

THREE-LEVEL DEFINITION

Michael Sussna

Linguistics 130
Semantics
Professor Suzanne Kemmer
December 6, 1991

Three-level definition describes entities in terms of *internal properties*, *surface properties*, and *external properties*. Dividing definitions into these particular facets allows us to make descriptions manageable and corresponds to our natural ways of partitioning the world. Familiar distinctions such as form/content, structure/function, intrinsic/extrinsic, and intension/extension can be seen to fit within this three-level framework. In addition, such tensions as those between checklist vs. prototype definition, essential vs. optional attributes, and dictionary vs. encyclopedic description fall into place and are resolved.

Internal properties include those which give the internal structure to an entity, whether it is a static object or a dynamic object such as a process. This internal structure includes the parts of the object and the pattern of interconnection of the parts. For a simple abstract entity such as a square, we might say that the parts are the four sides and that the pattern of interconnection is the set of requirements that the four sides have the same length, that the four angles at each vertex be right angles, and that the four sides lie in a plane relative to the perceiver's perspective. One could consider the interconnections or other structural information to be "parts" as well, but some researchers make a distinction between the two.¹ The internal structure of a process might include the sequence of events, the resources involved, and the dependencies between the two. For a complex object such as an automobile, the internal structure has vast intricacy. There are hierarchies of systems and subsystems containing structures and substructures, and detailed arrays of interconnections between structures. Notice that this notion of nested structures, where we have objects within objects, allows us to refer to components not only in terms of their internal structures (if we wish) but also in terms of the functions that they provide. In fact, we might often wish to describe the components only in terms of their function when we are discussing the system that they are components in. In other words, when describing an entity in terms of its component parts, we probably will tend to ignore the internals of those component parts and treat them as wholes. If we were to shift our focus to a particular component, then we might "open it up and see what's inside." Thus, depending on the context, we can sometimes view entities, which might be parts of some larger whole, as indivisible "black boxes" whose internals are invisible. At other times we can focus in on them and look into their internals.

Surface properties are those which concern an entity as a whole. This is exactly what we meant when we used the term "black box." We often treat entities as black boxes, not knowing or not caring about their internal details. We can distinguish the internal-properties perspective from the surface-properties perspective in the following way. Internal properties relate to *how* something is structured or *how* it does what it does, whereas surface properties relate to *what* something is (as a whole) or *what* it does. An entity's class, the kind of thing that it is, can be considered a surface property. To take our car example again, we have seen that a car has much internal structure, yet the average driver need not know much if anything about what that structure is in order to identify it as a car or use it for its intended purpose (while the automobile takes on a very different appearance to the auto mechanic). So we can refer to an

¹ Langacker (1988) considers the "content" of *go* and *gone* to be the same, but considers the relative emphasis ("profiling") of the parts in the pictures describing the two words to make the difference in their meanings. It is arguable that the relative emphasis is also a *part* of the entity, and thus is part of the content.

automobile's function, transportation, as a surface property. There are other surface properties of cars that we can think of. For example, we might label a car's color as an exemplary surface property. Less obvious but equally applicable would be such properties as sticker price, make and model, and year. What about such considerations as the material that a car is made out of, the number of passengers able to comfortably fit in the car, whether the sound system has a CD player, or the style of the wheel covers? What criteria allow us to reliably identify surface properties and distinguish surface from internal properties?

The type of material(s) that make up an entity is an internal property, so what a car is made from would be an internal property, although technically this would apply in a loose sense. These materials would more likely be what specific *components* of the car are made of. Topologically, the seating arrangement in a car is a surface feature (if the windows are open!). This may sound somewhat strange, but this point is driven home by recent "modular" vehicle designs where the seating can actually be continually rearranged by the owner. Is a minivan a car? For our purposes, yes. We would draw the line at RV's, since their acknowledged function is not just transportation, but overnight shelter as well. Whether a car has a CD player, or even a sound system at all, is superficial to its purpose as a means of transportation. Similarly, although having wheel covers technically affects the "parts list" of a car, as does having sound system components, certainly what the wheel covers look like does not constitute part of the car's internal structure. And please note, we are not strictly defining "internal structure" to be physically (and topologically) internal -- rather, we are considering parts or connections as internal if they significantly affect the car's intended functioning. So, the ignition system or emission system might be considered internal subsystems. Of course it *could* be argued that the emission control subsystem is peripheral and incidental, that the car would still be a car without it. In fact, following this line of reasoning, many of the car's subsystems could either be eliminated as unessential or replaced by functionally equivalent variants. For example, we could replace an automatic transmission with a manual one, or a piston engine with a rotary engine, etc.

