Improving NLP systems with Questions Under Discussion

Christopher Potts

Stanford Linguistics and the Stanford NLP Group

UnImplicit Workshop, ACL-IJCNLP 2021

Overview

- Linguistic motivation for Questions Under Discussion (QUDs)
- Bringing QUDs into NLP systems
- QUDs for image captioning
- Looking ahead

Linguistic motivation for QUDs

Groenendijk & Stokhof (and Wittgenstein)

Wittgenstein	We might very well write every assertion in the form of a question followed by an affirmative expression; for instance 'Is it raining? Yes!' Would that mean that behind every claim lies a question?
Groenendijk & Stokhof	Ja!

Did Bart pass? ↑ Bart passed.

Who passed? ↑ BART$_{\mathbf{F}}$ passed.

What did Bart do? ↑ Bart PASSED$_{\mathbf{F}}$.

Who wore what? ↑ The ROCKSTARS$_{\mathbf{F}}$ wore LEATHER$_{\mathbf{F}}$.

Wittgenstein 1953; Groenendijk and Stokhof 1984

Questions/Issues: What are they?

Discourse is structured by evolving abstract, implicit, issues about which the participants have only partial knowledge.

1. Questions present alternatives.
2. Questions are not necessarily linguistic objects, though some natural language sentences might identify some of them.
3. Questions can be partially ordered by some notion of resolution.

Conceptions of QUDs

Questions under Discussion (QUDs)

McCarthy 1980; Groenendijk and Stokhof 1984; Rooth 1985; Lewis 1988; Ginzburg 1996a; Roberts 1996; Büring 1999

Decision problems

Lewis 1969; Clark 1996; Merin 1997; Blutner 1998; Parikh 2001; Beaver 2002; van Rooy 2003; Benz et al. 2005; Franke 2009

Task-orientation

Perrault and Allen 1980; Allen 1991; Hobbs et al. 1993; Allen et al. 2007; Stone et al. 2007

Pragmatically required over-answering

Context: Homer calls a hotel.

Homer: Is Lisa Simpson in Room 10?
Clerk A: She's in room 20.
Clerk B: #No.

Roberts 1996; Ginzburg 1996a; Champollion 2008

Pragmatically required over-answering

Context: Homer calls a hotel.

Homer: Is Lisa Simpson in Room 10?
Clerk A: She's in room 20.
Clerk B: #No.

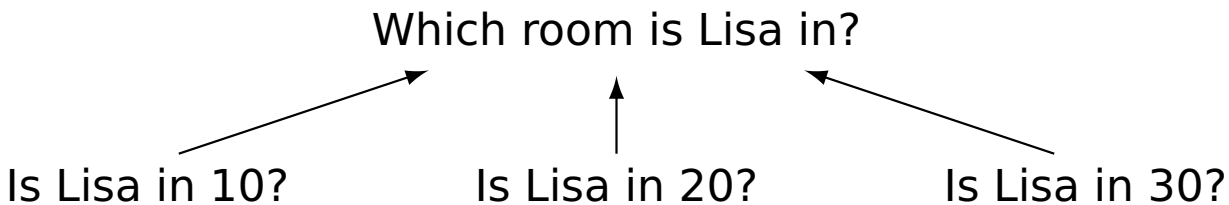

Roberts 1996; Ginzburg 1996a; Champollion 2008

Anaphora

Homer: Did Maggie shoot Burns?

Wiggum: She did.

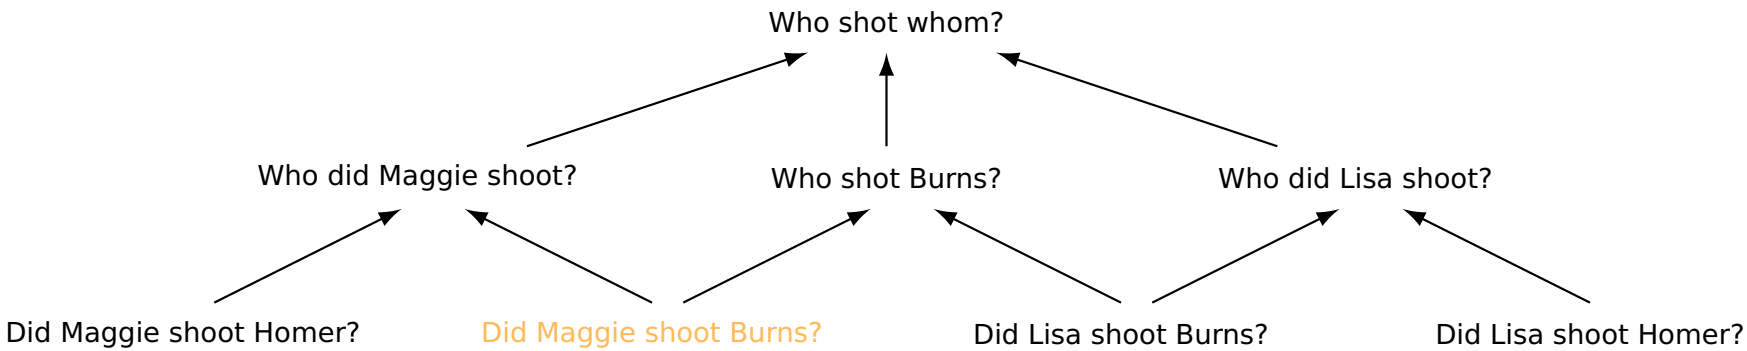

Clark and Parikh 2007; Schoubye 2009

Anaphora

Homer: Did Maggie shoot Burns?

