Coordinating taxonomical and observational meaning: The case of genus-differentia definitions

Bill Noble **Staffan Larsson** **Robin Cooper**
Centre for Linguistic Theory and Studies in Probability (CLASP)
Dept. of Philosophy, Linguistics and Theory of Science
University of Gothenburg
{bill.noble@, sl@ling, cooper@ling}.gu.se

Abstract

Genus-differentia definitions exhibit the dual nature of lexical semantic meaning—they incorporate both "hard" *X is a Y* relations between words, as well as "soft" aspects of meaning which can be supported or challenged by observation. Modeling such definitions as contributions in dialogue requires that we accommodate the fluidity of linguistic resources, while respecting the dual nature of the relations that hold between lexical items. In this paper, we use a Probabilistic Type Theory with Records (ProbTTR) to characterise genus-differentia definitions by describing the update they license to the common ground of a dialogue.

Metalinguistic dialogue is one way for speakers to align on the meaning of words. This is common, for example, between adults and child language learners (Clark, 2007):

(1) a. Naomi: *mittens*.
b. Father: gloves.
c. Naomi: *gloves*.
d. Father: when they have fingers in them they are called gloves and when they are all put together they are called mittens.

But such interactions also take place between adults engaged in a joint activity (Brennan and Clark, 1996):

(2) a. A: A docksider.
b. B: A what?
c. A: Um.
d. B: Is that a kind of dog?
e. A: No, it's a kind of um leather shoe, kinda preppy pennyloafer.
f. B: Okay, got it.

In both of these examples, the participants have a joint perceptual scene to help ground the meaning of the word, but that need not always be the case. Definition is also a common coordination strategy in *word meaning negotiations* that take place on text-based social media (Myrendal, 2019).

In this paper, we consider a particular definition paradigm known as a *genus-differentia* definitions. Consider the following (imagined) exchange between an expert ornithologist and aspiring birder:

(3) a. A: You know what a corvid is, right?
b. B: Yeah, sure. We have jays and crows in the garden sometimes.
c. A: A raven is a large black corvid.
d. B: Oh, okay.

Each of the above examples can be analysed as including a genus-differentia definition (Table 1). Furthermore, it seems reasonable to expect that each exchange results in some update to the *common ground* (Clark, 1996) of the participants.

Discussion of genus-differentia definitions can be traced back at least as far as Aristotle.[1] For Aristotle, each genus must be separated into species by some external *differentia*. Some species, acting as genera themselves, may be further differentiated into subspecies. We adopt some of the language of the Aristotelian tradition (genus, species, differentia), but rather than metaphysics, we are interested in genus-differentia definitions as a conventionalised resource for linguistic agents to coordinate on the meaning a word or phrase.

Genus-differentia definitions convey two kinds of information about the definiendum:

1. **taxonomical information** – A *X is a Y* relationship between the genus and the definiendum

2. **observational information** – One or more *features* that help to *differentiate* the definiendum from other species of the same genus

[1] See especially Books VI and VII of *Topics*.

Ex.	definiendum	genus	differentia
(1)	mittens	mittens $\vee$ gloves	fingers are all put together
(2)	docksider	shoe	leather
		pennyloafer	preppy
(3)	raven	corvid	large, black

Table 1: Three examples metalinguistic coordination analysed as genus-differentia definitions. While (3) fits neatly into the paradigm, the other two deviate somewhat. In (1), the genus is not explicitly stated, but can be taken to be a join type encompassing both *mittens* and *gloves* (see Cooper and Larsson, 2009). In (2), two alternative definitions are given, each with their own genus and differentia.

Marconi (1997) argues that there are two ways for speakers to be competent with the use of a word. *Referential competence* is the ability to map words to individuals or events in the world. If someone can identify a raven by sight (or by call, or by observing its behavior), they might be considered referentially competent with *raven*. This aspect of competence seems to be what is mainly at issue in the argument that at least some aspect of lexical semantic meaning may be associated with a perceptual classifier—a cognitive resource for identifying instances of a class, given some perceptual input (Larsson, 2013; Schlangen et al., 2016). On the other hand, *inferential competence* supports the ability to draw inferences based on the use of a word in context. In a community of bird watchers, one might be expected to infer from an utterance like *I saw a raven* that I saw a corvid. Someone who doesn't make that inference might be considered incompetent with the word *raven*, since part of the meaning of *raven* that they are corvids. Formal semantics in the Montagovian tradition, if it considers lexical semantics at all, focuses on inferential aspects of meaning, for example with meaning postulates (Carnap, 1952; Zimmermann, 1999).

