Encountering the Physical World

Bahar AAMERI [a], Michael GRÜNINGER [a]
[a] *Department of Mechanical and Industrial Engineering, University of Toronto, Ontario, Canada M5S 3G8*

Abstract. Manufacturing and product design are grounded in the physical world. The entire product lifecycle involves a wide range of integrated tasks that focus on the properties of physical objects, beginning with begins with the design of physical objects and the specification of the materials from which the physical objects are made. Thus, reasoning tasks within manufacturing and product design requires an ontology of physical world. In this paper, we present the current status of a project that is developing a suite of ontologies, which are modules of an overarching ontology called the PhysicalWorld Ontology. Each module of the PhysicalWorld Ontology captures a particular class of physical phenomena or property of physical objects, such as shape, location, connectedness, parthood, and kinetic and kinematic behaviour.

Keywords. manufacturing ontologies, physical objects, mereotopology, shape

1. Introduction

Although it may be obvious, manufacturing and product design are essentially grounded in the physical world. Manufacturing is concerned with processes that involve the creation of physical objects, such as assembly, joining, fastening, fabrication, machining, coating, and more recently in additive processes such as 3D printing. From a wider perspective, manufacturing begins with the design of these physical objects (i.e. products) and the specification of the materials from which the physical objects are made. The entire product lifecycle (conceptual design, detailed design, manufacture, maintenance, disposal) involves a wide range of integrated tasks that focus on the properties of physical objects. The supply chain of the manufacturing enterprises spans the sourcing of raw materials and the delivery of products in logistics. All of this is supported by a vast ecosystem of product data management software as well as international standards. It is instructive to consider the scope of ISO 10303, also known as STEP (Standard for the Exchange of Product Data): standard data definitions for geometry (wire frame, surfaces and solid models), product identification, product structure, configuration and change management, materials, finite element analysis data, drafting, visual presentation, tolerances, kinematics, electrical properties, and process plans [2].

Even with this cursory inspection, we can see that a rich set of ontologies about the physical world is needed for manufacturing. Within this paper, we present the current status of a project that is developing such a suite of ontologies, which are modules of an overarching ontology that we refer to as the PhysicalWorld Ontology. Based on the idea that solid physical objects are self-connected objects that are made of matter, have

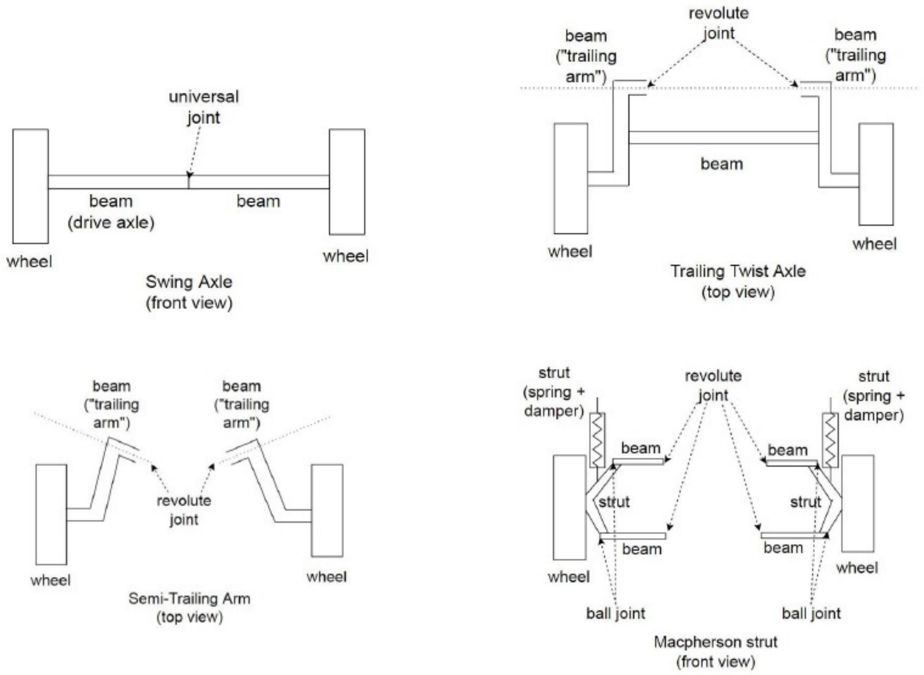

Figure 1. Schematics of swing axle, trailing swing axle, semi-trailing arm, and Macpherson strut suspension systems [1]

a shape with boundaries, and are located in space, the ontologies will support reasoning about physical objects, their behaviors and interaction. Each module of the PhysicalWorld Ontology captures a particular class of physical phenomena or property of physical objects.

