A Spatio-Semantic Model for Agricultural Environments and Machines

Henning Deeken[1,2,3], Thomas Wiemann[1], and Joachim Hertzberg[1,3]

[1] Knowledge-Based Systems Group, University of Osnabrück, Germany
firstName.lastName@uni-osnabrueck.de
[2] CLAAS E-Systems KGaA mbH & Co KG, Dissen a.T.W., Germany
firstName.lastName@claas.de
[3] DFKI Robotics Innovation Center, Osnabrück Branch, Germany
firstName.lastName@dfki.de

Abstract. Digitization of agricultural processes is advancing fast as telemetry data from the involved machines becomes more and more available. Current approaches commonly have a machine-centric view that does not account for machine-machine or machine-environment relations. In this paper we demonstrate how to model such relations in the generic semantic mapping framework SEMAP. We describe how SEMAP's core ontology is extended to represent knowledge about the involved machines and facilities in a typical agricultural domain. In the framework we combine different information layers – semantically annotated spatial data, semantic background knowledge and incoming sensor data – to derive qualitative spatial facts about the involved actors and objects within a harvesting campaign, which add to an increased process understanding.

Keywords: semantic mapping, environment modeling, ontologies, agriculture

1 Introduction

Digitization of agricultural processes currently concentrates on recording and processing telemetry data from individual machines to support precision farming. This implicitly leads to a machine-centric view on the ongoing processes. But many agricultural processes are complex, cooperative orchestrations of multiple machines. Automatic decision support in harvesting campaigns is still limited in assistance systems, as representations of cooperative agricultural processes and tools to analyze inter-machine relations are mostly missing.

Information on the whole process can not be derived from a single machine's telemetry data, but is covert in the combined telemetry of multiple machines. To embed this abstract data from different machines in the context of the ongoing process, machine data has to be fused with additional knowledge and information about the environment and the process itself. Most importantly, symbolic representations of the spatial relations between agricultural machines and their environment are needed to identify and monitor process states and associated events. Analyzing the geo location of individual machines and processing of spatial relations between them is therefore a valuable contribution to automated process managing in agriculture. Modern agricultural machines already provide a geo-referenced stream of telemetry data, based on RTK-GPS. The positional data is often used to inspect the containment of machines in polygonal boundaries representing fields and farms, to spatially locate machines at those facilities. Such a quantitative, geometric analysis already extracts a lot of relevant information, but does not account for qualitative relations between the machines and facilities nor for knowledge representation and reasoning on a semantic level.

Representing such spatial relations in terms of a well-defined semantic terminology allows to infer complex facts, built up from basic spatial relations to take a process-centric view on harvesting campaigns. This requires a machine-readable environment model that can be paired with geo-referenced telemetry-data from agricultural machines to geolocalize individual machines and derive spatial relations between machines and their environment, respectively. To meet these requirements, we use the semantic mapping framework SEMAP [1] to represent an agricultural domain. We show how to create a semantic environment model for agricultural environments and machines and how to connect it to the underlying geometric model. We illustrate how to ground qualitative spatial relations between a static environment and a set of dynamic vehicles with SEMAP.

In an application example, we replay telemetry of a harvesting campaign to continuously update the spatio-semantic environment model to derive symbolic facts about the ongoing process. Via rule-based inference we analyze the domain-specific spatial relations of a maize harvesting campaign to detect events such as the correct positioning of a transport vehicle next to the harvester for overloading.

2 Related Work

State of the art solutions in digital agriculture allow to record and process telemetry data of agricultural machines like position, velocity, and internal parameters like fuel consumption or mass throughput [2]. This data is used in precision farming to optimize the application of fertilizers or herbicides, and collected in farm management information systems to aggregate telemetry data to analyze the performance of agricultural machines [3, 4]. They also help to plan agricultural operations by maintaining information about crop rotations [5] or by creating field boundaries and sub-plots based on GPS data [6] to support the application of fertilizers and herbicides tillage strategies [7]. Automated scheduling of entire harvesting campaigns is also possible [8]. Usually, these solutions operate on centralized systems with web-based front ends [9]. This often causes severe latencies due to connectivity issues in remote or rural areas [10].

