TOWARDS A SENSOR-BASED INTERNAL EPISODIC MEMORY FOR AUTONOMOUS MOBILE ROBOTS

Tristan Schnell, Sven Vogel, Lennart Puck, Timothée Buettner, Arne Roennau, and Rüdiger Dillmann

Department of Interactive Diagnosis and Service Systems (IDS), FZI Research Center for Information Technology, Haid-und-Neu-Straße 10–14, 76131 Karlsruhe, Germany.

ABSTRACT

Mobile robots are used in more and more difficult scenarios with increasing degrees of autonomy. Especially working autonomously in unknown environments requires robots to quickly adapt to unforeseen circumstances to avoid potentially harming themselves. When these dangerous situations cannot be fully predicted ahead of time, it is vital for the robot to detect these when they occur and further learn to avoid them or adapt to them in the future to ensure safe long-term operability. To learn these adaptations, we propose an explicit, long-term episodic memory using robot sensor data for mobile robots, inspired by the human episodic memory. This is implemented by continuously recording all robot sensor data as a short-term memory during operation, while a data-based anomaly detection system monitors the robot for any unexpected events to then create an explicit memory. This allows the robot to recognize similarities between its current situation and relevant previous episodes by associative memory recall. We evaluated a prototype using exclusively internal robot sensor data on the ANYbotics ANYmal C robot, by damaging the robot after it walked up a ramp and testing its ability to recall the generated episode in a similar situation. Overall, we showed that an explicit long-term memory system based on the human episodic memory could be a feasible solution for a robot to efficiently learn to adapt to unknown situations.

Key words: Autonomous robots; Machine Learning; Walking Robots.

Figure 1. The high-level concept idea of an ANYmal robot with episodic memory. When in a situation the robot has previously been in and teamed important, for example due to detection of some error or anomaly , it is able to actively remember and replay this memory. This allows the robot to evaluate its situation based on this explicit past experience.

1. INTRODUCTION

Successfully completing complex and dangerous missions often requires humans to use a vast amount of experience. When faced with a difficult challenge, we are able to not only utilize learned behaviors but also actively remember and critically reflect on similar situations we previously encountered. This ability to continuously adapt and learn is something we have long tried to incorporate into robotic systems. Especially in recent years, machine learning approaches have produced promising results in variety of fields, such as machine vision, locomotion and manipulation. All of these techniques, though, are based on implicitly learning tasks by continuously adapting parameters using large amounts of data or experimentation. The complementary type of learning in the form of explicitly gathered and recalled memories, which is especially important for rare and potentially dangerous situations where example data cannot easily be collected for training of machine learning models, is significantly less utilized in robotics so far.

Episodic memory as a concept was initially introduced by psychologist Endel Tulving in [1] as a differentiation to semantic memory. It contains explicit, autobiographical memories and allows humans to remember their experiences in an active manner often referred to as *mental time travel* [2]. The recall of these memories is associative, based on sensory stimuli or spatial and temporal references to other events.

Several approaches to episodic memories intended for robots have been developed over time, especially as part of general cognitive architectures [3]. The ISAC (Intelligent Soft-Arm Control)[4] architecture implements different types of memory in a multi-agent approach, including an episodic long-term memory. Episodes in this model contain external sensor data, and explicit modeled internal states and events, such as the robot's goals and tasks [5]. To describe the relevance of an event that occurred, a salience level is introduced that declines with time, but can also be boosted when the associated or similar memories are recalled [6].

The Soar architecture, which is a rule-based general purpose cognitive architecture [7], similarly implements different types of memory. For these, explicit methods of encoding, storage and retrieval are implemented for comprehensive activation of memory [8]. A working memory holds these currently relevant and interesting experiences, which are removed depending on a decay function [9]. Built on this, with Soar-EpMem a task independent episodic memory is implemented [10]

The Deep Episodic Memory [11] implements a LSTM-based encoding for image series of tasks, allowing for inference of similar recorded episodes. Finally, EPIPROME [12] offers an abstract, high-level episodic memory of modeled events to improve planning based on past experiences.

