

The capacity to keep track of specificity: the genuine vs. the linguistic

Tamás Pólya – László Tarnay

polya@ektf.hu; tarnay@btk.pte.hu

Department of Communication; Department of Philosophy
Eszterházy Károly College; University of Pécs, Hungary

Abstract for the workshop "Direct Reference and Specificity
at ESSLLI 2003 in Vienna

'Specificity' is an ambiguous term. Roughly, for linguists, it refers to uniqueness in a domain of interpretation; for some other people in cognitive science – like we are –, it refers to uniqueness in the world. At first glance, it is by no means clear how the two uses relate to each other: whether they clash or differ, and if they differ, whether they can be made compatible. We are inclined to embrace the second option, emphasizing the incompatibility of the two uses – and in what follows we try to give a proper argument in favor of it.

1. Genuine specificity

Some one and a half year ago we started a research into the status of mental representations (as traditionally conceived, à la Fodor 1987) and the role they play in social cognition, namely whether mental representations are sufficient for an organism like man to cope with social situations. We ended up with four distinct criteria that constitute a disjunctively exhaustive set for introducing representations into one's theoretical talk (see Tarnay and Pólya, forthcoming). However, as our work progressed, we came across a host of cases of processing information from animal cognition to human social communication that did not sit well with how we originally conceptualized the conditions for representations. For while representations are commonly thought of as a processing aid \bar{n} mental, conceptual, linguistic, or what not \bar{n} to *generalize* over previous, or to *project* onto new, instances of stimulus information \bar{n} i.e., they are of particular *types* \bar{n} , the newly found data comprised cases when certain species like rodents, birds, whales, primates and, last but not least, humans, distinguish among family or group members. That is, they make *within-category discriminations* by detecting special physico-chemical substances (e.g. tastes, secretions, odors, sounds) or morphological parts (e.g. head-and-neck, faces, gait) released by or appertaining to particular conspecific members. Since such substances or '*proper parts*' as we call them are causally linked to a given organism and its physical existence, they are indicative of its singular and unique character, i.e., they are *indexes* of individuals. We claim that such individual specificity is inherent in the variability of the environment, or in what we call *environment specificity*. We then define specificity recognition as a case when an organism is endowed with a sufficiently fine-grained sensorial and processing system \bar{n} i.e. *organ specificity* \bar{n} to pick up idiosyncratic information. The latter is the lower \bar{n} sensorial \bar{n} bound of recognition of physical uniqueness. Thus, if an organism is sensitive to such single substances, it has the capacity to individuate, or to *keep track* of, the specific members the said substances belong to. A capacity like this is, we submit, sharply different from a capacity, say, to tell an edible fruit from a non edible one, a territorial intruder from a group member, a predator from a prey, etc. The latter are all cases of categorical perception or discrimination: they imply distinctions *among categories* like edible/non edible, intruder/home mate, predator/prey, etc. Thus we are going to call the former capacity, the one to individuate individual members (of a given category) as *recognition of specificity* and contrast it with *categorical discrimination*.

It seems that the former cannot be reduced to any of the possible modes of the latter. To be clear, we do not want to deny that in many cases (possibly unique) category membership is sufficient to identify an individual, provided that certain epistemic and pragmatic conditions obtain. But identifying is not recognizing individuality. For instance, having long, brownish curly hair may be identifying in a context, but it is not sufficient to constitute an individual *across contexts*, or *a-contextually*, simply because there could well be other individuals exemplifying the very same property. That it constitutes a unit set (in some or more contexts) is purely contingent fact. And it remains so irrespective of how complex a property we would wish to choose.

2. Language and specificity

Approaching language with the assumption that it is tailored to convey genuine specificity is at least controversial. But if we accept that categorical thinking predates language, and as a consequence, language can only contribute to further refining existing categories or creating new, more precise ones, the case for a specificity-specific language or linguistic ability could seem to be made.

Traditionally speaking, in theory the specificity/non-specificity contrast can be treated at either a syntactic, or a semantic, or a pragmatic level. Note that with the rise of psycholinguistics a separate cognitive level has been added. Historically, the problem appears as the referential capacity of language: reference (denotation, designation, etc.) is distinguished from meaning (sense, intension, conceptual content, etc). It seems to be most proper to keep the two corresponding problems distinct. For the first demands an inquiry into the structure of specificity itself; it implies an

ontological bias. Whereas the second implies a theoretical investigation of the level at which specificity could be most adequately treated; it asks about the expressive power of language.

