Towards a model of incremental composition[1]

Sigrid BECK — *Universität Tübingen*
Sonja TIEMANN — *Universität Tübingen*

Abstract. This paper reviews some recent psycholinguistic results on semantic processing and explores their consequences for a cognitively plausible model of incremental composition. We argue that semantic composition is neither strictly incremental (in the sense that every incoming word is composed immediately) nor global (in the sense that composition only proceeds when the entire syntactic structure is available). We conjecture that incremental composition is type driven: elements in the same type domain (e.g. temporal <i>) are composed immediately; elements that concern different type domains (e.g. temporal <i> vs. event <v>) cause delayed processing.

Keywords: Incrementality, composition, semantic processing.

1. Introduction

The central question explored in this paper is how a theory of semantic composition can be combined with processing results arguing that interpretation has incremental properties. Semantic theory takes as its starting point the principle of compositionality: The interpretation of a complex expression is determined by the interpretations of its parts and the way they are combined. This is usually implemented in terms of assigning a compositional interpretation to the **complete** syntactic representation of the sentence to be interpreted (see standard introductions, e.g. Heim and Kratzer, 1998; Chierchia and McConnel-Ginet, 2000; Zimmermann and Sternefeld, 2014; Beck and Gergel, 2014). However, certain results from psycholinguistic research fairly clearly show that people begin to compose meanings **before** the end of the sentence has been perceived, so no complete tree is available. A suggestive data point is the familiar garden path effect (Bever, 1970), illustrated in (1) (Ferreira, Christianson and Hollingsworth, 2001).

(1) While Anna dressed the baby spit up on the bed.

So far, neither compositional semantics nor psycholinguistics has established a psycholinguistically plausible model of how compositional interpretation proceeds incrementally (cf. also e.g. Chater et al., 2001; Bott and Sternefeld, to appear). This is a problematic gap in linguistic theory because results on semantic processing don't get integrated into a theory of the semantic parser or aligned with the theory of composition. This paper is a contribution towards closing this gap. Our plot is to use a standard semantic framework (concretely a Heim and Kratzer, 1998, type theory) as our starting point, and revise it according to a set of concrete processing results we take to be exemplary, as a step towards a model of incremental composition.

[1] We are very grateful to Gerhard Jäger, Detmar Meurers, Lyn Frazier, Florian Schwarz, David Beaver, Mark Steedman, the members of the SFB 833 'Construction of Meaning' and the audience at Sinn und Bedeutung 21, Edinburgh, for their helpful comments on this complex and daunting enterprise. We thank the DFG grant to the SFB 833 'Construction of Meaning' for financial support.

Section 2 defines our task and outlines some relevant existing work in semantics and in psycholinguistics. In section 3, we present the evidence for incremental composition from a set of processing studies and an incremental analysis of those data points. Section 4 combines the results of section 3 and generalizes towards a model of incremental interpretation. Our conclusions are presented in section 5.

2. Specifying the task

2.1. What is needed?

Standard semantic theory defines an interpretation function [[.]] recursively. [[.]] maps LF trees to meanings (e.g. Heim and Kratzer, 1998). Compositional interpretation proceeds as sketched in (3) for a simple example. We call this **global interpretation**.

(2) T -- [[]] --> $\cup D\tau$ (τ a semantic type)

(3) a. John invited Bill.
b. structure: [$_{IP}$ John [$_{VP}$ invited Bill]]
c. [[[$_{IP}$ John [$_{VP}$ invited Bill]]]] = 1 iff (2x Function Application FA)
[[invited]] ([[Bill]]) ([[John]]) = 1 iff (Lexicon + λ-conversion)
John invited Bill.

What if we start interpreting on the basis of a partial structure (4a)? Plausibly we anticipate (4c). A theory that can predict (4c) requires the concepts defined in (5), (6) and (7).

(4) a. partial structure: [$_{IP}$ John [$_{VP}$ invited …
b. from the lexicon: {[[John]], [[invited]]}
c. projected meaning: λy. [[invited]](y)([[John]]) =
λy. John invited y

(5) a. Let Θ be the set of syntactic structures produced by the human parser.
Each $T_i \in \Theta$ is a possibly partial syntax tree.
b. Let Σ be the set of interpretations produced by the corresponding human interpretive processor. The elements of Σ are sets of meanings, i.e. each $S_i \in \Sigma$ is a set whose members are elements of $\cup D\tau$ (τ a semantic type).
c. A pair $<T_i, S_i>$ is a stage reached in sentence processing.

(6) Incremental processing is a series of mappings $<T_i,S_i> \to <T_{i+1},S_{i+1}>$ ($1 \leq i \leq n$) such that
(i) T_n is an LF tree;
(ii) each mapping $T_i \to T_{i+1}$ is a matter of parsing (not our concern here);
(iii) each S_i is a set of meanings from $\cup D\tau$;
(iv) card(S_n)=1 (i.e. everything is composed into one meaning in the end);
(v) $<T_n,S_n> \in$[[.]].

(7) Incremental composition is the derivation of S_{i+1} from S_i.
Define a function $[[.]]_h$ ('heuristic interpretation'):
Suppose at stage i, the processor receives the structure σ as input, leading to T_{i+1}.
$[[.]]_h$ defines a mapping $<T_{i+1}, S_i, [[\sigma]]_h> \rightarrow S_{i+1}$.
On the basis of the new tree, the available set of meanings plus the new meaning, a new semantic stage is reached.

We can think of the function $[[.]]_h$ as an interpretive heuristic. It makes predictions about the meaning of partial trees, yielding a projected or anticipated meaning (which could be proven wrong by further input). A model of incremental composition is a recursive definition of the function $[[.]]_h$. For each stage that the parser may reach, $[[.]]_h$ defines the accompanying stage of the interpreter.