Fleet overview applications inform the operators about an on-going harvest operation by exchanging telemetry information between machines in real time and display vehicle positions on a static 2D map. Process-related decision making is still completely in the operator's hands, as these assistance systems do not provide a context-dependent and process-oriented analysis. To automatically detect relevant situations that give insight into the agricultural process – e.g., an empty transport vehicle arriving at the field ready for overloading – is a key feature to increase process transparency, which is necessary for improving agricultural efficiency through more process-oriented decision support systems.

To solve these problems, existing approaches from semantic mapping in robotics can be transferred to this application domain. Semantic maps are representations that in addition to spatial data provide assignments to known concepts for the mapped entities, such that semantic background knowledge can be used to reason about the environment [11]. Recent advances in semantic mapping are concerned with constructing general models of multi-modal environment data that can be flexibly queried for task-specific data in individual applications, see [12] for an overview.

Being able to analyze spatial relations in terms of qualitative predicates is important in data retrieval and reasoning. To fully utilize qualitative spatial reasoning, it is necessary to derive qualitative symbolic data from quantitative metric information. In [13], Wolter and Wallgrün pointed out that this process of qualification is essential for qualitative spatial reasoning in practical applications, but still rarely seen. The lack of qualification is also apparent when working with semantic maps. Tools for performing spatial analysis on quantitative metric data are also seldom used in semantic mapping. In our previous work [1], we showed the advantages of maintaining environment data in form of a generalized and persistent model, from which task-specific semantic maps can be extracted, rather than maintaining and aligning several different layers of semantic, geometric and topological information in parallel. We proposed to pair spatial databases and declarative knowledge bases to combine ontological and logical rule-based inference with spatial querying and analysis capabilities and called it the semantic mapping framework SEMAP.

In this paper, we integrate an ontology for agricultural processes into SEMAP to make knowledge about harvesting campaigns accessible for automatic analysis. We use this knowledge together with SEMAP's spatial reasoning capabilities to recognize relevant events in an maize harvesting process. In the presented experiment we were able to detect the correct positioning of an overloading vehicle based on recorded telemetry in an real life harvesting campaign.

3 The SEMAP Framework

The SEMAP framework is designed to represent and manage spatio-semantic environment data. Its purpose is to provide information about the objects and the environment in a specific application domain. It connects conceptual knowledge about the environment and factual knowledge about present object instances with their geometric representations to hold a combined spatio-semantic model that allows spatial analysis as well as semantic inference. To manage the fundamentally different structure of semantic and spatial information, SEMAP internally separates environment data into two dedicated databases to ensure op-

