

Super Linguistics: an Introduction*

Pritty Patel-Grosz^a, Salvador Mascarenhas^{bce}, Emmanuel Chemla^{cde}, Philippe Schlenker^{bcd}

Abstract We argue that formal linguistic theory, properly extended, can provide a unifying framework for diverse phenomena beyond traditional linguistic objects. We summarize key achievements of syntax, semantics and pragmatics, pertaining to constraints on possible structures and meanings, the role of variables, and the typology of linguistic inferences. We then show how syntax and semantics can help illuminate diverse non-standard objects. First, we discuss structure and meaning in animal communication, beginning with birdsong, which has been analyzed as exhibiting sophisticated finite-state grammars. This is followed by a discussion of the semantics of primate calls – where we state precise hypotheses about their lexical semantics, argue that no non-trivial rules of call composition are currently needed, and suggest that rules of competition among calls might play an important pragmatic role. Turning to human language, new developments make it possible to ask how iconicity interacts with logical semantics in spoken language, in sign language and in gestures. Going beyond language, syntax and semantics have recently been extended further, to pictorial meaning, to the syntax and semantics of visual narratives, and even to the syntax and semantics of music and dance, using a mix of new methods and standard tools of formal linguistic theory (e.g. variables). Finally, semantics is increasingly used to illuminate constraints on logical and non-logical concepts in the 'language of thought', which is investigated independently of its superficial aspects supported by language, leading to new accounts of reasoning. The aim of this overview is to show how the application of formal linguistic methodology (and methodology inspired by formal linguistics) to non-standard objects yields non-trivial insights, thus opening up a promising line of inquiry.

Keywords: formal theoretical linguistics, super linguistics, primate linguistics, sign language grammar, gestural grammar, picture semantics, visual narratives, dance, music, iconicity, concepts, reasoning

Word count (main text): **22, 279 words** [entire document, including footnotes, references, and table of contents: 29, 374]

* Authors' affiliations:

a: University of Oslo

b: Institut Jean-Nicod (ENS - EHESS – CNRS)

c: Département d'Etudes Cognitives, Ecole Normale Supérieure

d: PSL Research University

e: LSCP (ENS - EHESS - CNRS)

f: Department of Linguistics, New York University

1	Introduction	4
2	Achievements of formal linguistics	4
2.1	Syntax: where does human language fit in the hierarchy of formal languages?	5
2.1.1	<i>From the syntax of human language to the syntax of non-standard objects</i>	5
2.1.2	<i>Constituency and Grouping</i>	7
2.2	Semantics: what are the elementary operations by which language produces information?.....	8
2.2.1	<i>Expressive power and constraints on possible meanings</i>	8
2.2.2	<i>Variables</i>	9
2.2.3	<i>Compositionality</i>	10
2.3	Pragmatics: the inferential typology of language.....	11
2.3.1	<i>Implicatures</i>	11
2.3.2	<i>Presuppositions</i>	11
2.3.3	<i>Supplements</i>	12
2.4	Psycholinguistics of meaning: processing and acquiring meaning.....	12
2.5	Interim summary.....	12
3	Formal approaches to animal communication.....	12
3.1	The formal approach.....	12
3.2	Animal syntax: the example of birdsongs	13
3.3	Animal meanings: the example of monkey communication.....	14
3.4	Call meaning vs. pragmatic enrichment: Campbell's monkeys	14
3.5	Call combination vs. context change: Titi monkeys	16
3.6	New questions about the evolution of animal meanings	17
4	Integrating logical and iconic semantics.....	18
4.1	Pictorial semantics	18
4.2	Logical visibility and iconicity in sign language.....	19
4.3	Logical and grammatical structures in gestures.....	20
5	Linguistics beyond words I: iconic elements within the inferential typology	22
5.1	Iconic enrichments in the typology of linguistic inferences: sign with iconicity vs. speech with gestures.....	23
5.2	Iconic replacements in the typology of linguistic inferences: replicating the typology without words	26
5.2.1	<i>Iconic implicatures</i>	26
5.2.2	<i>Iconic presuppositions</i>	26
5.2.3	<i>Going further</i>	27
6	Linguistics beyond words II: visual and musical narratives.....	28
6.1	Visual narratives.....	28
6.1.1	<i>Syntax of visual narratives</i>	28
6.1.2	<i>Semantics of visual narratives</i>	29
6.2	Musical narratives.....	30
6.2.1	<i>Basic music syntax</i>	30
6.2.2	<i>Basic music semantics</i>	33
6.2.3	<i>Adding variables</i>	36
6.3	Dance	37

6.3.1	<i>Dance syntax</i>	37
6.3.2	<i>Dance semantics</i>	39
6.4	Interaction between visual and musical narratives.....	42
7	Semantics without phonology	45
7.1	Content concepts and logical concepts.....	45
7.2	Alternatives and attention	46
8	Conclusion	48
	References	49

1 Introduction

Semiotics was initially conceived as a general theory of natural signs (see Morris 1938), be they human or animal – or neither, as in the case where spots "mean" measles (Grice 1957:377). The focus has largely been on their meaning, though Morris (1938:6-7) explicitly lays out the distinction between a *syntactical*, *semantical* and *pragmatical dimension of semiosis*, in connection with the goal of achieving a general theory of signs. By contrast, formal linguistics investigates form and meaning in human language with sophisticated tools from formal language theory and logic, with rich descriptive and experimental data from numerous languages (e.g. Maienborn et al. 2011, Gutzmann et al. 2021). Contemporary formal linguistics consists of two branches: the study of form (phonology, morphology and syntax), and the study of meaning (semantics and pragmatics, broadly called 'formal semantics'). The goals of formal linguistics, as opposed to the more general aims of semiotics, have traditionally been restricted to human language, in fact to spoken languages, followed by a gradual and relatively recent integration of sign languages phonology, then morphology and syntax, and more recently semantics (e.g. Stokoe 1960, Klima and Bellugi 1979, Sandler and Lillo-Martin 2006, Pfau et al. 2012, Quer et al. 2021).

Several recent extensions of formal linguistics, which in part amount to integrating formal linguistics and semiotics into a kind of formal semiotics (see Greenberg 2021, Schlenker 2019b), suggest that it might contribute to a general theory of signs as originally envisioned by Morris (1938). In so doing, formal linguistics may come to encompass aspects of human communication (such as gestures, facial expressions) that were traditionally left outside its purview, as well as non-linguistic systems such as animal communication, visual and pictorial narratives, music and dance. It might even extend to reasoning and concept manipulation, where the form (the 'language of thought') must be inferred on indirect grounds. The aim of pursuing a general theory of signs that incorporates the tools and technologies of formal linguistics is to generate results to which formal and experimental sciences will make sustained reference in the future.

In this introduction, we sketch initial results of this broader field of 'Super Linguistics' (using 'super' in the original Latinized sense 'beyond', e.g. 'beyond language'), with special attention to structure and meaning phenomena. In brief, Super Linguistics makes it possible to investigate diverse representational forms with increased precision; it offers an interesting typology of structure and meaning operations in nature; and it draws surprising new connections among these domains, including between human words, animal concepts, body movement and music and general reasoning abilities. Before we lay out these results, however, we must first introduce key achievements of formal linguistics; each will serve to illuminate empirical domains that belong to Super Linguistics.

2 Achievements of formal linguistics

The Chomskyan revolution of the 1950's and 1960's showed that human languages can be analyzed as formal languages; this was a major step towards finding rules that systematically explain how sentences of natural language are put together and convey information. The next step, pioneered by Montague (1970a, 1970b, 1973), was to treat human language as a formal *interpreted* language, and thus to systematically associate syntactic rules with semantic (i.e. meaning) rules. But what is a semantic rule? The key insight, from the philosophy of language (e.g. Davidson 1967), was that *to know the meaning of a sentence is to know under what conditions it is true*. As a result, the key problem was to systematically associate truth conditions to sentences of arbitrary complexity. This reduced the problem to one that had


been solved much earlier for (far simpler) logical languages: Tarski (1935, 1944) had shown how to define systematic truth conditions for mathematical logic. The challenge was thus to assess whether, and how, these logical methods could be extended to human languages.

This multi-generation enterprise has yielded systematic accounts of form and meaning in human language. Four questions are particularly relevant for Super Linguistics: (1) where does human language fit within the hierarchy of formal languages? (syntax) (2) how is information about the world encoded in words? (semantics in the strict sense) (3) how is semantic information enriched by reasoning based on the beliefs of the speech act participants? (pragmatics) (4) how are semantic and pragmatic operations acquired and processed in real time? (psycholinguistics of meaning).

2.1 Syntax: where does human language fit in the hierarchy of formal languages?

2.1.1 From the syntax of human language to the syntax of non-standard objects

Generative syntax started with a bang with the publication of Chomsky's *Syntactic Structures* in 1957. Treating human language as a formal language, it included a result about what it could *not* be: with plausible assumptions about which sentences are and are not acceptable in English; it showed that human language could not be generated by a finite state machine, as illustrated by (X) for a grammar with the sentences *the old man comes*, *the old men come*, *the old old old old man comes*, etc. Chomsky (1957) shows that English has properties such as center-embedding (ab , $aabb$, $aaabbb$, ..., or, more generally $a^n b^n$), which make it impossible to derive all utterances of English by simply transitioning consecutively from one word in a sentence to the next (Chomsky 1957:18-23, see also Jäger & Rogers 2012 for recent discussion).



(illustration from Chomsky 1957:19)

This gave rise to a long line of research that sought to establish the following two points, which we discuss in turn: (i) the detailed properties of natural language syntax, both in English and across languages¹; and (ii) where in the hierarchy of formal languages ('regular' languages with finite state grammars, context-sensitive languages, etc) natural languages lies.

Notably, when Chomsky laid out the goals of linguistic investigation in *Syntactic Structures*, human language may well have been the object of study. But the associated questions and tools (notably, the hierarchy of formal languages) were fruitfully applied to further areas, such as animal signals, including in cases in which their formal properties are very different from human language (e.g. where they lie much lower on the 'Chomsky hierarchy' of formal language complexity). Similar extensions were natural in further areas: constituent analysis originated in studies of language (see Section 2.1.2) but turned out to be useful in the analysis of visual narratives and even music; variables originated in logic and semantics (see Section 2.2.2) but illuminated aspects of visual narratives.

Focusing on the study of human language for a bit longer, Chomsky (1957) established the following. Some arguments, such as the above-mentioned argument that finite state grammars cannot generate human language, can be made solely on the basis of the set of

¹ See Wiltschko 2018 for a discussion on the cross-linguistic developments within this research program.

sentences that are or are not acceptable in the language (the 'weak generative capacity' of a grammatical formalism). But linguists are generally interested in the question of which types of *structures* a grammatical formalism predicts (its 'strong generative capacity'), i.e. hierarchical constituent structures, which are taken to form the basis of compositional semantic interpretation. The shift from an earlier focus on 'weak generative capacity' (acceptable sentences) to a later focus on 'strong generative capacity' (possible structures) is generally taken to be articulated in texts such as Chomsky 1980:220. This distinction is also important for the study of objects outside of natural language, as we can probe for possible/predicted structures in animal communication, music, dance, visual narrative, and so forth.

When we compare human language to other potential objects of study, there is general agreement about the general part of the hierarchy of formal languages, in which natural language syntax lies, illustrated in (1): definitely beyond finite-state (as discussed above), and certainly beyond context-free grammars (i.e. phrase-structure grammars, such as Python code, Jäger & Rogers 2012:1959). Natural languages have been argued (e.g. see Jäger & Rogers 2012, Rohrmeier et al. 2015, amongst others) to be within a 'mildly context-sensitive' region between context-free languages (where rules require one non-terminal category as their input) and context-sensitive languages (where rules can have terminals and non-terminals in their input).

When we proceed to compare natural language to other potential grammars, such as a grammar of animal languages, music, dance, or pictorial representation, a clear division breaks down. On the one hand, not all parts of language seem to require powerful grammars: phonological rules are usually thought to lie within the finite-state bound (Heinz and Idsardi 2011). Similarly, most animal languages that have been described are well within the finite-state bound, including sophisticated birdsongs, which we briefly discuss below. On the other hand, there is an ongoing search for animal languages that clearly go beyond this bound (see e.g. Fitch & Hauser 2004, Gentner et al. 2006, van Heijningen et al. 2009, Kershenbaum et al. 2014, Ferrigno et al. 2020).

Things are different again when one studies *human* representational systems that are non-linguistic in nature. Here structure, which derives from grouping/constituency, is often one of the first things that formal studies ask about, as has been the case in the study of music, dance and visual narratives. As we will see below, in all three objects of study, sophisticated principles of structural grouping have been posited, and in the case of Western tonal music there are precise arguments (Rohrmeier 2011) to the effect that music syntax goes beyond the finite-state bound, as is illustrated in (1).

An important difference between human language and representational systems and animal languages is that in the case of the latter, one has no choice but to assess grammatical formalisms on the basis of animal behavior data, as opposed to a direct evaluation of stimuli that can be directly communicated. For the study of human language and representational systems, researchers can generate and study 'minimal pairs', i.e. minimal modifications of one and the same object,² such as a melody in the study of music syntax, see section 6.2.1.

² The method of 'minimal pairs' in linguistics corresponds to that of 'controlled experiments' in experimental fields. In both cases, what counts as 'minimal' or 'controlled' depends on an implicit theory of which properties may be relevant for the question at hand: Spanish *alguno* 'any.masc.sg' and *alguna* 'any.fem.sg' form a minimal pair for grammatical gender, but the modification also entails changes unrelated to gender, such as the phonology (change in the final vowel). The pair is 'minimal' only to the extent that the phonology doesn't matter.

- (1) Human language, animal vocalizations and music in the hierarchy of formal languages (Rohrmeier et al. 2015).




Figure 3. A Venn diagram of the Chomsky hierarchy of formal languages with three extensions annotated with a comparison of the hypothesized classifications of human languages, (human) music and animal vocalization. The areas marked with ‘question’ signs indicate that further research is required to settle examples for the respective class of complexity in these domains. (Online version in colour.)

2.1.2 Constituency and Grouping

A key notion of natural language syntax is that of constituency and grouping; for instance, in (1), the group of words *the wild deer* form a unit with a distinct syntactic behavior, a so-called constituent. Similarly, *feed the wild deer* forms a constituent; by contrast, *feed the wild* does not.

- (1) I will [feed [the wild deer]].

Constituency diagnostics allow us to determine such groups in natural language; for instance, constituents can be topicalized, as in (2a) and (2b) vs. (2c) (for the sentence in (1)), or (since topicalization may seem archaic to some speakers) occur in the predicative position of a pseudocleft created with the material of the sentence to be probed, as in (3a) and (3b) vs. (3c).

- (2) a. [The wild deer], I will feed.
 b. [Feed the wild deer], I will.
 c. * [Feed the wild], I will deer.
- (3) a. What *I will feed* is [the wild deer].

- b. What *I will do* is [feed the wild deer].
- c. * What *I will do deer* is [feed the wild].

Syntactic grouping of this kind (amounting to constituency) can be found in many objects of study outside of natural language, including music and dance, as we will see below. A deeper understanding of those objects of study will thus benefit from applying the analytical tools and devices established in formal linguistic analysis.

2.2 Semantics: what are the elementary operations by which language produces information?

To investigate semantics, a key notion is that of *compositionality*, according to which the meaning of a sentence is computed on the basis of the meaning of its elementary parts and the way they are put together (in other words, their syntax). Meaning is thus derived from two components, pertaining to the meaning of elementary expressions, and their mode of composition.

2.2.1 Expressive power and constraints on possible meanings

Elementary expressions, minimal meaning-bearing units of language, may be closed-class (such as the morphemes ‘some’, ‘the’, ‘in’ and ‘-ed’) or open-class (such as the morphemes ‘dog’, ‘laugh’ and ‘happy’), but in both cases there are systematic constraints on their semantic behavior. Formal semantics has particularly studied constraints on closed-class (= grammatical/logical) expressions, such as the determiner *some* in (2)a. Its contribution is partly similar to that of the quantifier *for some x* (written as $\exists x$) in logic. In this case, the truth conditions of the English sentence come close to those of the logical formula in (2)b, which belongs to Predicate Logic, the most common logic used to formalize mathematics and science in general.

- (2) a. Some senator is honest.
- b. $\exists x(\text{senator}(x) \text{ and } \text{honest}(x))$

But as soon as further cases are considered, this analysis falls short. *Most* as in *most senators are honest* cannot in general be analyzed in terms of Predicate Logic: natural language is a more expressive system, one with 'generalized' quantifiers – not just *some* and *every* as in Predicate Logic, but also *most*, *no*, *exactly one*, etc. Natural language quantifiers thus differ from quantifiers of Predicate Logic in being more diverse. In addition, a noun phrase can also *restrict* a natural language quantifier: for instance, *most senators are honest* has an approximate meaning where there are more honest senators than there are dishonest senators (see e.g. Barwise & Cooper 1981). This cannot be defined in terms of *most things in the universe satisfy certain properties*, which is what Predicate Logic would require. In other words, it won't do to treat it as: *most things are senators and are honest* – an obvious falsehood; other paraphrases inspired by Predicate Logic fail as well. As outlined above, work in formal syntax following the Chomskyan revolution sought to situate precisely the syntactic complexity of human language within a hierarchy of formal languages of increasing sophistication. The same enterprise has been conducted in logical (formal) semantics, but on the meaning side: semanticists can provide a reasonable approximation of the expressive power of human languages relative to all sorts of formal languages (e.g. Keenan and Westerstahl, 2011).

A panoply of constraints on natural language semantics have been discussed. One is that natural language expressions differ in their entailments (captured by the terms *positive*

and *negative* in Barwise and Cooper 1981), as follows. The sentence *Sam is in Paris* entails a modified sentence generated by replacing *in Paris* with the less restrictive property *in France*: this is the hallmark of a *positive* (or ‘upward entailing’) expression (here: *Sam is* ___). By contrast, the expression *Sam is not* ___ is *negative* (or ‘downward entailing’) in that it licenses an entailment in the opposite direction: from the sentence *Sam is not in France*, one can infer that *Sam is not in Paris*, as seen in (3):

- | | | | |
|-------------------------|-----------------------------|----|----------------------------|
| (3) Positive operators: | Sam is in Paris | => | Sam is in France |
| Negative operators: | Sam is not in France | => | Sam is not in Paris |

From such a perspective, some elementary grammatical (closed-class) words can be classified as positive or negative, and in particular so-called quantifiers: *some* as in *some senator* ___ is positive, *no* as in *no senator* ___ is negative. By contrast, *exactly one* ___ is neither positive nor negative: *exactly one archbishop is in France* doesn't entail that *exactly one archbishop is in Paris* (e.g. there might be zero in Paris and one in Lyons); and conversely, *Exactly one archbishop is in Paris* doesn't entail that *Exactly one archbishop is in France* (e.g. there might be one in Paris and another one in Lyons). From this restricted data, one could see emerging a constraint according to which only quantifiers that can be classified as positive or negative can be elementary expressions: *some* and *no* are elementary expressions but *exactly one* (two words) isn't. This indeed illustrates one of the simplest forms of a semantic universal, which puts constraints on the possible meanings of natural language elementary expressions.

Constraints on the meaning of open-class words like *chair* have been investigated as well. Gärdenfors (2004, 2014) proposes that one constraint on lexical expressions is connectedness. Formally, the denotation of a word is connected if for any two elements in it, all elements that are between those two are also in the denotation. Take the word *animal*, with the assumption that the duck and the badger are instances of animals. If one has reason to think, at any level of reasoning or conceptualization, that the platypus is intermediate between ducks and badgers, one will infer by connectedness that the platypus is an animal. Interestingly, this notion has recently been unified with a constraint on quantifiers: the property of being positive (*some*) or negative (*no*) in the above sense is subsumed under the term *monotonicity* (with quantifiers of the *exactly*-type being non-monotonic). Connectedness can then be seen as a generalization of monotonicity in open-class expressions.³ Strikingly, we will see that these constraints have precursors in non-human animals; this shows how the application of formal linguistic tools to research on non-human animals can benefit our understanding of animal minds.

2.2.2 Variables

Relative to its logical machinery, there are further respects in which natural language both resembles and differs from Predicate Logic. Variables, often notated as x , y , z in mathematics, are thought to be real yet unexpressed in spoken language.⁴ To illustrate, *Sarkozy told Obama that he would be elected* is ambiguous between a reading on which *he* refers to Sarkozy and one on which *he* refers to Obama. Many linguists have argued on indirect grounds that disambiguation is effected by way of invisible variables, as illustrated in (4)a: *he* may refer to $Sarkozy_x$ or $Obama_y$, and disambiguation is effected because *he* is mentally processed with the variable x or y . (4)b displays the same type of disambiguation, but the

³ Technically speaking, this can be formally stated as follows: a quantifier Q is monotonic if Q and its negation $not\ Q$ are connected (Chemla et al. 2019a), and conversely a quantifier Q is connected if it is a conjunction of monotonic quantifiers (Thijsse 1983).

⁴ Note that there exist alternative, 'variable-free' analyses, for both spoken and sign language (e.g. Barker and Jacobson 2007; Kuhn 2016).

variables are now dependent on ('bound by') quantifiers. Variables can also be used to model cross-sentential anaphoric dependencies between an indefinite and a pronoun, (4c).

- (4) a. **Sarkozy**_x told **Obama**_y that he_x / he_y would be elected.
 b. [**A representative**]_x told [**a senator**]_y that he_x / he_y would be re-elected.
 c. [**Some senator**]_x is honest. She_x opposed the proposal.

