Pragmatic Reasoning and Semantic Convention: A Case Study on Gradable Adjectives

Ming Xiang, Chris Kennedy, Weijie Xu, and Timothy Leffel
University of Chicago

Scales, degrees and implicature:
Novel synergies between semantics and pragmatics
University of Potsdam
28 May, 2021

Threshold semantics for gradable adjectives

(1) $[\![\text{long}]\!] = \lambda x.\textbf{length}(x) \geqslant \theta$

Threshold semantics for gradable adjectives

(1) $[\![\text{long}]\!] = \lambda x.\mathbf{length}(x) \geqslant \theta$

(2) "Compositional" θ

a. That pole is two meters long.

b. That pole is longer than this knife.

c. That pole is too long to fit in the rack.

Threshold semantics for gradable adjectives

(1) $[\![\text{long}]\!] = \lambda x.\textbf{length}(x) \geqslant \theta$

(2) "Compositional" θ

a. That pole is two meters long.

b. That pole is longer than this knife.

c. That pole is too long to fit in the rack.

(3) "Contextual" θ

a. That pole is long.

b. That knife is long.

c. That rope is long.

Threshold uncertainty

(4) a. That pole is made of local$_{to\ place}$ bamboo.
b. That pole is ready$_{for\ purpose}$.
c. That pole isn't$_{property}$.

Threshold uncertainty

(4) a. That pole is made of local$_{to\ place}$ bamboo.
b. That pole is ready$_{for\ purpose}$.
c. That pole isn't$_{property}$.

(5) a. That pole is long$_{\theta}$.
b. That pole is expensive$_{\theta}$.
c. That pole is heavy$_{\theta}$.

Threshold uncertainty and adjective class

(6) **Relative:** $\theta \approx mid+$

a. That pole is long.

b. That pole is expensive.

c. That pole is heavy.

Threshold uncertainty and adjective class

(6) **Relative:** $\theta \approx \mathit{mid}+$

a. That pole is long.

b. That pole is expensive.

c. That pole is heavy.

(7) **Absolute maximum:** $\theta \approx \mathit{max}$

a. That pole is straight.

b. That theater is empty.

c. That countertop is flat.

Threshold uncertainty and adjective class

(6) **Relative:** $\theta \approx \mathit{mid}+$

a. That pole is long.

b. That pole is expensive.

c. That pole is heavy.

(7) **Absolute maximum:** $\theta \approx \mathit{max}$

a. That pole is straight.

b. That theater is empty.

c. That countertop is flat.

(8) **Absolute minimum:** $\theta \approx \mathit{min}$

a. That pole is bent.

b. That door is open.

c. That shirt is striped.

Two theories of thresholds

Semantic accounts

Thresholds are associated with scalar endpoints, when such endpoints are available, as a matter of semantic convention. (Paradis 2001; Kennedy and McNally 2005; Rotstein and Winter 2004; Kennedy 2007; Toledo and Sassoon 2011; Burnett 2016; etc.)

Two theories of thresholds

Semantic accounts

Thresholds are associated with scalar endpoints, when such endpoints are available, as a matter of semantic convention. (Paradis 2001; Kennedy and McNally 2005; Rotstein and Winter 2004; Kennedy 2007; Toledo and Sassoon 2011; Burnett 2016; etc.)

Bayesian pragmatic accounts

LG Thresholds are inherently uncertain, but both thresholds and utterance content are a function of iterated pragmatic reasoning taking into account cost, informativity and prior distribution of a comparison class (Lassiter and Goodman, 2014)

QF Like LG, except that threshold posteriors eventually evolve into linguistic conventions (Qing and Franke, 2014)

Two questions

(S) What are the truth conditions of utterances of (9a-c) in particular contexts?

(P) What is the information communicated by utterances of (9a-c) in particular contexts?

(9) a. That pole is long.
  b. That pole is straight.
  c. That pole is bent.

Two kinds of answers

Semantic accounts

- Provide answers to (S)
- Don't say much about (P), and what they do say appears to be wrong

Two kinds of answers

Semantic accounts

- Provide answers to (S)
- Don't say much about (P), and what they do say appears to be wrong

Bayesian pragmatic accounts

- Provide answers to (P)
- Don't directly answer (S), but can do so if supplemented by appropriate linking hypothesis

This study

1. Elicit truth value judgments and posterior degree judgments from human subjects, and use these data to evaluate the predictions of the different approaches.
2. Elicit empirical priors from human subjects and use these to compute the model predictions, instead of artificial priors.

