Ontologies in Category Theory: A search for meaningful morphisms

Cauã Antunes, Mara Abel

Instituto de Informática – Universidade Federal do Rio Grande do Sul (UFRGS)
Caixa Postal 15.064 – 91.501-970 – Porto Alegre – RS – Brazil

crantunes@inf.ufrgs.br, marabel@inf.ufrgs.br

Abstract. *In order for the entire potential promised by the semantic web to be achieved, applications must be able to integrate knowledge from multiple sources. This requires reliable ways of relating disparate ontologies to be found. Category theory and its morphisms represent a possible solution to this problem, but careful definition of morphisms is required. This paper analyzes several category theoretical approaches to ontologies, focusing on the chosen morphisms and the semantic consequences thereof. Criteria for semantically evaluating morphisms are then discussed and standing challenges are identified.*

1. Introduction

The Semantic Web was envisioned as a mean to allow computers to access, process and make inferences over decentralized information [Berners-Lee et al. 2001]. In order to achieve this, the concepts and the relations between them must be specified formally in computational objects called ontologies.

Individual ontologies, however, are not enough to provide the full functionalities conceived by the semantic web, given the distributed nature of the information and intended applications. These applications must access knowledge spread across multiple ontologies. In order for this to be possible, semantically sound methods must be found to integrate information from different ontologies with heterogeneous specifications.

Category theory is a mathematical formalism that focuses on the relations (referred to as morphisms in related literature) between entities (objects) rather than emphasizing the entities themselves. Through this emphasis in relations, category theory abstracts from the representational aspects that hinder the integration of disparate ontologies. Thus, it provides a formal and sound foundation upon which relations between ontologies can be studied. Category theory also allows operations over these relations, as investigated by [Zimmermann et al. 2006], easing the effort necessary to associate indirectly related ontologies. Analogously, [Seremeti and Kougias 2013] indicate the composition of morphisms as a powerful tool to easily assimilate a new ontology into a previously existing network of ontologies.

Nevertheless, this expressiveness power depends on the morphisms used to express the relations between ontologies. An adequate selection of morphisms must be made in order to guarantee that the knowledge contained in the ontologies is not rendered useless by translation failures. This selection can only be done properly if criteria can be found for the identification of suitable morphisms. In this paper we discuss a possible set of guidelines to fulfill this role.

The remainder of this paper is organized as follows: Section 2 introduces the fundamentals of category theory, Section 3 discusses works from the literature that deal with category theoretical approaches to ontology and their respective morphisms, Section 4 presents an analysis on what it means for a morphism definition to be good in a semantic sense and propose criteria for evaluating such definitions, Section 5 discusses challenges that are still faced in the definition of ontology morphisms, and Section 6 contains a brief conclusion.

2. Category Theory Fundamentals

According to [Adámek et al., 1990], a category is a quadruple A = (O, hom, id, $\circ$), consisting of:

1. a class O, whose members are called A-objects,
2. for each pair (A, B) of A-objects, a set $hom(A,B)$, whose members are called A-morphisms from A to B,
3. for each A-object A, a morphism id_A: $A \rightarrow A$, called the A-identity on A,
4. a composition law associating each A-morphism $f : A \rightarrow B$ and each A-morphism $g : B \rightarrow C$ to an A-morphism $g \circ f : A \rightarrow C$, called the composite of f and g,

subject to the following conditions:

a) composition is associative, i.e., for morphisms f: $A \rightarrow B$, g: $B \rightarrow C$ and h:$C \rightarrow D$, the equation $h \circ (g \circ f) = (h \circ g) \circ f$ holds,

b) A-identities act as identities with respect to composition, i.e., for any morphism f: $A \rightarrow B$ we have $id_B \circ f = f = f \circ id_A$,

c) the sets $hom(A,B)$ are pairwise disjoint.

A diagram in a category A is a functor (a structure preserving morphism between categories) $D: I \rightarrow A$. A source for a diagram is a pair $(A,(f_i)_{i\in I})$, consisting of an object A and a family of morphisms $f_i : A \rightarrow A_i$ with domain A, indexed by I. A cone (also called a natural source) is a source such that for each I-morphism $d{:}i \rightarrow j$ the triangle $A \rightarrow D_i \rightarrow D_j \leftarrow A$ commutes, i.e., $d \circ f_i = f_j$. A limit is a terminal cone, i.e., it is a cone $(A, (f_i)_{i\in I})$ such that for every other cone $(A^j,(f_i)^j_{i\in I})$ there is a unique morphism $g_j : A_j \rightarrow A$ in order that the resulting diagram commutes. A limit for a diagram with two objects and no morphisms (other than identities) is called a product; a limit for a diagram of the form $B \rightarrow A \leftarrow C$ is called a pullback; and a limit for a diagram of the form $A \rightrightarrows B$ is called an equalizer.

