

# Stereotypical Encounters of the Third Kind

Colin Atkinson<sup>1</sup>, Thomas Kühne<sup>2</sup>, and Brian Henderson-Sellers<sup>3</sup>

<sup>1</sup> University of Kaiserslautern, Kaiserslautern, Germany  
colin.atkinson@ieee.org

<sup>2</sup> Darmstadt University of Technology, Darmstadt, Germany  
kuehne@informatik.tu-darmstadt.de

<sup>3</sup> University of Technology, Sydney, Australia  
brian@it.uts.edu.au

**Abstract.** As one of the UML’s main extension mechanisms, stereotypes play a crucial role in the UML’s ability to serve a wide and growing base of users. However, the precise meaning of stereotypes and their intended mode of use has never been entirely clear and has even generated much debate among experts. Two basic ways of using UML stereotypes have been observed in practice: one to support the classification of classes as a means of emulating metamodel extensions, the other to support the classification of objects as a means of assigning them certain properties. In this paper we analyze these two recognized stereotype usage scenarios and explain the rationale for explicitly identifying a third form of usage scenario. We propose some notational concepts which could be used to explicitly distinguish the three usage scenarios and provide heuristics as to when each should be used. Finally, we conclude by proposing enhancements to the UML which could support all three forms cleanly and concisely.

## 1 Introduction

In the context of object-oriented modeling, the notion of stereotypes was first introduced by Rebecca Wirfs-Brock [1]. She recognized that it is useful to classify objects according to the responsibilities they fulfill within a certain system. In addition to being classified by a natural domain specific type (e.g., `ATMKeypad`), such an object would also be classified as an “interface object”, others as “coordinator objects” etc. The categories for this *secondary* classification are drawn from the roles the objects play within the system (see Fig. 1). Names of natural types (i.e., class names) would typically be derived from the problem domain whereas the names of stereotypes would typically be derived from the solution structure. Thus, stereotypes can essentially be understood as a way of “branding” objects of a particular primary (natural) type with additional roles or properties. Since Wirfs-Brock used stereotypes primarily for explanation purposes she did not map the stereotype concept to a mechanism such as multiple inheritance.


This ability to add additional classification information to a concept in a concise and straightforward way made Wirfs-Brock’s stereotype mechanism ideal for the designers of the UML who were looking for a way of allowing users to

add additional modeling concepts to the UML without having to extend the metamodel directly. Since the introduction of the UML, therefore, stereotypes have primarily become known as the “lightweight” extension mechanism of the UML, allowing users to tailor the language to their own needs without having to directly deal with the UML metamodel. The choice of stereotypes as a (restricted) language extension mechanism was not only motivated by usability concerns, however, but also as a concession to many tool vendors who at the time were unable to support direct metamodel specialization.

In the sense that it allows model elements to be “branded” with additional classification information, the stereotype mechanism defined in the UML standard [2] is entirely faithful to the original concept of Wirfs-Brock. However, the “official” UML extension mechanism differs from the original usage of stereotypes in one very important way. Whereas Wirfs-Brock used stereotypes to brand objects, the UML standard defines stereotypes as a way of branding classes. In the terminology of the UML’s four level model hierarchy, the original stereotype mechanism was used to brand elements at the  $M_0$  level, while the UML standard presents stereotypes as a means for branding  $M_1$  level elements.

In the previous paragraph, we were very careful to refer to the “official” *definition* within the UML standard, because in practice many UML users have applied the UML stereotype mechanism in the way intended by Wirfs-Brock, that is, as a way to brand  $M_0$ -level elements. Thus, in practice one finds that users of the UML use stereotypes in both styles—to brand  $M_1$ -level elements and to brand  $M_0$ -level elements. It seems that the notational economy of stereotypes makes them an attractive alternative to the traditional ( $M_1$ -level) inheritance mechanism as a way of branding object, even when the latter captures the real-world scenario more accurately. One can therefore observe in the literature a large number of stereotypes applied in a very loose style. The result has been an intense debate among researchers and experts about how stereotypes should and should not be used, e.g., whether the original Wirfs-Brock usage should be considered valid in the context of UML [3].

It is overly simplistic, however, to assume that the two aforementioned basic usage scenarios span the full space of options. On closer analysis, it turns out that many real world scenarios are not naturally represented by either of these usage styles, but that a third interpretation that simultaneously brands a class and all its objects is the most appropriate. In other words, often, properties of real world concepts naturally imply both forms of classification at the same time.



