Trees and Forests: The Formal Semantics of Collectivity

Yoad Winter Seggev

Inaugural speech made following appointment to the professorial position "Computational Semantics and Artificial Intelligence"

Utrecht University, 10 April 2017

Sir, Rector Magnificus,

Dear attendees,

Introduction

Language gives structure to our thoughts. When I say "I saw that bird", I convey a different thought than when I say "that bird saw me". The different order of the words enables me to express the different roles of the players in these similar thoughts. What we see is a tight connection between linguistic ordering of words and mental life. This connection between *form* and *meaning* is present in the most common sentences that we use every day. In this way, language helps us to express important parts of our mental life and convey them to others.

The linguistic structures we use are immensely complex. In the Workbench of Academic Skills of the Free University of Amsterdam (VU) we find the following example:

(1) *Hoewel de meeste studenten, die vaak, ook al praten ze behoorlijk veel tijdens colleges, best geïnteresseerd lijken, zou het nog veel interessanter zijn als ze af en toe, wanneer het zo uitkwam, zwegen.*

"Although most students, who often, even if they tend to chat a lot during classes, look quite interested, it would still be much more interesting if they, here and there, when it fits, keep quiet."

The Workbench of the VU urges us to avoid such sentences in writing. This indeed may be a good advice. However, at the same time, this example of "bad" language use illustrates something quite marvelous about human language: the complexity of linguistic structure allows us to communicate complexities in our internal thoughts. This is what makes our communicative skills so unique among the big primates. Could you imagine an old alpha male chimpanzee expressing such a thought about young chimpanzees that keep interrupting his distinguished speech?

Because of the structured nature of human language, it functions as a window to our cognition: language provides cognitive scientists with a peephole through which they can look into some of the intricate workings of our mind. In this lecture I concentrate on one domain where language reflects our mental activity: the arrangement of objects into *collections*, or sets of objects. The structures that language supports allow us to form assemblages of objects like mountains, trees, and family members, or abstract notions like feelings, letters and numbers. Think of expressions like "the Sopranos", "Pauw en Witteman" or "the Milky Way". These expressions all refer to collections. How does language support the arrangement of objects into collections? What does it tell us about our mental life? How can we accurately model the ways we talk about collections?

To illustrate these puzzles, let me start with a short story.

On the way to the Rocky Mountains

While you are preparing for a trip to Denver, you meet a nature-loving friend. She is terribly excited to hear about your plans. "Lucky you!" she exclaims, "you surely know that Denver is only 100 km. from the magnificent Rocky Mountains!" Prompted by your friend's enthusiasm, you decide to check her claim on the map. Doing that, you instinctively look up those parts of the Rockies that are closest to Denver, and ignore other parts. This "zooming in" reflects your intuitive understanding of the expression "100 km. from". Looking for more information in Wikipedia, you read: "The Rockies were formed from 80 million to 55 million years ago". Now, no specific part of the Rockies is relevant. Rather, the sentence directs your attention to consider the Rockies holistically, as if they were an atomic entity without distinguishable parts. As you keep on reading, you learn that "the Rockies vary in width from 100 to 500 km." Now you again zoom in on parts of the Rockies, but your attention is drawn to several distinct parts.

In this story we classify three different ways in which language supports reference to collections:

(2) Denver is 100 km. from the Rockies
    - *only specific parts of the Rockies are relevant: the nearest ones*
(3) The Rockies were formed from 80 million to 55 million years ago
    - *the Rockies are conceptualized as a whole: an abstract object*
(4) The Rockies vary in width from 100 to 500 km.
    - *different parts of the Rockies are relevant: the narrowest and the widest ones*

Collective reference and collective categorization

In the Rockies example, a short linguistic expression invokes in our mind a complex object consisting of many distinguishable parts. The three sentences in (1)-(3) all talk about the same collection of mountains. However, each of these sentences is used for *categorizing* that collection using a different aspect, referring to location, age, or dimension. In sentences (2) and (4) we categorize the Rockies using their location and dimensions. This must be done by focusing on one or more specific parts of the Rockies, and by using a fair amount of abstraction: the Rockies are too big an object to gauge directly. Sentence (3) categorizes the Rockies in a different way: according to their age. This categorization does not involve any specific part of the mountain range: it concerns the collection of mountains as a whole. In this sentence, we need yet a greater level of abstraction to be able to talk about the inconceivable age of that huge collection of mountains.

To summarize the general point, examples (2)-(4) all demonstrate our human ability to *refer* to collections and *categorize* them. In simple sentences of the structure "the Rockies … (are faraway, vary in width)" we see the two abilities at work: the subject "the Rockies" *refers to* a collection, and the adjective or verb *categorizes* it. *Collective reference* is our ability to use a short linguistic expression when referring to an assemblage that has many distinguishable members or parts. *Collective categorization* is our ability to classify the collections we refer to, taking into account their internal structure and the properties of their parts.

Collective reference, collective categorization and their emergence in grammar

When referring to a collection, speakers have different grammatical options to choose from. For instance, when describing a collection of trees, we may use a grammatically singular term "this forest", or choose the plural phrase "these trees". This choice immediately affects the workings of collective categorization, hence the message we convey to the hearer. The following examples illustrate this:

(5) a. These trees are healthy.
    b. This forest is healthy.

In both (5a) and (5b), the concept "healthy" is used for categorizing a collection of trees. Sentence (5a) categorizes individual trees using criteria like color, quality of metabolism, functionality of tree parts etc. By contrast, in (5b) the health criteria are somewhat different, and apply to the collection as a whole. As the website of the North Carolina Forest Service states: "a healthy forest can have unhealthy trees".

