

Proceedings of the Paris Institute
for Advanced Study

Rethinking Compositionality

Sauerland, Uli ¹

¹ *Leibniz-Centre General Linguistics, Berlin, Germany*

TO CITE

Sauerland, U. (2026). Rethinking Compositionality. In *Proceedings of the Paris Institute for Advanced Study* (Vol. 27). <https://paris.pias.science/article/rethinking-compositionality>

PUBLICATION DATE

13/09/2026

ABSTRACT

Compositionality is the property of a communication system that the meaning of complex signals can be composed of the meanings of its parts. This paper first presents the distinction between trivial and non-trivial compositionality and how the latter is aligned with the concept of sentence in human language. For the account of non-trivial compositionality, the classical proposal in the field is based on the notion of function and its generalization in lambda calculus. But this approach would replace trivial compositionality and makes few, if any, predictions of how human compositionality is restricted. I present an alternative proposal for non-trivial compositionality analyzing it as the application of trivial composition followed by minimization. One advantage of the new, more restrictive, proposal is that it predicts that transitive verbs must be decomposed into different meaningful pieces, which had been observed before empirically, but not ever been predicted from a general theory.

Acknowledgement

I am grateful to Henry Davis, Benjamin Spector, Philippe Schlenker, my colleagues within the LeibnizDream project, the other fellows at the Paris Institute for Advanced Study during my writing residency in November 2025, and audiences in Prague and Paris for discussions of this work. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 856421). This article benefited from a fellowship at the Paris Institute for Advanced Study (France), with the financial support of the French State, programme "Investissements d'avenir" managed by the Agence Nationale de la Recherche (ANR-11-LABX-0027-01 Labex RFIEA+).

Introduction

One of the most important questions of linguistics is how language relates to other systems of communication, especially those found elsewhere in the animal world. Language has, since antiquity, been recognized as distinct from animal communication, but characterizing and demonstrating the distinctness of human language in formal terms has remained work in progress. One intuition has been that the human system allows an infinite array of distinct signals – a discrete infinity (Chomsky 2013) – while in the communication of non-human animal systems are either finite or have continuous components. An example of the former is the alarm call system of vervet monkeys as described by Seyfarth et al. (1980), with three calls transcribed as 'Leopard', 'Snake', and 'Eagle' warning of the presence of the respective predator. An example of the latter is the language of the bees as described by von Frisch (1974), with the length and direction of the waggle section of their dance indicating the distance and direction towards a food source. Looking at individual words, however, it is not apparent how human language differs from these systems other than by degree – the vocabulary of any human language is larger than the three calls of the vervet monkey. A fundamental difference, if there is any, must therefore be found in the way humans can put together individual signals into larger units – sentences – or putting it differently, how sentences can be *composed*. The question of where language stands in comparison to other communication systems is therefore intimately related to *compositionality*.

In this article, I introduce in the first section an existing proposal for how the compositional system of human language differs from compositionality in other


systems and some challenges for this view. In the second section, I then introduce a novel proposal and outline some of its advantages before concluding.

1. Compositionality

As we saw above, the question of how humans and other animals differ quickly takes us to the question of how humans can compose words into a sentence. Sentences have a specific grammatical form, but also are associated with a meaning. In the following, I am primarily interested in the composition of meaning. The focus on meaning is at least in part justified by the fact that sometimes a focus purely on form does not provide us with the full picture of what the units are. For example, in English, the word 'that' can mark the beginning of subordinate clauses but can also be omitted. As a result, the two sentence sequence in (1-a) can be in the spoken form the same (or at least very similar) to the single sentence in (1-b). However, the meanings of (1-a) and (1-b) are clearly distinct: only (1-a) allows us to infer that Mary is lying on the couch.

- (1) A. Mary was dreaming. She was lying on the couch.
 B. Mary was dreaming (that) she was lying on the couch.

To talk about meaning composition, first note that generally, the point of communication involves sharing information with each other, and sending messages or signals. For example, a red traffic signal (2-a) communicates the information that you should wait at this intersection. What you do with that information is, of course, up to you. Often, we find that communication is not one single signal, but complex, so made up of components that each carry meaning. For example, a red cyclist signal (2-b) has both the red color and the cyclist shape. The color communicates that you should wait, and the shape communicates that the signal addresses cyclists.



Here we have a first case of meaning composition. The two components of a traffic signal, shape and color, can vary independently, but they are always combined in the same way, namely by conjoining them with "and". The green cyclist (2-c) communicates that the signal is addressed at cyclists, AND that you are allowed to go.