Experiment 1: Degree priors

Goal: establish a prior probability for distribution of objects in a comparison class are relative to different scalar dimensions

Constraint: should be independent of facts about language users' experience with words used to talk about these dimensions

Method: use restricted judgment of likelihood as a proxy for prior degree probability: *"Which of these objects is most likely?"*

Reality check: two categories of objects, *artifacts* and *shapes*

(inspiration: Foppolo and Panzeri, 2011)

Image sets

24 artifacts; 24 shapes

Artifacts

Scalar dimension: HEIGHT

Scalar dimension: BEND

Shapes

Scalar dimension: HEIGHT

Scalar dimension: BEND

Experiment 1: Example stimuli

97 participants

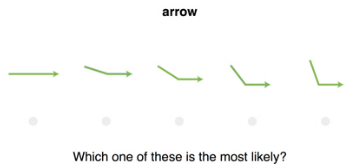

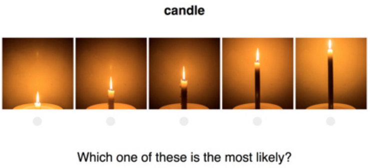

Experiment 1: Results

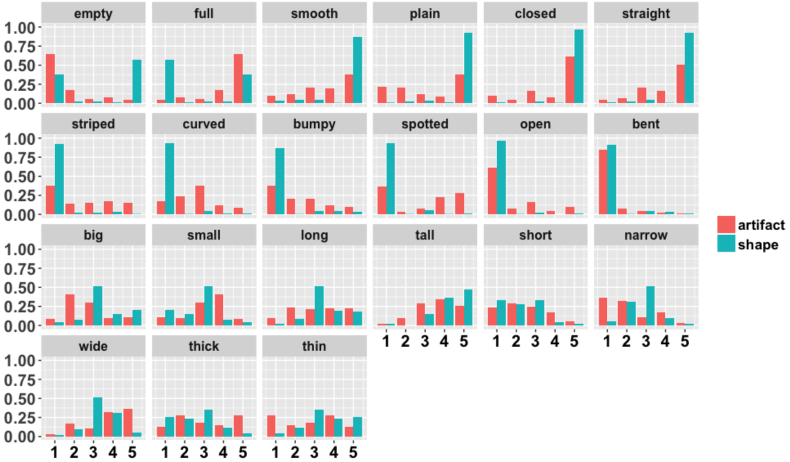

Experiment 2: Truth value judgments

Three **adjective classes**:

- **Relative**: *big*, *small*, *long*, *tall*, *short*,* *narrow*, *wide*, *thick*, *thin*
- **Maximum**: *straight*, *closed*, *plain*, *smooth*, *empty*, *full*
- **Minimum**: *curved*, *open*, *striped*, *spotted*, *bent*, *bumpy*

Two **image types**:

- **Shapes**
- **Artifacts**

Two groups of 58 participants for each image type, who each saw 48 items (24 image sets x 2 adjectives)

Experiment 2: Sample stimuli

"Please click on the checkbox beneath the image or images that the sentence appropriately describes."

It is striped

It is striped

Experiment 2: Results

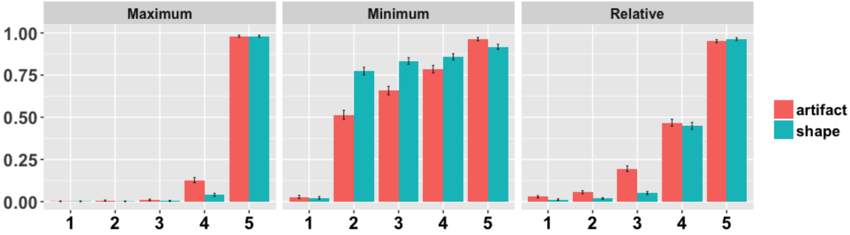

- Overall effect of adjective class as reported in literature
- Overall effect of image type
- Interaction of image type and adjective class

For statistical analysis fit Bayesian hierarchical models to data using R package "brms" (Bürkner, 2017))

Experiment 3: Posterior degree estimations

Same images as Experiments 1 and 2; new procedure:

- **Speaker:** *"I saw a candle. It was tall."*
- **Prompt:** *"Make a guess: how tall was the candle that the speaker saw?"*

Participants allowed to choose exactly one item from same item sets used in Experiments 1 and 2.

Experiment 3: Sample stimuli

Speaker: I saw a cylinder. It was tall.

Make a guess: How tall was the cylinder that the speaker saw?
Choose one image from below.

Speaker: I saw a nail. It was straight.

Make a guess: How straight was the nail that the speaker saw?
Choose one image from below.

