An Ontology of Robotics Science

John Hallam[1] and Herman Bruyninckx[2]

[1] Mærsk Institute, University of Southern Denmark, Campusvej 55, DK–5230 Odense M, Denmark john@mip.sdu.dk
[2] Dept Mechanical Engineering, Katholieke Universiteit Leuven, Celestijnenlaan 300B, B-3001 Leuven, Belgium. Herman.Bruyninckx@mech.kuleuven.be

Summary. This paper describes ground-breaking work on the creation of an *ontology* for the domain of *robotics as a science*. An ontology is a collection of terms, concepts and their inter-relationships, represented in a machine-usable form. An ontology in a particular domain is useful *if* the structure it adds to the domain is simple enough to be understood quickly and intuitively, *and* rich enough to increase insight into the whole domain to a level where this increased insight can lead to innovation and increased efficiency in scientific and practical developments.

This paper presents an ontology for *the science of robotics*, and not for *robots as objects*: the latter ontology describes the physical and technical semantics and properties of individual robots and robot components, while the ontology of the science of robotics encodes the semantics of the *meta-level concepts and domains* of robotics. For example, *surgical robotics* and *industrial automation* are two concepts in the ontology of the science of robotics, while the semantics of robot kinematics and dynamics, or of a particular robot control algorithm belong to the ontology of robots as objects.

The structure in the presented ontology for the science of robotics consists of two complementary sub-structures: (i) the *robot agent* and *robot system* models (i.e., what components are required in a robot device, and in a robotic application, respectively), and (ii) the *Context Space* (i.e., ordinal and categorical *relationships* between physical and computational aspects of robot agents and systems in which the sub-domains of robotics can be mapped out). The implementation of the Context Space concept using standard ontology tools is explained.

The paper illustrates its expected usefulness with examples of sub-domains of robotics expressed as contexts in the context space, and with two *use cases* for the ontology: (i) the classification of conference or journal paper submissions, and (ii) the guidance of new researchers into the domain of robotics.

1 Introduction

1.1 Robotics as a Science

Robotics is to a large extent a *science of integration*, constructing (models of) *robotic systems* using concepts, algorithms and components borrowed from

H.I. Christensen (Ed.): European Robotics Symposium 2006, STAR 22, pp. 1–14, 2006.

various more fundamental sciences, such as physics and mathematics, control theory, artificial intelligence, mechanism design, sensor and actuator technology.

The function and properties of a robotic system depend on the components from which it is made — the specific sensors, actuators, algorithms, mechanism — but, beyond that, they depend on the way those components are integrated. Furthermore, a full description of a robotic system must include information about the task it is to perform and the environment in which the task is to be undertaken. As we shall see, this view of robotics (i.e., the system formed by the robot agent, its task and the environment in which it has to perform that task) is crucial in developing a full ontology for robotics.

1.2 Ontologies for Robotics

An ontology is a formal definition of concepts, terms and relationships appropriate to some domain of knowledge, generally expressed in a formalism that allows machine-usability of the encoded knowledge. (The Wikipedia contains a good introduction to the concept of an ontology, [3].) Ontologies typically serve two purposes:

1. providing agreed and unambiguous terminology for a domain, with the goal of helping *humans* express, transform and transfer their knowledge more effectively and accurately.
2. allowing *automatic* use of that knowledge, for instance in exchange or linking of data between processes, or translation of terms between languages.

The latter purpose is sometimes better known under the name *Semantic Web*, [1], and the W3C have defined standards for computer-readable representation of ontologies, [2].

Robotics, like other sciences, has need of suitable ontologies. For instance, an agreed ontology for sensor data would make possible much greater interoperability between sensing and other modules in a robotic system. Individual modules could label the data they generate using terms from the agreed ontology; other processes would then have available a definite semantics for the data. Similar applications can be envisaged in the domains of mechanical/electronic and control/software engineering in robotics. Such ontological knowledge would also be of use for designers of robotic systems when selecting and matching components.

We describe this kind of ontology as an *ontology of robots as objects*. While a complete ontology of robotics requires such knowledge, the source of that knowledge is largely in the sciences and technologies on which robotics builds, viz. physics, control engineering, and so on. Such an ontology may be extremely valuable, but is not the focus or purpose of this paper. However, the ontology presented in this paper could be connected to these object ontologies without problem.

