Online Inference of Robot Navigation Parameters from a Semantic Map

Benjamin Kisliuk[1], Christoph Tieben[1], Nils Niemann[1] Christopher Bröcker[1], Kai Lingemann[1] and Joachim Hertzberg[1,2]

[1]*German Research Center for Artificial Intelligence, Plan-based Robot Control Group, Osnabrück, Germany*

[2]*Universität Osnabrück, Knowledge Based Systems Group, Osnabrück, Germany*

Keywords: Semantic Map, Navigation, Robot Autonomy, Semantic Navigation.

Abstract: Agriculture is becoming one of the key application fields for mobile robots. At the same time it poses serious challenges for true autonomous systems due to its heterogeneous and dynamic nature. To act robustly and reliably, robotic behaviour needs to be controlled by an intelligence, making explainable and informed decisions based on knowledge of its surroundings. However, this knowledge cannot only be derived from sensor data but has to be based on prior knowledge and external sources as well to comprehensively represent a robots deployment site. By representing this knowledge in formal and thus machine readable way, automated inference improves the handling of the complex nature of these requirements. In this paper, we show how quantitative and qualitative control parameters regarding a mobile robots navigation can be derived from a manually modelled semantic map of an agricultural deployment site. Also we describe how such a system can be integrated into a typical ROS system architecture. By making the derived knowledge easily available, the robotic system is enabled to dynamically adapt route planning on an agricultural deployment site and to switch between different local planning algorithms according to situational and prior knowledge.

1 INTRODUCTION

In agriculture, robots are widely considered an important building block to the future of a more sustainable source of food and primary production with various projects demonstrating success, mostly in singular applications (Bergerman et al., 2016).

Automation can help streamlining processes, and autonomous machines can enable processes which are too complex to be managed efficiently by humans. Current research projects often focus on multi-robot systems, which act and decide partially or completely without human supervision (Shamshiri et al., 2018).

As of today, most commercially available robots are able to execute certain agricultural processes, but still lack the capability of an integrated approach (Bergerman et al., 2016; Shamshiri et al., 2018). This can be partially attributed to the complexity of the domain and the multitude of conditions and processes required to run an agricultural business; but also the agricultural sites, i.e. farms, are usually very heterogeneous spaces, which are – from a roboticist's point of view – far from well structured in a robotics sense and dynamic by nature (Kunze et al., 2018; Egerstedt et al., 2018). Thus, a robot which is supposed to run autonomously will be challenged to deal flexibly with a lot of problems and steep constraints regarding reliability and robustness while behaving and deciding in a comprehensible and explainable way (Langley et al., 2017).

In practice, integrating robotic software often means that it is up to the roboticist to decide which approach or existing implementation of a certain technology – like navigation or task-planning and their respective parameters – fit the overall constraints best. A flexible approach allowing the robot itself to automatically choose between different approaches according to which fits best in a specific situation can improve the ability of a robot to interact with a complex environment enormously like the *Move Base Flex* framework for navigation (Putz et al., 2018). From that flexibility however, the question arises how to choose the best algorithm or control parameters for a given situation or context.

This paper demonstrates the usage of a reasoning engine over a semantic map of the geospatial environment as a key component to tackle challenges of autonomous robot control. We will show how semantic knowledge about the area the robot is deployed in – the *deployment area* – is modelled by annotating

Kisliuk, B., Tieben, C., Niemann, N., Bröcker, C., Lingemann, K. and Hertzberg, J.
Online Inference of Robot Navigation Parameters from a Semantic Map.
DOI: 10.5220/0010790200003116
In *Proceedings of the 14th International Conference on Agents and Artificial Intelligence (ICAART 2022) - Volume 1*, pages 156-163
ISBN: 978-989-758-547-0; ISSN: 2184-433X

a geospatial map with semantic knowledge. The resulting *semantic map* in combination with information sources like weather data, the current time of day and the robots position can be used to infer control parameters online. For representation we use a custom software called *SEMPR* with an integrated rule based reasoner using the RETE pattern-matching algorithm (Forgy, 1982).

We will show as an example how 2D navigation cost maps can be generated from the inferences and used for path planning in a robotic system.