Genus-differentia definitions are interesting to consider from the perspective of interaction because describing the result of grounding an utterance like (3-c) requires a framework that accounts for the dual nature of lexical meaning. We have essentially two desiderata for the shared meaning of *raven* that results from grounding (3-c):

D1 Raven is a species of the genus corvid.[2] This means two things: First, there is an intensional inferential relation from species to the genus. That is, there is no situation (actual or hypothetical) in which something might not be a corvid given that it is a raven, since the definition stipulates that being a corvid is part of *what it means* to be a raven. Second, being a raven is mutually exclusive with each of the sibling species of corvid.[3]

D2 Given that something is a corvid, being large and black (relative to corvids) is positive evidence for being a raven. However, this does not mean that ravens are a *type of* black thing. Any inference from *raven* to *large and black* is defeasible (for example, the speakers may entertain the possibility of albino raven, even if it happens to be extensionally true that all ravens are black). Furthermore, our account should accommodate the possibility that some differentia are interpreted in a way that is sensitive to the context given by the genus. For the sake of example, we will assume that this is the case for *large* but not for *black*.

Our analysis of (3) and therefore these desiderata is admittedly *ad hoc*. Indeed, the use of genus-differentia definitions as a metalinguistic resource is probably a source of variation across different communities of practice. The analysis that leads to these desiderata is partly motivated by the very fact that it requires us to distinguish between taxonomical and observational information about the meaning of *raven*.

We will come back to these desiderata in Section 4 after developing some formal machinery that we can use to express them more precisely. Section 1 introduces Probabilistic Type Theory with Records (ProbTTR). Section 2 describes a way of representing multiclass classifiers in ProbTTR, and Section 3 describes *classification systems*, a kind of ProbTTR type system that encodes a taxonomy

[2] Since we are interested in lexical meaning, the taxonomical information relevant to us is information about *folk taxonomies*, which are a resource for a particular *community of practice* (Gumperz, 1972). Among botanists, a banana is a species of berry while a strawberry is not. The opposite may hold among cooks or in ordinary discourse.

[3] Exactly what the sibling species are may be underspecified in the common ground. In this case, it includes at at least *jay* and *crow*, given the context of (3-b). In other cases, the relevant sibling species may be inferable from the differentia.

with types that refer to multiclass classifiers for their witness conditions. Finally, in Section 4, we will put these tools together to give an analysis of example (3).

1 Probabilistic Type Theory with Records

Probabilistic Type Theory with Records (ProbTTR) is a type system that allows for probabilistic type judgments of the form

$$p(a : T) = r, \tag{4}$$

where $r \in [0, 1]$ is a real number. In settings where the type system is a resource for (or models cognitive processes of) an agent, (4) is taken to mean that the agent judges entity a to be of type T with probability r.[4]

Possibilities and witness conditions In ProbTTR, *witness conditions* are used to compute the probability that a given entity is of a given type. For basic types, $T \in \mathbf{BType}$, witness conditions assign probability dependent on a *possibility* external to the type system. A possibility can be a set theoretic model (in which case the witness conditions for basic types is one of set membership) or it can, as in this paper, be based on a collection of classifiers (see Section 3.2). Thus, we write

$$p(a :_M T) = r \tag{5}$$

to mean that a is of type T with probability r in possibility M. Statements like (4) should only be used for judgments that hold regardless of possibility, or as a shorthand where it is clear that only one possibility is being considered.

We have not explicitly introduced a probability space underlying type judgments. In general, this may not be formally necessary (see Scott and Krauss, 1966). However, if we did, the sample space would be the set of all possible sets of pairs of basic types and entities:

$$\Omega = \mathcal{P}(\mathbf{BType} \times \mathit{Ind})$$

where, for $A \in \Omega$, $\langle T, a\rangle \in A$ would mean that a is of type T in outcome A.

As long as both $\mathbf{BType}$ and Ind are countable (for the purposes of this paper, we may assume they are finite), the distribution is discrete and there is no difficulty in talking directly about the probability of events.

[4] See Cooper et al. (2015) for a more complete introduction to ProbTTR.

A key point that is elucidated by considering the sample space of basic type judgments is that probabilistic dependencies between type judgments on basic types are entirely determined by M.

Conditional probability We may speak of the *conditional probability* that an entity a is of type T_1 given that it is of type T_2, written $p(a : T_1 \mid a : T_2)$. If we wish to express the probability (in general) that something is of type T_1 given that it is of type T_2, this is written $p(T_1 \| T_2)$. The use of the double stroke is to distinguish this expression from the probability that something *exists* of type T_1, given that something *exists* of type T_2, which is written $p(T_1 \mid T_2)$. These conditional probabilities are understood extensionally, specific to a particular *possibility*. If, for example, we know that penguins only live in Antarctica, we would, for the types $\mathit{Penguin}$ (the type of situation in which there is a penguin) and $\mathit{Antarctica}$ (the type of situation in Antarctica), judge $p(\mathit{Antartica} \| \mathit{Penguin})$ to be 1 (or close to 1) on the basis of this contingent fact.