2.2. What has been proposed in semantics?

Several linguistic frameworks have made proposals towards incremental interpretation, prominently including the categorial grammar tradition. A representative is Combinatory Categorial Grammar CCG (e.g. Ades and Steedman, 1982; Steedman, 2000; Steedman and Baldridge, 2011) and a simple example is given below. The syntax of CCG allows the incremental parse in (8b) — we are still looking for an NP to complete the sentence — and the semantics corresponds to this, employing Function Composition (9) to compose the meaning of the subject and the meaning of the verb as in (8d).

(8)
a. John invited Bill.

b. basic syntax:	John	invited	Bill
	S/(S\NP)	(S\NP)/NP	NP
		S\NP	
c. incremental parse:	John	invited	
	S/(S\NP)	(S\NP)/NP	
		S/NP	

d. semantics:
$[[\text{John}]] = \lambda P_{<e,t>}.P(\text{John})$
$[[\text{invited}]] = \lambda y.\lambda z.z \text{ invited } y$
$[[\text{John}]] \bullet [[\text{invited}]] = [\lambda P_{<e,t>}.P(\text{John})] \bullet [\lambda y.\lambda z.z \text{ invited } y]$
$= \lambda x. [\lambda P_{<e,t>}.P(\text{John})]([\lambda y.\lambda z.z \text{ invited } y](x))$
$= \lambda x. \text{ John invited } x$

(9) Function composition:
If g is a function: A->B and f is a function: B->C then
$f \bullet g$: A->C is the composition of f and g with $f \bullet g = \lambda x.f(g(x))$

The example illustrates what we call **strict incrementality**. Each new element that is parsed is added immediately to the tree and composed immediately with the semantics already available. At each stage i, $card(S_i)=1$. In section 3, we will reject strict incrementality as a property of the semantic processor. There is a tendency in the CCG tradition towards strict incrementality (recently e.g. Kato and Matsubara, 2015), though details vary and Steedman's

(2000) Strict Competence Hypothesis SCH does not lead one to always expect strict word-by-word incrementality (see e.g. Demberg, 2012; Ambati, 2016 for discussion).

Outside CCG, our most immediate predecessor in the search of a model of incremental composition is Bott and Sternefeld (to appear). Bott and Sternefeld differ from us in two important respects: (i) they aim for strict incrementality, and (ii) they use a different framework (namely a dynamic Neo-Davidsonian continuation semantics with unconstrained λ-conversion), which we will not present here. But they point out the same gap in linguistic theory (cf. their paper also for further references), they consult an overlapping set of psycholinguistic results to inform their model of incremental composition, and they develop an incremental perspective on complex semantic analyses e.g. of tense and aspect. We return to their paper below.

2.3. What has been done in processing?

First, a cautionary note: We want to know when **composition** of meanings in complex structures occurs; not all results to do with immediate semantic processing are therefore of relevance for us (e.g., a finding could be based on immediate lexical access but not immediate composition; see e.g. Altmann and Kamide, 1999; or Frazier, 1999 for an overview). In recent years, psycholinguistic research on semantics has produced a lot of results on how different phenomena are processed (e.g. quantifiers, presuppositions etc. - a recent overview is given in Pylkkänen and McElree, 2006). The findings indicate particular properties of semantic processing and define certain constraints on it. What has not been established is a semantic processing model in the sense of heuristic composition $[[.]]_h$, i.e. there is no model that we know of (with the exception of Bott and Sternefeld) which describes how actual incremental composition works.

There is of course more work on syntactic processing, and this is important as the input to compositional interpretation (see e.g. discussion in Crocker, 2010). Resnik (1992), building on earlier work, argues for an *arc-eager* left corner parser, i.e. a variant of a left corner parser in which nodes that are predicted bottom-up can be immediately composed with nodes that are predicted top-down. We assume that the syntactic processor continuously integrates new material in a roughly left corner parser fashion. At any point during processing, a partial tree structure like (10) is projected by the syntactic parser. In this paper we simplify in that only one parse tree will be entertained as a possible structure at a time. (Ideally, we would adopt whatever proposal about the parser is best motivated.) Importantly, the tree T_i is the LF structure (the input to compositional interpretation). The terminal nodes in T_i include the words heard so far (in their proper places in the structure). The tree is the projected syntax.

(10) a.


b.

Another input to compositional interpretation is lexical meaning. There is evidence that the interpretation of lexical terminal nodes is available immediately from the lexicon (e.g. Frazier and Rayner, 1990). Hence we assume that these meanings are added to S_i (the set of meanings made available by the parse so far). And finally, compositional interpretation depends on the values assigned to variables. Free variables get their value from the context via the salient variable assignment function g_c. When all goes well, they are assigned their values immediately (e.g. Carreiras and Clifton, 1993). The resulting meanings are also added to S_i. When there is no salient referent, binding of the variable is preferred; this may lead to a revision of the LF tree in such a way as to include a binder, or to optimize the chances of including a binder (that is, it can lead to delayed semantic interpretation, see section 3.2.; Bott and Schlotterbeck,2013).

The anticipated lexical meaning and anticipated contextual reference are the recursion basis for the function $[[.]]_h$: If α is a terminal element, $[[\alpha]]_h=[[\alpha]]$.

As an **interim summary**, we note that two interpretive strategies (11), (12) are made readily available by existing theories of interpretation. As a preview of what is to come, we argue that neither type of approach is the desired model of incremental composition. (Of course Global interpretation is not claimed to be a model of the semantic processor in the first place).

(11) **Global interpretation**:
Assume a syntactically **complete** parse tree T (i.e. no "…").
The meanings in S (here, the terminal nodes in T) are composed by [[.]]
according to T and the standard composition principles.

(12) **Strictly incremental interpretation**:
Assume an incrementally generated **partial** tree T, and a set of available meanings S.
Whenever card(S)>1, compose the meanings in S according to T and some
combinatory heuristic $[[.]]_h$.

3. Some psycholinguistic findings on incremental composition

In section 3.1. we collect a set of experimental results that argue for immediate composition in certain sentence contexts, and offer an incremental analysis of these cases. In section 3.2. we consider several cases of delayed composition, i.e. experimental evidence that there is no strictly incremental composition. The section summary sets the scene for our generalizations in section 4.