After this introduction to the problem the second chapter will provide information about the theoretical background and state of the art. In the third chapter we describe the implementation of the necessary components and show how they integrate into a ROS robot architecture in chapter four. Lastly, we summarize the system and show prospective future avenues.

2 BACKGROUND

In modern times mobile robot control architectures commonly base a robots behaviour on the input of the sensors and the results of some deliberation process. However, not every relevant information can or will be sensed by the robot: In the agriculture domain properties like ground moisture, dust or the quality of daylight often will not be sensed due to the lack of specific sensors or appropriate software modules. Some types of information cannot be derived from on-board sensory data at all, e.g. the legal situation regarding ownership of a plot or a weather forecast. This motivates the use of an environment representation which covers information about the deployment site which can be enriched with prior knowledge. To further enhance the autonomy of the robot control software, this knowledge should be expressed in a formal way which allows it to be interpreted by a reasoning software.

In (Hoellmann et al., 2020) a simple approach to handle different contexts which cannot or will not be sensed by the robot is described: By dividing an overall heterogeneous deployment site into subareas with in themselves higher homogeneity like single fields, yard areas, buildings or pasture areas, *zones* with more specific constraints were defined. Based on the context defined by a zone, control parameters and navigation algorithms were chosen according to predefined parameters. However, the limitations of this approach became clear quite early. While the context based approach was able to tackle the heterogeneity of the deployment site, it is not flexible enough to regard dynamic changes or complex interconnections. A more comprehensive approach can be taken by predefining only singular facts, making them changeable at runtime and using a rule based reasoning system to infer the control parameters for the robot at run time. This can be understood as three main components: The semantic map which incorporates geospatial anchored knowledge, a set of formalised rules and a reasoning algorithm.

From intuition, this approach offers multiple advantages: Independent knowledge can be modelled independently. This allows for discussing relevant information about a deployment area with experts in the domain using non-technical terms in the interviews. Also the complexity of the system architecture can be reduced drastically, as new facts can easily extend the knowledge base as do new rules or relations. Lastly, every decision that such a system makes can be easily explained to a human user by extracting the chain of rules and facts which led to an inference.

2.1 Semantic Maps

Enriching spatial maps for robots goes back at least to the works of Kuipers, Buschka in the early 00s (Kuipers, 2000; Buschka, 2005). Galindo et al suggest a hierarchy of map representations (spatial, semantic) and ”anchor” semantic meaning to objects and places (Galindo et al., 2005). Nüchter and Hertzberg showed how semantic annotation can be automatically applied to 3D SLAM Maps by reasoning over previous knowledge in form of a constraint network to enhance scene interpretation and automate annotation (Nüchter and Hertzberg, 2008). Pronobis reflects on place classification and the generation of semantic maps and the enhancement of object search. The mapping process is defined broadly as associating spatial concepts (e.g. ”kitchen”) with spatial entities (e.g. a polygon on map) (Pronobis, 2011). Lang and Paulus formally define a semantic map as a hybrid map of spatial and semantic information with the stress of the semantic information being represented in a way which allows for inference (Lang and Paulus, 2014). At the same time Kunze et al use reasoning to enhance object recognition (Kunze et al., 2014). In a survey Kostavelis and Gasteratos claim that most works regarding semantic mapping describe robotic indoor scenarios, note however that the use of semantic maps for robotic applications increases (Kostavelis and Gasteratos, 2015). Deeken et al generate low level occupancy grids from a semantic map (Deeken et al., 2015) while other works apply semantic maps to navigation in form of land mark recognition (Cosgun and Christensen, 2018). In recent years, Kunze et al anticipate a strong yet growing trend for semantic

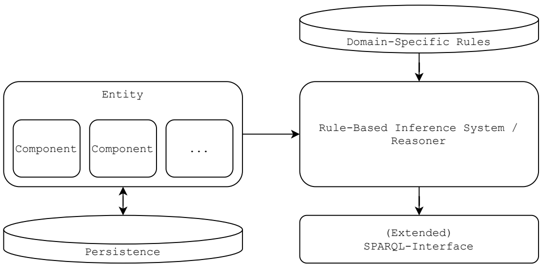

Figure 1: Overview of the system architecture of SEMPR.

maps for long term autonomy in their survey (Kunze et al., 2018). A rule based reasoner is used by Deeken et al to infer process phases in agricultural machines (Deeken et al., 2019). Recently, Crespo et al surveyed a growing number of approaches to semantic navigation by using semantic maps with reasoning methods. The focus, however, still lies on using semantic information for human-robot-interaction and exploiting high level information for room recognition and classification (Crespo et al., 2020).

