

Journal of the Association for Information Systems

JAIS 

Research Article

An Experimental Study of the Effects of Representing Property Precedence on the Comprehension of Conceptual Schemas*

Jeffrey Parsons

Memorial University of Newfoundland
jeffreyp@mun.ca

Abstract

Conceptual modeling is the process of using a grammar to construct abstractions of relevant phenomena in a domain. The resulting conceptual schemas are intended to facilitate understanding of and communication about a domain during information systems requirements analysis and during design. Despite keen practitioner interest in conceptual modeling, there is general agreement that the modeling constructs comprising grammars lack theoretical foundations pertaining to what the constructs are intended to represent, which, in turn, inhibits our understanding of whether and why they are effective. This research contributes to our understanding of conceptual modeling grammars by proposing a theoretically-grounded approach for modeling an important aspect of the nature of properties of the phenomena of interest in a domain. Specifically, conceptual schemas typically fail to express explicitly the semantics that, when things possess particular properties, they must also possess certain other properties. This research uses Bunge's ontological notion of property precedence as the theoretical rationale for explicitly modeling this dependence in conceptual schema diagrams. We examine several forms of precedence, and propose an approach to representing one form in conceptual schemas. We present the results of a laboratory experiment that tests the impact of explicitly representing precedence on how well participants comprehend the semantics conveyed by a conceptual schema. The results indicate that modeling precedence explicitly improves the comprehension of domain semantics expressed in a diagram's structure, but has varying effects on subjects' confidence in their comprehension.

Keywords: *Conceptual Modeling, Property Precedence, Ontology, Experimental Design, Comprehension*

* Ananth Srinivasan was the accepting senior editor. This article was submitted on 1st February 2010 and went through two revisions.

Volume 12, Issue 6, pp. 441-462, June 2011

An Experimental Study of the Effects of Representing Property Precedence on the Comprehension of Conceptual Schemas

1. Introduction

Diagrams are important modeling tools in information systems analysis and design. In addition to modeling the implementation of an information system (e.g., data and software structure), they are used to create conceptual models of the application domain or subject matter that is to be supported by the information system (Wand & Weber, 2002). The kinds of information that can be expressed using the grammar of a diagrammatic modeling technique (the symbols available and rules for assembling them) constitute its expressiveness (ter Hofstede & van der Weide, 1993). In conceptual modeling, the “expressiveness” of a modeling grammar is the extent to which it contains constructs (or combinations thereof) that represent relevant domain semantics. Often, the scope of relevant domain semantics is determined by an ontology that describes the nature of the phenomena to be represented. Since conceptual modeling deals with representing aspects of the perceived real world (including organizational and social activities) that will be supported by an information system, ontology is arguably an appropriate theoretical foundation for identifying the phenomena that might need to be supported by a conceptual modeling grammar (Wand & Weber, 2002).

A recent survey of users of the Unified Modeling Language (UML)—a set of diagrammatic grammars originally developed for software modeling—found that several of these grammars (notably class diagrams, activity diagrams, and sequence diagrams) were frequently used for a variety of purposes (Dobing & Parsons, 2006; 2008). More than 70 percent of respondents in that study used these diagrams for aspects of systems design (clarifying understanding of the domain among members of the design team, developing specification for programmers, and maintenance), and more than half of respondents reported the same diagrams to be useful during early phases of systems development to communicate about and verify information systems requirements with clients or users. This is practical evidence of the perceived value of diagramming techniques in validating conceptual models at all stages of the development process. In a similar vein, Davies, Green, Rosemann, Indulska, and Gallo (2006) and Fettke (2009) examined conceptual modeling use in practice in Australian and German contexts, respectively. Although the studies found somewhat different patterns of use (perhaps due to cultural difference between the nations in which the studies were conducted and the elapsed time between the studies), both found that ER diagrams were the most widely used technique. Fettke (2009) further found that UML use had increased significantly over time among respondents and was the second most used technique.

The fact that these diagrams were found to be frequently used for communication and validation (Wand & Weber, 2002) points to the importance of *understanding whether and how the constructs used in the grammars express domain semantics clearly*. If they do, there is the potential to improve communication and validation during systems analysis and design. If they do not, communication and validation might be impaired. Since errors in requirements analysis have the greatest impact on systems outcomes such as cost and quality (Boehm, 1981; Byrd, Cossick, & Zmud, 1992), improvements in the ability of conceptual modeling grammars to express domain semantics has the potential to improve the degree to which systems meet client requirements.

Existing conceptual modeling grammars or diagram types lack theoretical foundations for modeling the domain or subject matter to be supported by an information system (Dobing & Parsons, 2000). For example, the Entity-Relationship model (Chen, 1976) distinguishes entities from relationships, but allows both to be described by attributes. From a theoretical perspective, some researchers have argued that this practice should be avoided, as it is ontologically unsound (Wand, Storey, & Weber, 1999; Burton-Jones & Weber, 1999, 2003). Therefore, it is useful to understand why the practice persists and is commonly used in textbooks that introduce conceptual modeling.

In this paper, we examine the value of directly expressing an ontological construct called *property precedence* (Bunge, 1977) in conceptual schemas, based on the premise that an explicit representation will improve the ontological clarity of the conceptual schema – the extent to which one modeling construct corresponds to a single ontological construct (Wand & Weber, 1993). ER and UML class diagrams focus on representing information about the properties of, and associations

among, objects in a domain of interest. Therefore, they are understandably among the most used in systems analysis and design (Davies et al., 2006; Dobing & Parsons, 2006; Fettke, 2009). Property precedence refers to a relationship between properties in which some can be considered specializations of others (Parsons & Wand, 2003). We present an experiment that tests the impact on comprehension of expressing precedence semantics explicitly in class diagrams. The results of this experiment demonstrate that modeling precedence explicitly can improve the comprehension of relevant domain semantics among readers of a conceptual schema diagram, and can under specific conditions lead to higher or lower confidence in that comprehension. We conclude by summarizing the substantive and methodological contributions of this research, and discussing future research opportunities.

2. Conceptual Modeling from an Ontological Perspective

A fundamental premise in conceptual modeling is that the scripts (e.g., diagrams) resulting from applying a conceptual modeling grammar represent human conceptions of aspects of the real world, as well as of the social world that emerges from human conceptions of interactions of material objects. We take the view that ontology – the branch of philosophy dealing with the nature of reality – can shed light on the real world phenomena that need to be represented in conceptual models. In particular, we follow Wand and Weber's position (1990, 1993, 2002; Weber, 1997) that the ontology of Mario Bunge (1977, 1979) provides a formal and useful framework for understanding the real world phenomena that may need to be captured during conceptual modeling. Implicit in this approach is the expectation that greater ontological fidelity in a conceptual modeling grammar will improve the understandability and/or usability of scripts generated using that grammar.

2.1. Ontological Foundations

We begin with two basic ontological assumptions. First, the world consists of things, or substantial individuals with a material existence. Second, things possess properties. The second assumption provides the underlying ontological justification for including in a conceptual modeling grammar a construct to represent properties; namely, conceptual modeling scripts represent the properties of things. In view of this, it is not surprising that the attribute construct is widely used in conceptual modeling to represent properties (Weber, 1997).

According to Bunge (1977), properties can be divided into two kinds. Intrinsic properties can be understood in terms of a single thing in isolation. For example, the weight of some object is generally independent of other objects in the world (at least with respect to the earth's surface). Mutual properties, however, have meaning only in the context of two or more things. For example, "owns" is a mutual property that is meaningful only in the context of at least two things (the owner and the owned). In a traditional conceptual modeling grammar such as the Entity Relationship model, intrinsic properties can be represented as attributes of entity types, while mutual properties can be represented as relationships or associations between entity types (Evermann 2005; Wand et al., 1999).

2.2. Ontological Understanding of Conceptual Modeling Grammars

Recent experimental research has endeavored to improve the effectiveness of conceptual modeling grammars by analyzing the ontological fidelity of specific conceptual modeling practices and evaluating the effects of following ontologically justified modeling prescriptions. Until now, this work has focused on several problems in conceptual modeling that become apparent when applying principles of Bunge's ontology, including: (1) the modeling of optional properties and (2) the modeling of properties of relationships/associations.

