

A Knowledge-driven Approach for Human-Robot Collaborative Manufacturing

Francesca Stramandinoli¹, Amit Bhatia¹, Wenping Zhao¹, Brigid Blakeslee¹, Alex Shilov¹, Jeff Mendoza¹, Pradeep Rajendran² and Satyandra Kumar Gupta²

¹Raytheon Technologies Research Center, East Hartford, Connecticut, US

²Center for Advanced Manufacturing, University of Southern California, Los Angeles, California, US

Abstract

Human-Robot Collaboration (HRC) is expected to have a broad impact on manufacturing when human workers and robots can work together safely and efficiently. For HRC to be safe, no unintentional contact should occur between the human and robot. In this paper we focus on a human-robot collaborative manufacturing application, that is, composite sheet layup, as a context for presenting a knowledge-driven approach for ensuring safe and effective HRC. A knowledge-based system is developed to offer a single endpoint that fuses multiple sources of information from various sub-systems (i.e., perception, planning, and execution) into a coherent knowledge representation that can be stored and queried.

Keywords

Ontologies in the perception-action loop, human-robot collaboration, human early action prediction, robotic planning, execution monitoring, knowledge representation and reasoning

1. Introduction

In the manufacturing domain there is a plethora of tasks that require human labor (e.g., wire harness, electronic or aircraft assembly, composite fabrication, etc.). Some of these tasks can be tedious or strenuous for a human to perform. Many of these tasks are difficult – and in some cases, too expensive – to fully automate due to advanced dexterity and flexibility requirements yet increases in production volume and cost remain challenging. By combining robot repeatability and precision with human intelligence, flexibility and dexterity, human-robot collaborative systems can increase productivity and quality in production lines and reduce the physical and cognitive load of human operators. At present, nearly 90% of composite layup processes are performed by manual labor. Therefore, employee safety is a primary concern for manufacturers. Moreover, damage of robots or tools due to errors and collisions incurs significant cost and causes unacceptable production delays. The state-of-the-art safety techniques in Human-Robot Collaboration (HRC) are structured either in a stop-and-go fashion or to slow down the robot movements. Development and implementation of HRC frameworks have the potential to disrupt composite manufacturing with expected benefits of part consistency, quality and cost, and ergonomics. In effective collaborations, team members are not only capable of understanding

RobOntics 2023: Workshop on Ontologies in Autonomous Robotics, August 28, 2023, Seoul, South Korea

✉ francesca.stramandinoli@rtx.com (F. Stramandinoli); amit.bhatia2@rtx.com (A. Bhatia); guptask@usc.edu (S. K. Gupta)



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



CEUR Workshop Proceedings (CEUR-WS.org)

what their partners are currently doing, but they can also estimate the intentions and needs of each other and plan complementary actions accordingly. They also need to maintain appropriate situational awareness about the environment and the overall system. In this paper, we describe a knowledge-driven human-robot collaboration framework for smart manufacturing. The proposed approach combines statically compiled knowledge about the operational environment with runtime perceptual information to estimate the current and future class of an action being performed by a human. A knowledge-based system is developed to offer a single endpoint that fuses multiple sources of information from various sub-systems into a coherent knowledge representation that can be stored and queried. The information on current and future actions is fed to a task planning module, which selects the robot action that best suits the estimated present and future human states and can also inform the collision avoidance system to ensure safety for the human. An ontological framework is implemented to maintain situational awareness about the environment and the system.

2. Related Work

2.1. Human Early Action Prediction

Understanding actions and events from visual temporal data (e.g., RGB/IR/NIR/thermal video or depth map sequences) is valuable in many fields, including surveillance, safety monitoring, autonomous navigation, protocol adherence verification/enforcement, and human-robot collaboration. While extensive literature exists on action and event recognition from full sequences (e.g., data that captures the full action or event), early action and event detection (e.g., from incomplete or non-existing data on the action or event) is less studied due to its higher technical complexity. However, early detection is usually more valuable (and often required) in many applications: for example, for human-robot collaboration to be most effective, team members should be capable of not only understanding what their partners are currently doing, but also estimating their intentions and needs, and planning complementary actions accordingly [1, 2, 3]. Anticipation is crucial in scenarios where an autonomous system needs to react before an action is finalized. We propose an automated method for early action and event detection from incomplete or non-existing visual data, which effectively and efficiently exploits prior knowledge and historical observations. The bulk of the existing literature on action and event detection relies on the availability of full sequences [4, 5, 6, 7, 8, 9]. More recent work deals with online action and event detection, which relies on partial observations (typically of the start of the action or the event) in order to carry out real-time (albeit not predictive) recognition [10, 11]. Representative work on action anticipation includes [12, 13].