**Fig. 1.** Secondary Classification

Our goal in this paper is to investigate this phenomenon by fully characterizing the three forms of stereotype usage scenarios and describing a strategy for working out which form should be used for a particular real-world scenario. In the following section, we first analyze and give examples of the two basic usage scenarios and then elaborate on the third kind. We subsequently equip modelers with a litmus test to identify which kind of stereotype they are facing, provide guidelines as to when a certain form should be used, propose notational measures to make the intent of a particular stereotype explicit, and finally outline how the three kinds could be most naturally accommodated in a clean multilevel approach.

## 2 Generally Recognized Stereotypes Usage Scenarios

In this section we illuminate and discuss the two main types of stereotype usage scenarios which have been recognized and debated in the literature [3]. Then, we introduce the third form which—although not explicitly recognized as such in the literature to date—is possibly one of most common real world modeling scenarios.

When discussing how stereotypes are used in practice, particularly when analyzing examples of stereotype applications, it is important to bear in mind that the original intent of the modeler can never be known with absolute certainty. There is always an element of uncertainty in our interpretations of what the modeler originally had in mind. This can be regarded as being both a strength and a weakness of stereotypes in the UML. It is a strength because it means that users are able to exploit stereotypes without getting bogged down in detailed semantic issues, but it is a weakness because if the modeler had a clear meaning in mind this knowledge is lost. Consequently, in this and the following section, when we classify stereotype usage scenarios we do so on the basis of our interpretation of the real world situation that the modeler is trying capture.

### 2.1 Type Classification

The UML 1.4 specification states that:

A stereotype is, in effect, a new class of metamodel element that is introduced at modeling time [2] (page 3-33).

Hence, the motivation for the use of stereotypes is essentially to extend the standard modeling vocabulary offered by the UML by facilitating the additional classification of types. The mechanism by which this is achieved within UML 1.4 is described as—

... a way of defining virtual subtypes of UML metaclasses with new metaattributes and additional semantics [2] (page 2-79).

In other words, a new type of modeling element becomes available that users can make use of at the  $M_1$  level. According to this definition, with the exception of “redefining stereotypes” [4], stereotyped elements behave as if they were instances of the base modeling element. Thus, stereotyping of this kind provides a way of branding types so that they can be associated with special semantics. If one attaches a stereotype (e.g., `<<commented>>`) to a class (e.g., `Worker`), (see Fig. 2) for example, the effect is the creation of a (virtual) subclass of the  $M_2$ -level metaclass `Class`. The stereotyped class (`Worker`) is then considered to be an instance of this virtual subclass.



**Fig. 2.** Stereotyped Type

Thus in our example the stereotype `<<commented>>` brands `Worker` not its instances. This becomes most evident when considering stereotypes such as `<<abstract>>` or `<<interface>>` which unambiguously refer to the classes they stereotype, since there are no direct instances of abstract classes or interfaces. Even when the stereotyped model elements have instances, e.g., objects of a framework instantiation, the stereotypes `<<application>>` and `<<framework>>` [5] clearly document the nature of the classes and have no bearing whatsoever on their instances.


**Tagged Values** Tagged values are intended to be used in tandem with stereotypes as a means of assigning values to user types. They are a way to supply additional information along with the branding (see Fig. 2, specifying the commenter). Tagged definitions correspond to (meta-)attributes, defining what type of value can be assigned to a type, and tagged values actually define the value. Thus, tagged definitions essentially correspond to attributes of (virtual) metaclasses, while tagged values correspond to slots of user types. Tagged definitions and tagged values do not have to be used in tandem with stereotypes but, when they are, they reinforce the model of a stereotype as defining a virtual metaclass, and a stereotyped element as representing an  $M_1$ -level type.

## 2.2 Instance Classification

The previous section described the “official” purpose of stereotypes as defined in the UML standard. As such, this has been characterized by some authors as the “correct” way of using stereotypes. However, there are many practical examples of stereotype usage in the UML literature that can only reasonably be interpreted as serving to classify objects at the  $M_0$  level. For example, Conallen uses `<<Applet>>` to stereotype a class `OnLineGame` [6] (page 169). It seems obvious that instances of `OnLineGame` should be considered to be applets, whereas `OnLineGame` itself is simply an ordinary class. An alternative way to obtain the assumed modeling scenario would therefore have been to make `OnLineGame` a subclass of an  $M_1$ -level superclass `Applet`. Note that the JAVA AWT library does exactly this and provides `Applet` as a class from which to inherit.