2.2.1. Ontological Status of Optional Attributes

Bunge's ontology includes the notion that failing to possess a property is not itself a property. For example, "not blue" is not a property of all the things that do not possess the property of "being blue." Moreover, things possess properties. Gemino and Wand (2005) used this idea to reason that classes in a conceptual schema should be modeled in terms of attributes that represent properties. They argued that optional attributes should not be modeled as attributes of a class in conceptual schemas; instead, such attributes should be shown as mandatory attributes for appropriate subclasses of the

class in question. They hypothesized that, since optional properties are inconsistent with ontological principles, performance on a problem solving task would be significantly higher for the mandatory properties group than for the optional properties group. Their results supported the proposition that grammars using subtypes with mandatory properties lead to a higher level of understanding than those using optional properties.

Bodart, Patel, Sim, and Weber (2001) also compared mandatory and optional properties in conceptual modeling and found similar results. They hypothesized first (second) that diagrams using optional properties (mandatory properties) would assist users undertaking tasks that required a surface level understanding (deep level understanding) of the domain better than those using only mandatory properties (optional properties). The results supported the hypotheses. The optional properties group outperformed the mandatory properties group in diagram reconstruction and comprehension questions; however, the mandatory properties group outperformed the optional properties group in problem-solving questions.

2.2.2. Ontological Status of Attributes of Relationships

Burton-Jones and Weber (1999, 2003) investigated the practice of representing properties of properties in conceptual modeling as attributes of relationships in conceptual schema diagrams. Based on Bunge's ontological position that "properties do not have properties," they hypothesized that modeling relationships with attributes would decrease readers' understanding of conceptual models and confidence in their ability to interpret model semantics (2003), and decrease problem-solving performance with respect to information in the models and perceived ease of understanding of diagrams (1999). They found support for the hypotheses related to comprehension, but found no effects on confidence (2003), as well as partial support for hypotheses regarding problem-solving performance (1999) – in an unfamiliar domain, ontological clarity (relationships without attributes) had a positive impact on problem-solving performance, but no effect on perceived ease of understanding of the diagrams.

These and other studies have provided useful guidance for constraining the use of conceptual modeling grammars in order to improve the ontological fidelity of resulting schemas and, thereby, improve the ability of humans to read and understand these diagrams. Next, we consider an additional important ontological principle that might contribute to more effective use of grammars.

2.3. Property Precedence

Burton-Jones and Weber (1999, 2003) based their studies on the ontological position that "properties do not have properties" (Bunge, 1977). While we agree with this position, we believe that a more precise understanding of semantics can be obtained by applying Bunge's notion of property precedence to some situations that might at first glance appear to reflect "properties of properties," such as relationships with attributes. Next, we outline the concept of property precedence, and use it as the basis for predicting that modeling this phenomenon explicitly will have a positive impact on diagram comprehension.

Bunge's ontology recognizes that intrinsic or mutual properties at one level can sometimes be understood in terms of properties expressed at another, more general, level (and vice versa). The connection between specific and generic properties is formalized via property precedence. Specifically, a property P1 is said to **precede** another property P2 if and only if every thing that possesses P2 also possesses P1. Any property that must be possessed by a thing possessing a property P is said to be a **preceding property** of P. If property P1 precedes property P2, this can be expressed by the equivalent statement: "P2 is preceded by P1." The following are examples of precedence applied to intrinsic and mutual properties:

Intrinsic: Having been "born May 1, 1970" is preceded by having been "born";

Mutual: Having a "contract termination date between a customer and supplier" is preceded by having a "contract between a supplier and a customer."

In many cases, the semantics of a preceded property can only be understood by including the preceding property. For example, the property "address" precedes the properties "mailbox number" and

“street number,” since every thing that has a mailbox number and every thing that has a street number has an address. In other words, the set of things that have a mailbox number is a subset of the set of things that have an address. In this example, while all things having a mailbox number have an address (address is necessary to understanding mailbox number), not all things having an address have a mailbox number (mailbox is not necessary to understanding address, as there are other specifications of address such as street address). We contend that such cases are typical in understanding precedence relations among properties in conceptual modeling.

Typically, a class (or type) in a conceptual model is defined by a set of properties (often represented as attributes) possessed by all instances of the class. Precedence can be used to understand the relationship between instances and classes. Often, classes are described by preceding properties (e.g., street), while the instances of classes possess specific properties preceded by the generic ones (e.g., a specific street name and number). The latter are typically not shown in a conceptual schema constructed using a modeling grammar. We say that the specific property *manifests* the generic one, and that manifestation is a special case of precedence. Different instances might possess different manifestations of the same generic property, as illustrated next.

According to Parsons and Wand (2003), there are two main categories of manifestation, both related to intrinsic properties. First, *value manifestation* occurs when a given generic property is manifested by a specific value. For example, having a specific height (e.g., 170cm) is preceded by the property “has height.” In this case, the value has meaning only in the context of the generic property. Second, *specialization manifestation* occurs when the generic property can be refined in ways such that the individual cases do have meaning independent of the generic property. For example “moves on land” may be manifested as “crawls,” “walks,” “hops,” or “runs.” Animals moving in any of these ways can be members of a class that includes in its definition the property “moves on land.” In both cases, there is a dependency between the preceded and preceding intrinsic properties, in which possessing the preceded property implies possessing the preceding property.

We propose a further category of precedence, termed *mutual property precedence*, which we believe can contribute significantly to the understanding of domain semantics in conceptual modeling. Mutual property precedence arises when a mutual property (e.g., has contract) either has several elements or aspects (e.g., contract start date, contract termination date, contract amount), or precedes other distinct mutual properties (e.g., makes [contract] payment). Such precedence relations commonly arise in conceptual modeling, as representing the associations or relationships between classes of things is a fundamental aspect of most conceptual modeling grammars (Wand et al., 1999). In the remainder of this paper, we examine whether the greater ontological fidelity achieved by representing mutual property precedence (hereafter, simply termed precedence) explicitly in conceptual schema diagrams affects the ability of readers of these diagrams (users) to verify the precedence semantics conveyed in the diagram.

Prior research has shown that greater clarity in conceptual schema diagrams can lead to greater diagram understanding (Gemino & Wand, 2005; Bowen, O’Farrell, & Rohde, 2006). The underlying argument in support of this finding is that clarity supports local reasoning about diagram semantics, based on the cognitive theory of multimedia learning (Mayer, 2001). We contend that representing precedence explicitly provides greater ontological clarity.

Based on the notion of precedence as put forth by Bunge (1977) and elaborated by Parsons and Wand (2003) and above, we propose:

Proposition 1: Readers of conceptual schema diagrams that express precedence explicitly will *better comprehend* domain semantics related to precedence than readers of semantically equivalent diagrams that do not express precedence explicitly.

Prior research has also examined the impact of theory-based manipulations of independent variables on subjects’ confidence in the correctness of their responses (e.g., Allen & March, 2006; Allen & Parsons, 2010; Burton-Jones & Weber, 2003), although there is a lack of theory in understanding how certain conditions might increase or decrease confidence. In the context of representing property

precedence explicitly, we expect that the greater ontological clarity provided by representing precedence explicitly in conceptual schema diagrams will lead to more confidence in one's understanding of precedence-based domain semantics expressed in a schema.

Proposition 2: Readers of conceptual schema diagrams that express precedence explicitly will be *more confident* in the correctness of their comprehension of domain semantics related to precedence than readers of semantically equivalent diagrams that do not express precedence explicitly.

These propositions offer an alternative interpretation to the earlier-mentioned position of Burton-Jones and Weber (2003) that “properties should not have properties.” While their interpretation is ontologically founded (according to Bunge, only things have properties), we propose that instead of necessarily interpreting attributes of a relationship in a conceptual schema diagram as “properties of properties,” in some cases, they should be interpreted as representing preceded properties of a more general preceding property. Thus, our work contributes to the ongoing body of research seeking to improve our understanding of conceptual modeling grammars by applying ontology-based reasoning.

3. Expressing Mutual Property Precedence in Conceptual Schemas

UML class diagrams and the ER model provide constructs (associations and relationships, respectively) that have been interpreted as appropriate for representing mutual properties (Evermann, 2005; Wand et al., 1999). For example, in UML class diagrams, associations are depicted using labeled lines that connect classes, indicating links between instances of the connected classes. If additional information about the association needs to be captured, an association class (which may have attributes) may be linked to the association.

Figure 1 illustrates the use of associations and association classes with attributes. This notation differs somewhat from UML in that we attach labels and multiplicities to the various lines linking the object and association classes to better describe the links. For example, we use distinct labels for the association (Takes) and the association class (Section). Moreover, in Figure 1 the label Takes describes a generic association between Student and Course, while Enrolls in describes more specifically the nature of the Takes association between Students and Courses via the association class Section (the intended semantics refers to enrollments in sections at a particular time). Following Bunge, we do not model Section as a class of things (Evermann & Wand, 2005). Rather, the semantics of the association class indicates that a Section has no meaning independent of the students and courses it connects (i.e., in the example, it represents enrollment). The cardinality constraints imply that an object on the left side of the association is associated with (lower, upper) objects on the right side of the association. For example, a Student Enrolls in one to five Sections, and a Section has zero to many Students enrolled in it.