Once a framework with variables was in place, semanticists uncovered some differences between natural language variables and those of Predicate Logic. In logic, a quantifier such as $\exists x$ in $\exists xF$ can only control variables x that appear in the immediately following formula F . This is compatible with our analysis of (4)b. By contrast, in (4)c, the quantifier [*some senator*]_x controls a variable x (on *she*_x) that appears in another sentence altogether. To allow for such dependencies, non-standard logics called 'dynamic logics' have been developed for natural language (e.g. Kamp 1981, Heim 1982).

Positing abstract variables was initially motivated indirectly by their explanatory force. Remarkably, variables can arguably be overt in sign languages, namely as positions in signing space, called 'loci'. In (5), *past senator* co-occurs with a pointing sign (or 'index') *IX-a* towards position *a*, on the right, and *current senator* co-occurs with pointing sign *IX-b* towards position *b*, on the left. The two pronouns in the second sentence can then be fully disambiguated: *IX-a* refers to the former senator, *IX-b* refers to the current senator. The positions *a* and *b* play the role of the variables x and y in (4), disambiguating the meaning of pronouns and making overt the relation between a variable and an existential quantifier in a previous clause; the latter point also strengthens the case of dynamic logics (Schlenker 2011).

- (5) IX-1 KNOW [**PAST SENATOR PERSON**] IX-a IX-1 KNOW [**NOW SENATOR PERSON**] IX-b. IX-b SMART BUT IX-a NOT SMART.
 'I know a former senator and I know a current senator. He [= the current senator] is smart but he [= the former senator] is not smart.' (ASL; 4, 179; Schlenker 2014)

Strikingly, we will see that variables (in their linguistic version) have been argued to play a semantic role in pictorial narratives, in narrative dance, and possibly even in music; this shows how the study and understanding of such human artifacts or art forms can benefit, from the application of formal linguistic tools, such as variables.

2.2.3 Compositionality

How can elementary meanings be combined? This question has often received a relatively uniform semantic answer: a combination of two expressions is interpreted by treating one expression as a function and the other as an argument. For instance, in the sentence [*some senator*][*is honest*], [*is honest*] has as its meaning the set of people who are honest. And [*some senator*] has as its meaning a function f from sets x to truth values, with $f(x) = \text{true}$ just in case x contains at least one senator.

This is a versatile framework that extends to other quantifiers, including ones that are not found in Predicate Logic. To illustrate, the meaning of *most senators* can be taken to be a function g for which $g(x) = \text{true}$ just in case x contains more than half of the senators. What matters for our purposes is that this function-argument mechanism is very different from the mere addition of meanings by way of concatenation (= juxtaposition), of the sort we could get for: *It's raining. It's cold*. Here we can treat each sentence as making an independent contribution: the two meanings can be conjoined without further ado. This strategy won't normally work when function-argument combinations are involved (e.g. it won't do to treat *some senator is honest as: There is some senator. There is honesty*). Importantly, we will see that concatenation appears to suffice to combine monkey calls.

2.3 Pragmatics: the inferential typology of language

Contemporary semantics has discovered that language doesn't just convey information but simultaneously comments on its status relative to the beliefs of the speech-act participants, whence a surprisingly rich typology of inferences and a new question: how is information divided among the slots of this inferential typology?

2.3.1 *Implicatures*

To illustrate, let us go back to *Some senator is honest*. It makes two contributions: its literal meaning is that at least *one senator is honest*. But in addition, the sentence strongly suggests that *not all are*. This cannot be due to the literal meaning of *some* because it is no contradiction to say: *Some senator is honest – in fact all are*. The inference from *some* to *not all* is defeasible, and sometimes absent (e.g. an answer *no* to *Is some senator honest?* denies that *any* is, not that *some but not all are*). A standard analysis (following Grice 1975) involves competition between *some* and *every*: the statement *every senator is honest* is more informative, hence if the speaker didn't use it, chances are that this was because she took it to be false. This inference, called an implicature, is crucially based on an Informativity Principle:

(6) If S and S' are competing utterances, and if S' is more informative than S, prefer S' over S unless S' is false.

As we will see, implicatures are found throughout human communication (including gestures). More surprisingly, implicatures (or implicature-like inferences) have been argued to explain parts of the meaning of monkey alarm calls in several species (see section 3.4).

2.3.2 *Presuppositions*

The typology of linguistic inferences doesn't end there. *Mary stopped/continued/regretted smoking* triggers the inference that Mary smoked before. This is called a presupposition because it is presented as already established in the discourse (see for instance Beaver, Geurts & Denlinger 2021 for a recent survey). Unlike the 'at issue' entailment from *Mary is in Paris* to *Mary is in France*,⁵ the presupposition is preserved in questions and in negative statements, as shown in (7):

- (7) a. Did Mary stop/continue/regret smoking? => Mary smoked before
b. Mary did not stop/continue/regret smoking. => Mary smoked before

Presuppositions are ubiquitous. If I ask about someone in the distance: *Is she approaching?*, I am asking whether the person is approaching but presupposing that this person is female (unlike in the question: *Is this a woman approaching?*). Thus the gender specifications of pronouns generate presuppositions (see Cooper 1983). While presuppositions are often thought to be hard-wired in the meaning of words, we will see that new gestures and even visual animations can generate presuppositions. This suggests that a productive rule lies at their source: presuppositions need not be encoded in words. Notably, while we have already highlighted cases where non-standard objects such as animal communication, music or dance benefit from the application of formal linguistic tools, this is a clear case where formal linguistic theory benefits directly from the inclusion of non-standard objects.

⁵ We follow standard practice in using the term 'at issue' pre-theoretically to refer to the *main point* of an utterance (see e.g. Simons 2007 for discussion of the term *main point*, see Potts 2015:168 for discussion of *at-issueness* as a pre-theoretic concept). We do not adopt a technical (QUD-based) definition of 'at-issueness', as proposed, for example, by Beaver et al. (2017:280).

2.3.3 Supplements

A fourth inferential type pertains to supplements, which are the meaning inferences typically triggered by non-restrictive relative clauses (see e.g. Potts 2005, Schlenker 2010). Unlike presuppositions, they are presented as informative. But unlike ‘at issue’ entailments, they trigger the same inferences as independent clauses even when they are embedded under logical operators. They thus contrast with embedded conjunctions, hence (8)a,b but not (8)c trigger the inference that *if Ann lifts weights, this will adversely affect her health*.

- (8) a. If Ann lifts weights, which will adversely affect her health, we should talk to her.
b. If Ann lifts weights, we should talk to her. This will adversely affect her health.
c. If Ann lifts weights and this adversely affects her health, we should talk to her.

We will see that supplements too can be generated with new gestures or visual animations, ones that follow the words they specify.

2.4 Psycholinguistics of meaning: processing and acquiring meaning

The last question is how semantic operations are put to actual use, and how they are acquired by children. In its early days, formal semantics repudiated questions of mental reality (Thomason 1974). Later developments disavowed this position. Here the case of implicatures offers an enlightening case study. Their formal description suggests that they should be an add-on to the literal meaning of words, obtained by applying the Informativity Principle. This could be expected to take time, and processing studies have shown that this is indeed the case. In addition, their derivation requires a comparison between an utterance and its competitors, and finding the competitors might be expected to be difficult. This arguably accounts for an important finding: young children often fail to compute implicatures where adults do, although when given an explicit choice among alternatives they avoid the under-informative one (Chierchia et al. 2001). The key lesson is that semantic theories make predictions about language acquisition and processing, highlighting the fact that semantics is part of cognitive science.

2.5 Interim summary

In sum, contemporary linguistics offers a precise analysis of the form and meaning of sentences, treated simultaneously as formal and as cognitive systems. One key question on the syntactic side is where natural language lies in the hierarchy of formal languages; the same question can and should be asked about further systems such as animal call sequences and human artifacts such as dance, music, and visual narratives. One key question on the meaning side is which aspects of semantic information are encoded in words, and which are due to productive rules – such as the productive rules of presupposition generation discussed in section 2.3.2; investigating new, sometimes invented meaning-bearing forms (such as gestures) will help address the productivity question. More broadly, we would want to determine how these semantic operations compare to those found in other semantic systems in nature, and what we can learn from the comparison.

3 Formal approaches to animal communication

3.1 The formal approach

There are endless debates about appropriate terminology in animal communication, starting with the word 'language': are we even allowed to apply it to non-human systems? Debates about terminology are rarely fruitful, and for good reason in this case: human language is

characterized by so many distinct properties that any number of subsets of them could be taken to 'define' language, while extending differently to other systems. Formal approaches adopt a different method: starting from the notion of 'language' in formal language theory (e.g. any set of strings), they treat all sorts of things as languages, and focus their attention not on the term but on the properties of these systems with respect to form and sometimes meaning.

3.2 Animal syntax: the example of birdsongs

In the calls and gestures of primates, little syntactic sophistication has been demonstrated. Bird calls, which serve to convey information about the environment (e.g. the presence of predators), have not been shown to have a sophisticated syntax either. But things are different in the area of bird songs, which arguably differ from calls in not conveying information about the environment, and in possibly just advertising the singer's presence and general qualities. It is thus conceivable that birdsongs have a syntax but no clear semantics. However, the syntax of bird songs is interesting and worth discussing, bearing in mind that it is mostly contained within the bound of Finite State (= regular) languages as discussed in Section 2.1.1 (see Berwick et al. 2011). To illustrate, Bengalese finch songs have been analyzed with the finite-state grammar posited in (9) (where states are connected with transition probabilities).

(9) Syntax of Bengalese finch song




Figure 3. Probabilistic finite-state transition diagram of the song repertoire of a Bengalese finch. Directed transition links between states are labelled with note sequences along with the probability of moving along that particular link and producing the associated note sequence. The possibility of loops on either side of fixed note sequences such as *hh* or *lga* mean that this song is not strictly locally testable (see Box 3 and main text). However, it is still *k*-reversible, and so easily learned from example songs [35]. Adapted, with permission, from [75].

The type of finite-state grammar shown in Figure 3 is quite sophisticated in that it involves multiple states that can produce diverse strings; it exemplifies that such grammars are found in the animal kingdom. This leaves three important questions open. First, are there cases in which animal syntax goes beyond the finite-state bound? Second, do the finite-state rules found in animals have anything to do with human language? Fitch (e.g. 2018) proposes a "continuity hypothesis", according to which "humans share the processing capabilities required to deal with regular-level sequential processing, and thus phonology, with other animals, and these shared capabilities are implemented in homologous neural processing algorithms and circuitry" (p. 72); by contrast, human syntax goes beyond the finite-state bound. Third, do the finite-state languages found in animal communication occupy an interesting subpart of the set of finite state languages? Recent research (such as Berwick et al. 2011) has established a 'subregular hierarchy' (a hierarchy among the formal languages

within the finite-state bound), and the formal analysis of this subregular hierarchy has already started to be applied to animal syntax (e.g.; strikingly, this is also the case for human phonology (e.g. Heinz and Idsardi 2011)).

3.3 Animal meanings: the example of monkey communication

On the meaning front, it is with respect to monkey communication that the most precise formal tools have so far been used. Contemporary primatology has established that some primate calls and gestures are naturally produced in some situations but not in others; monkey alarm calls are produced in response to the presence of a predator (e.g. Zuberbühler 2009), great ape gestures serve purposes such as to solicit grooming from another ape (e.g. Byrne et al. 2017). Focusing on monkey alarm calls, field experiments have shown that the monkeys themselves know the correlation of calls with predation situations and thus derive information from the calls they hear. Under any name ("appropriate" vs. "inappropriate" or "true" vs. "false"), this means that these signals have a semantics – which is unsurprising since they clearly convey information.

Applying the methods of formal semantics to monkey communication has made two contributions: (i) it has clarified the space of possible theories by stating fully explicit analyses about (a) the meaning of elementary calls and (b) the ways they are combined; (ii) it has generated new potential explanations, especially about the division of labor between semantics and pragmatics, leading to the hypothesis that there are 'animal implicatures', cases in which the meaning of a call is enriched by the assumption that it was the most informative call that could be appropriately used. The main results, discussed in the remainder of section 3, are as follows:

- (10) a. Elementary calls can often be assigned simple and explanatory meanings.
- b. Call combinations can usually be analyzed without structure, by treating each call as a separate utterance reflecting the caller's information as it is produced: concatenation appears to suffice. The *-oo* component of Campbell's monkey calls is so far the only plausible case of morphology with compositional interpretation at the level of what we may dub 'call suffixes'.
- c. Call meaning might be enriched by an animal version of the Informativity Principle, although this hypothesis needs to be tested further.

3.4 Call meaning vs. pragmatic enrichment: Campbell's monkeys⁶

A rich example pertaining to the meaning of individual calls is afforded by Campbell's monkeys of the Tai Forest (= (11)). Male adults have a non-predation-related call, *boom*. They use a call *krak* to raise leopard alerts, and *hok* for raptor alerts. Remarkably, they also have suffixed calls: *krak-oo* is used for unspecific alerts, and *hok-oo* for non-ground disturbances. The challenge is to assign meanings to *boom*, *krak*, *hok*, and *-oo*. Further complexity is added by Campbell's call use on Tiwai Island, where leopards haven't been seen for decades: the Tai calls are used, but *krak* raises unspecific alerts (as does *krak-oo*), rather than leopard alerts. Should we conclude that call meaning is subject to a kind of dialectal variation – as it is for *pants* in American English (meaning "trousers") vs. British English (meaning "underpants")?

⁶ Sections 3.4 and 3.5 borrow from Schlenker et al. 2016b.

(11) **Male Campbell's monkey calls**

a. Description

b. Analysis

c. Results of call competition

CAMPBELLS MONKEYS

Call	Typical situations
boom boom	non-predation alert
hok	presence of an eagle
krak	<i>Tai</i> : presence of a leopard <i>Tiwai</i> : unspecific alert
hok-oo	alert from above
krak-oo	unspecfic alert


Literal meanings	
boom boom	non-predation alert
hok	non-ground alert
krak	alert
<i>R-oo</i>	weak <i>R</i> -alert

Informativity Principle
 "Prefer more informative expressions!"

Calls	Competitors	Enriched meanings
hok	hok-oo	serious non-ground alert
krak	krak-oo, hok	<i>Tai</i> : alert, serious, ground <i>Tiwai</i> : useless enrichment, hence literal meaning only

Schlenker et al. (2014) propose an analysis without dialectal variation, but with a crucial use of the Informativity Principle (see (6) in section 2.3.1). *Why not adopt the dialect analysis?*, a reader might want to know. First, the monkeys take *krak* to trigger general alerts, and *hok* to trigger non-ground alerts. Second, in order to analyze the meaning of the suffix *-oo*, they assume that if *R* is *krak* or *hok*, *R-oo* indicates a weak alert of the *R*-type. Thus *hok-oo* indicates a weak (*-oo*) non-ground (*hok*) alert – which is more informative than *hok*. Finally, they use the Informativity Principle, with the informativity relations in (12). We can then show why *hok* ‘there is a non-ground alert’ is understood to mean ‘there is a raptor’, and why *krak* ‘there is an alert’ in *Tai* is understood to mean ‘there is a leopard’.

(12) **Informativity relations among Campbell's calls (for two calls that are linked, 'higher' means 'more informative')**



First, *hok* competes with another call, namely *hok-oo*, and because *hok-oo* (pertaining to *weak* non-ground alerts rather than to any non-ground alert) is more specific, the meaning of *hok* is enriched to *hok but not hok-oo*: it only applies to aerial (= *hok*) non-weak (= *not hok-oo*) alerts – hence, *hok* is used for raptors.

By the same logic, the unspecific alert *krak* competes with two calls: with *krak-oo*, but also with *hok*. Due to this competition with two more informative calls, *krak* can only be used for *serious* (= *not krak-oo*) *ground* (= *not hok*) *disturbances* – as long as serious ground disturbances exist, which in the world of Campbell's monkeys means: leopards. This comes very close to the leopard uses observed in *Tai*. Crucially, if serious ground disturbances are taken out of the picture, as in *Tiwai*, where leopards are absence, then this inference disappears, and *krak* collapses into being interchangeable with *krak-oo*, which no longer differs in informativity; though both remain ‘*not hok*’, i.e. they are restricted to ground disturbances even in *Tiwai*. Importantly, while in human language the Informativity Principle is often taken to derive from more general rules of communicative rationality, it might conceivably be hard-wired and automatic in non-human primates.

3.5 Call combination vs. context change: Titi monkeys

The calls of the Titi monkeys of South America make a different point: not everything that looks like a complex syntax/semantics interaction requires complex composition rules.

The initial puzzle is that with two calls (A and B) re-arranged in various ways, Titi monkeys can provide information about both *predator type* (cat,⁷ raptor) and *predator location* (on the ground, in the canopy). Simplifying somewhat, and writing X^+ for a series of iterations of call X , B^+ is used for non-predation alerts and for situations involving a cat on the ground, while AB^+ is used in situations that involve a cat in the canopy. A raptor on the ground gives rise to an A^+B^+ sequence, while a raptor in the canopy triggers an A^+ sequence. The main patterns are summarized in (13). Should we conclude that these sequences have a complex syntax/semantics interface, as hinted by Cäsar et al. 2013? Or should they be treated as very long idioms, with no internal semantics? In sharp contrast to bird song, there is no evidence that combinations of monkey alarm calls exhibit syntactic structure; a concatenative approach that treats each call as a separate utterance appears to be sufficient.

(13) **Titi monkey calls**

a. Description

b. Analysis

c. Results of call competition

TITI MONKEYS	Call	Typical situations	Literal meanings A serious non-ground alert B alert	B ⁺ correspond to ground or weak alerts (or alerts that <i>have become</i> weak alerts) (Raptors in the canopy remain serious threats even after having been signaled by As)
	B ⁺	non-predation alert		
	B ⁺	cat on the ground		
	A B ⁺	cat in the canopy		
	A ⁺ B ⁺	raptor on the ground		
A ⁺	raptor in the canopy	Informativity Principle "Prefer more informative expressions!"		

Due to their length and slow time course, it is unlikely that these sequences are interpreted non-compositionally as idioms because hearers would need to wait for too long for the meaning of the message to be effective. A simpler analysis has been explored, in which each call is in effect an independent utterance and thus contributes its meaning independently from the others (Schlenker et al. 2016c). Since the B-call is used in predatory and non-predatory situations alike, one may take it to trigger an unspecific alert. In field experiments, the A-call triggers a 'looking up' behavior, and thus one can posit that it is indicative of *serious non-ground alerts*. While the 'looking up' behavior warrants the *non-ground* component, the *serious* component is motivated by the strong association with raptors – a serious arial threat for Titi monkeys. These assumptions explain why one finds B⁺-sequences (= series of B-calls) in 'cat on the ground' situations, and A⁺-sequences in 'raptor in the canopy' situations. Importantly, the Informativity Principle is important to explain why the unspecific B-call *fails* to be used at the beginning of sequences in the latter case, even though it would be applicable: it would simply be under-informative.

But why does one find A⁺B⁺ in 'raptor on the ground' situations? A remark about hunting techniques proved suggestive: raptors on the ground usually attack by *flying*, hence the serious non-ground alerts A⁺. Still, being on the ground isn't a typical hunting position, and after a while the alert stops being serious, which only leaves B as a possibility. In 'cat in the canopy' situations, one finds AB⁺ sequences, possibly because a serious non-ground danger is indicated, which then transitions to a weaker danger because a cat becomes less dangerous after detection (Zuberbühler et al. 1999).

⁷ Note that we use the term *cat* in the broad sense 'felid' – leopards, jaguars, cougars, etc. –, and not in the narrow sense 'domestic cat'.

On this view, which is still the subject of active debates,⁸ the apparent complexity of Titi sequences might reflect the interaction between simple meanings and the evolution of the contextual environment as the sequence unfolds, rather than a complex syntax/semantics interface or very long idioms.

In sum, the Campbell's and Titi studies illuminate several points: (i) they clarify the possible meanings of individual calls; (ii) in the case of Campbell's *-oo*, they make it plausible that there is a tiny composition rule at the level of a suffix; (iii) in both species, there seem to be general alert calls, but their use is constrained by the Informativity Principle; (iv) in all cases, a call can be treated as a separate utterance, and no non-trivial composition rules are needed at this point – and in particular no call-external application of functional application, as in human language.

3.6 New questions about the evolution of animal meanings

The investigation of primate signals has another benefit: in some cases, it yields insights into the evolution of animal meanings. As an example, *booms* are non-predation-related calls present, not just in Campbell's monkeys, but also in many cousin subspecies of the family cercopithecines. Inspection of their distribution is strongly indicative of their presence in the most recent common ancestor of entire subgroups: *booms* probably existed several million years ago (Schlenker et al. 2016a).

Turning to ape gestures, Byrne et al. 2017 argue that 36 gestures are shared among the great apes, and likely part of an innate repertoire. On the form side, Hobaiter and Byrne 2017 argue that the similarity of the repertoires is unlikely to be due to limitations of the articulatory possibilities for gestures. On the meaning side, Byrne et al. 2017 argue that partial similarities in gesture use among bonobos and chimpanzees are also unlikely to be due to chance.