3.1. Results supporting immediate composition

Subject + Verb: There are early effects indicating that before the object is encountered, the meanings of the subject and the verb are already put together (e.g. Kuperberg et al., 2003; Kim and Osterhout, 2005; Kamide et al., 2003; Knoeferle et al., 2005). We take this to mean that subject and verb are interpreted incrementally, before the sentence is finished. (13)

illustrates the relevant structure; # marks the point where studies have found an interpretive effect in processing.

(13) The soup greeted …
|
#

This invites the following interpretation: The parse tree contains (14a); this leads to a combination of the meanings of the verb and the subject as in (14b). (14b) can be derived by the heuristic rule in (15). If we suppose that the subject has the type <<e,t>,t> rather than <e>, the alternative formulation in (16) is applicable. (16) amounts to function composition (17).

(14) a. projected parse tree contains:

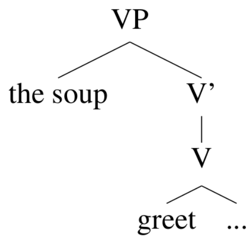

b. projected meaning:
$\lambda y.[[\text{greet}]](y)([[\text{the soup}]]) =$
$\lambda y.\text{the soup greet } y$

(15) Subject-Verb-Heuristic (SVH):
If $\alpha = [\ \beta_{subj}\ [\ \gamma_{verb} \ldots$
then $[[\alpha]]_h = \lambda y.[[\gamma_{<e,et>}]]_h(y)([[\ \beta_{<e>}]]_h)$

(16) Subject-Verb-Heuristic (<<e,t>,t> subject) (SVH'):
If $\alpha = [\ \beta_{subj}\ [\ \gamma_{verb} \ldots$
then $[[\alpha]]_h = \lambda y.[[\ \beta_{<et,t>}]]_h([[\ \gamma_{<e,et>}]]_h(y))$
SVH' defines $[[\beta]]_h \bullet [[\gamma]]_h$

(17) Function composition:
If g is a function: A->B and f is a function: B->C then
$f \bullet g$: A->C is the composition of f and g with $f \bullet g = \lambda x.f(g(x))$

Within DP: There is evidence that determiner and adjective are combined very early on (Sedivy et al., 1999). (18) illustrates this. Our interpretation is that the meaning of the determiner plus the meaning of the NP is incrementally interpreted as indicated in (19).

(18) Touch the tall …
|
#

(19) a. projected parse tree contains:

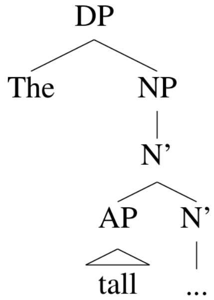

b. projected meaning:
$\lambda P.[[\text{the}]]([[\text{tall}]] \cap P)$

Part of predicting (19b) is the expectation that the meaning of the determiner is applied to a suitable argument, as modelled by the DP-Heuristic below. This heuristic defines predictive Function Application FA.

(20) DP-Heuristic:
If $\alpha = [_{DP}\ \beta_{Det}\ [_{NP}\ \gamma \ldots$
then $[[\alpha]]_h = [[\beta]]_h\ ([[NP]]_h)$

Subject + Adverb: Under certain circumstances, people anticipate a complex meaning given the input of a subject plus an adverb, here *wieder* 'again' (Tiemann, 2014; Tiemann et al., 2011). We interpret this as our participants anticipating that the adverb will modify some property attributed to the subject.

(21) a. context: Inge hat letzte Woche rote Handschuhe gekauft.
Inge has last week red gloves bought
b. Susanne hat wieder…
Susanne has again
|
#

(22) a. projected parse tree contains:

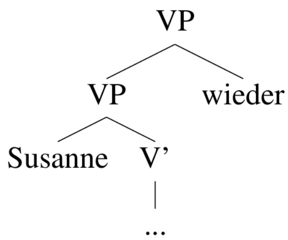

b. projected meaning:
$\lambda P_{<e,vt>}.\ [[\text{wieder}]](\ \lambda e.P(e)(\text{Susanne})) =$
$\lambda P_{<e,vt>}.\lambda e{:}\exists e'[e'<e\ \&\ P(e')(\text{Susanne})].P(e)(\text{Susanne})$

This projected meaning can be predicted by the heuristic rule in (23). Once more, if the subject is taken to be of type <<e,t>,t> rather than <e>, the heuristic is rephrased in such a way as to reveal it as function composition (24).

(23) <v,t>-Adverb-Subject Heuristic (AdvSH):
If α= [[β_{subj} …] γ_{adverb}] and γ is of type <vt,vt>,
then $[[\alpha]]_h = \lambda P_{<e,vt>}.[[\gamma]]_h(P([[\ \beta]]_h))$

(24) <v,t>-Adverb-Subject Heuristic (<<e,t>,t> subject) (AdvSH'):
If α= [[β_{subj} …] γ_{adverb}] and γ is of type <vt,vt>,
then $[[\alpha]]_h = \lambda P_{<e,vt>}.[[\gamma]]_h([[\ \beta]]_h(P))$
AdvSH' defines $[[\gamma]]_h\bullet[[\beta]]_h$

Temporal Adverb + Tense: Bott (2010) found that participants respond immediately to a mismatch between verbal tense and the meaning of an adverb, as in (25). Our take on what happens in processing is (26).

(25) Morgen gewann…
tomorrow won
|
#

(26) a. projected parse tree contains:

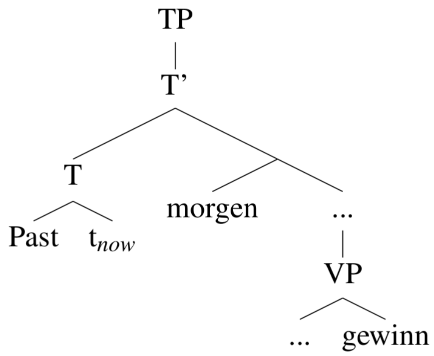

b. projected meaning:
$\lambda P_{<i,t>}.[[\text{PAST } t_{now}]](\lambda t'.t'\subseteq\text{tomorrow} \;\&\; P(t')) =$
$\lambda P_{<i,t>}.\exists t'[t'< t_{now} \;\&\; t'\subseteq \text{tomorrow} \;\&\; P(t')]$

The heuristic predicting this anticipated interpretation and the clash contained in it can be phrased as in (27) (assuming <i,t> type modifiers) or (28) (assuming <<i,t>,<i,t>> modifiers).