The existing general approaches explicitly model the internal states and potential actions of a robot, which causes two significant weaknesses. Firstly, this modelling is a complex and time-consuming task, which mostly has to be repeated for different robots or scenarios and requires a high degree of precision and completeness to allow for generation of accurate memories. Secondly, by defining these specific events and thus relevant elements of the memories before usage, the robot cannot extract all potentially important data from its memories. Details initially assumed to be minor, though, also have to potential to be of critical importance when reoccurring in a similar situation in the future. Due to this, we developed and implemented a prototype concept for an episodic memory, that directly operates on robot sensor data. It, thus, does not require manual modeling or expert knowledge, while also not assuming that important elements of memories can be reliably detected before or during their occurrence. Relevant episodes to be recorded can be chosen based on an anomaly detection system to avoid creation of a large body of memories of regular successful task executions. Recall of memories is possible by comparing the current internal situation of the robot with sensor data recorded in the episodic memory. The created model can serve as a basis for a system dynamically adapting behaviors or choosing mission strategies based on critical memories of the robot.

This paper is structured as follows: Section 2 describes the conceptual requirements defined for the developed memory, the robot setup used and the concept and implementation of the model. Section 3 presents the performed experiments and their results. Finally, in section 4 we summarize the work and describe the current and future work to extend the developed system.

2. EPISODIC MEMORY

The developed prototype episodic memory should enable a robot to remember dangerous situations it has previously been in. When an error occurs, the robot has to be able to record this experience. During operation, these memories should be recalled, once the robot is in a similar situation.

2.1. Requirements

Several conceptual requirements to be fulfilled by the prototype system were defined, based on the biological concept of an episodic memory.

- Sensor data of the robot should be used and recorded directly without interpretation, instead of via manually defined events.
- Generated episodes have to be saved for an arbitrary long amount of time.
- Remembering should be based on the sensor data itself to allow for associative recall.
- New episodes should be able to be generated at directly at runtime.
- Episodes are mainly based on internal sensor data, but should also contain contextual meta-information and should easily be extendable.

2.2. Robot setup

While the developed system is mostly robot independent, the reference platform during development and testing was chosen to be the ANYbotics ANYmal C [13] (seen in Figure 1). As a walking robot, locomotion changes potentially significantly based on internal and external factors such as the ground, allowing potential differences in the recorded internal data to be detected. Unlike other commercially available walking robots, the platform is open and different internal sensor values can be easily read. For this prototype, we decided to use the following sensor inputs: position, velocity, acceleration, torque and current for all 12 joints of the robot, resulting in a total of 60 sensor values for each timestep. While there are more relevant internal sensor values that could specifically be used to easier solve the experiment described later on, such as the robots IMU, we decided to only use the most general sensors to describe the robot's locomotion to ensure the transferability to other scenarios. While not directly including external sensor data, the episodic memory also utilizes the robot's position as coordinates in a map. For this, the VDB mapping framework [14] and its dependencies were used on the robot.

2.3. Concept

The episodic memory system is designed using the Robot Operating System (ROS) [15]. This allows for easy modularization and a relatively high degree of abstraction and thus hardware independence, and its structure is presented in Figure 2. The first component, implemented in the *preparingSensorData* node collects, synchronizes and normalizes all considered sensor data into a single message and is the only robot-dependent element.

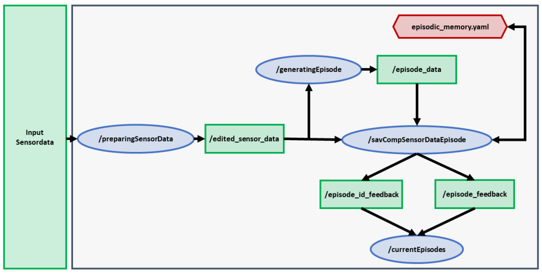

Figure 2. Overview of the episodic memory system built with ROS - with nodes or active system components in blue, the data they exchange as messages in green and files saved to the hard drive in red. The preparingSensorData node collects, synchronizes and normalizes sensor data of the robot. In the generatingEpisode node episodes are created to be saved whenever an anomaly is detected by an external system. The savCompSensorDataEpisode node stores memories and compares the current situation to all recorded experiences. Finally, in the currentEpisode node all memories currently considered relevant by the robot are presented.

The second component (the *generatingEpisode* node in Figure 2) is responsible for the creation of new memories. To accomplish this, a window of the last sensor values is continuously stored as a short-term memory in the form of a ring-buffer (Figure 3). Once an anomaly occurs, detected by a previously developed system detailed in [16], generation of an episode is triggered. This episode contains the short-term memory at that time and additional sensor values until no more anomaly is detected for a given time or a maximum episode duration is reached. Additionally, meta-information such as time and the location at the start of the memory are added. The created episode is then transferred to the next component of the system to be used and stored.