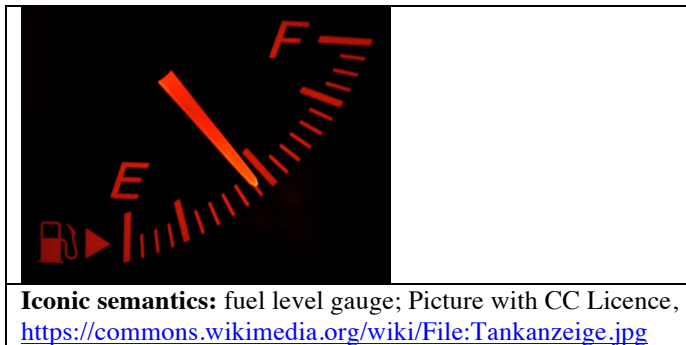
Strikingly, ape gestures are an area in which a genuine connection with human communication might be non-speculative, as argued by Kersken et al. 2018. They uncover a human infant gestural repertoire with 50 gestures (96%) that are shared with apes. A further step will be to connect these infant gestures with adult human gestures, and/or with signs. Yet another question will be to assess the extent of iconicity in ape gestures, discussed for pointing and pantomime by several groups (e.g. Genty and Zuberbühler 2014, Douglas and Moscovice 2015).

Call evolution raises questions of its own, in particular the following: under what conditions can the simultaneous presence of specific and general calls be stable? Steinert-Threlkeld et al. 2019 show that under broad modelling conditions, general calls are unstable: a Titi-style general B-call competing with a specific A-call meaning *serious non-ground alert* should lead in later generations to a division of lexical labor, with the B-call meaning not-A, i.e. *non-[serious non-ground] alert*. If so, the Informativity Principle won't have anything to operate on, since call meanings will never be in a specific-to-general relation. To block Steinert-Threlkeld et al's results, however, we might surmise that *non-[serious non-ground] alert* is not a possible concept because it does not correspond to a natural class of objects. But showing this requires an independent study of constraints on animal concepts, a point we revisit in Section 7.

⁸ Berthet et al. 2019a,b fail to replicate the AB⁺ sequences in 'cat in the canopy' situations, finding instead B⁺, with A's interspersed; they also develop a somewhat different theory.

4 Integrating logical and iconic semantics

Iconicity is an important way of producing meaning in spoken and especially signed languages. Informally, *iconicity* can be defined as resemblance-based meaning; in other words, an iconic sign resembles or imitates that which it denotes. More formally, Greenberg (2021) proposes that iconic semantics is attested whenever the *form* of the sign is also part of its meaning/denotation; in an adaptation of Greenberg's own example, a fuel level gauge has an iconic semantics in that the *angle* of the needle on the dial (which is part of its form) is mapped to the *amount* of fuel in the tank (which is part of its meaning).




A key challenge is to understand, both empirically and formally, how iconic and compositional semantics interact; this will have benefits for the analysis of meaning in speech, signs and gestures.

4.1 Pictorial semantics

Pioneering work by Greenberg (2013, 2019, 2021) on the formal semantics of pictures clarifies the workings of iconicity. The semantic content of a picture is obtained by asking which situations could give rise to the markings found on the picture given the projection method used. One among several is perspective projection, illustrated in (14).⁹

(14) Perspective projection (Greenberg 2019)



This naturally gives rise to a definition of truth for pictures: a picture P is true of those situations that can project onto P relative to the viewpoint and the system of projection used.


(15) A picture P is true in situation w relative to viewpoint v along the system of projection S iff w projects to P from viewpoint v along S , abbreviated as: $\text{proj}_S(w, v) = P$

Abusch 2013, 2015 uses this framework to account for the semantics of silent comics, a point to which we return in Section 6.1. But iconic semantics also interacts in interesting ways with compositional semantics in sign languages, to which we now turn.

⁹ Note that there are different systems of iconic semantics; acoustic iconicity in music semantics would have to be handled by different principles from perspective projection in picture semantics.

4.2 Logical visibility and iconicity in sign language

Sign languages have been argued to display the same general grammatical and semantic properties as spoken languages (e.g. Sandler and Lillo-Martin 2006), and sometimes to express aspects of Logical Forms in a more transparent fashion, as we saw in Section 2. But they also make greater use of iconicity than spoken languages to modulate the form of conventional words. A simple example is the verb *GROW* in ASL, which is (conventionally) realized as shown by the pictures in (16). But it can also be modulated in iconic ways by making the endpoints more or less broad, and by realizing the sign more or less quickly, with unmistakable semantic effects: the broader the end points of the sign, the larger the final size of the group; and the more rapid the movement, the quicker the growth process.



(16) POSS-1 GROUP GROW

'My group has been growing.' (ASL, 8, 263; Schlenker et al. 2013)

This phenomenon is not unique to sign language, however: the adjective *long* can be modulated by lengthening the vowel to evoke a greater duration (and doing the same thing with *short* is... odd). Here the rule of iconic modulation seems to be that *the longer the vowel, the greater the duration*. Due to the representational possibilities of the signed modality, these iconic modulations are just far richer in sign than in speech. In addition, it has been suggested that for *GROW* and *long* alike, iconic modulations make an at issue contribution. Thus *loooong* means very much the same thing as *very long*, and this is the content that gets denied by a negation, as in: *The talk wasn't loooong* (see Schlenker 2018b for ASL *GROW*).

Iconicity doesn't just interact with lexical material: grammatical expressions such as pronouns can be modulated in a highly iconic fashion as well. Just like the feminine specification of *she* triggers a presupposition that the denoted person is female, pronouns in ASL (American Sign Language) and LSF (French Sign Language) can have *high* specifications (realized by pointing upwards) that trigger the presupposition that the denoted person is tall, powerful or important. Furthermore, these modulations can display 'iconicity in action': one typically points towards the part of the representation corresponding to the head, with the result that when the denoted individual is presented as rotated in various position, the loci get rotated as well. The conclusion is that sign language loci can simultaneously be variables and simplified pictures of their denotations (see Liddell 2003, Schlenker et al. 2013, Schlenker 2014, Schlenker 2018c).

A striking case of interaction between grammar, logical semantics and iconicity involves repetition-based plurals. In ASL, one can optionally realize a plural, such as that of the word *TROPHY*, by repeating the noun. This yields standard readings of English-style plurals. But simultaneously, the shape of the repetitions can provide iconic information about the arrangement of the denoted objects (Schlenker and Lamberton 2019a): one may repeat *TROPHY* in a horizontal or triangular shape to signify that the trophies are arranged on a line or as triangle. But there is more: the repetitions can be easy to count, with clear separations between them ('punctuated'), or hard to count, without clear separations ('unpunctuated'). In the former case (*TROPHY TROPHY TROPHY*), they typically denote as many objects as there are iterations. In the latter case, three unpunctuated (and thus hard to count) iterations, transcribed as *TROPHY-rep3*, may mean *several* (typically at least three) *trophies*. Strikingly, the existence of repetition-based plurals as well as the distinction between punctuated and unpunctuated ones exists in several sign languages, and was described in a homesigner (deaf person who grew up without access to sign language) who appeared to have invented this device, as it was absent from the gestures of his hearing mother (Coppola et al. 2013).

Hearing non-signers might also understand related distinctions in repeated gestures (Schlenker, to appear a). Schlenker and Lamberton 2019b propose a modular analysis in which the iconic nature of repetition-based plurals as well as the distinction between punctuated and unpunctuated repetitions results from the interaction between four modules: (i) logical grammar makes available a mechanism of existential quantification over pluralities; (ii) pictorial semantics allows the arrangement and number of the repetitions to iconically represent various pluralities of objects; (iii) pictorial semantics can be made blurry so as to create instances of pictorial vagueness, with the result that in unpunctuated repetitions the number of depicted objects is vague; and finally (iv) pictorial vagueness can be exploited pragmatically so as to yield 'at least' readings (akin to standard plurals) for these unpunctuated repetitions.

In sum, the interplay between compositional semantics and iconicity is a key issue in sign language semantics/pragmatics (aspects of spoken language also benefit from iconic analyses, e.g. Dingemanse 2013). Strikingly, since sign languages have the same general grammatical and logical resources as spoken languages, but make greater use of iconicity, sign is along some dimensions more expressive than speech.¹⁰

4.3 Logical and grammatical structures in gestures¹¹

Gestures offer a prime example of iconicity in that their denotation often makes reference to their form; we can easily reproduce the fuel level gauge example with a gesture for how high the water was during the flood (where the angle of a stretched out arm corresponds to the height of the water). However, in contrast to the earlier example, gestures also have a compositional internal structure; for example, a single gesture for jumping downwards/upwards may be decomposed into a part that denotes *manner* ('jumping') and a part denotes *path* (downwards/upwards); compare Özçalışkan et al. (2016) for a related discussion. While they have nothing like the sophisticated grammar of sign languages, gestures sometimes have a proto-grammar that is reminiscent of sign language, e.g. a gesture that describes a telic event with an endpoint will differ from a gesture that describes an atelic event without an endpoint. In performing gestures in such a way, non-signers appear to know constraints that track some sign language rules standardly classified as 'grammatical', and play a particularly important role in linguistic semantics.¹² In order to abstract away from the special semantic issues raised by co-speech gestures, which are in some ways parasitic on the spoken expressions they enrich (a point we revisit in Section 5.1), we focus for the moment on gestures that fully replace some words (henceforth 'pro-speech gestures').

Two examples will make this line of research concrete. In (17), The expression *a mathematician* is pronounced with an open hand (palm up) on the right (transcribed as *IX-hand-a*, and preceding in the transcription the co-occurring expression, which is boldfaced), while the expression *a sociologist* co-occurs with an open hand on the left (transcribed as *IX-hand-b*). With these loci in place, a pointing gesture can fully replace a pronoun that would be expected after *pick*: an index towards the right (transcribed as *IX-a*) serves to refer to the mathematician, associated with locus a.

¹⁰ Note that we use the term *expressive* in the sense of 'expressive power' (see section 2.2.1), not in the more narrow sense of Potts (2005), where *expressive* meanings are contrasted with *descriptive* meanings.

¹¹ Our discussion follows Schlenker, to appear a.

¹² Needless to say, this implies in no way that sign languages are 'merely' gestural: their sophisticated grammars have been described in great detail, and share multiple properties with those of spoken languages (see Sandler and Lillo-Martin 2006). Rather, the argument is that despite their expressive limitations, gestures have a proto-grammar reminiscent of sign language.

- (17) Whenever I can hire IX-hand-a [**a mathematician**] or IX-hand-b [**a sociologist**], I pick IX-a.
Meaning: whenever I can hire a mathematician or a sociologist, I pick the former.

Importantly, if we were to translate (17) into spoken English with *him* or *her* as a substitute for XI-a, *him/her* could be ambiguous between the two antecedents; by contrast, the pointing gesture is not: the gesture is not just a code for a pronoun, it also disambiguates the referent.

A second case, involving grammatical agreement, was studied with experimental means. In ASL, some verbs involve a movement from one locus towards another locus as part of their realization. For instance, *I give you* could be realized with a movement going from the signer to the addressee (it is transcribed as *I-GIVE-2*); *I give him* starts from the signer's position and targets a third person locus, for instance *a* on the right – in which case it is transcribed as *I-GIVE-a*. For this reason, such verbs are often called 'agreement verbs'. The respective loci have been argued to display the behavior of agreement markers, although alternative analyses have been offered as well (Liddell 2003; Lillo-Martin and Meier 2011; Pfau et al. 2018 and Schembri et al. 2018). Irrespective of the final analysis, these agreement verbs appear to have gestural counterparts (Schlenker and Chemla 2018).


To have a point of comparison, let us consider the ASL examples in (18), constructed around the agreement verb *I-GIVE-2* or *I-GIVE-a*.

- (18) POSS-2 YOUNG BROTHER MONEY IX-1 **1-GIVE-a**.
 'Your younger brother, I would give money to.
 a. IX-2 IX-1 NOT.
 You, I wouldn't.'
 b. *IX-2 IX-1 NOT **1-GIVE-a**.

In the first sentence of (18), the verb *GIVE* displays object agreement with a third person locus *a*, corresponding to the younger brother, hence: *I-GIVE-a*. Now the continuation in (18)a involves a missing Verb Phrase, which in spoken and sign language alike usually entails copying an antecedent. But something happens in the copying process: the agreement markers can be disregarded, which explains why the continuation (18)a is acceptable even though that in (18)b (with overt copying) isn't. Similar rules of partial copying under ellipsis are attested in English, as in (19), where the third person features and the feminine features of *her* are disregarded under ellipsis, and thus *I did too* has a reading that *I too did her homework* lacks.



- (19) [Uttered by a male speaker] In my study group, Mary did her homework, and I did too.
can mean: I too did **my** homework

The ASL contrasts can be replicated with gestural verbs in English. As in ASL, a movement towards the side has to correspond to a third person object, and thus the second clause of (20)b is deviant because the object is second person but the object agreement is third person (the kisses ought to be sent towards the addressee). Strikingly, this problem doesn't really arise in (20)a: as the missing verb is copied, its third person object agreement is disregarded, just as is the case for object agreement in ASL. Importantly, there are no comparable cases of object agreement in English: subjects had to infer an ASL-style rule on the basis of no or extremely limited evidence.



- (20) a. Your brother, I am gonna SEND-KISSES-3_ then you, too.

b. *Your brother, I am gonna PUNCH-3_



, then you, I am gonna SEND-KISSES-3_

It is noteworthy that homesigners, who grow up without access to sign language, develop gestural languages that share some properties of sign languages, but are expressively far more limited (e.g. Abner et al. 2015, Goldin-Meadow 2003). In some cases, the reason homesigners discover such properties on their own might be that, more generally, non-signers 'know' them; a case can for instance be made for repetition-based plurals, discussed above, which (i) arise without input in some homesigners, and (ii) might be intuitively known by non-signers perceiving gestures (see Coppola et al. 2013, Schlenker and Lamberton, 2019b).¹³

An important question for future research is to determine how such instances of spontaneous grammatical learning of sign language features by non-signers can arise. Proponents of a Universal Grammar (UG) approach to human language may go as far as to argue that UG does not just specify grammatical rules in an abstract form, but also specifies part of the mapping between forms and grammatical/semantic content: a pointing sign / gesture might thus be intrinsically endowed with pronominal properties. A possible alternative view is that some signs/gestures are naturally associated with a fixed grammatical / semantic component for deeper cognitive reasons. This debate is currently open.

5 Linguistics beyond words I: iconic elements within the inferential typology

We turn to the place of iconic meanings in the typology of linguistic inferences. Here an important distinction must be drawn between iconic enrichments, which modulate the form and meaning of a word, and iconic replacements, which fully replace a word. Iconic enrichments, illustrated by the lengthening of the word *long* to *loooooong*, have a special status due to their parasitic nature as add-ons to the main message; if we say that a talk was *loooooong*, the main message is that it is long, while the precise meaning that we communicate includes an additional encoding of its duration by virtue of a lengthened vowel. Iconic replacements, where an element with iconic semantics (such as a gesture) altogether replaces a word that would communicate the same content, make a different but powerful point: they fall within the very same typology as regular words; this provides new insights about the cognitive origin of this typology.

To introduce the typological issue and the terminology, let us consider the examples in (21). (21)a involves a slapping gesture that co-occurs with the verb *punish*. In (21)b, a post-speech slapping gesture appears after the Verb Phrase it modifies.¹⁴ In (21)c, the slapping gesture fully replaces the verb and is thus a pro-speech gesture. Finally, in (21)d, a conventional word, *long*, is modified in an iconic fashion by way of an 'iconic modulation' (which by definition is always the modification of a conventional form). The same terminology is extended to sign language by replacing *-speech* with *-sign* (especially *co-sign* and *post-sign* gestures).

¹³ It has also been argued that several further properties of sign language grammar can be replicated with pro-speech gestures (Schlenker, to appear a). If confirmed, these results might be important to understand the historical origins of sign languages.

¹⁴ See Esipova 2019 for a broader discussion.

(21) a. Co-speech gestures (co-occur with the word they modify [boldfaced])



Asterix will **punish** his enemy.

b. Post-speech gestures (follow the word they modify)



Asterix will punish his enemy – .

c. Pro-speech gestures (replace a word)



His enemy, Asterix will .

d. Iconic modulations (modify the form of a conventional word)

The talk was loooooong.

We start by considering co-speech gestures, iconic modulations and post-speech gestures, which are iconic enrichments.¹⁵ We then turn to pro-speech gestures, which are iconic replacements.

5.1 Iconic enrichments in the typology of linguistic inferences: sign with iconicity vs. speech with gestures

One motivation for investigating semantic enrichments is to assess in greater detail our earlier conclusion that sign languages are, along some dimensions, more expressive than spoken languages because they make greater use of iconicity. This conclusion might be premature because it fails to take into account the means of iconic enrichment afforded to speech by gestures: sign with iconicity should be compared to speech with gestures rather than to speech alone (Goldin-Meadow and Brentari 2017). But *even* when gestures are taken into account, systematic differences arguably remain between sign with iconicity and speech with gestures. The reason has to do with the inferential typology discussed in Section 2. The most salient means of iconic enrichment in speech lies in co-speech gestures. But it was argued from the start (notably in pioneering work by Ebert and Ebert 2014) that co-speech gestures do not typically make at issue contributions. By contrast, several iconic modulations of signs seem to make at issue contributions. We already mentioned this point in connection with *GROW* above, and the shape modulations of repetition-based plurals apparently make at issue contributions as well (Schlenker and Lamberton 2019a).

Granting that gestures have an iconic component reminiscent of iconic signs, the key question is *how* these iconic enrichments are effected, and distributed across the inferential typology. As we will now see, none of the gestures in (21)a-c have quite the same properties as iconic modulations. *Co-speech gestures* trigger presuppositions of a particular sort (called ‘cosuppositions’, sometimes used as a mnemonic for ‘conditionalized presuppositions’¹⁶),

¹⁵ As we will see, post-speech gestures behave semantically as non-restrictive relative clauses, and could thus be taken to be gestures that replace certain words or phrases that are optional.




¹⁶ Note that the term ‘conditionalized presuppositions’ in connection with ‘cosuppositions’ was introduced by Tieu et al. 2017:6. In Schlenker 2018a:297 (written before Tieu et al. 2017), ‘co’ in ‘cosuppositions’ was justified differently, reflecting the technical implementation of the analysis: the terminology was “intended to suggest that a cosupposition triggered in a local context *c*’ is computed in tandem with (‘co’) an at-issue component in *c*’ (by contrast, a standard presupposition triggered in *c*’ is computed before (‘pre’) any at-issue component in *c*’).”

and thus fail to be at issue, unlike many iconic modulations found in sign language. *Post-speech gestures* trigger ‘supplements’, which is the type of meaning encoded by non-restrictive relative clauses (Potts 2005); for this reason, they too fail to be at-issue. *Pro-speech gestures*, for their part, make at-issue contributions, but unlike signs (including ones with iconic modulations) they are not conventional words at all, and are thus expressively limited for other reasons. The typology is illustrated in (22), and it turns out to be crucial to answer (in the negative, we think) the question whether speech with gestures ‘equals’ sign with iconicity.

(22) **Typology of iconic enrichments and replacements** (after Schlenker 2018b)

	Co-speech/co-sign gestures	Post-speech/post-sign gestures	Iconic modulations	Pro-speech/pro-sign gestures
Meaning	cosuppositions (= presuppositions of a special sort)	supplements (like non-restrictive relative clauses)	at-issue or not, depending on the case	at-issue, with an additional non-at-issue component in some cases
Speech	Asterix will  punish his enemy.	Asterix will punish his enemy  .	The talk was loooong.	His enemy, Asterix is going to  .
Sign	IX-arc-b NEVER  [SPEND MONEY]	IX-arc-b NEVER SPEND  MONEY], –	POSS-1 GROUP GROW_ 	[currently unclear]

To illustrate this typology, we note that the iconic enrichments in the positive sentences in (21) display radically different behaviors under negation, as seen in (23).

- (23) a. Asterix won't  punish his enemy.
=> if I were to punish my enemy, slapping would be involved
- b. #Asterix won't punish his enemy .
- c. His enemy, Asterix won't .
- d. The talk wasn't loooong.

Most centrally, co-speech gestures and post-speech gestures differ in their acceptability when they combine with a negated Verb Phrase.


(i) First, the co-speech gesture in (23)a is acceptable when accompanying a Verb Phrase that is in the scope of negation; it triggers an inference that gets inherited across negation, to the effect that *if I were to punish my enemy, slapping would be involved*. Crucially, further tests suggest that this inference behaves like a presupposition (for instance, embedding under *none of the girls* will give rise to a universal inference that holds for all of the girls); it has received a special name (cosupposition) because the inference is conditionalized on the meaning of the modified expression (here: *punish*). Experimental data (Tieu et al. 2017, 2018) have buttressed this conclusion by means of truth value judgment tasks and (more clearly) inferential judgment tasks.

(ii) By contrast, the post-speech gesture in (23)b is simply deviant after a negated Verb Phrase. Recent literature (Schlenker 2018b) has argued that this is because the post-speech gesture behaves like a non-restrictive relative clause and thus contributes supplement meaning, which is often deviant in such negative environments, as in (24)a.

- (24) #Asterix won't punish his enemy, which will involve slapping him.

Consideration of further examples strengthens the similarity with non-restrictive relative clauses: (25)a behaves like (23)a in suggesting that Asterix's action will involve slapping, unlike the control with a conjunction in (25)b.¹⁷ Such inferential contrasts were established with experimental means in Tieu et al. 2019, and they extend to cases in which a visual animation replaces the post-speech gesture – a point to which we return in Section 5.2.2.



- (25) a. If Asterix punishes his enemy – , I might scream.
=> if Asterix punishes his enemy, slapping will be involved
b. If Asterix punishes his enemy and this involves slapping him, I might scream.
≠> if Asterix punishes his enemy, slapping will be involved

(iii) Third, the pro-speech gesture in (23)c makes an at-issue contribution and yields neither a cosupposition nor a supplement. Importantly, however, pro-speech gestures are expressively limited because they are not based on conventional words, unlike the iconic modulations found in sign language: in LSF, the meaning of *UNDERSTAND* or *REFLECT* can be modulated by altering the speed with which part of the sign is realized (e.g. to indicate a quick or difficult understanding or reflection, Schlenker 2018c). It seems hopeless to represent modulations of such abstract concepts with pure gestures.