Wiggum: She did. Maggie shot Burns

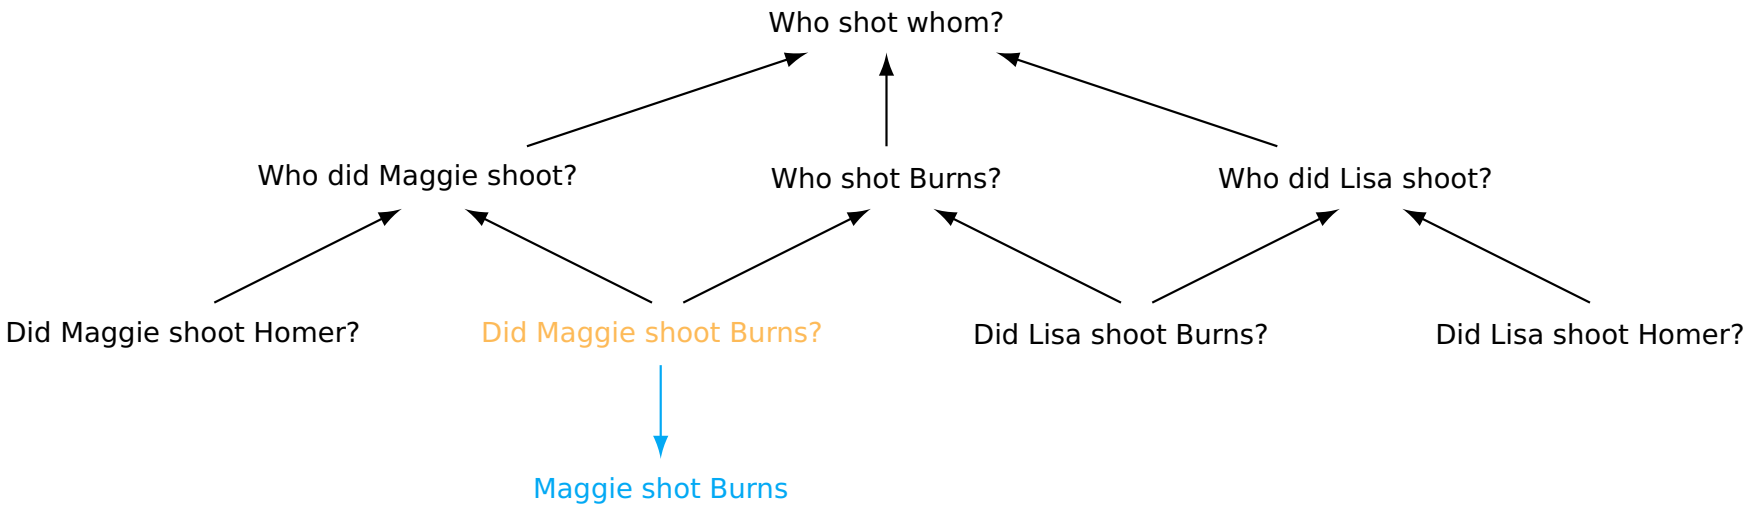

Clark and Parikh 2007; Schoubye 2009

Anaphora

Homer: Did Maggie shoot Burns?

Wiggum: She did. Maggie shot Homer

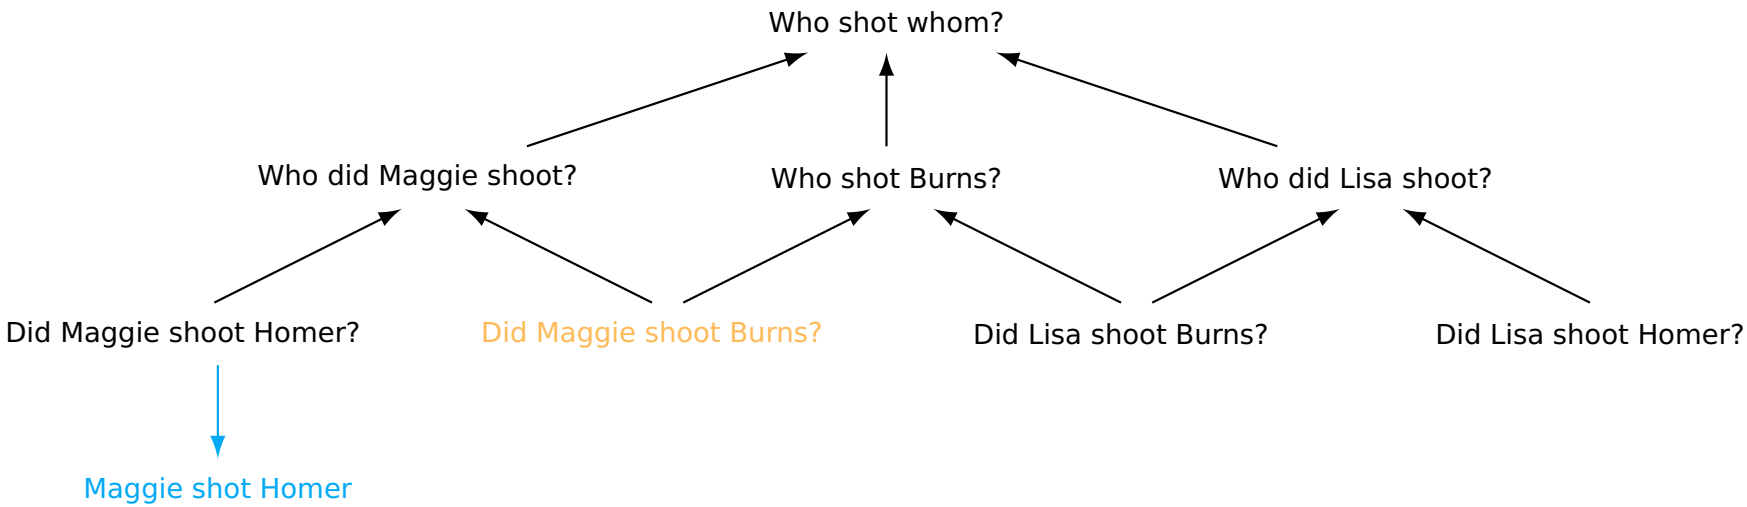

Clark and Parikh 2007; Schoubye 2009

Anaphora

Homer: Did Maggie shoot Burns?

Wiggum: She did. Lisa shot Burns

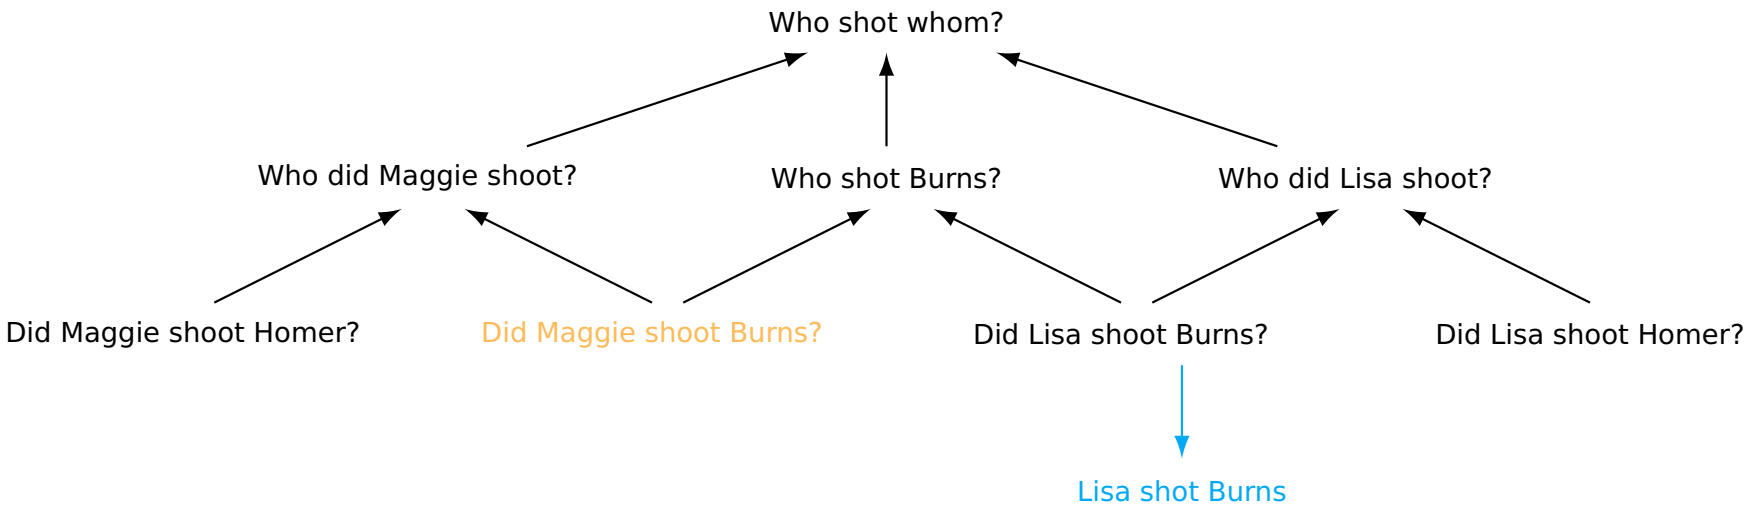

Clark and Parikh 2007; Schoubye 2009

Domain restriction

- Are there typos in my slides?
- Are there bookstores downtown?
- Are there cookies in the cupboard?
- . . .

