

A Best-Fit Approach to Productive Omission of Arguments

EVA H. MOK and JOHN BRYANT

University of California, Berkeley and International Computer Science Institute

0. Introduction

Construction grammars (Fillmore, Kay, and O'Connor 1988; Goldberg 1995; Kay and Fillmore 1999) provide a framework for describing grammaticality in terms of both form and meaning constraints. On the meaning side, cognitive linguists have contributed many great insights into the conceptual structures necessary for understanding language, but challenges remain in developing a precise, formal representation of construction grammar that supports both detailed linguistic analysis and computational use.

One such challenge arises from pro-drop languages, which allow productive argument omission where arguments can be omitted without any markedness in use. Mandarin Chinese is one such example where both the subject and the object can be freely omitted. This phenomenon is commonplace in conversations, but has also been shown in written text (Li 2004; Yeh and Chen 2004).

In actual use, different arguments of a construction tend to be omitted at different rates. Additionally, our preliminary data in Mandarin suggests that even semantically related arguments are omitted differentially in different constructions. This data is difficult to properly account for, both linguistically and computationally, by a general principle of omission that is applied uniformly to all constructions.

For a theory of grammar to accurately account for such data, the grammar must incorporate construction-specific rates of argument omission. Unfortunately, specifying such parameters in most theories of construction grammar requires enumeration of all possible argument combinations as separate constructions within the grammar. This leads to not only an explosion of constructions but also a loss of generality. However, situating the grammar within a *best-fit* processing model allows us to avoid such pitfalls.

Intuitively, a best-fit processing model finds the interpretation of an utterance that fits the utterance best. More precisely, the model incrementally builds up a set of competing interpretations for an utterance, attending only to those interpretations that are most likely given the probabilistic syntactic, contextual, and

semantic constraints. Upon completion of the utterance, the model returns the interpretation (analysis) of a sentence that is most likely.

In this paper, we show that a best-fit sentence analysis model simplifies the representation of argument structure constructions in languages that freely omit arguments. Instead of requiring enumeration of every combination of legal arguments or appealing to some imprecise, underspecified omission principles, we incorporate argument omission rates into the argument structure constructions, and use these rates within the best-fit model. This allows the model to accurately reflect the argument omission data and preserves generality within the grammar.

1. Productive Omission of Arguments in Mandarin Chinese

To illustrate the nature of productive argument omission and its implication for a grammar formalism, we rely on examples of the ditransitive construction in Mandarin Chinese, which follows this basic pattern: it has a verb of transfer, a subject (the giver) and two objects (the recipient and the theme).

Subj	Verb	Obj1	Obj2
↓	↓	↓	↓
Giver	Transfer	Recipient	Theme

We analyzed the basic version of this ditransitive construction which uses the verb *gei3* ('to give') and show below six utterances taken from the Tardiff Beijing Corpus (Tardiff 1993, 1996) in CHILDES (MacWhinney 2000), a corpus of caregiver-child interaction. These are used by the same caregiver in very similar situations, and some are essentially rephrasing of the prior utterance. Although (1) exhibits the construction without any argument dropping, all other argument combinations are also attested in the data, five of which are shown below. Brackets in the English gloss indicate the intended referent in the case of an omission.

- (1) ma1+ma gei3 ni3 zhei4+ge. 'Mother give you this (a toy).'
 mother give 2PS this+CLS
- (2) gei3 wo3 le ao. '[You] give me [the pen], alright?'
 give 1PS CRS EMP
- (3) ni3 gei3 yi2. 'You give auntie [the peach].'
 2PS give auntie
- (4) gei3 a1+yi2 yi2+ge qu4 ao. '[You] Give auntie one (peach).'
 give auntie 1+CLS away EMP
- (5) ao ni3 gei3 ya. 'Oh you give [auntie] [the peach].'
 EMP 2PS give EMP

- (6) *gei3*. ' [I] give [you] [some peach].'
 give

Across 97 *gei3* phrases from two speakers, XiXi's and HaoYu's mothers, only 6.1% of the phrases have all arguments present, as is exemplified by (1). On the other hand, phrases with all arguments omitted, such as Example (6), account for 30.6% of the data, excluding repetitions of a preceding child utterance. As shown in Figure 1, the giver is omitted in 78.4% of the phrases, the recipient 41.2%, and the theme 66.0%.