In this work we use rule based reasoning over a semantic map of an agricultural deployment site to infer not where the robot should go but how to act on the way. Like Deeken et al generated occupancy grids (Deeken et al., 2015), in our work we generate traversal costs for the whole deployment site which enables a more precise approach to path planning. In contrast to the descriptions of indoor scenarios the focus rests not on objects, but on the properties of free space areas.

3 SYSTEM DESIGN

3.1 Inferring Robot Control Parameters

The framework we developed to hold the representation and infer implications for robot control and navigation is called *Semantic Environment Mapping, Processing and Reasoning* (SEMPR)[1] and added as a library to the robots system.

Its **Architecture**, also shown in Fig. 1, comprises of a collection of *Entities* which make up the known facts about the environment, the domain specific rules which define how to infer new knowledge and a reasoning system to perform the automated inference. Interfaces to a persistence layer and a *SPARQL*-Query-Service enable further usability.

The **Entities** themselves are made up by a set of components which encode arbitrary data like geometries, semantic information or transformations. However, all semantic information is described according to the *Resource Description Framework* which is a well known format for machine understanding and reasoning (W3C, 2004).

The **Rules** facilitate inference in the form if-then. The inferred information is added to the knowledge base and can be used to activate robot behaviour, set control parameters or be used in chained rules. Knowledge of different types can be combined to enable reasoning over geospatial and semantic information alike as seen in the example below:

```
[robotInZone:
    (?robot <type> <Robot>),
    (?zone  <type> <Zone>),
    Geometry(?zone  ?zoneGeo),
    Geometry(?robot ?robotGeo),
    geo:intersects(?zoneGeo ?robotGeo)
    ->
    (?robot <inZone> ?zone)]
```

Which translates to *For each entity $R \in$ Robots and each entity $Z \in$ Zones, retrieve the respective geometries G_R, G_Z. For each combination where G_R intersects G_Z, add the fact that R is in zone Z.*

3.1.1 Reasoner

In order to decide which rules need to be triggered at a time, the inference system makes use of the RETE pattern matching algorithm (Forgy, 1982): The textual representation of the rules is parsed into a graph where every node implements a small check on the given data. If the check succeeds, the data is forwarded to its child nodes. The connections between the nodes thus construct the complex conditions as stated in the rules, while the terminal nodes implement the rules' consequences. By also inserting *memory-nodes*, the pattern matching implements a trade-off between performance and memory consumption, as they enable an iterative processing of changes to the knowledge base. Facts that match a set of conditions are stored in the memory nodes and can be used as partial matches in subsequent conditions and effects without re-evaluation. Furthermore, when retracting facts, only branches with memory nodes containing them need to be re-evaluated. The whole graph consists of two parts: The `alpha-network`, in which the basic data elements are inspected independently of each other, and the `beta-network`, in which multiple conditions get combined and more complex checks can be performed.

[1] https://github.com/sempr-tk/sempr

3.2 Robot Control Architecture

For our setup we implement a ROS node called `sempr_ros_bridge` from which a SEMPR instance is created and managed. The semantic map and rules are loaded from pregenerated files on startup. Fig. 2 shows, how the navigation modules `move_base_flex` and `waypoint_server` are connected. The `sempr_ros_bridge` node polls updates from the *SEMPR* instance in a specified time interval via the *SPARQL Protocol And RDF Query Language (SPARQL)* interface. The resulting information is then converted into typical ROS message formats and either published or introduced to the system in form of online reconfiguration requests. Additional information that the robot could not sense on its own, such as weather data, can be entered manually directly into the knowledge base to test the system.

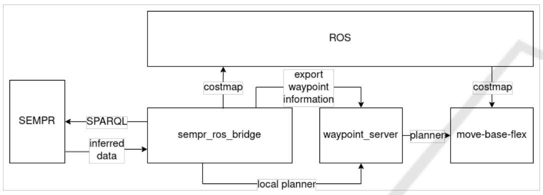

Figure 2: Integration of the reasoning module into a ROS based system.