Likewise, in Fig. 3 the song (i.e. instance) “Bohemian Rhapsody” can be regarded as a “soundStream” while this cannot be said of the type `Music`. Again, the

appropriate way of expressing that instances of `Music` are sound streams would be to define `SoundStream` as the superclass of `Music`. A shorthand way of expressing such a modeling scenario which uses angled brackets (e.g., `<SoundStream>`) instead of guillemets (`«soundStream»`) in the class compartment is proposed in [7].



**Fig. 3.** Stereotyped Instance

Even within the UML specification it appears that one can find stereotype usages that do not conform precisely to the definition of the stereotype mechanism. The stereotype `«parameter»`, which is applicable to associations, has the intended meaning of flagging the stereotyped association as being of temporary nature. However, this property applies primarily to the instances of associations—the links—rather than to the associations themselves. Whether the association itself is affected by this stereotype is a matter of debate and will be revisited in the subsequent section.

The aforementioned examples reflect the original intent of Wirfs-Brock in using stereotypes, where she observed “I find it useful to classify objects according to their primary purpose as well as their *modus operandi*.” [1]. The specific examples given by Wirfs-Brock for such *modus operandi*—Controller, Coordinator, Structurer, Information Holder, Interfacer— all refer to the roles of *objects* within the system.

Theoretically it is possible, given the current wording of the UML specification<sup>1</sup>, to achieve this form of classification by applying stereotypes directly to objects. For example, `bohemianRhapsody` could be stereotyped with `«informationHolder»` or `«soundStream»` to document the object’s role in addition to its type. However, there are few if any examples of this style of usage in practice.

End users of the UML normally attempt to achieve this form of classification (i.e., instance classification) by applying stereotypes to the type, as illustrated in Fig. 3. Without literally interpreting the UML standard they merely assume that they can use stereotypes in this way. However, attempts to capture the nature of instances, such as `bohemianRhapsody`, through the use of class stereotypes simply cannot work [8], since instantiation is not transitive. To illustrate this, let’s assume that the stereotype `«soundStream»` defines a tagged value “format” to be associated with stereotyped elements with values “mp3”, “wav”, “au”, etc. In our example, the tagged value (e.g., “{format = mp3}”) would be attached to `Music` and not to `bohemianRhapsody`. This obviously does not make any sense unless the stereotype `«soundStream»` was actually meant to introduce new `soundStream` types. In the latter case, all instances of `Music` would be “soundStream” objects conforming to the format specified at `Music`. In this case, however, the stereotype name should have been `«soundStreamType»`. Whatever special quality is to be expressed by the stereotype, it always refers first and foremost to the stereotyped element (`Music`) and not to its instances. Consequently, if one actually aims to obtain “soundStream” instances then one needs to stereotype their class as a

<sup>1</sup> The question is whether the term “model elements” includes both  $M_1$ -level types and  $M_0$ -level instances.

“soundStream type”. Yet, there are clear cut examples where the intent of class stereotypes is to brand objects instead of classes. For instance, Sauer and Engels explicitly state that—

Other stereotypes are used to distinguish different types of (semantic) application objects [9] (page 249).

Although stereotyping to achieve instance classification (as depicted in Fig. 3) seems to be used by many UML modelers it is described by some authors as the “wrong” way of using stereotypes in the UML, since it contradicts the stated purpose of stereotypes in the standard and is essentially redundant<sup>2</sup> because classification of  $M_0$ -level elements is covered by the traditional  $M_1$ -level specialization (inheritance) mechanism.

It is sometimes argued that traditional  $M_0$ -level element classification by using inheritance is not powerful enough to express certain advanced semantic properties [10]. For instance, the quality of being a clock (e.g., an active notion of ticking) could be captured with a stereotype so that a tool could generate appropriate code accordingly. This special instance behavior cannot be achieved by using a superclass `Clock`. However, this is an artificial result of the different levels of support given to stereotypes and subtyping by UML tools. Stereotypes do not inherently provide any better support for capturing such semantics. If a tool can exploit the existence of a certain stereotype to attach special semantics to modeling elements it could just as well exploit the fact that the model element is a subtype of a certain special supertype (e.g., `Clock`).