Robotics, in its guise as a science of integration, implies a second kind of ontology. Consider a typical robotics conference, at which one may find sessions or talks on "field robotics" or (perhaps) "factory automation". Roboticists use such labels as a taxonomy of their field, and have a more-or-less clear, but informal, understanding of the significant problems, opportunities and special requirements of each sub-field: of the *context* for the particular sub-domain. So, in the context of papers presented at conferences, even the short labels "field robotics" or "factory automation" are sufficient to help reviewers and readers focus their attention to the contributions of the paper in the topics that are supposed to be difficult and relevant in the state of the art in those domains. It takes students and practitioners in robotics quite some time to learn the links between the short labels on the one hand, and the scientific challenges and contributions they can expect in papers labelled in that way. This paper makes a first contribution in making these links more *explicit* and *structured*, in order to not only speed up the learning phase of human practitioners but also to facilitate (semi) automatic support for selection of reviewers, categorization of papers or research proposals, web or library searches, etc.

The starting point of the presented research is thus that there is knowledge about the *structure* of the robotics science itself — what sub-domains are there; how are they inter-related; how can we express the relationships between sub-domains in a principled way? Once again, the idea of an ontology can help. We call an ontology dealing with this knowledge an *ontology of robotics science*. The existence and formalisation of *that* ontology is the key insight and focus of the present paper.

1.3 Organisation of the Paper

In the following sections we motivate, illustrate and explore the idea of an ontology of robotics science. Section 2 considers robotic systems in general: comprising the agent, the task and the environment. A model of the components expected in a robotic system is presented. Section 3 then presents the key concepts in the ontology of robotics science, particularly the notion of scaling laws, which structure the domain of robotics. Section 4 investigates how the proposed ontology might be applied, and demonstrates its potential usefulness. Section 5 illustrates how the key concept of Context Space can be implemented using standard ontology tools. The final section summarises the paper's conclusions and indicates planned future work.

2 A Model of Robotics

Representing human knowledge about technological domains, such as robotics, is equivalent to defining sets of *models* with which humans can *analyse* the real world, *make calculations* about it, make *predictions*, etc. In other words, a model is a mathematical simplification of the real world, representing only those aspects of the real world that are relevant to a particular human purpose.

The purpose of this paper is *to bring structure* to human knowledge of robotics, but not to give detailed technological descriptions of the workings of robotics devices or algorithms. We believe that the graphical models in figure 1 and 2 provide us with the appropriate balance between detail and generality for the purpose of this paper.

Figure 1 gives a “robot-centric” model of robotics: it depicts a robot “agent” as an integration of mechanical hardware, sensing hardware and actuation hardware, with planning software, control software, sensor processing software and actuator drive software. A successful robot interconnects knowledge and hardware at “appropriate” scales and with “appropriate” interfaces. Later sections of the paper will indicate more clearly what “appropriate” really means, and make clear that terms such as “field robotics” or “surgical robotics” indicate particular sub-manifolds of the “robot-centric” continuum (and the “application-centric” continuum of figure 2, see below) that have proven sufficiently successful or interesting to attract the intense research efforts from large communities.

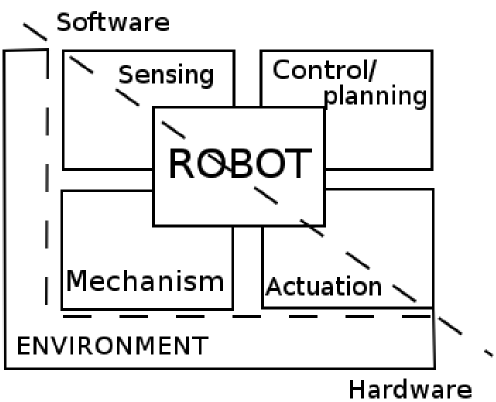

Fig. 1. The “robot-centric,” or *agent*-model, of robotics.

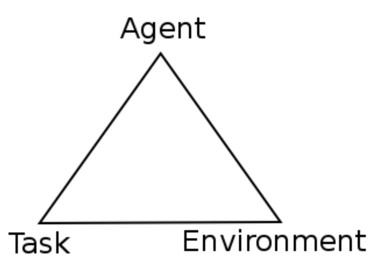

Fig. 2. Robotic systems as the manifold of robot *agents*, *tasks* and *environments*.