3.3 Semantic Parameters for Robot Control

In our work we concentrated on some central aspects of robot control and navigation to be inferred by the rule engine.

3.3.1 Global Path Costs

One very well researched aspect of mobile robots is the problem of planning a path from one spatial location to another in an optimal manner according to certain variables. Common criteria to be optimized by a path planning algorithm are the length of the path, power consumption and the time of traversal, but depending on the domain other considerations might be of interest: In agriculture, the distribution of strain to the soil by heavy machinery can be very important, as well as to route traffic in a specific way to ascertain smooth workflows. In robotic practice, different surface properties can have a huge impact on the drivability of the area. A common way to model these is the utilization of costmaps: Different cost values are assigned to discrete grid cells of an area as a function of the optimization criteria. By assigning different cost values to discrete grid cells of an area as a function of the optimization criteria, a shortest path planning algorithm such as *A** or *Dijkstra's* can easily optimize for lowest overall path costs instead of geometric length only. In that way, more costly areas such as grass or soil are only crossed by the robot if the routes are significantly shorter than routes only traversing better suited surfaces like pavement. Naively, it is possible to assign traversal costs in this way to all areas of the map and thus model the deployment environment of the robot. However, such a static approach meets its limits as soon as the input parameters to the cost function change. For example, the drivability of unsealed surfaces such as soil or grassland change drastically depending on whether the surface is wet or dry. By modelling static knowledge about the surface properties of the environment and according rules, we can infer the costs to traverse a certain area during runtime and thus generate a costmap dynamically with the logical rules modelling the function to calculate the costs of traversing a given area given the input variables at the time. In our experiment we defined traversal costs for discrete values of *very low*, *low*, *medium*, *high* and *very high*. Those values then would be assigned to regions of the map according to rules, taking into account the surface type and overall ground wetness to demonstrate the dynamic generation of a 2D costmap.

3.3.2 Fences and Gates

On a farm where animals are kept and moved between different places, fences and gates play a large role. While only few gates can be opened by the robot itself right now, it is very plausible that there might be automated systems opening the gates mechanically integrated in the future. At the moment, however, it would already be useful if the robot knew which gates were open or closed at a given time. In practice some gates' open/closed state or probabilities thereof can in fact be derived from certain knowledge like the time of day or information about the whereabouts of the animals. If a gate is known or reasoned to be open it can be represented in the static map layout, which updates whenever the state of a gate changes.

3.3.3 Local Path Planning Algorithms and Driving Speeds

Typical navigation stacks found in robots based on ROS use a local planning module in addition to the global path planning. It takes into account locally observed obstacles, the kinematics of the robot as well as strategies to account for smooth motor control. State of the art are multiple approaches of high quality local planners such as *dwa* (Fox et al., 1997), *teb* (Roesmann et al., 2012) and *eband* (Quinlan and Khatib, 1993). However, experience shows that different planners perform differently in different environments. So, in a heterogeneous environment it is obviously best to choose the planner depending on the local surroundings. To enable this choice, we modelled the criteria of *dynamics* which we define as a scalar between 0.0 and 1.0 to represent the presence of moving objects in the area the robot is going to cross. On the extreme value of 1.0, possibly contrary to intuition, it can make sense to choose an algorithm which sticks closer to the globally planned path and resolves situations of being blocked by waiting or reevaluation of the obstacle map instead of trying to locally evade obstacles which possibly move quicker than the robot itself. A dynamic value of 0.0 means the robot can expect to be the only moving object in the vicinity, which in turn means every obstacle sensed will probably be static and should be smoothly evaded using a fitting local planning algorithm.

As a second input parameter we defined the *freespace* value to represent the structuredness and openness of an area. For robotic practice a wide open field with few to no obstacles blocking sight means that a local planner can optimize for smooth movements freely without taking moving obstacles into account that could appear behind occlusions suddenly. Vice versa, a cluttered environment means that the robot should try to stick closely to the global path and possibly reduce driving speeds to accommodate the fact that obstacles might only become visible when already close to the robot. These control parameters, like maximum velocity and the used planning algorithm, can be inferred from the semantic state of the environment for the different zones with respect to their aforementioned properties. Their actual geometric shape and the robots location can also be taken into account during the reasoning process, as shown earlier.

Figure 3: Zones as modelled in QGIS.