In category theory, the concept of duality plays an important role. A dual for a categorical construct is obtained by reversing the domain and codomain of its morphisms. Thus, for a category A = (O, hom, id, $\circ$), the dual category of A is the category A^{op} = (O, hom^{op}, id, $\circ^{op}$), where $hom^{op}(A,B) = hom(B,A)$ and $f \circ^{op} g = g \circ f$. Similarly, the concepts previously introduced also present their own dual constructs. The duals of sources, cones, limits, products, pullbacks and equalizers are called, respectively, sinks, cocones, colimits, coproducts, pushouts and coequalizers. These categorical structures have different interpretations (with similar properties) over distinct domains.

Figure 1 shows a product (the object A and associated morphisms) and a cone (A' and associated morphisms) over a diagram containing only objects A_1 and A_2 (1), a coproduct and a cocone over the same diagram (2), a pullback and a cone over the diagram $A_1 \rightarrow A_0 \leftarrow A_2$ (3), a pushout and a cocone over the diagram $A_1 \leftarrow A_0 \rightarrow A_2$ (4), and an equalizer and a cone over the diagram $A_1 \rightrightarrows A_2$ (5).

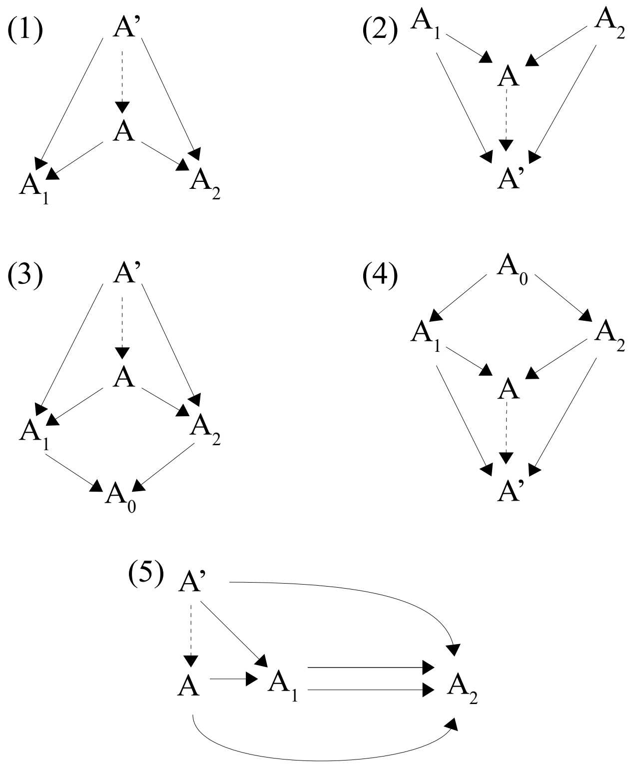

Figure 1. Categorical constructs: product (1), coproduct (2), pullback (3), pushout (4), and equalizer (5).

3. Related Works

In one of the earliest works to define a category of ontologies with its morphisms, [Bench-Capon and Malcom 1999] base their specifications on morphisms outlined by [Goguen and Meseguer 1992] between order-sorted theories, i.e., theories that classify the objects in the universe in different sorts (types) that may be related by sub-sorting. These morphisms are monotonic in relation to the ordered sorts, preserving the sub-sort relations. It should be noted that this definition of morphism does not account for any non-hierarchical relation in the ontologies. In the same work relations between two ontologies O_1 and O_2 are defined as structures composed of a third ontology O and morphisms $x_i : O \rightarrow O_i$ for i = 1,2. This structure is a cone for the diagram containing only O_1 and O_2, therefore a limit for the same diagram (a product) would represent the maximum possible alignment between the two ontologies.

[Zimmerman et al. 2006] expand on these relations, renaming them as V-alignments due to their shape and defining the operations of ontology merging (a pushout over the alignment), alignment composition, union and intersection. The V-alignment is found lacking the expressive power needed for complex alignments, such as between ontologies that contain concepts related not by equivalence but through subsumption relations. Two solutions are then proposed, the first using W-alignments, which are similar to V-alignments with the addition of a bridge ontology, and the second with the use of morphisms capable of expressing relations other than equivalence between concepts, namely strict inclusion, strict containment, disjointness and overlapping with partial disjointness. A composition table for such relations is also provided. Similar morphisms were later used by [Euzenat 2008]. V- and W-alignments were also used in posterior works by [Wang et al. 2008] and [Codescu et al. 2014].