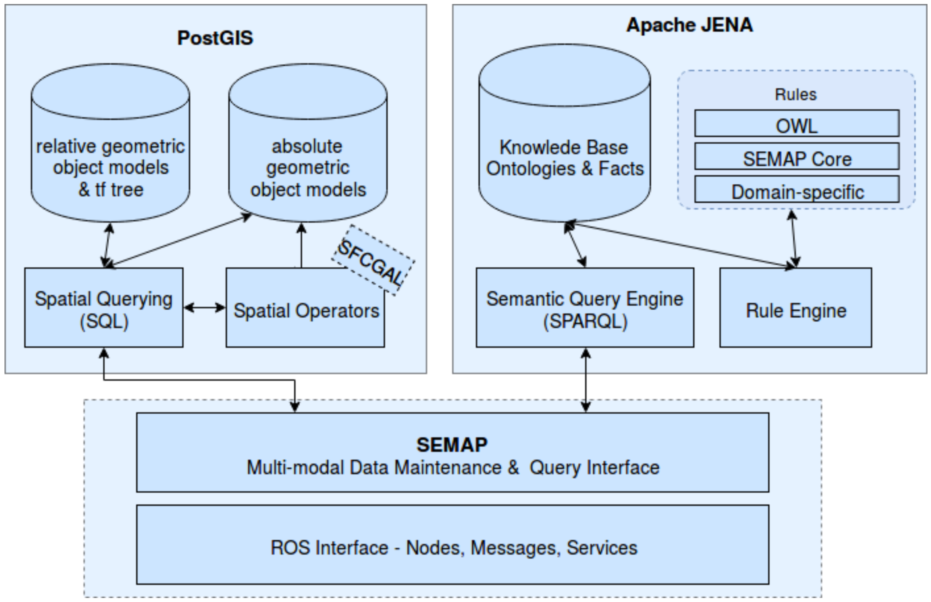

Fig. 1: SEMAP's architecture features a spatial database and a knowledge base system, which are combined by a multi-modal querying interface.

timized performance for each data modality especially in terms of data storage and retrieval. An outline of SEMAP's internal structure is given in Fig 1. The semantic part is represented by a knowledge base system component (KB) that is based on description logics with the obligatory separation into terminological and asserted knowledge. The environment's conceptual model and facts about the environment are represented in the Web Ontology Language (OWL) [14] and maintained in Apache JENA, which provides inference for ontological and rule-based reasoning as well as the capability to query the stored knowledge. The spatial part is a dedicated spatial database system (DB) that stores geometric primitives, and provides operators for quantitative spatial analysis and spatial querying. It is implemented as an extension to PostGIS using the SFCGAL plugin to create custom spatial operators, especially for detecting 3D spatial relations.

The framework's strength lies in combining both query systems to support combined queries with semantic and spatial aspects. In such queries, SEMAP utilizes the DB's spatial operators to ground qualitative spatial relations that are only stored implicitly in the geometric environment representation. Such relations are automatically inserted into the KB as facts for further inference. This approach enables rule-based reasoning and to construct complex spatial queries based on simpler deductions. This multi-modal query interface is advantageous in real-world applications, as it allows to answer complex questions about the positions, relations and roles of the stored objects in a natural way.

The framework's core components are designed to be domain-independent, yet extensible with domain-specific semantic models, rule-sets and geometries. A more detailed description of the SEMAP framework and its spatial querying capabilities is given in [1].

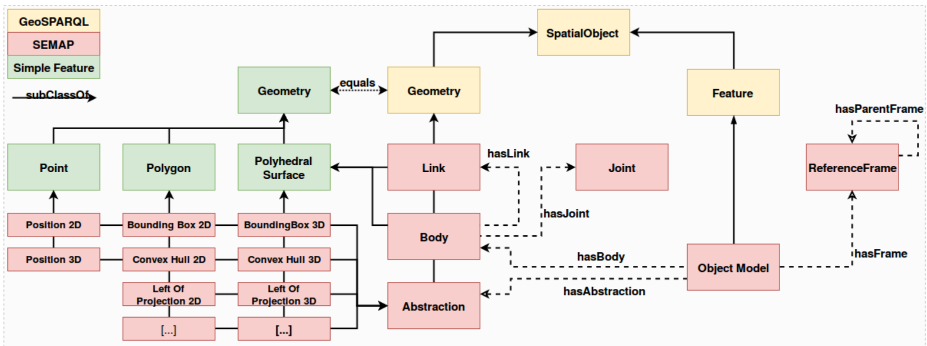

Fig. 2: An excerpt of the ontology that implements the semantics of SEMAP's environment model.

Fig. 2 sketches SEMAP's core ontology. It uses standards from the Open Geospatial Consortium (OGC), because these well-defined models of geo-spatial data are in alignment with PostGIS's data types, which were also defined by the OGC. GeoSPARQL's `SpatialObject` and the fundamental distinction between geometries and features are integrated in SEMAP's upper ontology.

