An Analysis of the Origin of Ontology Mismatches on the Semantic Web

Paul R. Smart[1] and Paula C. Engelbrecht[2]

[1] School of Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, United Kingdom
ps02v@ecs.soton.ac.uk
[2] School of Psychology, Shackleton Building, University of Southampton, Southampton, SO17 1BJ, United Kingdom
pce205@soton.ac.uk

Abstract. Despite the potential of domain ontologies to provide consensual representations of domain-relevant knowledge, the open, distributed and decentralized nature of the Semantic Web means that individuals will rarely, if ever, countenance a common set of terminological and representational commitments during the ontology design process. More often than not, differences between ontologies are likely to occur, and this is the case even when the ontologies describe identical or overlapping domains of interest. Differences between ontologies are often referred to as ontology mismatches and there is an extensive research literature geared towards the technology-mediated reconciliation of such mismatches. Our approach in the current paper is not to comment on the relative merits or demerits of the various technological solutions that could be used to resolve ontological differences; rather, we aim to explore the reasons why such differences may arise in the first place. In addition to a review of the various factors that contribute to ontology mismatches on the Semantic Web, we also discuss a number of focus areas for future research in this area. An improved understanding of the origins of ontology mismatches will, we argue, complement existing research into semantic integration techniques. In particular, by understanding more about the complex cognitive, epistemic and socio-cultural factors associated with the ontology development process, we may be able to develop knowledge acquisition and modeling tools/techniques that attenuate the impact of ontology mismatches for large-scale information sharing and data integration on the Semantic Web.

Keywords: semantic web, ontology, owl, semantic integration, ontology alignment, human cognition, ontology mismatches, ontology reconciliation, conceptual processing, knowledge acquisition techniques

1 Introduction

The Semantic Web [1] is a vision of the future potential of the World Wide Web (WWW) to provide a global infrastructure for the representation, dissemination and exploitation of human knowledge. Ontologies are a central element of this vision, providing the representational bedrock upon which advanced knowledge services and knowledge processing capabilities can be delivered. However, one of the problems associated with knowledge exploitation on the Semantic Web is the rather heterogeneous nature of the ontologies that are constructed to represent domain-relevant knowledge. The basic problem is that the open, distributed and decentralized nature of the Semantic Web tends to encourage the development of ontologies by developers who need not (and generally will not) countenance a common set of terminological and representational commitments during the ontology design process. Additionally, developers of ontologies will differ in their understanding of the meaning of a given term (and of the concept it denotes); which, in turn, will influence the kind of attributes and relationships they choose to include in the ontology. What all this means is that differences between ontologies are likely to occur, even in cases where the ontologies target the same domain of discourse or area of expertise. The result is an inability to integrate information across ontologies, thereby limiting the kind of information resources that can be exploited by knowledge processors. Such differences are often referred to as ontology mismatches [2], and the process of identifying and reconciling these differences is typically referred to as ontology reconciliation [2] or, sometimes, ontology mediation [3]. The ability to deal with the structural and semantic heteromorphism of ontologies on the Semantic Web is arguably of central importance with respect to the full exploitation of the Semantic Web as a system for information dissemination and automated knowledge processing.

A number of types of ontology mismatch have been identified in the course of previous research [2, 4, 5] and considerable efforts have been marshaled to support the technological resolution of such mismatches [e.g. 6]; very little research, however, has been devoted to understanding why ontology mismatches might arise in the first place. Two ontologies developed to represent the conceptual infrastructure of a common domain might be expected to show a number of differences, perhaps at the lexical or terminological level (i.e. the linguistic labels used to refer to specific concepts and relations), but why should ontologies often manifest more profound differences, differences that are grounded in the differential use of ontology modeling formalisms to express epistemic content? In this paper we present some factors that might contribute to the emergence of ontology mismatches; we then follow this up with a number of specific proposals for further research that could serve to improve our scientific understanding of the complex web of cognitive, epistemic and socio-cultural influences associated with the ontology development process.