A signal can also be complex in the temporal dimension. This is the case in language. I might say "Wave and smile" and thereby inform you that you should do two things. Beyond language, often we can't quite tell what the meaningful pieces are, actually. Think about music: individual notes don't seem to carry meaning, but a melody like 'so-so-so-miii' as in Beethoven's Fifth symphony or a crescendo or decrescendo as in the Ravi's Bolero seems to convey some meaning ([Schlenker, 2022](#)). And I can combine these two as indicated in (3):

(3) so-so-so-miii, SO-SO-SO-MIII, **SO-SO-SO-MIII**, SO-SO-SO-MIII, so-so-so-miii

So what's going on compositionally in this sequence? Ultimately, it seems to be also a conjunction. The meaning of melodies is not as conventionalized as that of traffic signals, but 'so-so-so-miii' might mean something like 'fate is knocking on the door' or 'that person has a Beethoven-hairstyle' depending on the situation. And the sequence of crescendo and decrescendo has a more fixed meaning, indicating that the cause of a sound passes by. Crucially, the two meanings are again composed by conjunction, so the little performance in (3) might serve to inform you that an opportunity to alter your fate or someone with a Beethoven-hairstyle passed by.

In human language, there are clear cases of compositionality that are not conjunctions. Negation is the prototypical case of non-conjunctive compositionality. What the message "It is not sunny" communicates has an empty intersection with what "It's sunny" does. The meaning of the conjunction in (4) clearly does not mean the same as "It's not sunny". What it actually means is

complex – in classical logic, (4) corresponds to a contradiction, but (4) actually doesn't seem to be completely uninformative. Alxatib et al. (2013) suggest that a repair strategy applies during meaning composition, resulting in the meaning that it's borderline sunny (see also Égré & Zehr 2018).

(4) It's sunny and it's not (sunny).

Not only negation, but also a sentence like "A girl is running" involves some non-trivial compositionality. If the pieces mean "There is a girl" and "Someone is running (something)", then the conjunction "There is a girl and someone is running" does not add up to "A girl is running". But generally, non-conjunctive composition seems to only occur within a single sentence. In (Sauerland, 2016), I suggest that we can formally characterize a sentence by the mode of meaning composition: Between sentence composition is always conjunctive, while within sentence composition can be something else. The sequence in (5-a) illustrates that it is impossible to deny a sentence once the sentence is complete. The sequence in (5-b) with the colon works, but the integrated intonation shows that it is one sentence, like example (5-c).

(5) a. #It's raining. That's not the case.

b. The following is not the case: It's raining.

c. That it is raining is not the case.

Combining meanings by conjunction is a type of compositionality, but in the field, it has recently been called *Trivial Compositionality* (Schlenker et al., 2023). This is justified because a conjunction just reflects what systems normally do when they remember information: Unless there is a conflict between the new information and the already remembered information, they remember both pieces of information. For example, suppose your doctor first gives you a rabies vaccine and then a flu vaccine. Your immune system, we hope, composes these into remembering immunity against both rabies viruses and against flu viruses. Your immune system does not need to know logic to do this.

The leading account of non-trivial composition is based on the concept of a mathematical function and dates back to the work of the mathematician and philosopher Frege (1879). In the basic case, the relevant function is a predicate that is either true or false of entities, i.e., a function from entities to truth values. In the function-based model, predicates would be used to capture the meaning of intransitive verbs and adjectives, such as *run* and *round*, respectively. Specifically, *run* is analyzed as a function mapping only those entities to *true* that are running. In the function analysis, non-trivial composition is captured by function composition. Negation can also be analyzed as a function, namely a function

mapping one truth value to the other, *true* to *false* and vice versa. If we furthermore assume that *raining* is either true or false, the meaning of '*It's not raining*' is captured by applying the negation function to the truth value of *raining*. If complete sentences must be truth values and therefore are never functions, it follows that function application can only apply internally to a sentence, while sentences must be composed trivially.

One further prediction of the function analysis is that meanings can be unsaturated. Because functions cannot be sentences, phrases denoting functions are predicted to require one or more arguments to make a complete sentence. This seems to be correct, at least initially: intransitive verbs like *run* are not complete as sentences, but only when an argument is present, as in *Sarah runs* are. And some transitive verbs, such as *like* even require two arguments: '*Sarah likes*' is incomplete, but '*Sarah likes Kim.*' is a well-formed sentence.

