happen as part of the iterative user-centric development cycle in the A-DRZ project, where the first responders periodically provide feedback on the implemented features and modules.

Our next steps are to continue with the implementation of the described components. One future challenge is the current lack of in-domain data for ASR, which is crucial, as the first responders use many abbreviations and special modes of expression. We are making arrangements with the German fire services in the A-DRZ project to obtain recordings of radio communications from exercises and real missions that we will annotate and use to expand the ASR models. Periodic collection of additional data is also needed for NLU development. We have seen in the TRADR project that the team communication changes as the robotic system develops, e.g., when robot capabilities and user interface functionalities are modified and new ones added. NLU needs to adapt accordingly to keep up with these changes.

An important challenge for the overall concept is that the system will not be able to catch all mission-relevant information by just monitoring the team-communication and sensor data. In their current work setup, German firefighters share at least 50% of the information in direct communication and not via radio. This is obviously a big concern for us, as we have no way to overhear these conversations without disturbing the first responders' work. An interesting research question is whether it is possible to (partially) infer the missing mission knowledge from the further course of the mission and from other sources, such as robot positions and inputs in user interfaces. This is a very ambitious task and should be considered a research project by itself.

V. Acknowledgements

Our work builds on many past and present collaborations. We thank our collaborators in TRADR, with whom we designed and developed the initial version of this framework, and our A-DRZ collaborators for valuable discussions. We thank the first responders who participated in experiments and discussions in NIFTi, TRADR and A-DRZ for their invaluable input. Our present work is funded by the German Ministry of Education and Research (BMBF) through the A-DRZ project [1], grant No. I3N14856.

References

[1] A-DRZ: Aufbau des Deutschen Rettungsrobotik-Zentrums (Setup of the German Rescue Robotics Center). https://rettungsrobotik.de. Accessed: 2019-04-24.

[2] Mumble - a voice chat utility for gamers. https://wiki.mumble.info/wiki/Main_Page. Accessed: 2019-07-08.

[3] NIFTi: Natural human-robot collaboration in dynamic environments. https://nifti.eu. Accessed: 2019-04-24.

[4] Nuance Mix NLU. https://developer.nuance.com/public/index.php?task=mix. Accessed: 2019-04-24.

[5] TRADR: Long-term human-robot teaming for disaster response. https://tradr-project.eu. Accessed: 2019-04-24.

[6] TRADR: Team communication corpus. https://talkingrobots.dfki.de/resources/tradr. Accessed: 2019-04-24.

[7] Anakina, T., and Kruijff-Korbayová, I. Dialogue act classification in team communication for robot assisted disaster response. In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue, Stockholm, Sweden, September 11-13, 2019* (to appear).

[8] Bunt, H., Alexandersson, J., Choe, J., Fang, A., Hasida, K., Petukhova, V., Popescu-Belis, A., and Traum, D. Iso 24617-2: A semantically-based standard for dialogue annotation. *Proc. of LREC-2012, European Language Resources Association* (2012).

[9] Chollet, F., et al. Keras. https://keras.io, 2015.

[10] de Greeff, J., Mioch, T., Vught, W. v., Hindriks, K., Neerincx, M. A., and Kruijff-Korbayová, I. Persistent robot-assisted disaster response. In *International Conference on Human-Robot Interaction (HRI-2018)* (2018), ACM.

[11] Fettke, P., and Loos, P. Perspectives on reference modeling. In *Reference modeling for business systems analysis*, P. Fettke and P. Loos, Eds. Idea Group Publishing, 2007, pp. 1–20.

[12] Kasper, W. Team monitoring and reporting for robot-assisted usar missions. In *International Symposium on Safety, Security, and Rescue Robotics (SSRR)* (2016), IEEE, pp. 246–251.

[13] Kittel, K., and Sackmann, S. Gaining flexibility and compliance in rescue processes with bpm. In *6th Int. Conf. on Availability, Reliability and Security (ARES), Vienna, Austria* (2011), IEEE, pp. 639–644.

[14] Konstantopoulos, S., Paliouras, G., and Chatzinotas, S. SHARE-ODS: an ontology data service for search and rescue operations. In *Hellenic Conference on Artificial Intelligence* (2006), Springer, pp. 525–528.

[15] Konstantopoulos, S., Pottebaum, J., Schon, J., Schneider, D., Winkler, T., Paliouras, G., and Koch, R. Ontology-based rescue operation management. In *Int. Workshop on Mobile Information Technology for Emergency Response* (2008), Springer, pp. 112–121.