The examples in (5) illustrate how singular/plural grammatical marking triggers different procedures of collective categorization. How does this triggering work? What are the general grammatical mechanisms for representing collections? How do these mechanisms interact with categorization? These questions are critical for designing computer algorithms for language processing, for translating concepts between languages, and analyzing linguistic processing in the human brain. Semanticists use introspection, corpus studies, and questionnaires to study these processes, as well as psycholinguistic and neurolinguistic experiments. Formal and computational semanticists study the mathematical theories underlying these meaning processes, and develop algorithms that model semantic processes and try to mimic them as accurately as possible. These efforts are at the heart of an interdisciplinary domain within Artificial Intelligence, involving principles from Linguistics, Mathematical and Philosophical Logic, Computer Science and Cognitive Psychology. From this huge domain of ideas, I will now capitalize on two general sub-domains.

Lexical semantics and the theory of meaning composition

In natural language semantics, we classify two different processes of assigning meanings to language expressions. One process is looking up meanings of words in our mental lexicon; another process is putting these meaning together in sentences. The assignment of meaning to words is what we see in common dictionaries. For instance, the Van Dale basic dictionary defines the meaning of the verb *zwijgen* ("keep quiet") as "not letting your voice be heard,

not talking". Like this verb, almost all words in natural language express ideas, or *concepts*. These ideas are the building blocks of language meanings, which linguists try to capture by developing theories of *lexical semantics*.

Another process of meaning assignment involves complex linguistic structures, like the example in (1). In this long sentence, the meaning of the word "zwijgen", although it is quite short, contributes something critical to the meaning of the long sentence. The meaning of this word amalgamates somehow with the complex structure that precedes it, and changes its meaning. When linguists study how word meanings amalgamate into sentence meanings, we say that they study the process of *meaning composition*. This process shows certain important regularities. In many linguistic textbooks, this is illustrated by the following example:

(6) We saw the man with the telescope.

Sentence (6) has two syntactic structures:

Structure A: "the man" forms a complex *noun phrase* with "the telescope" – "the man with the telescope". This complex NP is the object of the verb "saw".

Structure B: "the man" forms a complex *verb phrase* with the verb "saw" – "(we) saw the man". The additional expression "with the telescope" describes this activity.

In correspondence with these two structures, sentence (6) has two different meanings:

Meaning A: we saw a man; this man had telescope

Meaning B: we saw a man; to do that we used telescope

We explain these two meanings of sentence (6) as follows: the different structures A and B impose different orders of meaning composition on our understanding of the sentence. This classical analysis of *structural ambiguity* is based on analyzing meaning composition as illustrated in Figure 1.

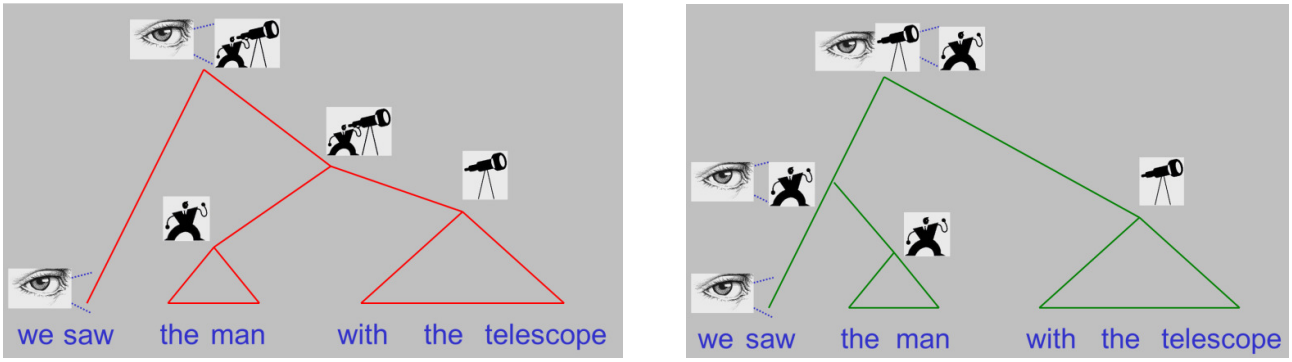

Figure 1: *The two structures and the two meaning compositions for the sentence "we saw the man with the telescope"*

The process of meaning composition is what allows us to take word meanings and a linguistic structure, and use them for understanding sentences, be them as simple as (6) or complex and cumbersome as (1). Analyzing meaning composition is critical for all parts of theoretical linguistics and AI. In a recent PhD thesis, Assaf Toledo (2015) has shown how these models of meaning composition are profitable when modeling how humans extract logical information from ordinary texts. This shows how formal models in semantic theory support current efforts in computational linguistics and AI to understand simple reasoning from texts.

In all our contemporary efforts in formal semantics, one of the greatest puzzles is how to combine models of meaning composition with models of lexical meaning. Much of my research in Utrecht University concerns the division of labor between word meaning and meaning composition, especially in relation to the way we use collections in language. This research involves much collaboration with other linguists, psychologists, and computer scientists. In a recent book that I edited with professor James Hampton from the City University of London (Hampton and Winter 2017), we collected some of the recent results that teams around the world have achieved in this area.

Composing collections and categorizing them

Collectivity is one of the major areas in language where we see lexical semantics and meaning composition work in tandem. As we saw, language supports our ability to *refer* to collections like *the Rockies* or *the trees*. And then,

language also allows us to *categorize* collections using verbs and adjectives. In the sentence "the trees are healthy" we see both processes very clearly:

- The *plural form* of the subject reflects the abstract idea of a collection. The meaning of the plural suffix "s" is composed to the noun "tree". This is how we know that the subject refers to a collection of trees, and not just to a single tree.
- Once we identify this collection, we categorize it according to the lexical meaning of the adjective "healthy".

This is a simple example where lexical meanings and linguistic processes of meaning composition support the way we think and talk about collections. Let me now move on to two more intricate problems of collections in language. We start out with a problem of meaning composition, and then to a problem of categorization and word meaning.

Collections and Meaning Composition: the semantics of natural language conjunction

When speakers want to refer to a collection, one of the simplest means they have is *conjunction*. The following sentence is about a famous collection: the Bryan twins, the most successful duo in the history of doubles tennis.

(7) Bob and Mike won 16 Grand Slam titles.