(6) a. *Sarah likes.

b. Sarah likes Kim.

While the basic intuition of the function-based analysis of non-trivial composition is easy to grasp, a more complete model is beyond the scope of this paper. The first model capturing a sizable fragment of English was provided by Montague (1974). Textbooks such as (Chierchia & McConnell-Ginet, 1990) and (Heim & Kratzer, 1998) provide accessible introductions. One important technique used in this line of work is the use of *Lambda-Calculus* (Church, 1932), which recursively generalizes the notion of function. Importantly, lambda-calculus is a very powerful model of computation, so powerful that the so-called untyped version of lambda-calculus includes all formally computable functions. This leads to the question of whether the lambda-calculus approach does not overgenerate as a model of languages (Janssen, 1983; Zadrozny, 1994). Some might say here that the use of lambda-calculus is just the mathematical meta-language used, not an actual model of human compositionality to derive. Two problems with this view are the following: on the one hand, that it does not answer the question of how to account for restrictions of compositionality, and on the other hand, it is inconsistent with the application of lambda calculus in the field: For example, the restriction of *like* to two arguments mentioned above is not predicted by the untyped lambda calculus since it allows functions that take either a single argument or a pair of arguments. Therefore, linguistic work since Montague has adopted a restriction to so-called *Typed* functions where functions taking one and functions taking two arguments belong to different types. So, in this case, the typed lambda calculus is taken to provide a better model of language than the untyped one. But for other restrictions in language, it is not clear how to account for them in a similar way within the function-based compositional system,

specifically, we will discuss below the requirement that complex predicates must be decomposable.

Another potential difficulty of the function-based analysis is that composition by function is discontinuous with composition by conjunction: a function cannot be conjoined with another function or with its argument. So, function application would need to be a trait that developed in a single jump. Such an evolutionary jump would be plausible if there was a single mutation that resulted in the ability to do function application. Nowak et al. (2000) present a model that they claim could explain such an evolution. They build on the intuition that a compositional language can be easier to learn because only the meanings of the units need to be learned, not every meaningful signal. For a species without composition, the number of word learning events during an individual's lifetime and the probability of learning a word in an encounter determine a maximum number of words $n_{\max}$ that can be learned by an individual. They compare this with a mutant with compositionality just sufficient to combine one noun and one verb. They assume that the probability of learning one of the two words from a communicative encounter is lower than it is for the original species learning the one word. But because once mutants learned n -many nouns and m -many verbs, the number of different states they can communicate is up to the product $n \cdot m$, they argue that composition can confer an evolutionary advantage if it increases the mutants' chance of success to be able to distinguish a large number of states.


But Nowak et al.'s model did not take into account trivial compositionality. Therefore, they assume that mutants can distinguish $n \cdot m$ states while non-mutants can distinguish $n + m$ many states. The right comparison, though, would be to assume that both non-mutants and mutants communicate two words, which the mutants combine with function composition, while the non-mutants apply conjunction. Assuming that the order does not lead to different meanings for the conjunction, non-mutants are predicted to be able to distinguish $(n + m)^2 / 2$. But this computation shows that non-mutants are for $n, m > 0$ predicted by Nowak et al.'s model to be able to communicate a greater number of states than mutants once trivial compositionality is taken into account.

In sum, the dominant view of the compositionality of human language for close to 150 years of research has been the function-based view. I have briefly mentioned two problems of the account, namely that it predicts only a few constraints on compositionality in language and requires an exceptional evolutionary jump from trivial to function-based compositionality.

2. Minimization Based Composition

My main new idea is to develop a new model of what non-trivial composition is, and that primarily in terms of restrictiveness does better than the function-based model. The central idea is that composition involves a process of minimization on top of conjunction.

Here is an illustration: Take "A girl is running", with its pieces "a girl" and "it's running". Let's say "A girl" requires that a girl be there. And let's say that "running" requires that someone be running. If we conjoin those two, that doesn't require that it be the girl that's running, but it allows it.



(7)

(picture generated by the author using picture generation software)

But the imaginable scenes where it's the girl that is running contain only one entity, while the others contain two entities, the girl and a runner. So, if we add minimality to the meaning, we require that it's the girl who is running.