Figure 2 gives an “application-centric” or “systems-centric” view of robotics complementary to the previous figure: it makes clear that success of a robot agent is not a property of the agent in itself, but is determined also by (i) the *task* the robot is expected to perform, and (ii) the *environment* in which that task is to be performed. Each successful robotic *systems* domain covers a certain subspace of the triangular continuum depicted in figure 2. In some cases, it is the robot device that contributes most to the success of the system, and its subspace is situated close to the *agent* vertex of the triangle. In other systems, it is the environment ("hardware") that is most important, for example in sub-sea or airborne robots where water or air provide the necessary support to the robot. In yet other systems, the task is the determining factor.

3 An Ontology of Robotics Science

The key to defining relationships and structure between the various subfields of robotics is an understanding of how the subfields come about. Recall that a particular subfield expresses similarities between a class of robotic systems, and implies relevant background knowledge, significant problems, useful properties and so on. What accounts for much of this variation is *scaling laws*.

Consider the difference in the flight of an air-plane and a bumblebee. Each exploits different aerodynamic properties accessible to it because of the physical scale (tens of metres vs. millimetres) in which it operates: air behaves quite differently in different physical circumstances. This difference is an expression of a physical scaling phenomenon wherein physical system properties depend on the characteristic length of the system. Analogues of this principle can be constructed in other physical scales, such as time or mass, and in more abstract scales such as computational complexity.

It is the placement with respect to a set of such scales that determines the properties, problems and possibilities of a given subclass of robotic system.

3.1 Context Space

The *Context Space* is a conceptual space spanned by the various scaling dimensions available to a robotic system (or to its designer). Since, as we have argued in section 2, such a system comprises a device or agent, an environment, and a task, each of these essential components contributes scale dimensions and constraints to the system. In this section we examine some of these dimensions.

The example dimensions we have chosen are intended to be illustrative rather than complete — space precludes an exhaustive listing of all possibilities, and such is not necessary to make our point.

Each essential component of a complete robotic system contributes scaling dimensions. Thus

device scales are the most straightforward to specify. The characteristic length and time constant of the robotic system, its mass, the number of components it comprises and the extent to which they interact with each other are important characteristics, as are the number of controllable degrees of freedom and the total degrees of freedom of the system.

environment scales define the type and properties of the environment in which the task is carried out. There are two classes of environment: bulk and interface. Examples of the former are vacuum (as in space robotics) or liquid (as in underwater robotics). Bulk environments can be ordered by their mechanical impedance. Examples of the latter, interface, type include solid-gas (e.g. a table top, or the outdoor land surface) or liquid-gas (e.g. the surface of the sea). Salient properties here include the size and spatial frequency of surface variation — how smooth the interface is.
The size of the environment (its characteristic length); the speed with which events occur (its characteristic time constant); and whether the environment is abiotic or biotic (or on the interface of the two) constitute other examples of relevant dimensions of variability.

task scales include the spatiotemporal properties of the task, for instance its characteristic precision and the lifetime of the robotic system when working on the task, and the risk cost, i.e. the cost associated with failure or error in the task. The latter can informally be defined as the cost of fixing a disaster caused by system failure. For laboratory robotics it is, relatively speaking, very small; for surgical or space robotics it is relatively large.

Many of these scales have an ordinal structure, that is, the axis representing the dimension has a 'small value' end and a 'large value' end, with intermediate values that may be arranged in order of size. As we shall see below in section 3.3, these ordering relationships along individual axes contribute to the rich structure of the context space.

3.2 Contexts

Contexts are the basic mechanism for using the ontology. They provide a means to distinguish groups of robotic systems — and so different subfields of robotics as a whole — whose properties are restricted to sub-ranges of certain scale axes. For example, nanorobots have a device length scale in the range 10^{-9}–10^{-7} metres. More generally, we can specify sets of systems whose properties lie within specific ranges on several axes simultaneously; most generally, a context may be a *union* of such convex regions in the space.

Conceptually, a context is a sub-class of the class of all possible robotic systems. Such classes will usually (but not necessarily) have intuitive names, for example "field robotics," "medical robots," "swarm robots," and each context implies a set of relevant technology, background knowledge, scientific problems and solutions following from the particular choices of scale which define the context.