Sentence (7) talks about a group achievement of the duo: neither of the Bryan twins can truthfully claim that *he* won 16 Grand Slam titles. This conception of "Kolkhoz Collectivity" was pushed forward by Utrecht Linguistics professor Henk Verkuyl (1994). This term Kolkhoz was used in Soviet Russia to describe collective ownership of a farm by its members (the way this idea developed later is another story). What Verkuyl explained is that in language, like in the ideal of a Soviet Kolkhoz or an Israeli Kibbutz, there is no individual ownership, and not even a "fair share" of the goods. It is only the collection of people that owns things. Similarly, according to Verkuyl's conception, in sentence (7) it is only the duo Bob and Mike *as a whole* that can claim to have won 16 Grand Slam titles. When analyzing simple sentences that show this "Kolkhoz Collectivity" we are facing one of the greatest puzzles for theories of collectivity and grammar: what is the function of the word *and*? How can it amalgamate two distinct names like *Bob* and *Mike* into a (very small) Kolkhoz, which has products, gains and achievements that no Kolkhoz member can claim for himself?

The simplest way to think of the meaning of the word *and* in (7) is as *summation*: a collection forming operator. In (7), the conjunction *and* intuitively creates a collection of entities **bob**+**mike** by combining the entities **bob** and **mike** together (Link 1983, Krifka 1990). However, despite its intuitive appeal, this summation analysis is quite limited. The reason is that in many cases, the function of the word *and* seems completely different from summation. Consider for instance the following examples:

(8) Bob Bryan was a Stanford student and a Triple Crown winner.

(9) Mike Bryan lives in Florida and likes to visit Disney World.

In such constructions, where conjunction occurs between verbs (8) and between sentences (9), its most natural interpretation is *set-theoretical intersection*. In (8), it is reported that Bob is in the intersection of two sets: the set of Stanford students and the set of Triple Crown winners. In (9), it is reported that Mike is in the intersection of the set of people who live in Florida and the set of people who like visiting Disney World. This intersective use of *and* is very different from the summation that we saw in sentence (9). The two functions of conjunction – as a summation operator, and as an intersection operator – are summarized in Figure 2.

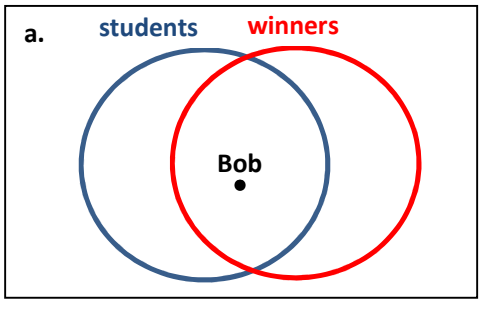

Bob was a student and a winner

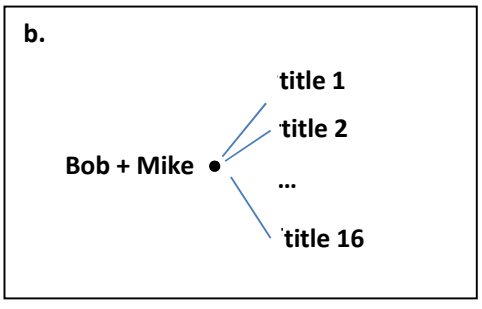

Bob and Mike won 16 Grand Slam titles

Figure 2: *Conjunction as summation, and conjunction as set intersection*

How can we explain the fact that conjunction has two different functions – a summation operator as well as intersection operator? One simple possibility is to assume that the word *and* in English is ambiguous (Hoeksema 1988): summation in sentence (7) and intersection in sentences (8) and (9). If this is the case, it means that one of the most productive ways that English has for producing collections, via the word *and*, has emerged by virtue of a lexical coincidence: ambiguity of *and* in English. However, it was discovered that conjunction in many different languages consistently shows precisely the same two functions: summation as well as intersection (Payne 1985, Haspelmath 2004). This makes the ambiguity idea improbable: why should conjunction be ambiguous in the same way in different languages that have no genetic connection with each other?

There is a viable alternative to ambiguity. This is obtained once we look at the behavior of conjunction in *generalized quantifier theory* (Montague 1973, Peters and Westerståhl 2006, Keenan 1996). In this theory, all conjunctions are treated as set intersection. What allows this to happen is the *uniform, set theoretic* treatment of noun phrases. In this treatment, all noun phrases represent complex sets: sets whose members are properties of entities. In this way, a simple proper name like *Bob* is also treated as a set: the set of properties that Bob has. The collection for *Bob and Mike* is now – surprisingly – treated as an intersection of two sets: the set of properties that Bob has and the set of properties that Mike has. Let us see how this work in the following sentence:

(10) Bob and Mike are tennis players.

Sentence (10) is treated as meaning "the property *tennis player* is in the set of properties that Bob has *and* in the set of properties that Mike has". This is illustrated in Figure 2. More generally: the collection for *Bob and Mike* is described by looking at the generalized quantifier that describes the intersection of their properties, which is described in Figure 3 (Winter 2001, Champollion 2016).

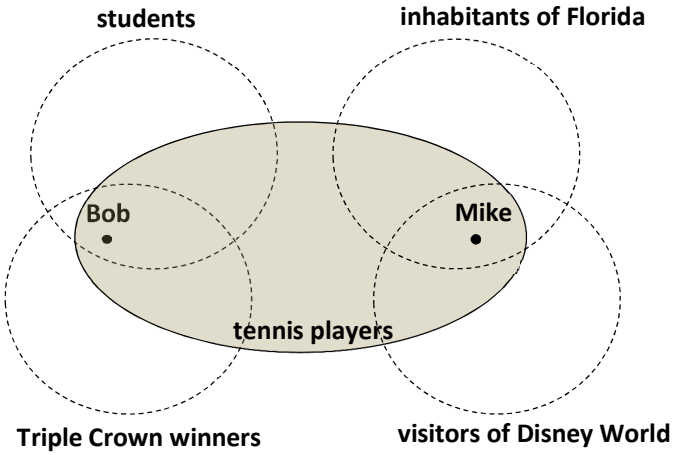

Figure 3: *Conjunction in sentences (8-9) functions as intersection of properties – in (8): students and Triple Crown winners; in (9): inhabitants of Florida and visitors of Disney World. Conjunction in (10) functions as intersection of sets of properties: the property "tennis players" is in the intersection of the set of Bob's properties and the set of Mike's properties.*