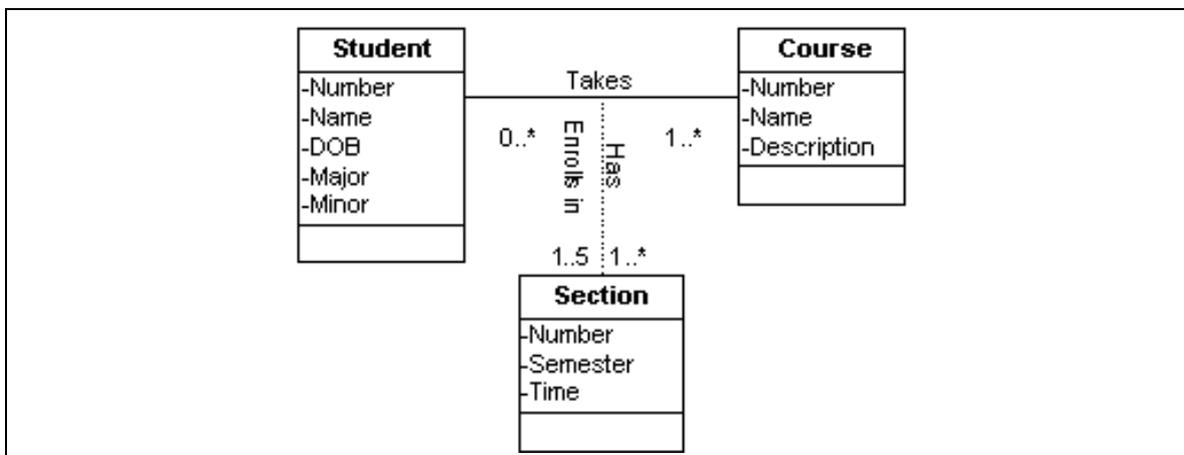


Figure 1. Associations and Association Classes

Mutual property precedence is expressed in Figure 1 by the attributes of the association class Section—Number, Semester, and Time—all of which are preceded by the existence of a Takes mutual property.

To allow a deeper hierarchy of mutual property precedence to be represented, we propose a notation to express precedence semantics by allowing association classes to be linked to other association classes. To illustrate, we extend the example in Figure 1 to the case where a Takes association (registration in a section of a course) has lab requirements, and a Lab is preceded by (or dependent on) a Section (see Figure 2). This reflects mutual property precedence as described in Section 2, since the participation in a Lab (i.e., the association of a Student with a Course through a Lab) implies the existence of a more general association between a Student and a Course through a Section. In other words, the student must be enrolled in the section in order to participate in the lab. Our proposal can be contrasted with an alternative approach that takes an associative entity perspective (Rob & Coronel, 2009), modeling the above in a class diagram via a series of associations by elevating Section and Lab from association classes to object classes.

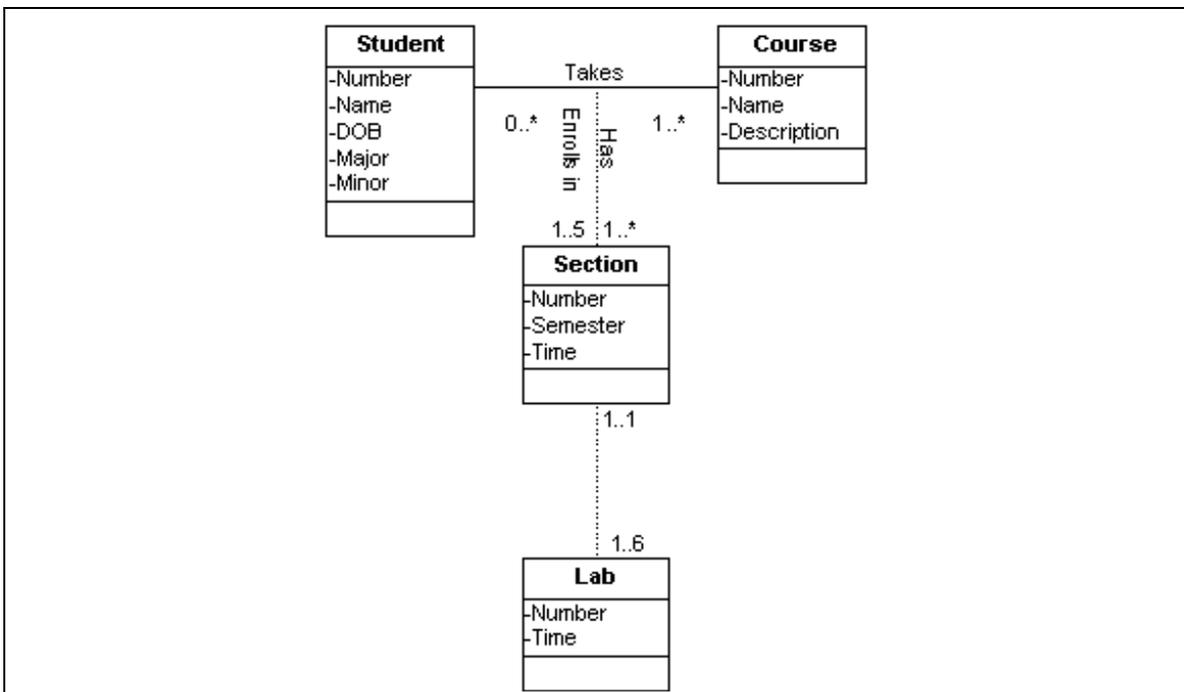


Figure 2. Expressing Specialization Precedence via Association Classes

Our research, then, focuses on whether expressing precedence explicitly, as in Figure 2, improves the understandability of precedence semantics communicated in a conceptual schema, relative to an informationally equivalent representation that provides only an implicit representation of precedence semantics (as will be explained later).¹ The next section describes an experiment designed to answer aspects of this question.

Although we have proposed a notation to represent mutual property precedence explicitly, our concern is neither with the form of the notation nor the particular choice of grammar (such as UML).² Alternate ways of expressing precedence might be devised. Instead, we are interested in using an explicit representation to compare whether this affects the ability of participants to

¹ In practice, there can likely be degrees of “explicitness,” such that it reflects a continuum, rather than a binary variable. More explicit treatments are possible in which the notion of precedence is explained and formally included in modeling notations, which we did not do. Thus, the treatments here are properly considered “more explicit” and “less explicit.” For simplicity, we use the terms “explicit” and “implicit.”

² For example, in UML 2.x, it is possible to model both Section and Lab of Figure 2 as association classes between student and course, and creating a “subset” or “specialization” between the two. This is similar to the approach proposed in what follows.

comprehend domain semantics expressed in a class diagram relative to an implicit representation. Such an extension could be applied to any conceptual modeling grammar that has constructs for representing properties.

4. Research Method

4.1. Experimental Design Issues

The proposition that diagrams expressing precedence explicitly will lead to better comprehension of domain semantics than diagrams that express precedence implicitly raises a number of issues with respect to designing an experiment to test the proposition effectively (Parsons & Cole, 2005). First, there might be many ways of expressing domain semantics explicitly. Our goal is not to compare alternative diagram formats that each express precedence explicitly (using different notation), since the underlying ontological framework does not speak to the question of which of several alternate graphical representations of precedence is expected to be better than others. Instead, we focus on precedence relations among mutual properties and use the approach outlined above as our primary treatment (explicit representation of precedence involving mutual properties). We use as the implicit representation a conceptual schema diagram in which ontologically mutual properties are represented as distinct object classes, and precedence semantics are implied by cardinality constraints (as discussed later). This approach is consistent with the use of associative entity types in data modeling (e.g., Hoffer, Prescott, & McFadden, 2008), and is typically used to represent situations in which relationships have attributes (Rob & Coronel, 2009).