(iv) Fourth, the iconic modulation in (23)d makes an at-issue contribution and triggers no conditionalized inference akin to cosuppositions.

Importantly, with the exception of pro-sign gestures (i.e. sign-replacing gestures, whose existence and status is still somewhat unclear), the same typology might hold in sign (Schlenker 2018b). A disgusted facial expression co-occurring with the Verb Phrase *SPEND MONEY* (as on the last line of (22)) yields the same cosuppositional behavior as the slapping gesture co-occurring with *punish*: 'if x spends money, this is disgusting'.¹⁸ The same facial expression could also follow the Verb Phrase, in which case it arguably behaves as a post-speech gesture and plausibly triggers a supplement. Finally, as argued above, the iconic modulations of *GROW* are best compared to those of *loooong*: at issue contributions are made in both cases.

Three conclusions are worth highlighting. First, iconic enrichments (co-speech and post-speech gestures as well as iconic modulations) can be handled by established semantic mechanisms, although one case (cosuppositions) requires adjustments to presupposition theory (by introducing conditionalized presuppositions). Second, there is no categorical difference between iconic enrichments in speech and in sign language: the same abstract typology is found in both modalities. Third, gestural enrichments do not make the same type of (at issue) contribution as iconic modulations. But iconic modulations are arguably rich in sign language, and impoverished in speech. This yields systematic differences between signing with iconicity and speech with gestures.

Importantly, the typology in (22) could be expected to apply to further types of enrichments, such as 'vocal gestures' (e.g. Schlenker 2018b). Pasternak 2019 and Pasternak and Tieu 2020 argue that sounds that co-occur with speech (e.g. sound effects in radio drama) behave like co-speech gestures and yield an inferential profile characteristic of gestural cosuppositions. Perhaps surprisingly, we will see in Section 6.4 that music that co-occurs with films or cartoons might behave like co-speech gestures in triggering cosuppositions.

¹⁷ These inferential facts matter because the deviance of (23)b could be analogized to (i) below rather than to (24). But in view of the difference between (25)a and (25)b, the analogy with (i) seems wrong-headed.

(i) #Asterix won't punish his enemy, and this will involve slapping him.

¹⁸ Importantly, this facial expression is not a grammatical marker, unlike other sign language facial expressions: it seems genuinely gestural.

5.2 Iconic replacements in the typology of linguistic inferences: replicating the typology without words

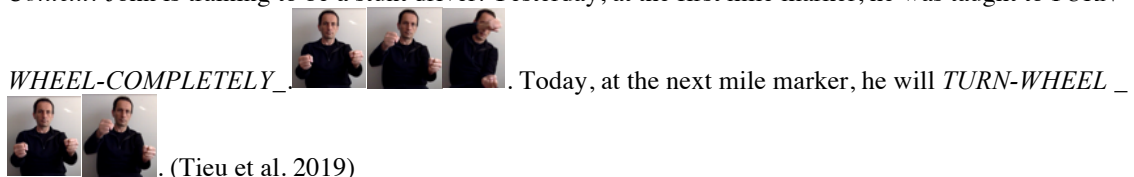
The semantic difference between co-, post- and pro-speech gestures is certainly due to the *manner* in which they are realized, namely as co-occurring, following or replacing a word (Schlenker 2018b and Esipova 2019 for possible explanations). But in addition, recent research suggests that pro-speech gestures alone neatly fill established categories of the inferential typology of language, which includes not just at issue entailments and supplements, but also implicatures and (standard, non-cosuppositional) presuppositions, among others: depending on their informational content, they may make *additional* contributions that reflect inferential types (and probably algorithms) that are found in normal words.

These gestural findings were obtained with experimental means in Tieu et al. 2019. Moreover, Tieu et al. 2019 goes one step further and replicates the typology in paradigms in which gestures are replaced with visual animations that the subjects could not have seen in a linguistic context before. We observe a high degree of productivity when it comes to dividing new semantic content 'on the fly' among established categories of the inferential typology. This argues for the existence of cognitive algorithms that make it possible to do so. For brevity, we discuss just two cases: iconic implicatures and iconic presuppositions.

5.2.1 Iconic implicatures

In some cases, the existence of inferences that arise 'on the fly' is expected by current theories. Consider the case of scalar implicatures. In (26), a gesture representing a partial wheel-turning is contrasted with a complete wheel-turning. It can be checked in separate examples that *not TURN-WHEEL* can mean 'not turn the wheel at all' (rather than 'not turn the wheel exactly as depicted', for instance). This suggests that the partial wheel-turning (i.e. *TURN-WHEEL*) can have a weak meaning, akin to 'turn the wheel'. But as soon as this gesture evokes (thanks to the context) a more informative alternative *TURN-WHEEL-COMPLETELY*, an implicature is derived: *John will TURN-WHEEL* is understood to mean that John will turn the wheel but not completely. This was shown in experiments where participants endorse the target inference in (27)a significantly more than the control inference in (27)b.

(26) *Context*: John is training to be a stunt driver. Yesterday, at the first mile marker, he was taught to *TURN-*



- (27) a. Target inference: John will turn the wheel, but not completely.
b. Baseline inference: John will turn the wheel completely.


Tieu et al. 2019 show in separate conditions that the gestures involved are unlikely to be mere codes for words because they have iconic implications that mere words would lack, for instance about the size of the wheel. In addition, similar results are obtained when gestures are replaced with artificial visual animations that subjects couldn't have seen before, with similar results. As expected, then, implicature derivation is a fully productive process.

5.2.2 Iconic presuppositions

In contrast with scalar implicatures, presuppositions are typically thought to be encoded in the lexical meaning of words (e.g. Heim 1983), although there have been various attempts to propose 'triggering algorithms' that *deduce* the presupposition of an expression from its

informational content (see for instance Abrusán 2011 and Schlenker, to appear b, for discussion). Strikingly, pro-speech gestures can trigger presuppositions, as can be illustrated by a modification of our *TURN-WHEEL* example: the question in (28) triggers the inference that Sally is behind the wheel, hence a significantly stronger endorsement of the inference in (29)a than in (29)b. Many further tests have been used in the literature to argue that this and other gestural examples genuinely trigger presuppositions (e.g. Schlenker 2019c).




(28) Jake and Lily are watching their four children ride bumper cars at the carnival. Each bumper car has two seats. As one of the bumper cars nears a bend in the track, the parents wonder:

Will Sally *TURN-WHEEL*  ?

- (29) a. Target inference: Sally is in the driver's seat.
 b. Baseline inference: Sally is in the passenger seat, not the driver's seat.

Going one step further, Tieu et al. 2019 carry out an experiment showing that the generation 'on the fly' of presuppositions from iconic material extends to stimuli that subjects could not possibly have seen before. For instance, take a visual animation, described in (30), which represents an alien changing color from their normal state (green) to their meditating state (blue); for participants in the experiment, a change from green to blue triggered an inference that the aliens were not initially meditating, which is in line with how natural language expressions and gestures behave. This inference is preserved in questions, a characteristic property of presuppositions (illustrated in (7)a in section 2.3.2). Further tests and further pictorial representations that represent changes of state were used to buttress the conclusion that visual animations too can trigger presuppositions.

- (30) **Pictures from Tieu et al.'s videos testing presuppositions generated by visual animations**
 (here: a change of state animation pertaining to an alien's antenna turning from green to blue; original video: <https://youtu.be/U6dfs-XI2-4>)

<p>Aliens are green. But when they are in a meditative state, their antennae are blue.</p>	<p>There is a meditation session in progress on the first floor of a business firm.</p>	<p>Bill is watching the union representative and says:</p>	<p>"Will the union representative's antenna..."</p>
			
(green)	(green+blue)	blue	

5.2.3 Going further

The productive nature of implicature and presupposition generation extends to further inferential types. For instance, the contrasts illustrated in (25), pertaining to supplements, were obtained with gestures and visual animations in Tieu et al.'s experiment. Still further inferential types (e.g. pertaining to so-called 'homogeneity inferences') follow the same logic (Schlenker 2019c, Tieu et al. 2019), suggesting that pictorial content can be productively divided among a rich inferential typology. Migotti and Guerrini 2019 argue that the main results can be extended to auditory stimuli, specifically to pro- and post-speech onomatopoeias and musical content, highlighting the generality of the phenomenon: the inferential typology can be replicated not just with gestures and visual animations but also with diverse non-linguistic sounds.

A key question for future research is whether this productive division of semantic content among established slots of the inferential typology only arises when the stimuli are embedded in sentences, or whether some of them might arise in visual narrations that are not

linguistically embedded (and are thus used in a more standard fashion) – a tantalizing if remote possibility.

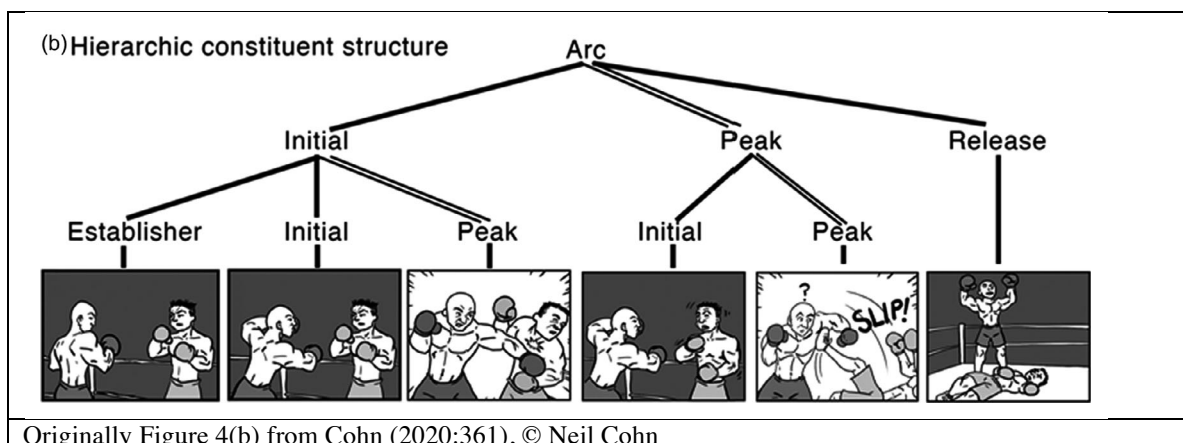
6 Linguistics beyond words II: visual and musical narratives

6.1 Visual narratives

6.1.1 Syntax of visual narratives

Turning to visual narratives, such as comics, recent literature (Cohn 2012, 2018, 2020) argues that an independent syntax of visual narrative is required in addition to a semantic analysis (see section 6.1.2). Initial studies of visual sequences, discussed in Cohn et al. 2012, replicate the core point behind Chomsky’s (1965) *Colorless green ideas sleep furiously*, to demonstrate the existence of syntactic structure in the absence of semantics. The experiments compared (i) comics that told semantically coherent narratives to (ii) comics that followed narrative structuring principles (i.e. syntax) but lacked semantic connection between the panels, as well as (iii) scrambled comics (which lacked narrative structure, thus lacking both syntax and a semantic connection). While the latter two (ii+iii) gave rise to increased N400 brain responses, indicating a problem with semantic processing, the meaningless narratives that obeyed structuring principles (ii) were closer to the coherent narratives (i) in response times.¹⁹

In line with a syntactic analysis of non-standard objects of study such as music and dance, higher level and lower level grouping are core tenets of this approach to visual narratives. Cohn proposes that their structure consists of narrative categories (e.g. *Establisher*, *Initial*, *Peak* and *Release*), grouping (constituency), and hierarchical structure, similar to those present in natural language. An application of syntactic analysis to comics is illustrated by the example below,²⁰ a comic strip depicting a boxing match, where one of the fighters slips (presumably on purpose) and forfeits the match, after having been struck just once. The depicted grouping structure reflects the ebb and flow of tension in the narrative (Initial-Peak-Release), modeled by the assignment of narrative categories to the panels; the four main narrative categories are defined in (1).



¹⁹ Cohn (2018) presents his syntactic approach as a critique of Bateman & Wildfeuer’s (2014) syntax-free semantic approach. The issue in that debate is that it seems self-evident that visual narrative must have a semantics, but less self-evident that it has a syntax. We discuss the semantics of visual narrative in section 6.1.2.

²⁰ The arrangement of panels in actual comics varies; for the purposes of analysis, the panels are given in a linear (left-to-right) fashion; see Cohn (2018) for discussion.

Each narrative category is determined on the basis of diagnostics (see Cohn 2015 for details). Although the syntax of visual narrative employs grouping and constituency, the narrative categories themselves encode discourse semantic properties of the panels.

- (1) *Basic narrative categories* (quoted from Cohn 2018:2-3)²¹
- a. Establisher (E) – sets up an interaction without acting upon it, often as a passive state.
 - b. Initial (I) – initiates the tension of the narrative arc, prototypically a preparatory action and/or a source of a path.
 - c. Peak (P) – marks the height of narrative tension and point of maximal event structure, prototypically a completed action and/or goal of a path, but also often an interrupted action.
 - d. Release (R) – releases the tension of the interaction, prototypically the coda or aftermath of an action.

While grouping structure is assumed to arise in the mind of a comic reader, a researcher can utilize diagnostics to generate and analyze it: the first step involves identifying the Peak(s) in the visual sequence, from which a researcher works backwards to identify the Initial and Establisher (building up to the Peak) and finally the Release²². Higher/intermediate level grouping applies the same narrative categories in a recursive manner; in the example above, the first Peak is less prominent than the second one, where the narrative tension culminates. This correlates with the observation that the narrative is disrupted more when the second Peak is deleted than when the first Peak is deleted.²³ Consequentially, at the higher level of structure, the constituent containing the second Peak functions as the (higher level) Peak, whereas the constituent containing the first Peak becomes the Initial.²⁴ Crucially, if deletion of the first Peak in a similar example had a bigger effect than deletion of the second Peak, then different narrative categories would emerge at the higher level: the constituent containing the first Peak would function as the higher-level Peak, while the constituent containing the second Peak would be part of the Release.

6.1.2 *Semantics of visual narratives*

A gulf separates the projection-based semantic of pictures introduced in Section 4.1 from actual pictorial narratives, spanning from comics to films: formal pictorial semantics is in its infancy. But important progress is being made on several fronts. First, pioneering work on the visual language of comics is getting the recognition it deserves, from sophisticated analyses of visual morphology (e.g. reduplication used to evoke movement) to a theory of narrative structure in comics (Cohn 2013). Second, Cumming et al. 2017 investigate in formal detail constraints on viewpoint shift in film, i.e. permissible changes of camera angle. Third, Abusch and Rooth 2017 have upgraded Greenberg's projection-based semantics 'from the ground up' to analyze visual narratives. The simplest example appears in (31), which represents "a short comic of two cubes moving apart" (Abusch and Rooth 2017).


²¹ See Cohn (2018) for further refinements of the system.

²² This is quite reminiscent of syllable structure, Onset-Nucleus-Coda, in phonology.


²³ Deletion is one of the diagnostics used to identify narrative categories, see Cohn (2015).

²⁴ The top-level (root) node *Arc* contains the entire narrative (story arc), and serves no further narrative role.

(31) Picture P1



Picture P2



The fact that the two pictures are arranged as a narrative sequence provides information about the situations described, as well as their ordering. In essence, a series of n pictures $\langle P_1, \dots, P_n \rangle$ is true of n situations $\langle s_1, \dots, s_n \rangle$ just in case the s_i 's are temporally ordered in the right way, and each s_i projects onto the corresponding P_i , as stated in (32), which is a temporal version of (15).

(32) **Picture sequences true of n situations (after Abusch)**

A picture sequence $\langle P_1, \dots, P_n \rangle$ is true of situations $\langle s_1, \dots, s_n \rangle$ **relative to viewpoint v** along the system of projection S just in case:

- (1) temporally, $s_1 < \dots < s_n$;
- (2) $\text{proj}_S(s_1, v) = P_1$ and ... and $\text{proj}_S(s_n, v) = P_n$.

While the situation-based analysis of picture sequences in (31) is a powerful tool to study visual narrative, Abusch convincingly argues that this framework is yet insufficient, in particular because narrative sequences give rise to ambiguities of cross-reference. Concretely: the second picture of (31) is most naturally interpreted as involving the same cubes as the first, but nothing fully excludes the possibility that the dark cube disappeared and was replaced by another dark cube a bit further away. A semantics that is based on temporal sequences of pictorial situation-descriptions, (31), doesn't suffice to derive the most salient reading (same cubes in both pictures), nor the ambiguity that permits less plausible readings (where one or both cubes have been replaced by identical cubes). Abusch 2013 suggests that the salient interpretation is obtained because visual representations contain abstract variables (which she relates to Pylyshin's 2003 indexes in vision). A projection-based semantics can thus be combined with variables to derive the meaning of visual narratives (see Abusch and Rooth 2017 for further pictorial operations possibly reminiscent of language; in terms of the semantic model, Abusch 2013 implements her approach in Discourse Representation Theory, an approach to pictorial narratives that is shared with Bateman & Wildfeuer 2014, Maier & Bimpikou 2019, and Maier 2019).

6.2 Musical narratives

6.2.1 Basic music syntax

The analysis of Western tonal music has led to the belief that music has a sophisticated syntax. Upon venturing into this domain, researchers find themselves with an embarrassment of plenty: different authors posit different syntactic systems, sometimes with different goals.

In their pioneering study of musical syntax, Lerdahl and Jackendoff 1983 posit four levels of structure. As summarized in Lerdahl 2001 (bold type added for emphasis), their theory proposes:

“four types of hierarchical structure simultaneously associated with a musical surface. **Grouping structure** describes the listener's segmentation of the music into units such as motives, phrases, and sections. **Metrical structure** assigns a hierarchy of strong and weak beats. **Time-span reduction**, the primary link between rhythm and pitch, establishes the relative structural importance of events within the rhythmic units of a piece. **Prolongational reduction** develops a second hierarchy of events in terms of perceived patterns of tension and relaxation.”

To illustrate, (33) provides the metrical structure (the square horizontal brackets beneath the musical score) and the grouping structure (the round horizontal brackets beneath the musical score) of the beginning of Mozart's sonata K. 331.²⁵ To begin with, a reader should focus on the round horizontal brackets (the grouping structure), which we discuss first.

- (33) Metrical structure [square brackets] and grouping structure [round brackets] of the beginning of Mozart's K. 331 piano sonata (Lerdahl and Jackendoff 1983)

POSSIBLE LINKS: https://www.youtube.com/watch?v=dJEwT_0UBmg&t=7s

The image shows a musical score for the beginning of Mozart's sonata K. 331. The score consists of two staves: a treble clef staff and a bass clef staff. Below the notes, there are two levels of brackets. The first level consists of square brackets indicating metrical structure, with each bracket spanning a single eighth note. The second level consists of round brackets indicating grouping structure, with brackets grouping notes into hierarchical levels: first, individual notes are grouped; then, pairs of notes are grouped; then, groups of three notes are grouped; and finally, the entire phrase is grouped under a single large round bracket.

As illustrated by the round brackets in (33), grouping yields a hierarchical, tree-like structure (the bottom-most round bracket including all other brackets, with two round brackets in the next line above it, which corresponds to a lower level grouping). Such **grouping structure** is "an auditory analog of the partitioning of the visual field into objects, parts of objects, and parts of parts of objects" (Lerdahl and Jackendoff 1983 p. 36). Principles of grouping thus derive from Gestalt principles of perception (Wertheimer 1923), with local preference rules that specify, for instance, that group boundaries are preferably created by (i) large pitch intervals, and (ii) pauses.²⁶ To illustrate (i), the third note in (34)a,b is preferably grouped with the first two notes because there is a large pitch interval between the 3rd and 4th note. Contrastively, the third note is preferably grouped with the last two notes in (34)c,d, because there is a large pitch interval between the 2nd and the 3rd note in (34)c,d.

(34) a. b. c. d.

The image shows four musical examples, labeled a, b, c, and d, each consisting of a single staff with five notes. Example a shows a sequence of notes with a large pitch interval between the third and fourth notes, and a smaller interval between the second and third notes. Example b shows a similar sequence but with a different interval structure. Example c shows a sequence where the third note is grouped with the last two notes. Example d shows a sequence where the third note is grouped with the first two notes.

To illustrate (ii), the pauses (rests, R) in (35)a-c create group boundaries, and in fact override the preference that might be expected from (i): in (35)b,c, despite the large pitch interval between the second and the third note, they are grouped together due to the pause that separates the first three notes from the last two.

(35) a. b. c.

The image shows three musical examples, labeled a, b, and c, each consisting of a single staff with five notes. Example a shows a sequence of notes with a pause (rest) after the second note. Example b shows a sequence of notes with a pause (rest) after the third note. Example c shows a sequence of notes with a pause (rest) after the fourth note.

Turning to **metrical structure**, this is given by the alternation of strong and weak beats, as in phonology, organized in a hierarchical structure: The initial level is given by the *tactus*, which is "the level of beats that is conducted and with which one most naturally coordinates foot tapping and dance step". It is given by eighth notes (i.e. ♪/♪) in (33) and (36), with occasional subdivisions below this level when smaller temporal units appear (as

²⁵ As Lerdahl and Jackendoff 1983 point out (p. 63), an alternative grouping is conceivable, with the last E of the 1st bar grouped with the 2nd bar, and with the last D of the 2nd bar grouped with the 3rd bar.

