

“We Need to Discuss the *Relationship*”: Revisiting Relationships as Modeling Constructs

Nicola Guarino¹ and Giancarlo Guizzardi^{1,2(✉)}

¹ ISTC-CNR Laboratory for Applied Ontology (LOA), Trento, Italy
guarino@loa.istc.cnr.it, gguizzardi@inf.ufes.br

² Computer Science Department, Federal University of Espírito Santo (UFES), Vitória, Brazil

Abstract. In this paper we propose a novel ontological analysis of relations and relationships based on a re-visitation of a classic problem in the practice of conceptual modeling, namely *relationship reification*. Despite the two terms ‘relation’ and ‘relationship’ are often used interchangeably, we shall assume a radical difference between the two: a relation *holds*, while a relationship *exists*. Indeed, the relation holds *because* the relationship exists. We investigate the ontological nature of relationships as *truthmakers*, proposing a view according to which they are *endurants*. Under this view, not only a relationship is responsible (with its existence) of the fact that a relation holds, but it also accounts (with its properties) of *the way a relation holds and develops in time*.

1 Introduction

“We need to discuss our relationship” is a statement that often occurs in our social every-day life. In such a context, a relationship is something that has its own life. Yet, in conceptual models relationships occur everywhere, but they tend to be considered as frozen entities. In this paper we argue for the need to (re)discuss the notion of relationship used in conceptual modeling, in order to better understand its ontological nature and to better capture the way it is used in describing application domains.

In his seminal paper [1], Chen describes relationships in E-R diagrams in the following manner: “*a relationship is an association among entities. [However,] it is possible that some people may view something (e.g., marriage) as an entity while other people may view it as a relationship. We think that this is a decision which should be made by the Enterprise Administrator*”. In the same passage, Chen defines an Entity as a “thing” and claims that “*a specific person, a company, or an event is an example of an entity*”. Further in the paper, he defines a *relationship type* as a “*mathematical relation*”, that is, a set of tuples, and a relationship as one of such tuples. Chen also admits that relationships can bear properties. In one of his original examples, he illustrates how a *project-worker* relationship can have the attribute *percentage-of-time* representing an intrinsic property of the relationship itself.

We here subscribe to many of Chen’s original intuitions. For example, as we argue in this paper, we agree with him on admitting that relationships can bear properties. However, in that paper, Chen does not take an explicit ontological stance on the very nature of relationships and, as a consequence, the real-world semantics of this

construct (a fundamental one in his notation and in conceptual modeling, in general) is left underspecified. In a subsequent paper [2], he elaborates on some of these ontological aspects (albeit in an indirect form) via an analysis of the English sentence structure of textual descriptions. He states that entity types are the E-R counterpart of common nouns and that relationships are normally expressed by transitive verb phrases (e.g., *owns*, *belongs to*, *loves*). In addition, he recognizes that “nouns converted from a verb” (i.e., *verb nominalizations*) correspond to relationships (e.g., the *shipping* of the product to a customer or the *assigning* of an employee to a machine). Now, since verbs are typically the language proxy for events (including actions), it seems at first that Chen is committed to the view that relationships are events. Yet, in the very same paper he brings examples of relationship names such as *location of* or *part of*, which may not obviously correspond to events.

In any case, events are ontologically very different from tuples, so a question-begging issue that seems to be present since these early seminal papers on Conceptual Modeling is: what is a *relationship* after all? Is it just a tuple? Is it an event? Or is it something else? Moreover, is the difference between Entities and Relationships only a matter of a pragmatic modeling, as Chen seems to suggest, or there are aspects of the intrinsic nature of “real-world” entities that would justify such distinction?

In this paper we shall address the issues above by proposing a novel ontological analysis of relations and relationships based on a re-visitation of a classic problem in the practice of conceptual modeling, namely *relationship reification*. Despite the two terms ‘relation’ and ‘relationship’ are often used interchangeably, we shall assume a radical difference between the two: a relation *holds*, while a relationship *exists*. Indeed, our idea is that the relation holds *because* the relationship exists. A relationship is therefore a *truthmaker* [3] for the corresponding relation. However, as we shall see, such a truthmaker, in our view, not only is responsible (with its existence) of the fact that the relation holds, but it also accounts (with its properties) for *the way a relation holds and develops in time*. This means that a relationship includes aspects whose existence *entails* that the relation holds, as well as aspects whose existence is a *consequence* of the fact that the relation holds. So, when a worker works in a project as a result of job assignment, a new object emerges, in addition to the two *relata* (the worker and the project), whose properties and internal structure reflect, at each time, not only the *obligation* the worker has (which entails that the relation holds), but also the actual effort invested, the degree of satisfaction or stress, and so on. The latter aspects are indeed a *consequence* of the obligation, and together contribute to describe *the way* the relation develops.

Our main claim is that relationships (at least those mostly interesting for us) are objects (*endurants*). The basic intuitions behind this claim were already present in Guizzardi’s early works [4,5], later advanced in [6] and with an alternative version in [7]. In such works, reified relationships are conceptualized as *relators*, which are endurants of a special kind, with the power of connecting (*mediating*) other endurants. In this paper we shall revisit the foundations of this early work, investigating and clarifying the very nature of relators in the light of the notion of truthmaking, and ultimately identifying them with relationships. We shall also discuss the subtle connections between relationships, relations and events, and their implications on the practice of conceptual modeling.