We conclude that conjunction operators may form collections by virtue of their uniform behavior in different syntactic categories (Keenan and Faltz 1985):

Hypothesis**:** *Conjunction represents set intersection in all syntactic categories – adjectival phrases, verb phrases, prepositional phrases and full sentences, as well as nominal phrases like "Bob and Mike".*

The ability to use this hypothesis for forming collections solves one of the major puzzles that lie at the heart of our understanding of collectivity in natural language: the way the conjunction *and* supports the formation of collections in our thoughts and conversations.

Collections and word meaning: the geometric semantics of spatial expressions

We have seen that collections can be seen as individuals, like the duo Bob & Mike. Along similar lines, collections can also be located in space like individuals are. In our daily communication, we often refer to the *spatial aspects* of collections: we say things like "the Rockies are *far away*", "the books are *inside* the box", "the soldiers are *surrounded* by the enemy". In all these cases we refer to the location of a collection in the same way that we talk about locations of ordinary entities (Zwarts 1997). When language is used to communicate about spatial collections, it shows some of the most challenging interactions between grammar and collectivity. The example below illustrates this challenge. In this example, research that was done in my work with graduate students at Technion and Utrecht University revealed a previously unfamiliar case of collective categorization: the spatial use of *indefinite descriptions.*

Example – a car race in the desert: Dan is participating in a car race in the desert. Suddenly, he notices that his car is running out of gas. In this context, we consider the following two sentences:

(11) a. Dan is <u>far from</u> a gas station.
b. Dan is <u>close to</u> a gas station.

As Grimm et al. (2014) demonstrate, the interpretations of (11a) and (11b) are dramatically different. Our study has shown that sentence (11a) preferably pertains to Dan's distance from <u>all</u> gas stations, a fact which may have fatal consequences for Dan's trip. However, in (11b), a similar sentence only requires Dan to be close to <u>one</u> gas station.

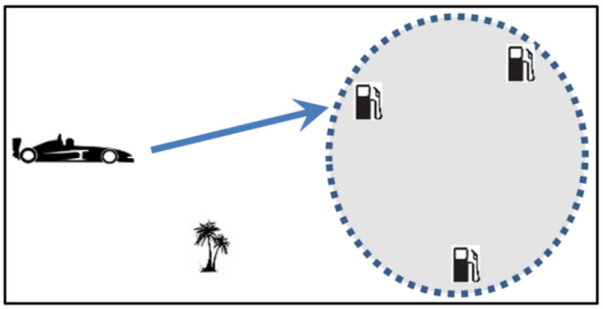

Figure 4: *Collective location of gas stations*

This distinction between the spatial expressions "far from" and "close to" is quite general with noun phrases like "a gas station", and it strongly depends on the grammar of these descriptions. In Mador-Haim and Winter (2015), we proposed that the meanings in (11a-b) appear because *far* and *close* express different geometric relations between Dan's location and the *collective location* occupied by gas stations. We analyzed reference to this collective location, as depicted in Figure 4, using the *property-based* meaning of noun phrases with the article "a" (Farkas and De Swart 2003). We then unified this approach with a view that allows us to talk about trees, even when only mentioning the forest (Scha 1981, Reinhart 1997). When we say that "the forest is nearby", it follows that some tree in it is nearby as well, though not all trees. However, when we say that "the forest is far away", *all trees* in the forest have to be far away. According to this unified analysis, in both (11a) and (11b) we look at the same collective region of gas stations, as if it were one "forest of gas stations". However, in (11a), the geometric features of the concept *far* entail that Dan is far from all gas stations in that location: being far from a region means being far from all of its parts. By contrast, in (11b), the geometric features of the concept *close* only entail that Dan is close to one gas station in that collective location. We reach the following hypothesis:

The geometric features of spatial concepts ("far", "close") *allow speakers to categorize a complex object in terms of its collective location.*

This principle is understood to be one of the building blocks of collective categorization with spatial objects. It makes a novel connection between geometric theories of spatial concepts (Gärdenfors 2004, Zwarts & Winter 2000) and formal semantic theories of collectivity (Winter & Scha 2015).

Intermediate summary

So far we have seen three major processes with collections in language:

- Collections are formed in language using conjunctions (*Bob and Mike*), or by plural number (*the tennis players*). Collections may also be referred to directly by dedicated lexical nouns like *forest* or *family.*
- We may treat collection as groups of individuals or as a quantifiers over individuals. When we report that people get something or achieve something together, or when we locate collections of mountains and gas stations in space, we choose to look at the forest as a whole, rather than at the trees. When we say that five women had a baby, or have the highest salary in some organization, we invariably mean that *each* of these women had a baby or earns a lot of money. In this case we choose to look at individual trees in the collection, rather than the forest as a whole.

These processes, which describe the formation and use of collections in language, have been studied extensively since the early stages of formal semantic theory (Bennett 1974) and have gradually led to a good understanding of the linguistic systems underlying collectivity (Schwarzschild 1996, Nouwen 2014, De Vries 2015, Winter and Scha 2015). I now move on to some advanced questions about collections that are studied in recent formal semantic research.

In the grey area between trees and forests

At a theoretical level, we analyze expressions like *the trees* or *Bob and Mike* as ambiguous between groups, a forest or a tennis duo, and quantifiers over individual trees or individual people. This distinction between groups and quantifiers is clear and useful. Works in formal semantics have analyzed groups and individuals in precise mathematical terms, and we have excellent logical languages and programming languages that model this distinction. However, in the linguistic practice, there are many circumstances where the distinction between groups and quantifiers is not clear cut. Sometimes we need to look at the trees and at the forest simultaneously. And sometimes, we even do not know if we better look at a collective language expression as representing a forest or different individual trees. This grey area between collectivity and quantification is one of our major challenges in semantic theory. I will now discuss three problems in this grey area, and mention some results that we have obtained over the last few years.