We begin in Section 2 by identifying semantic requirements for the representation of physical objects and their behaviour. We then explore the primary modules of the PhysicalWorld Ontology in Section 3 – Shape, Multidimensional Mereotopology, and Location. We finish with a look forward to the remaining ontology modules that will focus on physics and axiomatize the fundamental concepts required for representing kinetic and kinematic behaviour of physical systems.

2. Extracting Requirements for PhysicalWorld Ontology

We will begin by delineating the requirements for representing properties and behavior of physical objects. Throughout this section, we consider suspension design systems (see Figure 1) as our main use cases for extracting requirements. A suspension system consists of wheels, beams, struts, springs and dampers, related to each other by different types of joints.

2.1. Shape

Perhaps shape is the first feature that comes to mind when thinking about a physical objects; it is also a key concept in representing physical domains. As an example, consider the different suspension systems shown in Figure 1. One of the features that distinguishes these systems from each other is the shape of the beams between the two wheels. Later in Section 3.4.1, we will raise the problem of distinguishing between different rivet fastening methods, and we will show that a description of the shape of rivets is required in order to make this distinction.

The focus of the majority of the existing shape formalisms is on representing convexity and curvature (see [7,13]). In cases where information about convexity and curvature are not required, a shape can be described based on the adjacency and order of its points, edges, and surfaces. This approach has been taken in [10] for developing first-order ontologies for two-dimensional and three-dimensional shapes. We will discuss this approach further in Section 3.1.

2.2. Connection and Parthood between Physical Objects

Formalizing the part-whole relationship between a physical system (e.g., a suspension system) and its parts, as well as the physical relationships among the different parts of the system, requires a mereotopology for physical objects. A mereotopology is a formal theory which combines topology with mereology. The topological subtheory expresses connection relations between a set of individuals, while the mereological subtheory expresses parthood relations.

Mereotopological systems differ in their basic assumptions about supplementation, atomicity, extensibility, and closure under sum and product of spatial entities. Ground Mereotopology (MT) [6] is the weakest theory among the existing mereotopological theories, and does not take any of these assumptions. The signature of the MT theory consists of two primitive binary relations, parthood and connection. The axioms of the theory state that connection is a reflexive and symmetric relation, while parthood is a reflexive, transitive, and anti-symmetric relation. In addition, if one individual is connected to another, then the first one is also connected to any individual which the second is part of.

[9] shows that the MT theory is logically synonymous with a non-conservative extension of the RCC8 theory, called RCC8*, meaning MT and RCC8* are semantically equivalent, and only differ in signature (i.e., the non-logical symbols). In other words, MT is the mereotopology that underlies RCC8. [1]

The RCC8 relations have widely been used for describing spatial relationships within physical settings. This means that the underlying mereotopology (i.e., MT) used in such settings does not include any of the basic mereotopological principles (i.e., supplementation, atomicity, extensibility, and closure under sum and product). In the following, we provide examples of physical objects that do not satisfy atomicity, extensibility, and closure under sum. It remains, however, an open question whether physical mereotopologies should be supplemental and/or closed under product.

Many mereotopological theories, such as the Region Connection Calculus (RCC) [14], entail that domain entities are atomless. Within a domain, individuals are atomless

---

[1] RCC8 is a set of eight jointly exhaustive and pairwise disjoint binary relations representing mereotopological relationships between ordered pairs of individuals.

if every element has a proper part. However, physical objects are not necessarily atomless. For assembling a bookshelf, we do not care about proper parts of shelves and divider. In many applications we want to have a finite domain, meaning that the elements of the domain are not atomless. Thus, using atomless mereotopologies for representing physical objects is inappropriate as the additional unnecessary constraints result in the elimination of valid models. On the other hand, there might be domains which require an atomless representation of some classes of physical objects. Thus, a general physical mereotopology should not make any commitment about the atomicity of objects, and the atomicity assumption should be taken with respect to domain-specific requirements.