The paper is structured as follows. In section 2, we revisit the notion of truthmaking and use it to briefly outline a typology of relations. We isolate a class of relations particularly relevant for conceptual modelling –*extrinsic relations*– and motivate the practical need to reify the corresponding relationships as truthmakers. In section 3, we explore more in detail the complex nature of such truthmakers. First, on ontological and linguistic grounds, we reject *facts* and tuples as alternatives. Then we also reject perdurants (event-like entities) as an alternative. In section 4 we investigate further the nature of relationships. We show that these are indeed *thing-like entities* and that they behave very much like all other thing-like entities. In particular, they can maintain their identity in time while possibly changing in a qualitative manner. In other words, we claim that relationships are *full-fledged endurants*, and as such they are the natural bearers of modal properties. In section 5, we explore some consequences of this view on the practice of conceptual modeling. We show that, by taking relationships as endurants, we have that their types can be organized in taxonomic structures exactly like object types, according to well-known and tested ontology-driven conceptual modeling design patterns. We also revisit a well-known approach towards relationship reification in conceptual modeling and show how our framework supports an ontological analysis and conceptual clarification of that approach. Finally, section 6 presents some final considerations for the paper.

2 Relationships as Truthmakers (of a Certain Kind)

2.1 Truthmakers and Kinds of Truthmaking

To start our analysis on the ontology of relations and relationships, we introduce the notion of *truthmaker*. We say that propositions, such as ($p1$) “*a is an apple*” or ($p2$) “*a is red*”, are *truthbearers*, in the sense that they can be either true or false, depending on what happens in the world. A *truthmaker* for one of such propositions is an entity in virtue of whose existence the proposition is true. There are several attempts to capture formally the notion of truthmaking [3], but for our purposes we shall take it as primitive, as a fundamental relation linking what is true to what exists. More in general, a truthmaker for a property or a relation is an entity in virtue of whose existence that property or relation holds. Suppose that in ($p1$) *a* denotes a particular apple. The very existence of *a* is enough for making the proposition true, so *a* is a truthmaker of $p1$. For $p2$, in contrast, the mere existence of *a* is only indirectly responsible of $p2$ ’s truth, since it is in virtue of the specific *way* the apple is (and not just in virtue of the apple’s existence) that $p2$ is true. In particular, there is something *inhering* to *a*, which we may call *a’s redness*, which makes *a* red. Some philosophers term this redness an individualized property, or a *trope* (see also [8]). According to them, it is in virtue of such trope that the *red* property holds, so the trope (and not the apple) is the truthmaker. So, *being an apple* and *being red* are properties whose truthmakers are of a very different nature. We shall say that the former is a *formal property*, while the latter is a *qualitative property*. Formal properties account for *what* something is; qualitative properties account for *how* something is.

Let us shift now our focus from unary properties to (binary) relations. As we have anticipated, our first answer to the question discussed above, “*What are relationships, after all?*” is that *relationships are truthmakers of relations*. The nature and structure

of such truthmakers determine different kinds of relations. A first kind is that of *formal relations*, whose arguments jointly constitute a truthmaker. For instance, the *successor-of* relation between integers is a formal relation, since it holds because of very nature of its arguments. In contrast, *qualitative relations* are those whose truthmaking depends on the existence of something in addition to the relata. Among qualitative relations, we include relations such as *being-taller-than*, whose truthmaking depend on particular intrinsic properties of the relata (i.e., their individual heights). In other words, the truthmaking of these relations depends solely on their intrinsic height tropes, i.e., no *relational trope* is involved. We term these *intrinsic relations*. In this paper, we are interested in a class of qualitative relations that are of far greater relevance to conceptual modeling, i.e., the so-called *extrinsic (or material) relations* [5].

Differently from intrinsic relations, the truthmakers of *extrinsic relations* are entities that are truly relational in nature, i.e., that cannot be reduced to intrinsic properties of the relata. These entities can only exist while connecting the relata, i.e., they exhibit mutual dependency patterns involving all the arguments, as well as dependencies on external entities besides the arguments. *Married-with* is a prototypical example of an extrinsic relation. For this relation to hold between John and Mary, we need more than the existence of John, Mary and their intrinsic properties. We need the occurrence of a wedding event (or the signing of a social contract), which, in turn, bestows John and Mary with individual qualities (e.g., commitments, claims, rights, obligations) that are truly relational in nature. For instance, a particular commitment of John towards Mary is a quality of John (inhering in John) but which is also existentially dependent on Mary.

Recognizing extrinsic relations is of fundamental importance for conceptual modeling. Firstly, because most of the relations represented in conceptual models are of this sort. For instance, *enrollments*, *employments*, *purchases*, *employee allocation to projects* and *presidential mandates* are of the very same ontological nature as the *marriage* in the previous paragraph, in strong contrast to relations such as *being-older-than* or *successor-of*. Secondly, as discussed in depth in [5,6], extrinsic relations must be represented in conceptual models differently from intrinsic relations. In particular, extrinsic relations cannot be suitably represented just as “mathematical relations”, i.e., as sets of tuples. In the sequel, we elaborate on this issue.