3.3 Relationships Between Contexts

Contexts, as defined above, can be thought of as sub-classes of the class of robotic systems; sub-classes generated by restricting the values of certain scaling properties (e.g. device length, risk cost, number of components, and so on) to particular values or sets of values from their respective scales (e.g. millimetre-metre, EUR 100–10000, 5–10, etc.). As such, contexts can take part in the standard class relationships of set theory: they can be super- or sub-classes of other classes; intersections, unions and complements can be defined; and membership can be checked.

However, more interesting are the *relationships* implied by the ordinal structure of the scale axes. For instance, robot devices whose physical size (characteristic length) is in the centimetre-decimetre range are neighbours of those in the decimetre-metre range: they will have many properties in common. On the other hand, nanometre scale robotic devices will be quite different from centimetre ones, in some respects, since the dominant physical effects differ between those characteristic sizes.

Neighbouring contexts are often contexts between which there exist loose analogies or trade-offs. For instance, one could describe a 100 mobile robot collective as a group of 100 6-dof systems or a single 600-dof system: exchange physical dimensions for interconnected system piece complexity. Neighbourhood implies quasi-invariance between aspects of the relationships that describe the same system in different contexts. Change in one scaling dimension is compensated by a corresponding change in another dimension.

Sufficiently sophisticated reasoning mechanisms could exploit neighbourhood relationships to propose novel connections between robotic systems, or identify possible similarities between hithertofore disjoint subfields of the discipline. The neighbourhood relationships provided by the context scales also force human students in robotics to think more explicitly about the really unique, respectively shared, properties of specific robotics domains.

4 Using the Ontology

The proposed ontology is unusual in that it focusses on the structure of a field of science rather than on the objects with which that field is concerned. Nevertheless, such an ontology has considerable practical use. In addition, it is easier to start with the ontology of robotics as a science, because, as presented in this paper, this ontology can be explained and understood (almost) completely in a relative short amount of time. In this section we give three illustrative examples of practical usefulness of the ontology.

4.1 Structuring the Domain of Robotics Science

To illustrate the ideas presented so far, consider the various sub-fields of robotics typically found in the contents listings of conference proceedings: can we define them using context space ideas and axes?

Suppose we start with Field Robotics. We can say that the device scale (robot size) is in the metre range; the environment length scale is 10^2 to 10^{10}

metres; the device/mission lifetime ranges between 10^3 and 10^{10} seconds; the environment may be vacuum, gas, liquid or an interface.

So far so good; now consider underwater robotics. This is Field Robotics in a bulk liquid environment. Space Robotics is Field Robotics in a bulk vacuum environment, with environment length scale from 10^5 to 10^{10} metres, device/mission lifetime 10^6–10^{10} seconds, say (for space craft) and a high risk cost.

Consider Research Lab Robotics. In this case we have device lengths of 10^{-3}–10^0 metres; environment length scale 10^1 metres, say; gas or smooth gas-solid interface; lifetime in the 10^2–10^5 second range; and low risk cost.

Table-top robotics, for instance using the popular Khepera robot, is a specialisation of the Research lab context, with 10^{-2} metre device scale, 10^0 metre environment scale, very smooth solid-gas interface environment, very low risk factor, low environmental interaction and lifetime in the 10^2–10^3 second range.

As a final example, consider factory automation or industrial robotics. Here we have device scales in the 10^{-3}–10^0 range, environments in the 10^0–10^1 metre range, gas-solid interface, lifetimes from 10^5–10^7 seconds; medium-high risk cost; moderate system complexity and low environment interaction.

Although these examples are still somewhat vague — recall, for instance, 'medium-high risk cost' or 'moderate system complexity' — they nevertheless illustrate that the different sub-fields of robotics occupy specifiable parts of the context space and that related sub-fields occupy neighbouring or subset regions of the context space. Thus it is possible to *define* the subfields of robotics by labelling the appropriate regions of the context space and thereby infer relationships between subfields that might not immediately be apparent from just their names. In addition, being confronted with the different ordinal scales that are relevant in a particular robotics domain will stimulate students and researchers to think about how the fundamental physical properties of a particular context could change when "moving" to a neighbouring context, or they could get inspiration from successful solutions implemented in neighbouring domains that they had not thought of before.

