Compositionality

Part 1: Basic ideas and definitions

Dag Westerståhl
Stockholm University

EASLLC 2012
Southwest University
Chongqing, August 27–31, 2012

Outline of the course

Aim: To present the idea of compositionality and some of the discussion around it (for natural but also formal languages).

I start with some historical background, and then present a precise framework—or rather two related but slightly different frameworks, both due to Wilfrid Hodges—in which the idea of compositionality, and related ideas, can be expressed and studied.

Then I address some claims that compositionality is trivial or empty or simply false, and see what can be learnt from them.

At some point I will try to connect to Pauline Jacobson's notion of direct compositionality and things she does in her course.

Two important topics remain:

- how sentence meaning determines the meaning of words (Frege's Context Principle);
- how compositionality can deal with context dependence, with an application to the semantics of quotation.

Already posted slides

...

Semantics

Definition

A semantics is a partial function μ from $GT_{\mathbf{E}}$ (alternatively, from $\mathbb{E}$) to some set M of values ('meanings').

Thus, $dom(\mu) \subseteq GT_{\mathbf{E}}$: only grammatical terms are meaningful.

One may still want $dom(\mu)$ to be a proper subset of GT:

- Colorless green ideas ...
- μ is a semantics for a fragment.

Might one also want some ungrammatical terms or expressions to be meaningful?

(1) John don't sing.

But note that (a) these are often perceived as ungrammatical, (b) so we need the grammatical expressions anyway.

Synonymy

A semantics μ induces a sameness of meaning relation:

$t \equiv_\mu u$ iff $\mu(t), \mu(u)$ are both defined and $\mu(t) = \mu(u)$

$\equiv_\mu$ is a partial equivalence relation on $GT_{\mathbf{E}}$ (or on $\mathbb{E}$) — a synonymy — i.e. it is symmetric, transitive, and reflexive on $dom(\mu)$.

Any such synonymy $\equiv$ in turn induces a semantics $\mu_\equiv$:

$\mu_\equiv(t) = [t] =$ the equivalence class of t ($= \emptyset$ if $t \notin dom(\equiv)$)

$\mu_\equiv$ is the equivalence class semantics of $\equiv$.

Exercise. Show that the buck stops here, i.e., that

if $\equiv$ is a synonymy, then $\equiv_{\mu_\equiv} \; = \; \equiv$.

Compositionality, functional version

(i) A semantics μ for $GT_{\mathbf{E}}$, given by a grammar $(E, A, \alpha)_{\overline{\alpha} \in \Sigma}$, is compositional iff for each $\overline{\alpha} \in \Sigma$ there is an operation r_α such that whenever $\mu(\overline{\alpha}(t_1, \ldots, t_n))$ is defined,

$$\mu(\overline{\alpha}(t_1, \ldots, t_n)) = r_\alpha(\mu(t_1), \ldots, \mu(t_n))$$

(ii) A semantics μ for $\mathbb{E}$, relative to a constituent structure $(\mathbb{E}, \mathbb{F})$, is compositional iff for each $F \in \mathbb{F}$ there is an operation s_F such that whenever $\mu(F(e_1, \ldots, e_n))$ is defined,

$$\mu(F(e_1, \ldots, e_n)) = s_F(\mu(e_1), \ldots, \mu(e_n))$$

The idea is the same in both cases: the value of a complex expression is determined by the values of its parts and the mode of composition.

In the term algebra, we look at the immediate constituents. This notion is not in general available in constituent structures, so we need a separate condition for each frame.

NB Both versions require that the domain of μ is closed under constituents.

Compositionality, substitutional version

(i) A partial equivalence relation $\equiv$ on GT is compositional iff for each term $s[t_1, \ldots, t_n]$, if $t_i \equiv u_i$ for $1 \leq i \leq n$, and $s[t_1, \ldots, t_n], s[u_1, \ldots, u_n]$ are both in the domain of $\equiv$, then

$$s[t_1, \ldots, t_n] \equiv s[u_1, \ldots, u_n]$$

(ii) A partial equivalence relation $\equiv$ on $\mathbb{E}$ is compositional iff for each expression $F(e_1, \ldots, e_n)$, if $e_i \equiv f_i$ for $1 \leq i \leq n$, and $F(e_1, \ldots, e_n), F(f_1, \ldots, f_n)$ are both in the domain of $\equiv$, then

$$F(e_1, \ldots, e_n) \equiv F(f_1, \ldots, f_n)$$

In (i), $t_1, \ldots, t_n$ are disjoint, not necessarily immediate, subterm occurrences in the complex term.