2.2 The Need for Explicitly Representing Truthmakers of Extrinsic Relations

Take for instance the example discussed in [5] of a *treated-in* relation between *Patient* and *Medical Unit*. This is clearly neither a formal nor an intrinsic relation. For this relation to hold, it requires the existence of a third entity, such as an individual Treatment involving somehow a Patient and a Medical unit. So, independently of the nature of such third entity (which will be discussed in the rest of the paper), this is enough to classify this relation as an extrinsic relation. The presence of such third entity is at the origin of a specific practical problem, mentioned in [9] that affects the representation of extrinsic relations just as standard associations. In this particular example, let us assume we represent our relation as an association such that each patient can be treated in one-to-many medical units and that each unit can treat one-to-many patients. The problem is that these constraints are ambiguous, since many different interpretations can be given to it, including the following: (i) a patient is

related to only one treatment in which possibly several medical units participate; (ii) a patient can be related to several treatments to which only one single medical unit participates; (iii) a patient can be related to several treatments to which possibly several medical units participate; (iv) several patients can be related to a treatment to which several medical units participate, and a single patient can be related to several treatments. To disambiguate among these different interpretations, we need a way to express constraints concerning how many patients and medical units can interact at each instance of the relation. In the literature, such constraints have been called *single tuple constraints* [9]. However, these constraints do not concern the single tuple, but the material conditions that justify the presence of that tuple in the relation, i.e., its truthmaker. So, to solve the problem, we must explicitly represent this truthmaker, i.e., the *treatment* itself.

This was indeed the solution adopted in [5], shown in Fig. 1, where the treatment is reified as a *relator*. As discussed in [4,5], the relation stereotyped as *«mediation»* in fig.1 is a special type of existential dependence relation (a treatment only exists if both a patient and a medical unit exists). The figure also reports a dashed arc connecting the relator to the association between Patient and Medical Unit. This represents the *derivation* relation between a truthmaker and the tuples derived from it. In a nutshell, when a relation *R* is *derived* from a relator type *T*, then, for every *x*, *y*, *R(x,y)* holds if there is an instance *t* of *T* such that *mediates(t,x)* and *mediates (t,y)* hold. Both dependence and derivation are examples of formal relations.

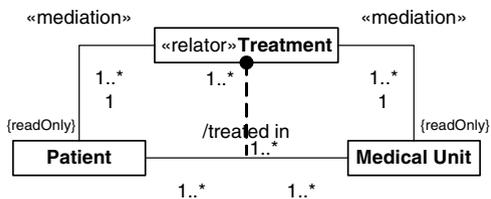


Fig. 1. Reifying Treatment as a relator solves the cardinality ambiguity problem (from [5])

In conclusion, we can see that by adopting the solution exemplified by Fig. 1 the cardinality ambiguity problem dissolves, since we can easily express the constraints on the number of patients and units involved in a single treatment. As demonstrated in [5], this problem manifests itself *for all extrinsic relations, and only for them*.

In summary, by recognizing extrinsic relations and reifying them as 'relators' we are able to address a number of classical problems related to the modeling of relations. These range from eliminating the ambiguity between *association specialization, subsetting and redefinition* [6] to solving the problem of transitivity of part-whole relations [4]. Independently of these modeling benefits, however, some fundamental questions still remain: what is the ontological nature of such relators? If they are truthmakers, what kind of truthmakers? How do they account for our idea of relationships as *ways* relations develop in time? This is the topic of next section.

3 What Kind of Truthmakers?

3.1 Are Relationships Facts?

One of the most common assumptions in conceptual modeling is that relationships are actually *facts*. Indeed, modeling approaches such as ORM [10] are normally called *fact-based approaches*. As we shall see, there are different ways of understanding relational facts, but none of them turns out to adequately account for the notion of relationship as truthmaker.

Suppose first that each instance of a relation (a tuple) directly represents a fact. So, for the relation *is-employed-by*, the tuple $\langle \text{John}, \text{CNR} \rangle$ represents the fact that John is employed by CNR. Take now the inverse relation *employs*. According to our assumption, the tuple $\langle \text{CNR}, \text{John} \rangle$ represents a different fact, since it is a different tuple. Indeed, the fact that John is employed at CNR is usually considered as different from the fact that CNR employs John, although, as argued in [11], in the world we have only one truthmaker, namely a state of affairs being described in two different ways. So, if we take that a relationship is just a syntactic instance of a relation, i.e., a tuple, then it cannot be a truthmaker.

A possibility to overcome the problem above is to associate specific roles to the arguments of the relation, so that the arguments' order is not crucial any more. This seems to be the position adopted by ORM and UML (with rolenames tied to association ends) and also entertained by Chen in his original article [1]. As discussed in depth in [11], this move solves the basic problem above, with two important caveats: (i) it makes *rolenames* as part of our basic ontological categories, instead of just a mere description tool; (ii) the solution does not really work for the case of symmetric relations, since in this case we would still have multiple possible entities standing for the same basic situation in reality. In other words, for asymmetric relations such as *is-employed-by* and *employs*, we avoid having as candidate truthmakers the distinct entities $\langle \text{John}, \text{CNR} \rangle$ and $\langle \text{CNR}, \text{John} \rangle$ by introducing a rolename for each argument position, say 'Employee' and 'Employer', and a unique entity that assigns the two relations to them $\{\text{John} \rightarrow \text{Employee}, \text{CNR} \rightarrow \text{Employer}\}$. However, for the case of symmetric relations such as *brother-of*, we would still have two distinct entities: $\{\text{John} \rightarrow \text{Brother1}, \text{Paul} \rightarrow \text{Brother2}\}$ and $\{\text{John} \rightarrow \text{Brother2}, \text{Paul} \rightarrow \text{Brother1}\}$ [11].

