A connectionist model of the retreat from verb argument structure overgeneralization.

Ben Ambridge

University of Liverpool

ESRC International Centre for Language and Communicative Development (LuCiD)

Ryan Blything

University of Manchester

RUNNING HEAD: Connectionist overgeneralization

Address for correspondence: Ben Ambridge, School of Psychology, University of

Liverpool, Eleanor Rathbone Building, Bedford St South, Liverpool, L69 7ZA. Email

Ben.Ambridge@Liverpool.ac.uk

A connectionist model of the retreat from verb argument structure overgeneralization.

A central question in language acquisition is how children build linguistic representations that allow them to generalize verbs from one construction to another (e.g., *The boy gave a present to the girl* → *The boy gave the girl a present*), whilst appropriately constraining those generalizations to avoid non-adultlike errors (e.g. *I said no to her* → **I said her no*). Although a consensus is emerging that learners solve this problem using both statistical and semantics-based learning procedures (e.g., entrenchment, pre-emption and semantic verb class formation), there currently exist few – if any - proposals for a learning model that combines these mechanisms. The present study used a connectionist model to test an account that argues for competition between constructions based on (a) verb-in construction frequency, (b) relevance of constructions for the speaker's intended message and (c) fit between the fine-grained semantic properties of individual verbs and individual constructions. The model was able not only (a) to simulate the overall pattern of overgeneralization-then-retreat, but also (b) to use the semantics of novel verbs to predict their argument structure privileges (just as real learners do) and (c) to predict the pattern of by-verb grammaticality judgments observed in adult studies.

Keywords: Connectionist model, retreat from overgeneralization, no-negative-evidence problem, DO-dative, PO-dative.

A connectionist model of the retreat from verb argument structure overgeneralization.

The centerpiece of Chomsky's (1957) landmark *Syntactic Structures* was the observation that speakers' knowledge of their native language does not consist solely of an inventory of rote-learned sentences; rather speakers acquire abstract generalizations which allow them to comprehend and produce utterances that they have never heard before. The question of how children acquire this quintessentially human ability occupies a core place in language acquisition research.

Braine (1971) pointed out a paradox that lies at the heart of this ability. On the one hand, even 2-3 year old children are adept at producing and understanding sentences in which a verb is used in a sentence-level *verb-argument structure construction* in which it has never appeared in the input. For example, when taught a novel verb in an *intransitive inchoative* construction (e.g., *The sock is weefing*), children are able to produce a *transitive causative* sentence with this verb (e.g., *The mouse is weefing the sock*; e.g., Akhtar & Tomasello, 1997; see Tomasello, 2003; Ambridge & Lieven, 2011, 2015, for reviews, and Gertner, Fisher & Eisengart, 2006; Noble, Rowland & Pine, 2011, for similar findings in comprehension).

On the other hand, children must somehow restrict this productivity to avoid producing ungrammatical utterances. While many verbs can appear in both the intransitive inchoative and transitive causative constructions (e.g., *The ball rolled / The clown rolled the ball*), and children must be able to generalize from one to another, they must also learn that certain verbs are restricted to the former (e.g., *The man laughed / *The clown laughed the man* [where * indicates an ungrammatical utterance]). Indeed, evidence from both diary and elicited production studies

demonstrates that many children pass through a stage in which they produce these

types of overgeneralizations, before "retreating" from error (e.g., Bowerman, 1988;

Brooks & Tomasello, 1999; Brooks, Tomasello, Dodson & Lewis, 1999; Brooks &

Zizak, 2002; Pinker, 1989). Analogous errors for the observed for the dative and

locative alternations are summarized in Table 1.

Table 1. Possible and attested verb argument structure overgeneralization errors.
Attested errors (all from Bowerman, 1988) are shown in **bold**, with the age of the child
(years;months) and a possible grammatical formulation using the alternate
construction. Reproduced by permission of Wiley from Ambridge, B., Pine, J.M.,
Rowland, C.F., Chang, F. & Bidgood, A. (2013). The retreat from overgeneralization
in child language acquisition: Word learning, morphology and verb argument
structure. *Wiley Interdisciplinary Reviews: Cognitive Science, 4*: 47–62.

	Transitive causative alternation	
	(a) (Inchoative) Intransitive	**(b) Transitive (causative)**
Alternating	The ball rolled	The man rolled the ball
(a) only	The man laughed Do you want to see our heads disappear? I don't want any more grapes; I'll cough I (didn't) giggle(d) Will I climb up there? Did it bleed? I always sweat [when I wear it] [They're nice enough that] I wish I had one	*The clown laughed the man ***Do you want to see us disappear our heads? (6;0)** ***I don't want any more grapes; they just cough me (2;8)** ***Don't giggle me (3;0)** ***Will you climb me up there (3;2)** ***Did she bleed it? (3;6)** ***It always sweats me** ***[They're nice] enough to wish me that I had one (5;8).**

	Dative alternation	
	(a) Prepositional-object (PO) dative	**(b) Double-object (DO dative)**
Alternating	The boy gave a present to the girl	The boy gave the girl a present
(a) only	The boy dragged the box to the girl The boy suggested the trip to the girl I said no to her Shall I whisper something to you?	*The boy dragged the girl the box *The boy suggested the girl the trip ***I said her no (3;1)** ***Shall I whisper you something? (7;8)**

	Locative alternation	
	(a) Contents (figure) locative	**(b) Container (ground) locative**
Alternating	The boy sprayed paint onto the statue	The boy sprayed the statue with paint
(a) only	The boy poured water into the cup Mommy, I poured water onto you I don't want it because I spilled orange juice onto it	*The boy poured the cup with water ***Mommy, I poured you [M: You poured me?] Yeah, with water (2;11)** ***I don't want it because I spilled it of orange juice (4;11)**
(b) only	*The boy filled water into the cup ***I'm gonna cover a screen over me (4;5)** ***Can I fill some salt into the bear [-shaped salt shaker]? (5;0)**	The boy filled the cup with water I'm going to cover myself with a screen Can I fill the bear with salt?

The problem of the retreat from overgeneralization has attracted a considerable amount of attention in the literature. Recent studies of all three types of overgeneralization error listed in Table 1 have provided support for three proposals.

Under the *entrenchment hypothesis* (e.g., Braine & Brooks, 1995), repeated presentation of a verb (regardless of sentence type) contributes to an ever-strengthening probabilistic inference that its use in non-attested constructions is not permitted. In support of this hypothesis, many studies have demonstrated a negative correlation between overall verb frequency (regardless of sentence type) and the relative acceptability and production probability of errors with that verb, in judgment and production tasks respectively. (Ambridge, 2013; Ambridge & Brandt, 2012; Ambridge, Pine & Rowland, 2012; Ambridge, Pine, Rowland & Chang, 2012; Ambridge, Pine, Rowland, Freudenthal & Chang, 2014; Ambridge, Pine, Rowland, Jones & Clark, 2009; Ambridge, Pine, Rowland & Young, 2008; Bidgood, Ambridge, Pine & Rowland, 2014; Blything, Ambridge & Lieven, 2014; Brooks, Tomasello, Dodson & Lewis, 1999; Stefanowitsch, 2008; Theakston, 2004; Wonnacott, Newport & Tanenhaus, 2008).

The *pre-emption hypothesis* (e.g., Goldberg, 1995) is similar, but with one important difference. Under entrenchment, a particular error (e.g., **Bart dragged Lisa the box*, where a PO-only verb is used in a DO-dative) is probabilistically blocked by any use of the relevant verb (e.g., *The man dragged the box; The boy dragged his feet; That movie really dragged on* etc.). Under pre-emption, errors of the form **Bart dragged Lisa the box* are probabilistically blocked only by uses that express the same intended message; i.e., PO-dative uses of that verb (e.g., *Marge dragged the package to Homer*). Thus this hypothesis predicts a negative correlation between the

acceptability/production probability of a particular error (e.g., DO-dative uses of *drag*) and the frequency of that verb in the single most nearly synonymous construction (e.g., PO-dative uses of *drag*). Although the two measures tend to be highly correlated, recent studies suggest that – to the extent to which they can be differentiated statistically – pre-emption plays a role above and beyond entrenchment (e.g., Ambridge, 2013; Ambridge et al, 2014; Ambridge, Pine & Rowland, 2012; Boyd & Goldberg, 2011; Brooks & Tomasello, 1999; Brooks & Zizak, 2002; Goldberg, 2011).

The *semantic verb class hypothesis* (Pinker, 1989) argues that learners form classes of verbs that are restricted to particular constructions only. For example, focusing on the dative alternation, verbs of accompanied motion and manner of speaking may appear in the PO-dative (e.g., *Marge pulled the box to Homer; Homer shouted the instructions to Lisa*), but are less than fully acceptable in the DO-dative (e.g., **Marge pulled Homer the box; *Homer shouted Lisa the instructions*). On the other hand, verbs of giving and illocutionary communication may appear in both constructions (*Lisa gave the book to Bart, Lisa gave Bart the book; Lisa showed the answer to Homer, Lisa showed Homer the answer*). Evidence for this hypothesis comes from production and judgment studies showing that if children are taught novel verbs, they use their notional semantic class membership to determine the constructions in which they can and cannot appear (Ambridge et al, 2008, 2009; 2011; Ambridge, Pine, Rowland & Chang, 2012; Bidgood et al, 2014; Brooks & Tomasello, 1999; Gropen, Pinker, Hollander & Goldberg, 1991a, 1991b; Gropen, Pinker, Hollander, Goldberg & Wilson, 1989).