Perhaps then we can distinguish between essential functional components and inessential ones. Within certain parameters a car *must* have some kind of engine. Whether it has an emission control subsystem is optional (definitionally rather than legally speaking). These are exactly the distinctions that Wierzbicka (1985) makes when referring to properties that people would say that an entity *would* have as opposed to those which they would say that the entity *could* have. Suppose we call these essential versus optional properties. Are optional properties always surface properties? Their presence or absence does not substantially affect what the entity is or does.² Yet we can easily imagine inserting some subsystem, say a driver's-side airbag, into the internal structure of the car which nevertheless makes no crucial difference as to whether overall the car still functions the same. In fact, some optional properties might neither be internal nor surface, as we shall see shortly in our discussion of external properties. Are essential properties necessarily internal properties? No. Although removing any of these properties changes the nature of the entity,³ these properties too can fall within the surface- or even

² Note that we are not talking about the property's particular instantiation (e.g. what type of emission control system is used) being optional but rather about the property itself being optional (that there is an emission control system at all).

³ But aren't these then necessary conditions, as in the truth conditional, formal semantics approach (e.g. Bierwisch 1970)? See the next paragraph in the text.

external-property sections of a definition. What do we mean by "within certain parameters"? What do we mean by "substantially"? And, have we satisfactorily determined what distinguishes internal from surface properties yet?

If we allow that there can be merely approximate conformity to a definition, as in the prototype approach (e.g., Fillmore 1975), then although an entity only has most of the essential properties of the concept being considered, the entity will be considered as an instance of that concept. Different entities might each have slightly different almost-complete subsets of the essential properties of the class, and so would form a cluster of closely related variants. The essential properties are not exactly necessary conditions, but rather highly likely ones. There is a sense in which a certain "critical mass" of properties must be met in order for an entity to qualify as exemplifying a concept. In other words, as long as "enough" essential properties are present, then the entity is an instance of the concept. This is what we meant by "within certain parameters" in the context of a car having an engine. There is some latitude allowed for variation in engine type, but when enough essential properties of "car engine" are missing, then we would most likely hesitate to categorize some mechanism of power production as a candidate car engine.

Determining whether something "substantially" affects what an entity is or does follows the same guidelines. Sometimes this is subjective. For some car owners the function of their car is not just to provide transportation, but to provide a place to listen to music in comfort as well. In that case, whether there is a sound system, and even whether there is a CD player may well affect what that car is and does for a given owner. Wierzbicka tries to dodge this issue somewhat in her discussion of what a cup is, claiming authoritatively that cups have such and such characteristics (e.g. that they come with saucers and in sets). This, however, is *her* conception (and may not conform to the conception of many other people). She does acknowledge in a footnote that such canonical definitions are known to vary significantly for "younger" people. After all, who is to say what the definition of an entity is?

What are we to do then? Should we throw up our hands in despair because it is ridiculous (and intractable) to determine for each individual what the definition of each entity is? No. Perhaps we can limit ourselves to the level of granularity of the subculture rather than the individual. Yet, we should acknowledge that the "meaning" of some entity really does vary from person to person, even if ever so slightly, based on that person's experiences and the associations that the entity has for a person. In any event, *if* a property of an entity can be essential for some individuals, we should allow for this. Things are relative. What is essential for some people might only be optional for others. But, this does not obviate the distinction between essential and optional, nor that between internal and surface. What is essential or optional cuts across the three levels of definition being used here. Perhaps we can say that a property is internal if it physically or logically forms part of the entity's internal structure or can affect the entity's behavior or identity, and that it is a surface property if it physically or logically lies on or applies to the surface of the entity in isolation or summarizes the behavior or identity of the entity in isolation.