### 2.3 Transitive Classification

In the previous two sections we identified two commonly seen ways of using stereotypes in practice: one “official” way that conforms to the UML standard (i.e., type classification, such as `<<interface>>`) and one “unofficial” way that does not conform to the standard way (i.e., instance classification, such as `<<soundStream>>`) and is thus strictly speaking incorrect. In line with the current literature we have until now implied that these are the only two options. However, the discussion in the previous section already pointed to the fact that this is an over simplification. In describing the `<<parameter>>` stereotype we pointed out that it is primarily intended as a way of branding instances, but also left open the question as to whether its also makes sense to think of the type (the association itself) as being branded by this stereotype.

In general, we make the following observations about the problem of mapping real world scenarios to the classification forms identified in the previous sections:

- There is no published systematic procedure to work out whether type classification or instance classification is the best match.

<sup>2</sup> When one actually wants to create new kinds of objects which possess features over and above slots and links, one indeed needs to resort to a (virtual or non-virtual) extension of the  $M_2$ -level. This, however, is almost never the case.


- Often both forms of classification apply.

The «parameter» example is a good case in point. Clearly, this scenario includes instance classification, since it is ultimately the links which exhibit a temporal property. But it is also the case that an association stereotyped with «parameter» is distinguished from other associations. The «parameter» label is essential information for a reader of the class diagram to understand that this particular association is not motivated by structural considerations but only indicates a temporal communication path.

Such a scenario clearly constitutes a third form of usage scenario, which we refer to as *transitive classification*. This usage scenario refers to situations in which the intent of the user is to brand both the type and the instances of the type. In other words, the user intends the branding of the type to *transitively* apply to its instances as well. Thus, in a sense this scenario represents a combination of the type and instance classification scenarios described in the previous sections.

In fact, when a modeler wishes to add a new natural type (i.e., add to the set of *primary* classifiers), the intent is usually to apply this third transitive form. When one removes secondary classification (such as «controller» or «commented») from a formerly decorated element the element will essentially remain unchanged. When, however, one removes a stereotype whose purpose is to extend the set of modeling element types (such as «parameter») the natural type of the model element will change from a special kind to a standard type, involving significant changes in semantics, including completely changing the nature of the instances. Whether their type is said to be «commented» or not is not relevant for instances, but whether their type is said to be, e.g., «active» certainly heavily influences their behavior.

As a typical example of “transitive classification” consider the situation in Fig. 4. It is obvious that instances of Account should be persistent and not the class Account itself. However, as well as requiring the instances to support a writeObject() operation, for the persistence mechanisms to work one must also require certain properties of class Account itself. Here, Account is required to realize the Serializable interface and also specify a



**Fig. 4.** Transitive Stereotype

format according to which its instances are to be serialized. The «persistent» stereotype as used in Fig. 4 therefore affects both type ( $M_1$  level) and instances ( $M_0$  level). Unlike in pure type classification (e.g., stereotype «commented») the intent is that distinct type properties transitively apply to instances as well.


As a second example, consider the use of design pattern role names as stereotypes. Sometimes such a role name represents type classification only but often its intent is best described by transitive classification. For instance, a class

carrying the stereotype `«composite»` is most naturally interpreted as having instances who play the role of a composite<sup>3</sup> supporting this role by the presence of certain slots (e.g., `components`) and being able to respond to certain messages (e.g., `addComponent`) and at the same time exposing certain class properties, such as realizing a `Component` interface.

It is important to reiterate that we are referring to the modeler’s intent when identifying this (and the other) stereotype usage scenarios. With the current definition of stereotypes and instantiation in the UML it is not possible to syntactically distinguish transitive classification from pure type classification.

### 3 Classification of Stereotype Usage Scenarios

Having motivated and introduced the different kinds of stereotypes usage scenarios, we now present a systematic classification of the different scenarios and describe concrete criteria for choosing between them. This classification is orthogonal to other classification approaches, such as that of Glinz et al. [4] that concentrate on categorizing usages conforming to the UML standard.



**Fig. 5.** Classification forms

Fig. 5 shows a practical example of the three different cases side-by-side. In Fig. 5(a) the modeler intends to brand class `Video` only, adding the information that this product type has a tax rate of 16% associated with it. The tax rate information is not accessible from instances of `Video`. Fig. 5(b) shows the opposite extreme in which the modeler wants to classify `Video` objects and give them each a separate tax slot. Without any additional constraints there is nothing to control the values assigned to the tax slots of individual instances of `Video`. These values may or may not be the same. Fig. 5(c) shows the intermediate situation

<sup>3</sup> Pattern role names are a good example for secondary classification. Typical primary types of objects playing the role of a composite would be `Group`, `Assembly`, etc. The classification as a composite object is secondary only and orthogonal to the primary type hierarchy.



in which the modeler wants a tax rate to be visible at the level of individual objects, but wants to fix the value of the tax rate for each object at the value of the type. Thus, the intent in Fig. 5(c) is to brand instances as well as the type. Notice that with current stereotype notation it is impossible to distinguish the situations shown in Fig. 5(a) and Fig. 5(c) from the stereotyped class alone.