Importantly, these semantic classes are not arbitrary. Rather, a particular class of verbs can appear in a particular construction only if there is sufficient overlap between the semantics of the verbs and the semantics of the construction. For example,

the DO-dative construction is associated with the meaning of "causing to have" (Pinker, 1989). Thus the reason that verbs from the *give* and *show* classes may appear in this construction is that they are consistent with this meaning (in the latter case, the possession transfer is metaphorical; a transfer of information). On this account, the reason that verbs from the *pull* and *shout* classes may not appear in the DO-dative construction is that they are not sufficiently consistent with this "causing to have" meaning. Instead, they are restricted to the PO-dative construction, because they are compatible with the meaning of this construction ("causing to go"). Ambridge et al (2014) found that independent ratings of the extent to which particular verbs were consistent with "causing to have" versus "causing to go" significantly predicted the rated acceptability of that verb in the DO- versus PO-dative construction (see Ambridge & Brandt, 2013; Ambridge, Pine & Rowland, 2012, for an analogous finding for the locative alternation).

In summary, previous studies have found support for the entrenchment, pre-emption and semantic verb class hypotheses. This raises the question of whether it is possible to posit a single learning mechanism that yields all of these effects (indeed, it is debatable whether any of these proposal indeed constitute mechanistic accounts of the retreat from error *per se*). One proposal for such a learning mechanism is the *FIT* account (Ambridge, 2013; Ambridge et al, 2014; Ambridge & Lieven, 2011; Ambridge, Pine & Rowland, 2012; Ambridge, Pine, Rowland & Chang, 2012;). The acronym captures the account's emphasis on the importance of the *Fit* between *Items* and (construction) *Templates* (in this case, with regard to semantics).

The central assumption of the account is that speakers maintain an inventory of argument structure constructions - each acquired by abstracting across concrete tokens of those constructions in the input – which, in production, compete to express the

speaker's desired message (e.g., MacWhinney, 2004). The activation level of each competitor is determined by three factors, illustrated here for the example message "MARGE CAUSED HOMER TO HAVE THE BOX BY PULLING THE BOX TO HOMER":

- ***Verb-in-construction frequency.*** The verb in the message (here *pull*) activates each construction in proportion to the frequency with which it has appeared in that construction in input sentences. This factor yields pre-emption effects because every input occurrence of *pull* in a PO-dative boosts the activation of this construction, at the expense of the DO-dative construction, in production. This factor yields entrenchment effects because every input occurrence of *pull* in *any* other construction (e.g., a simple transitive) boosts the activation of this construction at the expense of the DO-dative.

- ***Relevance***. A "relevant" construction is one that contains a slot for every item in the speaker's message. So, for the present example, both the PO-dative (yielding *Marge pulled the box to Homer*) and the DO-dative (**Marge pulled Homer the box*) are more relevant than, for example, the transitive (*Marge pulled the box*). The notion of relevance captures the intuition of the pre-emption hypothesis that the PO- and DO-dative are better competitors for one another than are other constructions such as the transitive.

- ***Fit***. The third factor is the compatibility (or *fit*) between the semantic properties of each item in the message (e.g., the verb) and the relevant slot in each candidate construction. The semantics of each slot are a frequency-weighted average of the semantics of each item which appeared in that position in the input utterances that gave rise to the construction. This factor is designed to capture the finding that ratings of the extent to which verbs exhibit semantic properties to do with "causing to have" and

"causing to go" predict acceptability in the DO- and PO-dative respectively (Ambridge et al, 2014).

•   A fourth factor, **overall construction frequency**, may also be important. That is, all else being equal, a speaker is more likely to select a higher frequency construction (e.g., an active transitive) than a lower frequency alternative (e.g., the passive). This factor may be necessary to explain, for example, why some alternations attract higher error rates than others. Although construction frequency is indirectly manipulated in the present study, since the simulation involves only two constructions, it is not possible to draw any conclusions regarding the importance of this factor.

Of course, as a verbal model, the FIT account is little more than a redescription of the experimental findings. Thus the aim of the present study is to instantiate the account as a computational model in order to investigate the extent to which it can simulate (a) the overall pattern of generalization to novel verbs, overgeneralization errors and subsequent retreat shown by children and (b) the verb-by-verb pattern of adult acceptability ratings for overgeneralization errors in judgment studies. Before introducing the computational model itself, it is important to consider the respects in which it differs from previous models that also simulate aspects of the retreat from overgeneralization. Note that – although not all include a role for semantics – all nevertheless share a degree of theoretical and implementation overlap with the present model, and thus should be considered complementary models, rather than radically different rivals.

Like the present model, the hierarchical Bayesian model of Perfors, Tenenbaum and Wonnacott (2010) learns item-based links between particular verbs and the PO- and DO-dative constructions. An important difference it that it generalizes to novel

verbs by additionally forming *overhypotheses* regarding the tendency of (a) all verbs

and (b) classes of distributionally similar verbs to occur in both constructions as

opposed to only one, rather than on the basis of semantic similarity (the same is true

for related models based on the notion of Minimum Description Length; e.g.,

Dowman, 2000, submitted; Hsu & Chater, 2010; Hsu, Chater & Vitányi, 2011, 2013;

Onnis, Roberts & Chater, 2002). However, although there is some evidence for the

importance of overhypotheses in artificial grammar learning studies with adults (e.g.,

Perek & Goldberg, in press; Wonnacott et al, 2008;), it remains to be seen whether this

procedure plays a crucial role in children's natural language learning. Indeed, the claim

that children make use of overhypotheses would seem to contradict the large body of

evidence suggesting that their early knowledge of language consists of holophrases and

low-scope formulae, and only later becomes more abstract (e.g., Ambridge & Lieven,

2015; Tomasello, 2003). Although some versions of the Perfors et al (2010) model

include verb-level semantic features, each feature has only three possible values -

corresponding to PO-only, DO-only and alternating verbs – and so does not simulate

fine-grained by-verb semantic effects observed for human participants (Ambridge et al,

2014).

The dual-path model of Chang (2002; see also Chang, Dell & Bock, 2006) is

able to simulate a wide range of language acquisition phenomena, including

generalization of novel verbs into unattested constructions and the retreat from

overgeneralization (including for the dative alternation; see p.638-640). The model

works by learning to sequentially predict the next word in a sentence (using a Simple

Recurrent Network), given a message (e.g., AGENT=*Marge*, ACTION=*drag*,

GOAL=*Homer*, PATIENT=*box*) and construction-level event semantics (e.g.,

CAUSE+MOTION+TRANSFER for the PO-/DO-dative). Due to its sequential nature,

the dual-path model constitutes a lower-level and hence more realistic approximation

of the task facing real language learners than any of the other models outlined here

(including the new model outlined in the present article). An important difference from

the present model is that the dual-path model does not represent verb-level semantics

(only construction-level event-semantics). Also, its use of artificially-generated

datasets means that the model does not make predictions regarding by-verb patterns of

adult grammaticality judgments.

Perhaps the model closest to the present simulation is that of Alishahi and

Stevenson (2008). This model receives input in the form of pairs of a scene and an

utterance (e.g, DRAG$_{CAUSE}$ MARGE$_{AGENT}$ BOX$_{THEME}$ TO HOMER$_{DESTINATION}$ +

Marge dragged the box to Homer), from which it extracts argument-structure frames

(e.g., [argument1] [verb] [argument2] [argument3]). Frames that are sufficiently

similar are collapsed into constructions, using an unsupervised Bayesian clustering

process. This model is similar to the present simulation in its use of verb and

construction semantics, which allows it to show both generalization to novel verbs and

overgeneralization with subsequent retreat, and also in its use of corpus-derived

verb+construction counts. An important difference is that Alishahi and Stevenson's

(2008) simulations did not investigate the relative importance of entrenchment, pre-

emption and verb semantics. Neither did these authors attempt to simulate the by-verb

pattern of adult grammaticality judgments. Indeed, in its present form Alishahi and

Stevenson's model is unlikely to be able to do so, since – like all of the previous

models discussed in this section - it does not represent verb semantics at a sufficiently

fine-grained level (see p.827 for discussion).

The present article introduces a new model that instantiates the key

assumptions of Ambridge and colleagues' verbal FIT account: competition between

constructions based on (a) verb-in construction frequency, (b) relevance of constructions for the speaker's intended message and (c) fit between the fine-grained semantic properties of individual verbs and individual constructions (or, more accurately, their [VERB] slot). The aim is to investigate the ability of the model (a) to explain generalization to novel verbs, overgeneralization error and subsequent retreat, (b) to model the pattern of by-verb grammaticality judgments obtained in adult studies and (c) to elucidate the relative importance of entrenchment, pre-emption and verb semantics, and explore one way in which these factors might be combined into a learning model.

Given that previous regression studies have already demonstrated that entrenchment, pre-emption and verb semantics play a role in the retreat from error – including for the dative constructions (Ambridge, Pine, Rowland & Chang, 2012; Ambridge, Pine, Rowland, Freudenthal & Chang, 2014) – a question arises as to how the present model adds to our understanding of the phenomenon. The main advantage is that, unlike a regression model, the present computational model instantiates – albeit at a relatively high level – a mechanistic account of a possible procedure for LEARNING verb argument structure restrictions. Nobody would argue that children use a single pass of an input corpus to calculate, for each verb, entrenchment, pre-emption and verb-semantic measures (i.e., "meta" or "macro" variables), which they then combine in a way analogous to a statistical regression. Rather, children – like the present computational model – use the semantic and statistical regularities that fall out of the raw input data (not meta variables that describe them) to incrementally learn probabilistic links between verbs and constructions. Thus a successful computational model, unlike a regression model, will simulate a period of overgeneralization followed by retreat, and allows for the investigation of factors that alter the trajectory

of this learning process. For example, one question we investigate is how learning is affected by the presence of arbitrary lexical exceptions.

Method

The problem is conceptualized as one of the speaker learning, via comprehension, verb+construction mappings that allow her, in production, to select the appropriate construction, given the verb that she intends to use (in this case the PO-dative, the DO-dative or Other). This is, of course, a relatively high-level conceptualization, and one that abstracts across the many factors other than verb-level properties that determine construction choice (e.g., information structure, the relative length of the theme and recipient NPs etc., see Bresnan, Cueni, Nikitina & Baayen, 2007). Neither does it address the issue of how either verbs or constructions are acquired in the first place (see Twomey, Chang and Ambridge, 2014, for one simulation of the acquisition of verb semantics). Nevertheless, the rather abstract and theory-neutral nature of this conceptualization renders it potentially compatible with any theoretical approach which assumes that adult speakers possess some kind of abstract knowledge of verb argument structure constructions.