Structured types The witness conditions of structured types are a function of the structure of the type and its components. For example, given types T_1 and T_2, the meet type $T_1 \wedge T_2$ has, witness conditions based on the Kolmogorov (1950) equation for conjunctive probability (Cooper et al., 2015):

$$\begin{aligned} p(a : T_1 \wedge T_2) &= p(a : T_1) \cdot p(a : T_2 \mid a : T_1) \\ &= p(a : T_2) \cdot p(a : T_1 \mid a : T_2) \\ &= p(a : T_2 \wedge T_1) \end{aligned} \tag{6}$$

In addition to types defined with $\wedge$, $\vee$ and $\neg$, ProbTTR defines *record types* as structured types—given a record s and record type R, $p(s : R)$ is a function of type judgments of the fields of s (see Cooper et al. (2015) for details).

1.1 Hard and soft relations between types

Subtype relation In TTR, T_1 is said to be a *subtype* of T_2, $T_1 \sqsubseteq T_2$ if and only if anything of type T_1 is also of type T_2 for any possibility M, (Cooper, forthc, p. 285). Extending this to ProbTTR, we can say,

$$T_1 \sqsubseteq T_2 \text{ iff } p(a :_M T_1) \leq p(a :_M T_2), \tag{7}$$

for any entity a and possibility M.

Naturally, it is not always necessary to check these conditions explicitly.[5] Subtype relations can be implicit in the structure of the types, as in the case of meet types. If $T_3 = T_1 \wedge T_2$, by the definition of the meet type we have $T_3 \sqsubseteq T_1$ and $T_3 \sqsubseteq T_2$.

In other cases, whether two types stand in a subtype relation may depend on what is meant by *all possibilities*. If we literally mean all possible assignments of probability to basic type-entity pairs, then two basic types will never stand in a subtype relation, since there will always be possibilities where $p(a :_M T_1) > p(a :_M T_2)$ and *vice versa*.

If, on the other hand, we restrict our attention to some class of possibilities $\mathcal{M}$, then subtype relations between basic types are possible. Witness conditions are one way to limit the possibilities under consideration and can therefore introduce probibalistic dependency between types.

Evidential relation We introduce a "soft" relation between types in ProbTTR, which captures the notion that T_2 is *evidence for* T_1 in the context of some type T^*. Two types stand in this relation with respect to T^* if learning that something is of type T_2 increases the probability that it is of type T_1:

$$T_1 \prec_{T^*} T_2 \text{ iff } p(T_1 \| T^*) < p(T_1 \| T_2, T^*) \quad (8)$$

This relation is also contingent, relative to a particular possibility.

1.2 Representing probability distributions

In the next section, we will define a type for probabilistic multiclass classifiers—that is, classifiers that compute the probability that a given entity belongs to each of several mutually exclusive classes. To that end, we must first encode discrete categorical probability distributions in TTR, since the output of the classifier takes that form.

Larsson and Cooper (2021) introduce a type theoretic counterpart of a random variable in Bayesian inference. To represent a single (categorical) random variable with a range of possible (mutually exclusive) values, ProbTTR uses a *variable type* $\mathbb{A}$ whose range is a set of *value types* $\mathfrak{R}(\mathbb{A}) = \{A_1, \ldots, A_n\}$. We might have, for example, $\mathfrak{R}(\mathit{Animal}) = \{\mathit{Bird}, \mathit{Reptile}, \ldots\}$.

[5] Indeed, it may not even be possible, depending on the notion of possibility since the "extension" of types with witness conditions based on classifiers is indeterminate (Larsson, 2020b).

We will use short-hands Animal, Bird etc, for the situation where some individual is an animal, bird, etc.:

$$\mathit{Animal} = \begin{bmatrix} \text{x} & : & \mathit{Ind} \\ \text{c} & : & \text{animal(x)} \end{bmatrix}$$

$$\mathit{Bird} = \begin{bmatrix} \text{x} & : & \mathit{Ind} \\ \text{c} & : & \text{bird(x)} \end{bmatrix}$$

For a situation s, a probability distribution over the m value types $A_j \in \mathfrak{R}(A), 1 \leq j \leq m$ belonging to a variable type $\mathbb{A}$ can be written (as above) as a set of Austinian propositions, e.g.,

$$\{ \begin{bmatrix} \text{sit} & = & s \\ \text{sit-type} & = & A_j \\ \text{prob} & = & p(s : A_j) \end{bmatrix} \mid A_j \in \mathfrak{R}(\mathbb{A})\} \quad (10)$$

However, we will also have use for an alternative representation of probability distributions, that indexes the probability assigned to each type with a unique label associated with the type:

$$\text{idx}(\{ \begin{bmatrix} \text{sit} = s \\ \text{sit-type} = A_j \\ \text{prob} = p(s : A_j) \end{bmatrix} \mid A_j \in \mathfrak{R}(\mathbb{A})\}) = \begin{bmatrix} \text{lbl}(\text{A}_1) & = & p_1 \\ \vdots & = & \vdots \\ \text{lbl}(\text{A}_\text{n}) & = & p_n \end{bmatrix}$$

where $p_j = p(s : A_j)$ and $\text{lbl}(A_j)$ is a unique label for $A_j \in \mathfrak{R}(\mathbb{A})$. This means that for a set of probabilistic Austinian propositions P_s, that concern a situation s, $\text{idx}(P_s).\text{lbl}(A_j) = p_j = p(s : A_j)$.

