Incremental understanding of conjunctive generic sentences

Michael Henry Tessler, Karen Gu, and Roger Levy

{tessler, karen, rplevy}@mit.edu

Department of Brain and Cognitive Sciences, MIT
Cambridge, MA 02139 USA

Abstract

Generic statements convey generalizations about categories, but how generic predications combine is unclear. “Elephants live in Africa and Asia” does not mean that individual elephants live on both continents. In addition, such conjunctive generics pose interesting questions for theories of incremental processing because the meaning of the sentence can change part-way through: “Elephants live in Africa” would imply most or all do, but “Africa and Asia” implies some live in each. We extend a recently proposed computational model of generic language understanding with an incremental processing mechanism that can begin to interpret an utterance before a speaker has finished their sentence. This model makes novel predictions about partial interpretations of conjunctive generic sentences, which we test in two behavioral experiments. The results support a strong view of incrementality, wherein listeners continuously update their beliefs based on expectations about where a speaker will go next with their utterance.

Keywords: semantics; pragmatics; incremental processing; generics; psycholinguistics

Introduction

Much of what we come to learn about the world comes not from direct experience but from knowledge conveyed to us from others, often in the form of linguistic utterances. “Elephants eat 300 pounds of a food in a day” succinctly conveys information extending beyond any particular moment in time or space: It could apply to any elephant, on any day of the week. Utterances that communicate generalizations are called *generic* utterances (Carlson, 1977; Carlson & Pelletier, 1995), and they are the foremost case study of rich, abstract knowledge conveyed in simple utterances (Gelman, 2009).

Generics are rife with philosophical puzzles that make it difficult to develop a unified, formal theory of their meaning (for useful reviews: Carlson & Pelletier, 1995; Nickel, 2016). One largely understudied puzzle concerns how generic predications combine. Consider the null hypothesis that generics convey information about the percentage of the category with the property—the *prevalence*—in a way analogous to how majority quantifiers (e.g., *most*, *all*) work (e.g., “Most elephants eat 300 pounds of food in day”). How can such an account treat a generic involving a conjunctive predication like “Elephants live in Africa and Asia”? No elephant actually lives on both continents; instead, the sentence should be understood as (*generically*) *elephants live in Africa* and (*generically*) *elephants live in Asia*, but this is impossible if each individual generic sentence means that the majority holds the property (i.e., it is impossible for more than half of elephants

to live in Africa and more than half of elephants to live in Asia; Nickel, 2008). The prevalence implied by a generic involving a conjunctive, mutually-exclusive predicate seems more lax than if only one of conjuncts were mentioned: If a speaker said “Elephants live in Africa”, you might think they all do.

The puzzle of understanding conjunctive generic sentences deepens when one considers that linguistic input is processed incrementally (e.g., Altmann & Kamide, 1999): Listeners ubiquitously form expectations about the intended meaning of a sentence before the speaker finishes it (e.g., Tanenhaus, Spivey-Knowlton, Eberhard, & Sedivy, 1995). For conjunctive generics about mutually exclusive properties, strongly incremental language understanding might produce non-monotonic belief updates: after the sentence prefix “Elephants live in Africa...”, a comprehender might infer a higher prevalence than after hearing the sentence completion “...and Asia”. If such non-monotonic updates occur, what types of linguistic input trigger them?

In this paper we show that a recently proposed model of generic language can accommodate these complex inferential patterns and we empirically test two predictions about generic interpretation that address these puzzles. The model of Tessler and Goodman (2019) treats generics as a kind of vague quantifier: interpretation of a generic depends on prior beliefs about properties. First, we show how when properties are likely to be mutually incompatible (as in *live in Africa and Asia*), listeners infer lower prevalences of each property following a conjunctive generic sentence than when the properties are compatible. Second, we show how the above model, when integrated with expectation-based probabilistic theories of syntactic processing (Hale, 2001; Levy, 2008), predicts that comprehenders update their beliefs about property prevalence not just when encountering a second, conjoined property, but immediately upon encountering evidence that a second, conjoined property is likely to be forthcoming. We test these predictions in two behavioral experiments that probe listeners’ understanding of conjunctive generic sentences at different points mid-sentence, analogous to gating paradigms in psycholinguistics (Grosjean, 1980). Our empirical data confirm both predictions, suggesting that generic language interpretation interacts jointly with world knowledge and strongly incremental syntactic processing according to principles of probabilistic inference under uncertainty.