2 Ontology Mismatches

Ontology mismatches are commonplace on the Semantic Web. Seldom, if ever, will any two ontologies developed for a common domain of discourse show strict (or even vague) isomorphism with respect to the representational formalisms used to represent and communicate domain-relevant knowledge. And such differences are not always limited to the realm of terminology. Often the mismatches between two ontologies can be so profound as to severely limit the prospect of discovering an effective ontology alignment and mapping solution. Consider, for example, the results of a small-scale pilot study that we undertook to explore the relative frequency of occurrence of different types of ontology mismatch. In this study, five individuals were asked to develop OWL ontologies to represent information in the domain of terrorist incidents. The specific details of the study need not concern us here, but what is interesting is the range of ontology mismatches that emerged in just this one small study. To take one example, we were interested in the ways in which subjects would elect to model information about suicide bomb attacks. All the subjects were provided with a sample set of reports about terrorist incidents, and at least some (five) of these incidents were about suicide bombings, i.e. the perpetrator of the attack voluntarily detonated an explosive device that resulted in his or her own death (irrespective of the death of innocent bystanders). The way in which this information was represented differed among all the ontologies, i.e. no two subjects settled on the same representational solution. Fig. 1 and Fig. 2 present the modeling approach adopted by two subjects in the pilot study. As can be seen from Fig. 1, OntologyA adopts an approach that is grounded in the use of a (datatype) property to indicate whether a terrorist attack is a suicide bomb attack. OntologyB, in contrast, represents the notion of a suicide bomb attack in a very different way; in this case, the notion of a suicide bomb attack is represented by casting the terrorist as a victim of the incident that they were the perpetrator of (see Fig. 2). Besides exemplifying one of the differences that can emerge between ontologies, this case is interesting because it is not entirely clear that the notion or concept of a suicide bomb attack, as represented in OntologyA, has any real counterpart in OntologyB. Even in cases where a viable interoperability solution can be seemingly derived [see 7], it is not always clear that such solutions preserve the intended meanings of the ontology authors. True, it is possible to represent instances of suicide bomb attacks using OntologyB, and such incidents can indeed be converted to instances in OntologyA using, for example, a rule that contingently creates instance data in OntologyA based on the data contained in OntologyB (see Fig. 3). But is it really appropriate to say that OntologyB *explicitly represents* the notion of a suicide bomb attack? It is difficult to say for sure. A particularly telling case is when we use the ontology to represent bombing incidents in which the suspect was killed during the attack; even though it was not clear that they intended to cause their own death (e.g. perhaps they were shot during the attack). In such cases, it seems, we confront incidents that stand a good chance of being 'successfully' converted from OntologyB to OntologyA without necessarily preserving the semantics implied by the original ontology (i.e. OntologyB in this case).

The kinds of ontology mismatches that are encountered during the course of semantic integration efforts have been investigated by number of authors [2, 4, 5]. Visser et al [5], for example, proposed a classification of ontology mismatches that distinguishes between conceptualization mismatches and explication mismatches (see Fig. 4). Conceptualization mismatches are mismatches involving the conceptualizations of a domain. They arise whenever we encounter differences in the kind of concepts that are identified in a domain and the way in which those concepts are related. Explication mismatches, in contrast, are mismatches in the way domain conceptualizations are defined, or explicated, in an ontology. Two ontologies could represent the same set of concepts, but because of explication differences they would appear to be quite different ontologies; in the extreme case, it may not even be clear that the ontologies are describing or representing the same things (the suicide bomb attack case, presented above, is a good example of this).

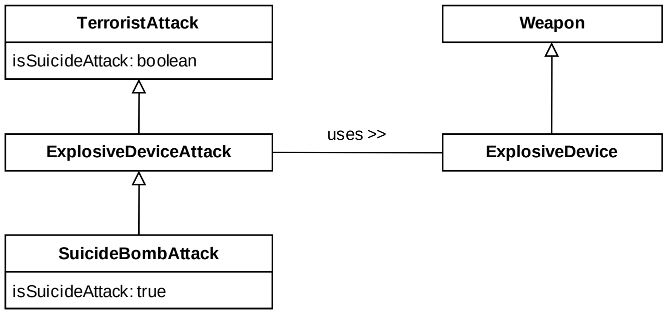

Fig. 1. Representation of a suicide bomb attack in the OntologyA

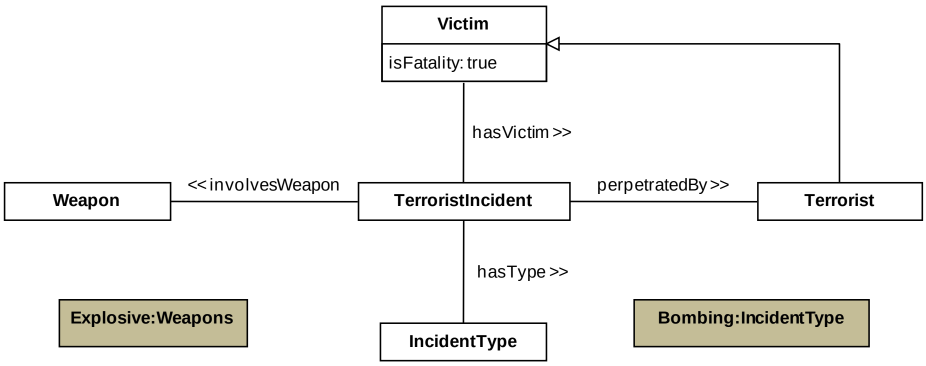

Fig. 2. Representation of a suicide bomb attack in OntologyB. The shaded boxes represent OWL individuals.

```
ontB:TerroristIncident(?x) ∧
ontB:hasType(?x, ontB:Bombing) ∧
ontB:involvesWeapon(?x, ontB:Explosive) ∧
ontB:hasVictim(?x, ?v) ∧
ontB:isFatality(?v, true) ∧
ontB:Terrorist(?v) ∧
ontB:perpetratedBy(?x, ?v)
 →
ontA:SuicideBombAttack(?t)
```

Fig. 3. SWRL rule expressing mapping relationship between the suicide bomb attacks in OntologyA (ontA) and OntologyB (ontB).