(27) Temporal adverb – Tense Heuristic (intersective modifiers):
If α= [β_{Tense} [γ_{adverb} …]] and γ is of type <it>,
then $[[\alpha]]_h = \lambda P_{<i,t>}.[[\beta]]_h(\lambda t'[[\gamma]]_h(t') \;\&\; P(t'))$

(28) Temporal adverb – Tense Heuristic (functional modifiers):
If α= [β_{Tense} [γ_{adverb} …]] and γ is of type <it,it>,
then $[[\alpha]]_h = \lambda P_{<i,t>}.[[\beta]]_h([[\gamma]]_h(P))$
The Temporal adverb–Tense Heuristic (shifted) defines $[[\beta]]_h\bullet[[\gamma]]_h$

Russian Aspect: Bott and Gattnar (to appear) found that in Russian a mismatch of aspect with an adverb (29) was detected immediately, suggesting incremental interpretation. This finding is especially interesting when compared to the processing of German aspect, discussed in the next subsection. The Russian results indicate that the compositional step in (30) is taken immediately:

(29) Celych tri casa vyigrala
Whole three hours win.pfv.Past …
|
#

(30) a. projected parse tree contains:

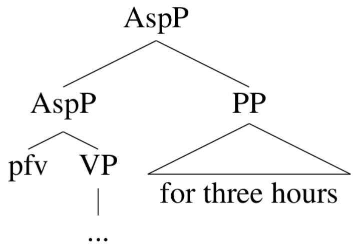

b. projected meaning:
$\lambda P_{<v,t>}. \exists e[[[pfv]](e) \;\&\; [[for\ three\ hours]](e) \;\&\; P(e)]$

No detailed analysis or heuristic is offered because the details of the analysis for Russian are not sufficiently clear to us (see Bott and Gattnar, to appear; Bott and Sternefeld, to appear, for discussion). It is clear however that there is an immediately perceived clash between the aspect information and the adverbial. Those two expressions must be part of a local tree in the AspP. We conjecture that the example is (abstractly) parallel to the temporal adverb–tense case above.

To sum up this subsection, we have identified five circumstances that showcase immediate composition of two semantic units that do not form a constituent in the LF. Note that in each case, the two units occur in the same LF domain (DP, VP, TP, AspP). Their combination may be understood as predicted function application or function composition. (See Bott and Sternefeld, to appear, for a different, but similarly incremental analysis of e.g. the tense and the aspect cases.)

3.2. Results supporting delayed composition

Quantifiers: Hackl et al. (2012) (see also Varoutis and Hackl, 2006; Breakstone et al., 2011; but cf. Gibson et al., 2014) argue that there is evidence for quantifier raising (QR) and delayed interpretation of quantifiers in object position (31). We illustrate our interpretation of this finding in (32): encountering the quantified determiner leads to a revision of the parse tree.

(31) John fed every dog.
|
no composition here.

(32) a. 1st projected parse tree contains:

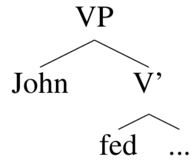

b. revised parse tree contains:

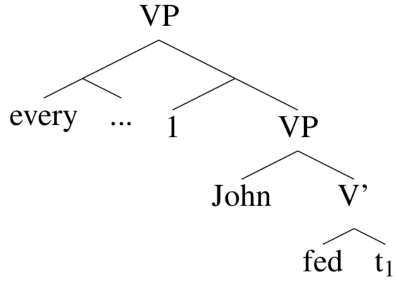

Note that consequently, it is not the case that the meaning calculated so far — presumably [λy. John fed y] — is combined with the meaning of *every* (yielding e.g. [$\lambda P_{<e,t>}$. for every y such that P(y), John fed y]). Hence this is a case that does not work according to strict incrementality. It seems extremely plausible that recovering from a garden path like (1) also involves such a revision (cf. e.g. Chater et al., 2001), i.e. in addition to throwing out the parse that turned out to be misguided, the corresponding interpretation is thrown out along with it.

German Aspect: Bott (2013) and Bott and Gattnar (to appear) show that aspectual mismatch in German is only processed when the verb has received its full argument structure, suggesting that the meaning of an adverbial ('for two hours') is not immediately combined with the meaning of a verb ('won'). Composition only happens later (in contrast to Russian).

(33)	Zwei Stunden lang	gewann	der Boxer	den Kampf.
	two hours for	won	the boxer	the fight
		|		
		no composition here.		

It seems plausible that the meanings of the available items are added to the set of meanings made available by the processor, but not composed. So this is an instance of delayed composition. In very general terms, the so-called sentence wrap-up effect (Just and Carpenter, 1980) may also be an indication of late processes in semantic composition.

Further candidates: We mention two further candidates that have been presented as indicators for late composition processes. The model in section 4 will not properly include them because they involve semantic issues we can't yet address (variable binding and presupposition projection) but they provide general support of our position. First, Bott and Schlotterbeck (2013) present an eyetracking study investigating the processing of inverse scope as in (34a) vs. (34b) without scope inversion. Their results suggest that scope inversion is only computed at the end of the sentence (this is also the interpretation of this finding in Bott and Sternefeld).

(34) a. Jeden seiner Schüler hat genau ein Lehrer voller Wohlwollen gelobt.
Each of-his pupils has exactly one teacher full-of goodwill praised
'A teacher praised each of his pupils full of goodwill.'
b. Jeden dieser Schüler hat genau ein Lehrer voller Wohlwollen gelobt.
Each of-these pupils has exactly one teacher full-of goodwill praised
'A teacher praised each of these students full of goodwill.'