However, even ignoring (i) and (ii) above, there are more cogent reasons for ruling out the assumption that relationships are facts. In [3], Moltmann discusses a distinction between different types of *nominalizations*. Nominalization is the linguistic process of turning expressions of various categories into nouns. Examples include *John's wisdom*, *Mary's beauty*, *John and Mary's marriage*. In the case of relational expressions (such as the latter), these nominalizations are equivalent to relation reification as treated in the conceptual modeling literature. One of the issues addressed by Moltmann concerns the semantics of such nominalizations: they refer to particulars, but what kind of particulars? She contrasts three options for these referents: (1) *facts*, (2) *events*, and (3) *qualities*. On the basis of linguistic evidence, she then rejects option 1.

According to option 1, *Clara's enrollment at UNITN* would mean *the fact that Clara is enrolled at UNITN*. The problem, however, is that facts, by their very nature, are completely determined entities. In order to illustrate this point, let us consider the model of Fig. 2 below, which depicts an enrollment relationship with two optional

attributes. If relationships are just facts, a possible instance of this model would be f_1 : *the fact that Clara studies at UNITN*. However, suppose that at a later point we learn that Clara has a 9.5 GPA. Now, f_2 : *the fact that Clara studies at UNITN with a 9.5 GPA* is clearly a different fact. Furthermore, if we learn later that Clara has accumulated 250 credits, we have f_3 : *the fact that Clara studies at UNITN with 250 credits*, as well as f_4 : *the fact that Clara studies at UNITN with a 9.5 GPA and 250 credits*. In other words, the creation of these multiple facts obscures the important issue that there is one single entity in reality, namely, Clara’s enrollment at UNITN, underlining these facts. In contrast, when the referent of a nominalization is an event (option 1) or a quality (option 2), we have a concrete object of reference whose properties can be progressively uncovered and described. For instance, we linguistically accept *Peter described Clara’s talk* and *Peter admired Clara’s beauty*, but not *Peter described the fact that Clara gave a talk*, nor *Peter admired the fact that Clara was beautiful*.

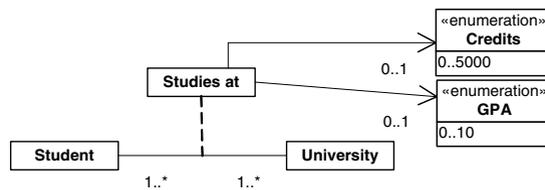


Fig. 2. Interpreting relationships as Facts

In conclusion, we think Moltmann brings enough evidence to discard facts as truthmakers of relations. The other two alternatives, events and qualities, deserve a bit of discussion. First of all, the use of the term ‘quality’ while referring to Moltmann’s work is our choice, since she uses the term ‘trope’, admitting however that her tropes “are not tropes as most commonly understood”. We believe that *individual qualities* in the sense of DOLCE [12] (or the so-called *moment persistents* in UFO [4]) are what she has in mind. Indeed, in her view qualities are understood as abstractions from tropes, which (differently from standard tropes) can change while maintaining their identity. For instance, in a sentence like *the color of the apple is changing*, we seem to make cognitive sense of an aspect of the apple that maintains its identity while changing in a qualitative way. After all, it is not ‘red’ that is changing; there is something there that changes while keeping existing as the referent for *the color of the apple*.

Now, coming to the difference between events and qualities from the point of view of their aptitude to act as truthmakers for relations, Moltmann highlights that qualities are capable of genuine changes while maintaining their numerical identity, while events are not. This latter point is further explored in next section.

3.2 Are Relationships Events?

At least since Chen’s paper, the idea that relationships are events has been often implicitly present in the conceptual modeling literature. After all, entities like marriages, contracts, enrollments, from a linguistic point of view are verbal nominalizations, whose most obvious semantics, as discussed above, would be that of events (note that we use ‘event’ here in its most general sense, i.e., as a synonym of ‘perdurant’).

Before investigating the viability of this idea, let us first discuss a particular type of event that can be a candidate for our purposes, namely, what we call an *episode*. Suppose that *John runs* between 8:00 and 8:45 tonight. Any temporal subpart of this event (up to certain granularity) is an instance of a running event. Having such multitude of entities can cause many problems for conceptual modeling, in particular in the unambiguous specification of cardinality constraints, in a way similar to the problems of modeling quantities and collectives discussed in depth in [4]. For this reason, we introduce the notion of *episode* as an event that is maximal given a certain unity criterion, i.e., given an episode e instantiating a property P , there is no event e' of which e is part such that e' also instantiates P . So, in our example, if John runs in a certain time interval there is at least one *running* episode that overlaps with that interval.