Given the potential complexity of ontology mismatches, coupled with their apparent importance in integrating and sharing information in the context of the WWW [see 8], it is perhaps not surprising that the weight of empirical research over the last decade has been devoted to the development of tools, techniques and algorithms that attempt to identify and resolve ontology mismatches [6, 9]. Relatively less effort, however, has been devoted to understanding why such mismatches occur in the first place. Clearly, there are some rather obvious (and perhaps not so obvious) candidates here, but until we begin to explore these factors in a scientific sense, our understanding of the relative significance of these factors to the emergence of (perhaps particular types) of ontology mismatch will always be somewhat limited. In the next section we begin to explore some of the factors that could (theoretically) contribute to the emergence of ontology mismatches on the Semantic Web.

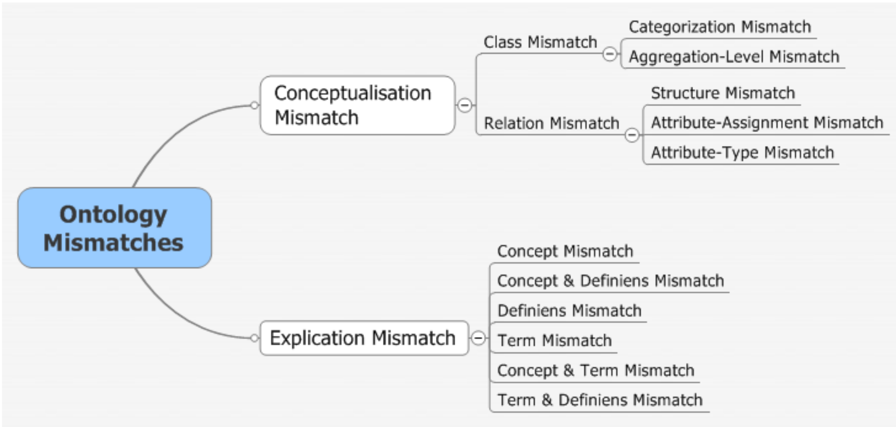

Fig. 4. Taxonomic hierarchy of ontology mismatch types as identified by Visser et al [5]. Concrete examples of each of these types of mismatch were provided by Hameed et al [2].

3 Factors Contributing to Ontology Mismatches

Why do ontology mismatches of the kind just described in the previous section occur, particularly in cases where individuals or agencies are attempting to model what is (or at least what seems to be) the same area of epistemic interest? We suggest that there may be a variety of reasons. They include the dynamic and context-sensitive nature of the human conceptual system, the influence of task-specific goals, previous experience with ontology authoring and the use of knowledge acquisition techniques that are differentially effective in eliciting particular types of knowledge. Each of these factors is explored in more detail below[1].

---

[1] It should be emphasized that we make no pretence about the completeness of our analysis as presented here. Clearly, a complete analysis of the factors that contribute to ontology mismatches is largely impossible given the current state of our scientific understanding in this area. Indeed, one of the motivations of this paper is to highlight the gaps in our current understanding and encourage a consideration of the opportunities for further research.

3.1 Dynamic Concepts

Perhaps the explanation that is *least* likely to come to mind when trying to account for ontology mismatches is the notion that the human conceptual system is, at root, a highly dynamic and context-sensitive system. The argument, first mooted (to our knowledge) by Barsalou [10], but then developed further by a number of theorists [e.g. 11], is that the computational substructure of the human conceptual system does not trade in anything that (*even vaguely*) resembles the stable, context-invariant, formalisms typically used to represent knowledge in the context of contemporary ontology engineering efforts. Rather, internal representations of concepts are argued to be the product of dynamic, context-dependent and (re)constructive processes. Barsalou [10] thus argues that concepts are actively constructed, rather than retrieved, and that "the *same concept* is rarely, if ever, constructed for a category" [10 pg. 101, our emphasis]. Barsalou's claim is supported by experimental findings that illustrate how people's conceptualizations of categories vary as a function of context, and that they are inherently unstable across relatively brief periods of time [10, 12, 13]. For example, it has been demonstrated that category construction can be based on the taxonomic (e.g. a sparrow is a bird) or thematic (e.g. cake and candles belong to the birthday party schema) relationships that exist between concepts, and that these constructive processes can be influenced by features of the task environment [13].

Another line of evidence in favour of a highly dynamic and context-infected conceptual processing system comes from Elman's work with recurrent neural networks [14]. What Elman discovered was that the processing of stable, symbolic lexical items was supported by a computational and representational system (a recurrent neural network) that traded in highly dynamic representations: "…the context always makes up an important part of the internal representation of the word, Indeed, it is somewhat misleading to speak of the hidden unit representations as word representations in the conventional sense, since these patterns also reflect the prior context. As a result it is literally the case that *every occurrence of a lexical item has a separate internal representation.*" [14 pg. 353, our emphasis]. Inasmuch as neural networks can be considered to capture the essential dynamics of the computational substructure undergirding much of human (and indeed biological) cognition, it would seem unlikely that the human brain trades in anything that can be called a stable and context-invariant conceptual structure.