2.2. Knowledge Representation and Reasoning

Effective, safe, and compliant human-robot collaboration for complex tasks requires endowing the robots with advanced reasoning capabilities. The reasoning capabilities of an HRC system need to consider not only the nominal operating conditions but also the off-nominal situations that could arise due to faults within the system itself or unforeseen changes in the surrounding environment. For the off-nominal scenarios, the reasoning subsystem within the HRC system

needs to also understand the root-cause of a contingency and take appropriate actions in off-nominal scenarios. The reasoning capabilities depend on the collection of facts and assumptions that are available to reason over. Examples of useful information to encode within a knowledge base are rules of physics, safety constraints, strategies for contingency management, etc. The twin problems of encoding a variety of world knowledge into a computationally amenable form and reasoning over it is collectively referred to knowledge representation and reasoning ([14, 15, 16]). Various paradigms to encode knowledge and reason over it have been proposed over the decades ([16, 17, 18, 15]). Knowledge base construction (KBC) is the process of populating a database with information from data such as text, tables, images, or video. Some prominent examples of large-scale, high-quality knowledge bases (KBs) are Freebase [19], YAGO [20], IBM Watson [21], PharmGKB [22], and Google Knowledge Graph [23]. The KBC process can be human-expert-driven (e.g., the CYC project [24], WordNet project [25], PaleoBioDB [26]) and it can also be automated by leveraging ongoing advances in natural language processing and understanding techniques (e.g., KnowledgeVault [27], DeepDive [28], MinIE [29], NELL [30], Alexandria [31], Fonduer [32]). We refer the reader to [16] for discussion on open challenges in creation of high-quality knowledge bases for various safety critical applications. In the context of robotics, a variety of knowledge-driven robotics frameworks have been proposed over last several years – KnowRob [17], RoboBrain [18], RoboEarth [33], PMK [34], and openEASE [35]. Ontologies offer a powerful paradigm for modeling and encoding knowledge into reusable knowledge pieces [36, 37]. An ontology encompasses a representation, formal naming and definition of the categories, properties and relations between the concepts, data and entities that substantiate one, many or all domains of discourse [36]. Towards the goal of formally defining an ontology, several formats exist. The W3C OWL 2 Web Ontology Language (OWL) is a Semantic Web language designed to represent rich and complex knowledge about things, groups of things, and relations between things [38]. OWL is a computational logic-based language such that knowledge expressed in OWL can be reasoned with by computer programs either to verify the consistency of that knowledge or to make implicit knowledge explicit [38]. OWL 2 is a knowledge representation language, designed to formulate, exchange, and reason with knowledge about a domain of interest through the modeled ontology [38]. A variety of ontologies have been proposed in robotics for autonomy related applications ([35, 39, 40, 41, 42, 43, 33, 18, 44]).

2.3. Planning

Several approaches exist for different kinds of planning problems [45, 46, 47, 48, 49, 50]. For trajectory planning, there exist sampling-based incremental search approaches ([45, 46, 48]), and graph search-based point to point planning techniques [51] as also planning under process constraints [52]. For grasp planning, large number of researchers have worked on grasp planning problems ([53, 54, 55]). The problem of task-allocation and sequencing has been studied with respect to homogeneous agents (only robots) using state-transition diagrams [56] and with respect to heterogeneous agents (humans and robots) [57]. There also exist several high-level task planning formalisms to support planning for human-robot teams [58]. This allows rapid contingency handling by refining plans on the fly if human operator decides to follow a different sequence than the one generated by the system.




Figure 2: Hybrid cell setup: a) design; b) physical setup.

ensured by integrating human early action prediction generated by a visual perception framework into the robot motion planner so that the planner can produce safe motions proactively, instead of relying on frequent re-planning.

4.2. Deep Generative and Discriminative Framework

We propose a perception-based framework (shown in Figure 3) that couples generative and discriminative models, which estimate the current and future classes of an action or an event taking place. The analytics leverage a sequence of skeletal joint positions extracted from the sequence. The generative model ingests current and previous data points and outputs a sequence of predicted future data points. A Long Short-Term Memory (LSTM) network is used, trained by optimizing the mean squared error between the predicted output frames and actual future frames. Once trained, the predicted data from the generative model along with the current and previous data from the sensor is made available to the discriminative model. The discriminative model implements a Convolutional Neural Network (CNN) model and a LSTM network with a SoftMax layer and produces a vector of probabilities indicating the likelihood that the future action or event belongs to a class among a set of classes being considered. The CNN uses dilated convolutional operations to capture relations between pairs of joints, and pairs of pairs of joints. The discriminative model is trained by optimizing a cross-entropy loss between the predicted class probabilities and the actual class categories. The predicted data from the generative model is made available to the discriminative model both during training and at inference.