Second, Schwarz and Tiemann (to appear) conducted experiments on the processing of sentences with unfulfilled embedded and unembedded presuppositions (35a,b). Whilst presupposition failure in the unembedded cases (35b) was immediately detected in online processing, there was no such effect in the embedded conditions. We take this to mean that the composition of embedded presuppositions does not happen strictly incrementally, otherwise presupposition failure should result in immediate processing effects as they do in the unembedded cases.

(35) a. Heute war Tina nicht wieder schlittschuhlaufen.
Today was Tina not again ice-skating
'Today, Tina didn't go ice skating again.'
b. Heute war Tina wieder nicht schlittschuhlaufen.
Today was Tina again not ice-skating
'Today, once more Tina didn't go ice skating.'

In sum, we have evidence that semantic units are not always composed immediately. Predictive combinatory mechanisms do not seem to be explored to exhaustion to calculate a composed meaning under all circumstances. This is why we depart from strict incrementality (as developed e.g. in Bott and Sternefeld).

Interim Conclusion: If the above view is correct, neither global interpretation nor strict incrementality seems to be the right model of semantic processing. Composition in semantic processing has incremental properties, but it also seems to require certain units to be built before composition proceeds. The required model needs to employ what we might call **enlightened incrementality**: sometimes composition is immediate, but under other circumstances it is delayed. What would be a useful hypothesis about when the processor applies which type of strategy? The next subsection addresses this question.

4. First steps towards a general framework

This section generalizes from the concrete incremental compositional analyses in section 3. The desired outcome is (the beginnings of) a framework for theories of semantic parsing: a definition of a function $[[.]]_h$ ('heuristic interpretation') as anticipated in section 2. Naturally, we are far from being able to propose a complete model for this mapping. But we can distill some generalizations from the case studies in section 3. We propose that a realistic semantic processor sometimes composes 'early' and sometimes 'late', depending on the linguistic input. Our evidence indicates the general possibility of four cases: (i) wait and see, (ii) revision of LF, (iii) predictive Function Application (FA), (iv) predictive Function Composition (FC). These are generalizations over the interpretive strategies that section 3 provides evidence for. Subsection 4.1. examines the 'late' strategies, subsection 4.2. the 'early' composition strategies. In subsection 4.3. we develop a hypothesis as to when the semantic processor employs which type of strategy.

4.1. Delayed composition

Beginning with 'late' composition strategies, section 3 provides evidence for (i) wait and see. The example indicative of this strategy is German aspect (and also (34), (35)). Stages of the processor are sketched in (37).

(i) **wait and see**
Given <T,S> and input σ, map to <T',S'>,
where T' is the modification of T derived by the syntactic parser and S' is defined by:
$[[.]]_h$: <T',S,$[[\sigma]]_h$> -> S∪{$[[\sigma]]_h$}

(36) Zwei Stunden lang gewann der Boxer den Kampf.
two hours for won the boxer the fight
|
no composition here.

(37) a. T=[CP [PP zwei Stunden lang] …]]
S={[[for 2h]]}
b. T'=[CP [PP zwei Stunden lang] _ [TP Past …]]]
S'={[[for 2h]], [[Past]]}

The second case of non-incremental interpretation we have seen is (ii) revision of LF. The example for this strategy from section 3 is quantifiers in object position (and also (1)).

(ii) **revision of LF**
Given <T,S> and input σ, map to <T',S'>,
where T' is the revision of T derived by the syntactic parser, and S' is defined by:
$[[.]]_h$: <T', S, $[[\sigma]]_h$> ->{x: x is the meaning of an atom in T}∪$[[\sigma]]_h$>

(38) John fed every dog.
|
no composition here.

At this point we digress a little in order to explain more fully our take on what happens in (38). Revision of the parse tree from T to T' would be compatible with keeping the meanings composed so far and adding the new meaning, according to the (i) wait and see strategy, as sketched in (39). We conjecture, however, that the processor also reconsiders the store of meanings. Our motivation comes from examples that have, in addition, a quantifier in subject position, (40). If the processor kept the meanings composed so far, we would get (41). Continuing processing on this basis would in our framework (i.e. without type shifting or further scope mechanisms) lead to the inverse scope reading (40b). It seems implausible that the processor smoothly generates the intuitively harder reading. It is more plausible that the processing of the doubly quantified example involves the steps in (42) — the composition of subject and verb is thrown out. This motivates our assumption that (ii) is operative in this case.

(39) a. T= [VP John [V' fed …]]
S={[[John fed]]} = {λy. John fed y}
b. T'=[VP [NP every …][1[VP John [V' fed t1]]]
S'= {λy. John fed y, [[every]]}

(40) Some guy fed every dog.
a. ∃x[∀y[dog(y) -> x fed y]] (surface scope)
b. ∀y[dog(y) -> ∃x[x fed y]] (inverse scope)

(41) a. T= [VP [NP some guy] [V' fed …]]
S={[[some guy fed]]} = {λy. ∃x[x fed y]}
b. T'=[VP [NP every …][1 [VP [NP some guy] [V' fed t1]]]]
S'= { [[every]], λy. ∃x[x fed y]}

(42) a. T=[IP _ [I' [VP [NP some guy] [V' fed …]]]]
S={[[some guy fed]]} = {λy. ∃x[x fed y]}
b. T'=[IP [NP some guy] [2[I' [VP [NP every …][1[VP t2 [V' fed t1]]]]
S'= {[[some guy]], [[every]], [[fed]]}

In the (ii) revision of LF case, therefore, the processor performs a revision of the parse tree and throws out a predicted meaning in S as well, reconsidering composition.

4.2. Incremental composition.

Let's next turn to 'early' composition. Section 3 anticipates (iii) predictive Function Application.