Why are notions of conceptual instability and conceptual dynamism of potential relevance to an understanding of the mechanisms by which ontology mismatches emerge. One answer is that inasmuch as the human conceptual system does indeed trade in dynamic and context-sensitive conceptual structures – structures that are (re)created on-demand rather than stored and retrieved as syntactically distinct entities – the potential for differences to emerge in the modeling of domain conceptualizations looms large. It looms large because in such cases it seems unlikely that we genuinely confront a system (the human conceptual system) that ever possesses the same concept on more than one occasion, let alone shares that concept with other (biological/non-biological) conceptual processing systems. Each time a concept is encountered it will be different, albeit in perhaps minor ways, and such differences will be exacerbated as the context for conceptual processing changes over time. If it is not possible to possess the same concept on more than one occasion, if indeed every occurrence of a concept is genuinely different, then it seems unlikely that our static and context-independent models of the conceptualizations associated with a domain could ever be more than mere approximations to the true reality of human knowledge representation and conceptual processing. Moreover, inasmuch as such (knowledge-level) models do attempt to reflect the kind of concepts we have about a domain of discourse, then it is not particularly surprising that we should encounter significant differences in the way in which such models are designed, the kinds of things they represent and their ability to stand up to scrutiny in a multiplicity of contextually and temporally disparate processing/evaluation[2] contexts.

3.2 Concept Representation

Concept representation is the natural bedfellow of conceptual dynamism. It is the claim that ontology mismatches arise because the nature of the representational system being used by the Semantic Web is radically different from that used by human agents in formulating concepts and reasoning about a domain. This claim clearly has much in common with the argument that the human conceptual system is, at root, highly dynamic. The key difference between the two claims is that we could admit that domain concepts are highly dynamic and context-sensitive without necessarily committing ourselves to the view that extant approaches to knowledge representation are wholly inadequate as a model of human conceptual processing; it is just that such models lack the dynamic, context-sensitive and reconstructive elements associated with much of human cognition. The claim from the perspective of concept representation is, however, somewhat different. It argues that the way in which conceptual knowledge is organized, represented and used within the human conceptual system is radically different to that suggested by contemporary approaches in ontology engineering.

Many of the representational commitments of the Semantic Web are consistent with a classical or defining-attribute view of concepts [15]. The classical view of concepts posits that the kind of knowledge we have when we grasp a concept such as '*bird*' is knowledge of some necessary and sufficient defining conditions for category membership (e.g. '*anything that has feathers is a type of bird*'). The classical view of concepts, however, has been undermined by findings that suggest we often judge category membership in probabilistic or scalar terms [16]. That is to say, an object may be viewed as falling under the concept '*bird*' to a greater or lesser extent. It has

---

[2] The notion of evaluation is significant here because it may shed light on the reasons why some ontology authors elect to develop a new ontology rather than (re)use an existing one (this is a common phenomenon in our experience as ontology developers). It is not entirely clear how widespread this phenomenon is, but clearly there is much to be said for empirical studies that explore the factors underpinning people's positive and negative evaluations of pre-existing ontologies; the whole notion of knowledge re-use is, of course, founded on the expectation that will favourably evaluate pre-existing ontologies and find them acceptable with respect to their own representational needs and concerns. If it is the case that ontologies always reflect the conceptual idiosyncrasies of the person who developed the ontology, then it may be that all pre-existing ontologies seem somewhat deficient with regard to our own conceptual models of the world.

also been pointed out, of course, that in many cases one cannot define category membership in terms of necessary and sufficient conditions (e.g. Wittgenstein's example of the concept '*game*').

What all this seems to suggest is that the representational substructure of the human conceptual system is not organized along the same lines as the various representational commitments embraced by the Semantic Web community. As such, the representational underpinnings of the Semantic Web may constitute a psychologically implausible model of human conceptual representation. Inasmuch as this is the case it is perhaps not surprising that we find it difficult to shoehorn our everyday conceptualizations into the logico-deductive apparatus of the Semantic Web. And in so doing, we suggest, all manner of morphosyntactic variations may arise with respect to the representation of conceptual knowledge. Note that our claim is not that ontologies are wholly inadequate with respect to the representation of (*at least some* types of) human knowledge; it is more that ontologies are, in general, (congenitally?) ill-equipped to adequately reflect the true nature of human knowledge representation.

3.3 Knowledge Elicitation Techniques

A further explanation for ontology mismatches emerges out of the literature on knowledge acquisition techniques. Knowledge acquisition techniques, it is suggested, are differentially effective at supporting the elicitation of particular types of knowledge [17-19]. So, for example, the use of standard interview techniques may be suitable in cases were knowledge is explicit and verbalizable, but may not be so good when it comes to more implicit knowledge for which the domain expert has no direct conscious introspective access. In these cases, other techniques, such as repertory grid and card sorting techniques may be more effective. Furthermore, it is not just that such techniques are differentially effective at eliciting knowledge along the implicit/explicit dimension (the so-called differential access hypothesis [17]), they may also be differentially-suited to the elicitation of particular types of knowledge construct, e.g. concepts, attributes, tasks, rules, etc. So, for example, laddered grids tend to be good at eliciting and capturing information about concept hierarchies, but they are perhaps not as good as card sorting techniques when it comes to the elicitation of attributes and their associated values.