The primary node (*saveCompSensorDataEpisode* in Figure 2) has two main roles. On the one hand, it is responsible for storage of episodes, both in memory and in a persistent manner on the robot's hard drive. Its more complex function on the other hand is the ability to recall similar memories when given a context, either continuously

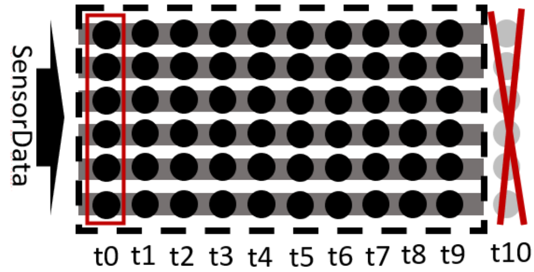

Figure 3. The short-term memory of the robot, implemented as a ring-buffer. For the last n (here 10) timesteps, the synchronized and preprocessed data values of all sensors are stored.

by using its current short-term memory or when triggered directly with some context information. To accomplish this, it calculates and updates a qualifier value for all existing memories, which is a weighted combination of 3 specific values: a data-based qualifier, a location-based qualifier and a time-based qualifier.

The data-based qualifier compares the given sensor data directly to the sensor values in each episode. As the internal sensor data of mobile robots, especially legged ones, has a strong periodic component data is not only compared directly, also in the frequency domain via Fast Fourier transform with zero-padding. This heavily reduces variance dependent on the precise moment of the locomotion cycle a memory was created in. This data-based qualifier is the most important factor when looking for similarities to previously encountered episodes and is evaluated in simulation and on a real robot in the following chapter.

The location-based qualifier is intended to describe the similarity between the current location and the environments the episodes occurred in. As for this prototype implementation no external data was used directly, this was simply implemented as a difference in coordinates in the robot's map, using the squared distance between locations and a configurable cut-off (as seen in Figure 4). This only allows the robot to remember episodes that occurred at the same location, but in the future we plan to extend this with a similarity measure based on the environmental data.

Finally, the time-based qualifier is used to measure the importance of memories based on how often and how recently they or episodes around the same time frame are recalled. As a baseline it is modeled as a slow exponential decay. Whenever the memory is recalled, its time-based qualifier is increased, as is demonstrated in Figure 5. Additionally, whenever a recall occurs, all other episodes that were recorded at a similar time also get a smaller boost in their time-based qualifier. This allows the chaining of memories, which is useful as once an episode is deemed important, their full context as known also potentially increases in significance.

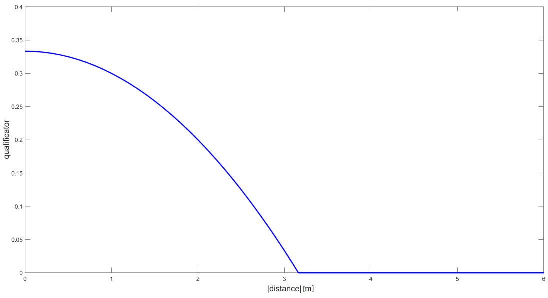

Figure 4. The location-based qualifier used when considering the relevance of a recorded episode. A penalty is applied using the squared distance of the location the episode was recorded at the current position of the robot. As similarities should also be detected at different locations, no negative qualifier is applied for large distances.

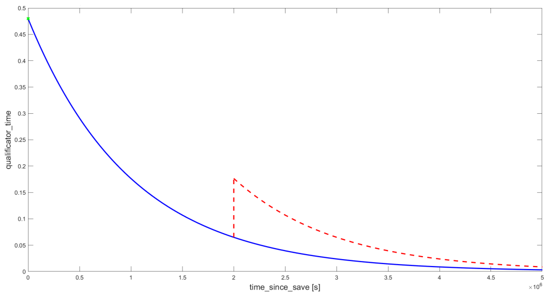

Figure 5. The time-based qualifier used when considering the relevance of a recorded episode, inspired by [6]. Using exponential decay, it is continuously reduced when the memory is not in use (blue). When it or other memories that were recorded at similar times are recalled, it is increased (red), causing easier recall of important memories.