All simulations used the OXlearn MATLAB package (Ruh &Westermann, 2009). The learning task was instantiated in a three-layer feed-forward backpropagation network with seven input units (representing the verb), three hidden units, and three output units (representing PO-dative, DO-dative and Other). The structure of the network is summarized in Figure 1. Both the output and hidden layers used a sigmoid activation function (learning rate 0.01), and received input from a bias unit. Seven input units were used in order to allow each verb to be represented as a

vector across seven composite semantic features taken from Ambridge, Pine, Rowland, Freudenthal and Chang (2014), roughly speaking: *causing to go* (two predictors), *causing to have*, *speech, mailing, bequeathing* and *motion*. In this previous study, participants rated each verb for the extent to which it exhibits each of 18 semantic features relevant to the alternation, with these features condensed to seven using Principal Components Analysis (PCA). Each verb was represented in terms of its mean rating on each of these seven composite semantic features.

It is important to acknowledge that this implementation sidesteps the extremely difficult question of how learners acquire verb semantics, assuming – in effect – that learners have perfect knowledge of the semantics of every verb from the very first time that they hear it. In fact, real-life acquisition of verb semantics no doubt requires a considerable amount of experience, and presumably proceeds, at least in part, on the basis of syntactic and lexical distributional information (e.g., Gleitman, 1990; Pinker, 1994; Twomey, Chang & Ambridge, 2014). However, it is important to point out that this problem is shared by all current models of the acquisition of verbs' argument structure restrictions; both the verbal models outlined in the introduction, and the computational models outlined in the previous section. Thus this shortcoming is no reason to disregard the present model in favour of its contemporaries.

Figure 1. Architecture of the network
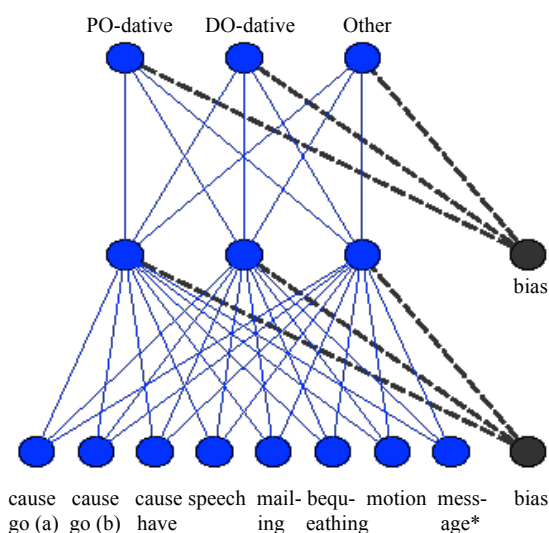
* = Pre-emption simulations only

Twenty four verbs were used; the core set from Ambridge, Pine, Rowland & Chang (2012), half PO-only, half alternating (note that this first simulation did not include DO-only verbs – e.g., *bet* and *wager* - which are of very low type and token frequency, particularly in speech to children, and so constitute a marginal phenomenon).

The PO-only verbs were drawn from two semantic classes: *pull*-verbs (*pull, drag, carry, haul, lift* and *hoist*) and *shout*-verbs (*shout, screech, whisper, hiss, scream* and *shriek*). The alternating verbs were drawn from two further classes: *show*-verbs (*show, teach, ask, pose, tell* and *quote*) and *give*-verbs (*give, hand, send, mail, throw* and *toss*).

For each training trial, a verb was presented to the network, along with its target construction (i.e., the target activation of the PO-dative, DO-dative or Other output unit was set to 1, with the other two output units set to 0). Verb+construction pairs were presented to the model in proportion to the log frequency with which the verb occurred in that construction in the British National Corpus (counts taken from Ambridge et al, 2014), as shown in Table 1 below. "Other" counts include all non-dative uses of the relevant verb, including – for example – simple transitives (*He pulled the rope*) and single word utterances (e.g., *Pull!*).

Table 1. Training set

Verb	PO-dative	DO-dative	Other				
pull	4	0	8	show	5	7	10
drag	4	0	7	teach	4	6	8
carry	5	0	9	ask	3	7	10
haul	3	0	6	pose	5	3	7
lift	5	0	8	tell	4	9	10
hoist	2	0	5	quote	2	3	7

shout	3	0	7	give	7	9	10
screech	1	0	5	hand	5	6	7
whisper	3	0	7	send	7	6	9
hiss	1	0	6	mail	2	2	5
scream	2	0	7	throw	5	5	8
shriek	1	0	5	toss	3	3	6
				Total	86	66	177

Thus a single training sweep consisted of 329 verb+construction pairs (86 PO-datives, 66 DO-datives and 177 Other constructions). The overgeneralization pressure on the model arises from the fact that half the verbs it encounters activate both the PO- and DO-dative output units (though only one or the other on any given trial), while the remainder activate only the PO-dative unit (with varying frequency).

For each test trial, the frozen model was presented with a verb - either a familiar verb from the training set or a novel verb (described below) - and the corresponding activation of the DO-dative output unit recorded. The activation of this output unit is taken as the model's "grammaticality judgment" for a DO-dative sentence with the relevant verb. The model's judgments were compared against those obtained from adult participants (from Ambridge et al, 2012). This method is preferable to simply investigating the model's ability to learn the training set to some error criterion (which is trivial, given the present set-up). It is important to emphasize that the model did not receive any information regarding participants' grammaticality judgments; as described above, target activations of output units were determined solely on the basis of corpus frequency.

Novel verbs were created in order to test the model's ability to generalize; that is, to determine the grammaticality or otherwise of previously-unseen verbs in the DO-dative construction on the basis of their semantics. Ambridge et al (2012) found that adults displayed this ability, though children aged 5-6 and 9-10 did not. Novel verbs

were created by averaging across the semantic ratings for all of the verbs in the relevant semantic class, excluding the target verbs. This resulted in the creation of four novel verbs, two PO-only (novel *pulling* and novel *shouting*), two alternating (novel *showing* and novel *giving*). Ambridge et al (2012) found that adults rated DO-dative uses of the former, but not the latter, as ungrammatical.

Results

Semantics+Entrenchment model. The model as described above implements entrenchment, but not pre-emption, as all non-DO-dative uses of a particular verb, whether PO-dative or Other, have an equal impact in causing the model *not* to activate the DO-dative output unit for this verb (pre-emption is added to a subsequent model). The model implements a role for verb semantics, by virtue of the fact that each verb is represented as a vector of seven semantic feature scores. The model was trained for 100,000 sweeps (each consisting of 329 verb+construction pairs) and its output recorded every 10,000 sweeps. All results presented here and subsequently average across ten runs of the model with different random seeds. Figure 2 shows the familiar and novel-verb results for the Semantics+Entrenchment model, averaging across the six verbs in each class.

Figure 2. *Semantics+Entrenchment* **model**
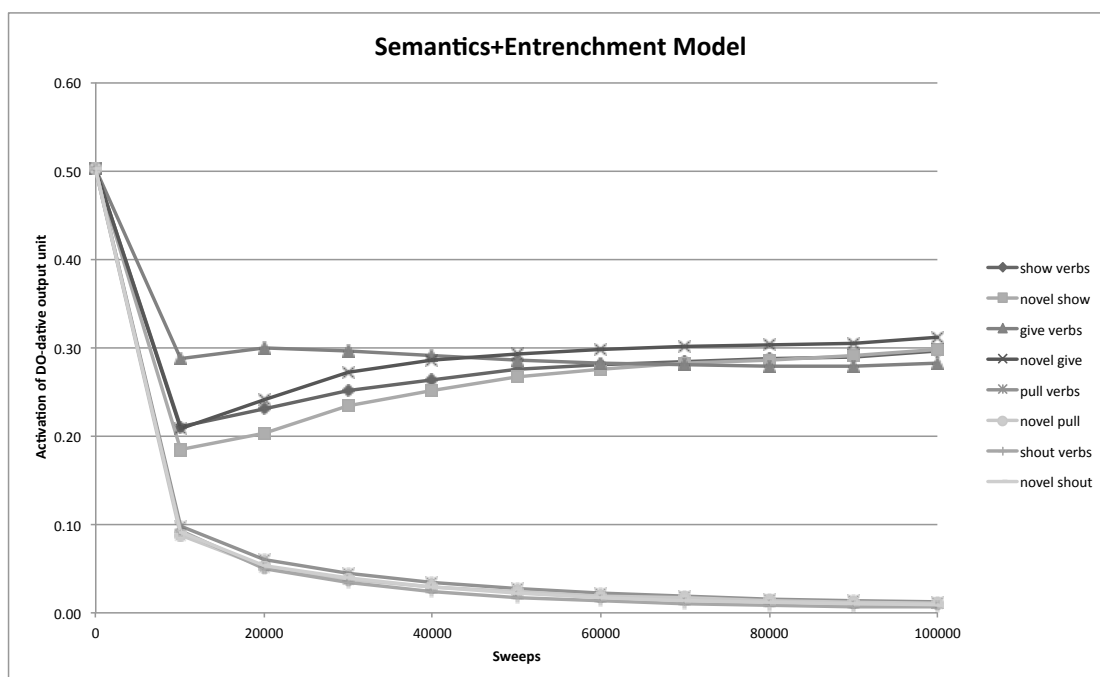
The model rapidly learns that the DO-datives are acceptable (i.e., to activate this output unit) for the *show* and *give* verbs, but not the *pull* and *shout* verbs. (The reason that the activation of the DO-dative output unit drops to around 0.3 even for verbs that are grammatical in this construction is that the model is learning that PO-dative and Other uses [e.g., transitives; single-word uses] are also possible constructions for this verb). The model also generalizes this pattern to the four novel verbs. This latter finding demonstrates one way in which it is possible for a model that includes no hard-wired discrete verb classes to show class-type generalization behaviour. A cluster plot of the hidden units (Figure 3) demonstrates that the model achieves this behaviour by forming representations in the hidden layer that map semantically-similar verbs onto the same output unit.