4.2 Use Case: Guided Reading for New Roboticists

All teachers in robotics have experienced the following evolution in their students: the first couple of weeks or months in their study of robotics, students are reading tons of papers, working hard to see the forest for the trees, but are mostly unable to evaluate the essential contributions (or lack thereof) of every individual paper they read. Good teachers also know that they must organise regular reading sessions with their students, in order to guide this process of getting acquainted with a particular robotics research domain.

We believe that a key role of the teachers in this coaching phase is to indicate to their students what are the relevant relationships between the details of the science explained in particular papers, and the goals of making a better robot system (i.e., the combination of agent, task *and* environment) in the particular domain of interest to the student. We believe that teachers (most often) *implicitly* convey to their students the kind of ontology that this paper presents, and that making these ontological relationships *explicit* will help the students in their learning. In addition, by using the ontology of this paper in real-world reading classes, the ontology will become more complete and refined. In addition, it could possibly become the basis for a series of textbooks that cover the whole domain of robotics in the most efficient way.

4.3 Use Case: Classification of Papers and Reviewers

A second application of these ideas is to paper classification. Imagine that authors submitting papers to a robotics journal or conference (or, indeed, a funding agency) are invited to indicate approximately where, on figure 2 and figure 1, their paper content falls. Then, depending on their answer, they are invited to specify where it lies on various appropriate scale axes. For example, the author of a paper on inertial guidance for underwater robots would indicate ‘agent-centred’, ‘sensors’, and choose ‘bulk liquid environment’, ‘environment size 10^2–10^3 metre’, ‘lifetime 10^3–10^4 second’, and so on. Reviewers could similarly indicate their areas of expertise. Software based on the ontology could then match reviewers to papers depending on their distance in the contextual space.

For most papers, this would produce the expected results: a paper on space robotics, for instance, would go to space robotics referees. But interesting desirable effects are possible. A paper on collective robotics might be offered to a reviewer interested in the control of very highly redundant systems, on the arguably sensible grounds that 100 6-dof robots have some similarity with one 600-dof one — the two subjects are neighbours in (some dimensions of) the context space.

5 Implementing the Ontology of the Science of Robotics

To be useful, it must be possible to implement the ontology of the science of robotics using the standard ontology tools of the semantic web.

The World Wide Web Consortium has selected OWL as its standard for the representation of ontological knowledge [4]. OWL is built on top of standard infrastructure (RDF and XML) and provides mechanisms for representing concepts (classes), instances, properties and relationships. It comes in three flavours — OWL-Lite, OWL-DL and Full OWL — which represent different compromises between representational power and computational tractability. The ontology described here is implemented using OWL-DL.

5.1 Implementing the Context Space

The key concept to implement is the Context Space. To do that, we identify a context with the set of all robotic systems falling into that context; i.e. 'Space Robotics,' in terms of the ontology, is defined as the class of robotic systems studied by that field. The class of such systems is actually defined using more fundamental properties, such as the environment, spatial and temporal scales in which the systems operate, and so on.

The basis for implementing the Context Space is then to define the class of all robotic systems, sub-classes of which represent individual contexts. Contexts are defined using standard class operations of union, intersection, complement and property value restriction (restricting the legal values of a specific instance property to a known class).

In line with the discussion of section 2, we therefore assert that each robotic system (an instance of the robotic *System Class*) comprises a device (the agent), an environment and a task. We further define three classes, which we call *Aspect Classes*, consisting respectively of all devices, all environments and all tasks, as the value ranges of the three instance properties of a system. Associated with the instances of each aspect class are scaling properties, for example device-characteristic-length or environment-impedance, the values for which are drawn from *Scale Classes* representing the various scaling axes. Notice that different aspects may have properties relating to the same scale: device, environment and task characteristic lengths or times would be an obvious example.

Figure 3 illustrates this general scheme. Technically, each instance of the robotic Systems class has three properties, whose values are restricted to lie each in one of the three aspect classes; and each instance of an aspect class (for example, a particular environment) has scaling properties whose values are restricted to particular sub-classes from given scale axes. This construction has the crucial advantage that the class structure is defined intensionally — we are not required to enumerate all robotic systems, all environments, etc. — while also allowing specific systems, devices and so on to be represented as instances.