Figure 3: Schematic view of proposed perception framework.

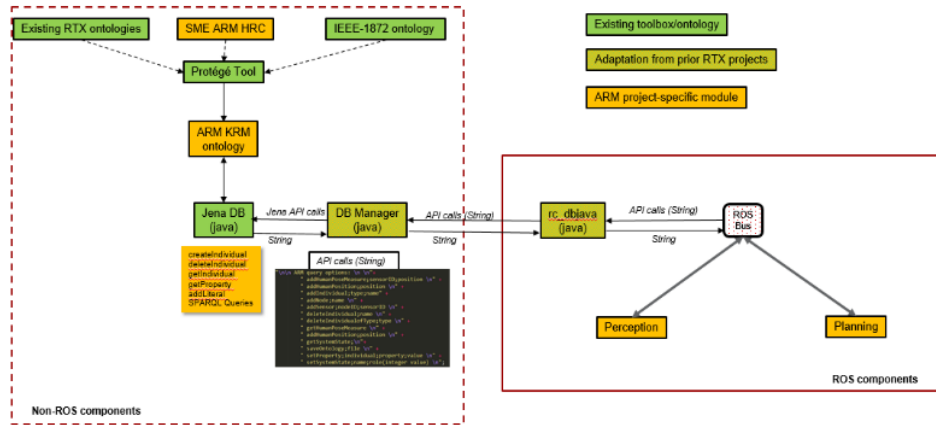


Figure 4: Overall ARM KRM Implementation Architecture.

4.3. Knowledge Representation and Reasoning

We developed a knowledge-reasoning based framework to maintain situational awareness about the environment and the system. An ontology-driven modeling framework for describing concepts, instances and relationships was developed and an ontology-specific reasoning engine, Jena [59], was used to reason over the instantiation of the ontology; A simplified user front-end to formulate and parse the queries was also developed. Towards the goal of defining the ontology, the software architecture, and the reasoning techniques, the team has taken the approach of re-using components and know-how from previous relevant project within the RTX organization. The system is named as The ARM Knowledge Reasoning and Management system (ARM KRM). In the next few paragraphs, we detail and discuss each of the aspects of the ARM KRM. The implementation architecture is shown in Figure 4. The ARM KRM system is used to integrate the perception module with the planning module to enable synchronization on situational awareness. In addition to the ontology itself, there are few other key pieces in the implementation architecture. The Apache Jena (referred to as Jena DB in the implementation architecture shown in Figure 4) is an open-source java library that provides a variety of APIs for instantiation, modification and reasoning over a given ontology [59]. However, the use of Jena requires an expert-level understanding of the API calls and the library itself. As a result, direct interaction with the ontology instance using Jena is prone to errors and requires a steep learning curve. For ARM KRM, we have adapted an existing RTX java implementation, called as the DB Manager that serves as a bridge between the Jena API and the ARM KRM relevant functionalities that are used by planning and perception module. The last key module in ARM KRM implementation architecture is rc_dbjava. This is a ROS implementation of a client-server architecture that sends queries (from Planning and Perception modules) and redirect responses from DBManager to the initiator of a query. For the scope of the project, several existing ontology implementations were studied for reusability and identifying salient features [43, 60, 61]. This included previously developed ontologies within RTX for other use-cases and projects. The team also consulted recently proposed standards specifically for manufacturing robotics [43, 61]. In the end it was determined that the existing ontologies are useful as a guidance but not readily re-usable for the

project. They were either too specific to their end use-case or too abstract. It was instead decided to adapt from prior existing works developed for RTX-internal projects. Various classes and their hierarchies are shown in Figure 5 below.

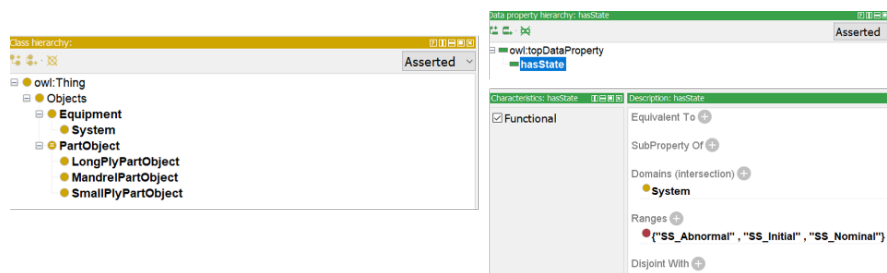


Figure 5: OWL2-based ontology for the system: concepts about state, system and part objects.