(iii) **predictive Function Application (FA)**
Given <T,S> and input σ, map to <T',S'>,
where T' is derived by the syntactic parser and
if there is a δ∈S such that (a) $[[\sigma]]_h(\delta)$ or
(b) $\delta([[\sigma]]_h)$ is defined, then,
(a) $S'=S\backslash\delta \cup \{[[\sigma]]_h(\delta)\}$ or
(b) $S'=S\backslash\delta \cup \{\delta([[\sigma]]_h)\}$ (whichever is defined).

The example from section 3 is immediate compositional interpretation in the DP, (43). Predictive FA would similarly be involved in (44). This proposal could be further tested by data like (45), which we give as a suggestion for future research.

(43) Touch the tall …

(44) a. Every dog…
b. S={[[every]]}, S'={[[every]]([[dog]])}

(45) a. Every dog that greeted its master was fed.
b. Every dog was fed that greeted its master.

The second ‘early’ composition mechanism from section 3 is (iv) predictive Function Composition (assuming the higher types for the heuristics). Examples from above were the data types in (46).

(iv) **predictive Function Composition (FC)**
Given <T,S> and input σ, map to <T',S'>,
where T' is derived by the syntactic parser and
if there is a $\delta \in S$ such that (a) $\delta \bullet [[\sigma]]_h$ or
(b) $[[\sigma]]_h \bullet \delta$ is defined, then,
(a) $S' = S \backslash \delta \cup \{\delta \bullet [[\sigma]]_h\}$ or
(b) $S' = S \backslash \delta \cup \{[[\sigma]]_h \bullet \delta\}$ (whichever is defined).

(46) a. The soup greeted… (subject–verb)
b. Morgen gewann… (tense–adverb)
tomorrow won …
c. Susanne hat wieder… (subject–adverb)
Susanne has again …

4.3. When is composition ‘early’ and when ‘delayed’ – a possible generalization

A model of semantic processing in the sense of enlightened incrementality should be an optimal compromise regarding two conflicting demands: (a) a low load on working memory: it is unrealistic that we carry around a large number of separate meanings until the end of an utterance; (b) reliable predictions: it is undesirable to randomly compose word meanings when the confidence that this is the actual interpretation is low. We offer the conjecture below for what this compromise could look like.

(47) **Enlightened Incrementality Conjecture**:
Units in the same LF domain (DP, VP, TP, AspP). are composed incrementally.

The idea is that there is incremental (‘early’) composition, but it is limited to a local LF domain. LF domains are defined by semantic type. E.g., we predictively combine the verb with its arguments within the VP <e,t> (‘the soup greeted…’). We predictively combine tense with temporal adverbials within the TP layer <i,t> (‘Morgen gewann…’) and event-level adverbials with expected event descriptions just above the core VP <v,t> (‘Susanne hat wieder…’). It appears that predictive composition occurs in layers. (This does not mean that you have to finish a layer before you start the next one, cf. ‘Morgen gewann…’.) The tree below illustrates the LF architecture this proposal is based on (e.g. von Stechow and Beck, 2015).

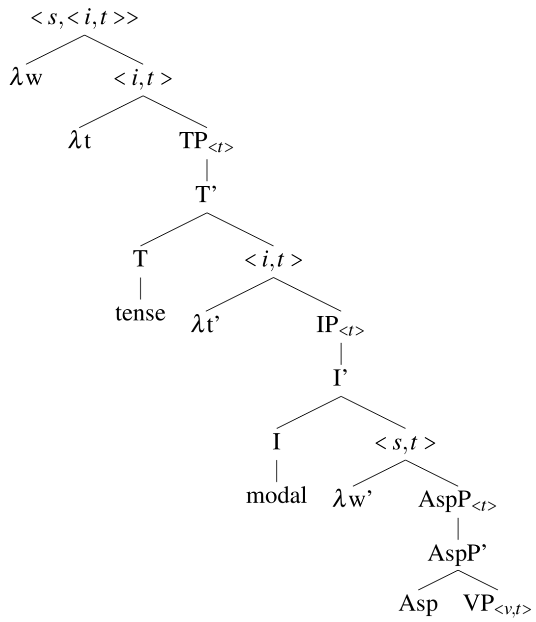

The examples that we have seen for 'delayed' composition, e.g. German aspect/Aktionsart ('Zwei Stunden lang gewann…'), concern material that in the LF is scattered over several layers (TP, AspP, VP). Quantifiers in object position also concern more than one layer: QR takes a quantifier above aspect as illustrated in (48) (e.g. von Stechow and Beck, 2015).

(48)