Figure 1: Model’s sequential interpretation of “Elephants live in Africa... and Asia.” A: Correlated priors reflected in the joint probability distribution over two features result in a mutual exclusivity inference. When the model hears only “... live in Africa”, it believes that probably all live in Africa (middle facet); when it hears they live in Asia as well, the model non-monotonically updates its beliefs about how many live in Africa. B: The mutual exclusivity inference holds for priors of different shapes and never holds if the prior knowledge about the two features is uncorrelated. Points show means of distributions. U-shaped priors were Beta(0.1, 1); multimodal priors were an equal mixture of a Beta(1, 100) and a Beta(25, 1). Correlated priors were created by adding an additional factor that decreased the probability of a prevalence-level if the sum of the prevalence of the two features exceeded 100%.

Computational Model

We extend a model for interpreting generics to incorporate an incremental processing mechanism that allows a listener to understand partial utterances. The original model of Tessler and Goodman (2019) interprets a generic utterance predicating a property of a category (“Elephants eat 300 pounds of food in a day”) as meaning that the prevalence (or probability) x of the property given the category— $P(\text{eats 300 lb. of food in a day} | \text{is an elephant})$ —is greater than an *a priori* uncertain threshold θ . The literal meaning of the generic—an uncertain threshold function, with uniform uncertainty over the threshold $P(\theta)$ —combines with a listener’s prior knowledge about the prevalence of the feature $P(x)$ within a relevant set of alternative categories (e.g., other animals) to compute a posterior distribution over prevalence x :

$$P(x | u) = \int_{\theta} P(x, \theta | u) d\theta \propto P(x) \cdot P(\theta) \cdot \delta[\llbracket u \rrbracket_{(x, \theta)}] \quad (1)$$

where $\delta[\llbracket u \rrbracket_{(x, \theta)}]$ is the Kronecker delta function assigning a value of 1 for utterances that are literally true (in the case of a generic: where $x > \theta$) and 0 for utterances that are false.

To interpret a generic with a conjunctive predicate such as “Elephants live in Africa and Asia”, we assume the semantic representation contains a conjunction of two generic statements: $[Gen(\text{elephant})(\text{live_in_Africa})] \wedge [Gen(\text{elephant})(\text{live_in_Asia})]$, where the *Gen* operator acts according to the belief-updating rule of Eq. 1 (see Nickel (2008) for supporting arguments of this semantic parse). A listener starts with a joint prior over the prevalence of the two properties (we denote variables associated with *living in Africa* with subscript r and *Asia* with s): $P(\mathbf{x}) = P(x_r, x_s)$,

which is incrementally updated with each successive generic. The model can then interpret multiple generics in succession, using the posterior distribution over prevalence $P(\mathbf{x} | u)$ (Eq. 1) as the prior for the next utterance.

$$P(\mathbf{x} | u_r, u_s) \propto \int_{\theta_s} \int_{\theta_r} P(\mathbf{x}, \theta | u_r) \cdot \delta[\llbracket u_s \rrbracket_{(x_s, \theta_s)}] d\theta_r d\theta_s \quad (2)$$

where $P(\mathbf{x}, \theta | u_r)$ is the posterior that results from hearing “Elephants live in Africa” given by Eq. 1.

The predictions for a sequential understanding of “Elephants live in Africa and Asia” are shown in Fig. 1. Upon hearing the first part of the utterance, the model believes that almost all elephants live in Africa (simulations assuming a uniform prior shown in Fig. 1A). What happens next depends upon the correlational structure of the prevalence prior: If the listener has prior knowledge suggesting the properties (living in Africa, living in Asia) are mutually exclusive (Fig. 1A top), they interpret the next part of the utterance (“...and Asia”) as indicating that some (perhaps half) of elephants live in Africa and other ones live in Asia. Without this correlation in the prior, the model ends up believing that most or all elephants live both in Africa and in Asia (Fig. 1A bottom). These inferences are robust to a variety of different prevalence prior distributions, so long as the prior has the necessary correlational structure (Fig. 1B shows predictions assuming a uniform, U-shape Beta, and mixture-of-Beta distributions).