Consider now *relational* episodes, i.e., episodes involving multiple disjoint participants being in a certain relation. Clearly the existence of an *episode of being in relation* R entails that R holds for the time the episode lasts, so episodes (differently from *facts*, as discussed above) can genuinely be taken as truthmakers of extrinsic relations. For instance, a *marriage* episode can be taken as a truthmaker of the *married-to* relation holding between John and Mary. Notice however that, if we take marriage in the institutional sense (excluding *de facto* marriages), then *married-to* is an historical relation, which, by definition, holds at a certain time as a consequence of a specific episode (say, a wedding) occurred in the past. So, it is the wedding episode, and not the marriage episode, that is the ultimate (minimal) truthmaker of the *married-to* relation. Which of the two truthmakers shall we pick up as a candidate for the *married-to* relationship? To avoid the impasse we need to introduce a very plausible principle: a relationship is a truthmaker of a relational proposition *that lasts as long as the proposition holds*. In other words, a relationship is a *synchronous* truthmaker of a relation.

The principle above reconciles our general assumption that a relationship is a truthmaker of a relation with the hypothesis that a relationship is an episode. The question remains, however, whether this latter hypothesis suits our intuitions on relationships. Consider again the *marriage* episode. In our everyday talk (as the title of the present paper suggests), we say that a marriage relationship changes while it develops in time: it becomes more litigious or more passionate, new obligations arise because of children, and so on. Now, what is the subject of such changes? Clearly it cannot be the marriage episode, which, according to classical theories of events, is a 'frozen' extensional entity defined by the sum of its parts [13,14]. Indeed, in the traditional literature, a key difference between endurants and events is that the former can genuinely *change* in time while maintaining their identity [12], while the latter can just *vary* in time by exhibiting different properties for their temporal parts. So, if we model the marriage as an event (an episode), if it is peaceful at t_1 and litigious at t_2 it has two different temporal parts that bear otherwise incompatible properties. There is nothing that is entirely present throughout the duration of the marriage, and this is especially problematic while the episode is ongoing, since we have no way to say that the two parts at t_1 and at t_2 belong to *the same* marriage. A further problem concerns modal properties, which we often ascribe to relationships: could a marriage have been different from what it is, allowing us to consider counterfactual situations (e.g., *would their marriage have been different had they moved to Australia*)? Again, according to classical theories of events, we are forced to answer *no* to these questions. An event could not have been different from what it is.

In conclusion, conceiving relationships as episodes captures only partially our intuition that a relationship accounts for the way a relation holds and develops in time, because of the difficulties of properly modeling change. Yet, the explicit introduction of relational episodes in conceptual modeling can be extremely effective as a way to reify relationships in all those cases where we are only interested in keeping track of such episodes, without being interested in modeling change phenomena within them. For instance, the explicit introduction of a commitment episode for modeling services turned out to be very useful [15], although to describe complex service dynamics we adopted a different approach [16], based the notion of relator as described below.

4 Relationships as Full-Fledged Endurants

In the following, we shall elaborate on a different notion of relationships that conceives them as truthmakers of a special kind, namely endurants consisting of bundles of individual qualities, revising and extending Guizzardi's previous work.

Going back to one of Chen's original examples, consider the following proposition: a worker w *works-in* project p , as a result of an assignment event a . What are the truthmakers of such proposition? As we have seen, there is a minimal (asynchronous) truthmaker, the event a , and a synchronous truthmaker, a working episode e . Both are perdurants. There is however another synchronous truthmaker, which is an endurant: the *obligation* (o) the worker w has to work in p . This can be seen as a *relational quality* inhering in w . Its very existence at a certain time entails that the relation *works-in* holds at that time. Notice however that the permanence of the obligation to work in a certain project does not necessarily mean that the worker actually works there, so let's assume that 'works in' actually means 'is supposed to work in'. Indeed, we can consider another relational quality inhering in w , namely w 's actual amount of *labour* currently (l) spent to fulfill o in the context of p . Clearly, l is dependent on the existence of o , besides being existentially dependent on p . Another relational quality depending on o may be w 's degree of satisfaction while fulfilling o in the context of p . Other relational qualities existentially depending on o may emerge later, such as the collaboration attitude towards a new co-worker. Moreover, we have also qualities on the side of p depending on o , such as the amount of work assigned to w .

We see that, collectively, all these qualities (which are endurants, changing in time) describe in a very fine-grained way how the *works-in* relation develops in time, which is exactly our view of a relationship. Within these qualities, we can isolate a *nucleus*, in virtue of whose existence the relation holds, and a *shell* that is existentially dependent on the nucleus. We define a relationship as *an endurant that, at each time the relation holds, is constituted by the mereological sum of all these qualities* (the nucleus plus the shell). Thanks to this construction, we can see a relationship as a truthmaker of a certain kind, which accounts of the way a relation holds. In short, *a relationship is the particular way a relation holds* for a particular set of relata.

The summary of this discussion is illustrated in Fig. 3. Fig. 3a illustrates the inherence relation (i) between a worker, John, and those qualities of him that contribute to constitute the relationship with project $P\#1$. We distinguish a nucleus of qualities that are directly existentially dependent (ed) on $P\#1$, and a shell of qualities which, depending on the nucleus, are indirectly existentially dependent on $P\#1$. Fig. 3b shows how the relationship, being a truthmaker of the *works-in* relation, mediates (m) between John and $P\#1$.