Table 1: Example data after the postprocessing steps.

id	groundType	freespace	dynamic	animals
1	gravel	0.25	0.4	false
2	gravel	0.7	0.8	false
3	gravel	0.8	0.75	true
4	gravel	0.35	0.2	false
5	gravel	0.15	0.2	false
6	gravel	0.6	0.75	true

3.4 Modelling of the Semantic Map

To model such a semantic map we model the knowledge as a collection of facts about the domain and a set of rules which defines allowed inferences. On a symbolic level this means to have a set of entities representing physical or virtual objects of the environment. Those are associated primarily with geospatial geometric information as a lot of relevant information in the agricultural domain revolves around spatial locations, areas and their extents. We used the open source tool *QGIS*[2] to model zones as geospatial vector polygons in ESRI shape-layers shown in Fig. 3. Each zone constitutes a primary semantic entity. The file format allows to annotate a layer with a table as seen in Tab. 1 with each row representing a zone. These are used to model semantic facts which make up RDF triples with the row as subject, the column name as predicate and the value of the cell corresponding to row + column as object. Rules are defined as a pair of two lists: One for the preconditions and a second for the assertions. The rules are then stored in a separate, human readable text file. The shapes and rules are then imported into the semantic representation and reasoning module described in section 3.1.

As mentioned before in 3.1, the rules can be considered to represent logical functions of the input values which means they can be derived from expert knowledge using conventional acquisition techniques like case studies, simulations or structured and unstructured interviews. An advantage of using a rule based approach can be seen in the fact that rules can be naturally formulated as in *"traversal costs increase for certain surface types when the individual wetness threshold is reached"* and then easily translated into the machine-readable form.

[2]www.qgis.org

4 RESULTS

Fig. 4 shows two costmaps generated online by the system. On the top is a generated costmap of the areas if the condition *dry* is met as a global condition. On the bottom a costmap is generated for a *moist* condition met globally.

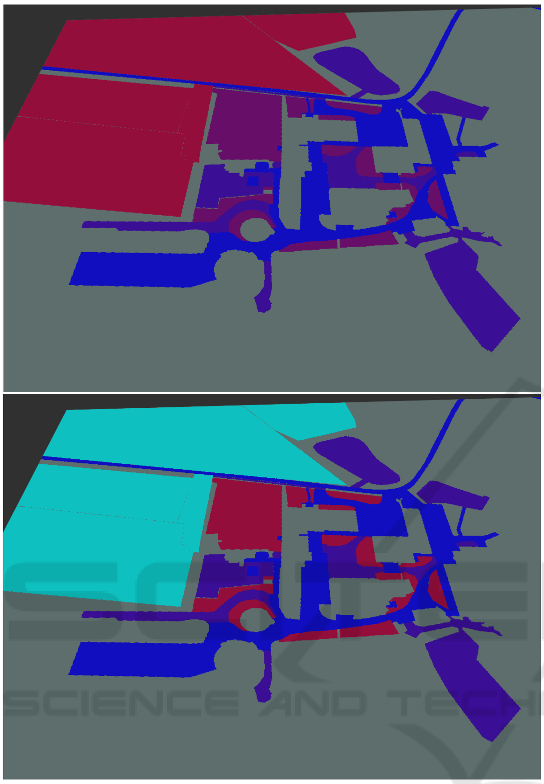

Figure 4: Inferred Costmaps: blue, purple, red, turquoise (low – high). Top: Costs for dry farm condition. Bottom: Costs for wet farm condition.

As can be seen in fig. 5 the generated costs can be incorporated in the global costmap of the ros navigation system and used online for path planning. This way the navigation behaviour of the robot can be adapted online according to updates from the reasoning process.

Furthermore, the navigation stack of this robotic system is able to switch between different local planning algorithms online which additionally offer an interface to update certain parameters such as maximum speed and acceleration online via *dynamic reconfigure*. For example, the data used by the context navigation algorithm shown in (Hoellmann et al., 2020) can be generated by the semantic control setup. This way it is possible for the robot to directly infer important information dynamically during its deployment. The ability to generate costmaps dynamically and thus react to changes in weather or other conditions can represent real world conditions much more accurately than any static approach. Likewise can the inference of choices like local planning algorithms or movement parameters open up new approaches of intelligent robot behaviour. To use a semantic reasoning system instead of implicit knowledge encoded in the programming of the robot allows to identify rules and facts with the fields respective experts using naturally formulated rule syntax. Decisions made by the robotic system in this way can be easily explained and documented due to the explicit modelling of the knowledge base.