Figure 1. Percent omitted out of 97 *gei3* phrases, from two speakers combined.

Not only are arguments of the ditransitive construction omitted at different rates, but the patterns of omission also change with constructions. For contrast, we compared the basic ditransitive construction with a basic caused-motion construction, which has a causer and a mover argument. Taking 52 additional phrases from XiXi's mother with the verb *ge1* ('to put / to place'), we compared the rates of omission of the subject (agent) and the direct object (theme). While the subject (agent) is omitted at similar rates (80.0% and 73.1%, for the ditransitive and caused-motion phrases respectively), the direct object (theme) is omitted at significantly different rates (58.0% and 78.9%, respectively).

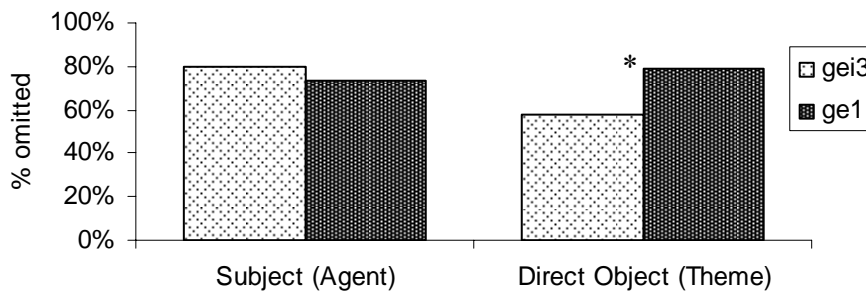


Figure 2. The omission rates of subject and direct object for 50 *gei3* phrases (ditransitive) and 52 *ge1* phrases (caused motion) from the same speaker.

Taken together, this data suggests that varying patterns of omission exist for the same construction and that the relative frequency of these patterns vary

between constructions. Although a larger corpus study is required to warrant a strong conclusion, argument omission appears to be a construction-specific phenomenon even in a language which allows productive omission of arguments. The goal of capturing these construction-specific regularities without resorting to the enumeration of constructions motivates the model we present in the remainder of this paper.

2. Embodied Construction Grammar (ECG)

In order to support both detailed linguistic analysis and computational use within the construction grammar framework, we rely on Embodied Construction Grammar (ECG) (Bergen and Chang 2005), a unification-based and computationally precise construction grammar formalism, for our work. Meaning is represented in ECG using *embodied schemas*, such as image schemas, frames, and action schemas; *constructions* link forms to these embodied schemas. The ECG schema and construction formalisms include mechanisms for expressing type constraints, ordering constraints, identification (unification) constraints, constituency and dependency relations.

Figure 3 shows how the ditransitive construction is represented in ECG, with ECG keywords in bold. The construction, named Active-Ditransitive-VP, has three **constituents**, which are a Verb and two NPs represented by the local names *v*, *obj1*, and *obj2* respectively. The **subcase** relation, which states that this Active-Ditransitive-VP is a subtype of Active-VP, is important in preserving compositionality between constructions: it enables a separate sentential or clausal construction to compose the Active-VP with a subject NP.