When processing a conjunctive phrase, listeners may form expectations about the complete utterance even before the sentence is over. For example, when a speaker reaches the word *Africa* in “Elephants live in *Africa*”, she has many syntactically distinct options available to her to complete the sen-

Figure 2: Incremental parse trees and syntactic expectations for upcoming conjunct properties in generic predication. The string prefix “Elephants live in Africa...” is compatible with a variety of continuations, including the four listed above. The next word, “and”, rules out the continuations in the top row (depicted in gray) and sharpens expectations around a conjunct at potentially different structural levels (light green). Probabilistic renormalization implies that an upcoming conjunct mutually exclusive with the first conjunct becomes more likely when “and” is encountered, driving the strong incremental predictions depicted in Fig. 3.

tence (Fig. 2 shows four possibilities). One such possibility is that she continues with a NP-coordination that includes a mutually exclusive property (e.g., *and Asia*; bottom-left tree). When the listener encounters the word *and* in “Elephants live in Africa *and*”, he knows he is entering into a coordination and the relative probability of a forthcoming mutually exclusive property increases. Such a continuation would yield a different inference about the prevalence of elephants in Africa than a continuation with a non-mutually exclusive property (e.g., with a verb phrase such as “and eat bugs”).¹ If listeners parse and interpret utterances incrementally at the level of individual words, then we would expect their inferences about the prevalence of elephants in Africa to represent a mixture of the inferences derived from different possible continuations, which can be represented by conditional probabilities of the full utterance u' given the sentence fragment heard f :

$$P(x | f) = \sum_{u'} P(x | u') P(u' | f) \quad (3)$$

If, however, listeners do not derive incremental interpreta-

¹For illustrative purposes, we assume a correlation between NP vs. VP coordination and mutually exclusive vs. non-mutually exclusive predicates. Of course, it is possible to continue with a verb phrase about a mutually exclusive property such as “...live in Africa and live in Asia” as well as continue with a noun phrase about a non-mutually exclusive property (e.g., “...eat figs and nuts”). The crucial fact is that the probability of a forthcoming mutually-exclusive property increases when the comprehender encounters the word *and*.

Figure 3: A model that incorporates syntactic expectations at the level of individual words (*strong incremental*) predicts intermediate mutual-exclusivity inferences part-way through the conjunction (at “and”), whereas a model that waits for content-words (*weak incremental*) does not show a difference in expected prevalence at the word “and.”

tions at each moment, but instead wait for meaningful pieces of an utterance (e.g., content words like *Asia*) to compute interpretations, then we would not expect such an intermediate degree of interpretation: “Elephants live in Africa and...” should mean the same thing as “Elephants live in Africa...” (Fig. 3). We test this prediction in Expt. 2 using a gating paradigm in the spirit of Grosjean (1980).

Experiments

We design two experiments to test the mutual exclusivity (ME) and incremental predictions. Expt. 1 tests the ME prediction that “Elephants live in Africa and Asia” means roughly that half live in Africa and half live in Asia; this experiment also serves to validate the gating procedure we employ in the second experiment. Expt. 2 is a pre-registered study that uses the gating paradigm to test the fine-grained incremental predictions of the model. The experiments and a full list of materials can be viewed at tinyurl.com/elephants-cogsci.

Experiment 1: Mutual exclusivity inference

Participants We recruited 27 participants through Amazon’s Mechanical Turk. Participants were restricted to those with verified U.S. IP addresses and at least a 95% work approval rating. The study took about 10 minutes and participants were compensated \$1.50.