2 Multiclass Classifiers in ProbTTR

In this section we extend the TTR classifier defined by Larsson (2013) to give probabilistic type judgments in multiclass setting.

Larsson (2013) shows how perceptual classification can be modelled in TTR and Larsson (2020a) reformulates and extends this formalisation to probabilistic classification. Adapting the notation of a probabilistic TTR classifier to the current setting, a probabilistic perceptual (here, visual) classifier $\kappa_\mathbb{A}$ corresponding to a variable type $\mathbb{A}$ provides a mapping from perceptual input (of type $\mathfrak{V}$ e.g., a digital image) onto a probability distribution over value types in $\mathfrak{R}(\mathbb{A})$, encoded as a set of probabilistic Austinian propositions.

We also want to explicitly parametrise our classifier. A classifier $\kappa_{\mathbb{A}}$, would thus be a function of type:

$$\Pi \to \mathit{Sit}_{\mathfrak{V}} \to \tag{11}$$
$$\{ \begin{bmatrix} \text{sit} & : & \mathit{Sit}_{\mathfrak{V}} \\ \text{sit-type} & : & \mathit{RecType}_{A_i} \\ \text{prob} & : & [0,1] \end{bmatrix} \mid A_i \in \mathfrak{R}(\mathbb{A}) \}$$

where Π is the type of the parameters needed by $\kappa_{\mathbb{A}}$, and $\mathit{Sit}_{\mathfrak{V}}$ is the type of situations where perception of some object yields visual information, and where $\mathit{RecType}_R$ is the (singleton) type of records identical to R, so that e.g.,

$$T : \mathit{RecType}_{\mathit{Bird}} \text{ iff } T : \mathit{RecType} \text{ and } T = \mathit{Bird}$$

We take classifiers to be part of word meanings. We associate a word like "bird" with a type Bird which is in turn associated with lexical entry in the form of a TTR record:

$$\mathit{Lex}(\mathit{Bird}) =$$
$$\begin{bmatrix} \text{bg} & = & \mathit{Sit}_{\mathfrak{V}} \\ \text{par} & = & \pi \\ \text{intrp} & = & \lambda r : \text{bg} \,.\mathit{Bird} \\ \text{clfr} & = & \lambda r : \text{bg} \,.\kappa_{\mathit{Animal}}(\text{par}, r) \end{bmatrix} \tag{12}$$

Assuming we have a function Lex that looks up the lexical entry related to a type (associated with a word), we also define a lookup function that gives us the classifier corresponding to a type:

$$\mathit{Clfr}(T) = \mathit{Lex}(T).\,\text{clfr}$$
$$\mathit{Intrp}(T) = \mathit{Lex}(T).\,\text{intrp}$$

Let us assume a s_{123} situation where a speaker points to a bird a and says "Bird!" (meaning "that is a bird"). We want to classify a perceived situation as being of the type Bird or not, or in the probabilistic case, compute the probability of the judgment.

Now, to judge the probability with which a situation s is of a type Bird (to continue with our example), the agent looks up the related classifier and applies it to s, which produces a probability distribution over different subtypes of Animal. The agent then looks up the probability associated with Bird. The general method for doing this can be written as:

$$p(s : T) = \text{idx}(\mathit{Clfr}(T)(s)).\,\text{lbl}(\mathit{Intrp}(T)(s))$$

In our case:

$$p(s_{123} : \mathit{Bird}) = \text{idx}(\kappa_{\mathit{Animal}}(\pi, s_{123})).\,\text{lbl}(\mathit{Bird})$$

3 Classification systems in ProbTTR

To represent both taxonomical and observational relations between types, we will embed a *classification system* in ProbTTR. A classification system has two components, a *taxonomy* (Section 3.1), which is a set theoretic object representing an ontological hierarchy, and a collection of *classifiers* (Section 3.2) associated with the taxonomy. Ultimately the classifiers will provide witness conditions for certain basic types and the taxonomy will be fully encoded in the type system, but first we define the structure in set theoretic terms so that we can create a ProbTTR system with the correct subtype relations.

3.1 Taxonomy

A taxonomy is a rooted tree structure defined by a tuple,

$$\mathbf{T} = \langle T, D, t^* \rangle, \tag{13}$$

where T is a set of *taxons*, $D \subseteq T \times \mathcal{P}(T)$ is a set of *distinctions* on T, and $t^* \in T$ is the root taxon.