These three qualifiers can be weighted depending on their importance as defined by the user. Initial testing showed the usefulness of all three, and their ability to cover each other's shortcomings. For example, when initially the data comparison was poorly tuned and did not allow for reliable recall alone, the addition of the other qualifiers allowed the system as a whole to function as planned. This, though, also showed that potential weaknesses in the main focus of this prototype implementation, the comparison of internal robot data, could be missed due to being covered by the other elements. Due to this, we decided to disable the time- and location-based qualifiers for the small-scale evaluation presented in this work, to be reevaluated in a larger test case at a later point.

Finally, the *currentEpisodes* node stores all currently active memories based on their qualifiers. When used to assist an external mission control instead of an autonomous robot system, this node is run on the base station and allows the operator to replay previous scenarios the robot considers relevant to its current situation. In full autonomy, these episodes are used by the robot to adapt and inform its decisions and could be utilized to critically reflect its past experiences based on its current situation, similar to the human *mental time travel*.

3. RESULTS AND EXPERIMENTS

The goal of the experiments is the evaluation of the robot's ability to recognize a dangerous situation it previously encountered before the time of the actual damage. Specifically, we deployed the prototype episodic memory on the ANYbotics ANYmal C and, after traversing even ground for some amount of time, walked up a slight ramp. Once the top of the ramp was reached, we manually caused the robot to fall down to simulate a critical error. This should be recognized as an important, anomalous event and a memory should be created. After, we again traversed the area while continuously querying the episodic memory using the current short-term memory. During locomotion on flat ground, the similarity to the generated memory should be small enough to not trigger recall. Once the robot is again directed to move up the ramp, it should be able to recall the prior critical event before reaching the top of the ramp, thus avoiding a repeat of the fatal error. To ensure that the sensor-based recall is working, we disabled the location and time-based components of the memory qualifiers, requiring the robot to detect the similarity purely based on raw internal sensor data.

If recall is performed successfully, we show that the core component of the prototype episodic memory is able to recognize previously encountered situations and would then be able to adapt its behavior before a similar error can occur. Additionally, the system's ability to only generate memories of relevant experiences can be demonstrated if the only recorded memory is of the triggered error.

The creation of new memories performed exactly as planned. No new memories were developed during the regular behaviors, but once the error was triggered a new episode of that moment was reliably created. As the component responsible for this is based on prior work [16], we did not examine its reliability further in these experiments. The recall of memories using only internal sensor data was initially tested on simulated data, the results of which can be seen in Figure 6. After 30 seconds of locomotion on flat ground, the robot is commandeered to walk up the ramp, which corresponds with a clear increase in the qualifier shortly thereafter. Once at the top, the robot stands still, causing the qualifier to quickly drop again. Additionally, while there is some variation in the qualifier based on the robot's exact movements, during regular operation there are no similar spikes.

For the second experiment on the real robot, locomotion was performed significantly slower. Once again, for a shorter period of 10 seconds in the example presented in Figure 7, the robot walked on even ground, causing no increase in the qualifier. Once the robot slowly reaches the ramp, the qualifier increases significantly, until a short

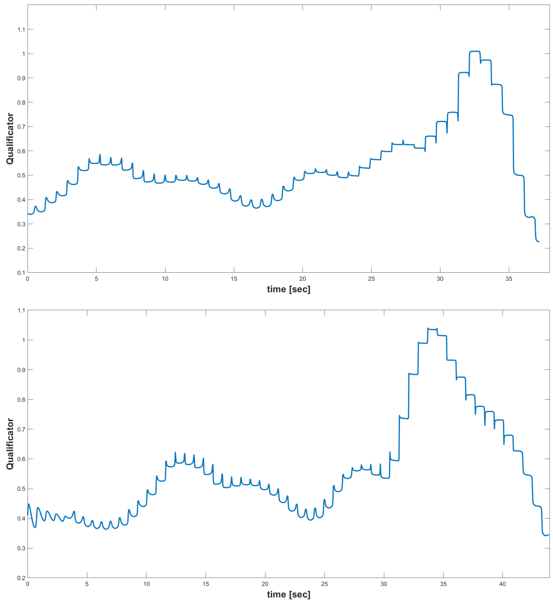

Figure 6. The data-based qualifier describing the similarity of the robot's situation to a recorded episode in which the robot crashes after fully walking up a ramp - tested on simulated data. For the initial 30 seconds the robot walks on even terrain, and no significant increases in the qualifier occur. Once the robot starts walking up the ramp, it drastically increases, showing that the robot is able to remember the recorded episode in both runs. After the robot stops, it quickly decreases again.

stop and turn has to be performed at about 20 seconds, again causing a reduced similarity to the recorded memory. Afterwards, the qualifier increases again and remains at a high level until locomotion is stopped.