If we consider the internal properties and surface properties of an entity as together making up what is *intrinsic* to it, then we can

see that this corresponds to looking at the entity in isolation. Participation in (i.e., being a part of) any larger structures can then be seen as participation in *extrinsic* relationships, which can be considered as *external properties*. Note that for definition purposes these instances of participation are abstract -- we are talking here about the *kinds* of things that the entity forms a part of, not the actual instances of those things. For example, if we are discussing power transformers and we want to list some of their external properties such as the kinds of things that they form part of, we might mention stereo receivers. Yet we are not referring to any specific stereo receivers. Similarly, although for some purposes we might want to enumerate instances of the entity itself (given that it can have instances, and is therefore some kind of thing), each instance in the entity's *extension* is not an external property. Rather, the extension as a whole could be considered one of the entity's external properties. And, the entire three-level definition corresponds to the *intension*.

If, for Wierzbicka, cups are generally associated with saucers and come in sets, then these essential properties can be seen to be external properties. A cup in itself has internal and surface properties such as the handle and the decoration respectively, but also forms part of larger structures, for example "a cup and saucer" or "a tea set." Are we arguing that perhaps such properties should not really be considered essential, and therefore should not be part of the definition of the entity, since they are in some sense not part of the entity? Far from it. Rather, we are encouraged by efforts such as Wierzbicka's to take into account many of an entity's attributes which might otherwise be overlooked as directly relevant to knowing about that entity. All too often we look at things in isolation. It is not only healthy but essential to enhance our understanding of entities by seeing them as they fit into larger structures. Doing so does not incur the exorbitant overhead of trying to list all of an entity's instances, but rather only the affordable cost of listing its *kinds* of participation.

So, let us summarize what we have so far. Internal properties tell us how something is internally structured or how it does what it does. Surface properties tell us about the entity as a whole, what it is or what it does. External properties tell us about the structures that an entity forms a part of, the kinds of relationships that it participates in. Let us illustrate this three-level definition technique with two further examples. Atoms have internal structure (for example protons, neutrons, and electrons in specific structures and configurations), have surface properties such as atomic weight and diameter and electronic valence, and have external properties in the typical molecules that they participate in. A computer program (especially of the modern object-oriented variety) has internal details which determine how it does what it does, it has its function as a module (what it does), and it often fits into larger systems. But these examples of cars, atoms, and programs don't cover the full spectrum of the kinds of things in the world, do they? Are there other types of entities for which this three-level definition approach applies, and indeed, are there types of entities for which this approach does *not* apply?

Whether entities are physical (concrete) or logical (conceptual, abstract), they generally have some kind of internal structure. As we mentioned earlier, this goes for processes as well as for (static) objects. Being entities, they also must in some way be delimitable, otherwise how would they come to have been named or conceived of? That is not to say that they must have hard, sharp-edged, crisp boundaries. Entities can have fuzzy boundaries and still be more or less identifiable. Is it possible that some entities have no internal details or external properties? Must all things have their own structure and in turn form part of some larger structure (either physically or

conceptually)? It would seem that the answer is "Yes" to both questions even if the participation in larger structures is only optional and potential. Even in the extreme cases of the most miniscule (what the atom once was, the indivisible) and the most all-encompassing (the universe), we have time and again been surprised to find that these alleged limiting cases were themselves respectively decomposable and only part of the story.

What about natural kinds versus artifacts? What about individual physical objects versus abstract concepts? Clearly kinds of things in nature have internal structure, but do they have functions? Even avoiding teleological debates about the meaning of life or the purpose of some organism's existence (although individual organs have functions), we don't need to rely on function for filling in our list of surface properties (or, for those who decry the checklist approach, our structured set of surface properties). Function is only one type of surface property, and may itself be optional. Function is an essential property for artifacts but not necessarily for other entities. In other words, applying our three-level approach to itself, perhaps we can identify certain properties of a definition as essential and others as optional, depending on the type of entity concerned. So, for some entities, function may not be applicable, thus making it an optional surface property. It is an open question whether there are any globally applicable properties at any of the three levels, except perhaps that all entities have participation in the universe (the set of all entities) as an external property.