Table 1 summarizes the three different scenarios in tabular form. The left hand column represents type classification, in which the intent is to brand the type only, and not the instances. The middle column represents regular instance classification in which the intent in using a stereotype is only to brand the instances of a type, not the type itself. As has already been pointed out, this situation should be handled more appropriately with regular inheritance (i.e., the classifier of the instances is defined to be a subtype of a supertype corresponding to the stereotype). Finally, the right hand column of Table 1 represents the transitive classification usage scenario in which the intent is to brand the instances of the type and the type itself.

**Table 1.** Classification scenarios

<i>Type Classification</i>	<i>Instance Classification</i>	<i>Transitive Classification</i>
e.g., «abstract»	e.g., «applet»	e.g., «persistent»
affects types only	affects instances only	affects types & instances

The difference between these three forms can also be explained in terms of their effects on the type and/or instance facets of the model element to which they are applied. As explained in [11], model elements representing types generally have two facets—a type facet which describes the properties of their instances (e.g., attributes and methods), and an instance facet which describes the element’s properties as an object (e.g., tagged values). In terms of these facets—

1. type classification purely affects the *instance facet* of types,
2. instance classification affects the *type facet* of types only, and
3. transitive classification affects *both* facets of types.

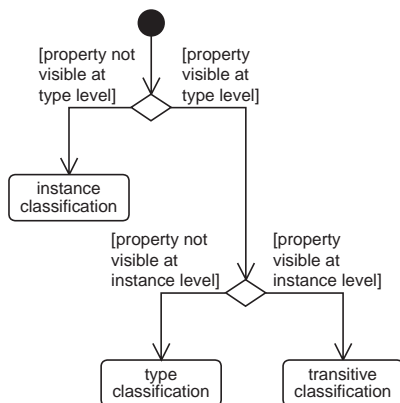
With this terminology we have an easy way to find out which category a certain stereotype belongs to. By checking which facets of a model element at the  $M_1$  level are influenced by the stereotype, one can immediately tell which kind of stereotype one is dealing with.

The systematic view just established is perhaps also helpful in explaining why modelers use stereotypes for instance classification even though the UML definition explains them to have the effect of type classification. As we already mentioned, in such cases the use of inheritance at the  $M_1$  level would be more appropriate, since inheritance is the most natural way to influence the type facet of a type. The type facet of a subtype is simply extended with that of

the supertype, which is exactly what is needed when one wishes to make sure that a set of instances carry certain properties. However, it is the case that a generalization relationship is a link between two types at the  $M_1$  level, i.e., such a link must be regarded as belonging to the instance facet of a type. Therefore, the use of a stereotype—which can only affect the instance facet of a type in the non-transitive case—seems understandable, since one can imagine that the modeler associates a designated type facet with a designated instance facet. One way to make this assumption true, i.e., truly make the stereotype have the effect of transitive classification, is to use a constraint which demands any type equipped with this stereotype to inherit from a particular class (e.g., `SoundStream`). Another, more cumbersome way is to use a much more complex constraint which spells out the full type facet required for “`SoundStream`” types.

While we have just established a way to actually make a stereotype realize transitive classification using plain UML only, we still need to be aware of the difference between a stereotype whose sole purpose is to classify instances and a stereotype which affects both type and instances, i.e., transitive classification. How to make the choice between the two cases and subsequently how to notationally distinguish them is the subject of the following two sections.

### 3.1 Usage Guidelines



**Fig. 6.** Supporting the decision process

nor does the type need to be distinguished from other types), instance classification is the obvious choice. In terms of the Video example, the property “price” is a good illustration (see Fig. 7), assuming the price can vary from video to video, this would be represented as an attribute at the type level, and be given different values for each instance. No information is required at the type level.


When there is some property that could reasonably be applied to the type, the situation becomes more difficult. The existence of a type-level property implies transitive or type classification. As specified in Fig. 6, one now needs to

As conceptual modeling mechanisms, with clearly distinct semantics (captured in Table 1), the difference between the three usage scenarios is easy to understand. What is much more challenging, however, is working out which scenario best fits a particular real world situation. A systematic process for making this decision is given in Fig. 6.