To illustrate the implementation strategy, consider systems using nano-scale robotic devices, which can be characterised by the fact that their device-length scale is in the nanometre range. Figure 4 illustrates how this is represented: a sub-class of the length scale class is defined to represent the nanometre range, and a sub-class of the devices aspect class is defined such that the device-length property is restricted to take values from the nanometre range class. Finally, a sub-class of the systems class with device property values restricted to the just-defined nano-device class completes the representation of nano-scale robotic systems. Notice that we have not restricted the environments in which these systems may operate nor the tasks to which they may be applied. The red (dashed) components in the figure illustrate the newly

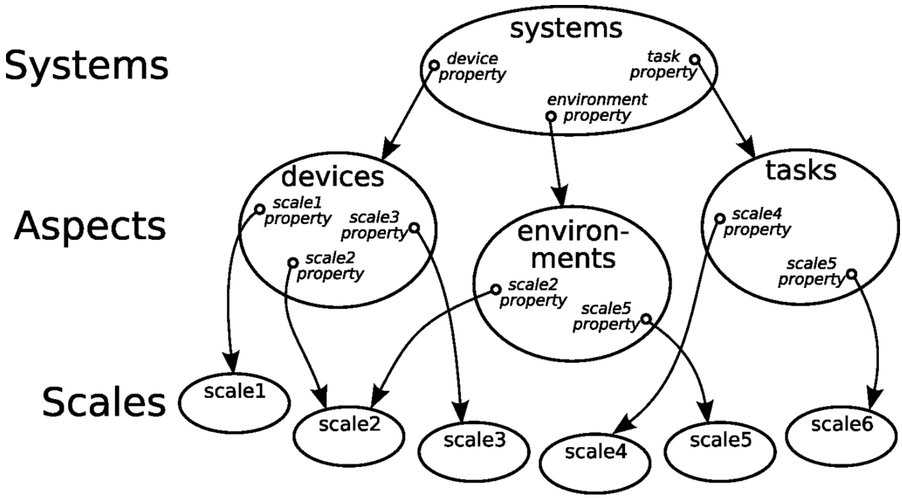

Fig. 3. Class hierarchy in the Robotics-as-Science Ontology: the universe is represented as the class of all robotic Systems, where each system (instance) has three Aspects — device, environment and task. Instances of the Aspect Classes have scaling properties, e.g. characteristic length, with values from the Scale Classes.

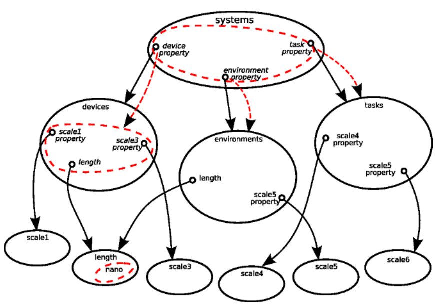

Fig. 4. Nano-Robotic Systems

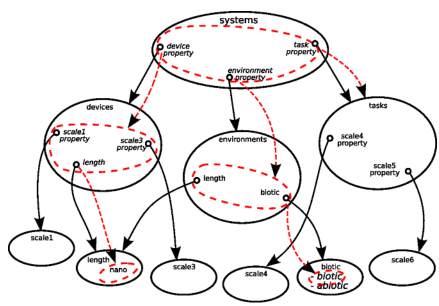

Fig. 5. Medical Nano-Robotics

defined sub-classes; the standard OWL reasoners are able to infer that these are indeed sub-classes of their parents.

Figure 5 takes the example one stage further. Consider medical nano-robotics. The principal distinction between this and nano-systems robotics, for our illustrative purposes, is the environment in which the systems must work — on a macro-molecular scale and in interaction with living matter. This is represented by the construction of a sub-class of biotic nanoscale environments to which the environment property of the system is restricted. The red (dashed) components in the figure illustrate these new definitions. Once again, standard OWL reasoners can infer that the medical nano-systems class is a sub-class of the nano-robotic systems or, equivalently, that medical nano-robotics is a sub-field of nano-scale robotics.

5.2 Context Relationships

The ability to represent contexts is crucial to the ontology, but so is the ability to represent or infer relationships. As we have seen, certain relationships such as class inclusion can be inferred by the standard OWL reasoners.