To elaborate, T is simply a finite set of labels and D provides the hierarchical structure of the taxonomy. *Distinctions* (elements of D) take the form $\langle g, S \rangle$, where $g \in T$ and $S \subset T$, and $|S| \geq 2$. We say that the taxons g and s stand in a genus-species relationship if there is some $\langle g, S \rangle \in D$ such that $s \in S$. Then s can be said to be a *species of* g. Alternatively, we can say that g is *the genus of* s.

This requires certain restrictions on $\mathbf{T}$. Namely, that it is:

- **Acyclic**: There are no cycles. I.e., no chain of distinctions $\{\langle g_1, S_1 \rangle, ..., \langle g_n, S_n \rangle\}$ such that $g_2 \in S_1, ..., g_n \in S_{n-1}$ and $g_1 = g_n$.
- **Rooted**: There is no distinction $\langle g, S \rangle \in D$ with $t^* \in S$.

- **Uniquely connected**: For every $t \neq t^*$ there is exactly one $\langle g, S\rangle \in D$ such that $t \in S$.[6]

Importantly, this still allows for multiple distinctions in which the same taxon acts as a genus. In other words, we can have $\langle g, S\rangle, \langle g, S'\rangle \in D$ where $S' \neq S$. For example, we might imagine a taxonomy in which both $\langle \mathit{Animal}, \{\mathit{Bird}, \mathit{Reptile}, ...\}\rangle$ and $\langle \mathit{Animal}, \{\mathit{Carnivore}, \mathit{Herbivor}, \mathit{Omnivore}\}\rangle$ are distinctions.

The *uniquely connected* constraint allows us to define a function

$$\mathit{Dist} : T \setminus \{t^*\} \to D \tag{14}$$

that gives, for each taxon, t (other than t^*), the distinction $\mathit{Dist}(t) = \langle g, S\rangle$ such that $t \in S$. For convenience we also define the functions Genus, and $\mathit{Siblings}$ such that

$$\langle \mathit{Genus}(t), \mathit{Siblings}(t)\rangle = \mathit{Dist}(t). \tag{15}$$

Note that under this definition, leaf taxons are those taxons for which there are no distinctions in D where the taxon appears as a genus.

3.2 Species Classifiers

In addition to the taxonomy, we have a collection of classifiers, $\mathbf{K}$ and parameters $\mathbf{P}$, each of which we index with elements of D, such that $\kappa_d \in \mathbf{K}$ is the classifier for distinction d provided with the appropriate parameters. This follows the intuition that a distinction in the taxonomy may be accompanied by an ability to *distinguish* among the relevant species. In general, we need only assume that we have classifiers for those distinctions that include at least one leaf taxon, since genus taxons can be defined as the join of their species in certain cases.[7] For now we will assume we have a classifier for each distinction in D.

3.3 The type system

Suppose we have a taxonomy $\mathbf{T} = \langle T, D, t^*\rangle$ and a collection of classifiers $\mathbf{K}$ on the distinctions of that taxonomy. Let Dom be a special type corresponding to the root of the taxonomy. We then define variable types $\mathbb{A}_d$ for each $d = \langle g, S\rangle \in D$ with $\mathfrak{R}(\mathbb{A}) = \{A_{s_1}, ..., A_{s_n}\}$ corresponding to $s_1, ..., s_n \in S$. Classifiers provide the witness conditions for the value types as described in Section 2. For a given entity a,

$$p(s : A_t) = \begin{cases} 1 & \text{if } t = t^* \\ \kappa_{\mathit{Dist}(t)}(a)(t) & \text{otherwise} \end{cases} \tag{16}$$

In other words, the probability assigned to A_t is 1 in the case of the root taxon, and otherwise determined by the classifier for the distinction corresponding to the variable in which A_t is a value type. These "auxiliary" value types we can give the witness conditions for the associated with the taxonomical categories as the product of the jugment of the genus and the axiliary type. For any object a,

$$p(a : T_t) = p(a : A_t) \cdot p(a : T'_t) \tag{17}$$

where

$$T'_t = \begin{cases} \mathit{Dom} & \text{if } t = t^* \\ T_{\mathit{Genus}(t)} & \text{otherwise} \end{cases}$$

This stipulates that the classifiers give us the probability that an individual is of each of the species types, *given* that it is of the genus type. Thus judgments about T_t correspond to an *absolute* judgment about belonging to the taxon.

Taken together, Equations 16 and 17 imply that for any a, $p(a : T_{t^*}) = p(a : \mathit{Dom})$. In situations where the root taxon corresponds to all individuals (i.e., where $\mathit{Dom} = \mathit{Ind}$), we have $p(a : T_{t^*}) = 1$ for any a. It is also possible, however, to embed a classification system in an existing type system, as long it provides witness conditions for Dom. For example, if the classification system is specific to birds, we might embed it in a larger system that gives witness conditions for Bird.