Constituent structures can model expressions with overlapping constituents, which allows a simpler formulation, and makes the second claim of the next fact trivial.

The first claim is also straightforward, but requires an argument by induction over the complexity of terms.

One notion of compositionality

Fact

If dom(μ) is closed under constituents then, in the syntactic algebra setting as well as in the constituent structure setting, μ is compositional iff $\equiv_\mu$ is compositional.

This is important, since it shows that there is just one basic notion of compositionality.

Thus, for any syntactic theory that satisfies the minimal requirement of having a reasonable notion of constituency, and for any proposed assignment of meanings to its expressions, the question whether this assignment is compositional or not has a definite answer.

Moreover, the only way of showing that such an assignment is not compositional, is to exhibit a complex expression that changes its meaning when some of its constituents are replaced by synonymous ones.

This basic notion has been called various things: local compositionality, strong compositionality, homomorphism compositionality, but it is the same concept.

Note on closure under constituents

That the domain of μ is closed under constituents means that every constituent of a meaningful expression has meaning.

That is one part of Pauline Jacobson's idea of direct compositionality.

But note also that the substitution version of compositionality does not have this requirement. It says: if meaningful constituents are replaced by synonymous one, then ...

So we do have a notion of compositionality for semantics where meaningful expressions have some constituents without meaning.

There are examples of such languages in logic: the original semantics for Hintikka's IF logic only gives meanings to sentences, not to formulas with free variables.

We will come back to this example. It is interesting because Hintikka claimed that you could not extend it to a compositional semantics for formulas.

But Hodges has shown that you can. This is a very general fact that we will come back to in Lecture 3 (or 4).

Note on 'triple grammars'

This note concerns the idea to let a grammar do syntax and semantics simultaneously, by generating triples, as in Pauline Jacobson's course.

Let an interpreted language L be a set of triples

$$\langle e, X, m \rangle$$

where $e \in E$ is a string, $X \in Cat$ is a category, and $m \in M$ is a meaning.

A grammar **G** for L is a set of partial functions generating L.

This combines syntax and semantics in one, so our earlier format for compositionality is not directly applicable. But it can be adapted.

First note that we no longer need to consider terms in term algebras:

A lexical ambiguity results in two different triples with the same string (and perhaps category), say, $\langle bank, N, m_1 \rangle$ and $\langle bank, N, m_2 \rangle$.

A structural ambiguity lets us derive two different triples with the same string (and perhaps category), say, for

(2) old men and women

Compositional 'triple grammars'

This format is quite powerful. One could imagine grammar rules where the string, the category, and the meaning together determined the result.

But that's not how the format is used: the syntax and the semantics run in parallel ('in tandem').

This is reflected in the definition of compositionality for such grammars (given by Markus Kracht): Let π_1, π_2, π_3 be projection functions:

$$\pi_i(x_1, x_2, x_3) = x_i, \quad i = 1, 2, 3$$

Then $\mathbf{G}$ is compositional iff for each function $\alpha \in \mathbf{G}$ there are operations $r_{1,\alpha}, r_{2,\alpha}, r_{3,\alpha}$ s.t.

$$\alpha(\sigma_1, \ldots) \ = \ \langle r_{1,\alpha}(\pi_1(\sigma_1), \ldots), r_{2,\alpha}(\pi_2(\sigma_1), \ldots), r_{3,\alpha}(\pi_3(\sigma_1), \ldots) \rangle$$

That is, each 'dimension' of the triple is compositional.