The work of [Cafezeiro and Haeusler 2007] defines morphisms that preserve relations, both taxonomical (being monotonic in this regard) and otherwise. These morphisms are understood as mappings between ontologies, and used to build ontology operations based on categorical constructions: the product and pushout represent, as in previous works, alignment and merge operations, respectively, while the pullback denote a search for similarities between two ontologies in the context of a third, broader ontology. It is also shown how categorical equalizers may be used to hide sensitive information in an ontology. These operations, with the exception of component hiding, were used in the work of [Seremeti and Kougias 2013] with the same categorical interpretations. These morphisms were later expanded by [Cafezeiro et al. 2014] to include mappings between the axioms of each ontology, where each axiomatic sentence is translated to its correspondent in the target ontology.

Several works ([Healy and Caudel 2006], [Zimmermann et al. 2006], [Cafezeiro and Haeusler 2007], [Seremeti and Kougias 2013]) indicate that morphisms are directed from the less informative ontology to the more informative one. This agrees with the intuitive notions behind the understanding of categorical product as ontology alignment, pushout as merge and pullback as similarity search.

Although no morphism definition is provided, [Hitzler et al. 2005] present possible conditions for suitable morphisms, however without much clarification or analysis. The presented conditions are:

1. The preservation of class hierarchies,
2. The preservation of types,
3. The taking into account of model-theoretic logical properties, if featured by the underlying ontology representation language
4. The taking into account of proof-theoretic properties, and
5. The preservation of language classes.

These criteria will be further analyzed in the next section, where we discuss what makes a good ontology morphism and propose our own set of conditions.

4. What is a Good Morphism?

The main goal that should be achieved by an ontology morphism is relating meaningfully the concepts present in source and target ontologies. This requires loss of information not being allowed, or at least severely restricted. A morphism definition that, for example, allows all concepts in a source ontology to be mapped to a single concept in the target ontology is not advisable, and also probably not useful at all.

The ability to represent complex relations is also desirable, for it entails greater expressiveness and semantic power. Without such capabilities, any alignment between two ontologies with related but not equivalent concepts is doomed to be represented as some particularly complicated categorical structure, such as the W-alignment. These intricate structures require greater effort to be operated on: the merge through a W-alignment, for example, is performed through three successive pushouts with the requirement of a bridge ontology, while the same operation is performed with a single pushout over a V-alignment.

From the criteria outlined by [Hitzler et al. 2005] and listed in the previous section, the condition (1), while satisfied by most ontology morphisms found in the literature due to monotonicity, is arguably not strong enough. By itself, it does not prevent multiple concepts of the same hierarchical level in the source ontology from being projected onto a single concept in the target ontology. This causes information to be lost when traversing the morphism, contradicting our initial assumptions on what a morphism should achieve, as well as the commonly found belief that morphisms should point to the more informative ontology.

Meanwhile, condition (2) needs to be better specified about what is the intended meaning of "type". One possible way to define these types is through the use of meta-properties (that is, properties of properties) such as the ones utilized by [Guarino and Welty 2009] in their OntoClean methodology, which bases its ontological analysis in the meta-properties of rigidity, identity, unity and dependence.

Conditions (3) and (5) are dependent on the ontology representation language used, and thus are not useful guidelines for more general, implementation independent morphisms. Condition (4) requires a deeper discussion on proof-theoretic properties that is beyond the scope of this paper. None of these five conditions account for the representation of complex relations.

In the following subsections, we discuss and propose desirable aspects for good morphisms, regarding the preservation of information and the capability of representing complex relations, and present our criteria for the identification of suitable morphisms.

4.1. Information Preservation

As previously discussed, one of the most common expectations on ontology morphisms is that they are directed from a less informative source ontology towards a more informative target ontology. This means that a good definition of ontology morphisms should not allow information to be lost when transitioning between ontologies through morphisms. This entails the preservation of multiple forms of ontological information. Bellow, we discuss three of these forms: the preservation of concepts, relations and meta-properties.