As mentioned earlier, the ARM KRM system is used to integrate the perception module with the planning module to enable synchronization on situational awareness. The DB Manager serves as a bridge between the Jena API and the ARM KRM relevant functionalities that are used by planning and perception module. Several functions have been implemented within the DBManager and we will not discuss all of them here. Instead, we highlight one below which provides key functionality regarding situational awareness. String inferState(): This method performs reasoning regarding the system state given current instance of the ontology. The reasoning is implemented through a combination of Jena Rules, Jena Generic Reasoner for forward chaining, SPARQL queries and the ontology instance itself. The Jena rule is a forward implication rule that expresses the condition “If there exist parts of type Small Ply, Large Ply and Mandrel, then the system state is nominal”. This method returns the inferred system state.

4.4. Planning

From planning perspective, there are three agents in the cell (Figure 2). UR5_A, UR5_B and Human Operator (HO) are considered agents in this work. UR5A belongs to the kitting station and the UR5B belongs to the final assembly station. The following items are the workpieces in this work namely, small ply, long ply, stacked-long ply, mandrel, wrapped mandrel. UR5A works with mandrels, long plies and small plies. And UR5B works with wrapped mandrels and stacked-long plies. The term workpiece buffer refers to an object that can store workpieces. In this work, there are many buffers in use. For example, in the kitting cell, a kitting buffer stores the small-plies and the mandrel. Another kitting buffer stores the long plies. The tray buffer stores the workpieces before they are taken by the HO. The sub-assembly buffer stores the wrapped mandrels and the stacked long plies. Finally, the cascade-frame buffer stores the assembled workpieces. For the goals for kitting and assembly cells, a set of primitive tasks needed to accomplish each of the goals was defined. For example, in the kitting cell, UR5_A needs to pick and place the workpieces into the tray. The pick and place task consists of the following primitive tasks: moving to a pose, moving to a grasp pose, moving along a path, moving to a configuration, grasping, and ungrasping. A special task known as wait for signal task is also needed to ensure that tasks can be executed after a certain event signal is received. Similarly, in the assembly cell,

UR5_B needs to transfer the wrapped mandrel from the sub - assembly fixture to the cascade frame. It also needs to collaboratively move the stacked-long ply from the sub-assembly buffer (or fixture) to the cascade frame with the help of the human operator. These operations involve the following primitive tasks: moving to a pose, moving to a grasp pose, moving along a path, moving to a configuration, grasping, and ungrasping. Therefore, both the cells share the same set of primitive tasks. The primitive tasks need to be sequenced together to accomplish the overall goals of each cell. The planning system is divided into two parts: the backend and the frontend. The backend was constructed using the C++ MoveIt API and the frontend was constructed using Python scripts and the Qt GUI framework. The backend is responsible for planning and executing the primitive tasks. The frontend is responsible for sequencing the primitive tasks and accepting user input from the GUI.

4.5. Evaluation

The developed technologies were demonstrated to support collaboration among robots and humans in the layup of the composite cascade prototype. The assessment of the system performance was done with respect to Key Performance Parameters (KPPs) that among other included (i) the ergonomic exposure of human operators, measured as the time spent on repetitive layup process, (ii) the direct labor required, as the number of human operators needed during layup process, and (iii) the throughput rate, as assembly time. The current manual process for the layup of thrust reverser composite cascades was considered as the baseline for the evaluation. Our assessment demonstrated that human ergonomic exposure and direct labor were reduced by 50%. About the layup assembly time, we considered that a human operator would take ~ 4 minutes to create a mandrel sub-assembly and ~ 9 minutes to create a long ply sub-assembly (i.e., baseline). Our objective was to achieve 30% improvement (i.e., 2.8 minutes for the mandrel sub-assembly and 6.3 minutes for the long ply sub-assembly). Through our approach we were able to achieve a total time for sub-assembly kit of 1.2 minutes for the mandrel sub-assembly and 1.6 for the long ply sub-assembly, only $\sim 42\%$ and $\sim 25\%$ of the time objective respectively.

5. Conclusion

In this paper, we have presented a knowledge-driven, human-robot collaboration framework for a smart manufacturing use-case – manufacturing of a thrust reverser composite cascade prototype for aircraft engine nacelles. The proposed framework for addresses the labor-intensive nature of sheet-based composite manufacturing associated with current manual processes for aircraft composite sheet layup. We have developed a visual perception analytics framework to estimate actions being performed by a human and developed a knowledge-based system to store and query knowledge. The planning architecture within the proposed framework determines suitable robot base positions to perform the task, selects the robot action that suits the estimated human actions, and generates low level instructions for the robot to precisely grasp and position objects. An execution monitoring layer tracks the state of the cell for preventing errors/collisions. The performance of the proposed framework has been evaluated with respect to selected Key Performance Parameters (KPPs).