A relation R is said to be extensional if it satisfies the following sentence

$$(\forall z)\,(R(z,x) \equiv R(z,y)) \supset x = y.$$

Mereotopologies like RCC assume that the connection relation is extensional. However, extensionality does not always apply to physical objects. Consider a model with two elements. The two elements are obviously connected to the same set of elements, but they are not identical. In fact, any model with finite number of elements (which is the case in many physical domains) may not be extensional.

The following axiom entails that if two (self-connected) entities are connected, they add up to a self-connected whole ($C(x,y)$ denotes 'x is connected to y' and $P(x,y)$ denotes 'x is part of y'):

$$(\forall x,y)\,C(x,y) \supset (\exists z)\,P(x,z) \wedge P(y,z).$$

However, there are physical domains that do not satisfy this axiom. For example, within the geospatial applications, two neighboring countries are connected, but their summation is not an entity in the domain. If a glass is placed on a desk, their sum does not make a new entity. With a similar example, we can also argue that physical domains are not necessarily closed under the summation of two underlapping objects.

2.3. Location

Mereotopologies alone are not sufficient for describing different configurations of physical objects. In standard mereotopologies, overlap relation between two individuals is defined with respect to their common parts; that is, two individuals a, b overlap if and only if there exists an individual which is part of both a and b. When two physical objects overlap, they do not necessarily have a common part. A book on a shelf, for example, overlaps with the shelf, however they do not have a part in common. In this case, the overlap relationship between physical objects should not be defined based on common parts; it should be defined based on a common abstract region that two physical objects occupy.

There are also mechanical objects with components which coincide but do not share parts. For example, while cartridges inside the cylinder of a revolver coincide with the cylinder, they are not part of the cylinder. Similarly, the ball of a ball joint coincides with the space that the hole of the joint surrounds. Definition of relations such as coincide requires a logical theory that axiomatizes relationships between physical objects and the spatial region they occupy. This is what is called a location ontology.

A location ontology is also required to distinguish between different types of spatial change. A moving object, for example, occupies a region that overlaps with the region the object originally occupies, while a shrinking object occupies a region that is a proper part of its original region. So, for distinguishing between movement and shrinkage of an object we need to know the relationships between the spatial region they occupy at each state.

2.4. Joints and Attachment

With topological connection it is only possible to express contact between objects. However, in many application domains we require to distinguish between being "in contact" and being "attached". A book on a shelf of a bookshelf, for example, is only in *contact* with the shelf, while the shelf is *attached* to the side panels of the bookshelf. In addition, in some domains, such as manufacturing and assembly, representations for different ways that physical objects can be join together are required.

There has been previous work on the ontologies for attachment. [12] suggests a set of first-order definitions describing mechanical joints. They use Smith's mereotopology [15] for describing the relationships between mechanical parts. The descriptions they have provided are, however, incomplete in the sense that it does not capture the intended specifications of different types of joints. Consider the joining methods depicted in Figure 2. [12] proposes the following definition for threaded fasteners:

$$\begin{aligned} J_{jf}(x,y) \equiv (\exists u,v)\,(X(u,x) \wedge X(u,y))\wedge \\ (T(v,x) \vee T(v,y)) \wedge (P(u,f_s) \wedge P(v,f_s)) \wedge X(f_s,j) \end{aligned} \tag{1}$$

and the following definition for fastening by rivets:

$$\begin{aligned} J_{rf}(x,y) \equiv (\exists u,v)\,(X(u,x) \wedge X(u,y))\wedge \\ T(v,x) \wedge T(v,y) \wedge P(u,f_s) \wedge P(v,f_s) \wedge X(f_s,j) \end{aligned} \tag{2}$$

Here, $P(x,y)$ denotes 'x is part of y', $X(x,y)$ denotes 'x crosses y', and '$T(x,y)$ denotes x tangents y'.