3.4 Feature classifiers

In addition to the distinction classifiers, a classification system may include some number of types based on feature classifiers. A feature classifier takes any entity $a : \mathit{Dom}$ as input, and receives its witness conditions from a classifier that results in a probabilistic type judgement. In general, feature and distinction classifiers need not interact explicitly though, considered as random variables, there may be probabilistic dependence between them. Distinction classifiers may be defined in terms of

[6] A weakness of insisting on a tree structure is that we cannot have taxons that appear in multiple places in the taxonomy, whereas in folk taxonomies it would appear this is common. We would either need to say that the apparently duplicated taxon is actually part of a distinction at a higher level that encompasses both, or that it corresponds to two senses of the same word.

[7] See Marconi (1997, ch. 6) on "subordinate concepts".

feature classifiers, for example as Bayesian classifiers that take the result of feature classifiers as their input (see, Larsson and Bernardy (2021)).

In general, some of these feature types may be dependent types. Consider a type like *Tall*. Whether or not an individual is tall may depend on a comparison class (for example, a type in the taxonomy). Following Fernandez and Larsson (2014), we define dependent feature types with classifiers that take a threshold function as a parameter. For example,

$$\theta_{Large} : Type \to \mathbb{R}^+ \tag{18}$$

This gives the classifier the following type:

$$\kappa_{Large} : (Type \to \mathbb{R}^+) \to Type \tag{19}$$

4 Combining the observation and taxonomical aspects of genus-differentia definitions

With this formal machinery in place, we return to the project of characterising the result of grounding (3-c). First, let's lay out what is shared among speakers A and B before (3-c) is grounded.

We will assume that A and B share a classification system with *Bird* at its root as part of their common ground. Utterance (3-d) establishes that a type for the lexical entry of *corvid*, for which we will use Cor, is a type in this system, and that there is a distinction on Cor such that $\mathfrak{R}(Cor) \supseteq \{Jay, Crw\}$, where Jay and Crw are the lexical entries for *jay* and *crow*—that is, for all species types of Cor given by the common ground , S (including at least Jay and Crw), $S \sqsubseteq Cor$. The witness conditions for each $S \in \mathfrak{R}(Cor)$ are given by a multiclass classifier κ_{Cor}. Since $Cor \sqsubseteq Bird$, we may also assume that $Dist(Cor)$ exists and that there is a classifier $\kappa_{Dist(Cor)}$, though it need not be common ground what the genus of Cor is.

Furthermore, we will assume we have types Lrg and Blk, whose witness conditions are given by feature classifiers. For the purposes of the example, we will assume that Blk is basic type that gets its witness conditions from a feature classifier, κ_{Blk}, whereas $Lrg : Type \to Type$ is a dependent type with a classifier that depends on threshold function θ_{Lrg}. Thus, the witness conditions for $Lrg(Cor)$ are given by $\kappa_{Lrg}(\theta_{Lrg}(Cor))$. This leaves open the question of exactly how θ_{Lrg} is defined, but we may assume that the value of $\theta_{Lrg}(Cor)$ depends in some way on the parameters of the classifier that defines the witness conditions for Cor, namely $\kappa_{Dist(Cor)}$.

Returning to our desiderata, we want to construct a type, Rav, such that:

$$\sum_{T \in Species(Cor) \cup \{Rav\}} p(T \| Cor) = 1 \tag{20a}$$

$$Rav \sqsubseteq Cor \tag{20b}$$

$$Rav \prec_{Cor} Lrg(Cor) \wedge Blk \tag{20c}$$

Here (20a) and (20b) formalise **D1** and 20c formalises **D2**.

4.1 Constructive approach

As discussed previously, one motivation for formalising this example and the interactive semantics of genus-differentia definitions in general is to expose some crucial distinctions in lexical semantics that are often overlooked. In this section, we give what is a rather straight-forward and intuitive solution to the challenge we have given ourselves, but one that fails to adequately make the distinction between taxonomical and observational lexical information.

In this solution, we attempt to directly construct a new type Rav out of the common ground types already available. The most straight-forward way to do this is with meet types:

$$Rav = Cor \wedge (Lrg(Cor) \wedge Blk) \tag{21}$$

This definition is intuitively appealing—(3-c) is saying that ravens are large *and* black *and* corvids. Furthermore, this definition does actually satisfy the desiderata stated so far.