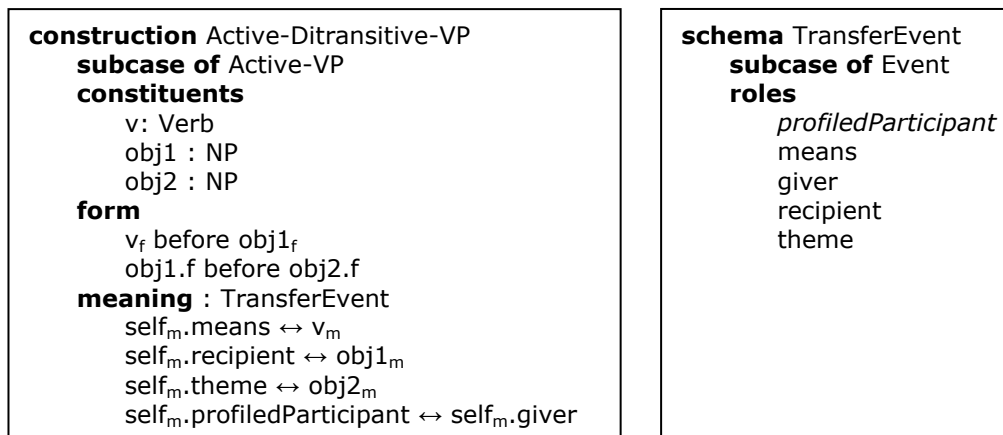


Figure 3. The active ditransitive VP construction and the transfer event

The ordering constraints among these constituents are expressed in the **form** block, which states that the constituent *v* must appear in the sentence before *obj1*, and *obj1* before *obj2*. The meaning of this construction is a TransferEvent, whose **roles** are shown on the right in Figure 3. The Active-Ditransitive-VP construction

links its constituents to roles in this TransferEvent in the **meaning** block through identification constraints denoted by \leftrightarrow . The means of transfer is identified with the meaning of the verb; the recipient is identified with the meaning of the first noun phrase (obj1), and the theme is identified with the meaning of the second noun phrase (obj2).

The composition of constructions in ECG is illustrated by Figure 4 with a simplified clausal construction combining an NP with a VP, of which Active-Ditransitive-VP is a subtype. The meaning of the resulting Subject-VP construction is an Event which is identified with the meaning of the verb phrase (vp); the profiledParticipant of this event is identified with the subject noun phrase (subj).

construction Subject-VP
subcase of S
constituents
subj: NP
vp : VP
form
subj _f before v _f
meaning : Event
self _m \leftrightarrow vp _m
self _m .profiledParticipant \leftrightarrow subj _m

Figure 4. Composing an NP and a VP to form a Subject-VP construction in ECG.

3. Best-Fit Sentence Analysis

While Embodied Construction Grammar is consistent with prevailing notions of construction grammar, what allows us to avoid enumerating constructions to handle the argument omission phenomenon is the leverage we get from our processing model. We have built a psychologically plausible construction-based sentence analyzer (Bryant 2004), called the analyzer for brevity. Given an ECG grammar, the analyzer takes a sentence in context and generates a meaning representation built out of embodied schemas, frames and the bindings between their roles. We refer to this interconnected set of schemas and frames as the *semantic specification* or *semspec*.

The analyzer is designed to be compatible with compelling psycholinguistic evidence (Narayanan and Jurafsky 2001) that people employ several quantitative measures in deciding among competing analyses. Once seeded with a grammar written in ECG, the analyzer incrementally processes an input sentence, integrating the situational and discourse context. Going word by word, the analyzer builds up competing interpretations by tracking syntactic, contextual, and semantic bindings, maintaining only the most likely interpretations that satisfy the constraints specified by the grammar. In the end, the analysis that not only satisfies all the constraints but also has the best fit is selected as the best interpretation of the sentence.

3.1 Defining Best-Fit

The best-fit score of an analysis given an utterance is a probabilistic metric which combines syntactic, contextual, and semantic factors. The syntactic component incorporates the combination of chosen constructions, their constituency relations and the argument omission probabilities. The argument omission probabilities are approximated by the omission rates obtained from corpus data, as we have done for the basic Active-Ditransitive-VP construction with the *gei3* ('to give') sentences and the basic Active-Caused-Motion-VP construction (not shown) with the *ge1* ('to put/ to place') sentences in Section 1.