The description for rivet fastening relays on the shape of the rivet (denoted by v): a part like the one shown in Figure 2 would satisfy both Definition 1 and Definition 2. In fact, to have a sound and complete description for rivets fastening, and so be able to distinguish it from threaded fasteners, we need to be able to describe the shape of the rivet. That is, in addition to mereotopological relations, a complete ontology of mechanical joints requires an ontology for specifying shapes of parts involved in joining methods.

2.5. Boundary

Specifying properties of mechanical joints and physical attachments requires a formal representation of the notion of boundary.

Consider, for example, a ball joint. Each of the ball and the hole of a ball joint have their own boundary surfaces, and one of the boundary surfaces of the ball is connected (in the topological sense) to one of the boundary surfaces of hole. However, if we weld

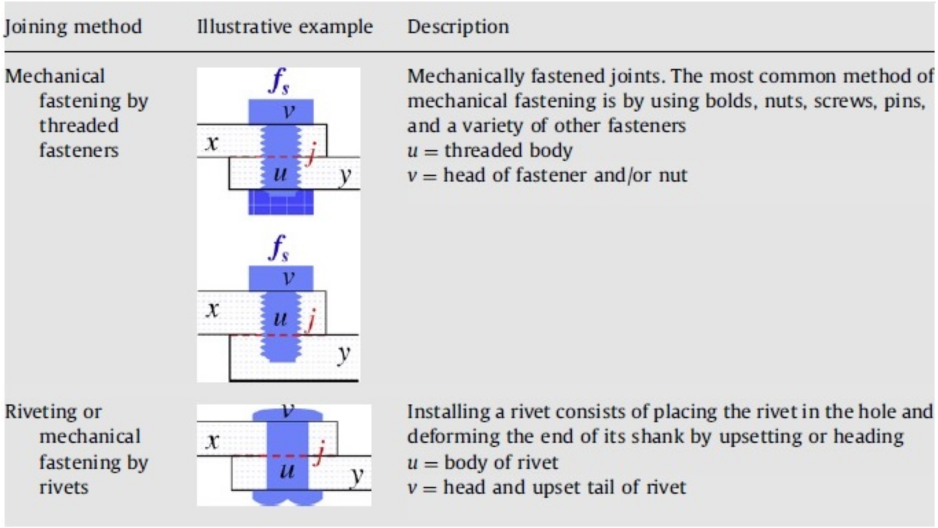

Joining method	Illustrative example	Description
Mechanical fastening by threaded fasteners		Mechanically fastened joints. The most common method of mechanical fastening is by using bolds, nuts, screws, pins, and a variety of other fasteners u = threaded body v = head of fastener and/or nut
Riveting or mechanical fastening by rivets		Installing a rivet consists of placing the rivet in the hole and deforming the end of its shank by upsetting or heading u = body of rivet v = head and upset tail of rivet

Figure 2. Examples of mechanical fastening methods [12].

two three-dimensional objects, the welded surfaces of the objects will be transformed into a single surface, and more importantly, the surface will not be a boundary surface anymore.

[15] defines boundary based on the interior of entities, using the closure operators. In Smith's theory a boundary is a region which has empty interior. That is, unlike other existing approaches, boundaries are not considered as lower-dimensional entities. Moreover, each boundary is a part of the region it bounds, and is a boundary of itself.

The alternative approach, adopted in GFO-Space theory [3] and CODIB [11], is to consider boundary as a lower-dimensional entity which is part of the bounded entity. A model of the GFO-Space theory is partitioned into four categories: *space regions*, *surface regions*, *line regions* and *point regions*, corresponding to three-, two-, one-, and zero-dimensional space entities, respectively. A boundary is a lower-dimensional entity which does not exist independently of the entity it bounds. Moreover, a boundary always bounds an entity with a higher dimension. A boundary does not necessarily fully cover the entity it bounds, and in that case, is part of another boundary which covers more of the entity. Within the GFO-Space theory it is assumed that boundaries are not connected (in topological sense) to other entities (including other boundaries). Rather, two boundaries may be coincide, meaning that they are congruent and there is no distance between them. [11] takes a similar approach, but provides a stronger specification of properties of boundaries, and their relationships to the corresponding bounding entity

Note that the multi-dimensional approach for axiomatizing boundary requires a multidimensional mereotopology. Since physical objects are multidimensional themselves, it seems (even without considering which approach for representing boundary is taken) that using multidimensional mereotopologies is more adequate than equidimensional mereotopologies.