I didn't see any.

Roberts 1996; Ginzburg 1996a; Malamud 2006

Granularity

Where are you from?

- *Connecticut.* (Issue: birthplaces)
- *The U.S.* (Issue: nationalities)
- *Stanford.* (Issue: affiliations)
- *Planet earth.* (Issue: intergalactic meetings)

Groenendijk and Stokhof 1984; Ginzburg 1996b

Mention-some/mention-all

Who has a lighter?

Mention-all

- **Context**: Speaker needs to ensure that no one in the group is going to get stopped by airport security.
- **Resolvedness condition**: List of everyone who has a lighter.

Mention-some

- **Context**: Speaker needs to light their cigar.
- **Resolvedness condition**: Just name one (friendly, willing, nearby) person with a lighter.

Beck and Rullmann 1999; Schulz and van Rooij 2006; Benz 2005

Structured domain restriction

What cards do you have?

Wide domain

- **Context**: Speaker dealt the cards and noticed that some were missing.
- **Resolvedness condition**: List everything you're holding.

Narrowed, structured domain

- **Context**: Speaker folds and wants to know why they lost.
- **Resolvedness condition**: Just name the good cards.

Beck and Rullmann 1999; Schulz and van Rooij 2006; Benz 2005

Identity and issue resolution

Who is Cassius Clay?

1. Muhammed Ali.
2. The person over there [pointing].
3. The greatest heavyweight boxer in history.

Aloni 2000; van Rooy 2003; Aloni and Port 2015

Others

- Intonational meaning: Rooth 1985; Büring 1999; Büring 2003; Roberts 1996
- Discourse particles: Roberts 2006; Beaver and Clark 2008; Kratzer and Matthewson 2009; Davis 2011; Rojas-Esponda 2015
- Presuppositions: Stone et al. 2007; Malamud 2006; Schoubye 2009; Tonhauser et al. 2013
- Connectives: Merin 1997; Toosarvandani 2010
- Negation and negative polarity: Fauconnier 1975; Anscombre and Ducrot 1983; Israel 2001, 2004; Potts 2011
- Ellipsis resolution: AnderBois 2010; Barros 2014; Weir 2014; Kotek and Barros 2018

Summary

- Good evidence that QUDs (broadly construed) are a factor in resolving context dependence.
- Growing body of quantitative and corpus exploration of the idea within linguistics and cognitive psychology: Cooper and Larsson 2001; DeVault 2008; DeVault and Stone 2007; Ginzburg and Fernández 2010; Goodman and Lassiter 2015; Kao et al. 2014; Hawkins and Goodman 2020
- This is helping us better understand where and how QUDs come into play, and how speakers represent discourses, issues, and lexical and constructional meanings.

Bringing QUDs into NLP systems

Application areas

Application areas

- Core language understanding

Application areas

- Core language understanding
- Dialogue

Application areas

- Core language understanding
- Dialogue
- Question generation

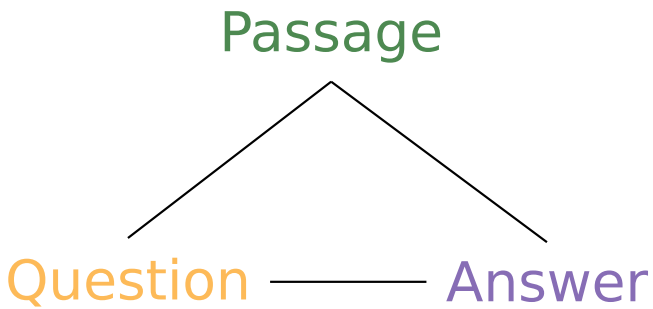

Application areas

- Core language understanding
- Dialogue
- Question generation
- Task-oriented question answering

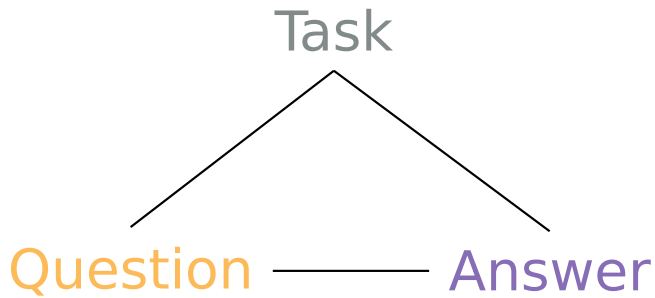

Application areas

- Core language understanding
- Dialogue
- Question generation
- Task-oriented question answering
- Image captioning

Application areas

- Core language understanding
- Dialogue
- Question generation
- Task-oriented question answering
- Image captioning

Which celebrities make the most money?

Application areas

- Core language understanding
- Dialogue
- Question generation
- Task-oriented question answering
- Image captioning

Which celebrities make the most money?

Lionel Messi is among the highest paid athletes in the world.

Application areas

- Core language understanding
- Dialogue
- Question generation
- Task-oriented question answering
- Image captioning

Famous Argentines

Lionel Messi hails from Rosario.

Application areas

- Core language understanding
- Dialogue
- Question generation
- Task-oriented question answering
- Image captioning

Famous Argentines

Lionel Messi is known for his elaborate hairstyles.

Application areas

- Core language understanding
- Dialogue
- Question generation
- Task-oriented question answering
- Image captioning

What's the deal with soccer players' hair?

Lionel Messi is known for his elaborate hairstyles.

Application areas

- Core language understanding
- Dialogue
- Question generation
- Task-oriented question answering
- Image captioning
- Image description

Example of Alt Text with Various Contexts

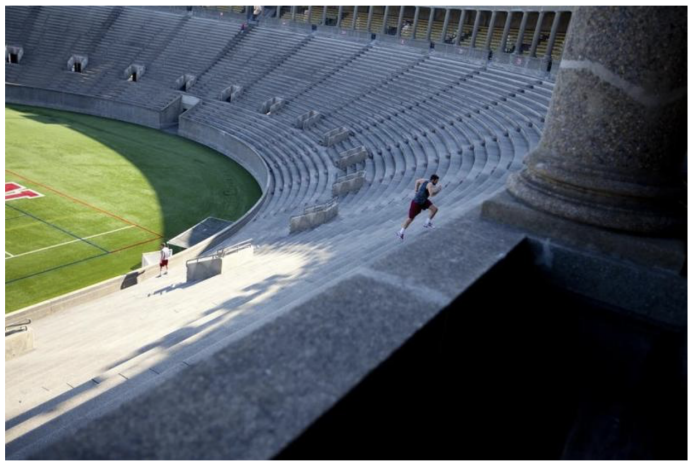

Alt-text with no context:

A mostly empty stadium.