4.1.1. Concept Preservation

If an ontology models two concepts as separate entities, this knowledge should not be muddled by the morphism. That is, two different concepts in the source ontology must be mapped to two different concepts in the target ontology, i.e., the morphism should be injective in regard to concepts. Mathematically, given a morphism $f: O \rightarrow O'$ and C_1, C_2 concepts in O,

$$f(C_1) = f(C_2) \rightarrow C_1 = C_2$$

4.1.2. Relation Preservation

Similarly, if an ontology models two relations separately, the ontology morphism should translate the relations from the source ontology to different relations in the target ontology, even if both relations share domains and codomains. That is, given a morphism $f: O \rightarrow O'$ and two relations R_1, R_2 in O,

$$f(R_1) = f(R_2) \rightarrow R_1 = R_2$$

Additionally, morphisms should not confuse domains nor codomains of relations. Thus, given a morphism $f: O \rightarrow O'$, a relation R relating two concepts C_1 and C_2 in O and a relation R' in O' such that $f(R) = R'$,

$$R(C_1, C_2) \rightarrow R'(f(C_1), f(C_2))$$

4.1.3. Meta-property Preservation

While concept and relation preservation should cover a great part of the information contained in ontologies, there are still pieces of ontological knowledge that could be confused if morphisms are modeled without additional care. This is the case of ontological arranges that are similar structurally, but carry deeply different semantic meanings. One example of such structure (shown in figure 2) would be the relations between a Car and the Factory where it was built and a Student and the University where he or she studies. While both the Car and the Student depend respectively on the Factory and on the University, this dependence takes a different form for each of these cases. The Student's dependence to the University is a relational one, i.e., the Student can only exist as such while there exists a University to which he is related. Otherwise, the Car's dependence to the Factory is purely historic, meaning that for the Car to exist, there once must have existed a Factory to build it, but if the Factory ceases to exist in any posterior moment it makes no difference to the existence of the Car [Guizzardi and Wagner 2010].

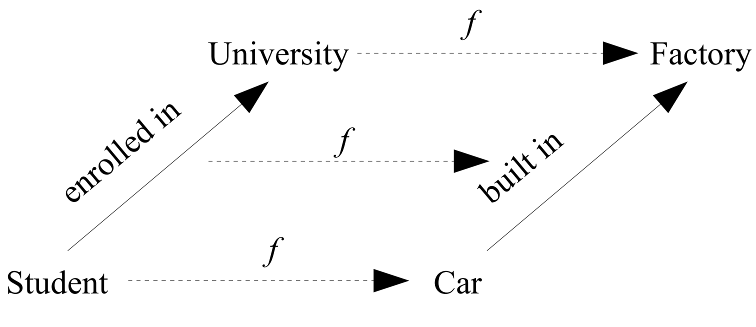

Figure 2. A morphism *f* relating concepts Student and Car

Another difference between the two concepts is due to the meta-property of rigidity: the concept Car is rigid, while Student is not. This means that a Car cannot cease to be a Car without ceasing to exist, while someone who is a Student can simply terminate their relation to the University and stop being a Student while keeping its existence.

Mismatches such as the one presented here can be avoided with the use of morphisms that preserve meta-properties of concepts.

4.2. Representation of Complex Relations

As previously noted, the limitation of associating concepts only through equivalence relations results in additional efforts when dealing with ontologies that do not share equivalent concepts, but contain ones that are otherwise related. [Zimmerman et al. 2006] perceived this issue and proposed a solution through the expansion of the definition of morphisms to accept also relations of strict inclusion, strict containment, disjointness and overlapping with partial disjointness.

This solution, nevertheless, is limited only to subsumption relations, which, despite covering the main taxonomical backbone upon which ontologies are built, disregards many other possible relations. It would be desirable to define morphisms that deal with mereological, dependence or even temporal relations, in addition to the already covered subsumption. For this, a set of relations must be chosen to be represented by the morphism, and the composition operation needs to be defined between each of these relations.

4.3. Criteria for Suitable Morphisms

The reasoning presented thus far can be summed up by five rules that should ideally be satisfied by adequate morphisms:

1. Concepts should be associated injectively;
2. Relations should also be associated injectively;
3. Domains and codomains of relations should be preserved;
4. Meta-properties of concepts should be preserved;
5. Morphisms should be able to represent relations other than equivalence between concepts.

The inspected morphisms do not fulfill most of these conditions. In particular, conditions (1), (2) and (4) are not satisfied in any of the reviewed works. Condition (3) is completely met only in the work of [Cafezeiro and Haeusler 2007] while partially met in the remainder works, where it is only true for taxonomic relations. Condition (5) is contemplated only by [Zimmermann et al. 2006] and [Euzenat 2008], but even in these works the expressiveness of morphisms has been expanded only to subsumption relations.