Reciprocal expressions

The following sentence illustrates a classical case where quantification and collectivity operate simultaneously.

(12) The girls are pinching each other.

The expression *each other* in sentence (12) involves a *reciprocal activity*: each girl is acting separately, but she is acting on other members of the group. Thus, reciprocal activities are somewhere on the middle of the scale between group activities and individual activities: on the one hand, the reciprocal sentence (12) reports about the collective activity of the group of girls; on the other hand, there is no way to understand this group activity without considering the individual activities of the girls: groups do not have hands to pinch with or cheeks that can be pinched. When we try to analyze what sentence (12) reports about individual girls, we discover some strange effects. In a linguistic experiment, where speakers were asked to give their preference for one of the illustrations in Figure 5 relative to sentence (12), two thirds of the participants preferred the complicated situation in Figure 5a, and about one third preferred the simpler situation in Figure 5b. Despite these different preferences, and despite the oddity of these reciprocal activities, very few speakers doubted that sentence (12) is true in both Figure 5a and Figure 5b.

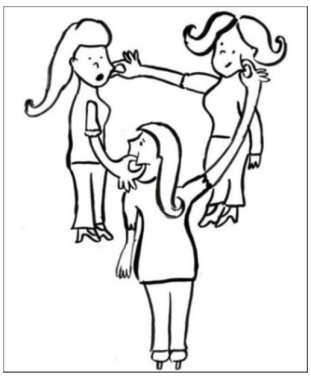

a

b

Figure 5: *Two situations where three girls are pinching each other* (drawings by Ruth Noy Shapira)

However, the results changed dramatically when we tested reactions of speakers to sentences like (13).

(13) The girls know each other.

Here, about half of the participants rejected sentence (13) altogether in the situation of in Figure 6, where each girl only knows one other girl. When testing reciprocal sentences with different verbs (Kerem et al. 2009, Poortman 2017) we then discovered a general phenomenon: reciprocal sentences with verbs like *pinch*, *bite* and *caress* are all accepted in situations where every agent acts on only one patient (Figure 6). However, verbs like *know*, *envy* and *hate* strongly prefer situations where everyone knows, or envies or hates, everyone else. We know this because many speakers reject completely reciprocal sentences in situations like Figure 6.

Figure 6: *a scheme where each agent stands in a uni-directional relation with only one patient*

We describe the difference between reciprocal sentences with the two groups of verbs using the following general principle:

Maximal Typicality Principle: *A reciprocal sentence holistically reports a property of a group. At the same time, speakers are able to use a holistic reciprocal property to infer about activities or states of individual within the group. When doing that, they take the most typical arrangement of the individuals as their point of departure.*

Using linguistic experiments with different verbs, we have shown that for verbs like *pinch* and *bite*, the most typical reciprocal arrangement is as in Figure 6: such cases do not require people to pinch or bite two other people simultaneously. However, with verbs like *know* and *hate*, this arrangement is often rejected. The conclusion is that reciprocal sentences require speakers to perform a complex task of collective categorization. They categorize a collection according to properties of the individuals in it, but when doing that, they pay close attention to which properties these are. The process of collective categorization does not only involve counting properties of individuals. It also involves a decision under uncertainty, where speakers also ask themselves question like "how many people are *likely* to be pinched simultaneously?", "how many people can a person know in a *typical* situation?". Such questions have so far remained outside the realm of formal semantics. However, when we see their relevance to our results, it leads us to a picture where logic and common sense interact with each other in language, much more closely that has been anticipated before.

Collectivity and conflicting actions

There are more cases where language deals with collections according to common sense knowledge. Eva Poortman in her dissertation (Poortman 2017) has conducted experiments where she tested speaker judgements on sentences like (14a-b) in situations where the members of groups perform different activities, as in Figure 7.

(14) a. The women are swimming and crawling.
b. The women are swimming and laughing.

Figure 7: *A situation in which a group of women is doing two different things* (drawing by Ruth Noy Shapira).

In the drawing of Figure 7, there is a group of four women, where two women are swimming and the other two women are crawling and laughing. The question Poortman asked was whether the activities that sentences describe matter for the semantic strategy speakers adopt when hearing them. Her experiments showed a systematic pattern. 100% of the participants accepted sentence (14a) in Figure 7. However, only 24% of the participants accepted sentence (14b) in Figure 7. This is surprising for traditional theories, because the crawling women in Figure 7 are also smiling. Why should it matter how we describe them? Poortman's explanation of is pattern is that swimming and crawling are completely incompatible: you cannot swim when crawling or crawl when swimming. This incompatibility boosts the acceptability of sentences like (14a) in situations like Figure 7, where the group is doing two different things. By contrast, to imagine people laughing while swimming requires no effort at all. This compatibility between the two activities makes sentence (14b) quite unacceptable in Figure 7.

Poortman went further than that, and checked similar effects with activities that are only partially incompatible, like swimming and reading. These activities are quite incompatible, but not completely. With some effort, we can imagine people reading while swimming (if you try it later, please use waterproof reading materials!). Now what happens when Poortman tested sentence (16) in the situation of Figure 8?

(15) The women are swimming and reading.

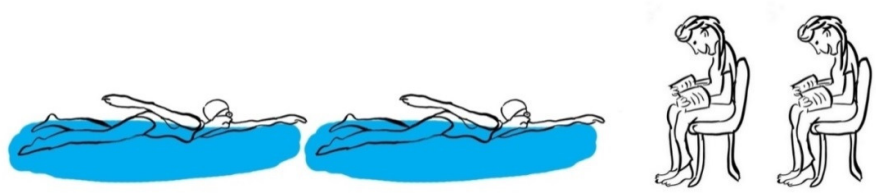

Figure 8: *Another situation in which a group of women is doing two different things* (drawing by Ruth Noy Shapira).