Materials Participants read a storybook with chapters about creatures on a faraway planet. Each chapter contained a short paragraph presented across 2–4 screens, with a button to “turn the page” (Fig. 4A). A chapter introduced one or a few novel categories (e.g., *wugs*) and semi-novel properties (e.g., *live on the continent of Caro*). Critical trial chapters ended with a generic sentence about conjunctive properties, which differed only in whether the second property was mutually exclusive with the first (*conjunct type*): “Glippets live on the continent of Caro and *on the continent of Este* (ME) /

<p>Chapter 3: Wugs</p> <p>Wugs are large creatures, quite intelligent, with</p> <p>1 of 3</p>	<p>a lifespan of about sixty years. They live in Africa and</p> <p>2 of 3</p>	<p>What percentage of wugs do you think live in Africa?</p> <p>0% 100%</p> <p>What percentage of wugs do you think live in Asia?</p> <p>0% 100%</p> <p>[-?-] 2 of 3</p>	<p>eat bugs.</p> <p>3 of 3</p>
--	---	--	--------------------------------

B		Experiment 1	[-?-]
INTERRUPTED	ME	[Wugs...] [Wugs live in Africa] [-?-] [and eat bugs.]	1. % AFRICA 2. % ASIA
INTERRUPTED	NME	[Wugs...] [Wugs live in Africa] [-?-] [and drink water.]	1. % AFRICA 2. % EAT BUGS
UNINTERRUPTED	ME	[Wugs...] [Wugs live in Africa and Asia.] [-?-]	1. % AFRICA 2. % ASIA
UNINTERRUPTED	NME	[Wugs...] [Wugs live in Africa and eat bugs.] [-?-]	1. % AFRICA 2. % EAT BUGS
		Experiment 2	
INTERRUPTED	A	[Wugs...] [Wugs live in Africa] [-?-] [and eat bugs.]	1. % AFRICA 2. % OTHER CONTINENT
INTERRUPTED	A&	[Wugs...] [Wugs live in Africa and] [-?-] [eat bugs.]	1. % AFRICA 2. % OTHER CONTINENT
INTERRUPTED	A&B	[Wugs...] [Wugs live in Africa and Asia] [-?-] [which are warm.]	1. % AFRICA 2. % OTHER CONTINENT
UNINTERRUPTED	A&B	[Wugs...] [Wugs live in Africa and Asia.] [-?-]	1. % AFRICA 2. % OTHER CONTINENT

Figure 4: Overview of experiments. A: Example book chapter from Expt. 2, depicting the *Interrupted* A& condition. “Africa and Asia” property is shown for illustration; actual stimuli used novel names for properties (“Caro and Este”). B: Overview of conditions for Expts. 1 and 2. [-?-] denotes point in the sentence at which the question appeared. Highlighting shows which properties were mentioned before the question, and what was asked about. See main text for full description of conditions.

enjoy the sunshine there (not ME: NME).”

The earlier content of the chapter supported the mutually exclusive interpretation of the properties when the properties were not *ipso facto* mutually exclusive. For example:

Krens are a tribe of the aliens that live on the continent of Benli. Animals like stups, four-legged creatures with large antlers, are a resource for the Krens. Stups roam all over the windy highlands of Benli, far from the oceans. Krens are stup-herders and (fishermen / incorporate stups into their religion).

Conjunct type (ME vs. NME) was manipulated within participants and items. There were 14 filler chapters with content similar to the critical chapters but using explicit quantifiers (*most*, *all*, *none*) to describe the properties of categories.

Procedure Participants were told they would be reading a storybook with a question in each chapter. Questions were all of the same type, an *implied prevalence* question (Gelman, Star, & Flukes, 2002; Cimpian, Brandone, & Gelman, 2010): “What percentage of Ks do you think F?”, where K represents a category and F a feature. Responses were recorded using a slider with endpoints labeled 0% and 100%, with the exact value selected appearing above the slider. Participants were familiarized with the response variable in a practice trial, where they were asked to report how many *dogs bark*, *birds are male*, *cats get cancer*, and *lions lay eggs*. These questions encouraged participants to use the full range of the response scale as well as served as a comprehension check.

In each critical chapter of the storybook, two questions appeared either at the end of the chapter (*Uninterrupted* condi-

tions) or interrupting the chapter right before the final page (*Interrupted* conditions; Fig. 4). In the *Interrupted* conditions, the question came in the middle of a conjunctive generic sentence, but before the conjunction so the reader was unaware the sentence would continue with a conjunction. The questions asked about the mentioned property (e.g., *Africa*) and either a mutually exclusive property (*Asia*; ME conditions) or a nonmutually exclusive property (e.g., *eats bugs*; NME conditions); the chapter then concluded with a conjunction about an unmentioned, nonmutually exclusive property (Fig. 4B), so as to not give the impression that the participant was being tricked by being asked about a property that we would eventually reveal. In the critical trials, the second question was asked about the second property mentioned (ME vs. NME). Filler trials asked about two properties described in the chapter using quantifiers (i.e., *all*, *most*, or *none*). The order in which the two questions appeared on the screen was randomized on each trial.