Artifacts are by nature conceived and constructed for some purpose. This applies not only to tangible, familiar entities such as tea cups but also to arcane, abstract entities such as tensors. Again, we are not limiting our treatment to static objects, but incorporate processes as well. So, we can just as easily have used tea ceremonies and gauge theory renormalization as our examples of concrete versus abstract artifacts. But are all abstract concepts artifacts? It can be argued that many are not. If we perceive patterns in the world, whether in the "physical" world or in the "mental" world, and we generalize over the instances' common denominators (performing logical induction), then we have not "made up" the concepts but only perceived them. These concepts are therefore not artifacts. If, on the other hand we have devised labels for these concepts and have conventions between us for structuring, externalizing, and transmitting these labels (in speech and writing, for example), then this is artifice. Thus, artifacts have internal properties, at least have some functional surface property and probably have additional surface properties, and may or may not be perceived to form parts of larger structures.

As far as individual physical objects (or processes), the same conclusion applies. They have internal properties, surface properties, and very possibly acknowledged external properties. A tree in isolation can be characterized by its internal structure and its outer appearance and behavior. But, we can also look at the tree in its context and see that it might form relationships with certain animals, providing food and shelter, and with the atmosphere, processing carbon dioxide into oxygen.⁴ Similarly, people have internal structure, surface aspects, and external relationships. What about inanimate objects? We have already discussed atoms and cups. We can apply this three-level breakdown to buildings, clouds, cities, oceans, and anything else one cares to name. Let us emphasize that we

⁴ We are indebted to the 1991 movie *Mindwalk* for this example of interconnection.

are not saying that we must exhaustively exhume and enumerate all actual and potential properties of an entity and fit them into the three partitions. Rather, we are simply saying that the three compartments exist and are there for us to use if we wish. This process of description can be left open-ended. We may build up a significant definition for an entity, with entries in any or all of the three departments, yet know that this task may be unfinished. The definition is significant when it is definite enough to allow us to recognize instances of a concept or class and to distinguish them from those of other categories, yet it is flexible enough to provide some leeway, as in the prototype approach. Furthermore, we do not reject outright the possibility of using necessary and sufficient conditions to define an entity. In some circumscribed contexts, where we have consciously devised and defined some entity, this is perfectly appropriate. For example, we can unequivocally define the concept "triangle" in terms of necessary and sufficient conditions.

How does the tension between dictionary and encyclopedic orientations become resolved here? Although this tension may become moot with ongoing trends in information access technology leading to ad hoc full-spectrum information services, which we will discuss later, this tension may already be largely moot because the two orientations are largely incommensurate. Dictionaries give knowledge about words while encyclopedias give information about topics, specifically "people, places, and things." Dictionaries are specifically constrained to *not* include people and places (proper nouns). Words (and the concepts that some of them represent) can be topics themselves, but these are only a subset of the topics about things in an encyclopedia. The "names" for many things consist of more than one word. An encyclopedia also goes into considerably more depth concerning facts and associations about topics, while dictionary entries often include linguistic information not found in encyclopedia entries. So what sense are we to make of this tangle -- what is the tension and how can the three-level approach to definition help resolve it?

Traditionally, dictionary definitions have primarily aimed at digging out and presenting some of a word's concept's salient internal or surface properties, possibly roughly its essential properties within the internal and surface domains. Wierzbicka and other modern semanticists, besides stressing the functionality of artifacts as crucial to defining them, and the interconnection of their features, have attempted to bring to the surface the importance of other salient associations for a word's concept, some of which we would classify as surface properties and others as external properties. Thus, the gate has been opened for a flood of encyclopedic facts and descriptions to be included in a definition. After all, most concepts do not exist in isolation and words' senses overlap and form hierarchies and networks. The resulting "definitions" might be considered unwieldy, however. If we can accept that a dictionary entry should be brief, whereas an encyclopedia entry can be considerably longer, then we can limit dictionary entries to only the most salient internal, surface, and external properties. The encyclopedia would be the place for more extended descriptions, which could include many of the associations and facts surrounding a topic. The "definition" and "meaning" of a concept would embrace *both* approaches.