What this literature on the apparent relationship between knowledge acquisition techniques and elicited knowledge suggests is that the use of a particular technique (or set of techniques) can have a significant impact on the *kind* of knowledge that is ultimately elicited. The differential use of knowledge acquisition techniques could therefore contribute to differences in the knowledge that gets represented (i.e. it may lead to a conceptualization-type mismatch [see 5]). Another, not altogether unrelated, concern is that the complex socio-technical relationships that exist between the domain expert, knowledge engineer and the tools, techniques, props, aids and artifacts associated with the knowledge acquisition process, may play an active role in shaping the internal representation of knowledge as possessed by the domain expert. At first blush this might seem quite a radical claim, but recall that if theories about the dynamic and context-sensitive nature of the human conceptual system are correct (and there is now considerable evidence from the study of human memory to suggest that mnemonic processes are indeed highly constructive in nature [20]) then the context provided by the knowledge acquisition process itself may contribute to the kind of concepts (and associated knowledge) that gets *created* or *constructed* in the head of the domain expert. Rather than see knowledge elicitation as a process of eliciting rather static bodies of domain knowledge from the long-term memory stores of a domain expert, it may be the case that both expert and engineer, in conjunction with a range of external resources (tools, props, aids, artifacts, etc.), may be working to actively reconstruct conceptual knowledge in highly context- and task-specific ways.

Other psychological evidence suggests that the kind of knowledge acquisition technique used by a knowledge engineer could (potentially) influence the kind of judgments an expert makes about domain-specific objects or object categories. Boroditsky [21], for example, reviews evidence that indicates how object similarity judgments can be influenced by object comparison procedures: similar objects appear *more similar* following comparison procedures, even if the comparison task involves the identification and enumeration of object differences. Such studies invite us to consider the impact of knowledge elicitation techniques (e.g. triadic elicitation, repertory grid or card sorting tasks, all of which involve the comparison of objects or object categories) on expert's perceptions of object similarity. Could such procedures trigger comparison-dependent shifts in object similarity judgments, judgments that subsequently make certain kinds of conceptual categories harder to discriminate? Clearly, without experimental studies it is difficult to say for sure. What is clear, however, is that there is at least a prima facie case to wonder whether the knowledge engineer (as well as the methodological and technological accoutrements of the knowledge acquisition process) could play a far greater role in influencing the kind of knowledge that gets elicited and the way in which it ultimately gets represented. Inasmuch as knowledge elicitation techniques and strategies differ from one acquisition context to another, we would expect such techniques to play a major role in contributing to the emergence of ontology mismatches.

3.4 Task Context

The task context in which ontology development occurs is, in all likelihood, a major influence on the scope, content and structure of the resulting ontology. Clearly, if the requirement of an ontology is simply to provide a taxonomy of domain terms, perhaps for the purposes of manual resource classification, then it is unlikely to be as semantically-rich as an ontology in the same domain that is intended to support more sophisticated forms of analysis. Often the representational focus for an ontology is determined by the purpose to which it is eventually put, and this means that ontologies may end up being somewhat tightly geared to specific task contexts. Task contexts are, however, not the only determinant of ontology mismatches. In the pilot study mentioned in Section 2, subjects were specifically asked *not* to anticipate the potential use of their ontologies for a particular application; nevertheless, a rich set of ontology mismatches still emerged.

3.5 Ontology Engineering Expertise

One of the factors that appeared relevant to the emergence of ontology mismatches in our small pilot study (see Section 2) was the relative expertise and experience of the ontology developer with respect to semantic technologies. It was clear that the models developed by experienced ontology engineers were different in a number of ways compared to those of less experienced subjects. For example, experienced ontology engineers tended to make greater use of local property restrictions, to avoid the assertion of property domains and ranges (global property restrictions), to rely on subsumption reasoners to dynamically compute taxonomic hierarchies, to create necessary and sufficient conditions that supported the automatic assortative processing of OWL individuals, and so on. Individuals with relatively low levels of ontology engineering experience tended to avoid using the full range of semantic axioms available as part of the OWL specification. In some cases this resulted in modeling outcomes whose status as a true ontology seemed somewhat questionable[3].

Despite these differences it is unclear to what extent relative expertise in ontology development is a significant factor in terms of the emergence of ontology mismatches. Not all of the aforementioned differences between the models of experienced and inexperienced ontology engineers would seem to have significant repercussions for ontology mismatches. The extent to which authors rely on subsumption reasoning, for example, does not necessarily affect the kind of concepts that actually get represented (something that might otherwise introduce a number of conceptualization-type mismatches[4]). In any case, it is unclear whether the variability between experienced ontology engineers (as a group) is any less than the variance between inexperienced ontology engineers when it comes to ontology mismatches. Ultimately, these issues can only be resolved with controlled empirical studies.

3.5 Domain Expertise

As with ontology engineering expertise, domain expertise might be expected to contribute to significant differences in the content, scope and structure of domain ontologies. Firstly, evidence from experimental psychology suggests a number of differences exist between experts and novices when it comes to the conceptualization of domain-relevant material. Experts, for example tend to have concepts whose boundaries are somewhat more fuzzy compared to those seen in novices [22]. Furthermore, experts tend to think about concepts differently because they possess more complex theories about a domain relative to novices. This difference in theoretical understanding enables experts to see relationships between concepts that are not immediately apparent to a novice. For example, whereas expert physicists are

---

[3] One example here is the notion of 'SuicideBombAttack' in Fig 1. The use of a Boolean attribute ('isSuicideAttack) is not sufficient to capture what it means for something to be a member of the class of 'SuicideBombAttack', and in this respect, the representation of the class may be deemed to be semantically impoverished.