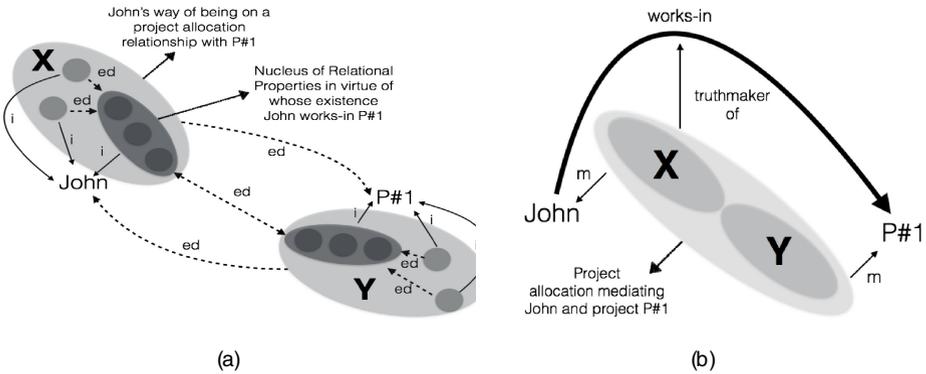


Fig. 3. (a) Relators (relationships) and their constituent relational qualities; (b) relators (relationships) as truthmakers of relations by mediating their relata

Let us now discuss the innovation with respect to the former work on relators. First, in the most recent formalization [7], relators were defined as sums of tropes. This played an important role in supporting reasoning in languages such as OWL. However, as we have seen (Section 3.1), classically speaking, tropes cannot change. In the present account (which makes justice of the original intuition [4]), since relationships are mereological sums of qualities and not tropes, they can change in all aspects in which their constituent qualities can change. For instance, an enrollment can change its accumulated credits or its GPA, and a medical treatment can change its cost. Note that, being an endurant, each relationship is an instance of a *kind*, which accounts for its identity conditions through time, determining the admissible extent of such changes. These kinds can be specialized and organized in various ways around distinctions such as those proposed by OntoClean and OntoUML (see discussion in section 5).

A further limitation of the original theory of relators was that it did not explicitly allow having relators that lose and gain qualities. Here we have defined a relationship as *constituted*, at each time it exists, by a mereological sum of qualities. Hence a relationship is a *variable embodiment* [17], which may have different constituent qualities at different times. So not only such qualities can have different values (qualia) at different times, but also that some of these qualities (those belonging to the shell) can cease to exist or come into being during the life of the relationship, exactly like the non-essential parts of a mechanical assembly. For instance, a service contract can contingently be insured, acquiring a ‘level of insurance’ quality.

Another innovation with respect to the previous approach concerns the role of *founding events* in determining the nature of relators. In Guizzardi's previous work, the focus was on relationships that hold as a consequence of an event (like the marriage as a consequence of a wedding), and the joint dependence on a single founding event was the unity criterion used to hold all the relator's qualities (or tropes) together. In our present view, since the identity of the nuclear qualities is given by the relator's kind and the non-nuclear qualities are just those dependent on the nucleus, the existence of a founding event is not required anymore, so the idea of relationships as endurents also applies to instantaneous relationships, such a ball hitting the floor.

Concerning the generality of this view, however, we have to observe that relators exist only for *genuine* relationships (as opposite to so-called *Cambridge relationships* [18]), which somehow impact on the properties of at least one of the relata. This means that the historical relationship existing between John and, say, a street he inadvertently crossed several years ago still has an (uninteresting) episode as a truthmaker, but none of John's qualities (and none of the street's qualities) keep memory of this, so to speak, so there is no relator. Indeed, the absence of a relator would mark the fact that the relationship is an uninteresting one.

On a final note, we should briefly comment on the connection between relationships and events (episodes). As we take (in general) events as manifestations of qualities, relational episodes are manifestations of relationships. The way the relation's arguments *participate* to the episode depends on the qualities constituting the relationship. For example, the *marriage episode* between John and Mary is the manifestation of the qualities constituting the marriage relationship [14]. Using a different terminology, we say that when an event is a manifestation of certain qualities, such qualities are the *focus* of that event. So *a relationship is the focus of a relational episode*.

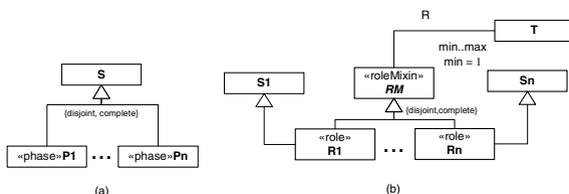


Fig. 4. Examples of OntoUML patterns for modeling taxonomic structures involving object types: (a) the phase pattern; (c) the role mixin pattern [20]

5 Practical Implications to Conceptual Modeling

5.1 Ontological Distinctions Applied to Relationship Types

Modal meta-properties play a fundamental role in conceptual modeling and ontology engineering. Indeed, in methodologies such as OntoUML [4] and OntoClean [19], there are a number of (meta)categories of types which are differentiated in terms of their modal meta-properties. For instance, while a *kind* such as *person* applies to its instances necessarily, a *phase* such as *teenager* or a *role* such as *student* apply to its instances only contingently. In other words, an instance of a person cannot cease to instantiate that type without ceasing to exist. In contrast, instances of a *phase* (teenager, living person) or *role* (student, employee) can move in and out of the extension of these classes without any impact of their identity. The ontology of types that underlies these approaches (containing these distinctions among many others) has proven to be one of the most productive conceptual tools in ontology engineering [4, 19]. However, this ontology of types has been mainly used so far to structure *object* types (i.e., types of independent entities), and not *moment* types (i.e., types of dependent entities like qualities or relators). The choice of considering relationships as endurants, admitting genuine changes and modal properties, allows us to adopt for them the same general patterns that have proved very useful for modeling objects in these approaches.