Interestingly, at the second coarsest grain size (the coarsest being PO-only and alternating verbs) the four clusters essentially correspond, with few exceptions, to Pinker's (1989) classes of "manner of speaking (*shout, screech, whisper, hiss, scream*

and *shriek*)" (p.112), "continuous imparting of force in some manner causing

accompanied motion (*pull, drag, carry, haul, lift* and *hoist*)" (p.110), "illocutionary

verbs (*show, teach, ask, pose, tell* and *quote*)" (p.112) and "verbs of giving (*give, hand,*

send, mail, throw and *toss)*" (p.110). Crucially, however, the model also groups verbs

at a finer level. For example, Figure 3 shows that the model conceptualizes the two

members in the pairs *hiss+screech*, *lift+carry*, *hoist+haul*, *ask+tell*, *throw+toss*, and

send+give as more similar to one another than to other verbs in the same Pinker class.

The ability to instantiate semantic similarity at a more fine-grained level is presumably

key if the model is to simulate the graded pattern of judgments shown by human

participants (explored in detail in subsequent simulations). Although it is beyond the

scope of the present investigation to implement and test a Pinker-style class-based

model directly, it seems unlikely that a model that can assign only one of four discrete

values (corresponding to the classes) will be able to simulate this graded pattern.

Figure 3. Cluster tree of hidden units in *Semantics+Entrenchment* model

The finding that the model does not link the familiar *pull* and *shout* verbs to the DO-dative is, in one sense, trivial: Because the activation of the output units sums to 1 and these verbs activate only the PO-dative and Other units during training, it is inevitable that the model will not activate the DO-dative unit for these verbs. In another sense, however, the triviality of this finding is exactly the point: A learner that probabilistically links verbs and the constructions in which they have appeared will inevitably show an "entrenchment" effect, even while retaining the ability to generalize novel verbs into unattested constructions on the basis of their semantics; there is no need to posit entrenchment as a special dedicated mechanism.

That said, the *Semantics+Entrenchment* model fails in two important respects. First, unlike children, it does not display a period of overgeneralization. For the *pull* and *shout* verbs – familiar and novel alike – the activation of the DO-dative unit drops rapidly to below 0.1 within the first 10,000 sweeps. Second, at no point during learning do the model's judgments of overgeneralization errors (i.e., DO-uses of PO-only verbs) correlate with those obtained from adult participants (see Figure 4, plotted at 60,000 sweeps, for comparison with the subsequent model).

Figure 4. No correlation between model and human ratings for
***Semantics+Entrenchment* model**
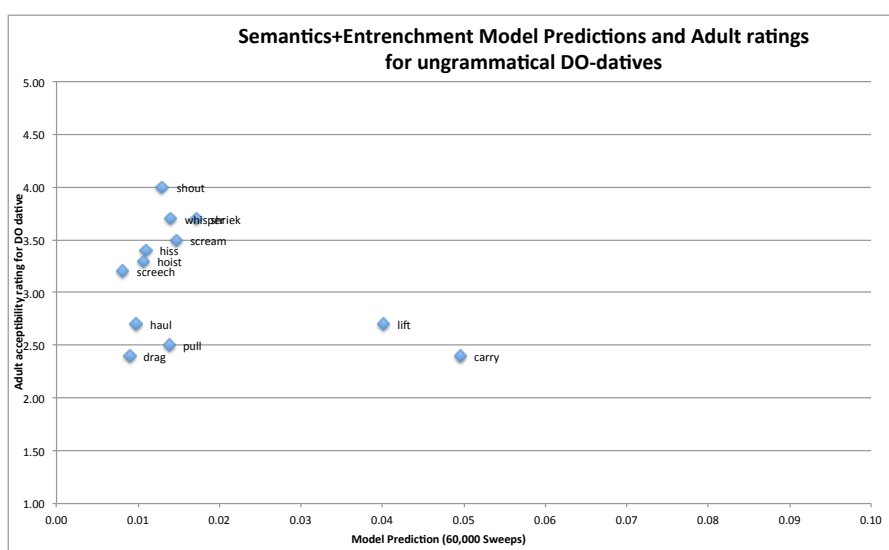
Semantics+Entrenchment+Pre-emption **model**. The pre-emption hypothesis

holds that overgeneralization errors are blocked not by *any* uses of the relevant verb (as

under entrenchment), but only – or, at least, especially - by uses that express the same

intended meaning (or "message"). For example, the error **Marge pulled Homer the*

box (*pull*+DO-dative) would be pre-empted by *pull*+PO-dative sentences (e.g., *Bart*

pulled the box to Lisa), as both sentences express a three-argument "transfer" message.

Such an error would not be pre-empted by simple transitives (e.g., *He pulled the rope*),

one-word utterances (*Pull*!), and so on, as such sentences do not express a transfer

message.

Pre-emption was instantiated in the model by adding an additional input unit to

encode the message. This unit was set to 1 when the target output unit was either the

PO- or DO-dative unit (= "transfer message") and 0 when the target output was Other

(= "non-transfer message"). In terms of real-word learning, the assumption is that

learners understand the speaker's intended message in comprehension (=the model's

learning phase) and have in mind their own intended message in production (=the

model's test phase). In all other respects, including the training set, the model was

identical to that outlined above.

This model (see Figure 5) addressed both of the shortcomings of the

Semantics+Entrenchment model. First, in contrast to this previous model, it showed a

protracted "overgeneralization" period (approximately 0-40,000 sweeps) in which

verbs that been presented solely in PO-dative sentences during training (the familiar

pull and *shout* verbs) activated the DO-dative output unit, with an activation strength

similar to that yielded by the alternating verbs (the familiar *show* and *give* verbs).

Presumably, this overgeneralization period is consequence of the fact that the "transfer

message" unit (which was always set to 1 during testing, since a DO-dative judgment

was being elicited) mapped to both the PO- and DO- output units during training. The

model retreats from overgeneralization as it learns *which* particular dative unit, PO or

DO, is appropriate for each verb; information that this model (like the previous model)

rapidly generalizes to novel verbs, presumably on the basis of semantic overlap with

familiar verbs (a later simulation tests this presumption).

Figure 5. *Semantics+Entrenchment+Pre-emption* **model**
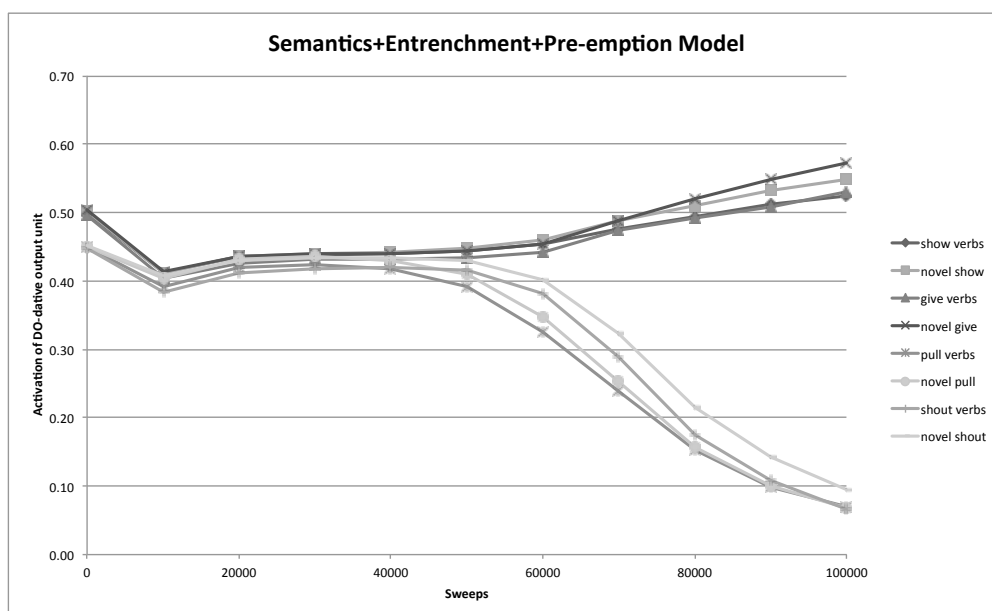
Addressing the second shortcoming of the previous model, the current model

simulated, at 50,000 and 60,000 sweeps, the by-verb pattern of adult grammaticality

judgment data for overgeneralization errors of PO-only verbs into the DO-dative

construction ($r=0.66$, $p=0.02$; $r=0.68$, $p=0.02$), see Figure 6. Beyond this point, no

significant correlations were observed, presumably because the model had over-

learned the solution, increasingly treating all verbs attested only in the PO-dative as

extremely – and equally – ungrammatical in the DO-dative.

Of course, the fact that a three-parameter model (*Semantics + Entrenchment + Pre-emption*) outperforms a two-parameter model (*Semantics + Entrenchment*) should surprise no-one. The point is that the Pre-emption mechanism is not a free parameter added for no reason other than to improve coverage of the data, but rather an implementation of a particular theoretical proposal that is well supported by previous empirical studies with children and adults. Presumably the reason that the Pre-emption mechanism (i.e., the "message" node) plays such a key role is that, without it, there is very little overgeneralization pressure on the model, which can simply map conservatively from input to output. This pressure arises only when a communicative need (i.e., the desire to express a transfer message) compels the learner to use a verb in a construction in which it has never (or very infrequently) been attested. Indeed, examination of children's errors (e.g., *I said her no*) suggests that these too are produced when the child's desire to express an intended message compels her to extend a verb into a construction which expresses that message, despite the fact that this combination is unattested in the input.