[4] It might, however, affect the way in which concepts are defined in the ontology, thereby contributing to a number of explication-type mismatches.

able to group physics problems according to the principles used in solving them, novices focus much more on the problems' surface features [23].

In contrast to the influence of ontology engineering expertise, we would expect differences in domain expertise to result (primarily) in conceptualization, rather than explication, type mismatches [5]. This is because conceptualization-type mismatches relate to the conceptualizations associated with a domain (the kind of conceptualizations established and the relationships between them), whereas explication-type mismatches relate to the way a conceptualization is specified (described or defined) in an ontology. Clearly, one would expect ontology engineering expertise to exert a greater influence on explication-type mismatches, as opposed to conceptualization-type mismatches, primarily because it is typically the ontology engineer that has primary responsibility for the formalization of domain-relevant knowledge. In cases where an individual plays the combined role of both ontology engineer and domain expert (a role that, in our experience, is not particularly uncommon in the ontology engineering community), then we would expect *both* types of ontology mismatch to occur. This suggests that even in cases where an individual is qualified to both create an ontology and serve as the source of expertise for that ontology, these roles should nevertheless be kept separate and distinct. To combine them arguably risks conflating two distinct (but potentially significant) sources of ontology mismatch.

3.8 Domain Experts

Differences in the level of domain expertise are no doubt significant when it comes to the elicitation of domain-relevant knowledge, but another factor for consideration relates to the differences between experts themselves, even experts that work within the same of domain of interest. Shadbolt and Burton [18], for example, identify three types of expert – the academic, the practitioner and the samurai – each of whose knowledge differs with respect to both its "internal structure and its external manifestation" [19. pg. 189]. Schreiber et al [19] suggest that differences between these expert-types motivates the use of different knowledge elicitation techniques, and that "Ignoring the nature of your expert is a potential pitfall in knowledge elicitation". Inasmuch as different experts pose different challenges for knowledge engineers *qua* the vagaries of their domain-relevant knowledge and conceptual models, then the characteristics of domain experts are clearly another factor for consideration when it comes to the causal processes associated with the emergence of ontology mismatches.

3.8 Conclusion

Clearly, the list of factors presented in this section does not exhaust those that could contribute to ontology mismatches[5] (and there are probably some factors

[5] Other factors not discussed here include the role of knowledge source materials (which may contribute to different perspectives on a knowledge domain or provide access to different bodies of domain knowledge), as well as cultural and linguistic influences.

presented here whose significance is overstated). Nevertheless, we now have a range of candidate factors to consider in accounting for ontology mismatches. In the next section we attempt to specify a series of studies that could be undertaken to evaluate the relative contribution of these factors to the emergence of various types of ontology mismatch.

4 Suggestions for Further Research

This paper has explored a variety of factors that may contribute to the emergence of ontology mismatches. It is clear, however, that much more empirical work still needs to be done to further our scientific understanding of the causal processes associated with the emergence of ontology mismatches. In this section we briefly describe a number of studies that could (or should) be undertaken as part of this research effort.

4.1 Formalization of Ontology Mismatches

Empirical analyses of ontology mismatches depend on an ability to identify and categorize ontology mismatches. Much progress has already been made in this area [e.g. 5], but it is not clear whether our current understanding of the types of mismatch that can arise in the course of ontology development is necessarily complete. In particular, it is not entirely clear whether we have a conceptual model of ontology mismatches that is not just adequate in terms of the types of mismatches that can be identified, but whether we also have a model that is adequate with respect to the *characterization* of those mismatch types. The value of this research is twofold. Firstly, a deeper characterization of ontology mismatches enriches the set of variables that can be studied as part of future empirical analyses. Secondly, formal specifications of ontology mismatches may contribute to semantic integration techniques that attempt to adapt the parameters of particular ontology alignment algorithms in order to optimize their performance particular situations [24].

4.2 Frequency Analysis of Ontology Mismatch Types

An understanding of the relative frequency of occurrence of ontology mismatches is of considerable practical importance because it suggests that certain forms of ontology reconciliation may be more important than others. Inasmuch as certain types of ontology mismatch occur with greater frequency, and inasmuch as certain types of semantic integration technique are differentially suited to resolving particular types of ontology mismatch [see 24, 25], then the relative frequency of ontology mismatches may serve as a guidepost in countenancing the most profitable area for subsequent semantic integration research. If it could be demonstrated, for example, that terminological differences were the most common type of ontology mismatch, then any advance in the state-of-the-art with respect to the reconciliation of terminological differences would have the greatest impact in terms of our prospective semantic integration capability[6].