Acknowledgments

Research was sponsored by the ARM (Advanced Robotics for Manufacturing) Institute through a grant from the Office of the Secretary of Defense and was accomplished under Agreement Number W911NF-17-3-0004. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Office of the Secretary of Defense or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

References

- [1] N. Sebanz, H. Bekkering, G. Knoblich, Joint action: bodies and minds moving together, *Trends in Cognitive Sciences* 10 (2006) 70–76. URL: <https://api.semanticscholar.org/CorpusID:1781023>.
- [2] J. Mainprice, D. Berenson, Human-robot collaborative manipulation planning using early prediction of human motion., in: *IEEE International Conference on Robots and Systems (IROS)*, IEEE, 2013, pp. 299–306. URL: <http://dblp.uni-trier.de/db/conf/iros/iros2013.html#MainpriceB13>.
- [3] H. Wang, C. Schmid, Action recognition with improved trajectories, in: *2013 IEEE International Conference on Computer Vision*, 2013, pp. 3551–3558. doi:10.1109/ICCV.2013.441.
- [4] I. Laptev, On space-time interest points, *International Journal of Computer Vision* 64 (2005) 107–123. doi:10.1007/s11263-005-1838-7.
- [5] A. Gaidon, Z. Harchaoui, C. Schmid, Actom sequence models for efficient action detection, in: *CVPR 2011*, 2011, pp. 3201–3208. doi:10.1109/CVPR.2011.5995646.
- [6] M. Jain, J. C. van Gemert, H. Jégou, P. Bouthemy, C. G. M. Snoek, Action localization with tubelets from motion., in: *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE Computer Society, 2014, pp. 740–747. URL: <http://dblp.uni-trier.de/db/conf/cvpr/cvpr2014.html#JainGJBS14>.
- [7] B. Fernando, E. Gavves, J. O. Oramas M., A. Ghodrati, T. Tuytelaars, Rank pooling for action recognition, *IEEE Trans. Pattern Anal. Mach. Intell.* 39 (2017) 773–787. URL: <https://doi.org/10.1109/TPAMI.2016.2558148>. doi:10.1109/TPAMI.2016.2558148.
- [8] H. Bilen, B. Fernando, E. Gavves, A. Vedaldi, S. Gould, Dynamic image networks for action recognition, in: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 3034–3042. doi:10.1109/CVPR.2016.331.
- [9] R. de Geest, E. Gavves, A. Ghodrati, Z. Li, C. G. M. Snoek, T. Tuytelaars, Online action detection, in: *European Conference on Computer Vision*, 2016. URL: <https://ivi.fnwi.uva.nl/isis/publications/2016/GeestECCV2016>.
- [10] Z. Shou, J. Pan, J. Chan, K. Miyazawa, H. Mansour, A. Vetro, X. G. i Nieto, S.-F. Chang, Online action detection in untrimmed, streaming videos - modeling and evaluation, *ArXiv abs/1802.06822* (2018). URL: <https://api.semanticscholar.org/CorpusID:3648714>.
- [11] J. Gao, Z. Yang, R. Nevatia, RED: reinforced encoder-decoder networks for action anticipa-