2.6. Kinematic and Kinetic Behaviour

In addition to the static properties of a physical system, one might be interested in the kinematic and kinetic behaviour of a system. Consider again a suspension system. Axiomatizing the behaviour of springs and dampers requires an axiomatic representation of force, which in turn requires axiomatic theories of mass, acceleration, velocity, time, and displacement.

3. Design of the PhysicalWorld Ontology

The PhysicalWorld Ontology is being developed in an ongoing project that aims to axiomatize concepts and properties required for representation and reasoning about physical domains. The PhysicalWorld Ontology consists of five main modules, namely the Multi-Dimensional Mereotopology, the Occupation Ontology, the Shape Ontology, the Attachment Ontology, and the Physics Ontology. Each of these ontologies have their own modules, and captures one or more of the required concepts described in Section 2. Figure 3 shows the relationship between modules of the PhysicalWorld Ontology.

3.1. The Shape Ontology

The Shape Ontology is a qualitative representation of shape of physical objects. The Shape Ontology is an extension of the BoxWorld Ontology presented in [10], which is based on Hilbert's axiomatic theory of geometry. Hilbert's theory consists of three subtheories: the first subtheory axiomatizes properties of the incidence relation; the second one is a theory of betweenness; and the third one describes congruence relationships. The focus of the Shape Ontology is on the incidence and betweennes relations, and ignores geometrical notions such as length and relative alignment of lines, or curvature and areas of surfaces.

The Shape Ontology consists of three main modules: CardWorld, BoxWorld, PolyWorld. Within the domain of a model of the Shape Ontology there are four disjoint categories of entities – *points*, *edges*, *surfaces*, and *boxes*, where they respectively correspond to zero-, one-, two-, and three-dimensional objects. CardWorld captures the relationship between points, edges, and surfaces. Describing properties of a single box and its parts (i.e., its edges and surfaces) are the focus of the BoxWorld Ontology, whereas the PolyWorld Ontology axiomatizes the relationships between multiple boxes.

The signature of the Shape Ontology includes a binary relation, *part*, that captures the incidence relations between different categories of objects. A lower-dimensional entity cannot exist independently and is always part of a higher dimensional object. For each box, there exists at least one surface which is part of the box. Similarly, for each surface exists at least an edge, and for each edge exists at least one point.

The set of edges in a surface, and the set of surfaces in a box, form cyclic orderings. The edges in a surface are partitioned into disjoint cyclic orderings so that one of these orderings is formed by the outer edges of the surface, and the remaining cycles represent holes within the surface. For each surface, there exists a unique set of outer edges that are all elements of the same cycle.