4.3 Experimental Analyses of the Factors Underpinning the Emergence of Ontology Mismatches

In order to understand the relative contribution of a variety of factors to the emergence of ontology mismatches, it will be important to undertake controlled empirical analyses that investigate the relationship between independent variables such as level of ontology engineering expertise, level of domain expertise, use of particular knowledge elicitation techniques, etc., and a number of dependent variables, most notably the number of different types of ontology mismatch that are encountered in each experimental condition. Clearly, there are a number of challenges here, not least the operationalization of rather vague notions such as ontology engineering expertise; however, an ability to analyze the relationship between these variables in controlled empirical settings would lead to a number of advances in our understanding of the causal processes contributing to ontology mismatches. Specific examples of studies that might be undertaken include (but are obviously not limited to) the following:

- **Source materials.** What is the effect of access to different knowledge source materials on the type and relative frequency of ontology mismatches?
- **KA techniques and materials.** What effect does the use of different knowledge acquisition techniques have on ontology mismatches? How do the various outputs of knowledge elicitation – the laddered grids, repertory grids, interview protocols and so on – constrain, guide or otherwise influence the representational strategies adopted during subsequent ontology development?[7]
- **Ontology development environments.** There are a number of ontology authoring environments now available, e.g. Protégé, Swoop and TopBraid Composer. What is the impact of these ontology authoring environments on the way in which domain conceptualizations are represented in ontologies?
- **Knowledge processing capabilities.** What impact does a specific set of application requirements or knowledge processing capabilities have on ontology development? Are ontologies that are developed to support specific capabilities any more or less different (in terms of ontology mismatches) than ontologies that are agnostic with respect to their eventual exploitation?

An understanding of the relative significance of various factors with respect to the emergence of ontology mismatches could play an important role in reducing the

[6] This assumes, of course, that all types of ontology mismatch are equally significant with respect to their interference with information exchange and interoperability processes, an assumption that may ultimately prove to be unfounded.

[7] Of particular relevance here are the results of the various Sisyphus projects, especially Sisyphus III 26. Jansen, M.G., Schreiber, A.T., Weilinga, B.J.: Rocky III -- Round 1 A Progress Report. 11th Workshop on Knowledge Acquisition, Modeling and Management, Banff, Alberta, Canada (1998). Sisyphus III generated a rich set of knowledge acquisition products from domain experts in the geology domain.

incidence of ontology mismatches on the Semantic Web. At the very least, we would expect such an understanding to assist with ontology reconciliation processes. If, for example, it could be shown that differences in ontology engineering expertise was a significant factor in the origin of ontology mismatches then effort could be directed to the development of technologies that assist users with respect to the creation of highly expressive ontologies. Additional interventions could include the dissemination of best practice ontology engineering guidelines and the development of methodologies that facilitate the creation of 'high-quality' ontologies[8].

Interventions for other factors could include: encouraging the use of better documentation for ontologies and their constituent elements, perhaps extending existing ontology editing tools to make such documentation easier to add; encouraging the inclusion of rationale information for various ontology modeling decisions, e.g. encouraging users to explain why an particular element was included in the ontology (this could help, in some case, to assess the task-specificity of the ontology); advocating the use of best practice guidelines about when particular knowledge elicitation techniques should be used; ranking knowledge sources for particular domains and encouraging the use of a standard set of knowledge sources for ontology development. Of course, the relative success of these interventions will be predicated on the extent to which things such as task context, knowledge sources and knowledge elicitation techniques actually contribute to ontology mismatches; ultimately, this can only be resolved by the use of controlled empirical analyses.

5 Conclusion

A number of types of ontology mismatch have been identified in the course of previous research [2, 4]; however very little research has been devoted to understanding why ontology mismatches arise in the first place. Two ontologies developed to represent a common domain might be expected to show a number of differences, perhaps at the lexical or terminological level (i.e. the linguistic labels used to refer to specific concepts and relations), but why should ontologies often manifest more profound differences, differences that are grounded in the differential use of ontology modeling formalisms to express epistemic content? In this paper we have explored a number of potential explanations for the emergence of ontology mismatches during the course of ontology development. In the absence of empirical studies it is difficult to comment on the relevance of many of these factors; however, the lack of empirical studies in this case is significant. It serves to highlight the current gaps in our understanding of the complex psycho-cognitive, socio-cultural and socio-technical processes that contribute to the development of ontologies on the Semantic Web. If we are to deal effectively with the problems posed by semantic integration and interoperability, then we need to have a greater understanding of the origins of ontology mismatches. The weight of scientific research over the past

[8] All this presupposes, of course, that 'high-quality' ontologies will manifest fewer mismatches than their lower quality counterparts. Even if this is not the case, however, it may still be true that semantically-expressive ontologies are somewhat easier to align than semantically-impoverished ones.

decade has been devoted to investigating technological solutions to the challenges posed by ontology mismatches, and such approaches have, it is true, yielded considerable success [6, 9]. Perhaps, however, it is time to explore a different strategy, one that places greater emphasis on understanding the root causes of ontology mismatches on the Semantic Web. Such an approach will, of course, not displace technological approaches, but it may serve to complement them in a number of useful and interesting ways. Ultimately, we suggest, an improved understanding of why ontology mismatches occur could contribute to the development of ontology authoring environments and knowledge acquisition techniques that are sensitive to the complex interplay between the various cognitive, epistemic and socio-cultural factors associated with the ontology development process. Such tools and approaches may help to minimize the occurrence of ontology mismatches and make those that do occur much more amenable to extant ontology reconciliation techniques.