- tion, in: British Machine Vision Conference 2017, BMVC 2017, London, UK, September 4-7, 2017, BMVA Press, 2017. URL: <https://www.dropbox.com/s/s5yyf1mo8n2f1tq/0284.pdf>.
- [12] C. Vondrick, H. Pirsiavash, A. Torralba, Anticipating visual representations from unlabeled video, 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016) 98–106. URL: <https://api.semanticscholar.org/CorpusID:10533233>.
- [13] O. Oshin, E. Bernal, B. Nair, J. Ding, R. Varma, R. Osborne, E. Tunstel, F. Stramandini, Coupling deep discriminative and generative models for reactive robot planning in human-robot collaboration, in: 2019 IEEE International Conference on Systems, Man and Cybernetics (SMC), 2019, pp. 1869–1874. doi:10.1109/SMC.2019.8913974.
- [14] H. J. Levesque, Knowledge representation and reasoning, Annual Review of Computer Science 1 (1986) 255–287. URL: <https://doi.org/10.1146/annurev.cs.01.060186.001351>. doi:10.1146/annurev.cs.01.060186.001351. arXiv:<https://doi.org/10.1146/annurev.cs.01.060186.001351>.
- [15] D. Paulius, Y. Sun, A survey of knowledge representation in service robotics, Robotics and Autonomous Systems 118 (2019) 13–30. URL: <https://www.sciencedirect.com/science/article/pii/S0921889018303506>. doi:<https://doi.org/10.1016/j.robot.2019.03.005>.
- [16] A. P. Amit Bhatia, Automated construction of knowledge-bases for safety critical applications: Challenges and opportunities, in: AAI Spring Symposium: Combining Machine Learning with Knowledge Engineering, 2021.
- [17] M. Beetz, D. Beßler, A. Haidu, M. Pomarlan, A. K. Bozcuoğlu, G. Bartels, Know rob 2.0 – a 2nd generation knowledge processing framework for cognition-enabled robotic agents, in: 2018 IEEE International Conference on Robotics and Automation (ICRA), IEEE Press, 2018, p. 512–519. URL: <https://doi.org/10.1109/ICRA.2018.8460964>. doi:10.1109/ICRA.2018.8460964.
- [18] A. Saxena, A. Jain, O. Sener, A. Jami, D. Misra, H. S. Koppula, Robobrain: Large-scale knowledge engine for robots, ArXiv abs/1412.0691 (2014). URL: <https://api.semanticscholar.org/CorpusID:14700558>.
- [19] K. Bollacker, C. Evans, P. Paritosh, T. Sturge, J. Taylor, Freebase: a collaboratively created graph database for structuring human knowledge, in: SIGMOD '08: Proceedings of the 2008 ACM SIGMOD international conference on Management of data, ACM, New York, NY, USA, 2008, pp. 1247–1250. URL: <http://portal.acm.org/citation.cfm?id=1376746#>. doi:<http://doi.acm.org/10.1145/1376616.1376746>.
- [20] A. Amarilli, L. Galárraga, N. Preda, F. M. Suchanek, Recent topics of research around the yago knowledge base, in: L. Chen, Y. Jia, T. Sellis, G. Liu (Eds.), Web Technologies and Applications, Springer International Publishing, Cham, 2014, pp. 1–12.
- [21] D. A. Ferrucci, E. W. Brown, J. Chu-Carroll, J. Fan, D. Gondek, A. Kalyanpur, A. Lally, J. W. Murdock, E. Nyberg, J. M. Prager, N. Schlaefer, C. A. Welty, Building watson: An overview of the deepqa project., AI Magazine 31 (2010) 59–79. URL: <http://dblp.uni-trier.de/db/journals/aim/aim31.html#FerrucciBCFGKLMNPSW10>.
- [22] M. Hewett, D. Oliver, D. Rubin, K. Easton, J. Stuart, R. Altman, T. Klein, Pharmgkb: the pharmacogenetics knowledge base, Nucleic acids research 30 (2002) 163–5. doi:10.1093/nar/30.1.163.
- [23] Google, Google knowledge graph, 2012.

- [24] D. Lenat, Cyc: A large-scale investment in knowledge infrastructure, *Communications of the ACM* 38 (1995) 33–38.
- [25] C. Fellbaum (Ed.), *WordNet: An Electronic Lexical Database, Language, Speech, and Communication*, MIT Press, Cambridge, MA, 1998.
- [26] A. K. Behrensmeyer, A. Turner, *Taxonomic occurrences of suidae recorded in the paleobiology database*, 2013.
- [27] X. Dong, E. Gabrilovich, G. Heitz, W. Horn, N. Lao, K. Murphy, T. Strohmman, S. Sun, W. Zhang, Knowledge vault: A web-scale approach to probabilistic knowledge fusion, in: *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, ACM, 2014, pp. 601–610. URL: <https://cs.cmu.edu/~nlao/publication/2014.kdd.pdf>.
- [28] C. Zhang, C. Ré, M. Cafarella, C. De Sa, A. Ratner, J. Shin, F. Wang, S. Wu, Deepdive: Declarative knowledge base construction, *Communications of the ACM* 60 (2017) 93–102. doi:10.1145/3060586.
- [29] K. Gashteovski, R. Gemulla, L. del Corro, MinIE: Minimizing facts in open information extraction, in: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Copenhagen, Denmark, 2017, pp. 2630–2640. URL: <https://aclanthology.org/D17-1278>. doi:10.18653/v1/D17-1278.
- [30] T. Mitchell, B. Kisiel, J. Krishnamurthy, N. Lao, K. Rivard, T. Mohamed, N. Nakashole, E. A. Platanios, A. Ritter, M. Samadi, B. Settles, W. Cohen, R. Wang, D. Wijaya, A. Gupta, X. Chen, A. Saparov, M. Greaves, J. Welling, M. Gardner, Never-ending learning, *Communications of the ACM* 61 (2018) 103–115. doi:10.1145/3191513.
- [31] J. Winn, J. Guiver, S. Webster, Y. Zaykov, M. Kukla, D. Fabian, Alexandria: Unsupervised high-precision knowledge base construction using a probabilistic program, in: *Automated Knowledge Base Construction (AKBC)*, 2019. URL: <https://www.microsoft.com/en-us/research/publication/alexandria-unsupervised-high-precision-knowledge-base-construction-using-a-probabilistic-program/>, best Research Paper Award.
- [32] S. Wu, L. Hsiao, X. Cheng, B. Hancock, T. Rekatsinas, P. A. Levis, C. Ré, Fondue: Knowledge base construction from richly formatted data, *Proceedings of the 2018 International Conference on Management of Data* (2017). URL: <https://api.semanticscholar.org/CorpusID:3627801>.
- [33] M. Waibel, M. Beetz, R. D’Andrea, R. Janssen, M. Tenorth, J. Civera, J. Elfring, D. Gálvez-López, K. Häussermann, J. Montiel, A. Perzylo, B. Schießle, O. Zweigle, R. van de Molengraft, RoboEarth - A World Wide Web for Robots, *Robotics & Automation Magazine* 18 (2011) 69–82. URL: <https://ai.uni-bremen.de/papers/waibel11roboearth.pdf>.
- [34] M. Diab, A. Akbari, M. Ud Din, J. Rosell, Pmk—a knowledge processing framework for autonomous robotics perception and manipulation, *Sensors* 19 (2019). URL: <https://www.mdpi.com/1424-8220/19/5/1166>. doi:10.3390/s19051166.
- [35] M. Beetz, M. Tenorth, J. Winkler, Open-EASE – a knowledge processing service for robots and robotics/ai researchers, in: *IEEE International Conference on Robotics and Automation (ICRA)*, Seattle, Washington, USA, 2015. Finalist for the Best Cognitive Robotics Paper Award.
- [36] S. Staab, R. Studer (Eds.), *Handbook on Ontologies*, *International Handbooks on Informa-*