Clearly, we cannot offer an exhaustive analysis of even the most important philosophical and linguistic theories covering these phenomena here. Rather, we would like to concentrate on the impact dynamic semantics in general, and Hans Kamp's Discourse Representation Theory in particular have made on theoretical thinking about specificity and co-referential relations, the suitable account of which seems to lie at the heart of our dilemma. But first we have to explain what is meant by the kind of specificity ubiquitously applied in linguistic theories, domain specificity.

2.1 Domain specificity

As a representative example, consider Groenendijk and Stokhof's (1981) suggestions about linguistic individuation. They claim that specificity is a proper domain of pragmatics, that is, language use, so that a term is used specifically only if it refers to an individual the speaker's utterance (in a given situation and according to him/her) is about. In other words, specificity depends on what the speaker knows about the proper denotations of the terms used. Groenendijk and Stokhof take pains to stress that specificity in their sense is objective and thus differs from the subjective condition of having a particular individual in mind. The latter is independent from using terms to denote particular individuals. If one knows one's language, the denotations of its terms, he can use it intersubjectively, even when the different speakers of the language have different beliefs about the world (while they share the knowledge about the use of their language).

Now, is it true that this pragmatic analysis of specificity captures the idea of genuine specificity? If we give a closer look to the technical apparatus of the authors, it appears that the idea of knowing the denotation of a term falls back on a prior knowledge concerning how the world (the relevant situation) is. How else can we understand their requirement that one cannot use the sentence

(1) A friend called me up tonight.

specifically if there were two friends who phoned me that night? Specificity in this sense boils down to domain specificity, or to knowledge of the proper extensions of the terms of the language one uses. This may point to another important aspect of linguistic specificity, namely *uniqueness with respect to a given universe (or domain) of discourse*. In linguistic theories it is often assumed that once a variable (such as the one corresponding to 'a friend' in (1) is bound, that is, assigned a value, no misunderstanding or mistaken reference with regard to it can take place. That means the reference of the variable can with no extra effort be re-identified by the hearer. But the problem crops up, of course, at the onset of interpretation and value assignment when there is more than one value in the domain of discourse to be found that the variable may assume. It is exactly in such cases that (1) cannot be used specifically (e.g. when at least two friends called the speaker the given night). Thus (1) cannot be specific on its own.

The uniqueness claim which figures almost in every semantically or epistemically biased linguistic theory is undoubtedly an important corollary of specificity. But it may not be enough in many cases. For consider for a moment, as an analogy, visual indexical individuation. Imagine a series of snapshots of a movie depicting a story about a princess who comes to be enchanted by some evil power. Keeping track of an individual through and across domains (i.e. different snapshots) is independent of how many members a given domain has (viz. domain specificity). If I want to represent the tale about the princess who enjoys her life in one shot and who has become a frog in another, it requires that I can keep track of her through the shots, but it does not require that she should be the only available princess in the first shot (domain) or the only potential frog in the other.

2.2 File change semantics and Discourse Representation Theory

That the problem of cross-identification cannot be resolved by relying on domain specificity can be illustrated with file change semantics. (Cf. Heim 1982) Some of the theory's proponents (e.g. Corazza 1995) make a distinction between *external vs. internal* identity relations between files. They ground their distinction on the idea that indexicals cannot establish internal coreference between files which are separate clusters of information (about discourse referents) the speaker heaps up gradually as he/she proceeds in understanding linguistic discourse. In our terminology, this means that files cannot serve for keeping track of the same object; they only do so through perception (until the object is in sight); this is why Perry introduces a separate category called 'buffers' as distinct from files which store information. Thus, it follows that by conceptualizing the perceptual content one does not necessarily keep track of the object indexically referred to; one may well lose sight of it and establish two different files (or buffers) even while subsuming them (the different appearances of the object) under the same category.¹ Since files are packages of information, they need not coincide with categorical or conceptual knowledge. This allows for obtaining different relations among files. The strongest one is called coreference or internal identity when the content of a given file or buffer is incorporated into another one. But there are weaker relations, like transference, when someone imagines to be someone else without thereby identifying himself/herself with him/her. Another famous example is the identity between Tully and Cicero, which may or may not obtain internally, viz. between files, while it does obtain in the external world.