Figure 6. Correlation between model and human ratings for *Semantics+Entrenchment+Pre-emption* model
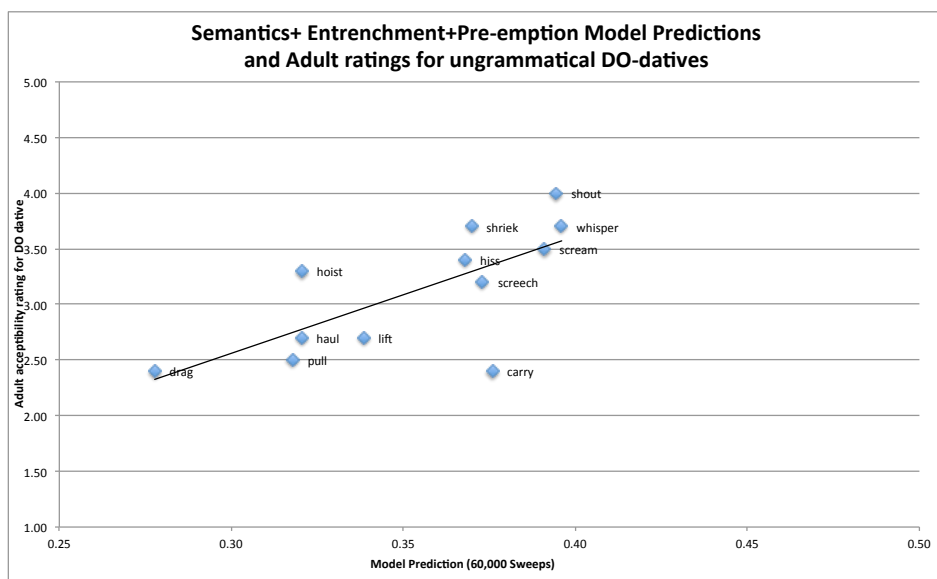
Lexeme-based Semantics+Entrenchment+Pre-emption **model**. The models

presented so far have – purely as a simplifying assumption – represented verbs solely

as bundles of semantic features, meaning that there is a very high degree of overlap

between verbs with similar semantics. However, this assumption is unrealistic, as real

learners encounter a (relatively) consistent phonological representation of each verb.

This lexeme binds together the particular bundle of semantic properties associated with

that verb and, crucially, differentiates this bundle from overlapping bundles associated

with other verbs. In order to instantiate this property, we added to the input layer a

further 28 input units, each representing an individual verb (24 familiar, 4 novel). The

training and test phases were the same as for the previous models, except that the input

unit representing the relevant verb was set to 1, with the remaining 27 units set to 0.

Compared with the previous *Semantics+Entrenchment+Pre-emption* model ,

this *lexeme-based Semantics+Entrenchment+Pre-emption model* (see Figure 7)

showed a slightly shorter period of "overgeneralization" with the predictions for PO-

only and alternating verbs beginning to diverge at around 30,000 sweeps (as opposed

to 40,000 for the previous model).

Figure 7. Lexeme-based *Semantics+Entrenchment+Pre-emption* model
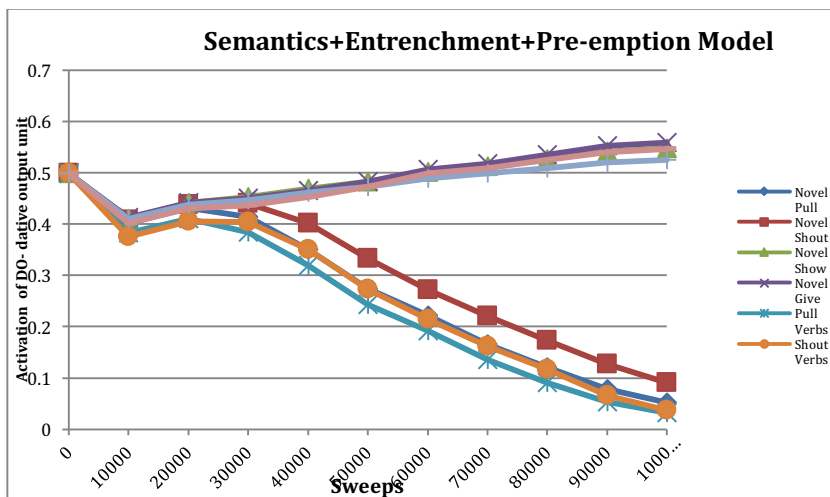
The lexeme-based model was also slightly better at predicting adults' judgments, with significant correlations observed both earlier (30,000 sweeps: $r = 0.59$, $p = 0.042$; 40,000: $r = 0.64$, $p = 0.036$; 50,000: $r = 0.58$, $p = 0.048$, see Figure 8) and later (70,000: $r = 0.61$, $p = 0.035$; 80,000: $r = 0.62$, $p = 0.03$) in development (though at 60,000 sweeps the correlation was not significant: $r = 0.56$, $p = 0.06$). Taken together with the fact that real learners encounter a binding lexeme for each verb presentation, the (albeit-modest) improvement in coverage shown by this model suggests that it is important for all models of this phenomenon to include both a semantic and lexical component (and we therefore retain this lexeme-based model for the remaining simulations).

Figure 8. Correlation between model and human ratings for the lexeme-based *Semantics+Entrenchment+Pre-emption* model
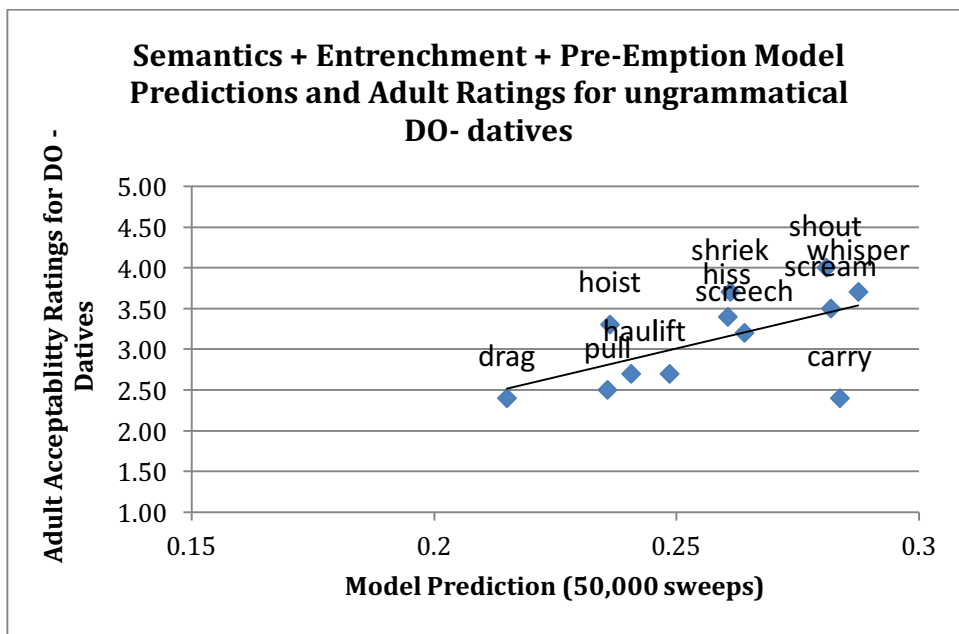
Semantics vs Pre-emption: Can pre-emption be used to learn arbitrary exceptions to semantically-based generalizations? It has often been suggested (e.g.,

Boyd & Goldberg, 2011; Goldberg, 1995; 2006) that pre-emption might be useful for learning exceptions to semantically-based generalizations (though, in the case of the dative, there has been some debate as to whether, given sufficiently fine-grained and probabilistic generalizations, such exceptions in fact exist; see Ambridge et al, 2014: 237). Consider, for example, a hypothetical verb that is semantically consistent with both the PO- and DO-dative (e.g., a novel *giving* or *showing* verb) but that, for some reason, happens to appear only in the PO-dative construction[1]. Would learners treat it as an exception (as pre-emption would predict), or would the exception be swamped by the semantic generalization?

In order to explore this question, the four novel verbs used previously in the test set only were added to the training set shown in Table 1. Each was presented 10 times per sweep, **always in PO-dative constructions** (hence they are referred to subsequently as *PO-only novel pull/shout/show/give*). *PO-only novel pull* and *PO-only novel shout* are best thought of as "control" verbs: Both semantics and pre-emption push the model in the direction of rejecting the DO-dative, which it would therefore be expected to do rapidly. *PO-only novel show* and *PO-only novel give* instantiate the thought-experiment outlined above, and pit semantics and pre-emption against one another. Each of these novel verbs is semantically similar to six verbs that appear in both the PO- and DO-dative during training; thus their semantics push the model in the direction of predicting the DO-dative for that verb. On the other hand, both are attested with very high frequency in the PO-dative only (10 presentations per sweep; a rate chosen to be higher than any other verb+construction pair in the dataset). Thus pre-emption pushes the model in the direction of rejecting the DO-dative for that verb (i.e., predicting the PO-dative instead). This new training set was given to the lexeme-based *Semantics+Entrenchment+Pre-emption* model outlined above.

Figure 9 plots these results for the novel verbs only (results for the familiar verbs were the same as for the previous model). Despite the high levels of pre-emption, the semantic information still holds considerable sway: From 10,000-60,000 sweeps activation of the DO-dative unit is higher for the semantically-alternating novel *show* and *give* verbs than for the semantically-PO-only novel *pull* and *shout* verbs. Nevertheless, slowly but surely, pre-emption wins out over semantics: By 70,000 sweeps, the semantically-alternating novel *show* and *give* verbs are indistinguishable from the semantically-PO-only novel *pull* and *shout* verbs, with activation of the DO-dative unit essentially zero by the end of the simulation at 100,000 sweeps. For comparative purposes, recall that the standard lexeme-based *Semantics+Entrenchment+Pre-emption* model showed the expected pattern for novel verbs from around 30,000 sweeps. Thus, although it takes some time, pre-emption can indeed be used to learn arbitrary exceptions to semantically-based generalizations (e.g., Boyd & Goldberg, 2011; Goldberg, 1995; 2006).