tion Systems, Springer, 2004.

- [37] F. Neuhaus, On the definition of ‘ontology’, in: S. Borgo, O. Kutz, F. Loebe, F. Neuhaus (Eds.), Proceedings of the Joint Ontology Workshops 2017 Episode 3: The Tyrolean Autumn of Ontology, Bozen-Bolzano, Italy, September 21-23, 2017, volume 2050 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2018. URL: http://ceur-ws.org/Vol-2050/FOUST_paper_13.pdf.
- [38] B. Motik, P. F. Patel-Schneider, B. Parsia, C. Bock, A. Fokoue, P. Haase, R. Hoekstra, I. Horrocks, A. Ruttenberg, U. Sattler, M. Smith, OWL 2 Web Ontology Language: Structural Specification and Functional-Style Syntax, Last Call Working Draft, W3C, 2008. URL: <http://www.w3.org/2007/OWL/draft/owl2-syntax/>, (to be published, may be superseded).
- [39] D. Marco, P. Levi, R. Janssen, M. Molengraft, A. Perzylo, A deliberation layer for instantiating robot execution plans from abstract task descriptions, in: International Conference on Automated Planning and Scheduling, Workshop on Planning and Robotics, 2013.
- [40] S. R. Fiorini, J. Bermejo-Alonso, P. Gonçalves, E. Pignaton de Freitas, A. Olivares Alarcos, J. I. Olszewska, E. Prestes, C. Schlenoff, S. V. Ragavan, S. Redfield, B. Spencer, H. Li, A suite of ontologies for robotics and automation [industrial activities], *IEEE Robotics & Automation Magazine* 24 (2017) 8–11. doi:10.1109/MRA.2016.2645444.
- [41] A. Freitas, D. Schmidt, A. Panisson, F. Meneguzzi, R. Vieira, R. Bordini, Semantic representations of agent plans and planning problem domains, in: Engineering Multi-Agent Systems: Second International Workshop, 2014. doi:10.1007/978-3-319-14484-9_18.
- [42] G. R., V. Uma, Ontology based knowledge representation technique, domain modeling languages and planners for robotic path planning: A survey, *ICT Express* 4 (2018) 69–74. URL: <https://api.semanticscholar.org/CorpusID:196000000>.
- [43] C. Schlenoff, E. Prestes, R. Madhavan, P. Goncalves, H. Li, S. Balakirsky, T. Kramer, E. Migueláñez, An iee standard ontology for robotics and automation, in: 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2012, pp. 1337–1342. doi:10.1109/IROS.2012.6385518.
- [44] R. Provine, C. Schlenoff, S. Balakirsky, S. Smith, M. Uschold, Ontology-based methods for enhancing autonomous vehicle path planning, *Robotics and Autonomous Systems* 49 (2004) 123–133. doi:10.1016/j.robot.2004.07.020.
- [45] S. M. Lavalle, *Planning Algorithms*, Cambridge University Press, 2006.
- [46] I. A. Sucan, M. Moll, L. E. Kavraki, The open motion planning library., *IEEE Robotics Autom. Mag.* 19 (2012) 72–82. URL: <http://dblp.uni-trier.de/db/journals/ram/ram19.html#SucanMK12>.
- [47] S. Karaman, E. Frazzoli, Sampling-based algorithms for optimal motion planning, *International Journal of Robotic Research - IJRR* 30 (2011) 846–894. doi:10.1177/0278364911406761.
- [48] A. Bhatia, M. R. Maly, L. E. Kavraki, M. Y. Vardi, Motion planning with complex goals, *IEEE Robotics & Automation Magazine* 18 (2011) 55–64. doi:10.1109/MRA.2011.942115.
- [49] C. Belta, A. Bicchi, M. Egerstedt, E. Frazzoli, E. Klavins, G. Pappas, Symbolic planning and control of robot motion, *Robotics & Automation Magazine, IEEE* 14 (2007) 61 – 70. doi:10.1109/MRA.2007.339624.