Some of the kinds of things that we can say about words do not apply to topics in general and conversely. Some words might not have concepts associated with them, but be merely syntactic and sub-semantic (e.g. speech act constructs such as "Well, ..."). There will be linguistic information associated with these words, but other kinds of content would not appear in their dictionary entries.

Conversely, some topics might have a number of non-linguistic cultural associations attached to them which would only appear in an encyclopedia entry. There is room for all of these kinds of information in a three-level definition. Depending on the type of entity under discussion and the type of discussion, various properties apply and various others do not. If we wish to limit a dictionary entry for a topic to a much smaller size than a corresponding encyclopedia entry, then besides predetermining the types of properties to be included, perhaps we can set a much smaller threshold to the "neighborhood" size (over the shared properties) for the dictionary definition. That is, if the word's concept is enmeshed in a web of meanings, we might define a neighborhood for the concept as the nearest meanings in the network. How close does a meaning have to be to be "nearby"? We can set that "distance" as the threshold. For encyclopedia versions of our concept description we can extend the boundaries out further, so that we pick up additional meanings and associations. What does this have to do with the three-level approach?

It is time to apply the three-level technique to itself. Most of the time we have been defining the three-level approach in terms of its internal structure -- the distinctions between internal, surface, and external properties. We have also occasionally touched on one of its surface properties -- that a definition can include associational and even subjective information in addition to the more intrinsic information concerning the item in isolation often found in lexicons (due to space restrictions). As we have said, the "meaning" of something can extend to one's experiences and associations with that entity. So, this surface property of the three-level approach itself, what we might call part of the spirit of the approach, allows us to view dictionary and encyclopedic information as close relatives which complement each other. They are special cases, just two of many subsets of points in a space of information about a topic. What might an external property of the three-level approach be, one might ask?

The three-level approach could be seen to fit into an overall orientation to knowledge processing. If we envision an information environment where we can query an online system for information on some topic, then depending on the nature of the topic, the types of information that apply would vary and the types of responses that we could get would vary. If our topic were a word, then we might expect to be able to get such information as its meaning, its pronunciation, its etymology, and so on. Perhaps we could also get examples of its usage. If the topic were a country, then perhaps we might get a map of the country displayed, its population and area figures, its history, its form of government, etc. But these are just our traditional dictionary and encyclopedia entries. The envisioned information system might be able to provide other services as well. Some of these would be relatively familiar, such as thesauri and technical glossaries, but others might be without precedent. For example, suppose that we wanted to know if there were a word which stood for a phrase that we had in mind. The system might be able to tell us that. Or suppose that we wanted to see a map of a "conceptual neighborhood." The possibilities would actually be quite open-ended. But how would such processing take place "behind the scenes"?

The system user would only be seeing the surface. Somehow, though, the information connections would need to be made. One current approach, that taken in the field of "information retrieval," is primarily statistical. Topics are "understood" as the distribution of terms occurring in documents about those topics. The law of averages allows us to conclude that this approach should more or less

succeed. But this approach does not work with the "content" of a topic, its semantics. The statistical approach works fairly well and is currently more efficient than existing semantic approaches. But there are times when the method breaks down, for example when the terms used in classifying documents are not those that the user supplies, even though the user's terms are appropriate. And what if we wanted to retrieve items smaller than whole documents. What if we wanted to retrieve paragraphs or even individual facts that satisfy the user's query. If we were able somehow to capture the semantics of both the query and the information in the "topic base," then we could compare them somehow for similarity.

But how do we measure such semantic similarity or relevance, or simply, semantic distance? Previous approaches (Rada et. al. 1989) have focused on taxonomies or "is-a hierarchies". These approaches are brittle. They do not take into account multiple types of relationships, for example cause-effect relationships combined with is-a relationships. Perhaps a more comprehensive semantic distance methodology can be devised that incorporates the three-level definition technique. After all, connections between entities can take many forms. They can have similar structures (sometimes giving rise to analogies and metaphors), they can have similar surface properties (such as function, color, age, location), and they can belong to similar structures (for example, two people belonging to the same family, organization, or discipline). Thus a comprehensive semantic distance measure might take all of these distinctions into account, perhaps weighting some factors more heavily than others. Also, the notion of defining a semantic neighborhood via a threshold, beyond which items are considered irrelevant, could use the distance measure or could itself be defined from first principles, using the three-level method. Perhaps notions such as "cue validity" could be useful in defining neighborhoods (Rosch 1978).⁵