Here, the concept `Geometry` describes any kind of spatial primitive and provides a semantic wrapper for all OGC data types and serves as a bridge to the well known Simple Feature Ontology. SEMAP's KB contains a corresponding `Geometry` sub-concept, for every geometric primitives stored in SEMAP's DB. The property `semap:hasDbId` is used to create an associative link between the geometric primitive and its semantic wrapper. SEMAP internally uses these associations to join spatial and semantic data.

The super-concept `Feature` is used for all things that can be described spatially like SEMAP's `ObjectModel`, which aggregates sets of semantically wrapped geometries to represent an object. For this, it uses the `geo:hasGeometry` property and its two specializations: `semap:hasBody` composes a set of geometries that constitute the object's actual body. In case of articulated objects, the `Link` and `Joint` concepts are used to describe the object's kinematics. `semap:hasAbstraction` provides a set of coarser representations, like oriented and axis-aligned bounding boxes and convex hulls. These abstractions are used for accelerated spatial processing and enable the analysis of directional relations like `left-of` or `above-of`, based on projection and half space geometries described in [15].

To create a spatio-semantic environment model for a particular application, domain-specific ontologies, knowledge bases and rule-sets can be imported into SEMAP. To describe domain-specific concepts spatially and reason about them as part of SEMAP's environment model, the respective entities can be associated with an `ObjectModel` via the `semap:hasObjectModel` relation, cf. Fig. 4 (b).

4 Applying SEMAP in Agriculture

In this section, we detail the process of customizing SEMAP for a specific application domain. Our goal is to create a spatio-semantic model of agricultural environments and machinery in SEMAP for spatial analysis and rule-based reasoning to derive more information about ongoing agricultural processes that involve multiple machines.

First, we present the description of the semantic model used to represent agricultural concepts, such as fields, farms and tractors in SEMAP's knowledge base. After that we discuss how spatial data is added to this model and how to continuously update the environment model by using telemetry data from actual agricultural machines. Finally, we make use of SEMAP's capabilities to ground spatial predicates to answer both spatial and semantic queries. We demonstrate how to analyze basic spatial predicates between agricultural machines and their environment and how rule-based reasoning is used to identify complex and domain-specific spatial relations. The demonstration scenario is the detection of the correct positions of multiple machines in the planned process, especially the correct positioning of a transport vehicle ready for overloading in a maize harvesting campaign.

4.1 The AgriCo Ontology

Our semantic model extends the logistics core ontology (LogiCo) by Daniele et al. [16]. This semantic model describes environments and resources in logistics. Since this domain is very similar to the general process of harvesting, we extended LogiCo with additional concepts needed to represent agricultural processes. We call this extended ontology AgriCo as depicted in Fig. 3.

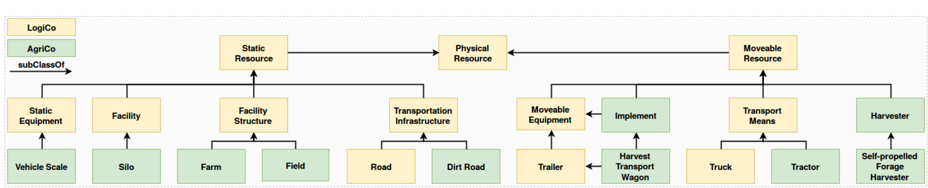

Fig. 3: Excerpts of the domain-specific model added to SEMAP. The LogiCo ontology (yellow) provides a model of static and movable resources, to which the AgriCo ontology (green) adds agricultural concepts like farms and tractors.