discusses the same example we have used above, namely the *works-in* relation. He assumes that, for each day a person works in a project, the number of hours worked is recorded (a day is conceptualized as a time point, i.e., as an atomic time interval). Moreover, for each convex time interval (i.e., continuous sequence of days) someone works in a project, he is assigned to a single task and has a single pre-fixed deadline. Finally, for the whole (non-convex) time interval the person works in a project, he has the same role and the same manager. Olivé then proposes three different types of temporal relationship reifications: (1) *per instant*: the relationship r is reified into a different entity e for each time point at which r holds. In this example, for each working day in a given project, we have a different entity e capturing the work in that day; (2) *per interval*: a relationship r is reified into a different entity e' for each (convex) temporal interval during which r holds. In this example, according to Olivé, e' can then capture properties such as the actual behavior with respect to deadlines and objectives; (3) *per life span*: a relationship r is reified into a single entity e'' , which is the same during the whole life span of r . In this example, e'' can then capture properties such as assigned role and manager.

In light of our discussion, an obvious question that comes to mind is: what kind of ontological entities these different types of reifications are intended to represent? If we take (1), in the solution presented by Olivé, the reified relationship (*workDay* by Olivé, although *work@Day* would be clearer) may have properties such as *hoursWorked* and (zero-to-many) *producedDeliverables*. Olivé highlights two meta-properties of this entity: it is instantaneous and can only exist in that instant. Given the chosen name (and these meta-properties), a salient interpretation is that the reified entity represents an *episode*. Since episodes cannot change in a qualitative way, then both the attributes *hoursWorked* and *producedDeliverables* are immutable. Now, if the instances of *workDay* are only recorded in a *posteriori* manner (which seems to be case since days are assumed to be instantaneous and mereologically atomic) then the only reason for representing them is to have a historical view on these relations. Thus, according to our framework, an instance of *WorkDay* is a fully determined entity and can be interpreted as an historical event or an historical fact derived from it (e.g., the fact that John worked 10 hours and produced deliverables d_1 and d_2 in March 20th, 2013). Also, in this case, of course the event only happens because of pre-existing durants (e.g., capacities, rights, obligations) of the employee. In other words, also here, the *WorkDay* event is the manifestation of the *person-qua-employee's* qualities.

Let us take now case (2). In this case, Olivé's solution produces an entity termed *assignment* (which could perhaps be named *continuous assignment*), connecting an Employee and a Project for a continuous period of time. In this second case, representing an *assignment* as an episode may have some limitations, exactly because the *assignment* relationship may have modal and temporal properties: e.g., assignments can be fulfilled before the deadline or be delayed; can be realized in different ways (more or less accurately), can change in time in a qualitative way (for instance, the number of actual worked hours per day can change), can be suspended (because of an illness) or re-negotiated. In fact, an *Assignment* can even fail to manifest at all (for example, if the employee fails to actually work in the project).

Finally, let us analyze case (3). In that case, Olivé's solution reifies the relationship by something termed *participation*. Unlike cases (1) and (2), however, a participation is not correlated to a convex time interval. In other words, a participation can be

active or inactive, being hence correlated with multiple disconnected time intervals. In this case, we believe that the most salient interpretation of this relationship is that of a complex bundle of commitments that can change qualitatively in many ways (e.g., in the number of working hours, or in the money paid per worked hour), can bear modal properties (e.g., being active or not because of medical leave) and can be manifested by a number of possible processes. In these different possible processes a person can have different task assignments, which can be fulfilled or not, with different performance evaluations, in different dates with different amounts of effort, etc.

The understanding of relationships proposed in this article allows us to analyze and clarify a number of interesting points regarding Olivé's proposal. In particular, reified relationships do clearly stand for episodes (events) only in the case of reification per instant. In the other cases, when the reified entity is possibly the subject of changes or modal properties, it is an endurant that is being represented. In any case, even when only a perdurant is represented, there are always some endurants (e.g., capacities, dispositions, qualities, commitments) that are responsible for the happening of that perdurant [14].

6 Final Considerations

In this paper, we have contributed to the theory and practice of conceptual modeling by addressing one of its most fundamental constructs, namely, *relationships*. On the theoretical side, we have first used the notion of truthmaking to isolate a particular class of relations that are of fundamental relevance for conceptual modeling, namely, extrinsic relations. By pointing to the fact that these relations rely on external truthmakers (i.e., truthmakers that are neither the relata themselves nor are reducible to their intrinsic properties), we have shown that these relations need a special representation treatment, i.e., the corresponding relationships need to be reified. Then, we have systematically investigated the nature of this reified entity. On ontological and linguistic grounds, we have eliminated facts as possible truthmakers of these relations, and argued for the limited practical utility of taking episodes as representatives of reified relationships, since they can't account for temporal and modal changes. So, we have defended a view according to which relationships are relational truth-makers *of a certain kind*: they are full-fledged endurants, i.e., thing-like entities that change in time while maintaining their identity, accounting for the *way* a relation holds and develops in time. Then we have explored the *internal structure* of such endurants, analysing them as bundles of qualities while revising and clarifying previous intuitions on relators. On the practical side, we have demonstrated how the interpretation of relationships as endurants enables the reuse of ontological patterns proven useful for the modeling of taxonomic structures involving other endurant types. Moreover, we have shown how this theoretical framework can support the ontological analysis and conceptual clarification of existing proposals for relation reification existing in the conceptual modeling literature.