Figure 9. Semantics vs Pre-emption
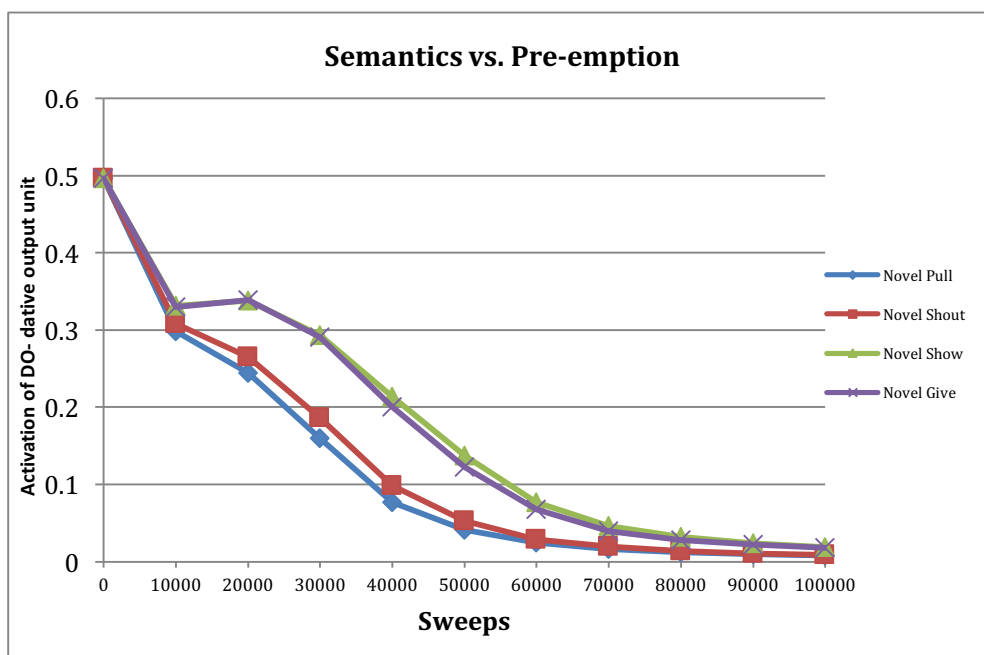
Extending the lexeme-based Semantics+Entrenchment+Pre-emption model.

The final set of simulations investigated whether the model would scale up to a larger

dataset: the full set of 301 dative verbs rated by adults in the study of Ambridge et al

(2014), comprising 145 alternating verbs, 131 PO-only verbs and 25 DO-only verbs.

Because this set was designed to be as comprehensive as possible, including every

English dative verb indentified in two major reference works on the topic (Levin,

1993; Pinker, 1989), it constitutes an appropriate test of the model's ability to scale up

to something like a life-sized dataset.

Table 2. Extended training set.

Verb	PO	DO	OT.	Verb				Verb			
Feed	4	4	8	Ship	3	0	6	Nudge	1	0	5
Give	7	9	10	Shunt	1	0	4	Pass	5	4	9
Lease	3	2	5	Slip	2	3	7	Pitch	1	0	6
Lend	6	5	7	Smuggle	2	0	5	Punt	0	0	3
Loan	2	3	5	Sneak	1	1	5	Shoot	2	4	8
Pass	5	4	9	UPS	0	0	0	Shove	2	1	6
Pay	7	6	9	Transport	4	0	6	Slam	1	0	6
Peddle	1	0	4	Deliver	5	0	8	Slap	1	0	6
Refund	1	1	5	Airfreight	0	0	1	Sling	0	0	5
Render	3	3	7	FedEx	0	0	0	Smash	1	0	6
Rent	3	2	6	Courier	0	0	1	Tap	1	0	7
Repay	2	2	6	Messenger	0	0	0	Throw	5	5	8
Sell	6	5	9	Bash	0	0	5	Tip	2	1	6
Serve	3	4	8	Bat	0	0	5	Toss	3	3	6
Trade	1	0	7	Bunt	0	0	0	Poke	0	0	6
Hand	5	6	7	Cast	2	3	7	Blast	1	0	6
Donate	4	0	6	Catapult	1	0	4	Propel	2	0	5
Contribute	5	0	8	Chuck	1	2	6	Release	3	0	8
Forward	3	2	5	Flick	2	2	6	Alley-oop	0	0	0
Hand	5	6	7	Fire	1	0	7	Lob-pass	0	0	0
Mail	2	2	5	Fling	3	2	6	Bounce	0	0	6
Pass	5	4	9	Flip	2	0	5	Float	1	0	7
Port	3	0	5	Hit	3	0	8	Move	5	0	9
Post	3	0	6	Hurl	2	0	6	Roll	2	0	7
Send	7	6	9	Kick	3	0	7	Slide	2	2	7
Shift	4	0	7	Lob	1	1	4	Carry	5	0	9
				Loft	0	0	3	Drag	4	0	7

Haul	3	0	6	Spare	0	4	6	Radio	0	0	3
Heave	2	1	6	Envy	0	4	5	Relay	3	0	5
Heft	0	0	3	Begrudge	1	3	3	Satellite	0	0	0
Hoist	2	0	5	Refuse	3	4	8	Semaphore	0	0	2
Kick	3	0	7	Ask	3	7	9	Sign	2	0	8
Lug	1	0	4	Save	2	5	8	Signal	3	1	6
Pull	4	0	8	Forgive	1	4	7	Telephone	2	0	7
Push	5	0	8	Deny	5	6	8	Telecast	0	0	1
Schlep	0	0	1	Bet	0	3	7	Telegraph	1	0	3
Shove	2	1	6	Bill	0	1	5	Telex	0	0	4
Tote	0	0	3	Charge	3	5	8	Wire	0	1	5
Tow	2	0	5	Fine	0	4	6	Wireless	0	0	1
Tug	2	0	6	Mulct	0	0	2	Babble	0	0	4
Drop	3	4	8	Overcharge	0	0	3	Bark	1	0	6
Hoist	2	0	5	Save	2	5	8	Bawl	0	0	4
Lift	5	0	8	Spare	0	4	6	Bellow	1	0	5
Lower	4	0	7	Tax	0	0	6	Bleat	0	0	4
Raise	5	0	9	Tip	2	1	6	Boom	0	0	5
Bring	7	6	9	Undercharge	0	0	1	Bray	0	0	4
Take	7	6	11	Wager	0	0	4	Burble	0	0	4
Advance	2	3	7	Ask	3	7	9	Cackle	0	0	4
Allocate	4	4	7	Cite	1	0	7	Call	4	0	9
Allot	3	2	5	Preach	2	0	6	Carol	0	0	2
Assign	5	4	6	Quote	2	3	7	Chant	0	0	5
Award	4	5	7	Read	4	5	9	Chatter	0	0	5
Bequeath	3	2	5	Relay	3	0	5	Chirp	0	0	3
Cede	3	0	4	Show	5	7	10	Cluck	1	0	4
Concede	3	2	6	Teach	4	6	8	Coo	0	0	3
Extend	4	0	8	Tell	4	9	10	Croak	0	0	5
Grant	5	6	8	Write	5	0	9	Croon	0	0	4
Guarantee	3	4	7	Pose	5	3	7	Crow	0	0	4
Issue	5	3	8	Spin	1	2	7	Cry	2	0	8
Leave	6	5	10	Explain	5	0	9	Drawl	0	0	3
Offer	6	7	9	Announce	3	0	8	Drone	0	0	4
Owe	6	6	7	Describe	4	0	9	Gabble	0	0	4
Promise	3	5	8	Admit	4	0	8	Gibber	0	0	3
Vote	2	2	8	Confess	3	2	6	Groan	0	0	6
Will	1	1	6	Repeat	3	0	8	Growl	0	0	5
Yield	3	3	7	Declare	3	0	8	Grumble	0	0	5
Refer	5	0	8	Recount	2	0	5	Grunt	0	0	5
Forward	3	2	5	Cable	1	2	4	Hiss	1	0	6
Guarantee	3	4	7	Email	0	1	2	Holler	0	0	3
Reserve	1	0	7	Fax	2	2	5	Hoot	0	0	4
Recommend	4	0	8	Modem	0	0	0	Howl	1	0	5
Permit	2	4	7	Netmail	0	0	0	Jabber	0	0	3
Cost	2	6	8	Phone	1	2	7	Lilt	0	0	3

Word				Word				Word			
Lisp	0	0	3	Twitter	0	0	3	Note	1	0	8
Moan	0	0	6	Wail	0	0	5	Observe	0	0	8
Mumble	1	0	5	Warble	0	0	3	Proclaim	2	4	6
Murmur	2	0	7	Wheeze	0	0	4	Propose	2	0	8
Mutter	3	0	7	Whimper	0	0	4	Reiterate	1	0	6
Purr	0	0	4	Whine	0	0	5	Relate	0	0	8
Rage	0	0	5	Whisper	3	0	7	Remark	1	0	7
Rasp	0	0	5	Whistle	1	1	6	Report	5	0	9
Roar	1	0	6	Whoop	0	0	4	Reveal	4	0	8
Scream	2	0	7	Yammer	0	0	1	Say	6	0	11
Screech	1	0	5	Yap	0	0	4	State	1	0	8
Shout	3	0	7	Yell	2	0	6	Suggest	3	0	9
Shriek	1	0	5	Yelp	0	0	4	Doubt	0	0	7
Sing	3	0	8	Yodel	0	0	2	Question	0	0	7
Snap	1	0	7	Admit	4	0	8	Credit	1	0	6
Snarl	0	0	5	Allege	0	0	7	Entrust	4	0	5
Snuffle	0	0	4	Announce	3	0	8	Furnish	2	0	6
Splutter	0	0	5	Articulate	2	0	6	Issue	5	3	8
Squal	0	0	2	Assert	0	0	7	Leave	6	5	10
Squawk	0	0	4	Blab	1	0	2	Present	6	0	8
Squeak	0	0	5	Blurt	0	0	5	Provide	6	0	9
Squeal	0	0	5	Claim	0	0	9	Serve	3	4	8
Stammer	0	0	4	Communicate	4	0	7	Supply	5	3	8
Stutter	0	0	4	Confess	3	2	6	Trust	3	0	7
Thunder	0	0	5	Confide	3	0	5	Reward	0	0	6
Tisk	0	0	0	Convey	4	0	7	Honour	0	0	6
Trill	0	0	3	Declare	3	0	8	Bestow	0	0	5
Trumpet	0	0	4	Mention	5	0	8				

This model was trained in exactly the same way as for the previous version, but with this larger training set. The model again showed a good fit to the adult data ($r=0.54$, $p<0.001$; see Figure 10), reaching asymptote at around 1,000 sweeps; considerably sooner that previous model (which is to be expected given the much larger dataset).