Overall, it can be seen, that in the simple tested examples both creation and recall of memories performed well. Memories were only generated when an anomaly occurred, avoiding cluttering of the episodic memory. The qualifier describing data-based similarity to existing episodes was continuously relatively low during locomotion on flat ground, with only some amount of variance. During traversal of the ramps, it increased significantly, showing the robots ability to recall its earlier memory of the same task.

It has to be noted, though, that the implementation and correspondingly the initial experiments are a proof-of-concept with reduced scope. The relative inefficiency of the recorded memories and the process of continuous comparison of the current state to all existing memories scale poorly to larger bodies of experience by the robot. Additionally, while the components not experimentally evaluated, the location- and time-based qualifiers, can easily be shown to work as intended when isolated, their effect on the whole system has to be reviewed in a more complex use-case with many memories over a longer time-period and area. Finally, to fully utilize the episodic memory, a more direct integration into the mission control system is necessary. The ability to recognize similar previously encountered situations in itself can be purely used as a warning system, but to utilize its full potential, the mission control has to be able to reflect on the relevant episodes and dynamically adapt its behavior accordingly.

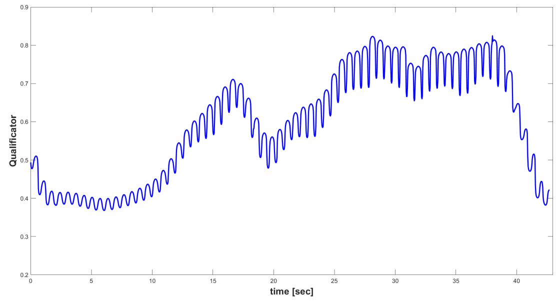

Figure 7. The data-based qualifier describing the similarity of the robot's situation to a recorded episode in which the robot crashes after fully walking up a ramp. For the initial 10 seconds the robot walks on even terrain, and no significant increases in the qualifier occur. Once the robot starts walking up the ramp, it drastically increases, showing that the robot is able to remember the recorded episode. After the robot stops, it quickly decreases.

The prototype episodic memory presented offers the baseline for a learning system using explicit knowledge through experiences. It can be shown, that without the manual modelling of events or states present in similar work, the robot is able to realize when in a potentially critical situation it has previously encountered. While further work to increase scope and efficiency is necessary, its principal ability to quickly improve the robot's knowledge base could be demonstrated in the considered experiments.

4. CONCLUSIONS AND FUTURE WORKS

We introduced a prototype model for an episodic memory fit for autonomous robots. It is fully based on direct internal sensor data of a robot and does not use any explicitly modelled events or states. The model was implemented and deployed on the ANYbotics ANYmal C robot and evaluated on both simulated and real data. We were able to show that using the episodic memory allows the robot to remember critical experiences that previously caused damage before the situation itself occurs again. When walking up a ramp that previously led to an error upon reaching the top, the robot was able to recall this memory during traversal of the ramp. This allows the robot to stop execution of the current mission or adapt its behavior accordingly, avoiding a repeated mistake. The similarity in these situations could be detected purely by utilizing internal sensor data of the robot, such as the joint states and currents, without introducing any external knowledge about the environment or the robot's position.

The developed episodic memory can be used as the basis for a complex prediction and mission adaptation framework. As relevant memories are recorded and recalled fully, the relevant elements of an experience do not have to be understood immediately by the robot or its operators - instead the robot can actively reflect on its experiences when necessary. Explicitly learning events instead of learning implicitly through adaptation of parameters also allows for full one-shot learning, which is especially critical when relevant situations occur only very rarely and are dangerous for the robot. Due to the direct usage of internal sensor data, instead of modelled internal states or events, no expert knowledge about the robot is necessary and the model can easily transferred to other robotic systems. The current prototype implementation uses full, unprocessed internal sensor data of the robot, which does not scale well with large numbers of held memories. Due to this, we are extending the data preprocessing component of the model to include an encoder, drastically reducing the amount of recorded and analyzed data, while only losing a small amount of precision in memories. Additionally, we will add external data to the memories of the robot to further increase the contextual information available.