Second, when comparing representation formats, we believe it is important that alternative conceptual schema diagrams generated to manipulate the independent variable be informationally equivalent with respect to the dependent measures (Gemino & Wand, 2004; Parsons & Cole, 2005). Recently, it has been argued that “informational equivalence cannot be verified conclusively” (Burton-Jones, Wand, & Weber, 2009, p. 504). Nevertheless, for the purpose of determining the clarity of possible modeling alternatives to represent specific domain semantics, it should be possible for participants to answer correctly questions involving comprehension of facts expressed in the diagram with any of the representational forms used as treatments in an experimental study. Otherwise, if one form provides enough information to answer selected questions correctly, while a second form does not, it would not be a valid test of the theory to find that participants receiving the first form outperform those receiving the second form on those questions. This does not mean there needs to be strict information equivalence between diagrams; however, with respect to questions asked related to the dependent variables, the representations should provide all information needed to answer the questions correctly. We ensured information equivalence by developing questions that could be shown to be answerable correctly from either representation. The Appendix contains the scripts used in one of the examples and explains how they are informationally equivalent with respect to the questions asked.

Third, domain knowledge can influence interpretations of the semantics expressed by a diagram's constructs (Khatri, Vessey, Ramesh, Clay, & Park, 2006; Parsons & Cole 2005). We proposed a notation to represent precedence explicitly and are interested in how respondents interpret the semantics conveyed by the notation. Using subject matter experts (SMEs) in tests of whether scripts convey semantics clearly can influence experimental results. For example, if we constructed a script related to the enrollment of students in courses at a university and conducted an experiment involving student participants (reasonably viewed as SMEs with respect to enrollment in courses at a university), participants might be able to answer some questions based on their experience at their own university, thereby mitigating any effects of differences in how semantics are expressed in modeling notation. Therefore, we wished to isolate the effects of structural semantics conveyed by the diagram syntax from additional semantics that might be inferred using background domain knowledge primed by the words used in a diagram. To do that, we used alternate forms of the experimental material. In one, we used English words from a domain with which most or all participants would be expected have some familiarity. In the other, we substitute Greek letters so that participants were able to answer questions only by using the semantics conveyed by the constructs and syntax in the diagram.

Finally, since we are interested in testing only the representation capacity of alternative forms of representation, we made scripts available to participants for review as they answer questions used to measure the dependent variables. As scripts normally would be available in a real-world communication/validation situation, there was no reason to remove diagrams in our study. In contrast, some prior work, focusing on using scripts to learn about the domain and integrate with deep knowledge structures, has used a protocol in which scripts are taken away from participants before measures of the dependent variables are collected (Bodart et al., 2001; Gemino & Wand 2003).

4.2. Task and Material

Given the objectives of the study, we used several small diagram segments focusing on precedence among mutual properties (instead of testing entire diagrams). We chose examples based on two domains: airline reservations and restaurant reservations. Figures 3 and 4 illustrate alternate representations for the airline reservation example. Both diagrams model logical associations (Flight and Reservation) or mutual properties between physical objects (Passengers and Aircraft). In Figure 3, the fact that Reservation is preceded by Flight, and that Flight is a mutual property linking Passenger and Aircraft, must be inferred by determining that a reservation involves exactly one flight and that a flight involves at least one passenger and exactly one aircraft. For example, to answer the question “must a reservation involve a specific aircraft?” one must first infer from the association cardinalities that a reservation involves exactly one flight, and a flight, in turn, involves exactly one aircraft.³ Thus, mutual property precedence is implicit or indirect in Figure 3 – it can be extracted only by reading the relevant cardinalities (in particular, lower bounds of at least one). In other words, the association construct is ontologically overloaded (Wand & Weber, 2002) in Figure 3, since it expresses aspects of both precedence and mutual property.

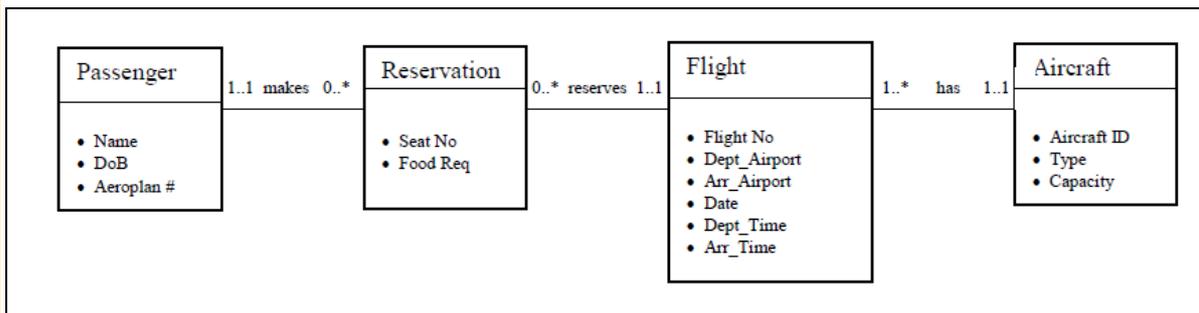


Figure 3. Implicit Representation of Precedence

Figure 4 introduces the additional notation we use to depict precedence explicitly. Although this is an atypical representation, the approach is consistent with earlier research using Bunge’s ontology to provide guidance for conceptual modeling, which has recommended using object classes in a conceptual model to represent only things or substantial objects (e.g., Evermann & Wand 2005; 2006). Thus, in Figure 4, Flight is shown as an association class and Reservation is shown as a second association linked to (i.e., preceded by) Flight. This means the existence of a reservation between a passenger and an aircraft implies the existence of a flight linking passengers and aircraft. In other words, the graphical notation implies that a reservation involves a flight, and a flight involves passengers and aircraft. In this respect, it is semantically equivalent to the precedence implicit in Figure 3; however, the graphical notation in Figure 4 makes the precedence explicit (via the dotted lines connected to association classes). Rather than relying on the cardinality construct to also express precedence, it uses association classes to represent the distinct ontological construct of property precedence. Note that there is some redundancy in this regard in Figure 4, since cardinalities are still shown on the links.

³ We are dealing here with semantics as expressed in the diagram. Note that applying “common sense” semantics could interfere here, as in many cases, an actual flight will involve different aircraft on different days. Thus, although one can distinguish actual and scheduled flights, the semantics of this example deals only with the former.

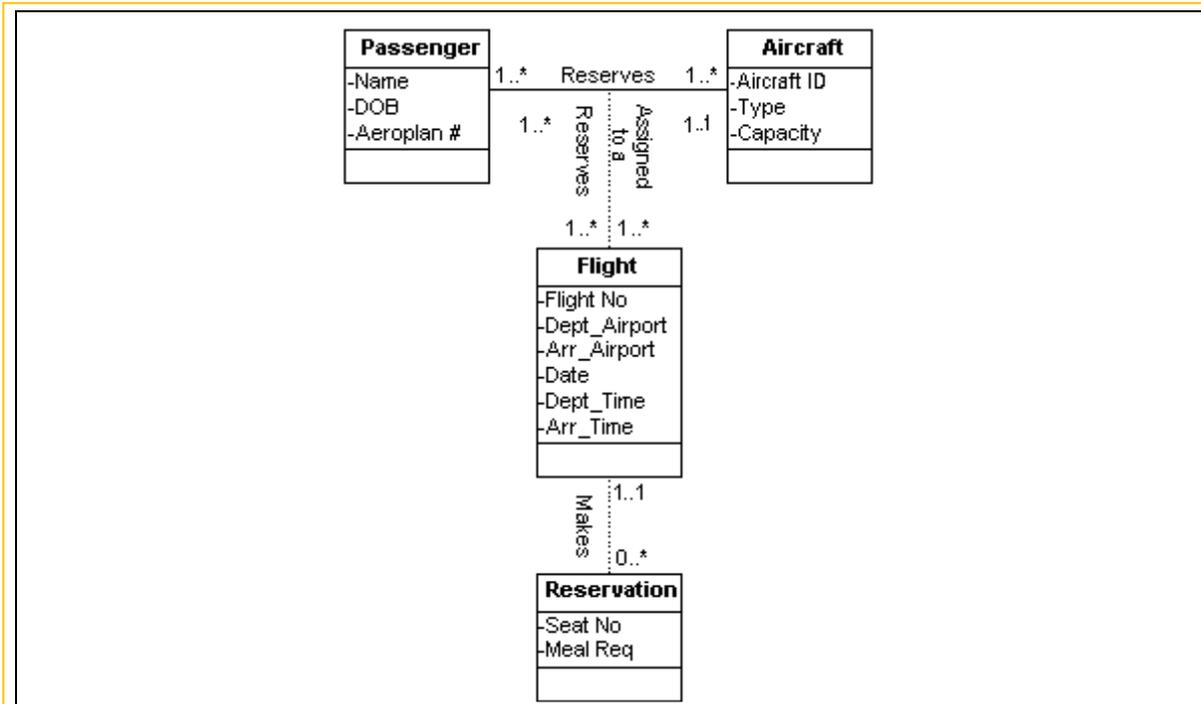


Figure 4. Explicit Representation of Precedence

We recognize that alternative equivalent representations are possible for this example. Indeed, the representation in Figure 4 may appear unnatural to those used to modeling logical associations (non-physical entities) as ER entity types or UML classes (e.g., Hoffer et al., 2007; Rob & Coronel, 2009). However, the approach in Figure 4 is consistent with prior research on ontological approaches to conceptual modeling (e.g., Evermann and Wand 2005; Wand et al., 1999).