Sentence (15) was more acceptable than sentence (14b), but still not for all speakers – only 76% of the participants in Poortman's experiment accepted it. The mild degree of incompatibility led to a mild degree of acceptability. Furthermore, Poortman discovered a *correlation*: the more incompatibility there is between the mixed actions of the group, the more speakers accept the report about the mixed action (Figure 9). This statistical correlation is another illustration to how our common sense knowledge affects our logical understanding of sentences in natural language. In a further MEG study at the NYU psycholinguistics lab, Poortman and Pylkkänen (2016) showed that the different semantic strategies are also reflected in brain activities.

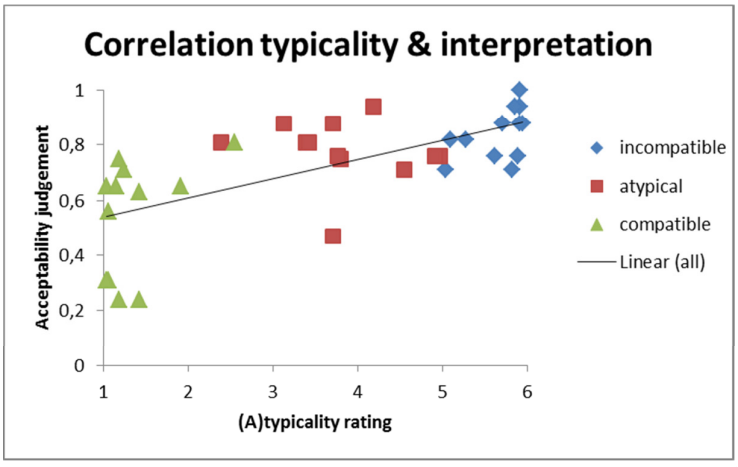

Figure 9: *Correlation between incompatibility and acceptability of mixed reports* (Poortman 2017, p.91)

The lesson we draw from all these facts is quite general: when a sentence mentions a collection, the way we decide to look at its different parts depends on what the sentence reports about those parts. Our level of attention to the trees within the forest is regulated by the meanings of the words and by the grammatical structures we use.

What is a *hug*?

As we have seen, reciprocal expressions like *each other* allow us to speak about groups and their members at the same time. This ability to refer to the forest and its trees simultaneously is also observed with many verbs. For instance, when we refer to an event as a *hug*, we may refer to one of two different things.

- A *hug* may be a reciprocal act, as in Figure 10a.
- A *hug* may be a one-direction act, as in Figure 10b.

Figure 10: *A reciprocal hug and a uni-directional hug* (drawings by Ruth Noy Shapira).

These two faces of the noun *hug* are also observed with the verb *hug*:

(16) a. The two people *hug*.
b. The drunk *hugs* the lamppost.

In relation to our two notions of a "hug", the verb *hug* has two different syntactic functions. In (16a) it is an intransitive verb like *sleep* (which only comes with a subject). And in (16b) *hug* is a transitive verb like *see* (which must also have an object). These two different syntactic functions were studied in detail in work by Alexis Dimitriadis (2008). The verb *hug* in (16a) is *intransitive*, and it holds of one couple with no specification of an active person and a non-active person. The verb *hug* in (16b) is transitive, and it reports an activity of a human agent on a non-animate, and hence surely passive, object – the lamppost. These two uses are clearly different: we cannot describe the strange behavior of the drunk by saying "the drunk and the lamppost are hugging". It just doesn't make sense!

We see that in relation to the two concepts of *hug* that we have, English, as well as Dutch, has two uses of the verb *to hug*. We see the same double-sided behavior of verbs – with collections and with individuals - with many other verbs in Dutch and English:

zoenen "kiss"
kussen "kiss"
botsen "collide"
roddelen "gossip"
verliefd zijn "be in love"
praten "talk"

And even the recent Dutch verb *appen "send WhatsApp messages"* behaves in this way. When we say that a person *appt* his friend about his late arrival, we mean something different than when saying that the two friends *appen*. Like more traditional talking activities, also sending WhatsApp messages may be uni-directional or reciprocal.

What do we actually mean when we say that two people are hugging, gossiping or "appen"? More generally: what do we mean when we refer to collective acts using a one word verb? This question leads us naturally to the area of lexical semantics – the meaning of words. What make verbs like *knuffelen* en *appen* so interesting for linguists is the fact that their tree-forest ambiguity at the semantic level has a nice reflection at the morpho-syntactic level: a distinction between transitive and intransitive usages of the same verb. How precisely does this reflection work? Can we understand what a collection is doing by looking carefully at what each of the members of the collection are doing?

The philosopher John Searle has answered this question in the negative (Searle 1990). As he pointed out, groups of people may have intentions that cannot be directly explained in terms of what each of the members in the collection is doing. In a football match, you cannot understand the goal of each team by just looking at what each player is

doing. Another philosopher, Yeshayahu Leibowitz, seemed to have disagreed with that, defining football as "22 hooligans running after a ball". Still, we all feel that Leibowitz's statement was sarcastic, and that Searle had a point when claiming that the whole may be bigger than its parts, especially when we are dealing with the way teams work together. Human teams are a bit like ant colonies in this respect: as a group, our acts sometimes have a logic of which no one of us is completely aware.

From this short digression about the philosophy of social behavior, let us get back to linguistics. In a recent work with Imke Kruitwagen and Eva Poortman (Kruitwagen et al. 2016) we were addressing the question of sentences like (16a): what does it take for two people to be *hugging*, or *conversing,* or *colliding*? We looked at Dutch examples like the following:

(17) Violet en Mark hebben geknuffeld. ("Violet and Mark hugged")

(18) Violet en Mark hebben gepraat. ("Violet and Mark talked")

In the most prototypical scenarios where sentences (17) and (18) are used, both individuals are active: *each* of them is hugging the other one, or talking to the other one. However, what happens when one of the parties does not reciprocate? In these cases we have one active individual who is hugging the other individual or talking to him, but the other individual remains passive. What happens in such active-passive situations? What can we say about sentences like (17) and (18)? In our experiments it turned out that what matters much for speakers' judgements is the amount of *collaboration* that the passive individual demonstrates. When the passive individual looked collaborative, about half of the participants accepted sentences like (17) and (18). When the passive individual was not collaborative, only between 10-25% of the participants agreed that the sentence is true. This is illustrated in Figure 11.