The contextual component scores how well the referring expressions are resolved to items from context and how easily the omitted arguments are recovered from context. In some cases the omitted antecedent is present in the preceding discourse or the immediate situational context, while in other cases the omitted antecedent is presumed (e.g. in English, *Did you eat?* presumes a generic meal as the omitted argument to *eat*). This phenomenon is well captured by the distinction made in FrameNet¹ between definite and indefinite null-instantiated arguments, in that some verbs, such as *eat*, only require indefinite null-instantiation of its omitted arguments (Fillmore, Johnson, and Petruck 2003). Our best-fit scoring of contextual fit takes into account the definiteness requirements of different verbs in deciding whether or not to penalize for an omitted argument which is not available in context.

Finally, the semantic component scores the semantic bindings essentially by evaluating how well the frame roles are being filled. For example, when the Active-Ditransitive-VP construction is used in an analysis, the meaning denoted by its constituents (e.g. a person or a pen) are identified with various roles of the TransferEvent frame. In this case, the semantic fit reflects whether each role of the TransferEvent frame is filled with likely fillers.

With this best-fit scoring of analyses, argument omission is accommodated straightforwardly by the model by selecting the best analysis in all three respects: it has to fit well with the sentence constructionally, the omitted arguments have to be recoverable from context, and the role fillers of the semantic frames have to establish coherent meaning, whether they are supplied by language or by context.

3.2 Simulation-Based Understanding

The analysis process is a central piece in the framework of simulation-based understanding in the Neural Theory of Language (NTL) project (Feldman 2006). The simulation hypothesis states that the same neural circuitry employed for carrying out an action is also used for the purpose of understanding sensory input such as seeing an action and other cognitive tasks such as processing language and general reasoning. Within a simulation-based understanding framework, the active process of simulation enables the use of both linguistic input and general knowledge in making inferences that enrich the understanding of a sentence.

¹ <http://framenet.icsi.berkeley.edu>

Upon encountering an utterance, it is the analyzer that initiates the simulation by providing the simulator with a semspec. The overall architecture of the analyzer is shown in Figure 5.

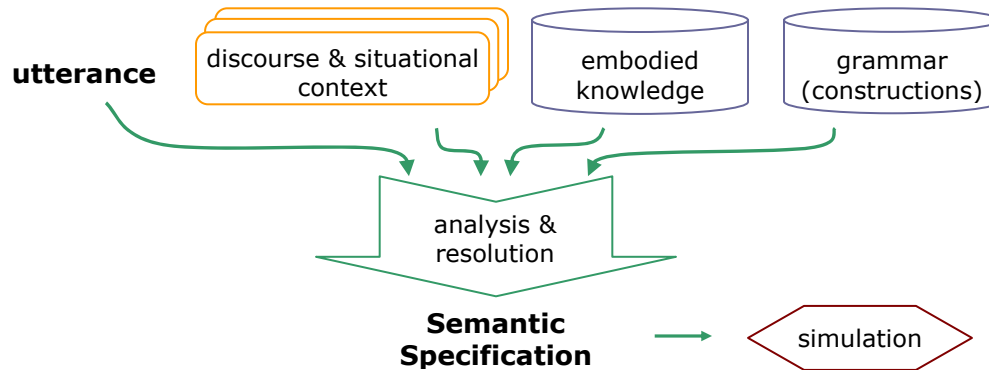


Figure 5. The analysis process finds the best semspec for a sentence in context given a grammar and embodied knowledge.

4. Example

We will illustrate our best-fit analysis model by working through the analysis of Example (3), given again below, which omits the theme of the ditransitive construction. There are a number of possible competing analyses of it using either the Active-Ditransitive-VP construction in Figure 3 or other constructions not given in this paper such as the simple transitive construction. For brevity, we will focus on two of the most likely competing analyses that use the ditransitive construction – one with the theme omitted and the other with the recipient omitted. We break down the comparison between of the two analyses into the three components of the best-fit score: syntactic fit, contextual fit, and semantic fit. For simplicity, we will largely forgo the numerical calculations and demonstrate the principles behind the best-fit analysis scoring.