Alt-text on a page about recent turnout for track tryouts:

Harvard Stadium with two lone runners bounding up the steps.

Alt-text on page about renovation projects:

Harvard Stadium with cracked concrete pillars.

(Kreiss et al. 2021)

Approaches

- Supervised learning on QUD annotations
- Symbolic encoding
- Distributed numerical encoding
- **Structured domains**

QUD annotations

. . . and then you realize the power you have. You can wire tap the President of the United States, you can wire tap a Federal Judge.

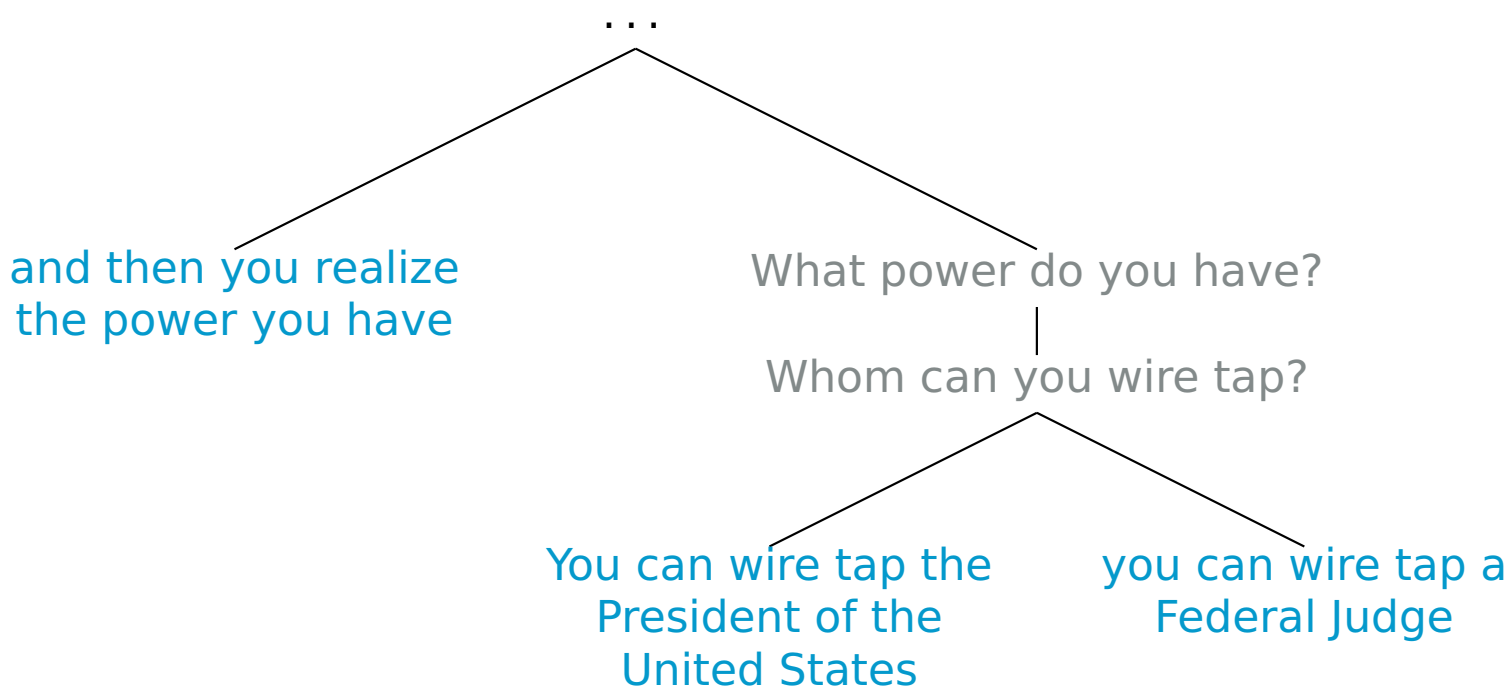

Riester et al. 2018

Symbolic encoding

ISSUE : $?x.\mathbf{price}(x)$
PLAN: ⟨
findout($?x.\mathbf{means\_of\_transport}(x)$),
findout($?x.\mathbf{dest\_city}(x)$),
findout($?x.\mathbf{depart\text{-}city}(x)$),
findout($?x.\mathbf{depart\text{-}month}(x)$),
findout($?x.\mathbf{depart\text{-}day}(x)$),
findout($?x.\mathbf{class}(x)$),
consultDB($?x.\mathbf{price}(x)$)
⟩

ISSUE : **?need_visa**
PLAN: ⟨
findout($?x.\mathbf{dest\_city}(x)$),
findout($?x.\mathbf{citizenship}(x)$),
consultDB(**?need_visa**),
⟩

S> `How do you want to travel?`

getLatestMove
integrateSysAsk
{ push(/SHARED/QUD, $?A$.how(A))

$$\left[\begin{array}{lcl} \text{PRIVATE} & = & \left[\begin{array}{lcl} \text{AGENDA} & = & \langle\rangle \\ \text{PLAN} & = & \left\langle \begin{array}{l} \text{findout}(?A.\text{dest\_city}(A)) \\ \text{findout}(?B.\text{dept\_city}(B)) \\ \text{findout}(?C.\text{month}(C)) \\ \text{findout}(?D.\text{dept\_day}(D)) \\ \text{findout}(?E.\text{class}(E)) \\ \text{consultDB}(?F.\text{price}(F)) \end{array} \right\rangle \\ \text{BEL} & = & \{\} \end{array}\right] \\ \text{SHARED} & = & \left[\begin{array}{lcl} \text{COM} & = & \{\} \\ \text{QUD} & = & \left\langle \begin{array}{l} ?G.\text{how}(G) \\ ?H.\text{price}(H) \end{array} \right\rangle \\ \text{LU} & = & \left[\begin{array}{lcl} \text{SPEAKER} & = & \text{sys} \\ \text{MOVES} & = & \{\ \text{ask}(?G.\text{how}(G))\ \} \end{array}\right] \end{array}\right] \end{array}\right]$$

Larsson 2002

Distributed encoding

GPT-3: “Professor” context

Sandy is Jesse’s professor. Sandy says, “Have you done the assignment?” Jesse replies, “No, I haven’t.” Sandy then says, “Can I see your grade? You’re failing.” Jesse is upset, but Sandy says, “Don’t worry. I’ll help you.” Jesse replies, “Thanks, but I can do it myself.” Sandy says, “You can’t do it yourself.” Jesse replies, “I can.”