All components of our model are based on `Physical Resource`s in the real world, which can be `Static` or `Movable Resource`s. Three sub-classes are used to describe static locations of interest: The `Facility` concept defines areas and structures designated for a specific purpose in the given domain and the `Facility Structure` defines aggregates of different facilities. In AgriCo, for example, `Farm` serves as an aggregate of agricultural facilities like `Silo`s. Additionally, the `Static Equipment` concept describes utilities available at a facility, e.g., a `Vehicle Scale` for weighing transport vehicles. Another important sub-class of static resources are the different kinds of `Transportation Infrastructure` to represent connections between locations. Since this important concept was missing in the LogiCo ontology, we added this concept and suitable sub-classes like `Roads` and `Dirt Roads`.

For movable resources, LogiCo gives concepts for `Transport Means`, i.e., trucks, and `Movable Equipment` such as trailers. AgriCo defines `Tractor`s as another kind of transportation and the `Implement` concept to account for various kinds of machinery that can be connected to a tractor for example plows, sowers or specialized `Harvest Transport Wagon`s. The latter inherit properties from the trailer and implement concept, e.g., to denote the volumetric capacity vie the `logico:hasCapacity` attribute or describe the interfaces use to control the active pickup systems and scraper floor via `agrico:hasISOBUSInterface`. Furthermore, we added the `Harvester` to represent combine and forage harvesters, which are directly derived from the `Movable Resource` concept, as they can not be used for transporting goods in a supply chain.

4.2 Instantiating the Environment Model

The semantic model presented so far provides the conceptual basis from which instances of agricultural facilities and machinery can be created and described. To link them to a spatio-semantic data sets in SEMAP, we proceeded as follows:

First, we imported the AgriCo ontology into SEMAP's KB component. Next, we allowed that the `hasObjectModel` property can map from instances of LogiCo's `Physical Resource` to SEMAP's `ObjectModel`s. This way, the domain-specific concepts and instances thereof can have a spatial representation in SEMAP. Finally, we instantiated the agricultural concepts and their spatio-semantic representation with an appropriate data set.

To setup static resources in our environment model, we used a set of polygonal boundaries to represent farms and fields and other facilities. Fig. 4 (a) shows an excerpt of the environment. It consists of the farm's grounds (blue), three silos (orange) and a vehicle scale (violet), as well as two fields (green). The data was modeled in Google Earth and automatically read into SEMAP's KB and DB components using a KML file importer. In Fig. 4 (b), the underlying semantic representation is depicted with three instances of AgriCo concepts related to their object representation using the `hasObjectModel` relation. Here `farm1` connects to `farm1_obj`. The polygonal boundary `farm1_boundary` is connected via the `hasConvexHull2D` property, which is a sub-property of `hasAbstraction`.

To add movable resources to the static environment, we created three dimensional and articulated object models of a tractor-trailer combination and a forage harvester as displayed in Fig. 6 (b). These objects are modeled in the Unified Robot Description Format, since SEMAP supports this format natively. The underlying semantic representation is a straight forward extension to the example in Fig. 4 (b), yet more complex due to the individual links and joints.

To introduce dynamics to our spatio-semantic model of farms and fields, we used telemetry data recorded on real agricultural machines to continuously update the position and articulation of the machines within it. We replayed the machines GPS signals and joint states in the Robot Operating System (ROS) and connected a bridge node to SEMAP, such that the environment model was updated accordingly.

(a) The spatial data used to represent a farm (incl. silos) and two fields.

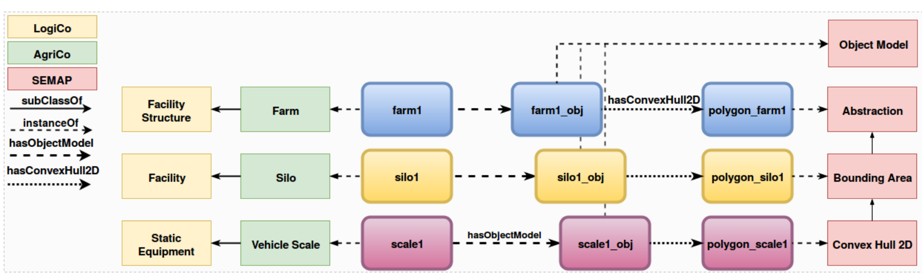

(b) The semantic representation within SEMAP's knowledge base.