Not so with neighbourhood relationships, however. The problem is that OWL reasoners contain no machinery for inferring 'distance' between classes. To illustrate this, consider the length scale. We can construct in OWL a partition of this scale into (amongst others) millimetre, centimetre, decimetre and metre lengths with the appropriate ordering relationships (e.g., centimetres are longer than millimetres and shorter than metres). But it is not possible with standard reasoners to infer that the metre sub-class is an indirect neighbour of the centimetre class two steps up the ordinal scale.

What this means is that neighbourhood relationships, which depend on the 'distances' between context properties in the ordinal scale axes, cannot be inferred by standard reasoners since the ordinal 'distances' themselves cannot be inferred by those reasoners. However, more specialist reasoners can be built to make such inferences, and one can imagine such reasoners working on the ontology enumerating neighbourhood relationships between the represented contexts and automatically adding to the ontology explicit representations of discovered neighbourhood relationships.

5.3 Summary

The ontology of the science of robotics depends on the key concept of contexts, which we can straightforwardly and intuitively represent using the standard ontology language OWL by defining a context (a sub-field of robotics) to be represented by the class of all systems belonging to that context (sub-field). Standard reasoners can determine simple class relationships between contexts. Neighbourhood relationships between contexts cannot be inferred using standard reasoners, but specialist discovery processes could be implemented to annotate the ontology with such relationships.

6 Conclusions and Future Work

6.1 Conclusions

The present paper has motivated the notion of an ontology of the science of robotics as opposed to an ontology for the objects of which robotic systems are composed. An ontology of robotics science allows us to *define* terms and concepts such as "surgical robotics," "field robotics," or "nanorobotics" in an objective way in terms of more primitive technological concepts. This allows us to give formal (machine-usable) meaning to these various terms and to explore the relationships between them. These meanings and relationships provide a *meta-level structure* in the scientific discipline of robotics, that could be exploited in various ways.

The formal setting in which such definitions are made is a *Context Space* established by considering the physical, environmental and task-related scaling laws and relationships that apply to robotic systems. New sub-fields of robotics, and new relationships between them and between existing subfields, may be discovered by considering the relationships between them in context space. The ontology has been implemented using the standard ontology language OWL-DL using the Protegé suite.

The proposed ontology is believed to be simple enough to be understood by humans within the span of a couple hours, while at the same time it is rich enough to bring non-trivial and hence useful structure to scientific domain of robotics. The ontology may be used, *inter alia*, as a tool for classification of robotic material, the matching of reviewers to conference, journal or funding agency submissions, or during teaching robotics to new students.

Although the description of the ontology in the paper is somewhat incomplete, this is natural: to our knowledge there is no state of the art *at all* addressing this kind of ontology for robotics and also no truly comparable work in other (technological or non-technological) domains. Hence the goal of the paper is to present in outline a novel but, we believe, valuable tool for roboticists.

6.2 Future Work

Before an ontology of robotics science is complete and fully usable, an amount of future work remains — both conceptual and practical.

Conceptual work includes elaborating the set of scale axes to identify a complete, minimal and consistent set; identification of the kinds of questions and use cases for which the ontology provides a useful resource; and wide consultation with the robotics community on the clarity, utility and completeness of the conceptual framework.

Practical work comprises completing the implementation of the concepts and relationships of the ontology tools conforming to the standards and recommendations established by the World-Wide Web Consortium; and the formal definition and implementation of suitable query and reasoning interfaces to the implemented ontology. The ontology must also be tested more intensely on the "educational workfloor" of teaching robotics to masters and PhD level students.

Acknowledgement The authors gratefully acknowledge the support of the EURON Research Atelier *Rose: a robotics ontology for the semantic web*, and the many stimulating discussions with members of EURON's Education and Training committee, which have, to a large extent, helped shape the ideas of the ontology presented in this paper. Wei Xu contributed many useful ideas to the described implementation of the ontology.

References

1. Grigoris, A. and Van Harmelen, F. *A Semantic Web Primer*, MIT Press, 2004.
2. W3C. *Semantic Web*, `http://www.w3.org/2001/sw/`, accessed on-line on September 30th, 2005.
3. Wikipedia. *Ontology (computer science)*, `http://en.wikipedia.org/wiki/Ontology (computer science)`, accessed on-line on September 30th, 2005.
4. W3C. *OWL Web_Ontology Language Overview*, `http://www.w3.org/TR/owl-features/`, accessed on-line on January 12th, 2006.