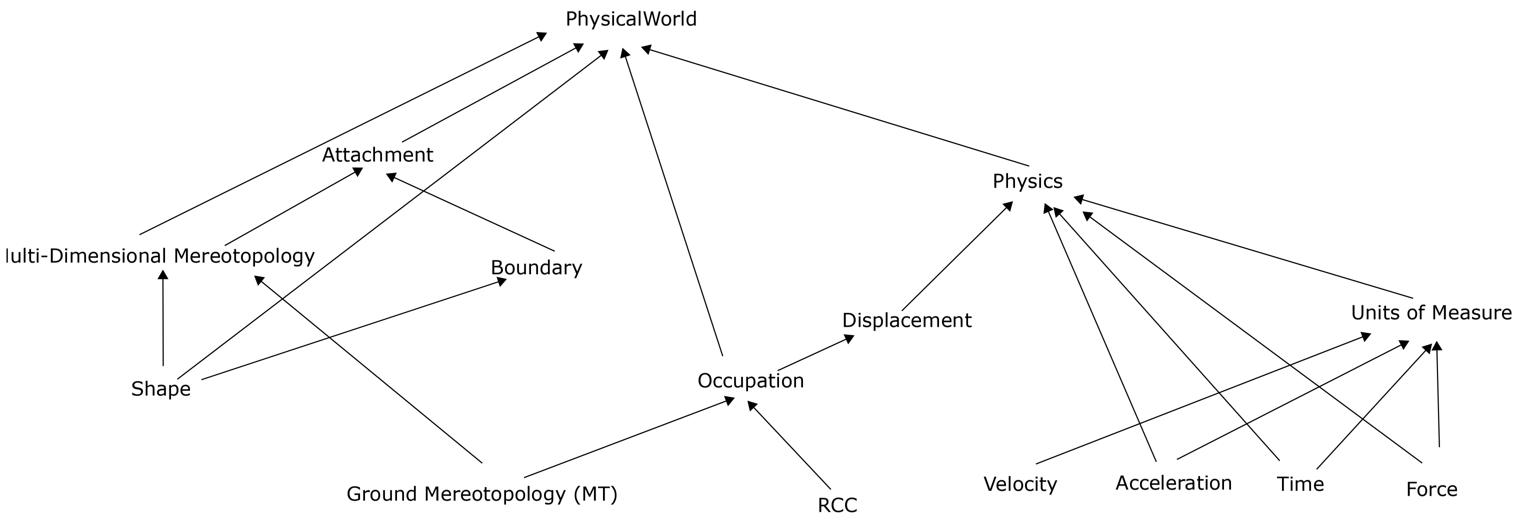

Figure 3.: Modules of the PhysicalWorld Ontology

Round objects (like circles) are the simplest two-dimensional object that can be described by the Shape Ontology. A Round object is a surface which has exactly one edge.

$$(\forall s)\, round(s) \equiv surface(s) \wedge (\exists e)\, edge(e) \wedge part(e,s) \wedge$$
$$(\forall e_1) edge(e_1) \wedge part(e_1,s) \supset (e_1 = e).$$

The Shape ontology, however, cannot represent the difference between a circle and an oval (i.e. curvature is not definable).

A box is a three-dimensional entity that contains at least one surface. For a box with multiple distinct surfaces, each surface will contain edges, called *ridges*, that are part of exactly two surfaces. In a polyhedron, every edge is a ridge. There are also models that are not polyhedra; in such models, there exist edges that are parts of unique surfaces. An edge that is part of a unique surface is a *border*.

Using the Shape Ontology, a sphere can be described as a box which has exactly one surface:

$$(\forall x)\, sphere(x) \equiv box(x) \wedge (\exists s)\, surface(s) \wedge part(s,x) \wedge$$
$$(\forall s_1) surface(s_1) \wedge part(s_1,x) \supset (s_1 = s).$$

And a cylinder can be described as a box that has three surfaces such that two of these surfaces are round objects, and the round objects do not have a common edge:

$$(\forall x)\, cylinder(x) \equiv box(x) \wedge (\exists s_1, s_2, s_3, l_1, l_2)\, surface(s_1) \wedge round(s_2) \wedge$$
$$round(s_3) \wedge edge(l_1) \wedge edge(l_2) \wedge part(l_1,s_1) \wedge part(l_1,s_2)$$
$$\wedge part(l_2,s_1) \wedge part(l_2,s_3) \wedge part(s_1,x) \wedge part(s_2,x) \wedge part(s_3,x).$$

3.2. Multidimensional Mereotopology

We use MT as the mereotopological theory for expressing connection and parthood between physical objects since MT is the weakest theory among the existng mereotopologies (recall from Section 2.2 that stronger mereotopologies impose constraints that may not be applicable to all classes of physical objects). However, as we explained in Section 3.1, there are four classes of physical entities in the PhysicalWorld Ontology, namely points, lines, surfaces, and boxes. Therefore, a multidimensional mereotopology is required.

Relationships between equidimensional individuals are captured by MT, while each class of object is mereotopologically independent of other classes. Individuals with different dimensions are only related by *part*, which is an incidence relation (see Section 3.1).

