Lattice Theoretic Relevance in Incremental Reference Processing

Julian Hough[1,2] **& Matthew Purver**[2]
[1]Dialogue Systems Group
Faculty of Linguistics and Literature
Bielefeld University
[2]Cognitive Science Research Group
School of Electronic Engineering and Computer Science
Queen Mary University of London
`julian.hough@uni-bielefeld.de`

1 Introduction

While there has been substantial work on referential communication tasks in psycholinguistics, computational and formal modelling (Dale and Reiter (1995), Krahmer and Van Deemter (2012), Frank and Goodman (2012) *inter alia*), the element we discuss here is incremental processing. Motivated by work in incremental generation of referring expressions (Guhe, 2007; Fernández, 2013) and incremental reference resolution in NLU (Kennington and Schlangen, 2014), we present a dialogue-motivated account which models the speaker and the hearer in reference identification games.

A central desideratum of an incremental account of reference identification tasks can be found in the evidence from Brennan and Schober (2001)'s experiments; namely that people reason at an incredibly time-critical level from linguistic information. They demonstrated *self-repair* can speed up semantic processing (or at least object reference) where an incorrect object being partly vocalized and then repaired in the instructions (e.g. "the yell-, uh, purple square") yields quicker response times from the onset of the target ("purple") than in the case of the fluent instructions ("the purple square"), with little effect on accuracy. We wish to model this faculty of repair processing, and also wish to model non-local repair processing of instructions such as "From yellow down to brown – no – thats red." (Levelt, 1989, via Ginzburg et al. (2014)), here using a syntactically simpler but illustrative alternative "the yellow square, no, purple".

We build on Hough and Purver (2014)'s integration of Knuth (2005)'s lattice-theoretic characterization of probabilistic inference to model interpretation of repaired instructions in a small reference domain.

2 TTR and probabilistic record type lattices

We assume a type-theoretic view of semantic processing following Cooper (2005)'s Type Theory with Records (TTR). In TTR, the principal logical form of interest is the *record type* ('RT' from here), consisting of sequences of *fields* of the form $[\, l : T \,]$ containing a label l and a type T.[1] RTs can be witnessed (i.e. judged as inhabited) by *records* of that type, where a record is a set of label-value pairs $[\, l = v \,]$. The central type judgement in TTR that a record s is of (record) type R, i.e. $s : R$, can be made from the component type judgements of individual fields; e.g. the one-field record $[\, l = v \,]$ is of type $[\, l : T \,]$ just in case v is of type T. Cooper et al. (2014)'s recent extension of TTR includes probabilistic type judgements of the form $p(s : R) = v$ where $v \in [0,1]$, i.e. the real valued judgement that a record s is of RT R.

In the case of the simple reference resolution game here, in interpretation, the challenge is to predict the reference situation judgement $s : R$ such that $\text{argmax}_R\, p(s : R)$ incrementally as instructions are heard word-by-word; this is a simplistic simulation of incremental linguistic processing. In generation, we model the task at the strategic level of selecting the most relevant

[1]See Cooper (2012) for a detailed formal description.