The notion of minimality we apply here can be formalized in different ways at this point. One way, following Barwise (1989), would be to assume that there are abstract entities – states (or scenes or scenarios) or events – that have other entities as parts, i.e., $x \sqsubseteq s$ for entity x is a part of states. Then the minimization operator **min** could apply as in (8) to restrict that set to those possible states that have the smallest number of entities as parts necessary to satisfy the sentence (Guasti et al., 2023).

$$(8) \text{ min } \{s \mid \Diamond[\exists x \sqsubseteq s \text{ girl}(x) \wedge \exists x \sqsubseteq s \text{ running}(x)]\}$$

A slightly different way would be to apply minimization to assignment functions (specifically, the cardinality of the image of the function g). Then, if it is

logically possible that $g(x_1)$ and $g(x_2)$ have the same value, the minimal assignments will be those where $g(x_1)=g(x_2)$.

$$(9) \text{min } \{g \mid \Diamond[g(\text{girl}(g(x_1))) \wedge \text{running}(g(x_2))]\}$$

For the following, the specific implementation of minimization does not matter. I adopt (8) for the following, but leave it for future work to determine whether both make the same predictions or whether there are cases where they can be distinguished.

With the **min** operator, we state non-trivial composition at least in the running-girl case, as **min** applying after trivial composition. If this were to generalize to all cases of non-trivial composition in language, it would make the evolutionary path from trivial to non-trivial composition less abrupt. Evidently, **min** itself combines in (8) and (9) via function application with an argument, but general function application might not be necessary to explain non-trivial compositionality in human language. Minimization itself could also apply to non-complex items, such as *breathing*, to provide the inference that if there is *breathing*, there must also be an individual with a lung present. Consequently, minimization might support reasoning and therefore be present independently of language. The step to non-trivial compositionality would then reduce to the ability to apply minimization to complex sequences, which themselves are composed trivially.

The minimization idea makes different predictions from the function-based approach. While some of these seem to be wrong at first glance, I think, though, that both show hidden virtues of the more constrained system of composition than lambda-calculus. The first concerns transitive relations. As I will show now, most transitive verbs are impossible to capture as primitive elements in the system. Consider the verb "fix" as in (10). Let us assume that (10) consists of at least three different pieces that might be combined into a sentence as shown in (10).

(10) A robot is fixing a computer

[A-robot]

[is-fixing a-computer]

Let's assume that the three pieces of this sentence mean respectively: "there is a robot", "there is something fixing something", and "there is a computer", which is similar to the case in (7). Minimization could then apply either to the smaller constituent – the verb and the object – first or only to the full sentence. The first option yields the minimal states that contain something fixing something and a computer. These are states where a computer is fixing itself because in these

states only a single entity is present – i.e., minimization results in reflexivization of the verb. When minimization must apply at the level of the complete sentence again, it cannot undo the reflexivization. Minimization at this level yields the minimal states that contain both a robot and, furthermore a computer that is fixing itself. Since robots that are also computers are possible, the only minimal states will be those containing one of the robot-computers fixing itself. What if minimization does not apply at the lower level? The outcome is the same because it is logically possible that the same entity serves as the argument of predication in four positions: the two of the nouns *computer* and *robot*, and both arguments of the verb *fix*. Furthermore, for any different structure from the one shown in (10), the same meaning is predicted: states of a single robot-computer fixing itself. A slight variation arises if the two nouns are logically incompatible, as in (11).

(11) A girl is fixing a computer.

The logical incompatibility arises in (11) because the lexical meanings of *girl* and *computer* are contrary, one requiring animacy and the other inanimacy. But logical incompatibility could also derive from an implicature or contrastive inference that, for example, would add to *computer* the inference that the argument is not a robot (or at least not certain to be one). In both cases, the order of composition affects the final meaning as follows: if we minimize the conjunction of *computer* and *fix* before composing the result with *girl*, the minimal states contain a self-fixing computer and a girl, while starting with the composition of *girl* and *fix* leads us to states containing a self-fixing girl and a computer. And applying minimization only once to the conjunction of *computer*, *fix*, and *girl*, results also in minimal states containing both a computer and a girl, where furthermore either the computer is fixing the girl, the girl fixing the computer, the girl fixing itself, or the computer fixing itself. While this set of states finally includes the states of a girl fixing a computer that (11) is true of, it still predicts that (11) should be true in states where it is intuitively false. Finally, one might assume that *fix* requires its two arguments to be distinct, i.e., means "an *x* fixes a *y* with $x \neq y$." Such an assumption would be incompatible with self-fixing robots, but even so, it would for (11) not quite predict the correct meaning since it would still allow both a girl fixing a computer and a computer fixing a girl. In sum, there is no way to derive the correct meaning for (11) using minimization and conjunction, assuming that it is decomposed into only three elements.