²⁶ Further principles help determine group boundaries at higher levels, e.g. pertaining to the optimal size of groups, or to optimal parallelism among groups that look similar.

with the second note of the 1st bar in (35), which is at the sixteenth note level, ♩). At any level, the strongest beat among a series of two or three beats can 'project' to the next level.²⁷ This projection is quite reminiscent of Peaks in the syntax of visual narratives, Section 6.1, indicating a formal commonality between two entirely different systems.

(36) Metrical structure of the beginning Mozart's K. 331 piano sonata (Lerdahl and Jackendoff 1983 p. 72)



The third type of hierarchical structure in music is **time-span reduction**.²⁸ To motivate it, Lerdahl and Jackendoff argue that their grouping structures are insufficient in that they fail to distinguish different levels of importance within musical groups. Lerdahl & Jackendoff thus propose that their tree structures are *headed*: each group at each level contains one musical event that is more important than the others and thus counts as its 'head'. In a nutshell, heads are events that are (i) rhythmically/metrical more prominent, such as strong beats, and/or (ii) harmonically more stable, in the sense that in the key of C major a C chord (= a tonic chord, notated in general as I) is more stable than a G chord (= a dominant chord, notated as V). Metrical structure helps select the most important notes at micro-levels, while at large levels, heads of musical groups are selected by a combination of both metrical and harmonic considerations. Thus a researcher can derive from a metrical and grouping structure as in (33) a time-span structure as in (37), where certain chords (notated with Roman numerals) are represented as the heads of the various groups. This serves as the connection between rhythm/meter and melody/pitch.

(37) Time-span reduction obtained from (33) by selecting in each the musical event which is metrically strongest/harmonically most stable (Lerdahl and Jackendoff 1983)



The outcomes of calculating a time-span reduction are headed trees that are somewhat reminiscent of natural-language syntax, but obtained from entirely different considerations (e.g. unlike syntactic constituents in language, grouping structure in music is derived from principles of general perception²⁹). Lerdahl and Jackendoff 1983 further posit that yet another structure, called **prolongational reduction**, is needed to provide a hierarchy of events "in terms of perceived patterns of tension and relaxation" (Lerdahl 2001), a specifically harmonic notion. Prolongational reduction may well be the hierarchical structure, out of the four that were introduced for music, that is most interesting from a purely musical

²⁷ A reader with training in ballroom dancing may find it useful to think of the 3/4 meter of waltz, which gives the waltz its distinct 1–2–3–1–2–3–... rhythm, associated with a **strong(1)**-weak(2)-weak(3) realization.

²⁸ The summary in this paragraph (including some of the wording) follows Schlenker (2019a).

²⁹ Katz 2018 has emphasized similarities between grouping structure in music and prosodic grouping in language (rather than syntactic constituency).

perspective; it is also the one that is most strongly tied to music in the sense that we would be hard-pressed to find it, say, in pictorial narrative (section 6.1) or in dance (section 6.3), where we do find counterparts of grouping structure.

Its musical importance is evident in more recent research on music syntax, where the analysis of 'harmonic syntax', which roughly corresponds to prolongational reductions, has given rise to active debates. While Lerdahl and Jackendoff 1983 stated all their structural rules as parsing rules (with preference principles guiding the listener towards preferred structures), Pesetsky and Katz 2009 propose to define the same system in terms of generation, highlighting the similarity with linguistic syntax. Departing even further from Lerdahl and Jackendoff's framework, Rohrmeier 2011 proposes phrase-structure rules to account for certain aspects of harmonic syntax in Western tonal music, while Granroth-Wilding and Steedman 2014 propose a combinatory categorial grammar (and parser) for jazz harmonic progressions. The latter is even associated with a semantics, one that encodes motion in tonal pitch space.

Strikingly, all these frameworks propose syntactic formalisms that go beyond the finite-state bound, and some (Rohrmeier 2011) have developed precise formal arguments to the effect that this is not just a matter of convenience: some properties of (Western tonal) music necessitate recursive mechanisms that are not finite-state. While the debate is still to some extent open (see Rohmeier et al. 2015 for discussion),³⁰ the overall picture is that music cannot be accounted for by purely finite-state methods.

6.2.2 *Basic music semantics*

While it is relatively uncontroversial that music has a sophisticated syntax, the existence of a music semantics is considerably more debated. By "music semantics", we will mean here a semantics in the usual sense, i.e. one that provides information about the extra-musical world.³¹ But is there such a thing?

To see that pitch alone can convey meaning, consider a musical variant of the fuel level gauge that we briefly discussed at the beginning of section 4. Imagine a scenario where a pilot pushes a button and hears a single tone in response; imagine that the pitch of the tone strictly corresponds to the amount of fuel in the tank (the higher the pitch, the more fuel in the tank); this would be a simple case of iconic semantics in the sense of Greenberg (2021).

Of course, music is much more complex than this toy example, but related (albeit more abstract) iconic inferences are broadly attested. There are numerous experimental and introspective investigations of inferences triggered by actual musical pieces about the extra-musical world. An initial list of systematic effects appears in (38); this list makes reference to an object (sometimes called a 'virtual source') that is described by the music. To illustrate (38)a(i), the French composer Saint-Saëns famously used a double bass to evoke an elephant, the virtual source, in his *Carnival of the Animals*. If a MIDI rendition is raised by three

³⁰ Beyond Western tonal music, Rohrmeier et al. 2015 write: "There are also analytical findings that suggest that principles of hierarchical organization may be found in classical North Indian music that is based on a tradition of extensive oral teaching. However, more cross-cultural research on other cultures and structure in more informal and improvised music is required [...]"

³¹ The terminology can occasionally be confusing: for several authors, musical meaning has to do with inferences about certain properties of the musical *form*, rather than about a music-external reality. For instance, Meyer 1956 writes that "one musical event (...) has meaning because it points to and makes us expect another musical event"; this gives rise to expectations and emotions that constitute what Meyer calls "embodied meaning". Huron 2006 argues that various emotions of a musical or extra-musical nature derive from general properties of expectation, or in other words of our attempts to anticipate what will come next, in music or elsewhere: the inferences are thus about musical form, although they have emotional consequences. As for the semantics of Granroth-Wilding and Steedman 2014, it pertains to movement in tonal pitch space, not to the extra-musical world.

octaves so as to be high-pitched rather than low-pitched, the inference is reversed: instead of a large object, a small one is evoked (original: <http://bit.ly/2mea8pQl>; three octaves higher: <http://bit.ly/2Cl6Xhk>). Potential denotations are diverse: they are not limited to movable objects such as animals, but could be landscapes described by the music, or even (in refinements of the analysis) emotions.

- (38) Examples of inferential effects (Schlenker 2019b Appendix II, with links to examples)
- a. Lower pitch may indicate that the denoted object (i) is larger, or (ii) is less excited/energetic.
 - b. Lower loudness may indicate that the denoted object is (i) less energetic, or (ii) further away.
 - c. Lower speed may indicate that the denoted object is slower.
 - d. Silence may indicate that an event is interrupted.
 - e. Lesser harmonic stability may indicate that the denoted is in a less stable (i) physical or (ii) emotional position.
 - f. A change of key may indicate that the denoted object is moving to a new environment.

How should these inferences from music be explained?³² In auditory or visual perception, agents seek to find information about the causal sources of their percepts. In auditory perception, certain sounds reach the human ear and, depending on their properties, give rise to information about the surrounding objects and events: one may hear a low-frequency call produced by an animal and infer that it is very large; or one may hear a car engine whose loudness decreases and infer that it is moving away. It has been suggested (e.g. Schlenker 2017, 2019a, to appear d) that the same general idea applies to music semantics: the listener seeks to draw inferences about certain objects (e.g. an imagined animal in Saint-Saëns' narrative) based on some properties of the musical sounds (such as low pitch). These acoustic properties usually trigger certain inferences about their causal source (for instance: low pitched sounds tend to be produced by large objects and animals), and these inferences are then applied to the fictional objects the music is 'about'.

Inferential rules of the type in (38) have been argued to be of two kinds. Some are derived from normal auditory cognition, as in the case of a low-frequency sound used to evoke a large object: the inferential rule in (38)a(i) presumably owes to the fact that large resonance chambers produce lower-pitched sounds. Similarly, all other things being equal a sound is perceived as softer when its source is further away, explaining the inferential rule in (38)b(ii).³³ But other inferences are more specifically musical in nature, particularly when it comes to harmonic notions. A key notion in music is that of consonance vs. dissonance, and it can produce powerful inferential effects. To illustrate, Saint Saëns uses a radically slowed down version of the French *Can Can* dance to evoke tortoises (see <https://youtu.be/6HQqaKEz4tg>). Later in the piece, dissonances were suggestive of... the tortoises tripping (see the simplified version in <https://youtu.be/UqUQQORfCMY>). (38)e(i) makes use of the distinct but related notion of harmonic stability, which is relative to a key and thus to a musical context.³⁵ Such harmonic inferences do not have counterparts in normal auditory cognition, and semantic inferences from harmonic properties of the music may even be culture-dependent, specific to a given musical tradition.³⁶

³² The discussion in this paragraph (including some of the wording) follows Schlenker (to appear c).

³³ One particularly fruitful line of research has focused on properties of animal and human vocal signals that can be used in music to trigger powerful emotional effects (e.g. Aucouturier et al. 2016, Liuni et al. 2020).

³⁴ Rewriting the music so as to remove the dissonances removes the inference as well (see the simplified version in https://youtu.be/0A2egp_OIVU, rewriting by A. Bonetto).

³⁵ For instance, in the key of C major (corresponding to the scale obtained by pressing the white keys on a piano, starting with C), C is the most stable note, and the C major chord (= CEG) is the most stable chord.

³⁶ For instance, McAuley et al. 2021:2 observe that an atonal piece of music triggered very different inferences in a cross-cultural study; participants in Arkansas and Michigan associated the atonality with horror and murder, while participants in Dimen, China interpreted the atonality very differently, "involving happy games with friends". These findings do not address the question of whether the harmonic inferences in (37) are themselves culture-specific; happy games may well be conceptualized as involving physical instability.

If music indeed triggers inferences about the extra-musical world, can we define the semantic content of music? This was done in a simplified framework by positing that a series of n musical events is true of an object undergoing n events in the (real or imagined) extra-musical world just in case the inferences triggered by the music are all satisfied by the corresponding events. Here it matters that all the inferential effects in (38) are of the following form: *If musical events M_1 and M_2 stand in relation R , their respective denotations s_1 and s_2 stand in relation R^** . A semantics can for this reason be defined by requiring that certain relations among musical properties be preserved by the events depicted by the music. This makes it possible to define a notion of truth for musical sequences, partly similar to the notion of truth for pictorial sequences in (32) (here we use the term 'events' rather than 'situations' because it is more intuitive, but the structure of the account is similar):

(39) **Musical sequences true of n events (after Schlenker 2017, 2019a)**

A musical sequence $\langle M_1, \dots, M_n \rangle$ is true of an object undergoing events $\langle s_1, \dots, s_n \rangle$ relative to auditory point v just in case:

(1) temporally, $s_1 < \dots < s_n$;

(2) certain preservation conditions are satisfied, for instance ones corresponding to (38)b,e:

a. If M_i is less loud than M_k , then (i) s_i has less energy than s_k ; or (ii) s_i is further from the auditory point v than s_k is.

b. If M_i is less harmonically stable than M_k , then s_i is less (i) physically, or (ii) emotionally stable than s_k .

In this simple proof of concept, it can already be seen that the inferences triggered will be very general, hence the abstract character of musical meaning: an excerpt will typically be true of lots of very diverse tuples of events (namely those that satisfy the preservation conditions). It can be highlighted by revisiting Leonard Bernstein's celebrated discussion of Richard Strauss's *Don Quixote (Variation II)* (link: <https://youtu.be/XFZ7wORtj2A>). As an instance of 'program music', this piece has an explicitly descriptive character. Yet Bernstein sought to convince his audience that *even* such a piece does not describe any real or imagined situation: instead, he argues that the true meaning of music is "the way it makes you feel when you hear it". To make his point, Bernstein claimed to tell the *wrong* story to fit the music, and showed that it worked just as well as the intended one. The original story is about Don Quixote (i) departing on his horse to conquer the world, (ii) approaching a flock of sheep baaing, which he mistakes for an army, (iii) charging at them and creating chaos in the process, and finally (iv) feeling proud about his knightly deed. Bernstein's *wrong* story pertained to Superman (i') departing on his motorcycle to free an unjustly imprisoned friend of his, (ii') approaching the jail, in which the prisoners are snoring, (iii') charging into the prison and wreaking havoc in the process, and finally (iv') triumphantly hurling his friend back to freedom. It is striking that Bernstein's two stories are, with minor exceptions, isomorphic to each other: there is a close correspondence to the events described in (i)-(i'), (ii)-(ii'), (iii)-(iii') and (iv)-(iv').

This is not for lack of imagination: the details of the music trigger inferences that must be preserved by both stories – which is unsurprising if music has an iconic semantics of the type discussed in section 4 (Schlenker, to appear c). For instance, the triumphant character of the beginning, (i)-(i'), is in part due to the upwards melodic movement, as heard in (40); importantly, when the music is rewritten in accordance with composition rules so as to invert the melodic movement, (40b), this triumphant character disappears, as expected by (38)a(ii). The sheep baaing, (ii), and the prisoners snoring, (ii'), create a somewhat chaotic effect in part because of dissonances in the music, as heard in (41); again, when the dissonant chords are replaced with consonant ones, (41b), the chaotic effect largely disappears, as expected by (38)e(i). Similarly, the sheep and prisoners appear to be approaching the perspectival center, (iii)-(iii'), because of a crescendo (increasing loudness) in the music, as heard in (42); once again, when the crescendo is replaced with a diminuendo (decreasing

loudness), (42b), the effect disappears (if anything, the denoted objects appear to move away), as expected by (38)b(ii).

- (40) a. Original: upwards melodic movement at the beginning (simplified MIDI)
<https://youtu.be/dSwjTMyzSM>
b. Rewritten: downwards melodic movement at the beginning (A. Bonetto)
<https://youtu.be/5e-39sxEKhk>
- (41) a. Original: dissonant chords (simplified MIDI)
<https://youtu.be/fKgJDy0wYk0>
b. Rewritten: consonant chords (A. Bonetto)
<https://youtu.be/EnhSaeMORCk>
- (42) a. Original: crescendo (simplified Midi) <https://youtu.be/mMA9dByPAw>
b. Rewritten: diminuendo (A. Bonetto) <https://youtu.be/kCfay6s4Igs>

In sum, one's ability to tell the 'wrong' story to fit the music doesn't show that the latter has no semantics, just that it has a fairly abstract one, i.e. that its meaning singles out rather diverse situations.

One key question for future research pertains to the interface between music syntax and music semantics. It has been suggested (rather speculatively) that Lerdahl and Jackendoff's grouping structure arises from an attempt to recover the structure of the events denoted by the music (Schlenker 2017, 2019a, Zaradzki 2021). In essence, the idea was that events come with a natural part-of structure, and that when further conditions are imposed, a tree structure can be derived, but with some occasional exceptions that Lerdahl and Jackendoff noted.³⁷ Similarly, the headed nature of groups might be reanalyzed in semantic terms, with some events being construed as more central than others. Whether this semantic reanalysis of grouping structure and time-span reductions is correct remains to be seen; and the interface between further music structures and music semantics has yet to be investigated.

6.2.3 Adding variables

The simple cases in section 6.2.2 involved excerpts with a single musical voice corresponding to a single virtual source; in more complex cases, a listener may infer several virtual sources, which raises a question analogous to the one Abusch 2013 asked about visual narratives: what are the cross-reference relations among them? Here it may be fruitful to enrich musical representations with variables, just as was argued for visual narratives by Abusch.

This issue is best illustrated by an example. The beginning of Chopin's Mazurka Op. 33 No 2 has an extremely simple structure that involves two melodies, A followed by B, which is then repeated. We can write the form of the piece as AB A'B' where A' and B' are the repetitions of A and B; that is, A' = A and B' = B. In the original score, (43)b, the first [AB] sequence is played loudly (*forte/f*), whereas the repetition [A'B'] is played softly (*pianissimo/pp*). Crucially, this affects a listener's inferences on the nature of the virtual source, i.e. the object that the music evokes. If we manipulate the music so that it is given a flat realization with constant loudness, as in (43)a, it is compatible with multiple interpretations: listeners can infer (i) a single object that is evoked by the entire piece, or, alternatively, (ii) one object that is associated with [AB] and a second object that is associated with [A'B']; possibly, a reader may even infer (iii) one object that is associated

³⁷ A particularly interesting case arises in structures of the form $[_1 a \ ^1 b_1 c^1]$, where b belongs both to the group $[a b]$ (= lower brackets) and $[b c]$ (= upper brackets). Lerdahl and Jackendoff 1983 argue that such cases arise in grouping in visual perception as well, for instance in case an object is naturally construed as being a part of two larger, overlapping objects. This is reinterpreted in Schlenker 2017, 2019a as a case in which, in the denotations of the music, a subevent is naturally construed as being part of two temporally overlapping events (neither of which is a subpart of the other). See Migotti and Zaradzki 2019 and Zaradzki 2021 for more subtle conditions on the syntax/semantics interface in music.

with A and A', and another one that is associated with B and B'. This ambiguity of (43)a, in effect, is an ambiguity in cross-reference, just as was the case with Abusch's cube in (31), or for that matter with (4)a in section 2.2.2.

- (43) a. Flat realization $[AB] [A'B']$ https://youtu.be/IMZrU_vEA7A
 b. Chopin's dynamics $[AB]_f [A'B']_{pp}$ <https://youtu.be/tyJ0ZQxG8rU>
 c. Anti-Chopin dynamics $A_f B_{pp} A'_f B'_{pp}$ <https://youtu.be/dYTIAdl4-B0>

We can now manipulate the piece in two ways that appear to disambiguate. When we add the dynamics of Chopin's original score³⁸ in (43)b, (iii) appears to become far less likely. The reason is that [AB] is played *forte*, while [A'B'] is played *pianissimo*. This can be made sense of in an interpretation such as (ii), in which the first denoted object is energetic or close-by and corresponds to [AB], while the second is less energetic or further away and corresponds to [A'B']. Importantly, one can also create an 'anti-Chopin' dynamics, (43)c, in which A and A' are realized forte, while B and B' are realized pianissimo: this may then suggest reading (iii), where A and A' correspond to one source, while B and B' correspond to another, as can be heard in (43)c.

In several orchestrations written for a ballet by Michel Fokine, the identity of the denoted objects is made more salient by the use of different timbres, associated with different instruments: Chopin's dynamics $[AB]_f [A'B']_{pp}$ (*forte*, then *pianissimo*) is reflected by the use of different timbres for [AB] and [A'B'], typically using the orchestra for [AB] and individual wind arguments for [A'B']. An example is Britten's orchestration in (44) (see Schlenker, to appear c for further examples of orchestration of the same piece that make related choices).

- (44) [Benjamin Britten's orchestration](#) (1941)³⁹: $[AB]_{orchestra} [A'B']_{oboe+flute}$ <https://youtu.be/aHUeA3526DY>

Finally, the importance of cross-reference becomes even more salient when a ballet co-occurs with the music. In Fokine's piece, [AB] corresponds to a movement of the main ballerina, [A'B'] to that of the other dancers.

- (45) Fokine's *Les Sylphides* (originally called *Chopiniana*), movement on Chopin's Mazurka Op. 33 No 2 [Performance from 1984](#), American Ballet Theatre, on an orchestration close to Britten's version <https://youtu.be/Lsuc3KUKKpQ>
 $[AB]_{main\ ballerina} [A'B']_{other\ dancers}$

In sum, an abstract notion of iconic semantics (based on virtual sources) may be used to develop a semantics for music, and it too might benefit from the use of variables to enrich iconic representations, just as was argued by Abusch for visual narratives.

6.3 Dance

In section 6.1, we discussed the syntax and semantics of visual narratives. Dance can be analyzed, at the very least, as a (3D) visual animation, but one with several peculiarities that require a special treatment: (i) it is produced by dancers, with physical and genre-specific constraints on permissible movements; (ii) depending on the genre, it may have a particularly abstract semantics (e.g. it need not be obvious what a ballet represents); (iii) it often stands in a particularly close relation to music.

6.3.1 Dance syntax

If one asks whether dance may have syntax (and thus potentially a 'grammar'), a reasonable point of comparison is the syntax of visual narrative (section 6.1.1), and the syntax of music


³⁸ Chopin's score can be found at: <https://drive.google.com/file/d/15Vz7lKgRDKWawUTxW9ZKWsmoaDT912A8/view?usp=sharing>

³⁹ Retrieved from <https://www.youtube.com/watch?v=DQLde6QJXvM> on December 13, 2019 Credits: Chopin: Les Sylphides (Arr. for Orchestra By Benjamin Britten, Mono Version) Joseph Levine, American Ballet Theatre Orchestra 1 January 1954.

(section 6.2.1). Building on the latter, Charnavel (2019) proposes that dance is naturally perceived as hierarchically structured in ways that are reminiscent of musical grouping. Specifically, she argues with experimental means that some grouping preference rules proposed by Lerdahl and Jackendoff 1983 for music have analogues in the grammar of dance. A central insight may thus be that grouping is a core concept that unifies all of the human modes of expression discussed so far: visual narrative, music, dance – and language, as discussed in section 2.1.2.