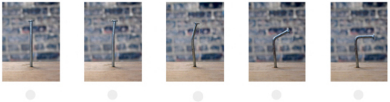

Experiment 3: Results

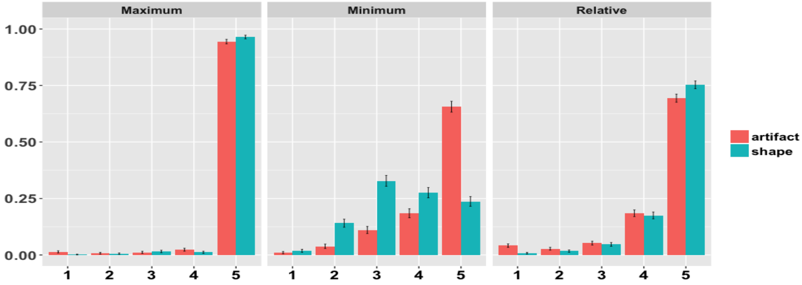

Same statistical analysis as Experiment 2, but results merely descriptive, since we don't have any initial hypotheses.

Model predictions: Procedure

Derive LG and QF model predictions for posterior object degree and posterior θ for each position on our 5-point scales:

"The candle is tall" $= \mathbf{height}(\mathbf{c}) \geqslant \theta$

- Use empirical priors collected in Experiment 1
- Assume uniform prior θ (LG)
- Set free parameters to λ=3 and *cost* = 2 (both models)

Degree posteriors: Models = humans

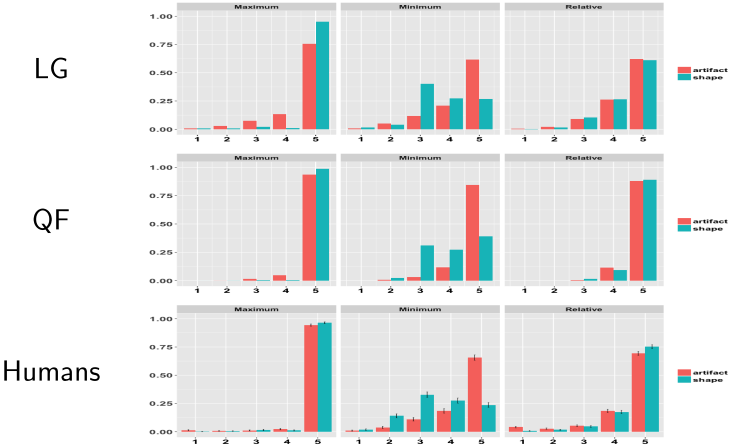

Linking θ posteriors to truth value judgments

As noted already, the Bayesian models do not directly predict truth value judgments.

But they do make predictions about θ posteriors, which can be combined with the following linking hypothesis to generate predictions about truth value judgments:

(10) $P(\text{“}x_i \textit{ is adj}\text{” is TRUE}) = P(\theta \leqslant i) = \sum_{\theta \leqslant i} P(\theta)$