GPT-3: “Classmate” context

Sandy is Jesse’s classmate. Sandy says, “Have you done the assignment?” Jesse replies, “Yes, I have done it.” Jesse is not a teacher, but he is an adult. Sandy is not an adult, but she is also not a child. Jesse and Sandy are third grade students.

Quick experiment

FWIW – In 16/20 cases, I was able to correctly guess the prompt from the continuation.

QUDs for image captioning

Issue-sensitive image captioning

Lionel Messi is known for his elaborate hairstyles.

Example of Alt Text with Various Contexts

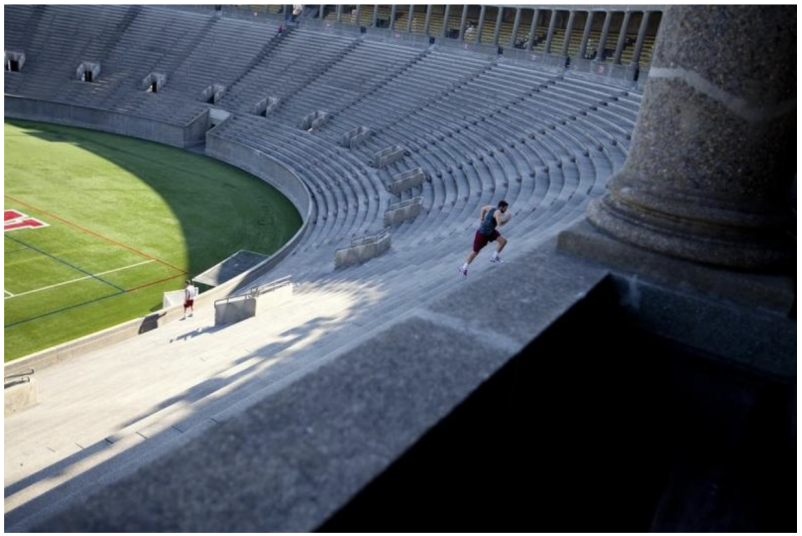

Alt-text with no context:

A mostly empty stadium.

Alt-text on a page about recent turnout for track tryouts:

Harvard Stadium with two lone runners bounding up the steps.

Alt-text on page about renovation projects:

Harvard Stadium with cracked concrete pillars.

Pragmatic Issue-Sensitive Image Captioning

Allen Nie[1] **Reuben Cohn-Gordon**[2] **Christopher Potts**[1,2]

[1]Department of Computer Science [2]Department of Linguistics
Stanford University
{anie, reubencg, cgpotts}@stanford.edu

Goals and approach

Goals

1. QUD-sensitive image/text pairs
2. No special annotation
3. No new datasets
4. No new model training

Approach

- Rational Speech Acts model (RSA) with QUDs (Goodman and Lassiter 2015; Kao et al. 2014; Hawkins and Goodman 2020)
- Neural RSA (Andreas and Klein 2016; Fried et al. 2018; Monroe et al. 2017, 2018)

Desired behavior

Issue/QUD	Target	Caption
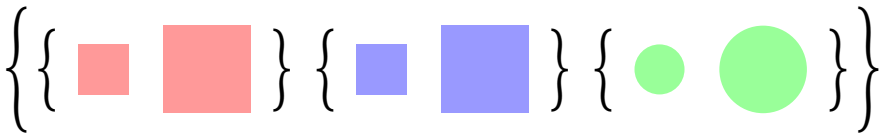		“A red square”
		“A small square”

Texts should describe the *cell* containing the target.

Some more issues

Some more issues

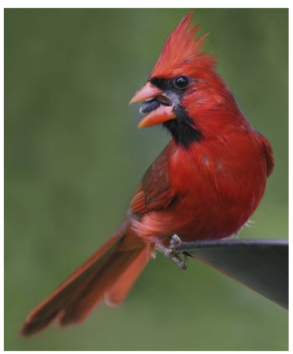



Some more issues

What color crown?

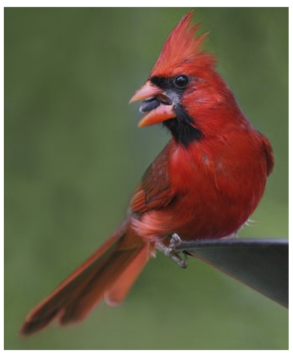

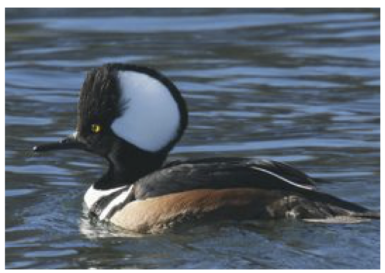


Some more issues

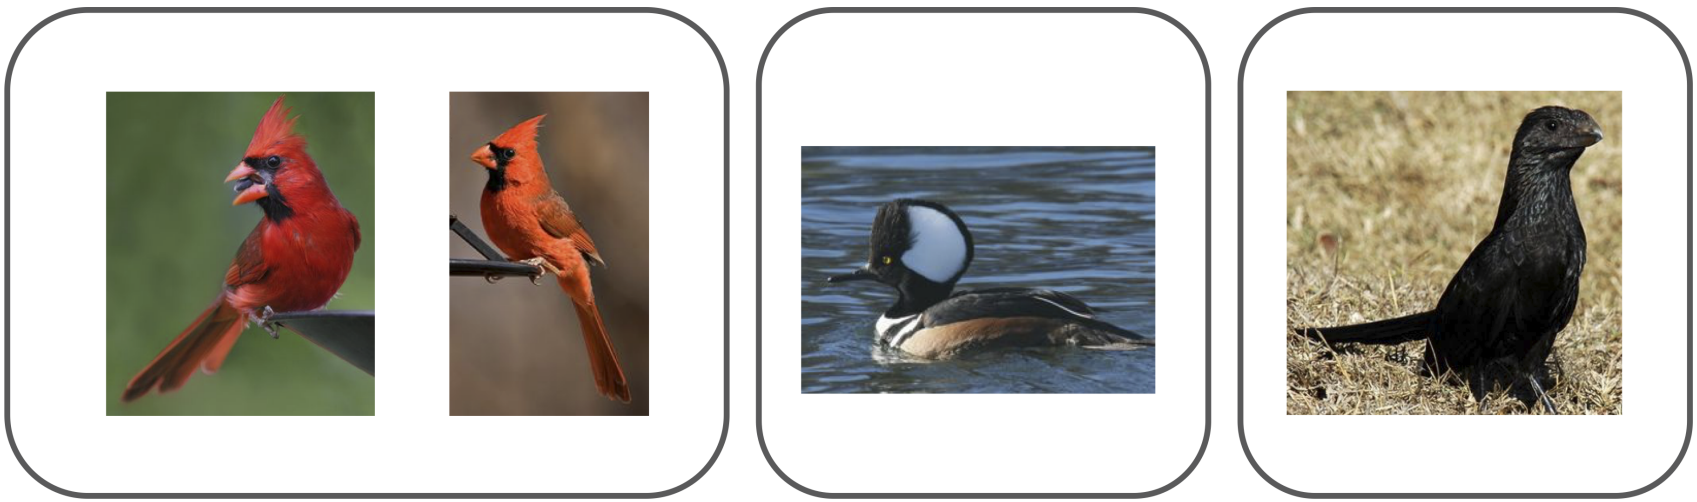

Some more issues

What color body?