Focusing on probing grouping preference rules in dance, Charnavel (2019) identifies six variables that result in the formation of a group boundary: *change of {body part, orientation, level, direction, speed and quality}*. Concretely, Fig. 1 illustrates three alternative ways of generating grouping structure based on the movement of a dancer (schematically seen from above in the shape of an arrow). Each p_i indicates a potential grouping boundary. The arrows indicate the path of the dancer, whereas the thickness of the arrows reflect the speed of the movement, for instance if the thicker arrow reflects a faster speed, the thinner arrow reflects a slower speed. Charnavel then considers three potential rules (labeled as A, B and C) according to which the movement could be sequenced. Rule A specifies that change of speed and change of direction each trigger a grouping boundary. By contrast, for Rule B, a single property, change of direction, triggers the appearance of a group boundary. For Rule C, a group boundary is created by change of speed and nothing else. To decide which rule is actually applied by onlookers, a partial answer can be found in an experiment that Charnavel (2019) carried out. In this experiment, participants had to split a video into two sequences, and each stimulus video contained two out of the six possible changes that could mark a grouping boundary. The results give some indication as to whether participants would choose Rule B or Rule C, though they do not bear on whether participants would consider Rule A. Charnavel’s findings were that some variables are weighted more than others in determining group boundaries: change of body part had a stronger effect than all other changes – orientation, speed, etc. For the variables in our illustration, there was no statistically significant preference for either change of direction (B) or change of speed (C); if such a preference is missing, this could be an indicator that participants would have preferred Rule A (which was not given as an option) or that they randomly chose between B and C.


(46) Figure 1 (cited from Charnavel 2019:5.)



Crucially, dance can also give rise to constituency and hierarchical grouping: when multiple variables that give rise to local grouping boundaries are applied to a given sequence, higher level grouping may be perceived by the onlooker, as illustrated by (47). Assuming that A is the grouping structure based on *both* change of speed and change of direction (Rule A in (46)), an onlooker may give more weight to one of these variables over the other, resulting in a higher-level group boundary as in A’; in this example, A’ would result from weighting change of direction more than change of speed. The larger rectangles in (47) denote higher level groups, contrastively the smaller rectangles mark the local grouping structure. This is

parallel to the notion of constituency present in language and music, which entails the usual properties of constituents, namely that they do not overlap with one another.

(47) Figure 2



Note that even in dance sequences that only contain one type of change, a higher level grouping structure may emerge due to manipulation of the *strength* of this change: for instance, if a dance sequence contains stronger and weaker changes of direction, all of them would give rise to a local grouping structure, out of which the strong ones (and only those) would trigger a higher level grouping structure; see Charnavel 2016 for details.

To conclude this discussion, we can ask if we find *headedness* in dance, related to what we find in visual narrative, where narrative Peaks can ‘project’ to a higher level (section 6.1.1), and in musical meter (illustrated in (36)) and time span reductions (illustrated in (37)), where a strong element projects to the next level (section 6.2.1). A main difference between dance and the other artforms is that we can detect *categories* in visual narrative (e.g. Peak vs. Initial) and music (e.g. a tonic I chord vs. a dominant V chord, or a strong vs. weak beat); this may be a precondition for groups being headed. As of now, it is an open question whether similar categories can be posited for dance as well.

6.3.2 Dance semantics

The semantics of non-figurative / non-representational dance (e.g. some varieties of classical ballet, but also purely expressive dance in a club or a ballroom) is as yet an open question. But some narrative dances designed to tell a story might be easier to approach with semantic means. Patel-Grosz et al (2018) probe for the presence of abstract meaning inferences in Bharatnatyam, a South Indian classical dance. Bharatnatyam is a type of narrative dance, thus closer in proximity to language and silent visual narratives than it is to other dance forms. In order to unearth how a dancer can encode meaning contrasts, Patel-Grosz et al investigate co-reference vs. disjoint reference in Bharatnatyam, through a series of motion capture production studies.

A sample stimulus is given below which consists of a context (= (48)) and target sentences (= (49)). (49)a and (49)b are identical except for the number of characters involved: in (49)a, the artist sees the same man twice, hence a relation of coreference between the first and the second sentence; this is replaced in (49)b with a relation of disjoint reference.

(48) *Context*: An artist has designed a statue for a temple. She is at the temple, watching how people interact with the statue; the room is full of people.

(49) *Item 1*

- a. The artist sees a strong man sitting on the ground.
Then she sees that *the same man* is holding a spear. (coreference)
- b. The artist sees a strong man sitting on the ground.
Then she sees that *another man* is holding a spear. (disjoint reference)

In a production experiment, the dancer performed a single, fluid movement in the coreferential condition, as shown in (50). In other words, co-reference between P11 and P14 is not explicitly marked through specific mechanisms. Note that Bharatanatyam clearly operates with an iconic semantics; the dance posture that conveys that a man is sitting, [P11], or holding a spear, [P14], resembles that which it describes.

(50) The artist sees a strong man [P₁₁ sitting on the ground]. Then she sees that *the same man* [P₁₄ is holding a spear].

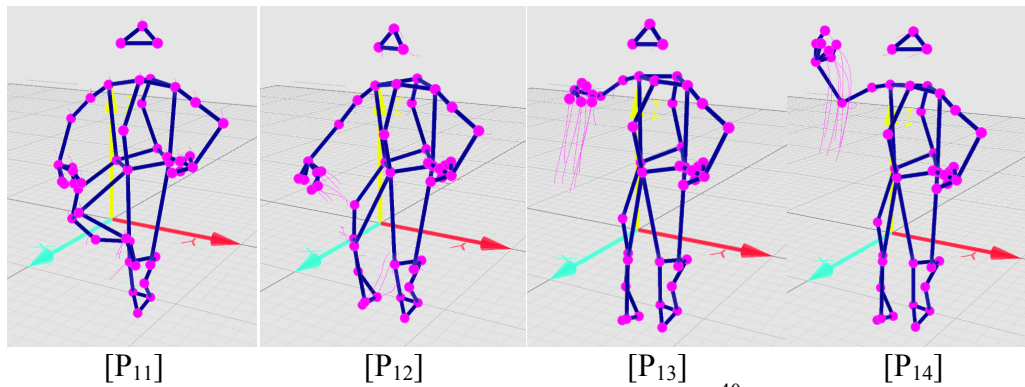


Figure 1: co-reference condition⁴⁰

By contrast with the coreferential case, disjoint reference is marked, as illustrated in (51). As expected, the dancer begins in a seated position, identical to the co-reference condition. The encoding of disjointness can be witnessed between [P21] and [P25].

(51) The artist sees a strong man [P₂₁ sitting on the ground]. Then she sees that [P_{22+P23+P24} another man] [P₂₅ is holding a spear].

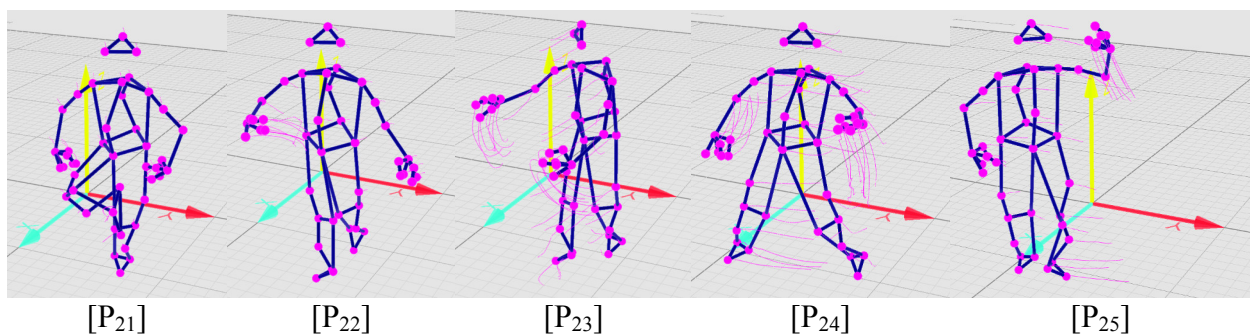


Figure 2: disjoint reference condition


The crucial observation, which Patel-Grosz et al (2020) replicate in two follow up studies, is that the dancer marks a new position in the visual space in [P23], which she then assumes in [P24], followed by the adoption of the spear-holding position in [P25]. The mechanism used by the dancer in [P23]-[P24], marking the space followed by moving into it, is reminiscent of the use of reference tracking via loci (cf. Lillo-Martin & Klima 1990, and see Schlenker 2017c for a recent survey article), and action role shift (where a signer shifts their body, head position and/or eye gaze in order to assume the perspective of another referent) in sign language (cf. Padden 1986, Lillo-Martin 1995, Quer 2005, Sandler & Lillo-Martin 2006, Herrmann & Steinbach 2009, 2012; see Davidson, 2015, for a recent discussion).

⁴⁰ Videos with stick figures, from which these screen shots are taken were generated by the Qualisys Track Manager software (QTM, Qualisys AB, Gothenburg, Sweden).

As discussed in sections 4.1 and 6.1.2, Greenberg (2011, 2013) and Abusch (2015) provide ways of defining truth in pictures and visual narrative, based on a generalized possible worlds model of information content (Abusch 2015:2), where informational entities are defined as entities (such as sentences or pictures) that rule out some possible situations and admit others. An illustration for a still from a dance sequence (e.g. [P₁₁], [P₂₁]) is given below. Truth in dance (as in other visual narrative) is defined in terms of how well a dance position [P_n] maps to a situation σ_n in the narrative; i.e. the dance position in (5) counts as satisfied by a fictional situation σ (i.e. “true” in σ) if a sitting activity is taking place in σ .

- (1) *satisfaction conditions for dance position that describes a sitting activity*

a situation σ satisfies



only if in σ a person is sitting.

A central (and yet unanswered) question concerns the exact nature of the mechanism that we witness in P₂₃ and P₂₄. There are at least two possible approaches (grouping-based co-reference vs. rudimentary *loci*), and the respective hypotheses are stated in (52).

- (52) *Hypothesis space for disjoint reference marking in dance* (Patel-Grosz et al. 2020)

- a. H₁ = grouping-induced disjointness (a type of visual iconicity) is sufficient to account for the change of position on part of the dancer (and no additional mechanisms, such as *loci*, are needed)
- b. H₂ = designated positions in space (*loci* of the type that we find in sign languages) are needed to account for the change of position on part of the dancer

(52)a assumes that due to the absence of change of direction/orientation in the co-reference condition, no grouping boundary arises. In the disjoint reference case, a grouping boundary is present due to the change of direction and orientation, which creates visual discontinuity. Unlike Charnavel (2019) who only assumes syntactic grouping (see section 6.3.1), Patel-Grosz et al (2018) propose that the syntactic groups are also interpreted as semantic groups, which correspond to events/sub-events. This is parallel to the grouping structure assumed in music semantics, as alluded to in section 6.2.2. Disjointness would then emerge by virtue of two referents contained in separate groups/events. It is worth noting that if (52)a is correct, the actual causal chain, the precise mechanism responsible for disjointness effects, is underspecified. A plausible approach could be that visual iconicity gives rise to syntactic grouping as a byproduct of perceived discontinuity; this syntactic grouping then maps onto a semantic grouping into two (or more) events, which triggers an inference of disjoint reference (non-identity) between characters in these events.

As an alternative, (52)b assumes that narrative dance uses a reference tracking mechanism that charts positions in space via rudimentary *loci*, reminiscent of the grammatical kind used in sign language (e.g., Lillo-Martin & Klima 1990, Liddell 1990). It remains to be seen whether visual iconicity engages in grouping structure to create disjointness, or whether *loci* are introduced to track multiple referents. In order to test for possible *loci* that could be responsible for disjoint reference, Patel-Grosz et al. (2020) develop the abovementioned production study and explore whether *loci* that occur earlier in

the narrative can be picked up again at a later stage, in a similar fashion to the way loci are used in sign language.

An example sequence from a follow-up study is given in (53). This example introduces two referents (an eating child and a child holding a spear). The first referent (the eating child) is in a different position (loc1, in P₃₂) from the second referent (loc2, in P₃₃). When the first referent is picked up later in connection with a directional predicate (*watch*), the artist assumes loc1 again and faces towards loc2. This was reproduced across items and in a follow-up study with three different referents.

(53) [P₃₁ The artist]_{loc0} sees [P₃₂ a child eating a mango]_{loc1} outside the temple.

Then she sees [P₃₃ another child holding a spear]_{loc2}.

[P₃₄]P₃₅ *The eating child*_{loc1} watches *the child with the spear*_{loc2}].

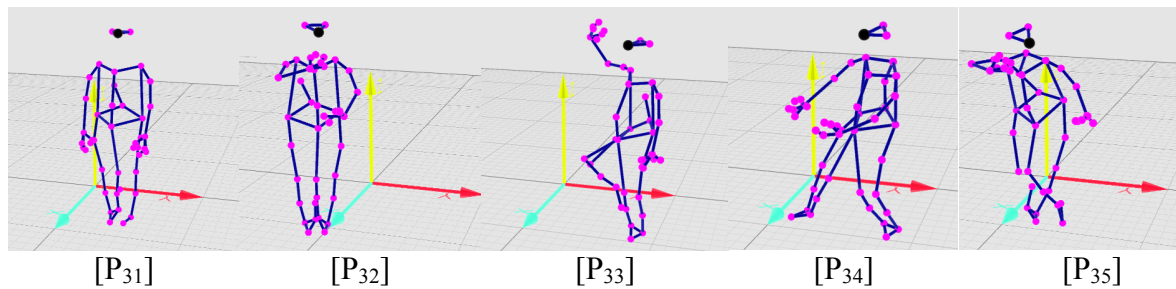


Figure 3: grouping in the disjoint reference condition

What remains unclear at this stage is whether these positions in space exhibit other properties of sign language loci; for instance, sign language loci do not obligatorily mark the actual (iconic) location of (fictitious or actual) physical objects or individuals in space, but can define arbitrary positions. This raises the question of whether audience members who pick up on loci identify them with the literal positions of entities (characters/referents) in the narrative, or just arbitrary positions that represent the entities.

Regardless of which analysis is correct, this exploration lends further support to something that we have seen for visual narrative in 6.1.2, and for music in 6.2.3. In all of these modalities, now also including dance, humans have the ability to track objects throughout a narrative, which can be modeled by virtue of variables. While grouping and constituency may be one property shared by these systems at the level of syntax, the tracking of denoted objects in a narrative may be one of their shared properties at the level of semantics.

6.4 Interaction between visual and musical narratives

So far, we have discussed visual narratives (section 6.1), music (section 6.2) and dance (section 6.3), where dance can also be treated as a type of visual narrative in some cases, such as narrative dance, and minimally as a type of visual animation in other cases. We established that all of these media involve grouping structure both at a lower level and at a higher level, and that there are semantic commonalities, such as a use of variables. We can then ask how they interact, and, specifically, how does music interact with visual narratives? When studying music co-occurring with dance, films or cartoons, it is tempting to take one medium to just 'repeat' what the other says (e.g. with dance 'interpreting' the content of the music). But this is unlikely to be the right approach: both media have their own semantics, and the challenge is to understand how they can get aggregated.

The simplest hypothesis is that these contents are conjoined. If we focus on the highly simplified case of n pictures aligned with n musical events, the similarity between the

semantic rules in (32) and (39) makes it possible to conjoin them (modulo some adjustments). We first need to take both kinds of sequences to be true of the same kinds of (n -tuples of) objects. So it will prove convenient to posit that pictures and musical events alike are true of eventualities, which may be states/situations or events. With this assumption, we can take a pictorial sequence $\langle P_1, \dots, P_n \rangle$ aligned with a musical sequence $\langle M_1, \dots, M_n \rangle$ as true of eventualities $\langle e_1, \dots, e_n \rangle$ just in case $\langle P_1, \dots, P_n \rangle$ is pictorially true of $\langle e_1, \dots, e_n \rangle$ and $\langle M_1, \dots, M_n \rangle$ is musically true of $\langle e_1, \dots, e_n \rangle$. On this view, then, each medium makes its own contribution. There might be respects in which the two media 'say the same thing', but this is by no means a necessity.

In view of the foregoing discussion, each medium will come with variables, also known as discourse referents, and thus anaphoric relations will serve to disambiguate cross-reference within and across media. An example might make this point concrete (Schlenker, to appear c). (54) displays a simplified version of the beginning of Disney's *Sorcerer's Apprentice*, treated for simplicity as a sequence of 4 pictures combined with 4 musical events (in fact, these are 4 visual scenes combined with 4 musical themes). We first see the sorcerer with a musical theme A, then the apprentice with a musical theme B, then the sorcerer with the genie he just conjured, co-occurring with a modification of A ($= A'$), then the genie alone, co-occurring with a modification of B ($= B'$). In a simple analysis, each musical theme comes with a discourse referent, hence the notations $A[v_1]$, $B[v_2]$, for instance.

(54) **Four pictures from Disney's *Sorcerer's Apprentice* (below), with four musical events (above)**
 (Disney, *Fantasia* 1940; figure from Schlenker, to appear c)

Video: <https://youtu.be/BR0Asfb2bxg>



The music alone might suggest that there is a total of two discourse referents, one associated with A and A', and one associated with B and B' (just because of the similarity between A and A', B and B'). But when the music is aligned with the cartoon, things become more interesting. A and A' are naturally associated with the sorcerer, hence it is natural to posit that the variable v_1 indeed appears in A and A', and also on the sorcerer in the first and third pictures. B is naturally understood to characterize the apprentice, hence a variable v_2 appearing in B and in the second picture. However, B' is not associated with the apprentice, as one might expect on the basis of the music, but rather with the genie, and thus despite the similarity between B and B', it makes some sense to posit that the cross-reference between the genie and B' is enforced by a third variable v_3 . In this case, then, variables play a non-trivial role in establishing and tracking coreference both within and across media.

An example of dance-music interaction was already discussed in section 6.2.3, examples (44) and (45). In Michel Fokine's ballet choreography for Chopin's Mazurka Op. 33 No 2, we saw that in an [AB] [A'B'] sequence, different timbres in the music (orchestra for [AB] vs. oboe+flute for [A'B']) are mapped onto a choice between movements of the main ballerina [AB] vs. the other dancers [A'B']. While this was a case of complete overlap, we can easily imagine a situation where the other dancers only dance A', whereas B' is performed by a new dancer on the stage. This would then replicate the pattern that we just

saw for the *Sorcerer's Apprentice*, in that the music would suggest two variables, whereas the dance sequence would introduce three. For the narrative dance sequences in 6.3, the question similarly arises whether an accompanying music would play a similarly supportive role in establishing and tracking variables.

While conjunction of visual and musical meaning might be a good starting point to analyze the interaction between visual and musical narratives, more sophisticated techniques from gesture semantics might prove useful as well. In a nutshell, it was proposed that, in some cases at least, music can enrich the meaning of visual sequences in the same way that gestures can enrich speech: by triggering cosuppositions.

The basic idea is illustrated in (55) by combining music with a simple gif representing Asterix drinking the magic potion, hitting a Roman soldier and leaving the room (Schlenker, to appear c).

(55) *Context*: Asterix had an earlier encounter with a Roman soldier. Now he is faced with him once again.⁴¹

What will happen next? Will Asterix...

a.  [WHISTLE](#)

<https://youtu.be/aOyr8yTS6uY>



b.  [WHISTLE](#)

<https://youtu.be/BGElapUY4nw>



A light-hearted whistling tune accompanies either the entire scene (= (55)a) or just Asterix's departure (= (55)b). The entire excerpt is embedded in a question so as to tease apart at issue content from presuppositions.

Unsurprisingly, the visual narrative provides information about the action and thus its main content is at issue. Strikingly, the music adds inferences to the scene, but different ones depending on where the music appears. When the whistling tune co-occurs with the entire scene, (55)a, the whole sequence of actions is presented as light-hearted. When the whistling scene only co-occurs with the section in which Asterix departs, (55)b, the inference is different, along the lines of: *if Asterix leaves the room after drinking the potion and hitting the Roman soldier, his departure will be light-hearted*. Crucially, this inference exists despite the presence of a question, and it is conditionalized relative to visual content it modifies. In these two respects, it behaves just like cosuppositions triggered by co-speech gestures. In the end, this behavior can be derived from a semantics for visual narratives along the lines of (32), combined with hypothesis that music semantics provides cosuppositional information about the part of the scene it co-occurs with (Schlenker, to appear c).

A final point worth mentioning is that the timing of the music may also interact with the hierarchical structure of the cartoon. In terms of the narrative categories discussed in section 6.1.1, the Asterix sequence in (55) can be mapped to Initial-Peak-Release, and the whistling either occurs during the Release, (55b), or throughout the entire Arc, (55a). In more complex narrative structures, e.g. with two Peaks, we can expect that the pairing of music with narrative categories is constrained in interesting ways, in that there may be Peak-specific types of music as opposed to, say, Release-specific types of music.

These examples raise two questions for future research. First, is it in general the case that, as in (55), when embedded in sentences, film and cartoon music makes the same kind of semantic contribution as co-speech gestures? Second, does this behavior only arise when film and cartoon excerpts are embedded in sentences, or also in normal (unembedded) films and cartoons? The latter possibility is tantalizing but remote – as was the case for the hypothesis

⁴¹ Music: Billy Mowbray, Uke and Whistle.

that visual content on its own gets productively divided among some established slots of the inferential typology.

7 Semantics without phonology

The systems we have reviewed so far all have an overt realization: human language, primate calls, music, and visual narratives are instantiated in utterances we can perceive and parse, music we can hear, drawings we can see. We close by showing that the same primitives can be found in systems without an overt realization.