Fig. 4: To represent a farm's facilities in SEMAP, we used the 2D polygonal boundaries, shown in (a), stored in the DB component. These spatial model are connected to instances of the domain-specific concepts of AgriCo via SEMAP's `ObjectModel` concept, as illustrated in (b).

5 Application Example

By moving the agricultural machines through the static environment in our experimental setup, the spatial relations between environment and machines and the machines themselves are changed continuously. SEMAP's spatial and semantic reasoning capabilities can be used to detect these spatial relations, which gives insight into the agricultural process underlying the machine activities.

For example, to detect where a movable resource is located topologically, we check whether its 2D position is spatially in a facility's boundary. The derived spatial predicate semap:isIn2D is used to infer that the topological relation `logico:isAt` holds, too. The reasoning takes place in two steps: First, we make use of SEMAP's qualification capabilities to ground spatial relations between agricultural machines and the environment or between pairs of agricultural machines. To perform such a quantitative spatial analysis, a suitable query is posed to SEMAP's DB backend. Fig. 5 (a) gives an example how to query for object pairs for whose 2D convex hulls a containment relation holds. The derived results are then inserted into SEMAP's knowledge base as qualitative semantic knowledge about the spatial relations. In case of our example, the objects pairs found by the query are inserted as facts over the `isIn2D` relation.

Second, we use the derived knowledge in order to reason about more complex spatial relations or to derive domain-specific information. An example for such rule-based inference is given in Fig. 5 (b). This rule identifies the topological relation of a movable resources being at a facility, by using the 2D spatial containment relation for grounding the `isAt` predicate.

```
rosservice call /containment_query
"reference_object_types: ['Facility'] reference_object_geometry_type: 'ConvexHull2D'
target_object_types: ['MovableResource'] target_object_geometry_type: 'ConvexHull2D'
fully_within: false insert_kb: true"
```

(a) SEMAP query to extract containment relations.

```
?machine rdf:type logico:MovableResource
?machine semap:hasObjectModel ?machine_obj
?machine_obj semap:hasPosition2D ?machine_abstr_pos2D
?facility rdf:type logico:Facility
?facility semap:hasObjectModel ?facility_obj
?facility_obj semap:hasConvexHull2D ?facility_abstr_ch2D
?machine_abstr_pos2D semap:isIn2D ?facility_abstr_ch2D
==>
?machine logico:isAt ?facility
```

(b) Rule to ground topological relations based on spatial relations.

Fig. 5: To geometrically ground spatial containment relations, we used the query shown in (a). The query results where extracted into SEMAP's KB as facts over the `isIn2D` relation and then used the rule (b) to derive that the topological relation `isAt` holds between machines and facilities.

While this seems a simple transition, it is important to note that this rule infers from a *spatial* predicate to a *topological* relation and that based on a grounding in the quantitative *geometric* data. Furthermore, the rule is generic for all instances of `Movable Resource` at any instance of `Facility` and its sub-concepts, which makes it applicable in a wide range of applications. The underlying spatial querying is also done automatically in SEMAP's multi-modal query interfaces, such that further queries to the environment model can be posed using the high-level relation `isAt`, without having to deal with the data transfer from DB to KB explicitly. This is convenient during application development.