- (3) ni3 gei3 yi2. ‘You give auntie [the peach].’
 2PS give auntie

4.1 Syntactic Fit

To determine whether the omitted-theme analysis or the omitted-recipient analysis is more syntactically likely for the example sentence, we employ the statistics obtained in Section 1. Recalling that across the *gei3* sentences, the giver is omitted in 78.4% of the phrases, the recipient 41.2%, and the theme 66.0%. These statistics are used to estimate the likelihood of different omission patterns showing up in a sentence.

The omitted-theme analysis of Example (3) matches *gei3* (‘give’) with the constituent *v*, *yi2* (‘auntie’) with the constituent *obj1* (i.e. recipient), and treats constituent *obj2* (i.e. theme) as omitted. The likelihood of an omitted *obj2* in the

ditransitive construction, as we have seen, is 0.66. The omitted-recipient analysis, on the other hand, matches *gei3* as the *v*, treats *obj1* (i.e. recipient) as omitted, and matches *yi2* with *obj2* (i.e. theme). The likelihood of an omitted *obj1* in the ditransitive construction is a lower 0.41². By the syntactic fit measure, then, the omitted-theme analysis has a higher score.

4.2 Contextual Fit

Intuitively, the best analysis for a sentence should make sense with respect to its context, and one aspect of this contextual match is that referring expressions in the sentence should be as grounded in the context as possible. This means that not only should overt arguments be resolved to elements in context, but omitted arguments should also be recovered from context as well. To calculate the contextual fit of a sentence, the analyzer must be supplied with both the discourse and situational contexts of each sentence.

In Example (3), the speaker (the mother) is directing the child to give a piece of peach to auntie (the investigator who did the recording). The three of them are the only people in the room, but among the child's surrounding is the peach that she may well be already holding in her hands and a table. While both the omitted-theme analysis and the omitted-recipient analysis contain the referring expression *yi2* ('auntie') which resolves well to the investigator in context, they differ in how well the omitted argument is recovered.

For the omitted-theme analysis, the peach is an obvious choice for the theme both because of its immediacy in the situation and its feasibility as a theme. As a result, this analysis thus scores rather high in its contextual fit. The omitted-recipient analysis, however, is ambiguous as to whether the child, the mother or the investigator should serve as the recipient. This ambiguity thus leads to a lower score for the omitted-recipient analysis.

4.3 Semantic Fit

The semantic component of the best-fit scoring reflects the semantic coherence of each analysis based on the likelihood of the frame role fillers. The two competing analyses of Example (3) both identify the addressee (the child) with the giver of the TransferEvent frame, but they vary with respect to the fillers for the theme and recipient roles.

The omitted-theme analysis picks the investigator as the recipient and recovers the peach from context as the theme. By using resources such as FrameNet, our best-fit model is able to determine that a person is a good recipient, and a peach (or physical object) is a good theme. In the omitted-recipient analysis, on

² In the model, the likelihood of each analysis is calculated by taking into account the presence versus omission probabilities of all three arguments – the giver, the recipient and the theme. This calculation shows that the omitted-theme analysis is in fact almost three times as likely as the omitted-recipient analysis:

The omitted-theme analysis has a likelihood of $(1-0.784)*(1-0.412)*0.66 = 0.08$

The omitted-recipient analysis has a likelihood of $(1-0.784)*0.412*(1-0.66) = 0.03$

the other hand, the investigator is constructionally chosen as the filler for the theme, and either the mother or the investigator is chosen from context as the recipient. Either of them is a good recipient, but a person is an unlikely theme. As a result, the omitted-recipient analysis is assigned a lower semantic fit score than the omitted-theme analysis.