Thought and reasoning allow us to manipulate representations, transforming models of what we know into models of conclusions we can draw from them, with crucial consequences for action. The tools of semantics can help illuminate this extraordinary human faculty. One starts with the description of a "language of thought" and a "logic of thought". That is, we specify what the elementary building blocks ("words") are and how they may be combined, and we can then derive new "sentences" from existing ones while preserving useful properties like truth. We outline concrete examples of insights that can be gained from this line of inquiry shortly. This semantics of thought doesn't just radically extend the program of linguistic semantics; it might contribute to an understanding of its cognitive roots.

Still, there is an important methodological difference between the semantics of language and of thought. In the latter case, the absence of an overt realization forces us to reverse-engineer these systems through more indirect evidence: behavioral consequences must be used to infer what the atomic elements (lexical items) are. Developing language-free diagnostics makes it possible to export this enterprise to many systems, such as the reasoning and concept-formation faculties of non-human animals. In this sense, the semantics of thought and reasoning has been pursued for a long time. Ongoing research argues that these results can be regimented with the formal tools of semantics.

7.1 Content concepts and logical concepts

Linguistic semantics starts with constraints on minimal lexical meanings: similar concepts are found across languages, while some concepts are never found. As we saw in Section 2, one such constraint is Gärdenfors's connectedness, which has the result that no word may mean *mushroom or table*. Similarly, no word in no language is known to mean *none or all, few or many*, and a logical version of connectedness can derive this constraint as well, and it generalizes the old insights that primitive words are either positive or negative.

Can similar constraints be displayed for language-free concepts? Consider the following thought experiment. You see sets of 5 objects each, and observe that groups with 2, 3 or 5 red objects behave in some way, while groups with 0 or 1 red object do not. It is natural to infer that groups with 4 red objects will behave like the former: connectedness constrains generalizations drawn on the fly. Similarly, a connectedness-compatible generalization involving groups with at least 2 red objects will be easier to infer than a connectedness-incompatible generalization involving groups with 0, 1 or 4 red objects. Using tasks of this sort, we can arguably investigate people's language of thought without relying on their languages (Piantadosi et al. 2016), with interesting conclusions – including that the language of thought might make use of bound variables (e.g. Overlan et al. 2017). Still, how can we ensure that subjects are not relying on an inner linguistic monologue – which would make the final analysis about words? Researchers have typically made participants work with concepts that are not expressed in their language: even in such cases, concepts that obey connectedness are inferred more rapidly than ones that don't.

Importantly, such language-free tasks can be extended to non-linguistic participants, such as pre-linguistic infants and non-human animals, with striking results. To cite but one,

baboons have been argued to obey a connectedness constraint just like human adults do: they infer connectedness-compatible generalizations faster than connectedness-incompatible ones (Chemla et al. 2019b). Thus a pervasive constraint in human content words and quantifiers arguably has roots in the conceptual behavior of rather distant non-human primates.

Going full circle, remember that in view of the complexity of field experiments, analyses of primate call meaning are greatly underspecified by available data. But constraints on animal concepts are plausibly constraints on call meanings as well, and they might help adjudicate among competing theories (e.g. the meanings that tend to replace general calls in Steinert-Threlkeld et al.'s results may not be connected, in which case the hypothesis of general calls might be saved in the end).

7.2 Alternatives and attention

There are further interesting restrictions on cognitive modules without a phonology like thought, reasoning, or vision. We focus here on the role of (i) alternatives, which have played a key role in recent semantics/pragmatics but might also hold the key to some puzzles about reasoning, and (ii) attention as a facilitator of tractability for hard problems, operating by selecting among these alternatives. The findings have ramifications for the typology of inferences discussed above.

First of all, recent years have seen an increased interest in applying linguistic methodologies to reasoning problems in an attempt to illuminate the line between interpretation and reasoning. It is a truism both in linguistics and the psychology of reasoning that it is extremely dangerous to diagnose *failures* of reasoning without a solid grasp on the precise interpretation of the premises at hand. Indeed, the quest for absolving interpretations of well-known reasoning tasks has been of interest to psychologists since at least the 80s.⁴² Semanticists have begun contributing substantively to this indispensable line of research, for example using semantic theories to explain framing effects (Geurts, 2013), or formal-pragmatics theories to explain rates of endorsement of apparent deductive fallacies (Picat, 2019).

Another key starting point is mental-model theory (Johnson-Laird 1983), one of the most influential theories of human reasoning. Focusing on deductive reasoning, mental-model theory proposes a representational system that distinguishes between categorical premises that correspond to *exactly one mental model*, and another class of premises that instead propose a set of *alternative mental models* for consideration. Alternative mental models are generated by premises involving natural language disjunction (e.g. English *or*), and by reasoning-internal processes that allow humans to flesh out the full space of possibilities given a set of premises, at a potentially considerable computational cost.

For present purposes, the crucial idea behind alternatives in mental-model theory is that they represent not the possible states of affairs that we *know* to be compatible with the premises, but those that we happen to *be attending to*. Imagine you had the following two pieces of information as premises in a deductive problem: P1: *Either John speaks English and Mary speaks French, or Bill speaks German*; P2: *John speaks English*. Can you conclude with certainty that *Mary speaks French*? In a study of structurally identical examples, Walsh and Johnson-Laird (2004) show that over 85% of participants draw this conclusion. Yet it is a fallacy, for the following possible state of affairs makes the premises true and the conclusion false: John speaks English, Mary *does not* speak French, Bill speaks German. On mental

⁴² The conjunction fallacy of Tversky and Kahneman (1983) has been a particularly fertile ground for these investigations by psychologists. For example, Dulany and Hilton (1991) apply Gricean pragmatics to articulate an alternative theory of the conjunction fallacy where ignorance inferences partly explain the phenomenon, and provide suggestive experimental evidence in favor of their account.

model theory, the fallacious conclusion is a result of the alternative possibilities we are attending to. We *know* that the counterexample just presented is possible and invalidates the fallacious conclusion, but we are not *attending* to that possibility. The only possibilities we consider by default are the two possibilities expressed by the disjunctive premise, and those possibilities suggest that the only way to make the second premise *John speaks English* true is by making *Mary speaks French* true as well.⁴³

Attention determines what alternative possibilities we consider when reasoning, and linguistic operators like disjunction drive attention. This is both a boon and a liability. On the one hand, it allows us to draw conclusions on the basis of a small space of alternative possibilities, reducing cognitive costs. On the other hand, it makes us vulnerable to fallacious reasoning.

What can semantics contribute to this research program? As mentioned, mental-model theory distinguishes between categorical premises and disjunctive premises that give rise to alternatives. Koralus and Mascarenhas (2013) show that this distinction is best regimented by incorporating semantic insights: semantic analyses of disjunction and several other operators have long made the same conceptual distinction as mental-model theory, and they thus offer off-the-shelf formal accounts of the representations of sentence meanings suitable to feed a mental-model-like reasoning module (e.g. the alternative semantics of Kratzer and Shimoyama, 2002; the inquisitive semantics of Groenendijk, 2008, and Mascarenhas, 2009; the truth maker semantics of Fine, 2012).

This discovery illustrates the fruitfulness of the interplay between semantics and reasoning. First, the formal landscape of theories of reasoning is enriched by this connection with linguistic semantics: semantic theories are characterized by their careful use of a broad palette of sophisticated formal systems, including non-standard logics developed in mathematics and computer science. The connection between semantics and reasoning allows us to import into reasoning brand new formal systems, offering new, useful, and insightful tools for theory building in the psychology of reasoning.

For example, Koralus and Mascarenhas's (2013) reimagining of mental-model theory offers a proof system that is associated with well-studied logical systems from the semantics literature. This proof system consists of a small number of operations that transform mental representations into other mental representations, much like say natural deduction systems provide rules for transforming formulas while preserving truth. Since Koralus and Mascarenhas (2013) were formalizing *intuitive* reasoning, their proposed operations on mental representations are not guaranteed to preserve truth in the general case. This proof-theoretic move allowed Koralus and Mascarenhas (2013) to prove meta-logical results with cognitive import. For example, they show that there is a derivation strategy (i.e. a reasoning strategy) in their system that guarantees logically valid conclusions, at the cost of exponential blowup of alternatives under consideration in the worst case. While the versions of mental-model theory due to Johnson-Laird and collaborators have typically come accompanied by open-source computer implementations, full-fledged soundness proofs of the kind just mentioned did not exist. The computer implementations of traditional mental-model theory did not constitute logics in the traditional sense, making soundness and completeness

⁴³ Reasoning problems like the *illusory inferences from disjunction* just exemplified have resisted accounts in frameworks other than mental models. For example, Sablé-Meyer and Mascarenhas (2021) prove that more recent probabilistic approaches to reasoning in the New Paradigm line (Oaksford and Chater, 2007) will only predict the observed fallacious conclusion to the extent that they predict an *unobserved* conclusion for the same deductive problem. Sablé-Meyer and Mascarenhas (2021) offer arguments in favor of unifying probabilistic approaches with the interpretive and model-theoretic commitments of work in the mental models tradition, but this direction of research is still in its earliest days.

theorems hard to formulate and prove. Formal methodologies from semantics opened the door to proper meta-logical results in mental-model theory broadly conceived.

Second, the empirical scope of existing theories of reasoning increases in interesting ways. For example, the effects of attention and alternatives in the representations of disjunctive sentences discovered by mental-model theorists have been replicated in empirical domains where semantic theories have posited disjunction-like alternatives: indefinites like *some* (Mascarenhas and Koralus, 2017) and epistemic operators like *might* (Mascarenhas and Picat, 2019; Johnson-Laird and Ragni, 2019). As the broader extent of the empirical phenomenon reveals itself, new and better constraints on our theories of reasoning emerge.

Finally, these results highlight the importance of cross-fertilization between theories of language and of thought, with the tantalizing possibility of broad unifications, as we have sketched here for the role of alternatives. Moreover, further cognitive modules can be brought into the fold in new ways. For instance, if effects of interpretation and reasoning with alternatives are at a deeper level about attention manipulation, then the need emerges to apply semantic methods to our study of visual attention as well.

It is interesting in this connection to return to the notions of implicatures, presuppositions, and supplements that we introduced before. One dimension of variation here concerns how attention is directed towards different parts of the world, or of a message in this case. We have argued that the division between these notions naturally arises beyond words (Section 5). In some cases at least, their pervasiveness and productivity in and beyond language might be due to how cognition is structured: the semantics of thought might yield the key to foundational issues in linguistic semantics.

8 Conclusion

This introduction outlined initial results of ‘Super Linguistics’, the application of tools and techniques from formal linguistics (and inspired by formal linguistics) to objects of study outside of natural (human) language. After reviewing the achievements of formal linguistics in Section 2 (with a focus on syntax, semantics and pragmatics), Section 3 demonstrated its applications to animal communication, where syntax, semantics and pragmatics can each provide important analytical insights. This discussion was followed by two Sections on iconic semantics (Sections 4+5), with a particular focus on speech with gestures, and sign language semantics as an important point of reference. Section 6 focused on three human art forms (visual narratives, musical narratives, dance, and interactions between them), demonstrating the insights that we can gain from a syntactic and semantic analysis; most importantly, what all three art forms seem to share is hierarchical grouping in syntax and variables as a referent-tracking device in semantics. Finally, Section 7 showed how Super Linguistics can contribute to theories of language-free concepts and reasoning, by applying the tools from linguistic semantics to the semantics of thought. The benefits of applying linguistic methodology to non-standard objects are reciprocal: linguistics can contribute to objects of study beyond natural language by revealing structure and meaning principles in areas as diverse as bird song, monkey alarm calls, and comics; at the same time, the study of these objects provides insights into fundamental properties of cognition in humans and beyond, which contribute to our understanding of human language.

References

- Abner, N., Cooperrider, K., & Goldwin-Meadow, S. (2015). Gesture for Linguists: A Handy Primer. *Language and Linguistics Compass*, 9(11), 437–449.
<https://doi.org/10.1111/lnc3.12168>
- Abrusán, M. (2011). Predicting the presuppositions of soft triggers. *Linguistics and Philosophy*, 34(6), 491–535. <https://doi.org/10.1007/s10988-012-9108-y>
- Abusch, D. (2013). Applying discourse semantics and pragmatics to co-reference in picture sequences. In *Proceedings of Sinn und Bedeutung* 17, 19–25.
- Abusch, D. (2015). Possible Worlds Semantics for Pictures. Manuscript, Cornell University.
- Abusch, D., & Rooth, M. (2017). The formal semantics of free perception in pictorial narratives. In Alexandre Cremers, Thom van Gessel & Floris Roelofsen (eds), *Proceedings of the 21st Amsterdam Colloquium*.
from <https://semanticsarchive.net/Archive/jZiM2FhZ/AC2017-Proceedings.pdf>.
- Aucouturier, J. J., Johansson, P., Hall, L., Segnini, R., Mercadié, L., & Watanabe, K. (2016). Covert digital manipulation of vocal emotion alter speakers' emotional states in a congruent direction. *Proceedings of the National Academy of Sciences*, 113(4): 948-53.
<https://doi.org/10.1073/pnas.1506552113>
- Barker, C., & Jacobson, P. I. (2007). *Direct Compositionality*. Oxford University Press.
- Barwise, J., & Cooper, R. (1981). Generalized quantifiers and natural language. *Linguistics and Philosophy* 4, 159-219.
- Bateman, J.A., & Wildfeuer, J. (2014). Defining units of analysis for the systematic analysis of comics: a discourse-based approach. *Studies in Comics* 5, 373-403.
https://doi.org/10.1386/stic.5.2.373_1
- Beaver, D. I., Roberts, C., Simons, M., & Tonhauser, J. (2017). Questions Under Discussion: Where Information Structure Meets Projective Content. *The Annual Review of Linguistics* 3, 265–284. <https://doi.org/10.1146/annurev-linguistics-011516-033952>
- Beaver, D. I., Geurts, B. & Denlinger, K. (2021). Presupposition. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2021 Edition). URL = <https://plato.stanford.edu/archives/spr2021/entries/presupposition/>
- Berthet, M., Mesbahi, G., Pajot, A., Cäsar, C., Neumann, C., & Zuberbühler, K. (2019a). Titi monkeys combine alarm calls to create probabilistic meaning. *Science Advances*, 5(5), eaav3991. <https://doi.org/10.1126/sciadv.aav3991>
- Berthet, M., Benjumea, J., Millet, J., Cäsar, C., Zuberbühler, K. & Dunbar, E. (2019b). Animal linguistics and the puzzle of titi monkeys alarm sequences. *Proceedings of the Amsterdam Colloquium 2019*. <http://events.illc.uva.nl/AC/AC2019/Proceedings/>
- Berwick, R.C., Okanoya, K., Beckers, G.J.L., & Bolhuis, J.J. (2011). Songs to syntax: the linguistics of birdsong. *Trends in Cognitive Sciences*, 15(3), 113-121.
<https://doi.org/10.1016/j.tics.2011.01.002>
- Byrne, R.W., Cartmill, E., Genty, E., Graham, K.E., Hobaiter, C., & Tanner, J. (2017). Great ape gestures: Intentional communication with a rich set of innate signals. *Animal Cognition* 20(4), 755–769. <https://doi.org/10.1007/s10071-017-1096-4>
- Cäsar, C., Zuberbühler, K., Young, R.J., & Byrne, R. (2013). Titi monkey call sequences vary with predator location and type. *Biology Letters* 9, 20130535. doi:10.1098/rsbl.2013.0535
- Charnavel, I. (2016). Steps Towards a Theory of Dance Cognition. Manuscript, Harvard University.
- Charnavel, I. (2019). Steps towards a Universal Grammar of Dance: Local Grouping Structure in Basic Human Movement Perception. *Frontiers in Psychology*, 10, 1364.
<https://doi.org/10.3389/fpsyg.2019.01364>

- Chemla, E., Buccola, B., & Dautriche, I. (2019a). Connecting content and logical words. *Journal of Semantics*, 36(3), 531-547. <https://doi.org/10.1093/jos/ffz001>
- Chemla, E., Dautriche, I., Buccola, B., & Fagot, J. (2019b). Constraints on the lexicons of human languages have cognitive roots present in baboons (*Papio papio*). *Proceedings of the National Academy of Sciences*, 116(30): 14926-14930. <https://doi.org/10.1073/pnas.1907023116>
- Chierchia, G., Crain, S., Guasti, M.T., Gualmini, A., & Meroni, L. (2001). The acquisition of disjunction: Evidence for a grammatical view of scalar implicatures. In *Proceedings of the 25th Boston University conference on language development*, 157-168.
- Chomsky, N. (1957). *Syntactic Structures*. The Hague: Mouton.
- Chomsky, N. (1980). *Rules and Representations*. New York, NY: Columbia University Press.
- Cohn, N. (2012). Structure, Meaning, and Constituency in Visual Narrative Comprehension. Doctoral Dissertation, Tufts University, Medford, MA.
- Cohn, N. (2013). *The visual language of comics: Introduction to the structure and cognition of sequential images*. London: Bloomsbury.
- Cohn, N. (2015). How to analyze visual narratives: A tutorial in Visual Narrative Grammar. Retrieved from: http://www.visuallanguagelab.com/P/VNG_Tutorial.pdf
- Cohn, N. (2018). In defense of a “grammar” in the visual language of comics. *Journal of Pragmatics*, 127, 1-19. <https://doi.org/10.1016/j.pragma.2018.01.002>
- Cohn, N. (2020). Your Brain on Comics: A Cognitive Model of Visual Narrative Comprehension. *Topics in Cognitive Science*, 12, 352-386. <https://doi.org/10.1111/tops.12421>
- Cohn, N., Paczynski, M., Jackendoff, R., Holcomb, P.J., & Kuperberg, G.R. (2012). (Pea)nuts and bolts of visual narrative: structure and meaning in sequential image comprehension. *Cognitive psychology*, 65(1), 1-38. <https://doi.org/10.1016/j.cogpsych.2012.01.003>
- Coppola, M., Spaepen, E., & Goldin-Meadow, S. (2013). Communicating about quantity without a language model: Number devices in homesign grammar. *Cognitive Psychology*, 67(1-2), 1–25. <https://doi.org/10.1016/j.cogpsych.2013.05.003>
- Cumming, S., Greenberg, G., & Kelly, R. (2017). Conventions of Viewpoint Coherence in Film. *Philosophers' Imprint* 17, 1.
- Davidson, D. (1967). Truth and meaning. *Synthese*, 17 (3): 304–323. <https://doi.org/10.1007/BF00485035>
- Davidson, K. (2015). Quotation, demonstration, and iconicity. *Linguistics and Philosophy* 38, 477-520. <https://doi.org/10.1007/s10988-015-9180-1>
- Dingemans, M. (2013). Advances in the cross-linguistic study of ideophones. *Language and Linguistics Compass*, 6(10), 654-672. <https://doi.org/10.1002/lnc3.361>
- Douglas, P.H., & Moscovice, L.R. (2015). Pointing and pantomime in wild apes? Female bonobos use referential and iconic gestures to request genito-genital rubbing. *Scientific Reports*, 5(1), 1-9. <https://doi.org/10.1038/srep13999>
- Dulany, D. & Hilton, D. (1991). Conversational implicature, conscious representation, and the conjunction fallacy. *Social Cognition* 9:1. <https://doi.org/10.1521/soco.1991.9.1.85>
- Ebert, C., & Ebert, C. (2014). Gestures, Demonstratives, and the Attributive/Referential Distinction. *Handout of a talk given at Semantics and Philosophy in Europe (SPE 7)*, Berlin, 28.
- Esipova, M. (2019). Composition and Projection in Speech and Gesture. PhD dissertation, New York University.
- Ferrigno, S., Cheyette, S.J., Piantadosi, S.T., & Cantlon, J.F. (2020). Recursive sequence generation in monkeys, children, U.S. adults, and native Amazonians. *Science Advances* 6, eaaz1002. <https://doi.org/10.1126/sciadv.aaz1002>

- Fine, K. (2012). A difficulty for the possible world analysis of counterfactuals. *Synthese*, 189(1), 29-57. <https://doi.org/10.1007/s11229-012-0094-y>
- Fitch, W.T. (2018). What animals can teach us about human language: the phonological continuity hypothesis. *Current opinion in behavioral sciences*, 21, 68-75. <https://doi.org/10.1016/j.cobeha.2018.01.014>
- Fitch, W. T., & Hauser, M. D. (2004). Computational constraints on syntactic processing in a nonhuman primate, *Science* 303 (5656), 377–380. <https://doi.org/10.1126/science.1089401>
- Gärdenfors, P. (2004). *Conceptual Spaces: The Geometry of Thought*. MIT Press. Cambridge, MA.
- Gärdenfors, P. (2014). *The Geometry of Meaning: Semantics Based on Conceptual Spaces*. MIT Press. Cambridge, MA.
- Gentner, T. Q., Fenn, K.M., Margoliash, D., & Nusbaum, H.C. (2006). Recursive syntactic pattern learning by songbirds. *Nature*, 440, 1204-1207. <https://doi.org/10.1038/nature04675>
- Genty, E., & Zuberbühler, K. (2014). Spatial Reference in a Bonobo Gesture. *Current Biology*, 24(14), 1601-1605. <http://doi.org/10.1016/j.cub.2014.05.065>
- Geurts, B. (2013). Alternatives in framing and decision making. *Mind and Language* 28:1. <https://doi.org/10.1111/mila.12005>
- Goldin-Meadow, S. (2003). *The Resilience of Language*. Taylor & Francis.
- Goldin-Meadow, S. & Brentari, D. (2017). Gesture, sign and language: The coming of age of sign language and gesture studies. *Behavioral and Brain Sciences*, 40. <https://doi.org/10.1017/S0140525X15001247>
- Granroth-Wilding, M., & Steedman, M. (2014). A robust parser-interpreter for jazz chord sequences. *Journal of New Music Research*, 43(4), 355-374. <https://doi.org/10.1080/09298215.2014.910532>
- Greenberg, G. (2013). Beyond Resemblance. *Philosophical review*, 122(2), 215-287. <https://doi.org/10.1215/00318108-1963716>
- Greenberg, G. (2019). Semantics of Pictorial Space. Manuscript, UCLA.
- Greenberg, G. (2021). The Iconic-Symbolic Spectrum. Manuscript, UCLA.
- Grice, H. P. (1957). Meaning. *The philosophical review*, 66(3), 377-388. <https://doi.org/10.2307/2182440>
- Grice, H. P. (1975). Logic and conversation. In Cole, P., and Morgan, J. L. (eds.), *Syntax and Semantics*, Vol. III: Speech Acts. New York: Academic Press, pp. 41– 58.
- Groenendijk, J. (2008). Inquisitive Semantics: Two possibilities for disjunction. In *Proceedings of the Seventh International Tbilisi Symposium on Language, Logic and Computation*.
- Gutzmann, D., Matthewson, L., Meier, C., Rullmann, H., & Zimmermann, T. E. (2021). *The Wiley Blackwell Companion to Semantics*. Oxford: Wiley.
- Heim, I. (1982). The Semantics of Definite and Indefinite Noun Phrases. Ph. D. thesis, University of Massachusetts.
- Heim, I. (1983). On the projection problem for presuppositions. In Barlow, M. and Flickinger, D. and Westcoat, M. (eds.), *Second Annual West Coast Conference on Formal Linguistics* (pp. 114-126). Stanford University.
- Heinz, J. & Idsardi, W. (2011). Sentence and Word Complexity. *Science*, 333(6040), 295-297. <https://doi.org/10.1126/science.1210358>
- Herrmann, A., & Steinbach, M. (2009). Visible context shifts: Quotation in sign languages. In *Talk presented at the SLE 42nd Annual Meeting of the Universidade de Lisboa, Faculdade de Letras, Portugal: Global Languages Local Languages*. Lisboa.