Instance classification is the most easy to match to real world situations, since it actually corresponds to an ordinary classification of instances, where one uses a type to describe the common properties of a set of objects. Thus, in scenarios where there is no type level information (i.e., no properties of instances are fixed at the type level

decide whether instances are affected or not. The difficulty in doing this arises because there is virtually no form of type-level information which cannot, in some form or another, be regarded as relevant to the instance level as well. Take the «commented» stereotype, for example. Normally, one would think that software objects do not care about whether their classes are commented. But on the other hand, one cannot exclude a situation in which someone wants to reject all objects from being dynamically added to his system unless they can testify that their classes have been developed with the use of comments.

Also, in the case of the video example, should the property of the tax rate be associated with the class only, with the objects only or with both (see Fig. 5)? Assuming that the tax rate is fixed for all videos, it is reasonable to view the tax value as belonging to the type only. However, from the perspective of clients of video objects, the tax rate is just a normal property that they might be interested in (like price). The fact that it is constant for all video instances might be irrelevant at the instance level. Since there is a case for viewing it as either a type-level property or an instance-level property, there is also a case for viewing it as both (i.e., transitive property).



**Fig. 7.** Priced Object

Thus, from these examples it can be seen that a decision for one of the classification scenarios purely based on the nature of concepts in a real world scenario cannot be made. The choice heavily depends on the needs and intent of the modeler. In the case of the video example, we are assuming that the real world scenario dictates that taxes are constant for videos of a particular kind. This, therefore, implies that “taxRate” is at least a type-level property—ruling out the pure instance classification scenario. The question then is whether or not “taxRate” should also be viewed as an instance-level property. This is where the purpose of the model and the needs of the modeler come in. The fundamental question at this point is whether the modeler wants tax rates to be visible to clients of the instances. In other words, should a client of a video instance be able to directly determine (i.e., see) its tax rate? If the answer is yes, the tax property is best thought of as transitive (i.e., belonging to the type and the instances) and the transitive classification approach should be chosen. If no, the tax property is best regarded as merely a type-level property, and the type classifications scenario should be chosen.

What remains to be discussed is how to make sure one uses a proper name for the stereotype. In section 2.2 we suggested «soundStreamType» instead of «soundStream» as the name for a stereotype so that the intent is better reflected. We now provide a litmus test for determining whether a stereotype name is appropriate or not. Assuming that the correct kind of stereotype has already been determined, and has been called “S”, the question to ask is:

Do I under any circumstance want instances to be understood as “S”-instances?

If the answer is yes, then the stereotype name is wrong. For example, if it is the case that `OnlineGame` instances ought to be regarded as “applet” instances then “Applet” is a good candidate for an `Applet` superclass name, but not for the stereotype. Just as class names spell out what their instances (i.e., objects) are (e.g., applets) stereotype names ought to spell out what their instances (i.e., types) are (e.g., `AppletType`). A stereotype name does not necessarily need to have the “-type” suffix, though. A stereotype name `<<component>>` makes sense since types branded with this stereotype are particular components, whereas their instances are usually referred to as component instances.

It cannot be overemphasized that if one wants instances to be regarded as “S”-instances then one is not really dealing with a proper stereotype name but with the name of an  $M_1$ -level supertype. What can be achieved by a stereotype `<<AppletType>>`, however, is that classes like `OnlineGame` inherit from a superclass `Applet`. In this way—exploiting the terribly imprecise meaning of “is-a” [8]—not only an `Applet` is-an applet (is-a  $\rightarrow$  instance-of) but `Applet` is-an applet (is-a  $\rightarrow$  kind-of) as well, which seems to be the primary motivation for (incorrectly) using an `<<Applet>>` stereotype in such cases.

## 4 Notational Support

Having identified the primary stereotype usage models and discussed the heuristics that can be used to choose between them, we now discuss how the scenarios could be distinguished notationally. Modelers not only need a systematic way of obtaining the kind of stereotype most appropriately fitting the situation in hand, they also need a way to unambiguously document their decision.

In the first subsection we discuss strategies for distinguishing them using the current version of the UML and then, in the following subsection, we discuss possible approaches that might be employed in future versions.