To maintain (20a), we can redefine each existing species type S as:

$$S' = S \wedge \neg Rav \tag{22}$$

We have $Rav \sqsubseteq Cor$, satisfying (20b), since by the Kolmogorov (1950) definition of the meet type (6), for any possibility M and any entity a,

$$\begin{aligned} & p(a :_M Rav) \\ =& p(a :_M Cor) \cdot p(a :_M Lrg(Cor) \wedge Blk \mid Cor) \\ \leq& p(a :_M Cor) \end{aligned}$$

Finally, (20c) holds since it follows from the definition of Rav that, $p(Rav \| Lrg(Cor) \wedge Blk, Cor) = 1$ and, assuming there are at least some non-large, non-black corvids, $p(Rav \| Cor) < 1$.[8]

However, the definition of the meet type (6) implies we also get $Rav \sqsubseteq Lrg(Cor)$ and $Rav \sqsubseteq Blk$. It does not make sense for Rav to be a *subtype* of large corvids or of black things (consider again the possibility of an albino raven). Put another way, it should be possible to construct a hypothetical possibility M and entity a such that:

$$p(a :_M Lrg(Cor) \wedge Blk) = 0 \text{ and } p(a :_M Rav) > 0 \quad (23)$$

In the next section, We will consider this a new desiderata along with the constraints in (20). Instead of constructing the type directly from existing types, we posit a basic type without explicit witness conditions, but with some constraints that are derived from by the genus-differentia definition.

4.2 Underspecifed approach

Cooper (forthc) treats types as having an existence independent of their witness conditions. Two types can share the same witness conditions, for example, and still play different roles in an agent's type system. Part of the motivation for doing this is that an agent can reason about a type and its relation to other types without specifying witness conditions for that type. This is in contrast to predicates in first-order logic, for example, which don't have any meaning independent of the model theoretic entities they are interpreted as.

We would like to interpret definitions like (3-c) as giving rise to an underspecified type; that is, a type without explicit witness conditions. Instead, we assert the following relationships between the new underspecified type Rav and other existing common ground types:

$$Rav \sqsubseteq Cor \quad (24a)$$

$$p(Lrg(Cor) \wedge Blk \| Rav) = 1 \quad (24b)$$

Notice that neither of these two conditions give us direct witness conditions for Rav. The first condition says that anything (in any possibility) that is a raven is also a corvid. The second condition says that anything that is a raven is, with probability 1, is large (for a corvid) and black. Note that (24b) is a constraint on the type's witness conditions given the current possibility, meaning that we can not infer $Rav \sqsubseteq Lrg(Cor) \wedge Blk$, since nothing prevents us from constructing a possibility in which (23) holds. In other words, albino ravens are still possible.

Clearly condition (20b) is satisfied by construction. This may be a bit unsatisfying, but it is worthwhile to consider that asserting $Rav \sqsubseteq Cor$ amounts to adding Rav as a witness condition to Cor. Put another way, for any entity a and possibility M, $P(a :_M Cor) \geq P(a :_M Rav)$.

In order to satisfy (20a), we need to redefine the witness conditions of the existing species types to "make room" in the probability distribution for Rav. How to do this depends somewhat on how completely the distinction is specified in the common ground. If there is an *other corvid* type, $Other$, we might just redefine the classifier for that type so that for any entity a, $\kappa'_{corvid}(a)(other) = \kappa_{corvid}(a)(other) - f(a)$, where f is such that $0 < f(a) < \kappa_{corvid}(a)(other)$. Alternatively, we might take some probability from each class. Either way, the solution should be a function of a that depends on the differentia, but exactly what that function is is not common ground since (24b) gives a unidirectional conditional—all ravens are large and black, but there may still be large, black, non-raven corvids.

It remains to be shown that (20c) holds. In the following, let $D = Lrg(Cor) \wedge Blk$ and S be the set of types representing each of the sibling species of Cor, including Rav.

$$\begin{aligned}
&p(Rav \| D, Cor) \\
&= \frac{p(Rav \| Cor) \cdot p(D \| Rav, Cor)}{\sum_{T \in S} p(T \| Cor) \cdot p(D \| T, Cor)} && (25) \\
&= \frac{p(Rav \| Cor) \cdot p(D \| Rav)}{\sum_{T \in S} p(T \| Cor) \cdot p(D \| T)} && (26) \\
&> p(Rav \| Cor) \cdot p(D \| Rav) && (27) \\
&= p(Rav \| Cor) && (28)
\end{aligned}$$

In the above, (25) follows from Bayes rule and the fact that $\sum_{T \in S} p(T \| Cor) = 1$, and (26) follows from $Rav \sqsubseteq Cor$. For (27), we must assume that

$$\sum_{T \in S} p(T \| Cor) \cdot p(D \| T) \leq 1.$$

This is the same assumption we made in the previous approach, which we argue follows from

[8] This assumption is justified by a pragmatic requirement of genus-differentia definitions that the differentia do at least some work to *differentiate* the definiendum from other species of the genus.

the pragmatics of genus-differentia definitions—namely that not all non-raven corvids are large and black. Finally, (28) follows directly from (24b).

In this approach, the type for *raven*, Rav is defined only in terms of its relationship to types corresponding to other terms in the utterance. A notable feature of this solution is that everything we learn from the definition can be stated in terms of witness conditions for types that already exist: In the case of *corvid*, we know that anything that witnesses the type Rav is a witness for the type Cor. This holds intensionally, meaning that it is true independent of possibility. In the case of *large* and *black*, we know *extensionally* that anything that is a raven will be large and black.