- Herrmann, A., & Steinbach, M. (2012). Quotation in sign languages. In Ingrid van Alphan and Isabelle Buchstaller (eds.), *Quotatives: Cross-linguistic and cross-disciplinary perspectives* (pp. 203-228). Amsterdam: John Benjamins.
- Hobaiter C, & Byrne, R.W. (2017). What is a gesture? A meaning-based approach to defining gestural repertoires. *Neuroscience & Biobehavioral Reviews*, 82, 3-12. <https://doi.org/10.1016/j.neubiorev.2017.03.008>
- Huron, D. (2016). *Voice Leading: The Science Behind a Musical Art*. MIT Press.
- Jäger, G., & Rogers, J. (2012). Formal language theory: refining the Chomsky hierarchy. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1598), 1956-1970. <https://doi.org/10.1098/rstb.2012.0077>
- Johnson-Laird, P. N. (1983). *Mental models: towards a cognitive science of language, inference, and consciousness*. Cambridge: Cambridge University Press.
- Johnson-Laird, P. N. & Ragni, M. (2019). Possibilities as the foundation of reasoning. *Cognition*, 193, 103950. <https://doi.org/10.1016/j.cognition.2019.04.019>
- Kamp, H. (1981). A theory of truth and semantic representation. In J. A. G. Groenendijk, T. M. V. Janssen and M. B. J. Stokhof (eds.): *Formal Methods in the Study of Language*, (pp. 277-322). Amsterdam: Mathematical Centre.
- Keenan, E. L. & Westerståhl, D. (2011). Generalized Quantifiers in Linguistics and Logic. In van Benthem and ter Meulen (eds), *Handbook of Logic and Language*, (pp. 859–910). Amsterdam, Elsevier.
- Kershenbaum, A., Bowles, A. E., Freeberg, T. M., Jin, D. Z., Lameira, A. R., & Bohn, K. (2014). Animal vocal sequences: not the Markov chains we thought they were. *Proceedings of the Royal Society B, Biological Sciences* 281, 20141370. <https://doi.org/10.1098/rspb.2014.1370>
- Kersken, V., Gómez, J.C., Liszkowski, U., Soldati, A., & Hobaiter, C. (2018). A gestural repertoire of 1-to 2-year-old human children: in search of the ape gestures. *Animal cognition*, 22(4), 577-595. <https://doi.org/10.1007/s10071-018-1213-z>
- Klima, E. S., & Bellugi U., eds. 1979. *The signs of language*. Cambridge, MA: Harvard Univ. Press.
- Koralus, P., & Mascarenhas, S. (2013). The erotetic theory of reasoning: bridges between formal semantics and the psychology of deductive inference. *Philosophical Perspectives*, 27, 312–365.
- Kratzer, A. & Shimoyama, J. (2002). Indeterminate pronouns: the view from Japanese. In *Third Tokyo Conference on Psycholinguistics*.
- Kuhn, J. (2016). ASL loci: Variables or features?. *Journal of Semantics*, 33(3), 449-491. <https://doi.org/10.1093/jos/ffv005>
- Lerdahl, F. (2001). *Tonal Pitch Space*. Oxford University Press.
- Lerdahl, F. & Jackendoff, R. (1983). *A generative theory of tonal music*. Cambridge, MA: MIT Press.
- Liddell, S. (1990). Four functions of a locus: Re-examining the structure of space in ASL. In C. Lucas (ed.), *Sign language research: Theoretical issues* (pp.176-198). Washington, DC : Gallaudet University Press.
- Liddell, S. K. (2003). *Grammar, Gesture, and Meaning in American Sign Language*. Cambridge: Cambridge University Press.
- Lillo-Martin, D. (1995). The point of view predicate in American Sign Language. In K. Emmorey and J. Reilly (eds.), *Language, gesture, and space* (pp.155–170). Hillsdale.
- Lillo-Martin, D., & Klima, E. S. (1990). Pointing out differences: ASL pronouns in syntactic theory. *Theoretical issues in sign language research*, 1, 191-210.
- Lillo-Martin, D., & Meier, R. P. (2011). On the linguistic status of ‘agreement’ in sign languages. *Theoretical Linguistics* 37(3-4), 95-141. <https://doi.org/10.1515/thli.2011.009>

- Maestripietri, D. (1997). Gestural communication in macaques: Usage and meaning of nonvocal signals. *Evolution of communication*, 1(2), 193-222.
<https://doi.org/10.1075/eoc.1.2.03mae>
- Liuni, M., Ponsot E., Bryant, G. A., & Aucouturier, J. J. (2020). Sound context modulates perceived vocal emotion. *Behavioural Processes*, 172, 104042.
<https://doi.org/10.1016/j.beproc.2020.104042>
- Maienborn, C., von Heusinger, K., & Portner, P. (2011). *Semantics: An International Handbook of Natural Language Meaning*. De Gruyter.
- Maier, E. (2019). Picturing words: the semantics of speech balloons. In: Schlöder, J. et al. (eds.): *Proceedings of the 22nd Amsterdam Colloquium* (pp. 584-592). Amsterdam: ILLC.
- Maier, E. & Bimpikou, S. (2019). Shifting perspectives in pictorial narratives. In: M. Teresa Espinal et al. (eds.): *Proceedings of Sinn und Bedeutung 23* (pp. 91–105). Barcelona: Universitat Autònoma de Barcelona.
- Mascarenhas, S. (2009). *Inquisitive Semantics and Logic*. Master's thesis, ILLC.
- Mascarenhas, S., & Picat, L. (2019). Might as a generator of alternatives: the view from reasoning. In *Proceedings of SALT 29, UCLA*.
- McAuley, J. D., Wong, P. C., Mamidipaka, A., Phillips, N., & Margulis, E. H. (2021). Do you hear what I hear? Perceived narrative constitutes a semantic dimension for music. *Cognition*, 212, 104712. <https://doi.org/10.1016/j.cognition.2021.104712>
- Meyer, L.B. (1956). *Emotion and Meaning in Music*. University of Chicago Press, Chicago.
- Migotti, L., & Guerrini, J. (2019). The music/language interaction: a semantic account. Talk given within the LINGUAE seminar, ENS, Paris, April 25, 2019.
- Migotti, L., & Zaradzki, L. (2019). Walk-denoting music: Refining music semantics. In *Proceedings of the Amsterdam Colloquium 2019*.
- Montague, R. (1970a). English as a Formal Language. In Visentini, B. et al. (eds.), *Linguaggi nella società e nella tecnica* (pp. 189-224). Milan: Edizioni di Comunità.
- Montague, R. (1970b). Universal grammar. *Theoria* 36: 373-398.
- Montague, R. (1973). The proper treatment of quantification in ordinary English. In K. J. J. Hintikka, J. M. E. Moravcsik, and P. Suppes (eds.), *Approaches to Natural Language* (pp. 221–242). Dordrecht: Reidel.
- Morris, C. W. (1938). Foundations of the Theory of Signs. In Otto Neurath (ed.), *International Encyclopedia of Unified Science*, vol. 1 no. 2. Chicago: University of Chicago Press.
- Oaksford, M. & Chater, N. (2007). *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford University Press.
- Overlan, M. C., Jacobs, R. A., & Piantadosi, S. T. (2017). Learning abstract visual concepts via probabilistic program induction in a Language of Thought. *Cognition* 168, 320–334.
<https://doi.org/10.1016/j.cognition.2017.07.005>
- Özçalışkan, Ş., Lucero, C., & Goldin-Meadow, S. (2016). Does language shape silent gesture?. *Cognition* 148, 10–18. <https://doi.org/10.1016/j.cognition.2015.12.001>
- Padden, C. (1986). Verbs and role-shifting in American Sign Language. In *Proceedings of the fourth national symposium on sign language research and teaching* (Vol. 44, p. 57). National Association of the Deaf Silver Spring, MD.
- Pasternak, R. (2019). The projection of co-speech sound effects. Manuscript, ZAS Berlin.
- Pasternak, R. & Tieu, L. (2020). Co-linguistic content projection: From gestures to sound effects and emoji. Manuscript, ZAS Berlin and U. of Western Sydney.
- Patel-Grosz, P., Grosz, P. G., Kelkar, T. & Jensenius, A. R. (2018). Coreference and disjoint reference in the semantics of narrative dance, In Uli Sauerland & Stephanie Solt (ed.), *Proceedings of Sinn und Bedeutung 22*, (vol. 2, pp.199-216). ZASPiL 61. Leibniz-Zentrum Allgemeine Sprachwissenschaft (ZAS).

- Patel-Grosz, P., Grosz, P. G., Kelkar, T. & Jensenius, A. R. (2020). Steps towards a formal semantics of dance. Manuscript, University of Oslo.
- Pesetsky, D. & Katz, J. (2009). The Identity Thesis for Music and Language. Manuscript, MIT.
- Pfau, R., Steinbach, M., & Woll, B. (eds). 2012. *Sign language: An international handbook*. Boston: Walter de Gruyter.
- Pfau, R., Salzmann, M., & Steinbach, M. (2018). The syntax of sign language agreement: Common ingredients, but unusual recipe. *Glossa: a journal of general linguistics*, 3(1): 107. 1-46. <https://doi.org/10.5334/gjgl.511>
- Piantadosi, S. T.; Tenenbaum, J.; Goodman, N. (2016). The logical primitives of thought: Empirical foundations for compositional cognitive models. *Psychological review*, 123(4), 392-424.
- Picat, L. (2019). *Inferences with Disjunction, Interpretation or Reasoning?* MA thesis, Ecole Normale Supérieure.
- Potts, C. (2005). *The Logic of Conventional Implicatures*. Oxford Studies in Theoretical Linguistics, Oxford: Oxford University Press.
- Potts, C. (2015). Presupposition and implicature. In Shalom Lappin & Chris Fox (eds.), *The handbook of contemporary semantic theory*, 2, (pp. 168-202). Oxford : Wiley.
- Pylyshyn, Z. (2003). *Seeing and Visualizing: It's Not What You Think*. MIT Press.
- Quer, J. (2005). Context shift and indexical variables in sign languages. In *SALT XV*. Ithaca, NY: CLC Publications.
- Quer, J., Pfau, R., & Herrmann, A. (2021). *The Routledge Handbook of Theoretical and Experimental Sign Language Research*. Routledge.
- Rohrmeier, M. (2011). Towards a generative syntax of tonal harmony. *Journal of Mathematics and Music*, 5(1), 35-53. <https://doi.org/10.1080/17459737.2011.573676>
- Rohrmeier, M., Zuidema, W., Wiggins, G.A., & Scharff, C. (2015). Principles of structure building in music, language and animal song. *Philosophical transactions of the Royal Society B: Biological sciences*, 370(1664), 20140097. <http://dx.doi.org/10.1098/rstb.2014.0097>
- Sablé-Meyer, M. & Mascarenhas, S. (2021). Indirect illusory inferences from disjunction: a new bridge between deductive inference and representativeness. Forthcoming in *Review of Philosophy and Psychology*.
- Sandler, W. & Lillo-Martin, D. (2006). *Sign Language and Linguistic Universals*. Cambridge University Press.
- Schembri, A., Cormier, K., & Fenlon, J. (2018). Indicating verbs as typologically unique constructions: Reconsidering verb ‘agreement’ in sign languages. *Glossa: a journal of general linguistics*, 3(1), 1-40. <https://doi.org/10.5334/gjgl.468>
- Schlenker, P. (2010). Supplements within a Unidimensional Semantics I: Scope. In Maria Aloni, Harald Bastiaanse, Tiki de Jager, and Katrin Schulz (eds.), *Logic, language and meaning* (pp. 74-83). Springer, Berlin, Heidelberg.
- Schlenker, P. (2011). Donkey Anaphora: the View from Sign Language (ASL and LSF). *Linguistics and Philosophy*, 34(4), 341-395. <https://doi.org/10.1007/s10988-011-9098-1>
- Schlenker, P. (2014). Iconic Features. *Natural Language Semantics* 22(4), 299-356. <https://doi.org/10.1007/s11050-014-9106-4>
- Schlenker, P. (2017). Outline of Music Semantics. *Music Perception: An Interdisciplinary Journal* 35(1), 3-37. <https://doi.org/10.1525/mp.2017.35.1.3>
- Schlenker, P. (2018a). Gesture Projection and Cosuppositions. *Linguistics & Philosophy* 41(3), 295–365. <https://doi.org/10.1007/s10988-017-9225-8>
- Schlenker, P. (2018b). Iconic Pragmatics. *Natural Language & Linguistic Theory* 36(3), 877–936. <https://doi.org/10.1007/s11049-017-9392-x>

- Schlenker, P. (2018c). Visible Meaning: Sign Language and the Foundations of Semantics. *Theoretical Linguistics*, 44(3-4), 123-208.
- Schlenker, P. (2019a). Prolegomena to Music Semantics. *Review of Philosophy & Psychology* 10(1): 35-111. <https://doi.org/10.1007/s13164-018-0384-5>
- Schlenker, P. (2019b). What is Super Semantics?. *Philosophical Perspectives* 32(1), 365-453. <https://doi.org/10.1111/phpe.12122>
- Schlenker, P. (2019c). Gestural Semantics: Replicating the typology of linguistic inferences with pro- and post-speech gestures. *Natural Language & Linguistic Theory* 37, 2: 735–784.
- Schlenker, P. (to appear a). Gestural Grammar. Accepted with minor revisions, *Natural Language & Linguistic Theory*.
- Schlenker, P. (to appear b). Iconic Presuppositions. *Natural Language & Linguistic Theory*.
- Schlenker, P. (to appear c). Musical Meaning within Super Semantics. Accepted with minor revisions in *Linguistics & Philosophy*.
- Schlenker, P. (to appear d). Iconic Presuppositions. *Natural Language & Linguistic Theory*.
- Schlenker, P., & Chemla, E. (2018). Gestural Agreement. *Natural Language & Linguistic Theory* 36(2), 587-625. <https://doi.org/10.1007/s11049-017-9378-8>
- Schlenker, P. & Lamberton, J. (2019a). Iconic Plurality. *Linguistics & Philosophy* 42, 45-108. <https://doi.org/10.1007/s10988-018-9236-0>
- Schlenker, P. & Lamberton, J. (2019b). Meaningful Blurs: the Sources of Repetition-based Plurals in ASL. Manuscript, Institut Jean-Nicod and New York University.
- Schlenker, P., Lamberton, J. & Santoro, M. (2013). Iconic Variables. *Linguistics & Philosophy* 36(2), 91-149. <https://doi.org/10.1007/s10988-013-9129-1>
- Schlenker, P., Chemla, E., Arnold, K., Lemasson, A., Ouattara, K., Keenan, S., Stephan, C., Ryder, R., & Zuberbühler, K. (2014). Monkey Semantics: Two 'Dialects' of Campbell's Monkey Alarm Calls. *Linguistics & Philosophy* 37(6), 439-501. <https://doi.org/10.1007/s10988-014-9155-7>
- Schlenker, P., Chemla, E., Schel, A., Fuller, J., Gautier, J-P., Kuhn, J., Veselinovic, D., Arnold, K., Cäsar, C., Keenan, S., Lemasson, A., Ouattara, K., Ryder, R., & Zuberbühler, K. (2016a). Formal Monkey Linguistics. *Theoretical Linguistics* 42(1-2), 1–90. <https://doi.org/10.1515/tl-2016-0001>
- Schlenker, P., Chemla, E., & Zuberbühler, K. (2016b). What do Monkey Calls Mean?. *Trends in Cognitive Sciences*, 20, 12, 894–904. <https://doi.org/10.1016/j.tics.2016.10.004>
- Schlenker, P., Chemla, E., Cäsar, C., Ryder, R., & Zuberbühler, K. (2016c). Titi Semantics: Context and Meaning in Titi Monkey Call Sequences. *Natural Language & Linguistic Theory*, 35(1), 271-298. <https://doi.org/10.1007/s11049-016-9337-9>
- Simons, M. (2007). Observations on embedding verbs, evidentiality, and presupposition. *Lingua*, 117(6), 1034-1056. <https://doi.org/10.1016/j.lingua.2006.05.006>
- Steinert-Threlkeld, S., Schlenker, P., & Chemla, E. (2019). Referential and General Calls in Primate Semantics. Manuscript, U. of Washington, IJN and LSCP.
- Stokoe, W. C. (1960). Sign Language Structure: An Outline of the Visual Communication System of the American Deaf. *Studies in Linguistics*, Occasional Papers, Vol. 8, Buffalo, NY: University of Buffalo.
- Tarski, A. (1935). The Concept of Truth in Formalized Languages. *Logic, Semantics, Metamathematics* (pp. 152-278). Indianapolis: Hackett 1983, 2nd edition.
- Tarski, A. (1944). The Semantic Conception of Truth and the Foundations of Semantics. *Philosophy and phenomenological research*, 4(3), 341-376. <https://doi.org/10.2307/2102968>
- Thijsse, E. (1983). On some Proposed Universals of Natural Language. In A. ter Meulen (ed.), *Studies in Model-Theoretic Semantics* (pp. 19-36). Foris, Dordrecht.

- Thomason, R. H. (1974). *Formal Philosophy. Selected Papers of Richard Montague*, New Haven: Yale University Press.
- Tieu, L., Pasternak, R., Schlenker, P., & Chemla, E. (2017). Co-speech gesture projection: Evidence from truth-value judgment and picture selection tasks. *Glossa* 2(1), 102. <http://doi.org/10.5334/gjgl.334>
- Tieu, L., Pasternak, R., Schlenker, P., & Chemla, E. (2018). Co-speech gesture projection: Evidence from inferential judgments. *Glossa* 3(1), 109. <http://doi.org/10.5334/gjgl.580>
- Tieu, L., Schlenker, P., & Chemla, E. (2019). Linguistic Inferences Without Words. *Proceedings of the National Academy of Sciences* 116(20), 9796-9801. <https://doi.org/10.1073/pnas.1821018116>
- Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: the conjunction fallacy in probability judgment. *Psychological Review* 90(4), 293–315. <https://doi.org/10.1037/0033-295X.90.4.293>
- van Heijningen, C. A. A., de Visser, J., Zuidema, W., & ten Cate, C. (2009). Simple rules can explain discrimination of putative recursive syntactic structures by a songbird species. *PNAS* 106, 20538-20543. <https://doi.org/10.1073/pnas.0908113106>
- Walsh, C., & Johnson-Laird, P. N. (2004). Coreference and reasoning. *Memory and Cognition* 32, 96–106.
- Wertheimer, M. (1923). Laws of Organization in Perceptual Forms. In *A Source Book of Gestalt Psychology*. W. D. Ellis, ed. London: Routledge & Kegan Paul, 1938.
- Wiltschko, M. (2018). Discovering syntactic variation. In *Syntactic Structures after 60 Years. The Impact of the Chomskyan Revolution in Linguistics*. N. Hornstein, H. Lasnik, P. Patel-Grosz and C. Yang, ed. Berlin: De Gruyter Mouton, 427-460. <https://doi.org/10.1515/9781501506925-431>
- Zaradzki, L. (2021). *Les évènements en sémantique linguistique et musicale*. Doctoral dissertation, Université Paris Diderot.
- Zuberbühler, K. (2009). Survivor signals: the biology and psychology of animal alarm calling. *Advances in the Study of Behavior* 40, 277-322. [https://doi.org/10.1016/S0065-3454\(09\)40008-1](https://doi.org/10.1016/S0065-3454(09)40008-1)
- Zuberbühler, K., Jenny, D., & Bshary, R. (1999). The predator deterrence function of primate alarm calls. *Ethology* 105(6), 477–490. <https://doi.org/10.1046/j.1439-0310.1999.00396.x>