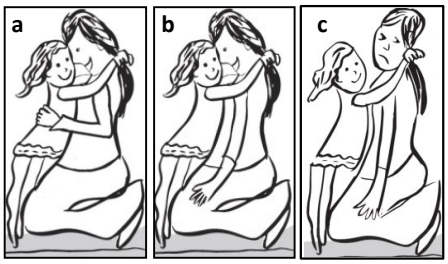

Figure 11: *Sentence (17) was accepted by 100%, 48% and 19% of the participants in a, b, and c, respectively* (drawings by Ruth Noy Shapira).

We conclude that language encodes the interaction between group-action and individual-action. However, as Searle expected, this encoding is not straightforward. A collective *hug* is not simply a uni-directional *hug* going both ways. There is more in the forest than just the individual trees that it contains!

Concluding comments

The recurrent message in my talk has been that language has rich ways to represent collections. To understand our ability to conceptualize collections in our mind, and to model this ability in computational semantics, we should analyze in exact detail how collections are represented in language. We have seen two major mechanisms that are responsible for that. One sort of group representation emerges within language itself, by means of *grammatical* mechanisms: especially *conjunction* ("Bob and Mike") and *plurality* ("the tennis players"), as well as indefinite descriptions in certain constructions ("Dan is far from *a gas station*"). These grammatical strategies for representing collections belong in our language faculty: the great majority of human languages use them productively. Another way of representing collections is through *lexical* concepts: think of nouns like *forest*, *family, galaxy* etc. The grammatical and the lexical strategies for representing collections are relatively well understood. What we are currently studying – and this is a major challenge for much on-going work – is how humans *categorize* collections. What do we actually mean when we say that "these people *admire each other*"? How do we understand a sentence like "these four cooks are standing and sitting"? How come the same verbs can categorize a collective act ("Mary and John are hugging") as well as a non-collective, directional act ("Mary is hugging John")? These questions about collective categorization are about one of the most fascinating interactions between grammar and the psychology of concepts. Our work in Utrecht continues to explore these questions, and I am looking forward to new discoveries to be made. The results I have described so far are not the end of the game, but they show us where some of the action

in is in the grey areas between precise, logical, claims about trees, and the fuzzier, common sense reasoning that we apply when categorizing forests.

This talk about forests and trees brings me to some final words about the relations between individuals and groups. These relations are always important in life, and especially in interdisciplinary research activities like those I am engaged in at Utrecht University. The title "Computational Semantics and Artificial Intelligence" of my position highlights the connection that this university sees between Linguistics and Computer Science. No model of language can be complete without attention to how that model could be implemented on a computer. And vice versa: no computer algorithm that processes language can be complete if it does not provide insights into how the human mind supports our use of natural language. I am privileged to be working on the problems I mentioned in this talk within the stimulating intellectual atmosphere of the Utrecht Institute of Linguistics and the excellent study programs in Linguistics and Artificial Intelligence. These grounds have created here a lively and diverse community of scholars and students. It is only thanks to this fertile "forest", if I may call it so, of people and ideas that interdisciplinary teaching and research can thrive. My own research has been made possible thanks to collaborations with many individuals and organizations in this big community. At the end we see a body of scientific "outputs", or "fruits" – articles, computer systems, books, reports. Just imagine how many forests are used when putting our academic outputs on paper! But these cannot be evaluated at their true worth without considering the work of many different individuals, each of them with his or her special qualities and skills, energies, and feelings of hope, success and disappointment that we all experience. In this talk I have been referring to the collaborations I have had in different groups as if it was one uniform body of research. I wish I could say more here about the individuals behind these collaborations – teachers, students, and colleagues – each of them with her unique individuality. I am thankful to all the students, teachers and colleagues, at Utrecht University, in Israel and in other countries, who taught me and have made my work possible in these collaborations. I could not possibly mention all of them by name here.

However, of all the good people I have met in academia there are two to whom I am especially indebted, and whose memory I always cherish. These are my late friends and teachers: Tanya Reinhart, Professor of Theoretical Linguistics at Utrecht University, New York University and Tel Aviv University, and Remko Scha, Professor of Computational Linguistics at the University of Amsterdam. At the time that I was beginning to wander about in the big forest of linguistic research, Remko and Tanya were around to help me find my way in it. I feel honored to be able to continue some of the work they started. From a different personal angle, I am always grateful to my dear family: especially Yael and Dana, my late father Micha, and my mother Riwka who came a long way from Jerusalem to be here today.

I also wish to thank you ladies and gentlemen for your kind attention, and for the great respect you have paid me by attending this inaugural speech. I am looking forward to continuing my work at Utrecht University.

I have spoken.

I am grateful to Elly Koutamanis, Martin Everaert, Simone van Weteringen and Joost Zwarts for their help with the Dutch translation of this speech. - Y.W.S.

Bibliography

Bennett, M. (1974), *Some Extensions of a Montague Fragment of English*, PhD thesis, University of California Los Angeles.

Champollion, L. (2016), `Ten men and women got married today: Noun coordination and the intersective theory of conjunction', *Journal of Semantics* 33(3), 561-622.

de Vries, H. (2015), *Shifting Sets, Hidden Atoms: the semantics of distributivity, plurality and animacy,* PhD thesis, Utrecht University.

Dimitriadis, A. (2008), Irreducible symmetry in reciprocal constructions, in E. König & V. Gast, eds, *Reciprocals and Reflexives: Theoretical and Typological Explorations*, De Gruyter, Berlin, pp. 375-410.