θ posteriors

LG

QF

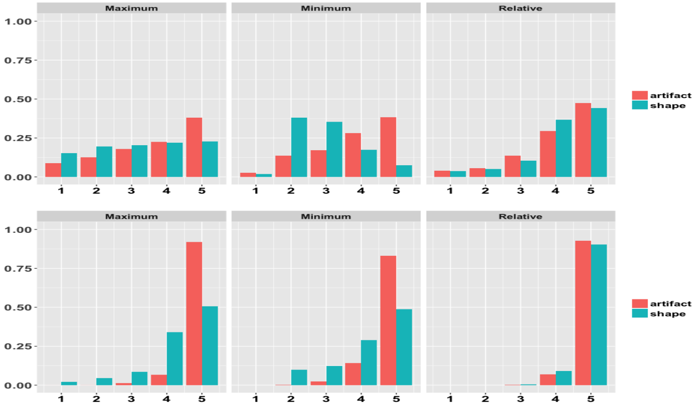

Truth value judgments: Models $\neq$ humans

LG

QF

Humans

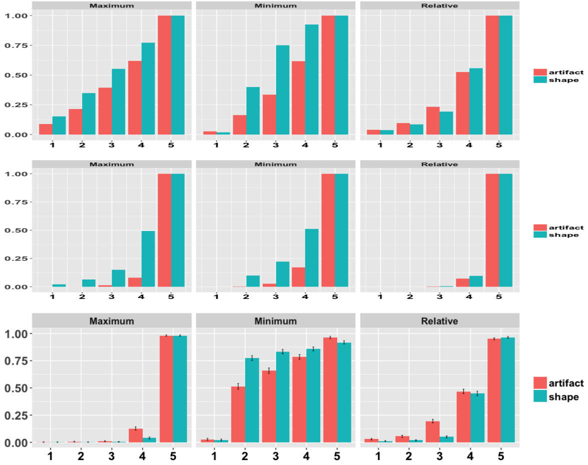

Discussion

- Bayesian models largely align human judgments about posterior degree estimations
- Bayesian models do not align with human truth value judgments, in particular for absolute adjectives
- Three possible explanations:
  1. Problem with our method for eliciting priors
  2. Problem with our parameter settings ($\lambda = 3$, $cost = 2$)
  3. Problem with model predictions about θ posteriors

Discussion

- Bayesian models largely align human judgments about posterior degree estimations
- Bayesian models do not align with human truth value judgments, in particular for absolute adjectives
- Three possible explanations:
  1. ~~Problem with our method for eliciting priors~~
  2. Problem with our parameter settings ($\lambda = 3$, $cost = 2$)
  3. Problem with model predictions about θ posteriors

Discussion

- Bayesian models largely align human judgments about posterior degree estimations
- Bayesian models do not align with human truth value judgments, in particular for absolute adjectives
- Three possible explanations:
  1. ~~Problem with our method for eliciting priors~~
  2. ~~Problem with our parameter settings ($\lambda = 3$, $cost = 2$)~~
  3. Problem with model predictions about θ posteriors

Discussion

- Bayesian models largely align human judgments about posterior degree estimations
- Bayesian models do not align with human truth value judgments, in particular for absolute adjectives
- Three possible explanations:
  1. ~~Problem with our method for eliciting priors~~
  2. ~~Problem with our parameter settings ($\lambda = 3$, $cost = 2$)~~
  3. Problem with model predictions about θ posteriors

Model fit by adjective class ($\lambda = 3$, $cost = 2$)

Evaluation of model fit corroborates qualitative assessment:

	LG (R^2)		QF (R^2)	
	TVJ	POSTERIOR	TVJ	POSTERIOR
Overall	0.65	0.84	0.6	0.88
Maximum	0.63	0.94	0.85	0.98
Minimum	0.7	0.67	0.39	0.71
Relative	0.91	0.81	0.81	0.87

- Both models less good on TVJs than degree posteriors
- Problems with TVJs driven by absolute adjectives

Reverse engineering θ distributions

What kind of θ distributions *would* derive empirical truth value judgments? We can use our linking hypothesis to recover them from human TVJs.

(10) $P(\text{“}x_i \textit{ is adj”} \text{ is TRUE}) = P(\theta \leqslant i) = \sum_{\theta \leqslant i} P(\theta)$

TVJs

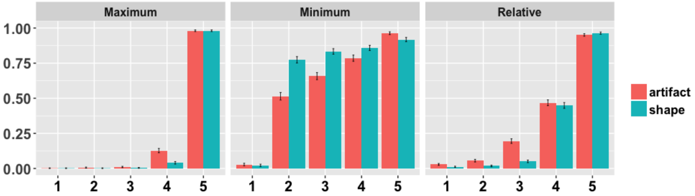

Reverse engineering θ distributions

What kind of θ distributions *would* derive empirical truth value judgments? We can use our linking hypothesis to recover them from human TVJs.

(10) $P(\text{“}x_i \textit{ is adj}\text{” is TRUE}) = P(\theta \leqslant i) = \sum_{\theta \leqslant i} P(\theta)$

θs

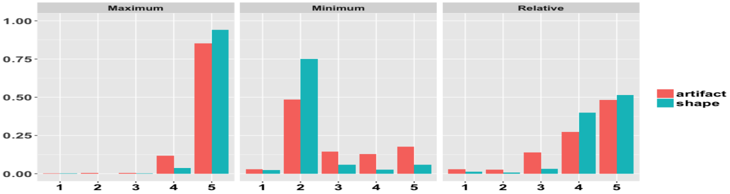

These are “semantic” θs!