- [50] S. Srivastava, E. Fang, L. Riano, R. Chitnis, S. Russell, P. Abbeel, Combined task and motion planning through an extensible planner-independent interface layer, in: IEEE International Conference on Robotics and Automation, 2014, pp. 639–646. doi:10.1109/ICRA.2014.6906922.
- [51] A. M. Kabir, B. C. Shah, S. K. Gupta, Trajectory planning for manipulators operating in confined workspaces, in: 2018 IEEE 14th International Conference on Automation Science and Engineering (CASE), 2018, pp. 84–91. doi:10.1109/COASE.2018.8560414.
- [52] A. M. Kabir, A. Kanyuck, R. K. Malhan, A. V. Shembekar, S. Thakar, B. C. Shah, S. K. Gupta, Generation of synchronized configuration space trajectories of multi-robot systems, in: 2019 International Conference on Robotics and Automation (ICRA), 2019, pp. 8683–8690. doi:10.1109/ICRA.2019.8794275.
- [53] A. Bicchi, V. Kumar, Robotic grasping and contact: A review, in: IEEE International Conference on Robotics and Automation 1, 2000.
- [54] J. Bohg, A. Morales, T. Asfour, D. Kragic, Data-driven grasp synthesis—a survey, IEEE Transactions on Robotics 30 (2014) 289–309. doi:10.1109/TRO.2013.2289018.
- [55] N. B. Kumbla, S. Thakar, K. N. Kaipa, J. Marvel, S. K. Gupta, Handling Perception Uncertainty in Simulation-Based Singulation Planning for Robotic Bin Picking, Journal of Computing and Information Science in Engineering 18 (2018) 021004. URL: <https://doi.org/10.1115/1.4038954>. doi:10.1115/1.4038954. arXiv:https://asmedigitalcollection.asme.org/computingengineering/article-pdf/18/2/021004/6102753/jcise_018_02_021004.pdf.
- [56] M. Bonfè, N. Preda, C. Secchi, F. Ferraguti, R. Muradore, L. Repele, G. Lorenzi, P. Fiorini, Distributed control architecture for automated surgical task execution with coordinated robot arms, IFAC Proceedings Volumes 47 (2014) 10213–10218. URL: <https://api.semanticscholar.org/CorpusID:55944553>.
- [57] B. Hayes, B. Scassellati, Challenges in shared-environment human-robot collaboration, Learning 8 (2013). URL: <https://api.semanticscholar.org/CorpusID:14217706>.
- [58] C. Morato, K. Kaipa, S. Gupta, Improving assembly precedence constraint generation by utilizing motion planning and part interaction clusters, Computer-Aided Design 45 (2013) 1349–1364. doi:10.1016/j.cad.2013.06.005.
- [59] A. S. Foundation, Apache jena, 2021.
- [60] S. R. Fiorini, Ieee1872-owl, 2022.
- [61] A. Olivares-Alarcos, D. Beßler, A. Khamis, P. Goncalves, M. Habib, J. Bermejo, M. Barreto, M. Diab, J. Rosell, J. a. Quintas, J. Olszewska, H. Nakawala, E. Pignaton, A. Gyrard, S. Borgo, G. Alenyà, M. Beetz, H. Li, A review and comparison of ontology-based approaches to robot autonomy, The Knowledge Engineering Review 34 (2019) e29. URL: <https://ai.uni-bremen.de/papers/olivares19ker.pdf>. doi:10.1017/S0269888919000237.