There are two cases, though, where it is possible to derive the correct semantics for a relational concept. One case is when the arguments are logically required to belong to different categories. For example, assume that there is a concept *anim-fix* that requires its object to belong to the class of inanimates, but furthermore

requires its subject to belong to the class of animates. If this verbal relationship is composed with *computer* and *girl*, regardless of the order, the girl will have to be the subject and the computer will have to be the object. Candidates for verbs expressing such relations are especially verbs relating concrete and abstract entities, such as *claimor wonder*. The other potential class is relations that require at least two entities and a reciprocal relationship between them, such as *differs*, which logically entails that its two arguments must be distinct. For example, for 'A differs from B', the meaning minimization derived is correct and predicted to be equivalent to 'A and B differ'. For example, with three nouns such as 'A differs from B and C' predicting the correct meaning requires the assumption that the referent of the noun phrase *B and C* must be treated as a single entity. Of these two cases, the first case can provide an initial way to represent relational concepts, but ultimately, I will favor a solution that also requires the second case. In earlier work ([Guasti et al., 2023](#)), we pursued a third approach to allow relational concepts to be expressed – namely to stipulate a finite set of functions – but I will not pursue this in this paper.

For the first approach to the relational concept problem, consider example (10) again. Assume for this approach that the transitive verb *fix* involves actually a third entity of a different category X, it is possible to represent the relational concept via two other relations. Namely, relation R1 that relates an entity and something of category X and relation R2 that relates also an entity and something of category X. If the constituencies are as shown in (12) and minimization only applies to the maximal constituent in (11), the predicted interpretation is that an x of category X stands in relation R1 to a robot and also in relation R2 to a computer.

(12) [[a-girl R1] [R2 a-computer]]

It is easy to see that the mechanism in (12) allows us to create any relation: namely for the *fix*-relation, we could assume that category X is that of ordered pairs $\langle x_1, x_2 \rangle$ such that x_1 stands in the *fix* relation to x_2 and we define R1 as an entity being the first member of a pair x and R2 as an entity being the second member of a pair x . However, there are at least two problems with (12). For one, (12) brings back functions via the assumption that sets of ordered pairs are part of the ontology. One could possibly address this by regarding category X as that of *events* and R1 and R2 as thematic relationships between events and individuals, as in the Neo-Davidsonian tradition in linguistics ([Barwise, 1989](#)). A second potential problem for (12) concerns the predicted minimization was to apply to either of the two subconstituents. For example, minimization of $*[a \text{ girl } R1]*$ might be understood to require the ordered pair to be the constant $\langle x, x \rangle$ where x is a robot, since a pair $\langle x, y \rangle$ with $y \neq x$ seems to contain an extra

individual. But we could also assume that the pairs are not transparent for minimization in the same way as we would assume for plural entities.

While (12) might look like a kluge, it is an attractive result from the perspective of linguistic theory. Linguists have noticed for a long time that frequently transitive verbs seem to be made up of pieces, or, looked at differently, verbs must be decomposed into pieces. Several linguistic theories propose that verbs must be decomposed across languages into smaller units under certain conditions, even if there is no morphological evidence for the decomposition ([Levin & Hovav, 1995](#); [Pietroski, 2018](#); [Ramchand, 2008](#)). However, this decomposition has usually been argued for empirically, while it needed to be stipulated as a constraint with the function-based theory of semantics generally assumed (Pietroski's account is slightly different from the function-based theory, but not fundamentally so; [Icard & Moss, 2023](#)). The account sketched in (12) is different: it is predicted by a general theory. We now need to verify that it can indeed account for the evidence that has motivated the decomposition of transitive and ditransitive verbs in linguistics.