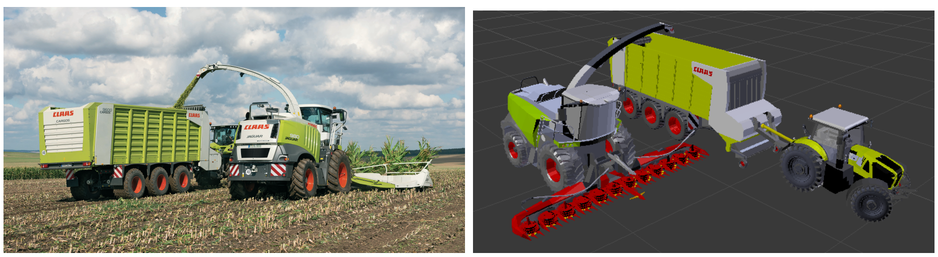

(a) Overloading in reality. (b) Overloading in RViz.

```
?sfh rdf:type agrico:Harvester
?sfh semap:hasObjectModel ?sfh_obj
?sfh_obj semap:hasLeftOfProjection2D ?sfh_proj_l2D
?tv rdf:type agrico:TransportVehicle
?tv semap:hasObjectModel ?tv_obj
?tv semap:hasConvexHull2D ?tv_abstr_ch2D
?tv_abstr_ch2D semap:isIn2D ?sfh_proj_l2D
==>
?tv agrico:positionedForOverloading ?sfh
```

(c) The rule for grounding the `positionedForOverloading` relation in SEMAP.

Fig. 6: We used telemetry data from an actual overloading procedure (a), to move and articulate the machines in ROS and visualize them in RViz (b). We also synchronized the telemetry with our SEMAP model and used the rule (c) to identify the correct spatial positioning of two machines for overloading harvested goods from a forage harvester onto a transport vehicle.

Fig. 6 exemplifies how to combine several basic spatial relations with domain-dependent knowledge to construct complex domain-specific relations.

For example, we used SEMAP to detect that a transport vehicle (TV) is correctly positioned for an overloading procedure, due to the directional relations of the self-propelled forage harvester (SFH). Fig. 6 (a) depicts such a situation in real life, whereas (b) shows visualization of a similar scene represented in SEMAP. It shows the object models subject of the rule shown in (c). To identify that the transport vehicle is properly positioned for overloading, the rule checks the trailer's 2D convex hull for containment in the harvester's `left-of` projection, to verify that the transport vehicle is left-of the harvester. If so, the relation `positionedForOverloading` is inferred to hold between the transport vehicle and the harvester. This is valuable information about the underlying agricultural process, which was previously covert in the telemetry data of both machines, but due to SEMAP's spatio-semantic processing is now explicitly available within SEMAP's KB, where it can be used for further processing.

6 Conclusion and Future Work

In this paper we used the SEMAP framework for combined spatial and semantic reasoning about machine-environment and machine-machine relations in an agricultural domain. To create a semantic model of agricultural environments and machines, we extended an ontological model from the logistics domain resulting in the agricultural core ontology AgriCo. Based on this semantic model, we instantiated a data set that combined factual knowledge with spatial data in our framework. Using recorded telemetry data, we moved and articulated several agricultural machines to replay a forage maize harvesting campaign. We used SEMAP's spatial operators for quantitative spatial analysis to classify containment relations between fields and machines. Using rule-based reasoning over the identified relations, we were able to detect process states relevant to analyze the harvesting process, namely that a transport vehicle is ready for overloading due to its position relative to the harvester.

Our approach demonstrated that the use of semantic mapping technology in agriculture is beneficial, as we were able to extract valuable information about the agricultural process out of the geo-referenced stream of telemetry data. The derived knowledge about machine-machine and machine-environment relations is validated in the geometric state of the environment and also available as machine-readable facts that adhere to a formal ontological model, which opens up possibilities for the further development of decision support systems.

To further improve SEMAP's spatio-semantic querying, temporal information must be included, too. Currently, the data model is updated continuously to represent the environment's current state, but provides neither a history of past states, nor methods to query about temporal change. This denies the possibility to detect events by querying the temporal sequence of certain relations and states. Adding a temporal information layer to SEMAP will be a necessary next step to realize temporal analysis. For this, stream reasoning approaches like the Continuous SPARQL framework (CSPARQL) [17] could be used.