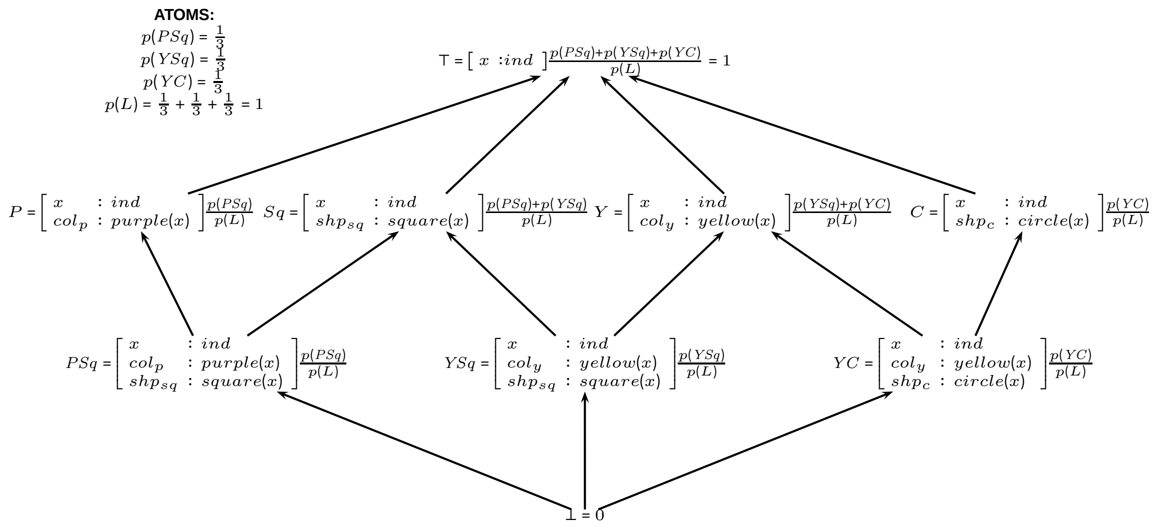

Figure 1: Record type lattice L with uniform atomic probabilities

goal concept (also a RT) to communicate the referent, a goal that may change incrementally during word-by-word surface realization (see (Hough, 2011)).

A domain of objects modelled as RTs can be decomposed into their constituent supertypes and compacted into a type lattice L ordered by the subtype relation $\sqsubseteq$ as in Figure 1. The *atoms*, the elements that cover $\perp$, are the exclusive referents, and we assume they all have a uniform probability before any instruction has been heard. The probabilities of the non-atomic type judgements can be calculated purely in terms of the meets (maximal common subtypes) and joins (minimal common supertypes) in L, as described by (Knuth, 2005) and (Hough and Purver, 2014). Every probability calculation has the total probability mass of the lattice $p(L)$ as its denominator, and all unconditional atomic probabilities will sum to 1.

As RTs become available from an incremental TTR parser (Purver et al., 2011), beginning from a situation of a uniform distribution of possible referents in the domain with the only available type judgement being $s : \top$, the lattice is descended until s is judged to be of one of the atom RTs, assuming for illustration with a conditional probability of 1.

Type-theoretic evidence becomes available incrementally through speech, e.g. after "the yellow" the judgement $s : Y$ can be made (i.e. "the referent is yellow"), and the atoms' probability distribution changes conditioned on this. The conditional probability judgements throughout the lattice are calculated as:

$$p(s : R_y \mid s : R_x) = \begin{cases} 1 & \text{if } R_x \sqsubseteq R_y \\ 0 & \text{if } R_x \wedge R_y = \perp \\ p & \text{otherwise, where } 0 \leq p \leq 1 \end{cases} \tag{1}$$

The final case can be calculated in terms of the probability of the meet of the two types in the lattice, with the equation following (Cooper et al., 2014):

$$p(s : R_y \mid s : R_x) = \frac{p(s : R_x \wedge s : R_y)}{p(s : R_x)} \tag{2}$$

For repair instances, where commitment to one judgement is revoked, it may be appropriate to condition on negative types. Given a type lattice will be distributive but not guaranteed to be complemented, we must derive $p(s : R_y \mid s : \neg R_x)$ by obtaining $p(s : R_y)$ in G modulo the probability mass of R_x and that of its subtypes:

$$p(s : R_y \mid s : \neg R_x) = \begin{cases} 0 & \text{if } R_y \sqsubseteq R_x \\ \frac{p(s:R_y) - p(s:R_x \wedge R_y)}{p(s:\top) - p(s:R_x)} & \text{otherwise} \end{cases} \quad (3)$$

3 Worked examples of repaired instructions

We assume a referent set of a purple square, yellow square and yellow circle $\{PSq, YSq, YC\}$ as shown in Figure 1. For our two test case utterances containing self-repairs, "the yellow, uh, purple square" and "the purple square, no, yellow", we show the probability distribution over the referent set at each word in Figure 2. The second row in each table also shows the incremental type judgements by which these values are calculated conditionally from equations (1)-(3). The type judgements are available from the maximal semantics from a word-by-word incremental TTR parser and also by the negation of the reparandum type judgement after self-repair detection (details omitted here, but see (Hough and Purver, 2014)).

The first example straight-forwardly simulates Brennan and Schober (2001)'s result of the hearer's use of the negation of the reparandum type judgement, i.e. $s : \neg Y$ that the referent is no longer a yellow object, upon repair detection. While $s : P$ would be sufficient conditioning evidence for $s : PSq$ to become the most likely referent situation immediately, $s : \neg Y$ needs to become available as extra conditioning information. We argue the faster resolution by human subjects upon the onset of "purple" could be due to $s : \neg Y$ becoming available before $s : P$, due to repair detection functioning before the full parse has been made. This requires strong interleaving of repair detection in the parser with the dialogue management module responsible for the lattice-based judgements for this to become possible.