Figure 10. Correlation between model and human ratings for Semantics+Entrenchment+Pre-emption model with extended dataset (301 verbs).
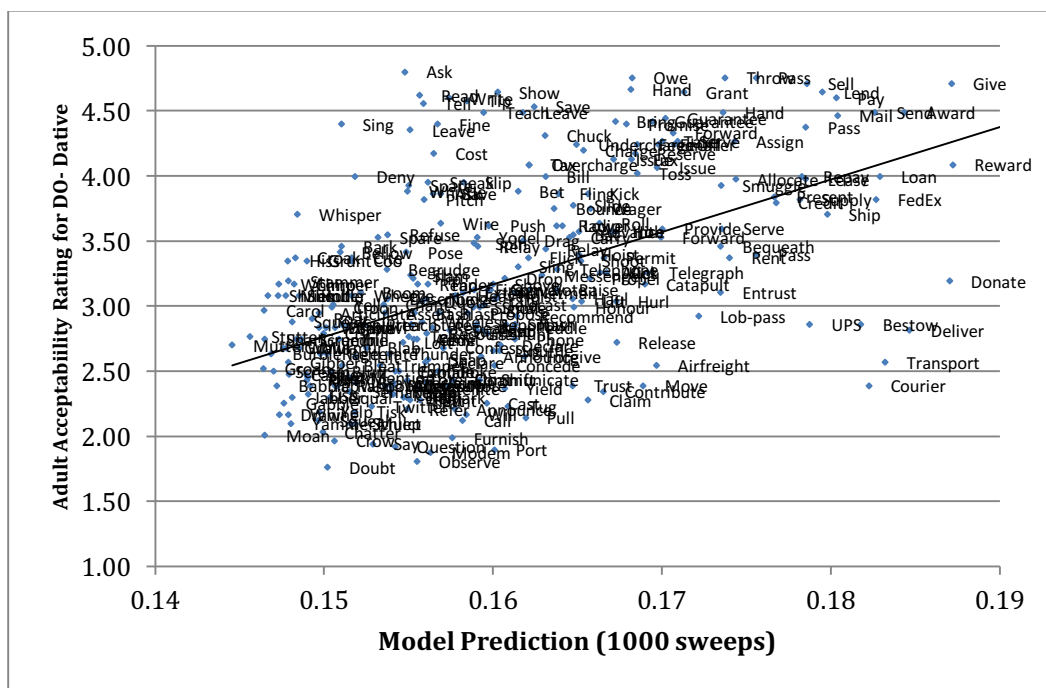
To summarise, the most successful model – the lexeme-based Semantics+Entrenchment+Pre-emption model – was successful in (a) simulating an overall overgeneralization-then-retreat pattern (b) predicting the correct dative argument structure for novel verbs on the basis of their semantics and (c) modeling the fine-grained pattern of by-verb grammaticality judgments observed in adult studies, including a large scale study that included almost all English dative verbs.

Discussion

A central question in the cognitive sciences is how children build linguistic representations that allow them to generalize verbs from one construction to another (e.g., *The boy gave a present to the girl* → *The boy gave the girl a present*), whilst appropriately constraining those generalizations to avoid nonadultlike errors (e.g. *I*

said no to her ➔ **I said her no*). Indeed, for the many children who pass through a stage in which they produce such errors, the question is how they learn to retreat from them, given the absence of consistent evidence regarding which of their utterances are ungrammatical.

Recently, a consensus has begun to emerge that children solve this "no negative evidence problem" (Bowerman, 1988), using a combination of statistical learning procedures such as entrenchment (e.g., Theakston, 2004) and pre-emption (e.g., Boyd & Goldberg, 2011), and learning procedures based on verb semantics (e.g., Ambridge et al, 2008, 2009, 2011, 2012a, 2012b, 2014). Despite this emerging consensus, there have been few attempts to propose a unitary account that combines all three approaches. One exception is the FIT account (Ambridge et al, 2011) which argues for competition between constructions based on (a) verb-in construction frequency, (b) relevance of constructions for the speaker's intended message and (c) fit between the fine-grained semantic properties of individual verbs and individual constructions.

The present study demonstrated that a simple connectionist model that instantiates this account can not only simulate the overall pattern of overgeneralization then retreat but also use the semantics of novel verbs to predict their argument structure (as in the human studies of Ambridge et al, 2008, 2009, 2012b; Bidgood et al, 2014) and predict the by-verb pattern of grammaticality judgments observed in adult studies (Ambridge et al, 2012b).

Although the model used is computationally extremely simple, there is an important sense in which this is its greatest strength. The success of the model suggests that statistical learning effects such as entrenchment and pre-emption need not make use of sophisticated Bayesian or rational learner algorithms to compute an inference from absence. Rather, these effects arise naturally from a learning mechanism that

probabilistically links verbs to competing instructions. Similarly, semantic effects need not rely on an explicit procedure for semantic class formation (e.g., Pinker, 1989), but fall naturally out of a model that learns which bundles of semantic features ("verbs") are predictive of which constructions.

Another advantage of the present model's simple and high-level approach is that it can easily be extended to other constructions for which children are known to make overgeneralization errors, and for which suitable semantic and statistical measures have been collected. These include the locative (e.g., Ambridge, Pine & Rowland, 2012), passive (Ambridge, Bidgood, Pine, Rowland & Freudenthal, in press) and reversative *un-* prefixation (e.g., Ambridge, 2013; Blything, Ambridge & Lieven, 2014). Future simulations using the same architecture could also investigate the semantic restrictions on construction slots other than VERB. For example, a (probabilistic) requirement of the DO-dative construction is that the first argument be a potential possessor of the second argument (e.g., **John sent Chicago the package; c.f.,* the PO-dative equivalent *John sent the package to Chicago*). It would also be possible to investigate other types of overgeneralization error (e.g., the [in]famous case of the English past tense) by using phonological rather than semantic representations at the input level (indeed, the present study essentially uses the same architecture as classic past-tense models such as Rumelhart & McClelland, 1986).

These advantages notwithstanding, it is important to acknowledge the ways in which the present simulations considerably simplify the task facing real learners. First, the precise semantic properties of individual verbs are known from the start. For child learners, acquiring verb meanings is a notoriously difficult task (Gillette, Gleitman, Gleitman, & Lederer, 1999), and one that presumably proceeds mostly in parallel with learning verb argument structure constructions (e.g., Twomey et al, 2014). Second, the

simulated learner is assumed to have already abstracted the necessary verb argument

structure constructions (e.g., PO- and DO-dative) from the input, and to be able to

correctly recognize all further instances of these constructions in the input (though the

semantic characteristics of these constructions are learned during the simulation). For

real learners, acquiring verb argument structure constructions is an extremely difficult

task; indeed, there are very few proposals for how this might be done (though see

Tomasello, 2003; Alishahi & Stevenson, 2008). Finally, the model does not – unlike

both real learners and more sophisticated computational models (e.g. Chang, 2002;

Chang et al, 2006) – produce sentences as sequences of temporally ordered words.

There is quite a leap to be made from (a) knowing the verb+argument structure

combination that one intends to use to (b) producing a well-formed sentence.

Indeed, under many accounts, verb argument constructions are not, in fact, seen

as entities that are abstracted from the input, then stored for subsequent retrieval.

Exemplar-based accounts (e.g., Bybee, 2013; Langacker, 2000) propose that learners

store nothing more than individual exemplars, (in this case, sentences), and that the

notion of – for example – "using a DO-dative construction" – is simply a shorthand

way of referring to a process of online generalization across stored DO-dative

sentences that meet some criterion (e.g., similarity to the intended message). In order

to instantiate such accounts computationally, we will need considerably more

sophisticated models that are able to use both semantic and distributional

commonalities to abstract, on the fly, across stored exemplars in a way that yields

something like conventional verb argument-structure constructions.

In the meantime, the present findings suggest that the traditional

conceptualization of entrenchment, pre-emption and the formation of semantic

generalizations as rival "mechanisms" may be unhelpful. Rather, all three are best

thought of as labels for *effects* that fall naturally out of a learning mechanism that probabilistically associates particular verbs (where each verb is a collection of semantic features) and particular argument structure constructions. The computational model outlined in the present article constitutes one possible account of how this can be done.

References

Akhtar, N., & Tomasello, M. (1997). Young children's productivity with word order and verb morphology. *Developmental Psychology 33*, 952-65.

Alishahi, A., & Stevenson, S. (2008). A computational model of early argument structure acquisition. *Cognitive Science 32*, 789-834.

Ambridge, B. (2013). How do children restrict their linguistic generalizations?: an (un-)grammaticality judgment study. *Cognitive Science 37*, 508-43.

Ambridge, B. & Brandt, S. (2013). "Lisa filled water into the cup": The roles of entrenchment, pre-emption and verb semantics in German speakers' L2 acquisition of English locatives. *Zeitschrift fur Anglistik und Amerikanistik 61*, 245-63.