The latter example clearly shows that files change semantics is not so much concerned with how internally obtained identity relations or coreferences correlated with external identities, but rather with descriptive or attributive cases à la Donnellan (1966). It can thus explain how it is possible that one identifies *X* as a spy even though he/she does

¹ There is no guarantee, however, that pace Corazza proper names are essentially different (being labels of files).

not know who *X* is (viz. Orcutt). What seems to be more problematic is the mirror-image of “splitting” concepts or categories. For when it comes to the question how one can track down the same individual in different files and buffers, the only thing the theory can say is that they get incorporated into each other. It simply presupposes a capacity to create coreferential links between stored information, and thereby it begs the question of how genuine specificity is possible. Since such links cannot be inherited from indexical (external) relations, there is an almost infinite source for creating internal links. This may explain imaginary cases, but it might also become overly prolific.

Yet, the problem is not so much with the number of possible identity relations between files, but rather with how the mechanism that co-indexes certain files work. When, for instance, we imagine that a princess is turned into a frog, do we incorporate all information within the princess-file into the frog-file? If not, can we go on saying that one file is incorporated into another, provided that files are packages of information indexed by proper names (or maybe by some other means)? It seems that if specificity is tantamount to clustering clusters of information (and so on *ad infinitum* perhaps), what we lose sight of is precisely the specificity of the individual (singular) object.

The proper grasp of genuine specificity seems to be beyond the capabilities of another, perhaps the most prominent, theory in the dynamic branch of semantics, the Discourse Representation Theory (DRT), devised by Hans Kamp. (Kamp & Reyle 1993) The problem itself is faced, as common in linguistic theories, when the analysis of de re sentences is attempted. However, DRT assumes that sentences can not only be regarded as abstract entities of language but rather as instruments of communication being actually used and interpreted by language users. Consequently, it maintains that a semantic analysis of sentences is greatly supported by importing theoretical constructs of a quasi-psychological nature, i.e. ones that model the knowledge state of the speaker and/or the hearer relevant to the interpretation of sentences. Thus, consider that, for example, a de re belief is a ternary relation between the believer, the property believed by the believer and the thing to which the believer attributes this property. Say, Oscar has a de re belief about one of his friends when he believes that Mary kissed his friend, and has a particular friend in mind. There, Oscar attributes the property of ‘having been kissed by Mary’ to the friend involved. (Clearly, a sentence describing Oscar’s belief state of the same content may easily be put in words, too.) So Bende-Farkas & Kamp (2001) note:

[It is important] that many of our thoughts are *about* other people or things – thoughts, in other words, through which we are related to *them*. And that, moreover, there are linguistic forms which seem to be describing just such entity-related thoughts. [...] We will assume that the kind of »relational thought« is possible only when the subject has an individual-representation for any individual with respect to which his thought is *de re* (or, in other words, the individual to which his thought relates him). And we will [...] take such individual representations to have the form of so-called anchors. (19)

Anchors (the notion is to be clarified in a moment) are argued to be especially useful when a person does not possess a single but rather a set of attitudes towards the very same entity.. Say, Oscar may doubt that George has kissed Mary but also hope that he will and that Mary hopes so, too – here one plus two anchors are needed to represent Mary in Oscar’s second and third attitude, respectively. In cases like this, Bende-Farkas and Kamp say that diverse attitudinal states share a discourse referent.

For Bende-Farkas and Kamp (see 2001: 17ff), anchoring is also an attitude of the cognizer, but not a propositional one. It is represented in their Discourse Representation Structures (DRSs) as a composite of a mode indicator – whose only job is to introduce the anchoring relation into the representational complex – and the specification of the information content through which the relevant entity is represented by the subject. Subsequent anchors will borrow their information content from that first specification, sort of referring back the cognizer to the content given at the onset. Such variables internal to the DRS are called *internal anchors*. Obviously they are construed as co-indexing elements, securing a referential content to be available in the context of the relevant attitudes only with respect to an original specification of content. I.e., they cannot reach outside the domain defined by the given Discourse Representation Structure, and hook onto someone or something in the external world. For that is what *external anchors* are for. Whenever a subject has an object in thought which a real-world entity does correspond to, an external, binary, statement will be added to the DRS at hand, so that the (set of) internal anchor(s) variable be bound to that entity. That is, according to Bende-Farkas & Kamp (2001: 20ff)

[t]he external anchor $\langle x, x \rangle$ confirms that the subject’s internal representation x is anchored to (and thus represents) the object x . (We also call x the *external anchor* for x in this case.) Note that the external anchor is not part of the representation of the mental state, but an external judgement, of which the subject himself is not capable. Internal anchors presuppose an external anchor: When there is an external anchor for an internally anchored discourse referent, then it is the thing for which the discourse referent stands. If there is no external for the internally anchored x , then the internal anchor is deficient, x suffers from a reference failure and all the representations in which x occurs fail to determine well-defined propositional content. (21)

Note that external anchors are no good for the subject in determining whether there is an external object he thinks is referring to, let alone in determining the specific referent of a mental or linguistic expression – they only establish, *unbeknownst to the subject*, a sort of ontological warranty for there occurring no vacuous references in the subject’s thoughts concerned (and utterances based on those thoughts). But *genuine specificity of an entity*, we maintain, *must be in some form available to the mind of the cognizer* dealing with it. That is decidedly not the case here.