ACKNOWLEDGMENT

The research leading to these results has received funding from the DLR Space Administration under grant agreement No. 50RA2026 by the German Bundestag.

REFERENCES

1. E. Tulving, "Episodic and Semantic Memory," *In E. Tulving and W. Donaldson, Organization of memory*, pp. 381–403, 1972.

2. T. Suddendorf and M. C. Corballis, "Mental time travel and the evolution of the human mind," *Genet Soc Gen Psychol Monogr*, vol. 123, no. 2, pp. 133–167, May 1997.

3. D. Stachowicz and G. J. M. Kruijff, "Episodic-like memory for cognitive robots," *IEEE Transactions on Autonomous Mental Development*, vol. 4, no. 1, pp. 1–16, 2012.

4. K. Kawamura, R. A. Peters II, D. M. Wilkes, W. A. Alford, and T. E. Rogers, "ISAC : Foundations in Human – Humanoid Interaction," *IEEE Intelligent Systems*, no. July/August, pp. 38–45, 2000.

5. W. Dodd and R. Gutierrez, "The role of episodic memory and emotion in a cognitive robot," *Proceedings - IEEE International Workshop on Robot and Human Interactive Communication*, vol. 2005, pp. 692–697, 2005.

6. T. Deutsch, A. Gruber, R. Lang, and R. Velik, "Episodic memory for autonomous agents," *2008 Conference on Human System Interaction, HSI 2008*, pp. 621–626, 2008.

7. J. E. Laird, P. Rosenbloom, and P. S. Rosenbloom, "The Evolution of the Soar Cognitive Architecture," Tech. Rep., 1994. [Online]. Available: https://www.researchgate.net/publication/2719989

8. A. Nuxoll and J. E. Laird, "A Cognitive Model of Episodic Memory Integrated With a General Cognitive Architecture," *Proceedings of the Sixth International Conference on Cognitive Modeling*, pp. 220–225, 2004.

9. J. E. Laird, A. Nuxoll, and M. R. James, "Comprehensive Working Memory Activation in Soar." Tech. Rep., 2004. [Online]. Available: https://www.researchgate.net/publication/221548319

10. N. Derbinsky and J. E. Laird, "Efficiently implementing episodic memory," in *International Conference on Case-Based Reasoning*. Springer, 2009, pp. 403–417.

11. J. Rothfuss, F. Ferreira, E. E. Aksoy, Y. Zhou, and T. Asfour, "Deep episodic memory: Encoding, recalling, and predicting episodic experiences for robot action execution," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 4007–4014, 2018.

12. S. Jockel, M. Weser, D. Westhoff, and J. Zhang, "Towards an episodic memory for cognitive robots," in *Proc. of 6th Cognitive Robotics workshop at 18th European Conf. on Artificial Intelligence (ECAI)*. Citeseer, 2008, pp. 68–74.

13. M. Hutter, C. Gehring, A. Lauber, F. Gunther, C. D. Bellicoso, V. Tsounis, P. Fankhauser, R. Diethelm, S. Bachmann, M. Blösch, *et al.*, "Anymal-toward legged robots for harsh environments," *Advanced Robotics*, vol. 31, no. 17, pp. 918–931, 2017.

14. M. G. Besselmann, L. Puck, L. Steffen, A. Roennau, and R. Dillmann, "Vdb-mapping: A high resolution and real-time capable 3d mapping framework for versatile mobile robots," in *2021 IEEE 17th International Conference on Automation Science and Engineering (CASE)*. IEEE, 2021, pp. 448–454.

15. M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, A. Y. Ng, *et al.*, "Ros: an open-source robot operating system," in *ICRA workshop on open source software*, vol. 3, no. 3.2. Kobe, Japan, 2009, p. 5.

16. T. Schnell, C. Plasberg, L. Puck, T. Buettner, C. Eichmann, G. Heppner, A. Rönnau, and R. Dillmann, "Robot health estimation through unsupervised anomaly detection using gaussian mixture models," in *2020 IEEE 16th International Conference on Automation Science and Engineering (CASE)*. IEEE, 2020, pp. 1037–1042.