Ambridge, B. & Lieven, E.V.M. (2011). *Child Language Acquisition: Contrasting Theoretical Approaches*. Cambridge: Cambridge University Press.

Ambridge B, & Lieven E.V.M. (2015). A Constructivist Account. In: MacWhinney B and O'Grady W, editors. Handbook of Language Emergence. Hoboken, NJ: Wiley Blackwell; 2015. pp. 478–510.

Ambridge, B., Pine, J. M, Rowland, C. F., Freudenthal, D., & Chang, F.

(2014). Avoiding dative overgeneralization errors: semantics, statistics or

both? *Language, Cognition and Neuroscience 29*, 218-43.

Ambridge, B., Pine, J. M., Rowland, C. F. & Chang, F. (2012). The roles of verb

semantics, entrenchment and morphophonology in the retreat from dative

argument structure overgeneralization errors. *Language 88*, 45–81.

Ambridge, B., Pine, J. M., Rowland, C. F., Jones, R.L, & Clark, V. (2009). A

semantics-based approach to the 'no negative-evidence' problem *Cognitive

Science 33*, 1301-16.

Ambridge, B., Pine, J.M. & Rowland, C.F. (2012). Semantics versus statistics in the

retreat from locative overgeneralization errors. *Cognition 123*, 260–79.

Ambridge, B., Pine, J.M. & Rowland, C.F., (2011). Children use verb semantics to

retreat from overgeneralization errors: A novel verb grammaticality judgment

study. *Cognitive Linguistics 22*, 303–23.

Ambridge, B., Pine, J.M., Rowland, C.F. & Young, C.R. (2008) The effect of verb

semantic class and verb frequency (entrenchment) on children's and adults'

graded judgements of argument-structure overgeneralization errors. *Cognition

106*, *87-129.*

Ambridge, B., Pine, J.M., Rowland, C.F., Chang, F. & Bidgood, A. (2013). The

retreat from overgeneralization in child language acquisition: Word learning,

morphology and verb argument structure. *Wiley Interdisciplinary Reviews:

Cognitive Science 4*, 47–62.

Bidgood, A., Ambridge, B., Pine, J.M. & Rowland, C.F. (2014). The retreat from

locative overgeneralisation errors: A novel verb grammaticality judgment

study. *PLoS ONE.*

Blything, R.P. Ambridge, B. & Lieven, E.V.M. (2014). Children Use Statistics and

Semantics In the Retreat From Overgeneralization. PLoS ONE.

Bowerman, M. (1988). The "no negative evidence" problem: how do children avoid

constructing an overly general grammar? In J. A. Hawkins (Ed.), *Explaining*

language universals (pp. 73-101). Oxford: Blackwell.

Boyd, J. K., & Goldberg, A. E. (2011). Learning what not to say: The role of statistical

preemption and categorization in a-adjective production. *Language 87*, 55-83.

Braine, M. D. S. (1971). On two types of models of the internalization of grammars. In

D. I. Slobin (ed.), *The ontogenesis of grammar: a theoretical symposium*. New

York: Academic Press.

Braine, M. D. S., & Brooks, P. J. (1995). Verb argument strucure and the problem of

avoiding an overgeneral grammar. In M. Tomasello & W. E. Merriman (Eds.),

Beyond names for things: young children's acquisition of verbs (pp. 352-76).

Hillsdale, NJ: Erlbaum.

Bresnan, J., Cueni, A., Nikitina, T. V., & Baayen, H. R. (2007). Predicting the dative

alternation. In G. Boume, I. Kramer & J. Zwarts (Eds.), *Cognitive foundations*

of interpretation (pp. 69-94). Amsterdam: Royal Netherlands Academy of

Science.

Brooks, P. J., & Tomasello, M. (1999). How children constrain their argument-

structure constructions. *Language 75*, 720-38.

Brooks, P. J., & Zizak, O. (2002). Does preemption help children learn verb

transitivity?. *Journal of Child Language 29*, 759-81.

Brooks, P. J., Tomasello, M., Dodson, K., & Lewis, L. B. (1999). Young children's

overgeneralizations with fixed transitivity verbs. *Child Development 70*, 1325-

1337.

Bybee, J.L. (2013). Usage-based theory and exemplar representation. In Thomas

Hoffman, T. and G. Trousdale (eds.) *The Oxford Handbook of Construction*

Grammar, pp. 49-69. Oxford University Press.

Chang, F. (2002). Symbolically speaking: A connectionist model of sentence

production. *Cognitive Science 26*, 609-651.

Chang, F., Dell, G. S., & Bock, K. (2006). Becoming syntactic. *Psychological Review 113*, 234.

Chomsky, N. (1957). *Syntactic Structures*. The Hague: Mouton.

Dowman, M. (submitted). Minimum description length as a solution to the problem of generalization in syntactic theory. *Machine Learning and Language*.

Dowman. M. (2000) Addressing the Learnability of Verb Subcategorizations with Bayesian Inference. In Gleitman, L. R. & Joshi, A. K. (Eds.) *Proceedings of the Twenty-Second Annual Conference of the Cognitive Science Society*. Mahwah, New Jersey, USA: Lawrence Erlbaum Associates.

Gertner, Y., Fisher, C., & Eisengart, J. (2006). Learning Words and Rules Abstract Knowledge of Word Order in Early Sentence Comprehension.*Psychological Science 17*, 684-91.

Gillette, J., Gleitman, H., Gleitman, L., & Lederer, A. (1999). Human simulations of vocabulary learning. *Cognition 73*, 135-76.

Gleitman, L. (1990). The structural sources of verb meanings. *Language Acquisition 1*, 3-55.

Goldberg, A., E. (1995). *Constructions: A Construction Grammar Approach to Argument Structure*. Chicago: University of Chicago Press.

Goldberg, A.E. (2006). *Constructions at work: The nature of generalization in language*. Oxford: Oxford University Press.

Goldberg, A. E. (2011). Corpus evidence of the viability of statistical preemption. *Cognitive Linguistics 22*, 131-153.

Gropen, J., Pinker, S., Hollander, M., & Goldberg, R. (1991a). Affectedness and direct objects: The role of lexical semantics in the acquisition of verb argument structure. *Cognition 41*, 153-195.

Gropen, J., Pinker, S., Hollander, M., & Goldberg, R. (1991b). Syntax and semantics in the acquisition of locative verbs. *Journal of Child Language 18*, 115-151.

Gropen, J., Pinker, S., Hollander, M., Goldberg, R., & Wilson, R. (1989). The learnability and acquisition of the dative alternation in English. *Language 65*, 203-257.

Hsu, A. S., & Chater, N. (2010). The logical problem of language acquisition: A probabilistic perspective. *Cognitive Science 34*, 972-1016.

Hsu, A. S., Chater, N., & Vitányi, P. (2011). The probabilistic analysis of language acquisition: Theoretical, computational, and experimental analysis. *Cognition 120*, 380-390.

Hsu, A. S., Chater, N., & Vitányi, P. (2013). Language Learning From Positive Evidence, Reconsidered: A Simplicity‑Based Approach. *Topics in Cognitive Science 5*, 35-55.

Langacker, R. (2000). A dynamic usage-based model. In *Usage-Based Models of Language*, Barlow, M. & S. Kemmer (eds.), 1–63. Stanford: CSLI Publications.

Levin, B. (1993). *English Verb Classes and Alternations: A Preliminary Investigation*. Chicago: University of Chicago press.

MacWhinney, B. (2004). A multiple process solution to the logical problem of language acquisition. *Journal of Child Language 31*, 883-914.

Noble, C. H., Rowland, C. F., & Pine, J. M. (2011). Comprehension of Argument Structure and Semantic Roles: Evidence from English‑Learning Children and the Forced‑Choice Pointing Paradigm. *Cognitive Science 35*, 963-82.

Onnis, L., Roberts, M., & Chater, N. (2002). Simplicity: A cure for overgeneralizations in language acquisition? *In Proceedings of the 24th annual conference of the Cognitive Science Society* (pp. 720–5). Mahwah, NJ: Erlbaum.

Perfors, A., Tenenbaum, J. B., & Wonnacott, E. (2010). Variability, negative evidence, and the acquisition of verb argument constructions*. *Journal of Child Language 37*, 607.

Pinker, S. (1989). *Learnability and Cognition: The Acquisition of Argument Structure.* Cambridge, MA: MIT Press.

Pinker, S. (1994). How could a child use verb syntax to learn verb semantics? *Lingua 92*, 377-410.

Ruh, N., & Westermann, G. (2009). OXlearn: A new MATLAB-based simulation tool for connectionist models. *Behavior Research Methods 41*, 1138-43.

Stefanowitsch, A. (2008). Negative evidence and preemption: A constructional approach to ungrammaticality. *Cognitive Linguistics 19*, 513-531.

Theakston, A. L. (2004). The role of entrenchment in children's and adults' performance on grammaticality judgement tasks. *Cognitive Development 19*, 15-34.

Tomasello, M. (2003). *Constructing a Language: A Usage-Based Theory of Language Acquisition.* Harvard University Press.

Twomey, K. E., Chang, F., & Ambridge, B. (2014). Do as I say, not as I do: A lexical distributional account of English locative verb class acquisition. *Cognitive Psychology 73*, 41-71.

Wonnacott, E., Newport, E. L., & Tanenhaus, M. K. (2008). Acquiring and processing verb argument structure: Distributional learning in a miniature language. *Cognitive Psychology 56*, 165-209.

Footnotes

1. Verbs such as *contribute* and *donate* are examples of such verbs (**I donated/contributed the appeal some money* vs *I donated/contributed some money to the appeal*). However, there is empirical evidence (Ambridge et al, 2012; 2014) that speakers treat such verbs as conforming to a morphophonological generalization not instantiated in the present model. Thus such verbs do not necessarily constitute exceptions to generalizations that must be learned by pre-emption alone.