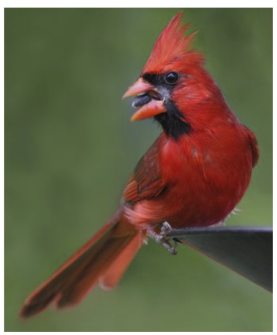

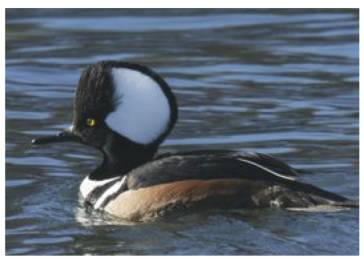


Some more issues

Some more issues

What position?

Some more issues

Some more issues

Gray pants?

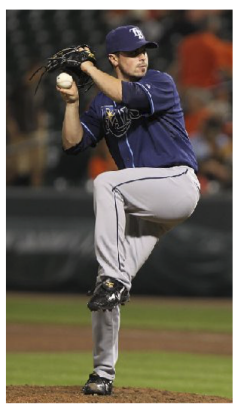

Rational Speech Acts model

Rational Speech Acts model

Base speaker

$S_0(msg \mid img)$ as given by a pretrained image2text model

Rational Speech Acts model

Pragmatic listener

$$L_1(img \mid msg) = \frac{S_0(msg \mid img)P(img)}{\sum_{img' \in \textbf{Images}} S_0(msg \mid img')P(img')}$$

Base speaker

$S_0(msg \mid img)$ as given by a pretrained image2text model

Rational Speech Acts model

Issue-sensitive pragmatic speaker

$$S_1^{\mathbf{C}}(msg \mid img, \mathbf{C}) \propto$$

$$\exp\left(\alpha \log\left(\sum_{img' \in \mathbf{Images}} \delta_{[\mathbf{C}(img)=\mathbf{C}(img')]} L_1(img' \mid msg)\right) + \log S_0(msg \mid img)\right)$$

Pragmatic listener

$$L_1(img \mid msg) = \frac{S_0(msg \mid img)P(img)}{\sum_{img' \in \mathbf{Images}} S_0(msg \mid img')P(img')}$$

Base speaker

$S_0(msg \mid img)$ as given by a pretrained image2text model

Rational Speech Acts model

Issue-sensitive pragmatic speaker

$$S_1^{\mathbf{C}}(msg \mid img, \mathbf{C}) =$$

$L_1(img \mid msg)$ at the level of details given by C

− message costs

Pragmatic listener

$L_1(img \mid msg) =$ base speaker × prior on images

Base speaker

$S_0(msg \mid img)$ as given by a pretrained image2text model

Example calculation

	small	red	circle	green
	1	1	0	0
	0	1	1	0
	1	0	1	1
	0	0	0	1

- $s_1^{\mathbf{c}}$
- L_1
- s_0

Example calculation

S_0	small	red	circle	green
	0.50	0.50	0.00	0.00
	0.00	0.50	0.50	0.00
	0.33	0.00	0.33	0.33
	0.00	0.00	0.00	1.00

- $S_1^{\mathbf{c}}$
- L_1
- S_0

Example calculation

L_1				
small	0.60	0.00	0.40	0.00
red	0.50	0.50	0.00	0.00
circle	0.00	0.60	0.40	0.00
green	0.00	0.00	0.25	0.75

- $S_1^{\mathbf{c}}$
- L_1
- S_0

Example calculation

S_1	small	red	circle	green
	0.55	0.45	0.00	0.00
	0.00	0.45	0.55	0.00
	0.38	0.00	0.38	0.24
	0.00	0.00	0.00	1.00

- $s_1^{\mathbf{c}}$
- L_1
- s_0

Example calculation

$S_1^{\mathbf{c}}$	small	red	circle	green
	0.27	0.45	0.27	0.00
	0.27	0.45	0.27	0.00
	0.22	0.00	0.22	0.56
	0.22	0.00	0.22	0.56

- $S_1^{\mathbf{c}}$
- L_1
- S_0

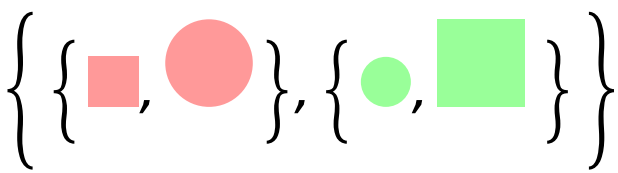

CalTech-UCSD Bird Dataset (CUB)

11,788 images for 200 species of North American birds, each annotated with 312 attributes using a system devised by ornithologists.

Attribute Annotation
Has_Bill_Shape::All-purpose
Has_Wing_Color::Brown
Has_Wing_Color::Rufous
Has_Back_Color::Brown
Has_Head_Pattern::Eyebrow
Has_Size::Small

Welinder et al. 2010

Some CUB system outputs

Issues	Target	Caption
What is the **color** of the bird? 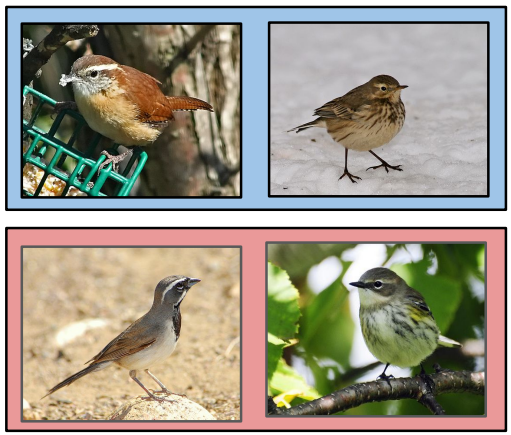	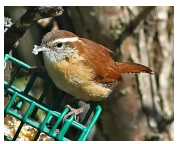	a small **brown bird** with a tan chest and a tan beak
What is the **head pattern** of the bird? 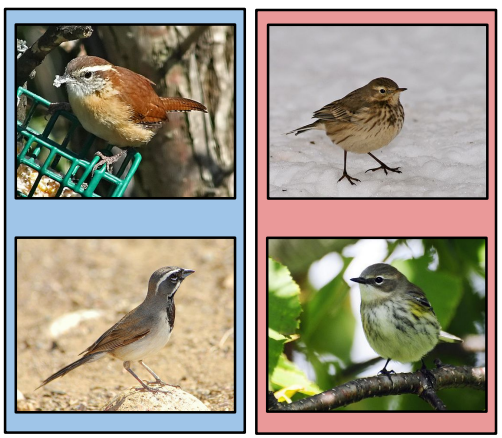	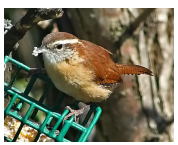	this bird has a brown crown a **white eyebrow** and a rounded belly

Pretrained model from Hendricks et al. 2016

Human evaluation

4

Question: **What is the beak shape?**

Caption: **this is a white bird with black feet and a pointy downward beak**

Select the answer conveyed by the caption, or indicate that the caption doesn't provide an answer:

- ○ **curved_(up_or_down)**
- ○ **dagger**
- ○ **hooked**
- ○ **needle**
- ○ **hooked_seabird**
- ○ **spatulate**
- ○ **all-purpose**
- ○ **cone**
- ○ **specialized**
- ○ **The caption answers the question, but not with one of the above options**
- ○ **The caption does not contain an answer to the question**

Human evaluation

Caption Source	Percentage	Size
S_0	20.9	273
S_1	24.5	273
$S_1^{\mathbf{C}}$	42.1	273
$S_1^{\mathbf{C}+H}$	**44.0**	273
Human	33.3	273

MS COCO and VQA 2.0

Inducing partitions

- VQA 2.0 contains triples (`image, question, answer`)
- Given `question`, find all associated `images`, and partition based on `answer`.