To control for the effects of background knowledge on how domain semantics are understood in conceptual schema diagrams, we constructed alternative representations for which the class and association labels were semantically void. Figures 5 and 6 contain segments corresponding to those in Figures 3 and 4, respectively.

As is clear from Figures 5 and 6, answers to any comprehension questions about this artificial domain can be correctly determined only in relation to the semantics conveyed via the modeling constructs and syntax in the diagram. In contrast, in Figures 3 and 4, correct answers might be generated based on prior knowledge about airline reservations.

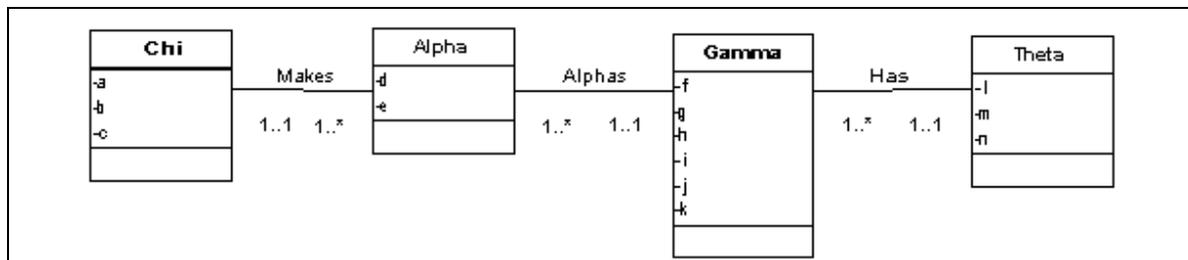


Figure 5. Semantically Void Equivalent of Figure 3

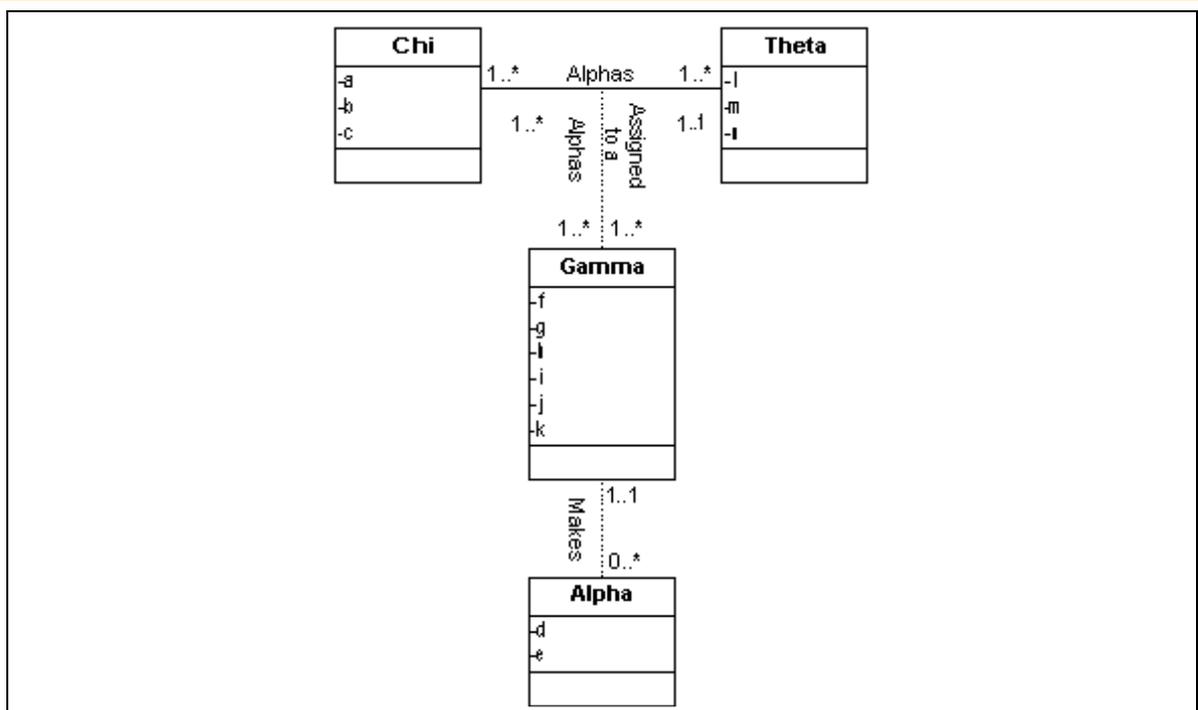


Figure 6. Semantically Void Equivalent of Figure 4

The dependent variables in this study consist of the number of correct answers to questions about the semantics conveyed in a diagram segment. Table 1 contains examples of the questions asked regarding the diagrams in Figures 3/4 and 5/6 (Appendix A contains the experimental material used in one of the treatments).

Table 1: Sample Questions

Semantic Segment (Figures 3/4)	Semantically Void Segment (Figures 5/6)
Can a Passenger be associated with a Reservation with an Aircraft that does not involve a specific Flight?	Can a Chi be associated with an Alpha with a Theta that does not involve a specific Gamma?
Does the diagram specify an upper limit on the number of Reservations associated with a specific Flight?	Does the diagram specify an upper limit on the number of Alphas associated with a specific Gamma?

The questions for both examples involve two aspects of semantics. The first question deals with whether a particular association can exist without another association. The semantics here involve the existence of precedence. In Figure 4, this is shown explicitly via the link between two association classes (participants were instructed on how to read the diagram notation); while in Figure 3, it has to be inferred through a chain of associations linking object classes, as well as the lower bounds on the cardinalities of these associations (see Appendix A for a detailed explanation). Although a correct answer can be generated from the information in either diagram, we contend that, compared to Figure 3, the depiction in Figure 4 is more explicit and does not overload the cardinality construct; therefore, we expect it to facilitate comprehension.

The second question in Table 1 deals with the semantics carried in cardinality constraints involving mutual properties. This information deals with the manifestation of preceding properties. In particular, it describes whether a particular property (e.g., Flight) can have multiple manifestations (e.g., more than one reservation for a particular flight) and whether an upper bound is specified on the number of manifestations. However, unlike questions about the existence of precedence, this information

appears the same (expressed as upper bounds on cardinality constraints) in both diagram segments. Hence, there is no reason to expect a difference in performance on questions of this type. Therefore, we separate this kind of question in our hypotheses and analysis.

For each question of each type, we coded a response 1 for a correct answer and 0 for an incorrect answer. The Restaurant segments contained one question about the existence of precedence and three questions about different aspects of cardinality related to mutual properties. The Airline segments contained two questions about the existence of precedence and one question about cardinality related to mutual properties. The second columns of both Tables 2 and 3 below describe this in more detail.

Participants were also asked to express their confidence in their response to each comprehension question on a 5-point scale ranging from “No confidence” to “Absolute confidence.”

Note that we do not consider problem solving questions as dependent variables, unlike the approach advocated by Gemino and Wand (2003) and used in previous studies (Bodart et al. 2001, Gemino & Wand 2005). As suggested by Gemino and Wand (2003), these questions go beyond the information provided in the model. Prior research has shown that problem-solving performance in this kind of task can be affected by contextual domain clues provided by the words used in a conceptual schema diagram (Parsons & Cole 2005).

The experimental materials consisted of four diagram segments (two subject domains, explicit vs. implicit representation of precedence), each with a series of questions. We also varied the use of segments that carried semantics in the words used in class and association labels (which we call Semantic segments) with the use of semantically void segments that used Greek letters as labels (which we call Void segments). We developed two versions of the material. In version A, the ordering of segments was: Void-Indirect; Void-Direct; Semantic-Direct; Semantic-Indirect. In version B, the ordering of segments was: Void-Direct; Void-Indirect; Semantic-Indirect; Semantic-Direct. In both versions, we presented the Void cases first to eliminate the possibility that participants could recognize that the semantically void segments were structurally identical and answer questions based on the wording in the Semantic segments. We also asked several demographic questions (age, gender) and questions about familiarity with several relevant information technology areas to be used as covariates.