"Semantic" vs. "pragmatic" θs

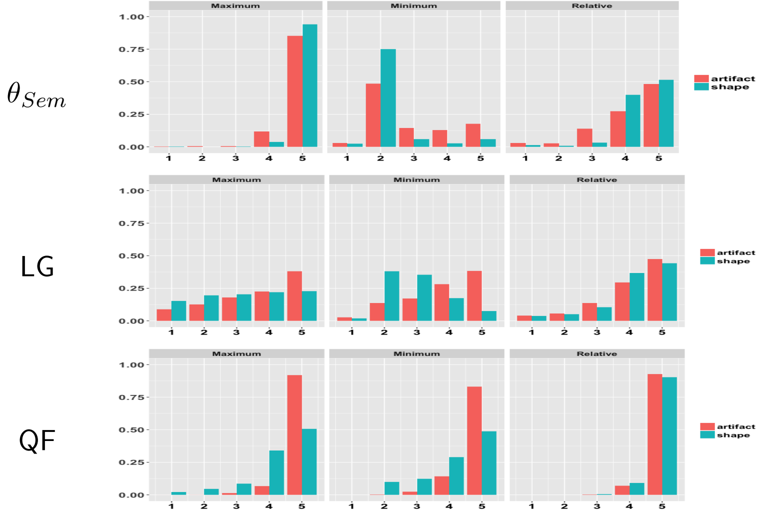

Comparison of poster degree predictions (QF only)

$\mathrm{QF}/\theta_{Prag}$

QF/θ_{Sem}

Humans

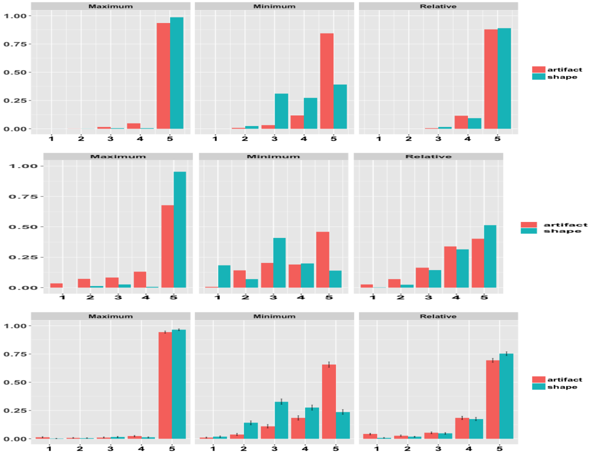

Conclusion

Results do not indicate that absolute adjectives require a "two stage" analysis (semantic thresholds; pragmatic content): empirical priors needed to explain different TVJs for shapes and artifacts.

Instead, results suggest that a single pragmatic model can account both for communicated content (posterior degree estimations) and TVJs (threshold estimations) as long as it is supplemented with reasoning based on lexical/conceptual content (scale structure).

This information is available in the input and attended to by children (Syrett, 2007), so why not give it to the robots too?

Burnett, H. (2016). *Gradability in Natural Language: Logical and Grammatical Foundations*. Oxford Studies in Semantics and Pragmatics. Oxford University Press, Oxford, UK.

Bürkner, P.-C. (2017). brms: An r package for bayesian multilevel models using stan. *Journal of Statistical Software, Articles*, 80(1):1–28.

Foppolo, F. and Panzeri, F. (2011). When *straight* means *relatively straight* and *big* means *absolutely big*. Paper presented at the 31st Incontro di Grammatica Generativa, Rome, Italy, February 24-26.

Kennedy, C. (2007). Vagueness and grammar: The semantics of relative and absolute gradable predicates. *Linguistics and Philosophy*, 30(1):1–45.

Kennedy, C. and McNally, L. (2005). Scale structure and the semantic typology of gradable predicates. *Language*, 81(2):345–381.

Lassiter, D. and Goodman, N. (2014). Context, scale structure, and statistics in the interpretation of positive-form adjectives. In *Proceedings of Semantics and Linguistic Theory*, volume 24, pages 587–610.

Paradis, C. (2001). Adjectives and boundedness. *Cognitive Linguistics*, 12(1):47–65.

Qing, C. and Franke, M. (2014). Gradable adjectives, vagueness, and optimal language use: A speaker-oriented model. In Champollion, L. and Szabolcsi, A., editors, *Proceedings of Semantics and Linguistic Theory 24*, pages 23–41.

Rotstein, C. and Winter, Y. (2004). Total adjectives vs. partial adjectives: Scale structure and higher-order modifiers. *Natural Language Semantics*, 12:259–288.

Syrett, K. (2007). *Learning about the Structure of Scales: Adverbial Modification and the Acquisition of the Semantics of Gradable Adjectives*. PhD thesis, Northwestern University.

Toledo, A. and Sassoon, G. W. (2011). Absolute vs. relative adjectives - variance within vs. between individuals. In Ashton, N., Chereches, A., and Lutz, D., editors, *Proceedings of Semantics and Linguistic Theory 21*, pages 135–154.