Moving to a trained VQA system

Given a question and a set of images, use the system to create a partition.

Lin et al. 2014; Goyal et al. 2017

Some MS COCO system outputs

Target	Issues	Partitions	Issue-sensitive Caption	Base Caption
	What color is the sky?		a **black and white** photo of an airplane in the sky	an airplane taking off from an airport runway
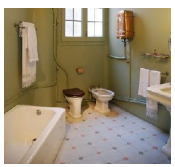	How many toilets are there?		a bathroom with **two toilets** and a tub	a bathroom with a tub and a toilet and a window

Summary

Goals

1. QUD-sensitive image/text pairs — Progress!
2. No special annotation — ✔
3. No new datasets — ✔
4. No new model training — ✔

Required ingredients

- Pretrained image-to-text model
- Method for structuring images into *issues* (partitions)

Looking ahead

1. QUDs are central to many aspects of language use
2. QUDs could benefit a wide range of NLP tasks
3. We can control text generation via simple QUD-like structures on the contexts (images):

$$S_1^{\mathbf{C}}(msg \mid img, \mathbf{C})$$

4. Can we achieve similar effects using language models?

$$\mathbf{LM}(msg \mid world, question)$$

Thanks!

References I

James F. Allen. 1991. *Reasoning About Plans*. Morgan Kaufmann, San Francisco.

James F. Allen, Nathanael Chambers, George Ferguson, Lucian Galescu, Hyuckchul Jung, Mary Swift, and William Taysom. 2007. PLOW: A collaborative task learning agent. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence*, pages 1514–1519. AAAI Press, Vancouver, British Columbia, Canada.

Maria Aloni. 2000. *Quantification under Conceptual Covers*. Ph.D. thesis, Universiteit van Amsterdam, Amsterdam. Published in the ILLC Dissertation Series, 2001-1.

Maria Aloni and Angelika Port. 2015. Epistemic indefinites and methods of identification. In Luis Alonso-Ovalle and Paula Menéndez-Benito, editors, *Epistemic Indefinites: Exploring Modality Beyond the Verbal Domain*, pages 117–140. Oxford University Press, Oxford.

Scott AnderBois. 2010. Sluicing as anaphora to issues. In *Proceedings of Semantics and Linguistic Theory 20*, pages 451–470, Ithaca, NY. CLC Publications.

Jacob Andreas and Dan Klein. 2016. Reasoning about pragmatics with neural listeners and speakers. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1173–1182. Association for Computational Linguistics.

Jean-Claude Anscombre and Oswald Ducrot. 1983. *L'Argumentation dans la lange*. Mardaga, Bruxelles.

Matthew Barros. 2014. *Sluicing and identity in ellipsis*. Ph.D. thesis, Rutgers The State University of New Jersey-New Brunswick.

David I. Beaver. 2002. Pragmatics, and that's an order. In David Barker-Plummer, David I. Beaver, Johan van Benthem, and Patrick Scotto di Luzio, editors, *Logic, Language, and Visual Information*, pages 192–215. CSLI, Stanford, CA.

David I. Beaver and Brady Zack Clark. 2008. *Sense and Sensitivity: How Focus Determines Meaning*. Wiley-Blackwell, Oxford.

Sigrid Beck and Hotze Rullmann. 1999. A flexible approach to exhaustivity in questions. *Natural Language Semantics*, 7(3):249–298.

Anton Benz. 2005. Utility and relevance of answers. In Benz et al. (2005), pages 195–219.

Anton Benz, Gerhard Jäger, and Robert van Rooij, editors. 2005. *Game Theory and Pragmatics*. Palgrave McMillan, Basingstoke, Hampshire.

Reinhard Blutner. 1998. Lexical pragmatics. *Journal of Semantics*, 15(2):115–162.

Daniel Büring. 1999. Topic. In Peter Bosch and Rob van der Sandt, editors, *Focus — Linguistic, Cognitive, and Computational Perspectives*, pages 142–165. Cambridge University Press, Cambridge.

Daniel Büring. 2003. On D-trees, beans, and B-accents. *Linguistics and Philosophy*, 26(5):511–545.

Lucas Champollion. 2008. The influence of goals on ambiguities in certain donkey sentences. Presentation at the 4th Formal Semantics in Moscow Workshop.

Herbert H. Clark. 1996. *Using Language*. Cambridge University Press, Cambridge.

Robin Clark and Prashant Parikh. 2007. Game theory and discourse anaphora. *Journal of Logic, Language, and Information*, 16(3):265–282.

References II

Robin Cooper and Staffan Larsson. 2001. Accommodation and reaccommodation in dialogue. In *Gothenburg Papers in Computational Linguistics*. Department of Linguistics, Göteborg University.

Christopher Davis. 2011. *Constraining Interpretation: Sentence Final Particles in Japanese*. Ph.D. thesis, UMass Amherst, Amherst, MA.

David DeVault. 2008. *Contribution Tracking: Participating in Task-Oriented Dialogue under Uncertainty*. Ph.D. thesis, Rutgers University, New Brunswick, NJ.

David DeVault and Matthew Stone. 2007. Managing ambiguities across utterances in dialogue. In *Proceedings of DECALOG 2007: Workshop on the Semantics and Pragmatics of Dialogue*.

Gilles Fauconnier. 1975. Polarity and the scale principle. In *Papers of the Eleventh Regional Meeting of the Chicago Linguistic Society*, pages 188–199, Chicago. Chicago Linguistic Society.

Michael Franke. 2009. *Signal to Act: Game Theory in Pragmatics*. ILLC Dissertation Series. Institute for Logic, Language and Computation, University of Amsterdam.

Daniel Fried, Jacob Andreas, and Dan Klein. 2018. Unified pragmatic models for generating and following instructions. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1951–1963, New Orleans, Louisiana. Association for Computational Linguistics.

Jonathan Ginzburg. 1996a. Dynamics and the semantics of dialogue. In Jerry Seligman, editor, *Language, Logic, and Computation*, volume 1, pages 221–237. CSLI, Stanford, CA.