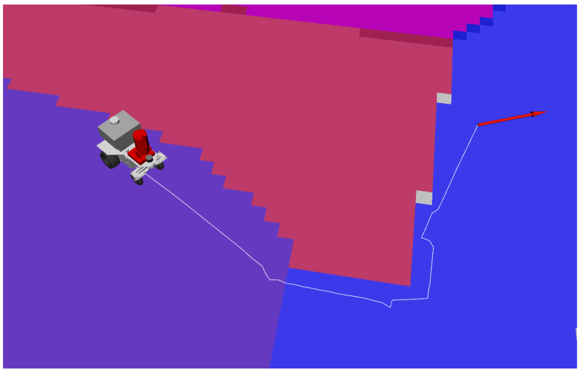

Figure 5: Robot planning a path according to costmap.

5 OUTLOOK

We showed how a semantic map, that is a representation of a geographical area with georeferenced subareas, can be annotated in a way to allow efficient inference using known tools like QGIS together with a reasoning system for geospatial and semantic information based on the well known RETE algorithm. With the integration into the ROS ecosystem, this can be used for the online change of parameters essential for robotic applications, thus increasing the flexibility with which a robot can adapt to changes in a complex environment like agriculture. Relevant rules and facts can be defined using conventional knowledge acquisition techniques with experts. In agriculture it makes sense to organize those rules and facts around geospatial information, thus being easily approachable to discuss with experts like farmers.

A further advantage of such a system is to define knowledge in an explicit way, enabling the robot to document an explanation for each of its decisions and profiting of sensing information as well as of prior knowledge and even inferring from knowledge that cannot or will not be sensed by the robot.

Future work will include the integration of the inferred information in a live robotic system, replacing existing solutions to look up context based control parameters with the demonstrated inference system. It appears promising to extend the control parameters to infer, e.g., covariances for localization filters depending on the traction to expect on certain surfaces or additional behavioral strategies like acoustic or optical signalling when expecting to act in the vicinity of humans or animals. Also strategies like not optimizing for shortest paths but instead following right hand rules when driving along paths or streets might be beneficial. In a wider perspective the extension towards probabilistic reasoning appears sensible to cope with information not easily conveyed with simple facts.

The generation of fine-grained navigation costmaps provides the foundation for further work: As a future avenue we plan to use the semantic representation to map detected obstacles and annotate additional information. In the long term there are many kinds of obstacles which might move in the scale of minutes, hours or sometimes days. Instead of just adding them to the costmap and remove them, once not seen anymore, it might be better to actively check once a certain amount of time has passed or not regard them for path planning.

Using semantic reasoning technologies can be an important contribution to add to the flexibility and thus robustness of robots expected to act in complex environments without human supervision. Making knowledge about the environment explicit can add to the explainability of artificial intelligent decisions made by robots as well as to the ease in identifying relevant rules and facts with the help of experts in the respective field like agriculture.

ACKNOWLEDGEMENTS

The DFKI Niedersachsen Lab (DFKI NI) is sponsored by the Ministry of Science and Culture of Lower Saxony and the VolkswagenStiftung. The paper describes work that has been developed in the context of the funded projects Experimentierfeld Agro-Nordwest (BMEL, 28DE103E18), DAKIS (BMBF, 031B0729B) and ZLA (NiMWK, Volkswagenstiftung, ZDIN 11-76251-14-3/19).

REFERENCES

Bergerman, M., Billingsley, J., Reid, J., and Henten, E. V. (2016). Robotics in Agriculture and Forestry. In Siciliano, B. and Khatib, O., editors, *Springer Handbook of Robotics*, chapter 56, pages 1463–1492. Springer.

Buschka, P. (2005). *An investigation of hybrid maps for mobile robots*. PhD thesis, University Örebro.

Cosgun, A. and Christensen, H. I. (2018). Context-aware robot navigation using interactively built semantic maps. *Paladyn, Journal of Behavioral Robotics*, 9(1):254–276.