Acknowledgements

This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon

References

1. Berners-Lee, T., Hendler, J., Lassila, O.: The Semantic Web. Scientific American 284 (2001) 34-43
2. Hameed, A., Preece, A., Sleeman, D.: Ontology reconciliation. In: Steffen, S., Studer, R. (eds.): Handbook on Ontologies. Springer Verlag, Berlin (2004)
3. de Bruijn, J., Ehrig, M., Feier, C., Martin-Recuerda, F., Scharffe, F., Weiten, M.: Ontology mediation, merging and aligning. In: Davies, J., Studer, R., Warren, P. (eds.): Semantic Web Technologies: Trends and Research in Ontology-Based Systems. John Wiley & Sons Ltd, Chichester, UK (2006)
4. Klein, M.: Combining and relating ontologies: an analysis of problems and solutions. International Joint Conferences on Artificial Intelligence (IJCAI): Workshop on Ontologies and Information Sharing, Seattle, USA (2001)
5. Visser, P.R.S., Jones, D.M., Bench-Capon, T.J.M., Shave, M.J.R.: An analysis of ontology mismatches: heterogeneity versus interoperability. AAAI 1997 Spring Symposium on Ontological Engineering, Stanford, USA (1997)
6. Kalfoglou, Y., Schorlemmer, M.: Ontology mapping: the state of the art. Knowledge Engineering Review 18 (2003) 1-31
7. Smart, P.R., Russell, A., Braines, D., Kalfoglou, Y., Bao, J., Shadbolt, N.R.: A Visual Approach to Semantic Query Design Using a Web-Based Graphical Query Designer. 16th International Conference on Knowledge Engineering and Knowledge Management (EKAW'08), Acitrezza, Catania, Italy (2008)

8. Smart, P.R., Shadbolt, N.R.: The Semantic Battlespace Infosphere: A Knowledge Infrastructure for Improved Coalition Inter-operability. 4th International Conference on Knowledge Systems for Coalition Operations (KSCO), Massachusetts, USA (2007)
9. Noy, N.F.: Semantic integration: A survey of ontology-based approaches. Sigmod Record 33 (2004) 65-70
10. Barsalou, L.W.: The instability of graded structure: Implications for the nature of concepts. In: Neisser, U. (ed.): Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization. Cambridge University Press, New York (1987)
11. Clark, A.: Associative Engines: Connectionism, Concepts, and Representational Change. MIT Press, Cambridge, Massachusetts (1993)
12. Roth, E.M., Shoben, E.J.: The effect of context on the structure of categories. Cognitive Psychology 15 (1983)
13. Lin, E.L., Murphy, G.L.: Thematic relations in adults' concepts. Journal of Experimental Psychology: General 130 (2001) 3-28
14. Elman, J.L.: Representation and Structure in Connectionist Models. In: Altmann, G.T.M. (ed.): Cognitive Models of Speech Processing: Psycholinguistic and Computational Perspectives MIT Press, Cambridge, Massachusetts, USA (1991)
15. Smith, E.E., Medin, D.L.: Categories and Concepts. Harvard University Press, Cambridge, Massachusetts, USA (1981)
16. Rosch, E.H.: Cognitive representations of semantic categories. Journal of Experimental Psychology: General 104 (1975) 192-233
17. Hoffman, R.R., Shadbolt, N.R., Burton, A.M., Klein, G.A.: Eliciting knowledge from experts: A methodological analysis. Organizational Behavior and Human Decision Processes 62 (1995) 129-158
18. Shadbolt, N.R., Burton, M.: Knowledge elicitation: A systematic approach. In: Wilson, J.R., Corlett, E.N. (eds.): Evaluation of Human Work: A Practical Ergonomics Methodology. Taylor & Francis, London, England (1995)
19. Schreiber, G., Akkermans, H., Anjewierden, A., de Hoog, R., Shadbolt, N.R., Van de Velde, W., Weilinga, B.: Knowledge Engineering and Management: The CommonKADS Methodology. MIT Press, Massachusetts (2000)
20. Schachter, D.L., Norman, K.A., Koustaal, W.: The cognitive neuroscience of constructive memory. Annual Review of Psychology 49 (1998) 289-318
21. Boroditsky, L.: Comparison and the development of knowledge. Cognition 102 (2007) 118-128
22. Murphy, G.L., C, W.J.: Changes in conceptual structure with expertise. Journal of Experimental Psychology: Learning, Memory and Cognition 10 (1984)
23. Chi, M.T., Feltovich, P.J., Glaser, R.: Categorization and representation of physics problems by experts and novices. Cognitive Science 5 (1981) 121-152
24. Ehrig, M., Sure, Y.: Adaptive Semantic Integration. In: Collard, M., Cavarero, J.-L. (eds.): Workshop on Ontologies-Based Techniques for DataBases and Information Systems (ODBIS) at the 31st International Conference on Very Large Databases (VLDB), Trondheim, Norway (2005)
25. Noy, N.F.: Tools for Mapping and Merging Ontologies. In: Steffen, S., Studer, R. (eds.): Handbook on Ontologies. Springer Verlag, Berlin (2004)
26. Jansen, M.G., Schreiber, A.T., Weilinga, B.J.: Rocky III -- Round 1 A Progress Report. 11th Workshop on Knowledge Acquisition, Modeling and Management, Banff, Alberta, Canada (1998)