Farkas, D. F. & de Swart, H. (2003), *The Semantics of Incorporation: from argument structure to discourse transparency*, CSLI Publications, Stanford.

Gärdenfors, P. (2004), *Conceptual spaces: The geometry of thought*, MIT press.

Grimm, R. M., Lee, C., Poortman, E. B. & Winter, Y. (2014), Evidence for non-existential readings of locative indefinites, in T. Snider, S. D'Antonio & M. Wiegand, eds, *Proceedings of Semantics and Linguistic Theory, SALT24*, pp. 197-212.

Hampton, J. A. & Winter, Y., eds (2017), *Compositionality and Concepts in Linguistics and Psychology*, Springer. Forthcoming.

Haspelmath, M. (2004), Coordinating constructions: an overview, in M. Haspelmath, ed., *Coordinating Constructions*, John Benjamins Publishing Company, Amsterdam/Philadelphia, pp. 3-39.

Hoeksema, J. (1988), `The semantics of non-boolean *and*', *Journal of Semantics* 6, 19-40.

Keenan, E. L. (1996), The semantics of determiners, in S. Lappin, ed., *The Handbook of Contemporary Semantic Theory*, Blackwell, pp. 41-64.

Keenan, E. L. & Faltz, L. (1985), *Boolean Semantics for Natural Language*, D. Reidel, Dordrecht.

Kerem, N., Friedmann, N. & Winter, Y. (2009), Typicality effects and the logic of reciprocity, in E. Cormany, S. Ito & D. Lutz, eds, *Proceedings of Semantics and Linguistic Theory, SALT19*, pp. 257-274.

Krifka, M. (1990), Boolean and non-boolean 'and', in L. Kálmán & L. Pólos, eds, *Papers from the Second Symposium of Logic and Language*, Akaemiai Kiado, Budapest.

Kruitwagen, I., Poortman, E. B. & Winter, Y. (2016), Reciprocal verbs as collective predicate concepts. *NELS 2016*, UMass, Amherst, to appear.

Link, G. (1983), The logical analysis of plurals and mass terms: a lattice theoretical approach, in R. Bauerle, C. Schwarze & A. von Stechow, eds, *Meaning, Use and Interpretation of Language*, De Gruyter, Berlin.

Mador-Haim, S. & Winter, Y. (2015), `Far from obvious: the semantics of locative indefinites', *Linguistics and Philosophy* 38, 437-476.

Montague, R. (1973), The proper treatment of quantification in ordinary English, in J. Hintikka, J. Moravcsik & P. Suppes, eds, *Approaches to Natural Languages: proceedings of the 1970 Stanford workshop on grammar and semantics*, D. Reidel, Dordrecht, pp. 221-242. Reprinted in Thomason, ed. (1974), *Formal Philosophy: selected papers of Richard Montague*, Yale, New Haven.

Nouwen, R. (2014), Plurality, in M. Aloni & P. Dekker, eds, *Cambridge Handbook of Semantics*, Cambridge University Press, Cambridge.

Partee, B. & Rooth, M. (1983), Generalized conjunction and type ambiguity, in R. Bauerle, C. Schwarze & A. von Stechow, eds, *Meaning, Use and Interpretation of Language*, De Gruyter, Berlin.

Payne, J. (1985), Complex phrases and complex sentences, in T. Shopen, ed., *Language Typology and Syntactic Description: complex constructions*, Vol. 2, Cambridge University Press, Cambridge.

Peters, S. & Westerståhl, D. (2006), *Quantifiers in Language and Logic*, Oxford University Press, Oxford.

Poortman, E. B. (2017), *Concepts and Plural Predication: The Effects of Conceptual Knowledge on the Interpretation of Reciprocal and Conjunctive Plural Constructions*, PhD thesis, Utrecht University.

Poortman, E. B. & Pylkkänen, L. (2016), `Adjective conjunction as a window into the LATL's contribution to conceptual combination', *Brain and Language* 160, 50-60.

Poortman, E. B., Struiksma, M., Kerem, N., Friedmann, N. & Winter, Y. (2017), When logic meets a prototype: Relational concepts and reciprocal reasoning. Unpublished ms., Utrecht University.

Reinhart, T. (1997), `Quantifier scope: how labor is divided between QR and choice functions', *Linguistics and Philosophy* 20, 335-397.

Scha, R. (1981), Distributive, collective and cumulative quantification, in J. Groenendijk, M. Stokhof & T. M. V. Janssen, eds, *Formal Methods in the Study of Language*, Mathematisch Centrum, Amsterdam, pp. 483-512.

Schwarzschild, R. (1996), *Pluralities*, Kluwer, Dordrecht.

Searle, J. R. (1990), Collective intentions and actions, in P. R. Cohen, J. Morgan & M. E. Pollack, eds, *Intentions in communication*, MIT Press, Cambridge, Massachusetts, pp. 401-416.

Toledo, A. (2015), *Semantic Modeling of Textual Entailment: Proof-Based Annotation in a Compositional Framework*, PhD thesis, Utrecht University.

Verkuyl, H. (1994), Distributivity and collectivity: a couple at odds, in M. Kanazawa & C. J. Piñon, eds, *Dynamics, Polarity and Quantification*, CSLI Publications, Stanford.

Winter, Y. (2001), *Flexibility Principles in Boolean Semantics: coordination, plurality and scope in natural language*, MIT Press, Cambridge, Massachusetts.

Winter, Y. & Scha, R. (2015), Plurals, in S. Lappin & C. Fox, eds, *Handbook of Contemporary Semantic Theory, Second edition*, Wiley-Blackwell, pp. 77-113.

Zwarts, J. (1997), `Vectors as relative positions: a compositional semantics of modified PPs', *Journal of Semantics* 14, 57-86.

Zwarts, J. & Winter, Y. (2000), `Vector space semantics: a model-theoretic analysis of locative prepositions', *Journal of Logic, Language and Information* 9, 169-211.