Each participant read a total of 21 chapters, which included 8 ME conjunctions, 4 NME conjunctions, and 6 quantifier fillers; for each of these categories, equal numbers of interrupted and uninterrupted were used. The experiment began with a chapter with no questions and 2 filler trials; the remaining trials were presented in a random order such that no two critical trials were presented back-to-back. Subjectively, the task is very difficult as participants learn about many different animals with lots of new names; in practice, however, participants only need to recall information from the previously en-

countered sentence to answer the trial questions. Following the storybook, participants completed a memory check where they had to select all the facts they had learned from a list of 10 (5 real, 5 distractor); in addition, participants were asked to explain what the experiment was about in broad terms.

Results 11 participants were excluded for failing to respond accurately to all of the practice trials or failing to respond accurately to at least 7 of the 10 memory check prompts (same exclusion criteria for Expt. 2). We describe the results using the running example of “Elephants live in Africa and Asia”, but the experimental stimuli used novel categories and relatively novel properties.

Figure 5: Experiment 1 results. Participants rate prevalence for mentioned property (% live in Africa) and either the mutually exclusive property (left facet) or non-mutually exclusive property (right facet), mid sentence (“Africa”) or after the sentence finishes (“Africa and X”). Error-bars denote bootstrapped 95% confidence intervals.

The results are visually apparent in Fig. 5. Reading only “Elephants live in Africa...” led listeners to believe, on average, that all elephants lived in Africa and that none lived in Asia, whereas if the sentence finished “...and Asia”, listeners inferred that roughly half live in Africa and half live in Asia. This inference is not categorical, however; there are a number of responses to ME conjunctive generics wherein participants infer that all or almost all have both properties (recall that many of our items are unfamiliar properties). A different pattern was observed for the NME properties, where hearing about the second property (“eat bugs”) only increased participants’ degree of belief in each property applying. Further, when answering about unmentioned properties, participants rated the prevalence of ME properties close to 0% whereas NME properties were rated as somewhat prevalent (green bars, “Africa”). The results replicate intuitions about how “Elephants live in Africa” should be interpreted in a context where the sentence is interrupted. The comparison with the non-mutually exclusive condition shows that the results cannot be attributed to the very act of being asked about two properties mentioned in a conjunctive generic sentence.

Experiment 2: Strong incrementality

In Expt. 1, we demonstrated that the mutually exclusive inference effects can be measured using a gating paradigm wherein participants are queried for their beliefs part-way through a sentence. Here, we exploit this paradigm to test the strong incremental processing predictions of the model, where syntactic expectations can modulate the interpretations of generic sentences in a fine-grained manner. Sample size, participant exclusion criteria, and analyses for this experiment were pre-registered on OSF osf.io/pjt9c.

Participants We recruited 108 participants through Amazon’s Mechanical Turk. Participants were restricted to those with verified U.S. IP addresses and at least a 95% work approval rating. The study took about 10 minutes and participants were compensated \$1.50.

Materials and procedure The materials and procedure followed those of Expt. 1 with the following exceptions. We modified the critical conjunctive generics to primarily involve the conjunction of two noun phrases (e.g., *ascribe to Cabooism and Daithism*) in order to strength the correlation between the NP-conjunction and mutual exclusivity.² The fillers were modified to introduce page breaks immediately before and immediately after conjunctions (“and”) in order to raise participants’ expectations that a sentence might be broken at a conjunction. We used additional examples of the Uninterrupted ME condition of Expt. 1 (“live in Africa and Asia.”) as fillers to raise participants’ expectations about ME continuations. Half of the filler trials had page breaks immediately before the “and” and half immediately after.