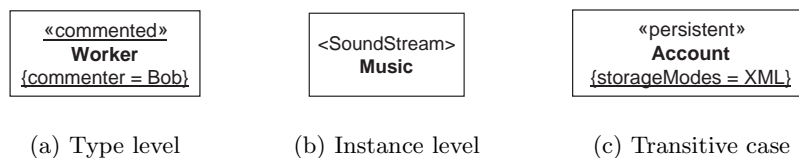
### 4.1 Current UML

Fig. 5 illustrates that, by considering class `Video` alone, it is not possible to deduce which of the two stereotype kinds was intended. We therefore propose slight notational enhancements to the current version of the UML so that all three kinds of stereotypes can be distinguished. Figure 8 depicts our suggestions for all three cases.

Note the use of an adjectival stereotype name in Fig. 8(c). It is equally valid to think of both `Account` itself and an `Account` instance as supporting persistence.

Also, note the underlined tagged values which indicate that this information is meant to be useful for the type level only. Whenever a tagged value needs to be visible to instances as well (as in “`taxRate = 16%`”) it should not be underlined.

Note, however, that while underlining tagged values works for conveying intent, i.e., restricting the visibility of a class feature (e.g., `commenter`) to be visible to the class only as opposed to all of its instances as well, is not suitable to control whether a class feature (such as `taxRate`) shall be a part of an object’s



**Fig. 8.** Notational support

feature list (as in Fig. 5(c)). In order to achieve the level of expressiveness to independently control visibility and feature extent it is best to assume a true multi-level modeling approach for UML modeling.

## 4.2 Advanced Support

In [12] it is suggested that the traditional two level ( $M_1$  and  $M_0$ ) modeling assumption be abandoned and UML modeling be based on a multi-level approach involving user metatypes in addition to user types and user instances.

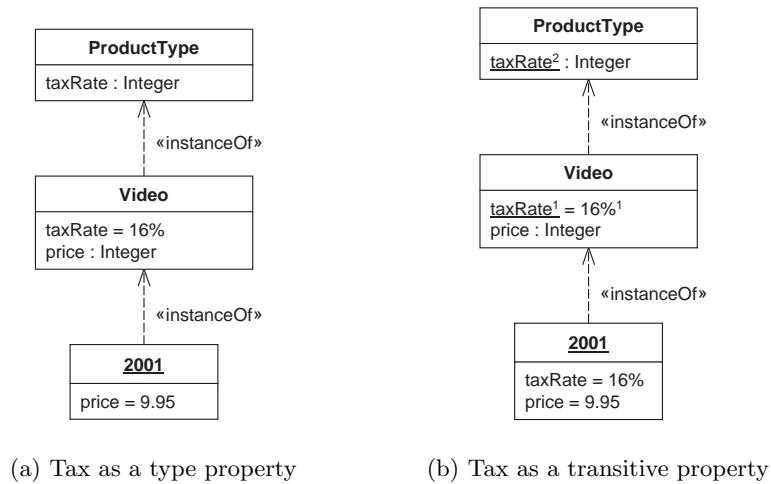
With the availability of a user modeling level above the normal  $M_1$  type level and in the presence of “deep instantiation” [12] it becomes almost trivial to control where the features involved in the scenario of Fig. 7 become visible and where not.

Figure 9 illustrates how one can control the visibility of “taxRate”. In Fig. 9(a) “taxRate” is just an ordinary (meta-)attribute, giving rise to a slot at type **Video** only. This is analogous to a tagged definition/tagged value pair in the UML. In contrast, Fig. 9(b) specifies “taxRate” to be of potency two, i.e., makes it appear as an attribute (rather than a slot) at type **Video**. Therefore, this attribute gives rise to a corresponding slot at instance 2001.

Note the underlining of “taxRate” in Fig. 9(b), meaning that “taxRate” is a so-called dual field rather than a simple field. In contrast to simple fields, dual fields may already be assigned a value before they become a slot. This allows the value of “16%” to already be assigned a value at the type level. Furthermore, we used “16%<sup>1</sup>” to indicate that this value is constant for the level below. Had we used “16%” instead, it would have meant that the type **Video** is associated with a tax rate of 16% but that any **Video** has a “taxRate” slot which may be assigned any individual value (analogue to the price slot).

Such a multi-level modeling approach not only makes stereotyping and tagged values redundant as mechanisms in their own right<sup>4</sup> but furthermore makes it trivial to control the scope of properties independently of the level where they are introduced. The modeler’s intent can be captured in a concise and intuitive manner without ever requiring constraints enforcing certain type facets to be present.