To put it sharply, what the theory accomplishes is to secure domain specificity, by virtue of internal anchors (no matter how complex domains the theory can cope with!), and to introduce some placeholder notion, viz. that of external anchors, as the one that should carry the burden of grasping what we call genuine specificity. However, it is unclear how *genuine specificity* could be preserved *mentally* once the cognizer starts considering *internal* anchors during linguistic understanding, *unless* by his/her reliance on the very same mechanisms that secured the grasp of genuine specificity at the level of *external* anchors in the first place. But to explain how those mechanisms work, we submit, demands a much more detailed scrutiny of how lower and higher order cognition can cope with the task – and that is where the notion of *proper part* and related notions, as explained in Section 1, may come in very useful (see also Tarnay and Pólya, forthcoming, for a book-long discussion). One may consider, as DRT theorists seem to do, the cognitive processes at work as responsible for, say, ‘external anchoring’, and label them accordingly, but that does not make those processes any more linguistic. On the contrary, proper parts are grasped via low-level cognition such as olfaction, facial and gait perception and so on, rather than by means of (sets of) linguistic predicates.

3. Conclusions

Although some may wish that our linguistic mechanisms be able to grasp genuine specificity, we have been arguing against such an impression. Categoriality involves domain specificity and genuine specificity simply cannot be equated with domain specificity. If you grasp something through categories (concepts or predicates), you will only end up with a vague, if not necessarily undetailed, draft or outline. What would you envisage, for example, if I tell you about my father’s favorite ‘black hat with a gray band’? You may picture it as an old-fashioned black hat with a gray band – and what? At best you also envisage it as wrinkled or smooth, shabby or shiny, stained or clean, and so on. Of course, you can come to know a couple of things about it. You may be able to judge if my father is wearing his favorite hat or not, find it in his closet, compare it to other hats, etc. But except for the case when, by sheer luck, you imagine the old piece of cloth exactly the way it is (a very rare case indeed), there will be a whole world of difference between the actual hat and what you will have grasped.² For words and concepts are aimed at capturing the world in a sufficiently abstract and concise manner so that generality – viz. commonality with other objects – is almost always preferred to hooking onto things in their specificity – viz. physical uniqueness – or, as in terms of the example above, one envisaged sample of ‘black hat’ would do just as well as another. But that is not how genuine, physical specificity works. For the fine-grained details of the latter are beyond, or more precisely, below, the rather rough categorial reach.

References:

- Corazza, E. (1995): Je est un autre. *Archives de Philosophie* 58, 199-212.
Donnellan, K. (1966): Reference and Identifying Descriptions. *Philosophical Review* 75, 281-304.
Fodor, J. A. (1987): *Psychosemantics: The Problem of Meaning in the Philosophy of the Mind*. Cambridge, Mass.: MIT.
Gronendijk, J. & Stokhof, M. (1981): A Pragmatic Analysis of Specificity. In: Heny (1981): 152-191.
Heny, F., ed. (1981): *Ambiguities in Intensional Contexts*. Dordrecht: Reidel.
Heim, I. (1982): *The Semantics of Definite and Indefinite Noun Phrases*. Amherst: University of Massachusetts.
Tarnay, L. and Pólya, T. (forthcoming): *Recognizing Specificity and Social Cognition. A Multi-faceted Approach to a Unitary Phenomenon..* Frankfurt am Main: Peter Lang Verlag

² It is important to note here that our argument does not turn on some realist or correspondence theory of truth. That is, it does not matter *how far* the envisaged mental ‘hat-object’ resembles, or corresponds to, some actually existing hat. It is exactly the reverse: provided that there exists some distinguishable (i.e., spatiotemporally distinct) object unique in every sense, does language offer us a means to grasp it in its specificity? What we argued for here is that it does not. But we also tried to suggest that it need not have to simply because language was not designed for that, and we should not blame it for something it is not expected to achieve.