3.3. The Occupation Ontology

All of the existing axiomatic theories of location (including [6,8,4,5]) use a mereotopology stronger than MT over non-region entities. Thus, as we discussed in Section 2.2, they are not desirable for representing locative properties of some classes of physical objects. Moreover, some of these theories ([6]) allow mereotopological relationship between abstract regions and non-abstract objects, which leads to the existence of models that physically do not make sense. To overcome these shortcomings, we developed a new location theory called the Occupation Ontology.

The Occupation Ontology specifies physical location. Within the Occupation Ontology space is consider as an abstract entity in which other elements are located. For example, Canada is an object which is located on the abstract region between Atlantic Ocean and Pacific Ocean. The following is the list of ontological commitments the Occupation Ontology satisfies:

- Spatial regions and physical objects are distinct entities.
- There is no mereotopological relationship between spatial regions and physical objects. That is, a physical object is not part of (or connected to) an spatial region, or vice versa. Instead, physical objects occupy spatial regions.
- Occupation is a relation between a physical object and an spatial region. In other words, we assume that a spatial region does not occupy itself, or other spatial regions.
- There is no mereotopological relationships between spatial regions and physical objects; that is, a physical object is neither part of nor connected to a spatial region.
- The mereotopological relations between physical objects must be mirrored in the mereotopological relations between the corresponding spatial regions. That is, if a physical object a is part of (connected to) another physical object b, then the region occupied by a is part of (connected to) the region occupied by b.

Considering these commitments, the Occupation Ontology consists of three modules: a mereotopology over abstract regions, namely the Region Connection Calculus (RCC) [14], the MT theory relativised to physical objects, and the following axioms that specify the occupation relationship between abstract regions and physical objects:

$$obj(x) \supset \neg region(x).$$

$$occupy(x,y) \supset obj(x) \wedge region(y).$$

$$occupy(x,y) \wedge occupy(x,z) \supset (y = z).$$

$$obj(x) \supset (\exists y)\, occupy(x,y).$$

3.4. Remaining work

In this section we discuss the design of modules of the PhysicalWorld Ontology that have not been axiomatized yet.

3.4.1. The Attachment Ontology

The Attachment Ontology consists of definitions, based on relations specified by the Shape Ontology and Multidimensional Mereotopology, for describing different types of physical attachments and joints.

Currently, two types of attachment are included in the ontology, namely strong attachment and weak attachment. We define two boxes to be strongly attached if they are connected and have a common surface (i.e., there exists an surface which is incident with the two boxes). Two boxes are weakly attached if they are connected but do not have a common surface.

$$(\forall x,y) strong\_attach(x,y) \equiv C(x,y) \wedge (\exists z)\, surface(z) \wedge part(z,x) \wedge part(z,y).$$

$$(\forall x,y) weak\_attach(x,y) \equiv C(x,y) \wedge \neg strong\_attach(x,y).$$

In order to specify properties of other types of joints, we require to incorporate the notion of boundary into the PhysicalWorld Ontology. The existing theories of boundary, discussed in Section 2.5, only consider boundaries in abstract regions, and cannot be applied for representing boundaries of physical objects. It is part of the remaining work to apply ideas from these theories, and develop an ontology of physical boundaries. In particular, we need to identify an axiomatic specification for boundaries of three-dimensional objects (i.e., boxes).

3.4.2. Physics Ontology

The Physics Ontology axiomatizes fundamental concepts required for representing kinetic and kinematic behaviour of physical systems. These concepts include time, displacement, velocity, acceleration, mass, and force. The Physics Ontology includes a module for each of these fundamental concepts. Considering the quantitative formulation of these concepts, the Force ontology depends on the Mass and Acceleration Ontologies, and the Acceleration Ontology is axiomatized using the Time and Velocity Ontologies. The Velocity Ontology itself is specified with respect to the Time and the Displacement Ontologies. Note also that for representing displacement we require a representation for physical location, that is, the Displacement Ontology depends on the occupation relation specified by the Occupation Ontology.

In addition to axiomatizing fundamental concepts, the Physics Ontology includes a module, called Units of Measure, that specifies how units of measure corresponding to each concept is manipulated. More specifically, the ontologies explicitly axiomatize how units can be added, subtracted, and multiplied. Moreover, the Units of Measure Ontology utilizes existing ontologies for time, mereotopology, location, and constitution to axiomatize the relationship between units of measure and the concept being measured.