<sup>4</sup> Notational shortcuts as presented in Fig. 8 are still valuable and should be made available.



**Fig. 9.** Choosing the desired tax scope

## 5 Conclusion

Thanks to their notational conciseness and convenience of use, stereotypes have earned a permanent place in the affections of UML users and have become an indispensable part of the UML’s extensibility mechanisms. However, as explained in this paper and many others, they are often deployed by users in ways that do not strictly conform to the UML standard and can therefore be ambiguous or misleading to readers of the model. Our goal in this paper has been to first identify and understand the different de facto ways in which stereotypes are used by UML modelers (whether they be “official” or “unofficial”) and then to consolidate practical guidelines for working out which scenario makes sense under which circumstances. We then elaborated on how the different stereotype usage scenarios could be distinguished notationally, both in terms of existing UML notation and in terms of possible future notations.

We observed that a suitable interpretation of the UML standard would actually allow both type and instance classification usages for stereotypes but beyond that established that a simplistic reduction to two basic modes is insufficient. We identified that many real world scenarios are actually best thought of as combining these two forms so that they simultaneously classify the type as well as its instances. For this form of usage scenarios we introduce the term “transitive classification”.

In fact, one of the possible reasons for the extensive debate on what stereotypes mean and how they should be used has been that the three forms have never hitherto been made fully explicit. It turns out that many examples of stereotype applications in the literature are best understood as transitive classification, rather than simply type (“official”) or instance (“unofficial”) classification. Since

transitive classification best matches the properties of many real-world situations, it will continue to be used in the future regardless of the intended “official” stereotype semantics and regardless of how well this case can be expressed with standard stereotypes. However, in the spirit of providing a language which is as expressive and accurate as possible we believe that it is important for future versions of the UML to take all three stereotype usage scenarios into account, and to provide users with explicit notation for specifying which scenario they are intending to use.

## References

- [1] Rebecca Wirfs-Brock, Brian Wilkerson, and Lauren Wiener. Responsibility-driven design: Adding to your conceptual toolkit. *ROAD*, 1(2):27–34, July–August 1994.
- [2] OMG. *Unified Modeling Language Specification*, 2000. Version 1.4, OMG document ad00-11-01.
- [3] Brian Henderson-Sellers. Notes on the use of subtypes and stereotypes in the UML metamodel. *Journal of Database Management*, 13(2):44–52, 2002.
- [4] Stefan Berner, Martin Glinz, and Stefan Joos. A classification of stereotypes for object-oriented modeling languages. In Robert France and Bernhard Rumpe, editors, *Proceedings of the 2<sup>nd</sup> International Conference on the Unified Modeling Language*, LNCS 1723, pages 249–264, Berlin, October 1999. Springer Verlag.
- [5] Marcus Fontoura, Wolfgang Pree, and Bernhard Rumpe. *The UML Profile for Framework Architectures*. Addison-Wesley, 2002.
- [6] Jim Conallen. *Building Web Applications with UML*. Addison-Wesley, Reading, MA, 2000.
- [7] Colin Atkinson and Thomas Kühne. Strict profiles: Why and how. In Andy Evans, Stuart Kent, and Bran Selic, editors, *Proceedings of the 3<sup>rd</sup> International Conference on the UML 2000, York, UK*, pages 309–323, LNCS 1939, October 2000. Springer Verlag.
- [8] Colin Atkinson, Thomas Kühne, and Brian Henderson-Sellers. To meta or not to meta – that is the question. *Journal of Object-Oriented Programming*, 13(8):32–35, December 2000.
- [9] Stefan Sauer and Gregor Engels. UML-based behavior specification of interactive multimedia applications. In *Proceedings of the IEEE Symposium on Human-Centric Computing Languages and Environments*, pages 248–255. IEEE, September 2001.
- [10] Bran Selic. Re: Ask about UML. email communication 10:15:56 -0400, September 2001. email to third author and others.
- [11] Colin Atkinson and Thomas Kühne. Meta-level independent modeling. In *International Workshop Model Engineering (in Conjunction with ECOOP’2000)*. Cannes, France, June 2000.
- [12] Colin Atkinson and Thomas Kühne. The essence of multilevel metamodeling. In Martin Gogolla and Cris Kobryn, editors, *Proceedings of the 4<sup>th</sup> International Conference on the UML 2000, Toronto, Canada*, LNCS 2185, pages 19–33. Springer Verlag, October 2001.