Crespo, J., Castillo, J. C., Mozos, O. M., and Barber, R. (2020). Semantic Information for Robot Navigation: A Survey. *Applied Sciences*, 10(2):497.

Deeken, H., Wiemann, T., and Hertzberg, J. (2019). A spatio-semantic approach to reasoning about agricultural processes. *Applied Intelligence*, 49(11):3821–3833.

Deeken, H., Wiemann, T., Lingemann, K., and Hertzberg, J. (2015). SEMAP - a semantic environment mapping framework. In *2015 European Conference on Mobile Robots (ECMR)*, pages 1–6. IEEE.

Egerstedt, M., Pauli, J. N., Notomista, G., and Hutchinson, S. (2018). Robot ecology: Constraint-based control design for long duration autonomy. *Annual Reviews in Control*, 46:1–7.

Forgy, C. L. (1982). Rete: A fast algorithm for the many pattern/many object pattern match problem. *Artificial Intelligence*, 19(1):17–37.

Fox, D., Burgard, W., and Thrun, S. (1997). The dynamic window approach to collision avoidance. *IEEE Robotics Automation Magazine*, 4(1):23–33.

Galindo, C., Saffiotti, A., Coradeschi, S., Buschka, P., Fernandez-Madrigal, J., and Gonzalez, J. (2005). Multi-hierarchical semantic maps for mobile robotics. In *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2278–2283. IEEE.

Hoellmann, M., Kisliuk, B., Krause, J. C., Tieben, C., Mock, A., Sebastian, P., Igelbrink, F., Wiemann, T., Martinez, S. F., Stiene, S., and Hertzberg, J. (2020). Towards Context-Aware Navigation for Long-Term Autonomy in Agricultural Environments. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Las Vegas.

Kostavelis, I. and Gasteratos, A. (2015). Semantic mapping for mobile robotics tasks: A survey. *Robotics and Autonomous Systems*, 66:86–103.

Kuipers, B. (2000). The Spatial Semantic Hierarchy. *Artificial Intelligence*, 119(1-2):191–233.

Kunze, L., Burbridge, C., Alberti, M., Thippur, A., Folkesson, J., Jensfelt, P., and Hawes, N. (2014). Combining top-down spatial reasoning and bottom-up object class recognition for scene understanding. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2910–2915. IEEE.

Kunze, L., Hawes, N., Duckett, T., Hanheide, M., and Krajnik, T. (2018). Artificial Intelligence for Long-Term Robot Autonomy: A Survey. *IEEE Robotics and Automation Letters*, 3(4):4023–4030.

Lang, D. and Paulus, D. (2014). Semantic Maps for Robotics. In *Proceedings of the Workshop on AI Robotics at the IEEE International Conference on Robotics and Automation*.

Langley, P., Meadows, B., and Sridharan, M. (2017). Explainable agency for intelligent autonomous systems.

In *IAAI '17 proceedings of the twenty-ninth AAAI conference on innovative applications of artificial intelligence*, pages 4762–4763.

Nüchter, A. and Hertzberg, J. (2008). Towards semantic maps for mobile robots. *Robotics and Autonomous Systems*, 56(11):915–926.

Pronobis, A. (2011). *Semantic Mapping with mobile robots*. PhD thesis, KTH.

Putz, S., Santos Simon, J., and Hertzberg, J. (2018). Move Base Flex A Highly Flexible Navigation Framework for Mobile Robots. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3416–3421. IEEE.

Quinlan, S. and Khatib, O. (1993). Elastic bands: connecting path planning and control. In *Proceedings IEEE International Conference on Robotics and Automation*, pages 802–807 vol.2.

Roesmann, C., Feiten, W., Woesch, T., Hoffmann, F., and Bertram, T. (2012). Trajectory modification considering dynamic constraints of autonomous robots. In *ROBOTIK 2012; 7th German Conference on Robotics*, page 74–79.

Shamshiri, R. R., Weltzien, C., A. Hameed, I., J. Yule, I., E. Grift, T., K. Balasundram, S., Pitonakova, L., Ahmad, D., and Chowdhary, G. (2018). Research and development in agricultural robotics: A perspective of digital farming. *International Journal of Agricultural and Biological Engineering*, 11(4):1–11.

W3C (2004). Resource Description Framework. https://www.w3.org/RDF/.