On critical trials, questions always interrupted the chapter right before the last page. On the question screen, the page number of the penultimate page remained on the screen to provide an additional cue that the chapter was not complete (Fig. 4A). There were three conditions corresponding to the point in the sentence at which the page break and prevalence questions occurred: “Elephants live in Africa__ and__ Asia__” (where __ denotes the page-break). In the two conditions where participants did not see the full conjunctive property before the question (INT A and A&), the sentence continued with a non-mutually exclusive property (e.g., *eat bugs*).

Finally, we changed the question about the second property (% live in Asia) to ask about “some other X”, where X was the kind of property that was mentioned in the first conjunct (e.g., *live on some other continent*). This change was introduced to raise the plausibility that a second, ME property was possible while not naming one explicitly, which would be pragmatically odd given that the property is unmentioned in the INT A and A& conditions. Participants saw 18 chapters. The story started with a chapter with no questions, then participants saw 4 fillers: 2 quantifiers with interrupting questions

²Of the 13 items in this experiment, 9 of them were NP-coordinated (the others used coordination of prepositional phrases and adjectives). In both experiments, critical conjunctive generics always involved conjunctions of the same syntactic types (e.g., *ascribe to the Caboo religion and the Daith religion*).

and 2 uninterrupted ME fillers, in a random order. Finally, participants saw 2 of each kind of critical trial with 4 uninterrupted ME fillers and 3 quantifier fillers interleaved to avoid back-to-back critical trials.

Figure 6: Experiment 2 results. Participants are interrupted at various stages of the sentence (after *Africa*, *and*, or *Asia*) to be asked about the prevalence of *living in Africa* and *living in some other place*, or asked at the end of the sentence (right-most bars). When participants are interrupted before the second conjunct (*Asia*), the sentence continues with a non-mutually exclusive property. Error-bars denote bootstrapped 95% confidence intervals.

Results 28 participants were excluded for failing at least one of the two attention checks.³ To test our main prediction, we constructed a Bayesian mixed-effects regression model predicting implied prevalence ratings for the first conjunct (e.g., % Africa) as a function of the point in the sentence in which participants were queried. We included by-item and by-participant random intercepts and slopes.⁴ The regression model was created in Stan (<http://mc-stan.org/>) accessed with the *brms* package using default priors (Bürkner, 2017).

Replicating the findings of Expt. 1, when participants only read that “Elephants live in Africa”, they tended to infer that almost all lived in Africa. When they read that “Africa and Asia”, they tended to infer that the distribution was close to 50%-50%. Finally, as predicted by a strong version of incremental processing, participants began to anticipate a mutually-exclusive conjunct when only the word “and” was mentioned, as evidenced by their implied prevalence ratings being substantially less for the “...live in Africa and...” condition than the “live in Africa” condition (posterior mean estimate and 95% credible interval: $\beta = -0.08$ (-0.13, -0.04)). In addition, these ratings were substantially higher than when the full conjunctive predicate was present “...live in Africa and Asia” ($\beta = 0.17$ (0.12, 0.23); Fig. 6). Thus, we find that participants’ implied prevalence ratings of how many elephants

live in Africa monotonically decreased as a function of how many words of the conjunctive predicate they were allowed to see. These results suggest that listeners begin to draw pragmatic interpretations of generics before the end of the sentence and even in the absence of additional content words.

It is notable that in the “Africa and Asia” conditions, participants on average infer greater than 50% prevalence for *Africa* and lower than 50% for *Asia*, a departure from the results of Expt. 1. This deviation may be due to participants forgetting what they have read and/or not making the inference that the second conjunct stands in a subset relation to the category in the second question (e.g., that *Asia* is a kind of “some other continent”). Explicitly asking about the conjuncts alleviates memory demands by allowing participants to merely recognize, rather than recall, that they have seen this conjunct mentioned. Asking about *some other continent* (Expt. 2) requires participants’ to recall the second conjunct and could lead to lower prevalence ratings in response to this question.