[The more permissive requirement would be

$$\alpha(\sigma_1, \ldots) \ = \ \langle r_{1,\alpha}(\sigma_1, \ldots), r_{2,\alpha}(\sigma_1, \ldots), r_{3,\alpha}(\sigma_1, \ldots) \rangle.]$$

'Triple grammars' vs. ordinary grammars + semantics

Each 'triple grammar' $\mathbf{G}$ for L corresponds to a grammar

$$\mathbf{E}_G = (L, A, \mathbf{G})$$

in our sense, where A is the set of atomic triples in L. Then one can show that, under a small condition,*

(3) $\mathbf{G}$ *is compositional (in Kracht's sense) iff each of* π_1, π_2, π_3*, seen as an assignment of values to triples, is compositional in the usual sense.*

Conversely, given a grammar $\mathbf{E}$ and a semantics μ for $GT_{\mathbf{E}}$ in our sense, one can construct a corresponding 'triple grammar' $\mathbf{G}_E$ such that, if $dom(\mu)$ is closed under subterms and has the Husserl property (see Lecture 3), then μ is compositional (in our sense) iff $\mathbf{G}_E$ is compositional (in Kracht's sense).

*The condition is: If $\alpha(\sigma_1, \ldots, \sigma_n)$ is defined, and $\pi_i(\sigma_j) = \pi_i(\tau_j)$, all i, j, then $\alpha(\tau, \ldots, \tau_n)$ is defined. This always holds if $\mathbf{G}$ is compositional.

Back to ordinary compositionality: a misguided objection

To see that we understand things, let us go through a putative counter-example (that sometimes occurs in the literature).

Adjective-noun compounds. Sometimes the extension of an [Adj N] phrase is the intersection of the extension of the Adj and the extension of the N:

(4) brown cow, male cat, prime number, Swedish woman

But sometimes it is not:

(5) white wine, red hair, yellow fever

Does this show that Adj-N combinations in English are not compositional?

– Explain why it doesn't!

Does it show that e.g. white means something else in white wine than it does in white paper?

– Explain why this has nothing to do with compositionality!

– What kind of counter-example would be needed to show that Adj-N combinations in English are not compositional?

Weak compositionality

The following is a weaker notion of compositionality (that one sometimes sees in the literature): The meaning of a complex expression is determined by the meanings of its atomic constituents and its syntactic structure.

A precise formulation (for syntactic algebras) is obtained by restricting the subterms mentioned in the substitutional version of compositionality to atomic ones.

This is the kind of compositionality that would apply if we interpreted

(6) Marie [likes Kelly]

without attention to structure (as in predicate logic), by

$$like(m, k)$$

as discussed in Pauline Jacobson's course. Then we would have to interpret

(7) Marie ate fish and ran

by a new rule, etc.

We would need one rule for each syntactic structure.

Weak compositionality, cont.

So weak compositionality can hardly explain how one can figure out the meanings of complex expressions, since there usually are infinitely many syntactic structures, that one would have to learn one by one.

But note also that the same can be said for the constituent structure version of ordinary compositionality, which has one semantic operation for every frame.

The constituent structure approach is not designed to model the generative feature of syntax and semantics.

Basic compositionality is a weak requirement

It is important to realize that compositionality is a very weak constraint.

First, it doesn't use anything about what meanings, according to μ, are.

It only requires sameness of meaning, i.e. $\equiv_\mu$, to be fixed.

Second, and most importantly, we have the following fact.

Fact

If a semantics μ is one-one, i.e. if each expression (term) has a different meaning, then μ is compositional.

This follows from the substitution version of compositionality, since $\equiv_\mu$ is then the identity relation.

The observation highlights the fact that the word "determine" in the intuitive formulation of compositionality just means 'is a function of': it doesn't mean that one is 'able to figure out' the meaning of complex expressions from the meanings of their parts.

For that, one must impose extra requirements, notably that the meaning operations are computable in some suitable sense.

Weak but not trivial

No doubt the computability aspect is also part of the intuitive motivation for compositionality.

Still, it makes sense to isolate a core meaning of 'compositionality', as in the above definitions.

It is the requirement expressed by standard formulations like Barbara Partee's:

> *The meaning of an expression is a function of the meanings of its parts and the way they are syntactically combined.*

And experience shows that in practice, the semantic operations used are easily computable.