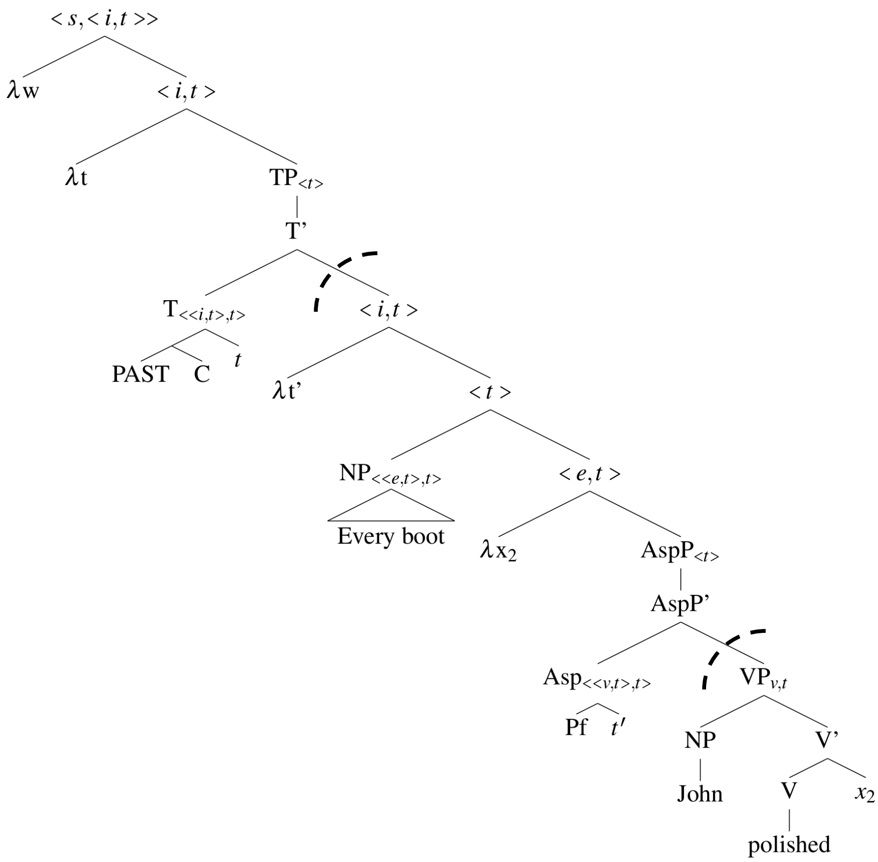

In sum, late composition facts mean that predictive FA and predictive FC cannot always apply. We conjecture that predictive composition happens in local LF domains, identifiable by semantic type, where the confidence that this is the correct composition is high.

Next steps: There are a couple of issues that need to be addressed for a more complete proposal. The QR data draw our attention to movement and the question of how Predicate Abstraction in standard composition transfers to incremental composition. Analyses are available in CCG (see e.g. Steedman, 2000; Demberg, 2012 for relevant discussion). Similarly, the tense and aspect data show that for an incremental analysis of a complete fragment, we need to think about the interaction of the several LF layers. A proposal is made in Bott and Sternefeld (to appear). We must leave an investigation of these issues, consideration of the available processing evidence and its integration into our proposal for future research.

5. Conclusions

We have seen that semantic processing has incremental properties (e.g. subject + verb seems to be composed immediately) but also 'global' properties, i.e. processing requires larger units (e.g. quantifiers). Standard theories of semantic composition do not model this because they require the whole LF tree and only assign meanings to constituents. Strictly incremental theories of semantic composition do not model this because every sentence prefix is assigned a meaning strictly incrementally. Hence the field is still in search of a model of incremental composition.

We formulate first ideas towards a definition of a heuristic interpretation function $[[.]]_h$ which models incremental composition (keeping as much as possible from standard semantic theories). Our goal is to offer the beginnings of a framework for theories of semantic parsing. Naturally, the question when and to what extent the semantic parser composes incrementally needs to be addressed for further phenomena (variables, decomposition phenomena, presupposition etc.).

Central to our proposal is the enlightened incrementality conjecture: incremental composition occurs within a local LF domain, when the confidence is high that the composition will prove correct.

Our model for a semantic processor concentrates on grammatically determined aspects of incremental interpretation. This is not to deny that other factors may enter into (incremental) understanding. An important factor is what we might call expectations, coming from e.g. frequency, contextual fit or background knowledge. It is clear that these factors affect processing, for example of garden path sentences (e.g. MacDonald, Perlmutter and Seidenberg, 1994) and even scope (Raffray and Pickering, 2010; Chemla and Bott, 2015). We take them to be relevant for our model as well: for instance, the very early effect on *again* noted in section 3 is plausibly due to the kind of context used in the experiment. One way this can be thought about is in terms of when to apply which heuristic rule. Hale (2003) models syntactic expectations by adding the likelihood of the application of a rule of the parser as a probability. A similar path would be open to models of the semantic processor. At any rate, we assume that a component handling such factors can and should be added to what we propose about the processor.

We find it important to model findings on processing in terms of a compositional semantic processor, even though our empirical knowledge in this area is still quite limited. We offer the heuristics in this paper as a framework for beginning this enterprise. If semanticists don't worry about incremental interpretation, and psycholinguists don't model the composition steps, there will be a regrettable gap in linguistic theory. Individual results on semantic processing remain isolated instead of contributing towards a theory of incremental interpretation.

References

Ades, A. E. and M. Steedman (1982). On the order of words. *Linguistics and Philosophy 4*, 517–558.
Altmann, G.T. and Y. Kamide. (1999). Incremental interpretation at verbs: Restricting the domain of subsequent reference. *Cognition 73*, 247–264.
Ambati, B. (2016). *Transition Based Combinatory Categorial Grammar parsing for English and Hindi*. Ph.D. thesis, University of Edinburgh.
Beck, S. and R. Gergel. (2014). *Contrasting English and German Grammar: An Introduction to Syntax and Semantics*. Berlin: de Gruyter.
Bever, T.G. (1970). The cognitive basis for linguistic structures. In J.R. Hayer (Ed.), *Cognition and the Development of Language*. New York: Wiley.
Bott. O. (2010). *The Processing of Events*. Amsterdam: John Benjamins.
Bott, O. (2013). The processing domain of aspectual interpretation. In B. Arsenijevic, B. Gehrke and R. Marín (Eds.), *Subatomic Semantics of Event Predicates*. Studies in Linguistics and Philosophy. Dordrecht: Springer.
Bott, O. and A. Gattnar (to appear). The cross-linguistic processing of aspect – An eyetracking study on the time-course of aspectual interpretation in German and Russian. *Language, Cognition and Neuroscience*.
Bott, O. and F. Schlotterbeck (2013). The processing domain of scope interaction. *Journal of Semantics 32*, 39–92.
Bott, O. and W. Sternefeld (to appear). An event semantics with continuations for incremental interpretation. *Journal of Semantics*.
Breakstone, M., A. Cremers, D. Fox and M. Hackl (2011). On the analysis of scope in comparative constructions. In N. Ashton, A. Chereches and D. Lutz (Eds.), *Proceedings of SALT 21*.
Carreiras, M. and C. Clifton. (1993). Relative clause interpretation preferences in Spanish and English. *Language and Speech 36*, 353–372.
Chater, N., M. Pickering and D. Milward (2001). What is incremental interpretation? Edinburgh Working papers in Cognitive Science 11.
Chemla, E. and L. Bott (2015). Structural priming to study representations and operations. To appear in *Linguistic Inquiry*.
Chierchia, G. and S. McConnell-Ginet. (2000). *Meaning and Grammar*. Cambridge, MA: MIT Press.