4.3. Participants and Procedure

Participants in the study were 80 students in an undergraduate business program. All participants had completed an introductory information systems course, but had no training in UML. As this research deals with basic cognitive processes in reading and understanding diagrams, rather than management or technical issues that would require practical experience in UML modeling, students are not anticipated to differ in important aspects of cognition from the general population of non-modelers who are expected to read and verify or understand conceptual modeling scripts. Thus, students are a suitable group from which to draw participants.

Students were recruited via solicitation in an introductory marketing course. Participants were paid \$10 for taking part in the study. A short training session introduced the basics of class diagram notation, and students participated in a self-test of their understanding of the concepts covered. After the self-test, the administrator reviewed the answers to the questions with the participants. On average, participants required about 30 minutes to complete the training and experimental task.

Each participant was randomly assigned to version A or version B (as described above) of the material. Participants were told that the purpose was to test how people read information from diagrams used in information systems development. Booklets containing the experimental material were distributed. Participants spent as much time as they wished completing the exercise before returning their booklets to the administrator.

4.4. Hypotheses

To test the general Proposition 1 outlined above, we developed the following specific hypotheses. Based on the discussion of the different types of semantics that can be associated with precedence (Table 1 and the discussion that follows it), we hypothesize differences in performance only for questions involving the existence of precedence.

H1a: *Scores on comprehension questions about existence of precedence will be higher when diagrams represent precedence explicitly than when they represent precedence implicitly.*

H1b: *Scores on comprehension questions about manifestation of precedence (upper bounds on cardinalities) will not differ when diagrams represent precedence explicitly versus when they represent precedence implicitly.*

These hypotheses are motivated by the importance of precedence in Bunge's ontology. We expect that diagrams that represent the existence of this concept explicitly will communicate semantics more effectively, as measured by diagram comprehension, than those that represent it implicitly.

As the representation is explicit, we believe it is ontologically clearer and that, as a result, participants will be more confident in the correctness of their answers to questions about the existence of precedence. Thus, to test the general Proposition 2 above, we developed the following hypotheses:

H2a: *Participants' confidence in the correctness of their answers to questions about existence of precedence will be higher when diagrams represent precedence explicitly than when they represent precedence implicitly.*

H2b: *Participants' confidence in the correctness of their answers to questions about manifestation of precedence (upper bounds on cardinalities) will not differ when diagrams represent precedence explicitly versus when they represent precedence implicitly.*

5. Results

First, we examined covariates such as gender and age and found that none of those measured had any significant impact on the results. To examine H1a and H1b, we used one-sided, independent samples t-tests and a Chi-squared test to compare results for each diagram segment between the Explicit and Implicit representation of property precedence (considering Semantic and Void material separately). We used the t-test to compare scores on question types having more than one question, and the Chi-squared test to compare the percentage of correct responses for question types consisting of only a single question. For each dependent variable, we tested for homogeneity of variance between the experimental and treatment groups and report results accordingly. Table 2 shows the results of this analysis.

Table 2 provides general support for Hypothesis 1. For Question Type 1 (existence of precedence), the results support H1a with significant differences (using a Bonferroni-adjusted alpha level of $.05/4=.013$) in the hypothesized direction in three of the four cases (Restaurant: Semantic and Void; Airline: Void) and a marginally significant p-value in the fourth case (Airline: Semantic). Questions of this type asked about what amounted to the existence of precedence between two mutual properties (associations).

Table 2: Hypothesis 1a, 1b Results (Explicit vs. Implicit Representation of Precedence)

Diagram Segment	Question Type	Semantics	Mean		SD		Sig.
			Expl.	Impl.	Expl.	Impl.	
Restaurant	Existence of precedence [one question]*	Void	0.75	0.38	0.44	0.49	.001
	Cardinalities related to preceding mutual properties [three questions]		2.10	1.93	0.900	0.949	n.s.
Airline	Existence of precedence [two questions]		1.33	0.60	0.701	0.391	< .001
	Cardinalities related to preceding mutual properties [one question]*		0.83	0.58	0.385	0.501	.01
Restaurant	Existence of precedence [one question]*	Semantic	0.93	0.60	0.267	0.496	.001
	Cardinalities related to preceding mutual properties [three questions]		2.38	2.38	0.838	0.806	n.s.
Airline	Existence of precedence [two questions]		1.28	0.98	0.679	0.800	.04
	Cardinalities related to preceding mutual properties [one question]*		0.80	0.55	0.405	0.504	.01

*These rows report Chi-squared test results

For Question Type 2 (Cardinalities), as expected, no significant difference was found in two cases (Restaurant: Semantic and Void). However, significant differences (in the direction hypothesized in H1a) were found in the remaining two cases (Airline: Semantic and Void). Note that the Airline example was more complicated than the Restaurant example, so it is possible that differences with respect to cardinality understanding hold only if a certain threshold of complexity is met.

Overall, these results provide reasonable support for H1. Representing precedence directly results in better comprehension of diagram semantics dealing with the existence of precedence relations between mutual properties. Results were mixed for questions related to the manifestation of precedence (upper bounds on cardinalities), favoring explicit representation of precedence in a more complex diagram segment. In addition, there was no difference in the pattern of results between Semantic and Void segments, indicating that the Explicit representation was equally effective (relative to Implicit) in both cases.

To examine Hypothesis 2, we used one-sided, independent samples t-tests to compare the confidence scores on each question type for each diagram segment between the Explicit and Implicit diagram segments. For each dependent variable, we tested for homogeneity of variance between the experimental and treatment groups and report results accordingly. In addition, we separated scores for the Explicit versus Implicit representation of precedence cases. Table 3 shows the results.

The pattern of results in Table 3 is interesting. Considering first the Void diagram segments, in the (more complex) Airline example there were significant differences in participants' confidence in their answers to both the existence of precedence and cardinality questions. However, these differences were in the opposite direction to that hypothesized – participants were more confident their answers were correct when they viewed the Implicit representation. Yet, as Table 2 shows, they produced more correct answers to these questions when they viewed the Explicit representation. In the restaurant example, there was a significant difference in confidence for the question on the existence of precedence, with participants being more confident when answering questions based on the Explicit representation of precedence. As the airline example was more complex than the restaurant example, it is possible that complexity is a moderating variable in this situation. In general, when these results are considered in the context of prior research (Allen & March, 2006; Allen & Parsons, 2010; Burton-Jones & Meso, 2008),

there is a common theme suggesting a gap between perceived understanding (proxied by confidence) and actual understanding. We believe there is clearly a need for theorizing and further research to understand what causes this gap, and what the consequences are for the design and use of IT artifacts.

Considering the Semantic diagram segments, there was partial support for the hypothesized difference in confidence for the existence of precedence in the restaurant example, but no significant difference in confidence for the (more complex) airline example. The differences in results for confidence between the Void and Semantic groups might be in part due to the ability of subjects to base their confidence ratings on the wording in the diagrams in the latter case (in other words, by using their domain knowledge), but not in the former.

Table 3: Hypothesis 2 Results (Confidence in Explicit vs. Implicit Representations)

Diagram Segment	Question Type (Confidence)	Semantics	Mean		SD		Sig.	
			Expl.	Impl.	Expl.	Impl.		
Restaurant	Existence of precedence [one question]	Void	4.25	3.70	0.889	0.966	.005	
	Cardinalities related to preceding mutual properties [three questions]*		11.76	11.05	2.159	2.164	n.s.	
Airline	Existence of precedence [two questions]*		6.72	7.63	1.849	1.580	.011**	
	Cardinalities related to preceding mutual properties [one question]		3.82	4.15	1.023	0.893	.006**	
Restaurant	Existence of precedence [one question]		Semantic	4.63	3.73	0.705	1.012	< .001
	Cardinalities related to preceding mutual properties [three questions]*			13.35	12.64	1.777	2.133	.05
Airline	Existence of precedence [two questions]*			7.63	8.05	2.121	1.663	n.s.
	Cardinalities related to preceding mutual properties [one question]			3.98	3.95	1.165	1.037	n.s.

*Confidence scores for multiple questions summed. **Results opposite direction to Hypothesis 2.

6. Discussion

This study provides evidence that representing property precedence explicitly has a positive effect on comprehension of the domain semantics expressed in a conceptual schema. The results are clear for questions about the existence of property precedence. We observed some differences for questions about the manifestation of precedence (cardinalities of associations in precedence relations). The experimental findings show that Bunge's notion of precedence can be used to modify conceptual modeling techniques to produce meaningful improvements in comprehension of conceptual schema diagrams. In particular, common approaches to conceptual modeling overload the cardinality construct, in some cases using it to also model the semantics of the existence of preceding mutual properties.