Discussion

Generic sentences exhibit extreme sensitivity to context that make it difficult to precisely define what a single generic conveys. “Elephants live in Africa and Asia” means neither that *most elephants live in (both) Africa and Asia* nor that *most elephants live in Africa, and most live in Asia*. Here, we empirically measured interpretations of generics about conjunctive predicates, building on the observation of Nickel (2008) of the range of troubling examples for quantificational views of generics. Notably, the uncertain threshold model of Tessler and Goodman (2019) accounts for such conjunctive generics seamlessly: An underspecified threshold can be updated as more information comes in and is sensitive to prior beliefs regarding compatibility of the conjunct properties.

We extended that model to include syntactic expectations and found evidence for the strongest form of incremental syntactic processing, wherein beliefs are continually updated based on expectations of how a sentence will continue. The fact that generic language understanding can be modulated simultaneously by correlations in background knowledge and by syntactic expectations calls for a tighter coupling between models of syntactic processing (Levy, 2008), pragmatic language understanding (Goodman & Frank, 2016), and intuitive theories (Tenenbaum, Kemp, Griffiths, & Goodman, 2011).

It remains an open question, however, how specific the effects observed in this paper are to generics rather than to quantification more generally. For example, it appears that, in some contexts, one can use *most* to convey similar mutually exclusive conjunctions: “Elephants are the largest land animal on Earth and are one of the gentlest creatures. Most live in Africa and Asia but are brought to other places for the entertainment of humans.”⁵ Further work is needed to determine the felicity and interpretation of such utterances.

Data, code, and links to experiments are available at <https://github.com/mhtess/elephants>

³6 failed slider check; 9 failed memory check; 13 failed both.

⁴model: rating ~ cond + (1 + cond | subj) + (1 + cond | item)

⁵Example from theodysseyonline.com/want-to-ride-an-elephant

Acknowledgments

This work was supported by the National Science Foundation Grant #1456081, Elemental Cognition, and the MIT Sense-time Alliance and Quest for Intelligence.

References

- Altmann, G. T. M., & Kamide, Y. (1999). Incremental interpretation at verbs: Restricting the domain of subsequent reference. *Cognition*, 73, 247–264. doi: 10.1016/S0010-0277(99)00059-1
- Bürkner, P.-C. (2017). brms: An r package for bayesian multilevel models using stan. *Journal of Statistical Software*, 80(1), 1–28.
- Carlson, G. N. (1977). *Reference to kinds in English*. Unpublished doctoral dissertation, University of Massachusetts, Amherst.
- Carlson, G. N., & Pelletier, F. J. (1995). *The Generic Book*. Chicago, IL: Chicago University Press.
- Cimpian, A., Brandone, A. C., & Gelman, S. A. (2010). Generic statements require little evidence for acceptance but have powerful implications. *Cognitive Science*, 34(8), 1452–1482.
- Gelman, S. A. (2009). Learning from others: Children’s construction of concepts. *Annual review of psychology*, 60, 115–140. doi: 10.1146/annurev.psych.59.103006.093659.LEARNING
- Gelman, S. A., Star, J. R., & Flukes, J. E. (2002). Children’s Use of Generics in Inductive Inferences. *Journal of Cognition and Development*, 3(2), 179–199.
- Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*.
- Grosjean, F. (1980). Spoken word recognition processes and the gating paradigm. *Perception & Psychophysics*, 28(4), 267–283. doi: 10.3758/BF03204386
- Hale, J. (2001, 2–7 June). A probabilistic Earley parser as a psycholinguistic model. In *Proceedings of the second meeting of the north american chapter of the Association for Computational Linguistics* (pp. 159–166). Pittsburgh, Pennsylvania. Retrieved from <http://acl.ldc.upenn.edu/N/N01/N01-1021.pdf>
- Levy, R. (2008). Expectation-based syntactic comprehension. *Cognition*, 106, 1126–1177. doi: 10.1016/j.cognition.2007.05.006
- Nickel, B. (2008). Generics and the ways of normality. *Linguistics and Philosophy*, 31(6), 629–648. doi: 10.1007/s10988-008-9049-7
- Nickel, B. (2016). *Between logic and the world*. Oxford: Oxford University Press.
- Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K. M., & Sedivy, J. C. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science*, 268(5217), 1632–1634.
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, 331(6022), 1279–1285.
- Tessler, M. H., & Goodman, N. D. (2019). The language of generalization. *Psychological Review*, 126(3), 395–436.