[16] Krieger, H.-U., and Willms, C. Extending OWL ontologies by cartesian types to represent n-ary relations in natural language. In *Language and Ontologies 2015* (2015).

[17] Kruijff, G., Kruijff-Korbayová, I., Keshavdas, S., Larochelle, B., Janíček, M., Colas, F., Liu, M., Pomerleau, F., Siegwart, R., Neerincx, M., Looije, R., Smets, N., Mioch, T., van Diggelen, J., Pirri, F., Gianni, M., Ferri, F., Menna, M., Worst, R., Linder, T., Tretyakov, V., Surmann, H., Svoboda, T., Reinštein, M., Zimmermann, K., Petříček, T., and Hlaváč, V. Designing, developing, and deploying systems to support humanrobot teams in disaster response. *Advanced Robotics 28*, 23 (2014), 1547–1570.

[18] Leemans, S., Fahland, D., and van der Aalst, W. Discovering block-structured process models from event logs containing infrequent behaviour. In *Business Process Management Workshops* (2014), N. Lohmann, M. Song, and P. Wohed, Eds., Springer, pp. 66–78.

[19] Long, M., and Murphy, R. A proposed approach to semantic integration between robot and agent systems. *Journal of Human-Computer Studies 59*, 6 (2005), 983–1024.

[20] Loos, P., Fettke, P., Walter, J., Thaler, T., and Ardalani, P. Identification of business process models in a digital world. In *BPM – Driving Innovation in a Digital World*, J. vom Brocke and T. Schmiedel, Eds. Springer, 2015, pp. 155–174.

[21] Mikolov, T., Grave, E., Bojanowski, P., Puhrsch, C., and Joulin, A. Advances in pre-training distributed word representations. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2018)* (2018).

[22] Milde, B., and Köhn, A. Open source automatic speech recognition for German. In *13th ITG Symposium on Speech Communication* (2018).

[23] Mioch, T., Peeters, M., and Neerincx, M. Improving adaptive human-robot cooperation through work agreements. In *Proceedings of the 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)* (2018), IEEE, pp. 1105–1110.

[24] Nardi, T. The drones and robots that helped save Notre Dame. https://hackaday.com/2019/04/17/the-drones-and-robots-that-helped-save-notre-dame. Accessed: 2019-03-05.

[25] Peinel, G., and Rose, T. Deploying process management for emergency services-lessons learnt and research required. In *Future Security-6th Security Research Conference, Berlin, Germany* (2010).

[26] Peinel, G., Rose, T., and Wollert, A. The myth of business process modelling for emergency management planning. In *9th International Conference on Information Systems for Crisis Response and Management (ISCRAM), Vancouver, Canada* (2012).

[27] Polyvyanyy, A., Smirnov, S., and Weske, M. Business process model abstraction. In *Handbook on Business Process Management 1*. Springer, 2015, pp. 147–165.

[28] Povey, D., Ghoshal, A., Boulianne, G., Burget, L., Glembek, O., Goel, N., Hannemann, M., Motlicek, P., Qian, Y., Schwarz, P., Silovsky, J., Stemmer, G., and Vesely, K. The kaldi speech recognition toolkit. In *Workshop on Automatic Speech Recognition and Understanding* (2011), IEEE.

[29] Rehse, J.-R., Fettke, P., and Loos, P. A graph-theoretic method for the inductive development of reference process models. *Software & Systems Modeling 16*, 3 (2017), 833–873.

[30] Rüppel, U., and Wagenknecht, A. Improving emergency management by formal dynamic process-modelling. In *24th Conference on Information Technology in Construction, Maribor, Slovenia* (2007), pp. 559–564.

[31] Saad, E., Hindriks, K., and Neerincx, M. Ontology design for task allocation and management in urban search and rescue missions. In *Proc. of the 10th Int. Conf. on Agents and Artificial Intelligence (ICAART 2018) - Volume 2* (2018), pp. 622–629.

[32] van der Aalst, W. *Process Mining: Data Science in Action*. Springer, 2016.

[33] Yazdani, F., Kazhoyan, G., Bozcuoğlu, A. K., Haidu, A., Bálint-Benczédi, F., Bessler, D., Pomarlan, M., and Beetz, M. Cognition-enabled framework for mixed human-robot rescue teams. In *International Conference on Intelligent Robots and Systems (IROS)* (2018), IEEE, pp. 1421–1428.