Acknowledgment

Work by Deeken is supported by the German Federal Ministry of Education and Research in the SOFiA project (Grant No. 01FJ15028). The DFKI Osnabrück branch is supported by the state of Niedersachsen (VW-Vorab). The support is gratefully acknowledged.

References

1. Deeken, H., Wiemann, T., Lingemann, K., Hertzberg, J.: Semap-a semantic environment mapping framework. In: European Conference on Mobile Robots (ECMR), 2015, IEEE (2015) 1–6
2. Steinberger, G., Rothmund, M., Auernhammer, H.: Mobile farm equipment as a data source in an agricultural service architecture. Computers and electronics in agriculture **65**(2) (2009) 238–246
3. Pfeiffer, D., Blank, S.: Real-time operator performance analysis in agricultural equipment. In: 73rd International Conference on Agricultural Engineering (AgEng). (2015) 6–7
4. Steckel, T., Bernardi, A., Gu, Y., Windmann, S., Maier, A., Niggemann, O.: Anomaly detection and performance evaluation of mobile agricultural machines by analysis of big data. In: 73rd International Conference on Agricultural Engineering (AgEng). (2015) 6–7
5. Dury, J., Garcia, F., Reynaud, A., Bergez, J.E.: Cropping-plan decision-making on irrigated crop farms: A spatio-temporal analysis. European Journal of Agronomy **50** (2013) 1–10
6. Lauer, J., Richter, L., Ellersiek, T., Zipf, A.: Teleagro+: Analysis framework for agricultural telematics data. In: 7th ACM SIGSPATIAL International Workshop on Computational Transportation Science. IWCTS '14, ACM (2014) 47–53
7. Sørensen, C.G., Nielsen, V.: Operational analyses and model comparison of machinery systems for reduced tillage. Biosystems engineering **92**(2) (2005) 143–155
8. Amiama, C., Pereira, J.M., Castro, A., Bueno, J.: Modelling corn silage harvest logistics for a cost optimization approach. Computers and Electronics in Agriculture **118** (2015) 56–65
9. Kaloxylos, A., Groumas, A., Sarris, V., Katsikas, L., Magdalinos, P., Antoniou, E., Politopoulou, Z., Wolfert, S., Brewster, C., Eigenmann, R., et al.: A cloud-based farm management system: Architecture and implementation. Computers and Electronics in Agriculture **100** (2014) 168–179
10. Mark, T.B., Whitacre, B., Griffin, T., et al.: Assessing the value of broadband connectivity for big data and telematics: Technical efficiency. In: 2015 Annual Meeting, January 31-February 3, 2015, Atlanta, Georgia, Southern Agricultural Economics Association (2015)
11. Nüchter, A., Hertzberg, J.: Towards semantic maps for mobile robots. Robotics and Autonomous Systems (2008)
12. Kostavelis, I., Gasteratos, A.: Semantic mapping for mobile robotics tasks: A survey. Robotics and Autonomous Systems **66** (2015) 86–103
13. Wolter, D., Wallgrün, J.O.: Qualitative spatial reasoning for applications: New challenges and the sparq toolbox. IGI Global (2010)
14. Bechhofer, S.: Owl: Web ontology language. In: Encyclopedia of Database Systems. Springer (2009) 2008–2009
15. Borrmann, A., Rank, E.: Topological operators in a 3d spatial query language for building information models. In: In Proc. of the 12th Int. Conf. on Computing in Civil and Building Engineering (ICCCBE). (2008)
16. Daniele, L., Ferreira Pires, L.: An ontological approach to logistics. In: Enterprise Interoperability, Research and Applications in the Service-oriented Ecosystem, IWEI 13, ISTE Ltd, John Wiley & Sons, Inc. (2013)
17. Barbieri, D.F., Braga, D., Ceri, S., Della Valle, E., Grossniklaus, M.: C-sparql: Sparql for continuous querying. In: Proceedings of the 18th international conference on World wide web, ACM (2009) 1061–1062