As readers are better able to comprehend diagram semantics when precedence is represented explicitly, this can reduce the likelihood that errors in diagrams are missed during schema validation and that designers have a better understanding of domain semantics when using class diagrams for communication within the development, programming, and maintenance teams. Dobing and Parsons

(2006) found that class diagrams were among the most widely used UML diagrams for analysis and design purposes. Thus, the improved comprehension resulting from representing precedence explicitly can help reduce the likelihood that misunderstandings are not detected during systems development.

The study also provides evidence that the value of representing precedence explicitly can hold whether or not the reader of a diagram is familiar with the domain, as the comprehension results were consistent whether or not the diagram segments contained words that could prime the participants' existing domain knowledge. This provides evidence that the modeling notation can be understood correctly without relying on other cues (e.g., class and association names) that might appear in schemas.

On measures of confidence, results were mixed, and no clear pattern emerges. In some cases, participants were more confident about the correctness of their responses (and were more correct) when precedence was represented explicitly than when it was represented implicitly. However, in one case, participants were more confident of the correctness of their responses when precedence was represented implicitly, even though their comprehension of domain semantics was better when precedence was represented explicitly. In particular, this occurred in the more complex of the two examples when the segment did not contain English words. In that case, respondents provided more correct responses when precedence was represented explicitly, yet were more confident of their response if precedence was represented implicitly. This is a particularly compelling finding, as it suggests that, for complex models, implicit representation of precedence may result in poor understanding coupled with a stronger (misplaced) belief that the reader correctly understands the semantics of a schema. One possible explanation for this is that representing precedence explicitly might result in a diagram that appears more complex, and other things being equal, higher complexity may result in lower confidence. This can occur even as an explicit representation of precedence improves comprehension. Since schemas are frequently complex, this strengthens the case for representing precedence explicitly. Note, however, that this conclusion is speculative, as the experiment did not formally manipulate complexity. Further research is needed to formally examine the role of complexity on both understanding and confidence when precedence relations are represented explicitly versus implicitly.

The implications of our findings clearly have to be weighed in the context of the limitations of our experimental task and setting. Notably, the use of students and a contrived task, while useful for testing comprehension of the semantics expressed in a conceptual schema diagram, does not help us understand the impact that differences in representation capability have on the practical use of conceptual modeling techniques. Future research is needed to extend the work to more realistic settings and more complex experimental tasks, including more complex conceptual schema diagrams.

An additional issue arising from the experiment is the need to better understand both the separate effects of, and the interaction between, the formal semantics implied by the modeling constructs used in a schema and the practical semantics implied in the choice of labels attached to diagram constructs. Our study attempted to separate these issues by providing participants with two structurally identical versions – one (Semantic) containing English words and the other (Void) containing Greek letters. However, we encountered some challenges in labeling relationships or associations between classes. Typically, associations are labeled with verbs, and we were inconsistent in labeling these in materials in creating the Void segments. We used generic verbs in some cases (“assigned to” and “makes” in Figure 6). Although there is no evidence to suggest that this affected responses, the issue draws attention to the need to understand how semantics is extracted from diagrams based on both the use of modeling constructs and prior knowledge triggered by the use of familiar terms in conceptual schemas. The focus in this study was on how semantics are expressed and interpreted using modeling constructs. Nevertheless, we believe that additional insights can be gleaned from further research to examine how these two sources of semantics interact to support or contradict each other. In one preliminary study based on cardinalities, Siau, Wand, and Benbasat (1997) found that modeling experts rely on structural information expressed via modeling constructs, while novices tend to base their interpretation of semantics on inferences associated with the labeling of conceptual schemas.

7. Conclusions and Implications

This is the first study to examine empirically the impact on diagram comprehension of representing precedence explicitly in conceptual modeling. The results show that representing precedence explicitly can improve comprehension of the semantics of a problem domain expressed in a conceptual schema. Moreover, representing precedence implicitly can reduce understanding (relative to an explicit representation) but lead to greater confidence in one's (mis)understanding in more complex diagrams. A proper understanding of the semantics conveyed in a schema is an important prerequisite for detecting and correcting errors before they become embedded in other artifacts created during information systems development.

This study provides a theoretical basis for the conceptual modeling practice of representing relationships with attributes. In particular, attributes of relationships can be understood as mutual properties that are preceded by more general mutual properties. The empirical results provide evidence that this approach to modeling produces diagrams that can be better understood than diagrams in which the semantics of the existence of mutual property precedence is hidden in a chain of binary relationships. Moreover, the theoretical basis for precedence suggests how modeling techniques can be extended to show multiple levels of precedence explicitly. That is, it provides a basis for extending modeling grammars (in ways consistent with an underlying ontology) to express domain semantics more clearly.

In a broader sense, the results provide further evidence that ontological principles can be used to enhance the effectiveness of conceptual modeling techniques, and illustrate precisely how increasing ontological fidelity can improve comprehension. Thus, this research contributes to a growing and cumulative body of laboratory research indicating that ontological analysis and evaluation can be used both to predict the effectiveness of conceptual modeling techniques and to improve the way in which these techniques capture real world semantics.

We believe this kind of research improves our understanding of the value of constructs in conceptual modeling, and contributes to a stronger foundation for effective practical modeling using grammars. While the findings ultimately need to be field-tested in real development contexts, the results of laboratory studies provide guidance, both on the modification of conceptual modeling grammars and on what kinds of issues merit attention when studying the effectiveness of modeling techniques in practice.

Acknowledgements

This research was funded in part by a grant from the Social Sciences and Humanities Research Council of Canada. I would like to thank Linda Cole for assistance in the preparation of experimental materials and in data collection. I would also like to thank participants in research workshops at the University of British Columbia and the University of Queensland for valuable comments on earlier versions of this research.