5. Future Challenges

Though many works in the area have been published, some challenges still remain to be faced when defining ontology morphisms. The first and most evident is relative to the definition of ontology on which the morphisms are based. Many of the works discussed here ([Bench-Capon and Malcom 1999], [Healy and Caudel 2006], [Cafezeiro and Haeusler 2007]) provide their own definitions of ontology and build their morphisms accordingly. These definitions are not necessarily compatible; hence the same is true for the morphisms. Additionally, these ontology definitions may not be fully reconcilable with available ontology representation languages.

Another challenge refers to these languages. Apart from the two conditions presented by [Hitzler et al. 2005] that explicitly deal with ontology representation languages, the problem surfaces also in relation to meta-property preservation, since it depends on the support of said meta-properties being provided by the underlying representation language.

A third challenge, and maybe of even greater concern, is related to another important yet often overlooked part of ontologies, the axioms. For an ontology morphism definition to be complete, it is necessary for it to describe how the morphisms deal with these axioms, yet such descriptions are seldom found in the literature.

Finally, there remains an open problem on the representation of complex relations, that is, the finding of a finite set of relations to be represented by morphisms and the definition of all possible compositions between them. A compromise must be reached, since the morphism expressiveness grows with the number of relations that can be represented, but so does the work needed to investigate the compositions between them.

6. Conclusion

While many works proposing category theoretical approaches to ontologies can be found in the literature, the studies on establishing meaningful morphisms have been scarce and rare. Several works have been presented which defined morphisms for ontologies, and a discussion was made on criteria for evaluating their semantic soundness and expressiveness. We have proposed five conditions to be met by suitable ontology morphisms. None of the morphisms analyzed fulfills satisfactorily these conditions.

Thus, much work is still needed on the search for meaningful ontology morphisms, and this work will face challenges related to the chosen definition of ontology, the underlying ontology representation language, how ontology axioms are to be dealt with and the composition rules on ontology relations.

References

Adámek, J., Herrlich, H., and Strecker, G. E. (1990). *Abtract and concrete categories: The joy of cats*. John Wiley & Sons, Inc., New York.

Bench-Capon, T. and Malcolm, G. (1999). Formalising ontologies and their relations. In *International Conference on Database and Expert Systems Applications*, pages 250-259. Springer.

Berners-Lee, T., Hendler, J., and Lassila, O. (2001) The semantic web. *Scientific American*, 284(5):34-43.

Cafezeiro, I. and Haeusler, E. H. (2007). Semantic interoperability via category theory. In *Tutorials, posters, panels and industrial contributors at the 26th international conference on Conceptual modeling-Volume 83*, pages 197-202. Australian Computer Society, Inc.

Cafezeiro, I., Viterbo, J., Rademaker, A., Haeusler, E. H., and Endler, M. (2014). Specifying ubiquitous systems through the algebra of contextualized ontologies. *The Knowledge Engineering Review*, 29(2):171-185.

Codescu, M., Mossakowski, T., and Kutz, O. (2014). A categorical approach to ontology alignment. In *Proceedings of the 9th International Conference on Ontology Matching-Volume 1317*, pages 1-12.

Euzenat, J. (2008). Algebras of ontology alignment relations. In *International Semantic Web Conference*, pages 387-402. Springer.

Goguen, J. A. and Meseguer, J. (1992). Order-sorted algebra I: Equational Deduction for multiple inheritance, overloading, exceptions and partial operations. *Theoretical Computer Science*, 105(2):217-273.

Guarino, N. and Welty, C. A. (2009). An overview of ontoclean. In *Handbook on ontologies*, pages 201-220. Springer.

Guizzardi, G. and Wagner, G. (2010). Using the unified foundational ontology (UFO) as a foundation for general conceptual modeling languages. In *Theory and Applications of Ontology: Computer Applications*, pages 175-196. Springer.

Healy, M. J. and Caudell, T. P. (2006). Ontologies and worlds in category theory: implications for neural systems. *Axiomathes*, 16(1-2):165-214.

Hitzler, P., Krötzsch, M., Ehrig, M., and Sure, Y. (2005). What is ontology merging? In *American Association for Artificial Intelligence*.

Seremeti, L. and Kougias, I. (2013). Networks of aligned ontologies through a category theoretic approach. *Journals of Emerging Trends on Computing and Information Sciences*, 4(10).

Wang, Q., Rong, L., and Yu, K. (2008). A pushouts based knowledge merging method: Knowledge reorganization in emergency decision-making support. In *4th International Conference on Wireless Communications, Networking and Mobile Computing, 2008. WiCOM'08*, pages 1-4. IEEE.

Zimmermann, A., Krötzsch, M., Euzenat, J., and Hitzler, P. (2006). Formalizing ontology alignment and its operations with category theory. *Formal Ontology in Information Systems: Proceedings of the Fourth International Conference*.