4. Summary

Any ontology that supports reasoning about the design and manufacturing of products must be rooted in a set of more foundational ontologies that represent the commonsense intuitions about the physical world. Starting with the idea that solid physical objects are self-connected objects that are made of matter, have a shape with boundaries, and are located in space, we have designed a suite of ontologies which are modules of an overarching ontology that we refer to as the PhysicalWorld Ontology. The current status of the development of the PhysicalWorld Ontology and its modules is summarized in Table 1.

Concept	**Ontology**	**Development Phase**
Connection and Parthood	Multidimensional Mereotopology	Axiomatized
Location	Occupation Ontology	Verified
Qualitative Shape	Shape Ontology	Axiomatized
Joints and Attachment	Attachment Ontology	Under development
Kinematic and Kinetic Behaviour	Physics Ontology	Under development

Table 1. Current status of the development of modules of the PhysicalWorld Ontology.

In addition to supporting automated reasoning about manufacturing and product design, the PhysicalWorld Ontology also provides a possible foundation for the ontological analysis of relevant existing standards and to integrate the ontologies within those standards, in particular ISO 18629 (PSL), OWL-Time, ISO 10303 (STEP), and ISO 15531 (MANDATE).

References

[1] Autodesk. Use cases and model examples from autodesk. 2016.

[2] Raphael Barbau, Sylvere Krima, Sudarsan Rachuri, Anantha Narayanan, Xenia Fiorentini, Sebti Foufou, and Ram D. Sriram. Ontostep: Enriching product model data using ontologies. *Computer-Aided Design*, 44(6):575 – 590, 2012.

[3] R. Baumann, F. Loebe, and H. Herre. Towards an ontology of space for gfo. In *FOIS*, pages 53–66, 2016.

[4] T. Bittner. Logical properties of foundational mereogeometrical relations in bio-ontologies. *Applied Ontology*, 4(2):109–138, 2009.

[5] S. Borgo, N. Guarino, and C. Masolo. An ontological theory of physical objects. In *Proceedings of Qualitative Reasoning 11th International Workshop*, pages 223–231, 1997.

[6] R. Casati and A.C. Varzi. *Parts and places: The structures of spatial representation*. MIT Press, 1999.

[7] A. Cohn. A hierarchical representation of qualitative shape based on connection and convexity. In *Proceeding of COSIT95, LNCS 988*, pages 311–326. Springer Verlag, 1995.

[8] M. Donnelly, T. Bittner, and C. Rosse. A formal theory for spatial representation and reasoning in biomedical ontologies. *Artificial Intelligence in Medicine*, 36(1):1–27, 2006.

[9] M. Gruninger and B. Aameri. A new perspective on the mereotopology of rcc8. In *COSIT*, 2017.

[10] M. Gruninger and S. Bouafoud. Thinking outside (and inside) the box. In *Proceedings of SHAPES 1.0: The Shape of Things. Workshop at CONTEXT-11*, volume 812. CEUR-WS, 2011.

[11] T. Hahmann. *A reconciliation of logical representations of space: from multidimensional mereotopology to geometry*. PhD thesis, PhD thesis, Univ. of Toronto, Dept. of Comp. Science, 2013.

[12] K. Kim, H. Yang, and D. Kim. Mereotopological assembly joint information representation for collaborative product design. *Robotics and Computer-Integrated Manufacturing*, 24(6):744–754, 2008.

[13] R. Meathrel and A. Galton. A hierarchy of boundary-based shape descriptors. In *Proceedings of IJ-CAI'01 - Volume 2*, pages 1359–1364. Morgan Kaufmann Publishers Inc., 2001.

[14] D. A. Randell, Z. Cui, and A. Cohn. A spatial logic based on regions and connection. In *Proceedings of the KR92*, pages 165–176. Morgan Kaufmann, 1992.

[15] B. Smith. Mereotopology: a theory of parts and boundaries. *Data & Knowledge Engineering*, 20(3):287–303, 1996.