Before approaching the evidence, I introduce a slightly different way of capturing relations that seems more promising to me than the one in (12). In particular, it seems unsatisfactory to introduce pairs of both agent and patient in the meaning of each nominal constituent. Singletons (and) provide a formally more straightforward way of capturing entities of a different category based on individuals than pairs do. And it would still be possible to assume that singletons like those suggested for pairs above are not transparent to minimization, such that both a state containing x and and one containing x and with $y \neq x$ can simultaneously be considered minimal. Actually, we at least require two different singletons based on the same individual x for the following, i.e., and $\langle \rangle$. Then a better solution to representing a transitive relation requires a third sortal type that invokes the pairing. Specifically, we might adopt the following new definitions for R1 and R2: R1 requires that there be an entity x and the singleton or $\langle \rangle$ and an entity y and the singleton or $\langle \rangle$. We furthermore assume that R2 means that there is a singleton b based on an individual z and a different singleton a such that the pair $\langle a, b \rangle$ stands in relationship R' . We furthermore assume that R' is a non-reflexive relationship on the set of singletons such that if and only if x fixes y , there are singletons a and b based on x and y such that the relation R' contains the pair $\langle a, b \rangle$. Since both and $\langle \rangle$ are singletons based on x , R' can represent a reflexive case of fix without being itself reflexive. With these assumptions, the structure in (13) captures a transitive relation as follows. The constituent [R1 a-computer] is satisfied by minimal states containing y and a singleton based on y , e.g., . At the next level, R2 requires a pair of two singletons $a \neq b$ with $\langle a, b \rangle$ in R' and an individual z that b is based on. The minimal way to satisfy the conjunction at this level is to assume that individual z is identical to the computer y and that, therefore, b is a singleton based on a

computer *y*. Finally, at the highest level, this combines with a condition that requires a girl *x* and a singleton *a* based on *x*. Since we assume that computers and girls are incompatible properties, two distinct individuals are required in the minimal states. Furthermore, there must also be two distinct singletons, but since one must be based on the girl *x* and the other on the computer *y*, the requirement that the two be distinct coming from *R2* is automatically satisfied. And because the earlier minimization already required that of the pair $\langle a, b \rangle$ in *R'*, the second member *b* must be based on the computer, the first member must be based on the girl.

(13) [[a-girl *R1*] [*R2* [*R1* a-computer]]]

The structure (13), therefore, provides a second way to account for a relational concept, for example, like (10), where the nominal predicates are logically incompatible. Note that, if a reflexive interpretation is possible, as it might be in (10), it is still required. The irreflexive relationship *R'* in such a case would need to be satisfied by a pair of singletons such as $\langle a, b \rangle$ that are both based on the same individual.

The schema in (13) still allows different ways to distribute aspects of the verb meaning over the *R2* and the two occurrences of *R1*. Also, the two occurrences of *R1* could also be two different conditions, especially for non-symmetric predicates. While (13) illustrates the general strategy, the structure in (14) is close to the proposals for decomposition mentioned above by Levin et al. (1995) and Ramchand (2008).

(14) []
 [a-girl *do*] [*cause*]
 [functioning a-computer]

To apply (13) to (14), we assume the two predicates *intent* and *function* relate entities to monadic events, i.e., *do* requires an entity *x* and an event *e* *x* carries out, while *function* requires an entity *y* and an event of functioning by *y*. But *cause* must be a relation between two events, namely one event *e* and a second event *e'* that involves individual *y* such that *e* causes *e'*. The meaning that results is that a girl does something that causes a computer to be functioning.

We can now see how this account predicts some properties that have motivated decomposition in the prior literature. One argument for the decomposition of verbs is derived from looking at the meaning of examples such as (15) (von Stechow 1996 and others). Specifically, observe that the use of *again* does not require the computer to have been fixed by some entity in the past, only that it

must have been working before it broke. It is sufficient for the truth of (15) that the girl returned it to the functioning state that it was in before. This shows *again* can apply to only a part of the relation that *fix* expresses, namely the result of the computer being in a functioning state.

(15) A girl fixed the computer again.

The interpretation of (15) can be captured within the account given in (14) by assuming that *again* applies only to the phrase that expresses that the computer is functioning as in von Stechow's original account.

The second argument in favor of decomposition relates to the question of how the abstract structures shown above in (12) to (14) relate to the written or spoken form of the sentence in (11). There must be some mapping, but at this point, it is debated how this mapping is determined. I adopt the minority view that the mapping is *Meaning First* following (Sauerland & Alexiadou, 2020), i.e., that the representation (14) is directly mapped to its articulation. As part of that, the complex sequence of *cause* and *function* is mapped to the verb *fixed*. The verb in English does not show any similarity to the two pieces it is claimed to articulate, and in particular, *cause* is not articulated directly. But the Meaning First approach assumes that the same structure as (14) underlies sentences expressing the same meaning in other languages. So it is predicted that some other languages should articulate *cause* and that even transitive predicates, such as *see* should show a bipartite structure. The Salish languages are examples of such languages where all transitive verbs must contain a morpheme called the *Transitivizer*, except for two exceptions (*leave* and *eat*) as Davis and Matthewson (2009: p. 1101) report. They specifically give examples from St'át'imcets (ISO code: lil), such as the following:

- (16) a. mayš 'to get fixed' (intransitive) máyš-ən 'to fix something'
(transitive)
- b. ʔacx 'to get seen' ʔacx-ən 'to see someone/something'

Similar data are reported from other unrelated languages, for example, from Tamil (ISO code: tam) by Paramasivam (1977).