Speaker B learns the type Rav and the constraints associated with it (24) based on the definition offered by A in (3-c). After (3-d), this type and the associated constrains are added to the common ground.

5 Conclusion

The main goal of this paper was to develop a framework that can deal with the distinction between taxonomical and observational lexical information. We argue that this distinction is one that speakers make in metalinguistic interaction, as in genus-differentia definitions. In order to account for this distinction, we use a type system in which intensional relations between types can be reasoned about independently of their witness conditions, which depend on facts about the world.

Our account has been agnostic to the implementation of the classifiers involved. This is justified, in part, by the fact that we describe updates to the conversational common ground, rather than individual agents' abilities. However, it may also be interesting to consider what effect a dialogue like (3) may have on speaker B's ability to recognise ravens. This is related to the machine learning task of *zero-shot classification*, in which an existing classifier is adapted to recognise instances of previously unknown classes based on external information (such as a natural language descriptions). Future work should consider how zero-shot classification can be analysed from an interactive perspective.

References

Susan E Brennan and Herbert H Clark. 1996. Conceptual Pacts and Lexical Choice in Conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(6):1482.

Rudolf Carnap. 1952. Meaning postulates. *Philosophical Studies*, 3(5):65–73.

Eve V. Clark. 2007. Young Children's Uptake of New Words in Conversation. *Language in Society*, 36(2):157–182.

Herbert H. Clark. 1996. *Using Language*. Cambridge University Press.

Robin Cooper. forthc. *From Perception to Communication: A Theory of Types for Action and Meaning*. Oxford University Press.

Robin Cooper, Simon Dobnik, Shalom Lappin, and Staffan Larsson. 2015. Probabilistic Type Theory and Natural Language Semantics. In *Linguistic Issues in Language Technology, Volume 10, 2015*. CSLI Publications.

Robin Cooper and Staffan Larsson. 2009. Compositional and ontological semantics in learning from corrective feedback and explicit definition. In *Proceedings of the 13th Workshop on the Semantics and Pragmatics of Dialogue - Full Papers*.

Raquel Fernandez and Staffan Larsson. 2014. Vagueness and Learning: A Type-Theoretic Approach. In *Proceedings of the Third Joint Conference on Lexical and Computational Semantics (*SEM 2014)*, pages 151–159, Dublin, Ireland. Association for Computational Linguistics and Dublin City University.

J Gumperz. 1972. The Speech Community. In Pier Paolo Giglioli, editor, *Language and Social Context: Selected Readings*. Harmondsworth : Penguin.

A. N. Kolmogorov. 1950. *Foundations of the Theory of Probability*. New York: Chelsea Pub. Co.

Staffan Larsson. 2013. Formal semantics for perceptual classification. *Journal of Logic and Computation*, 25(2):335–369.

Staffan Larsson. 2020a. Discrete and Probabilistic Classifier-based Semantics. In *Proceedings of the Probability and Meaning Conference (PaM 2020)*, pages 62–68, Gothenburg. Association for Computational Linguistics.

Staffan Larsson. 2020b. Extensions are Indeterminate if Intensions are Classifiers. In *SemDial 2020 (WatchDial) Workshop on the Semantics and Pragmatics of Dialogue*, page 10, Waltham, MA and online.

Staffan Larsson and Jean-Philippe Bernardy. 2021. Semantic Classification and Learning Using a Linear Tranformation Model in a Probabilistic Type Theory with Records. In *Proceedings of the Reasoning and Interaction Conference (ReInAct 2021)*, pages 14–22, Gothenburg, Sweden. Association for Computational Linguistics.

Staffan Larsson and Robin Cooper. 2021. Bayesian Classification and Inference in a Probabilistic Type Theory with Records. In *Proceedings of the 1st and 2nd Workshops on Natural Logic Meets Machine Learning (NALOMA)*, pages 51–59, Groningen, the Netherlands (online). Association for Computational Linguistics.

Diego Marconi. 1997. *Lexical Competence*. Language, Speech, and Communication. MIT Press, Cambridge, Mass.

Jenny Myrendal. 2019. Negotiating meanings online: Disagreements about word meaning in discussion forum communication - Jenny Myrendal, 2019. *Discourse Studies*, 21(3):317–339.

David Schlangen, Sina Zarrieß, and Casey Kennington. 2016. Resolving References to Objects in Photographs using the Words-As-Classifiers Model. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1213–1223, Berlin, Germany. Association for Computational Linguistics.

Dana Scott and Peter Krauss. 1966. Assigning Probabilities to Logical Formulas. In Jaakko Hintikka and Patrick Suppes, editors, *Studies in Logic and the Foundations of Mathematics*, volume 43 of *Aspects of Inductive Logic*, pages 219–264. Elsevier.

Thomas Ede Zimmermann. 1999. Meaning Postulates and the Model-Theoretic Approach to Natural Language Semantics. *Linguistics and Philosophy*, 22(5):529–561.