Crocker, M. (2010). Computational psycholinguistics. In A. Clark, C. Fox and S. Lappin (Eds), *Handbook of Computational Linguistics and Natural Language Processing*. London: Blackwell.
Demberg, V. (2012). Incremental derivations in CCG. In *Proceedings of the 11th International Workshop on Tree Adjoining Grammars and Related Formalisms: TAG+11*, pp.198–206.
Ferreira, F., K. Christianson and A. Hollingsworth (2001). Misinterpretations of garden-path sentences: Implications for models of sentence processing and reanalysis. *Journal of Psycholinguistic Research 30*, 3–20.
Frazier, L. (1999). *On Sentence Interpretation*. Dordrecht: Springer.
Frazier, L. (1979). *On Comprehending Sentences: Syntactic Parsing Strategies*. Ph.D. thesis, University of Connecticut.
Frazier, L. and K. Rayner (1990). Taking on semantic commitments: Processing multiple meanings vs. multiple senses. *Journal of Memory and Language 29*, 181–200.
Gibson, E., P. Jacobson, P. Graff, K. Mahowald, E. Fedorenko and S. Piantadosi (2014). A pragmatic account of complexity in definite antecedent contained deletion relative clauses. *Journal of Semantics 31*. doi: 10.1093/jos/ffu006.
Hackl, M., J. Koster-Hale and J. Varvoutis (2012). Quantification and ACD: Evidence from real-time sentence processing. *Journal of Semantics 29*, 145–206.
Hale, J. (2003). *Grammar, Uncertainty and Sentence Processing*. Ph.D. thesis, Johns Hopkins University.
Heim, I. and A. Kratzer (1998). *Semantics in Generative Grammar*. Oxford: Blackwell.
Just, M. and P. A. Carpenter (1980). A theory of reading: From eye fixations to comprehension. *Psychological Review 87*, 329–354.
Kamide, Y., G. T. M. Altmann and S. L. Haywood (2003). The time-course of prediction in incremental sentence processing: Evidence from anticipatory eye movements. *Journal of Memory and Language 49*, 133–156.
Kato, Y. and S. Matsubara (2015). Incremental semantic construction using normal form CCG derivation. *Proceedings of the Fourth Joint Conference on Lexical and Computational Semantics (*SEM 2015)*, pp. 269–278, Denver, Colorado, June 4–5, 2015.
Kim, A. and L. Osterhout (2005). The independence of combinatory semantic processing: Evidence from event-related brain potentials. *Journal of Memory and Language 52*, 205–222.
Knoeferle P., M.W. Crocker, C. Scheepers and M.J. Pickering (2005). The influence of immediate visual context on incremental thematic role-assignment: Evidence from eye-movements in depicted events. *Cognition 95*, 95–127.
Kuperberg, G.R., T. Sitnikova, D. Caplan and P.J. Holcomb (2003). Electrophysiological distinctions in processing conceptual relationships within simple sentences. *Cognitive Brain Research 17*, 117–129.
MacDonald, M., N. Perlmutter and M. Seidenberg (1994). Lexical nature of syntactic ambiguity resolution. *Psychological Review 101*, 676–703.
Pylkkänen, L. and McElree, B. (2006). The syntax-semantic interface: On-line composition of sentence meaning. In M. Traxler and M.A. Gernsbacher (Eds.), *Handbook of Psycholinguistics* (2nd edition), pp. 537-577. New York: Elsevier.
Raffray, C. N. and M. J. Pickering (2010). How do people construct logical form during language comprehension? *Psychological Science, 21*, 1090–1097. doi: 10.1177/0956797610375446

Resnik, P. (1992). Left-corner parsing and psychological plausibility. In *Proceedings of the Fourteenth International Conference on Computational Linguistics*, pp.191–197.

Schwarz, F. (2015). Presuppositions versus asserted content in online processing. In F. Schwarz (Ed.), *Experimental Perspectives on Presuppositions*, pp. 89–108.

Schwarz, F. and S. Tiemann (to appear). Presupposition projection in online processing. *Journal of Semantics*.

Sedivy, J. C., M. K. Tanenhaus, C. G. Chambers and G. N. Carlson (1999). Achieving incremental semantic interpretation through contextual representation. *Cognition 71*, 109–147.

Shan, C-C. and C. Barker (2006). Explaining crossover and superiority as left-to-right evaluation. *Linguistics and Philosophy 29*, 91–134.

Von Stechow, A. and S. Beck (2015). Events, times and worlds – An LF architecture. In C. Fortmann (Ed.), *Situationsargumente im Nominalbereich*, pp. 13–46. Berlin: de Gruyter.

Steedman, M. J. (2000). *The Syntactic Process*. Cambridge, MA: MIT Press.

Steedman, M.J. and J. Baldridge. (2011). Combinatory Categorial Grammar. In R. Borsley and K. Borjars (Eds.) *Non-Transformational Syntax: Formal and Explicit Models of Grammar*. Malden, MA: Wiley-Blackwell.

Tiemann, S. (2014). *The Processing of* wieder *('again') and other Presupposition Triggers*. Ph.D. thesis, University of Tuebingen.

Tiemann, S., M. Schmid, N. Bade, B. Rolke, I. Hertrich, H. Ackermann, J. Knapp and S. Beck (2011). Psycholinguistic evidence for presuppositions: On-line and off-line data. In I. Reich, E. Horch and D. Pauly (Eds.), *Proceedings of Sinn und Bedeutung 15*, 581–597.

Varvoutis, J. and M. Hackl (2006). Parsing quantifiers in object position. Presented at CUNY 2006.

Zimmermann, E. and W. Sternefeld (2013). *Introduction to Semantics – An Essential Guide to the Composition of Meaning.* Berlin: de Gruyter Mouton.