According to the approach we have proposed, transitive verbs must consist of multiple pieces. The function-based approach predicts that transitive verbs can be decomposed but can also be primitive. So the existence of languages like St'át'imcets does not refute the function-based approach. One way of distinguishing the two approaches is to look at children learning a language. Specifically, in other work of mine, we argue that sometimes children, unlike adults, pronounce pieces of complex words that adults leave unpronounced if there are morphological resources to do so (Guasti et al., 2023; Sauerland et al.,

2024). In the former paper, we report data from French children that indicate that they overuse causative marking as illustrated in (17). In this specific case, the child seemingly corrects its first utterance by inserting a causative verb.

(17) On va le chacher. On va le faire cacher (Madeleine 2;2).

We will it hide we will it make hide

'We will hide it. We will hide it.'

The developmental path indicated by (17) is not predicted by the function-based approach without introducing additional conditions. This approach predicts binary functions to be available as primitives, and unless children receive contrary evidence, they should analyze transitive verbs like *hide* in French as binary functions. The approach we developed, on the other hand, predicts that children and adults should be forced to analyze transitive verbs as composed out of multiple pieces. To then sometimes articulate some of these pieces, even when this doesn't occur in the adult grammar, is a natural error for children to make. For example, the children might simply assume that they did not hear fully what the adults were saying.

So far, we discussed cases where minimality based compositionality predicts restrictions while lambda-calculus-based approaches do not. In (6) above, I mentioned that the typed lambda calculus has been claimed to predict ungrammaticality in some case where an argument must be filled. The minimization bases proposal has no equally direct way of capturing such restrictions. But the relevant restrictions differ across languages. For example, while English can drop the object *something* in (18-a), the literal translation of the sentence requires the object to be present in Ewe (18-b) (ISO code: ewe). The variation across languages seems unexpected, given the function-based account that ties the obligatory presence of an argument to the semantic concept of a verb's meaning.

(18) a. The bird ate (something).

b. Xevi-a ɖu *(nú)

bird-the eat thing

In sum, we have seen that the conception of compositionality based on minimality predicts a restriction on primitive transitive verbs (or relational concepts) that the function-based approach does not predict. As a result, relational concepts are predicted to be necessarily decomposed into several pieces. The

predicted decomposition can account for arguments that have been made already in favor of a decomposition of transitive verbs.

Conclusions

Understanding the compositionality of human language should lead to deep insights into human nature and its place within biology. Outside of human language, trivial compositionality by conjunction is perhaps all we find. Within human language, trivial compositionality is also all that exists above the sentence level. But within the sentence, the compositionality must be richer than the trivial kind. Existing work is primarily based on mathematical functions in the generalized form of lambda-calculus that replace trivial composition. This paper has advanced a novel approach based on minimization that applies in addition to trivial composition.

The two proposals differ in their restrictiveness. Lambda-calculus as a universal theory of computation predicts only trivial restrictions on human compositionality. The novel proposal, however, predicts restrictions. Specifically, we show in section 2 that it predicts relations between two arguments of the same type cannot be captured as primitives unless they are irreflexive. This entails that transitive verbs can never correspond to the primitive units of meaning but must be decomposed out of multiple meaningful pieces, and also entities of a different category are required. As discussed in section 2, the prediction is borne out – transitive verbs must always be decomposed. The decomposition of verbs has been observed empirically for several decades, but an explanation of why decomposition must occur has been elusive in the past.