So it is a weak requirement.

But weak is not the same as trivial or empty.

PART 2: Trivial, empty,false?

It has been claimed that, for mathematical as well as philosophical reasons, compositionality is a trivial or empty property.

The mathematical claim is that every semantics can be made compositional with trivial manipulations.

The philosophical claim is rather that compositionality adds nothing to an account of linguistic meaning.

We will look at one typical example of each kind.

The conclusion is that none of these are convincing.

But there are (much) more interesting questions, about actual attempts to obtain compositionality when one is faced with a purported counter-example, and we will also look at some of these.

Mathematical triviality: Zadrozny

Zadrozny (1994) shows that given any semantics μ one can find another semantics μ^* with the same domain such that (a) μ^* is compositional; (b) μ can be recovered from μ^*.

(He also shows that, if one allows non-wellfounded sets in the meta-theory, the only composition operation needed is function application; that is another story.)

In fact, his semantics μ^* is one-one, so, as we just saw, its compositionality is indeed trivial.

But the claim that a semantics satisfying (a) and (b) exists is itself trivial: just let, for each $e \in dom(\mu)$,

$$\mu'(e) = \langle \mu(e), e \rangle$$

Then μ' is compositional (since it is one-one), and μ is easily recovered from μ': $\mu(e)$ is the first element of the pair $\mu'(e)$.

Clearly, this says nothing at all about the original semantics μ.

Mathematical triviality, cont.

This illustrates that if you let meanings be too fine-grained, e.g. if you include the expressions themselves, or their sounds, or the associations they evoke, then compositionality becomes trivial.

And then there seems to be little chance of systematically ‘figuring out’ the meanings of complex expressions from the meanings of their parts.

Conclusion: Compositionality is only interesting when there are many non-trivial synonymies.

“But isn’t it true that meanings are very fine-grained?”

— There are many notions of meaning (and synonymy). Choosing a coarser-grained one is not just a way to avoid complications.

On the contrary, abstracting away from certain details is often necessary to uncover underlying regularities.

This is done in the sciences all the time.

Philosophical triviality: Horwich

In "Deflating compositionality" (2005), Paul Horwich accepts compositionality but gives it no role in explaining the meaning of complex sentences. The idea is that the meanings of words (atoms) and the rules of syntax provide all the information needed:

The argument:

(a) That x means DOGS BARK consists in x resulting from putting together words whose meanings are DOGS and BARK, in that order, into a schema whose meaning is NS V.
(b) "dogs" means DOGS, "bark" means BARK, and "ns v" means NS V.
(c) "dogs bark" results from putting "dogs" and "bark", in that order, into the schema "ns v".
(d) Hence, "dogs bark" means DOGS BARK.

Horwich's conclusion is that compositionality holds as a direct consequence of what it is for a complex expression to have meaning.

Horwich, cont.

This is a philosophical argument, so harder to evaluate ...

I think the possible attraction of the argument is because the example is so simple that any meaning explanation will appear trivial.

But if you look closer, this impression disappears:

First, is the idea is that no other string of words can mean DOGS BARK? Then we have trivial compositionality because of a one-one meaning assignment. But that is not the reason offered.

Second, this is unclear is because we are not told what the meanings of DOGS or BARK are, or about the operation of concatenating them.

Is the notation a shorthand for an operation of combining the meaning of a bare plural with the meaning of an intransitive verb? Compositionality says that such an operation exists.

But the order of explanation is the reverse: after we have specified such an operation, we can conclude that compositionality holds

Horwich, cont.

Third, the example may look trivial but the compositionality claim still has content. It says that other sentences of the same form, e.g. "Cats meow", should be analyzed with the same semantic operation.

If you find that too trivial, you have an argument for compositionality!

Finally, the appearance of triviality fades with more complex cases:

(8) Everyone knows someone.

It is easy to specify schemas generating (8).

It is less trivial — but by now, well-known — to specify corresponding semantic operations that yield the intended meaning of (8).

To say that the meaning of (8) is EVERYONE KNOWS SOMEONE is completely uninformative until the semantic operations are specified.