Jonathan Ginzburg. 1996b. Interrogatives: Questions, facts, and dialogue. In Shalom Lappin, editor, *The Handbook of Contemporary Semantic Theory*, pages 385–422. Blackwell Publishers, Oxford.

Jonathan Ginzburg and Raquel Fernández. 2010. Computational models of dialogue. In Alex Clark, Chris Fox, and Shalom Lappin, editors, *Handbook of Computational Linguistics and Natural Language Processing*. Blackwell, Oxford.

Noah D. Goodman and Daniel Lassiter. 2015. Probabilistic semantics and pragmatics. In Shalom Lappin and Chris Fox, editors, *The Handbook of Contemporary Semantic Theory*, 2 edition. Wiley-Blackwell, Oxford.

Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the V in VQA matter: Elevating the role of image understanding in Visual Question Answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6904–6913.

Jeroen Groenendijk and Martin Stokhof. 1984. *Studies in the Semantics of Questions and the Pragmatics of Answers*. Ph.D. thesis, University of Amsterdam.

Robert X. D. Hawkins and Noah D. Goodman. 2020. Why do you ask? the informational dynamics of questions and answers. Ms., Princeton University and Stanford University.

Lisa Anne Hendricks, Zeynep Akata, Marcus Rohrbach, Jeff Donahue, Bernt Schiele, and Trevor Darrell. 2016. Generating visual explanations. In *European conference on computer vision*, pages 3–19. Springer.

Jerry R. Hobbs, Mark Stickel, Douglas Appelt, and Paul Martin. 1993. Interpretation as abduction. *Artificial Intelligence*, 63(1–2):69–142.

References III

Michael Israel. 2001. Minimizers, maximizers, and the rhetoric of scalar reasoning. *Journal of Semantics*, 18(4):297–331.

Michael Israel. 2004. The pragmatics of polarity. In Laurence R. Horn and Gregory Ward, editors, *The Handbook of Pragmatics*, pages 701–723. Blackwell, Oxford.

Justine T. Kao, Leon Bergen, and Noah D. Goodman. 2014. Formalizing the pragmatics of metaphor understanding. In *Proceedings of the 36th Annual Meeting of the Cognitive Science Society*, pages 719–724, Wheat Ridge, CO. Cognitive Science Society.

Hadas Kotek and Matthew Barros. 2018. Multiple sluicing, scope, and superiority: Consequences for ellipsis identity. *Linguistic Inquiry*, 49(4):781–812.

Angelika Kratzer and Lisa Matthewson. 2009. Anatomy of two discourse particles. Paper presented at Semantics of Under-Represented Languages of the Americas.

Elisa Kreiss, Noah D. Goodman, and Christopher Potts. 2021. Concadia: Tackling image accessibility with context. ArXiv:2104.08376.

Staffan Larsson. 2002. *Issue-based Dialogue Management*. Ph.D. thesis, Göteborg University, Sweden.

David Lewis. 1969. *Convention*. Harvard University Press, Cambridge, MA. Reprinted 2002 by Blackwell.

David Lewis. 1988. Relevant implication. *Theoria*, 54(3):161–174.

Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft COCO: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer.

Sophia A. Malamud. 2006. (Non)-maximality and distributivity: A decision theory approach. Paper presented at SALT 16, Tokyo, Japan.

John McCarthy. 1980. Circumscription — a form of non-monotonic reasoning. *Artificial Intelligence*, 13(1):27–39.

Arthur Merin. 1997. If all our arguments had to be conclusive, there would be few of them. Arbeitspapiere SFB 340 101, University of Stuttgart, Stuttgart.

Will Monroe, Robert X. D. Hawkins, Noah D. Goodman, and Christopher Potts. 2017. Colors in context: A pragmatic neural model for grounded language understanding. *Transactions of the Association for Computational Linguistics*, 5:325–338.

Will Monroe, Jennifer Hu, Andrew Jong, and Christopher Potts. 2018. Generating bilingual pragmatic color references. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2155–2165, Stroudsburg, PA. Association for Computational Linguistics.

Allen Nie, Reuben Cohn-Gordon, and Christopher Potts. 2020. Pragmatic issue-sensitive image captioning. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1924–1938, Online. Association for Computational Linguistics.

Prashant Parikh. 2001. *The Use of Language*. CSLI, Stanford, CA.

C. Raymond Perrault and James F. Allen. 1980. A plan-based analysis of indirect speech acts. *American Journal of Computational Linguistics*, 6(3–4):167–182.

References IV

Christopher Potts. 2011. On the negativity of negation. In Nan Li and David Lutz, editors, *Proceedings of Semantics and Linguistic Theory 20*, pages 636–659. CLC Publications, Ithaca, NY.

Arndt Riester, Lisa Brunetti, and Kordula De Kuthy. 2018. Annotation guidelines for Questions under Discussion and information structure. *Information structure in lesser-described languages. Studies in prosody and syntax*, pages 403–443.

Craige Roberts. 1996. Information structure: Towards an integrated formal theory of pragmatics. In Jae Hak Yoon and Andreas Kathol, editors, *OSU Working Papers in Linguistics*, volume 49: Papers in Semantics, pages 91–136. The Ohio State University Department of Linguistics, Columbus, OH. Revised 1998.

Craige Roberts. 2006. *Only*, presupposition and implicature. Ms., The Ohio State University.

Tania Rojas-Esponda. 2015. *Patterns and Symmetries for Discourse Particles*. Ph.D. thesis, Stanford University, Stanford, CA.

Mats Rooth. 1985. *Association with Focus*. Ph.D. thesis, UMass Amherst.

Robert van Rooy. 2003. Questioning to resolve decision problems. *Linguistics and Philosophy*, 26(6):727–763.

Anders J. Schoubye. 2009. Descriptions, truth value intuitions, and questions. *Linguistics and Philosophy*, 32(6):583–617.

Katrin Schulz and Robert van Rooij. 2006. Pragmatic meaning and non-monotonic reasoning: The case of exhaustive interpretation. *Linguistics and Philosophy*, 29(2):205–250.

Matthew Stone, Richmond Thomason, and David DeVault. 2007. Enlightened update: A computational architecture for presupposition and other pragmatic phenomena. To appear in Donna K. Byron; Craige Roberts; and Scott Schwenter, *Presupposition Accommodation*.

Judith Tonhauser, David I. Beaver, Craige Roberts, and Mandy Simons. 2013. Towards a taxonomy of projective content. *Language*, 89(1):66–109.

Maziar Toosarvandani. 2010. *Association with Foci*. Ph.D. thesis, UC Berkeley.

Andrew Weir. 2014. *Fragments and clausal ellipsis*. Ph.D. thesis, UMass Amherst.

Peter Welinder, Steve Branson, Takeshi Mita, Catherine Wah, Florian Schroff, Serge Belongie, and Pietro Perona. 2010. Caltech-UCSD Birds 200. Technical Report CNS-TR-2010-001, California Institute of Technology.

Ludwig Wittgenstein. 1953. *Philosophical Investigations*. The MacMillan Company, New York. Translated by G. E. M. Anscombe.