For information retrieval, relevance is the key. Relevance is not only determined by similarity, however. Any kind of connection, not just one of family resemblance, will be potentially relevant. Thus in one approach we might weight internal and functional properties most heavily, other surface properties less heavily, and external associations most lightly. Internal structure might be seen to be "inherited" in large part from the class that an entity belongs to (though the name of the class would be considered a surface property). Such internal properties might be relatively fixed in value for instances of the class. For example, our car will have only a few possible engine types or transmission types. Surface properties might be allowed more variation in their instantiation. For example, our car will have a color, but the particular color may vary widely. The binding of internal properties might be quite tight, that of surface properties only moderately so. With external properties the number of possibilities, although not endless, being constrained by the type of entity at hand (i.e. the internal

⁵ It is tempting to try to tie in Rosch's ideas about basic level categories with the three-level definition scheme for two reasons. First, we can map the default level of focus, the surface level, to Rosch's basic level. And we can map the internal and external levels to Rosch's subordinate and superordinate levels, respectively. Second, what constitutes a basic level is relative to expertise, the focus descending the class hierarchy with increasing expertise. Similarly, experts are more familiar with internal structure than laypersons and even tend to lose sight of the higher levels (e.g. systems programmers who have lost touch with the "end-user" perspective). Here again, focus has shifted down the hierarchy. Can Rosch's distinctions be seen as a special case of the three-level perspective, where we are limiting ourselves to categories in a taxonomy? For the mapping (from Rosch to here) to be coherent, we would somehow have to reconcile the subtypes under a basic category with what we have been calling internal properties. It does not seem too unintuitive to think of subtypes of a category as what the category "breaks down into." The same can be said for mapping superordinate categories to larger structures that an entity forms part of.

and surface properties), might nevertheless be wide open. It will require much more investigation, both through empirical studies and through analysis, to determine whether these initial intuitions are on target. Is this three-level approach any better or worse than the essential vs. optional properties approach, for example, at finding relevance? Again, more work needs to be done.

Can we accommodate form/content and structure/function? Yes. Form is internal structure (but can also refer to the appearance of an entity and thus its surface) and content consists of the participants in the structure. In this sense, "profiling" is separate from content. Structure and function clearly correspond to internal structure vs. the surface property called function. That traditionally there is this distinction between structure and function is evidence supporting separating them into two distinct levels within the three-level scheme.

In summary, we have seen the three-level approach to definition explained, and seen examples of its application. We have touched on such issues as prototypes and checklists, essential and optional properties, dictionary and encyclopedic descriptions, and the scope of the technique's applicability. We have also delved briefly into the realm of eventual full-service information systems and the possibility that some measure of semantic distance will be useful there. Finally, directions for future work have been identified.

References

Manfred Bierwisch. *Semantics, New Horizons in Linguistics*, ed. John Lyons, Penguin, Harmondsworth, Middlesex, England, 1970.

Charles J. Fillmore. *On the Organization of Semantic Information in the Lexicon*, excerpted in "Papers from the Parasession on the Lexicon," ed. Farkas, Jacobsen, and Todrys, Chicago Linguistic Society, U. of Chicago, 1978.

Ronald W. Langacker. *A View of Linguistic Semantics, Topics in Cognitive Linguistics*, ed. Brygida Rudzka-Ostyn, John Benjamins, 1988.

Roy Rada, Hafedh Mili, Ellen Bicknell, and Maria Blettner. Development and Application of a Metric on Semantic Nets, *IEEE Trans. on Sys., Man, and Cyb.* 19:1, Jan./Feb. 1989.

Eleanor Rosch. *Principles of Categorization, Cognition and Categorization*, ed. Rosch and Lloyd, Lawrence Erlbaum, 1978.

Anna Wierzbicka. *Lexicography and Conceptual Analysis*, Karoma Publishers, Inc., Ann Arbor, 1985.