Bibliography

- Alxatib, S., Pagin, P., & Sauerland, U. (2013). Acceptable contradictions: pragmatics or semantics? A reply to Cobreros et al. *Journal of Philosophical Logic*, 42, 619–634. <https://doi.org/10.1007/s10992-012-9241-7>.
- Barwise, J. (1989). *The Situation in Logic*. CSLI Publications.
- Chierchia, G., & McConnell-Ginet, S. (1990). *Meaning and Grammar: An Introduction to Semantics*. MIT Press.
- Chomsky, N. (2013). The Dewey lectures 2013: What kind of creatures are we? Lecture I: what is language? *Journal of Philosophy*, 110(12), 645–662. <https://doi.org/10.5840/jphil2013110121>.
- Church, A. (1932). A set of postulates for the foundation of logic. *Annals of Mathematics*, 33, 346–366.
- Davis, H., & Matthewson, L. (2009). Issues in Salish syntax and semantics. *Language and Linguistics Compass*, 3, 1097–1166. <https://doi.org/10.1111/j.1749-818X.2009.00145.x>.
- Frege, G. (1879). *Begriffsschrift. Eine der arithmetischen nachgebildete Formelsprache des reinen Denkens*. Neubert.
- Frisch, K. von. (1974). Decoding the language of the bee. *Science*, 185, 663–668. <https://doi.org/10.1126/science.185.4152.663>.
- Égré, P., & Zehr, J. (2018). Are paps preferred to gluts? A closer look at borderline contradictions. In E. Castroviejo & L. McNally (Eds.), *The Semantics of Gradability, Vagueness, and Scale Structure: Experimental Perspectives*. Springer. https://doi.org/10.1007/978-3-319-77791-7_2.
- Guasti, M. T., Alexiadou, A., & Sauerland, U. (2023). Undercompression errors as evidence for conceptual primitives. *Frontiers in Psychology*, 14. <https://doi.org/10.3389/fpsyg.2023.1104930>.
- Heim, I., & Kratzer, A. (1998). *Semantics in Generative Grammar*. Blackwell.
- Icard, T. F., & Moss, L. S. (2023). A simple logic of concepts. *Journal of Philosophical Logic*, 52(3), 705–730. <https://doi.org/10.1007/s10992-022-09685-1>.
- Janssen, T. (1983). *Foundations and applications of Montague Grammar* [Doctoral,]. University of Amsterdam. <https://eprints.illc.uva.nl/id/eprint/1842/>.
- Levin, B., & Hovav, M. R. (1995). *Unaccusativity: At the Syntax-Lexical Semantics Interface*. MIT Press.
- Montague, R. (1974). The proper treatment of quantification in ordinary English. In R. Thomason (Ed.), *Richard Montague: Selected Papers*. Yale University Press.
- Nowak, M. A., Plotkin, J. B., & Jansen, V. A. (2000). The evolution of syntactic communication. *Nature*, 404, 495–498.
- Paramasivam, K. (1977). *Effectivity and causativity in Tamil* [PhD Thesis,]. University of Chicago.

- Pietroski, P. M. (2018). *Conjoining Meanings: Semantics Without Truth Values*. Oxford University Press.
- Ramchand, G. (2008). *Verb Meaning and the Lexicon: A First-Phase Syntax*. Cambridge University Press.
- Sauerland, U. (2016). On the Definition of Sentence. *Theoretical Linguistics*, 42, 147–153. <https://doi.org/10.1515/tl-2016-0007>.
- Sauerland, U., & Alexiadou, A. (2020). Generative Grammar: A Meaning First Approach. *Frontiers in Psychology*, 11. <https://doi.org/10.3389/fpsyg.2020.571295>.
- Sauerland, U., Bassi, I., & Bill, C. (2023). An algebra of thought that predicts key aspects of language structure. *Proceedings of Sinn Und Bedeutung*, 27, 589–608. <https://doi.org/10.18148/sub/2023.v27.1088>.
- Sauerland, U., Meyer, M.-C., & Yatsushiro, K. (2024). The cum-sine pattern in German child language: An argument for antonym decomposition. *Language Acquisition*, 23, 1–13. <https://doi.org/10.1080/10489223.2024.2332452>.
- Schlenker, P. (2022). *What It All Means: Semantics for (Almost) Everything*. The MIT Press. <https://doi.org/10.7551/mitpress/12128.001.0001>.
- Schlenker, P., Coye, C., Leroux, M., & Chemla, E. (2023). The ABC-D of animal linguistics: Are syntax and compositionality for real? *Biological Reviews*, 98(4), 1142–1159. <https://doi.org/10.1111/brv.12944>.
- Seyfarth, R. M., Cheney, D. L., & Marler, P. (1980). Monkey responses to three different alarm calls: Evidence of predator classification and semantic communication. *Science*, 210, 801–803. <https://doi.org/10.1126/science.7433999>.
- Stechow, & Arnim. (1996). The different readings of wieder ‘again’: A structural account. *Journal of Semantics*, 13, 87–138.
- Zadrozny, W. (1994). From compositional to systematic semantics. *Linguistics and Philosophy*, 17, 329–342. <https://doi.org/10.1007/BF00985572>.