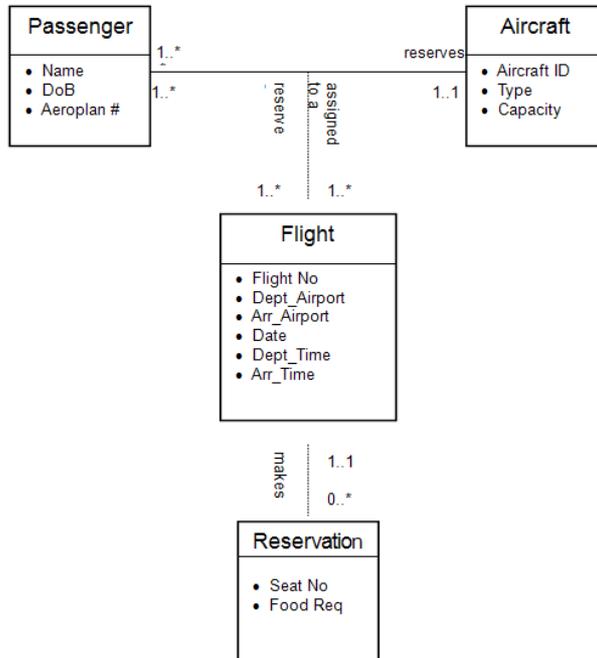
References

- Allen, G., & March, S. (2006). The effects of state-based and event-based data representations on user performance in query formulation tasks. *MIS Quarterly*, 30(2), 269–290.
- Allen, G., & Parsons, J. (2010). Is query reuse potentially harmful? Anchoring and adjustment in adapting existing database queries. *Information Systems Research*, 21(1), 56-77.
- Bodart, F., Patel, A., Sim, M., & Weber, R. (2001). Should the optional property construct be used in conceptual modeling? A theory and three empirical tests. *Information Systems Research*, 12(4), 384-405.
- Boehm, B. W. (1981). *Software engineering economics: Advances in computing science and technology*. Englewood Cliffs, NJ: Prentice-Hall.
- Bowen, P., O'Farrell, R., & Rohde, F. (2006). Analysis of competing data structures: Does ontological clarity produce better end user query performance. *Journal of the Association for Information Systems*, 7(8), 514-544.
- Bunge, M. (1977). *Treatise on basic philosophy (Volume 3), Ontology I: The furniture of the world*. Boston: Reidel.
- Bunge, M. (1979). *Treatise on basic philosophy (Volume 4), Ontology II: A world of systems*. Boston: Reidel.
- Burton-Jones, A. & Meso, P. (2008). The effects of decomposition quality and multiple forms of information on novices' understanding of a domain from a conceptual model. *Journal of the Association for Information Systems*, 9(12), 748-802.
- Burton-Jones, A., Wand, Y., & Weber, R. (2009). Guidelines for empirical evaluations of conceptual modeling grammars. *Journal of the Association for Information Systems*, 10(6), 495-532.
- Burton-Jones, A. & Weber, R. (1999). Understanding relationships with attributes in entity relationship diagrams. In P. De & J. DeGross (Eds.), *Proceedings of the Twentieth International Conference on Information Systems*, Charlotte, NC, 214-228.
- Burton-Jones, A. & Weber, R. (2003). Properties do not have properties: Investigating a questionable conceptual modeling practice. In D. Batra, J. Parsons & V. Ramesh (Eds.), *Proceedings of the Second Annual Symposium on Research in Systems Analysis and Design*, Miami, 14 pp.
- Byrd, T., Cossick, K., & Zmud, R. (1992). A synthesis of research on requirements analysis and knowledge acquisition techniques. *MIS Quarterly*, 16(1), 117-138.
- Chen, P. (1976). The entity-relationship model: Toward a unified view of data. *ACM Transactions on Database Systems*, 1(1), 9-36.
- Davies, I., Green, P., Rosemann, M., Indulska, M., & Gallo, S. (2006). How do practitioners use conceptual modeling in practice? *Data and Knowledge Engineering*, 58(3), 358-380.
- Dobing, B., & Parsons, J. (2000). Understanding the role of use cases in UML: A review and research agenda. *Journal of Database Management*, 11(4), 28-36.
- Dobing, B., & Parsons, J. (2006). How UML is used. *Communications of the ACM*, 49(5), 109-113.
- Dobing, B., & Parsons, J. (2008). Dimensions of UML diagram use: A survey of practitioners. *Journal of Database Management*, 19(1), 1-18.
- Evermann, J. (2005). The association construct in conceptual modelling – An analysis using the Bunge ontological model. In O. Pastor & J. Falcao e Cunha (Eds.), *Proceedings of the 2005 Conference on Advanced Information Systems Engineering (CAiSE'05), LNCS 3520*, 33-47.
- Evermann, J., & Wand, Y. (2005). Ontology based object-oriented domain modelling: Fundamental concepts. *Requirements Engineering Journal*, 10(2), 146-160.
- Evermann, J., & Wand, Y. (2006). Ontological modeling rules for UML: An empirical assessment. *Journal of Computer Information Systems*, 46(5), 14-29.
- Fettke, P. (2009). How conceptual modeling is used. *Communications of the Association for Information Systems*, 25(1), 571-592.
- Gemino, A., & Wand, Y. (2003). Evaluating modeling techniques based on models of learning. *Communications of the ACM*, 46(10), 79-84.
- Gemino, A., & Wand, Y. (2004). A framework for empirical evaluation of conceptual modeling techniques. *Requirements Engineering Journal*, 9(4), 248-260.
- Gemino, A., & Wand, Y. (2005). Complexity and clarity in conceptual modeling: Comparison of mandatory and optional properties. *Data and Knowledge Engineering*, 55(3), 301-326.

- Gemino, A., & Wand, Y. (2003). Foundations for empirical comparisons of conceptual modeling techniques. In D. Batra, J. Parsons & V. Ramesh (Eds.), *Proceedings of the Second Annual Symposium on Research in Systems Analysis and Design*, Miami, FL, 29 pp.
- Hoffer, J., Prescott, M., & McFadden, F. (2008). *Modern database management (Eighth Edition)*, Upper Saddle River, NJ: Prentice Hall.
- Khatri, V., Vessey, I., Ramesh, V., Clay, P., & Park, S. (2006). Understanding conceptual schemas: Exploring the role of application and IS domain knowledge. *Information Systems Research*, 17(1), 81–99.
- Mayer, R. (2001). *Multimedia learning*. Cambridge, MA: Cambridge University Press.
- Parsons, J., & Cole, L. (2005). What do the pictures mean? Guidelines for experimental evaluation of representation fidelity in diagrammatical conceptual modeling techniques. *Data and Knowledge Engineering*, 55(3), 327-342.
- Parsons, J., & Wand, Y. (2003). Attribute-based semantic reconciliation of multiple data sources. *Journal on Data Semantics 1*, LNCS 2800, 21-47.
- Rob, P., & Coronel, C. (2009). *Database systems: Design, implementation, and management*. Boston, MA: Course Technology.
- Siau, K., Wand, Y., & Benbasat, I. (1997). The relative importance of structural constraints and surface semantics in information modeling. *Information Systems*, 22(2-3), 155-170.
- ter Hofstede, A. & van der Weide, Th. P. (1993). Expressiveness in conceptual data modelling. *Data and Knowledge Engineering*, 10(1), 65-100.
- Wand, Y., Storey, V., & Weber, R. (1999). An ontological analysis of the relationship construct in conceptual modeling. *ACM Transactions on Database Systems*, 24(4), 494-528.
- Wand, Y. & Weber, R. (1990). An ontological model of an information system. *IEEE Transactions on Software Engineering*, 16(11), 1282-1292.
- Wand, Y., & Weber, R. (1993). On the ontological expressiveness of information systems analysis and design grammars. *Journal of Information Systems*, 3(4), 217-237.
- Wand, Y. & Weber, R. (2002). Information systems and conceptual modeling: A research agenda. *Information Systems Research*, 13(4), 363-376.
- Weber, R. (1997). *Ontological foundations of information systems*. Melbourne: Coopers & Lybrand.

Appendix. Material Used for Airline Domain (Semantic version)

Case 1. Direct Representation of Precedence



1. Can a Passenger be associated with a Reservation that does not involve a specific Flight?

NO [Reservation is preceded by Flight]

2. Does the diagram specify an upper limit on the number of Reservations associated with a specific Flight?

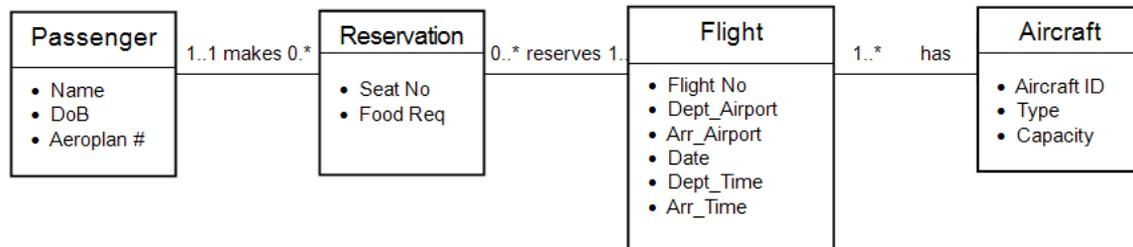
NO [The maximum cardinality associating a flight with a reservation is '*' – not specified]

If so, what is the limit? _____

3. If a passenger takes a flight on an aircraft, must he/she have a reservation?

NO [A flight need not precede a reservation. The lower bound on the cardinality connecting flight to reservation is '0,' while the lower bound on the cardinality connecting reservation to flight is '1.' This structure thereby preserves the semantics of precedence.]

Case 2. Indirect Representation of Precedence



1. Can a Passenger be associated with a Reservation that does not involve a specific Flight?

NO [The minimum cardinality associating a reservation with a flight is '1']

2. Does the diagram specify an upper limit on the number of Reservations associated with a specific Flight?

NO [The maximum cardinality associating a flight with a reservation is "*" – not specified]

If so, what is the limit? _____

3. If a passenger takes a flight on an aircraft, must he/she have a reservation?

NO [The lower bound on the cardinality connecting Passenger and Reservation is '0']

About the Author

Jeffrey PARSONS is Professor of Information Systems and Associate Dean (Research) in the Faculty of Business Administration at Memorial University of Newfoundland. His research interests include conceptual modeling, data management, business intelligence, and recommender systems; he is especially interested in classification issues in these and other domains. His research has been published in journals such as *Nature*, *Management Science*, *MIS Quarterly*, *Information Systems Research*, *Journal of Management Information Systems*, *Communications of the ACM*, *ACM Transactions on Database Systems*, and *IEEE Transactions on Software Engineering*. Dr. Parsons is a senior editor for the *Journal of the Association for Information Systems* and recently completed a term as an associate editor for *Information Systems Research*. He has also served as Program Co-chair for a number of conferences, including AMCIS, WITS, and the ER conference. He holds a PhD in Information Systems from the University of British Columbia.