As a next step in our research, we intend to re-engineer the OntoUML metamodel to accommodate the proposal put forth here. This shall have an impact not only in the language, but also in its supporting methodological and computational tools [20].

Moreover, in order to allow these distinctions to be formally reflected in the ontological semantics of the language, we need to extend its underlying foundational ontology (UFO). Thus, we intend to extend the formal theory of relations in UFO and characterize its connection to the formal theory of events presented in [14].

References

1. Chen, P.: The entity-relationship model: Towards a unified view of data. *ACM Transactions on Database Systems* **1**(1) (1976)
2. Chen, P.: English Sentence Structure and Entity-Relationship Diagrams. *Information Sciences* **29**(2–3), 127–149 (1983)
3. Moltmann, F.: Events, Tropes and Truthmaking, *Philosophical Studies* (2007)
4. Guizzardi, G., *Ontological Foundations for Structural Conceptual Models*, Telematics Instituut Fundamental Research Series, No. 015, ISSN 1388-1795, The Netherlands (2005)
5. Guizzardi, G., Wagner, G.: What’s in a relationship: an ontological analysis. In: Li, Q., Spaccapietra, S., Yu, E., Olivé, A. (eds.) *ER 2008*. LNCS, vol. 5231, pp. 83–97. Springer, Heidelberg (2008)
6. Costal, D., Gómez, C., Guizzardi, G.: Formal Semantics and Ontological Analysis for Understanding Subsetting, Specialization and Redefinition of Associations in UML, 30th Int. Conf. on Conceptual Modeling (ER 2011), Brussels, Belgium (2011)
7. Guizzardi, G., Zamborlini, V., Using a Trope-Based Foundational Ontology for Bridging different areas of concern in Ontology-Driven Conceptual Modeling. *Science of Computer Programming* (2014)
8. Guizzardi, G. et al.: In the Defense of a Trope-Based Ontology for Conceptual Modeling: An Example with the Foundations of Attributes, Weak Entities and Datatypes, 25th Int. Conf. on Conceptual Modeling (ER’2006), Tucson
9. Bock, C., Odell, J.: A More Complete Model of Relations and Their Implementation: Relations as Object Types. *Journal of Object-Oriented Programming* **10**(3), June 1997
10. Halpin, T., Morgan, T., *Information Modeling and Relational Databases*, Morgan Kaufman (2008)
11. Fine, F.: Neutral relations. *The Philosophical Review* **109**, 1–33 (2000)
12. Masolo, C., Borgo, S.: Foundational choices in DOLCE. In: Staab, S., Studer, R. (eds.) *Handbook on Ontologies*, pp. 361–381. Springer-Verlag, Berlin (2009)
13. Simons, P.: *Parts: A Study in Ontology*. Oxford University Press (1997)
14. Guizzardi, G., Wagner, G., de Almeida Falbo, R., Guizzardi, R.S., Almeida, J.P.A.: Towards ontological foundations for the conceptual modeling of events. In: Ng, W., Storey, V.C., Trujillo, J.C. (eds.) *ER 2013*. LNCS, vol. 8217, pp. 327–341. Springer, Heidelberg (2013)
15. Ferrario, R., Guarino, N.: Commitment-based modeling of service systems. In: Snene, M. (ed.) *IESS 2012*. LNBIP, vol. 103, pp. 170–185. Springer, Heidelberg (2012)
16. Nardi, J.C., de Almeida Falbo, R., Almeida, J.P.A., Guizzardi, G., Pires, L.F., van Sinderen, M.J., Guarino, N., Fonseca, C.M.: A Commitment-based Reference Ontology for Services, *Information Systems*, Elsevier (2015). doi:10.1016/j.is.2015.01.012
17. Fine, K.: Things and their Parts. *Midwest Studies in Philosophy* **23**(1), 61–74 (1999)
18. Mulligan, K., Smith, B.: A Relational Theory of the Act. *Topoi* **5**(2), 115–130 (1986)

19. Guarino, N., Welty, C.: An Overview of OntoClean. In: Staab, S., Studer, R. (eds.) *Handbook on Ontologies*, 2nd Edition. Springer-Verlag, pp. 201–220 (2009)
20. Guizzardi, G.: Ontological patterns, anti-patterns and pattern languages for next-generation conceptual modeling. In: Yu, E., Dobbie, G., Jarke, M., Purao, S. (eds.) *ER 2014. LNCS*, vol. 8824, pp. 13–27. Springer, Heidelberg (2014)
21. Olivé, A.: Relationship Reification: A Temporal View. In: Jarke, M., *Advanced Information Systems Engineering: Proceedings of 11th Int. Conf., CAiSE 1999*, pp. 396–410 (2009)