To say that language requires such operations to exist is to presuppose compositionality. But then it is anything but trivial.

Triviality, end

To me it seems more fruitful to regard it as an hypothesis about natural language meaning.

After all, it is easy to make up non-compositional languages.

So it is a substantial hypothesis, to which empirical evidence is relevant.

It may look 'deflated' with very simple examples, but it really isn't.

It is not only empirical but also theoretical, since it depends on how we construe the syntax-semantics interface.

So it depends on both facts and theory.

So do most scientific hypotheses.

That doesn't make them trivial or empty.

Real compositionality issues: Frege

The problems with compositionality discussed so far are not real problems, in my opinion. So let's discuss some real ones.

The first, and one of the most famous, compositionality problems was introduced by Frege. It is still a live issue in semantics: the treatment of belief sentences, or attitude reports.

Frege had two kinds of semantic values: Sinn (sense) and Bedeutung (reference).

The Bedutung (*bed*) of a name is an object and the Bedutung of a sentence is a truth value. The following differ in 'cognitive value':

(9) Hesperus is Phosphorus (factual)
(10) Hesperus is Hesperus (logically true)

Sinn (*sinn*, 'the mode of presentation of the referent') explains the difference. The Sinn of a sentence is a thought. The Sinn of a name is a part of a thought (roughly, a description of the referent).

Frege and belief sentences

Frege stated explicitly that the function *bed* is compositional. (He may have indicated that *sinn* was too.)

Also, Sinn determines Bedeutung. So there is a function *det* s.t.

$$det(sinn(e)) = bed(e)$$

Now Frege observed that the following can differ in Bedeutung:

(11) Mary believes that Hesperus is Phosphorus
(12) Mary believes that Hesperus is Hesperus

Let the form of these sentences be

$$Bel(a, S)$$

In the example we have $bed(S) = bed(S')$ even though

$$bed(Bel(a, S)) \neq bed(Bel(a, S'))$$

This is a counter-example to the compositionality of *bed*.

What to do?

Frege and belief sentences, cont.

It is debatable exactly how Frege ‘solved’ this problem.

What he explicitly says is: In a belief context, a sentence doesn’t have its ordinary Bedeutung. Instead, it has ‘indirect Bedeutung’, which in fact is its ordinary Sinn.

One interpretation is this: *bed* is in fact not compositional. But Sinn is: we still have

$$sinn(Bel(a, S)) = r_{Bel}(sinn(a), sinn(S))$$

And since Sinn determines Bedeutung (via the function *det*), we still have a treatment of the Bedeutung of belief sentences.

But note that Bedeutung is not determined from Sinn in a compositional way. I.e. the function

$$det(sinn(e)) = bed(e)$$

is not compositional.

Frege and belief sentences, end

Another way to interpret Frege (more true to what he writes, I think), is that we need both functions *sinn* and *bed* in a compositional account.

So when we compute *bed* of a belief sentence, we need to use the other function, *sinn*, for one of the parts.

This is an extension of the notion of compositionality, that Peter Pagin and I call general compositionality.

I come back to this when we talk about compositionality and context.

Note that both accounts need to do something more in order to deal with iterated belief sentences:

(13) Henry believes that Mary believes that Hesperus is Phosphorus.

In any case, this is a good example of a real problem about compositionality, and suggestions for real (non-trivial) solutions.

A simple linguistic example

The following example is from Higginbotham (2007). Consider deontic

(14) Mary may not leave. (Contrast: Mary may not know John.)

According to H., the VP has the form [may [not leave]], and not leave denotes the complement of *leave*, and may indicates permission.

So a compositional analysis gives the meaning that Mary is permitted not to leave. That's not what (14) means. It means that Mary is forbidden to leave.

What to do? Several things are possible, e.g.

(A) Say (14) is ambiguous; the VP also has the analysis [[may not] leave]. Cf.

(15) Mary will not leave.

Does (15) describe something Mary will do, or something she will not do?

(B) Or say the LF is [not [may leave]], and try to explain the difference wrt

(16) Mary must not leave.

 in terms of so-called polarity items.

(C) More elaborate semantics of deontic modals.