4.4 Combining the Scores for the Best Fit

The overall score of each analysis is obtained by combining all three component scores. Between the two competing analyses that we have considered here, the omitted-theme analysis is the obvious choice due to its higher scores in all three of the syntactic, contextual, and semantic components. Often, however, the choice among competing analyses is not so obvious. One analysis may have the best semantic fit, and yet a different analysis may have the best contextual fit. One major strength of our probabilistic best-fit scoring mechanism is in allowing the analyzer to weigh these different factors in determining the best analysis.

5. Implications and Conclusion

Our running example of the ditransitive constructions in Mandarin Chinese demonstrates how a difficulty in grammar representation involving argument omission is resolved by situating the grammar within a best-fit processing model. Our data is particularly relevant because it is suggestive of construction-specific variations in omission patterns, which are difficult for traditional theories of grammar to represent precisely and elegantly. We thus put forth a view of grammar that sees the grammar as synergistic with process so that the burden of describing linguistic phenomena is shared by both.

Fundamental to the analysis process is the best-fit scoring mechanism that determines the appropriateness of an analysis in terms of syntax, the context, and semantics. With a minimal addition to the grammar representation, i.e. the argument omission probabilities for each construction, we let the analysis process determine the best analysis for each sentence and avoid having to enumerate in the grammar all the possible omission patterns and their individual contextual constraints.

In addition to using cross-linguistic data to further validate our model of best-fit analysis, two directions in our current work are applying best-fit analysis to our model of language understanding and to our model of grammar learning in pro-drop languages. With our current efforts we hope to better understand how the argument omission probabilities interact with compositionality, and how they can be learned.

References

- Bergen, Benjamin. K., and Nancy Chang. 2005. Embodied Construction Grammar in simulation-based language understanding. In *Construction Grammars: Cognitive Groundings and Theoretical Extensions*, edited by J.-O. Östman and M. Fried. Philadelphia, PA: John Benjamins.
- Bryant, John. 2004. Recovering coherent interpretations using semantic integration of partial parses. Paper read at 3rd workshop on Robust Methods in Analysis of Natural Language Data, at Geneva, Switzerland.
- Feldman, Jerome A. 2006. *From Molecule to Metaphor: A Neural Theory of Language*. Cambridge, MA: MIT Press.
- Fillmore, Charles J., Christopher R. Johnson, and Miriam R.L. Petruck. 2003. Background to Framenet. *International Journal of Lexicography* 16 (3):235-250.
- Fillmore, Charles J., Paul Kay, and Mary Catherine O'Connor. 1988. Regularity and idiomacity in grammatical constructions: The case of let alone. *Language* 64:501-538.
- Goldberg, Adele E. 1995. *Constructions: A construction grammar approach to argument structure*. Chicago: University of Chicago Press.
- Kay, Paul, and Charles Fillmore. 1999. Grammatical constructions and linguistic generalizations: the What's X doing Y? construction. *Language* 75 (1):1-33.
- Li, Wendan. 2004. Topic chains in Chinese discourse. *Discourse Processes* 27 (1):25-45.
- MacWhinney, Brian. 2000. *The CHILDES project: Tools for analyzing talk*. Third Edition. Mahwah, NJ: Lawrence Erlbaum Associates.
- Narayanan, Srini, and Daniel Jurafsky. 2001. A Bayesian Model Predicts Human Parse Preference and Reading Times in Sentence Processing. Paper read at Neural Information Processing Systems, at Vancouver, BC.
- Tardif, Twila. 1993. Adult-to-child speech and language acquisition in Mandarin Chinese, Yale University.
- . 1996. Nouns are not always learned before verbs: Evidence from mandarin speakers' early vocabularies. *Developmental Psychology* 32:492-504.
- Yeh, Ching-Long, and Yi-Chun Chen. 2004. Zero anaphora resolution in Chinese with shallow parsing. *Journal of Chinese Language Computing*.