To explain the second example, we require (Knuth, 2005)'s notion of lattice-theoretic *relevance*. Relevance measures are derived from the degree of inclusion of questions to the *central issue*, which in this case is $s : PSq \vee s : YSq \vee s : YC?$ (≈ "which of the three objects is it?"). The questions are not directly on L, but rather on its dual $Q(L)$, the *question lattice*, which is homomorphic to L as it is the lattice of the down-sets of each RT element in L, each of which can be characterized as possible answers to the questions. The down-set elements of $Q(L)$ are ordered by the set-inclusion relation and its join-irreducible elements form a sub-lattice that is isomorphic to L – see (Knuth (2005), p.13) for details. The relevance of a question Q on $Q(L)$ to the central issue (which we will abbreviate I) can be defined analogously to conditional probability as in (1) above, but is denoted $d(I \mid Q)$. This can be read as the degree to which I is answered by Q, or in lattice terms, the degree to which I includes Q:

$$d(I \mid Q) = \begin{cases} 1 & \text{if } Q \subseteq I \\ 0 & \text{if } I \wedge Q = \bot \\ d & \text{otherwise, where } 0 \leq d \leq 1 \end{cases} \quad (4)$$

In the spirit of optimal relevance (Sperber and Wilson, 1986), we assume the hearer interprets each word as an answer to the most relevant question under discussion incrementally, an approach similar to (Ginzburg et al., 2014). The interregnum "no" is interpreted as an answer to the most relevant question to the central issue answered in the context so far: in this case this is the elementary question $s : PSq?$ (≈ "is it the purple square?") with the down-set of possible answers only including two possibilities $\{s : PSq, s : PSq \rightarrow \bot\}$. In answering $s : PSq?$ with $s : PSq \rightarrow \bot$,'no' causes $s : \neg PSq$ to be inferred, following the characterization of negative type judgements from (Cooper et al., 2014), and the probability distribution adjusts accordingly. The equally most relevant questions to the central issue left unanswered after "no" include $s : YSq?$ and $s : YC?$, both with relevances of 1 as they are included in I.

Next, "yellow" explicitly answers the question $s : Y?$ but adds no new information in terms of the probability distribution of possible referents. If the previous judgement that this was a

	the	yell-	uh	purple	square
conditioning type judgement $s:$	$\top$	Y	$\top$	$P \wedge \neg Y$	PSq
$p(s : PSq)$ (purple square)	$\frac{1}{3}$	0	$\frac{1}{3}$	1	1
$p(s : YSq)$ (yellow square)	$\frac{1}{3}$	$\frac{1}{2}$	$\frac{1}{3}$	0	0
$p(s : YC)$ (yellow circle)	$\frac{1}{3}$	$\frac{1}{2}$	$\frac{1}{3}$	0	0

	the	purple	square	no	yellow
conditioning type judgement $s:$	$\top$	P	PSq	$\neg PSq$	$Sq \wedge Y$
$p(s : PSq)$ (purple square)	$\frac{1}{3}$	1	1	0	0
$p(s : YSq)$ (yellow square)	$\frac{1}{3}$	0	0	$\frac{1}{2}$	1
$p(s : YC)$ (yellow circle)	$\frac{1}{3}$	0	0	$\frac{1}{2}$	0

Figure 2: Probability distributions for the objects given maximal incremental semantic information

square $s : Sq$ is incorporated however (given there is no evidence in the repair this has changed to $s : \neg Sq$), then $s : Y \wedge Sq$, equivalent to $s : YSq$, answers a more relevant question $s : YSq?$, an elementary question of $Q(L)$, having the simple down-set $s : \{YSq, YSq \rightarrow \bot\}$ as its possible answers. $s : YSq$ is of course defeasible if 'circle' were to be the following word, however at this point the hearer has assumed optimal relevance in terms of the speaker answering the central issue on $Q(L)$ as efficiently as possible, so the probabilities in Figure 2 represent the most likely interpretation (top hypothesis).

4 Conclusion

We have presented a view of reference processing in terms of incremental probabilistic interpretation of an utterance given a scene (assuming full observability of words and referents) which includes an account of self-repaired referring expressions. The ellipsis resolution account in the second example is driven by interpreting fragments as answers to the most relevant outstanding question for reference resolution. In generation, the producer incrementally attempts to answer the jointly most relevant outstanding question for reference resolution and most relevant question for identifying the referent they intend to communicate as efficiently as possible. This question-based account has the same spirit as Ginzburg et al. (2014), though here we incorporate probabilistic judgements.

The implementation is ongoing. We have a TTR parser with adequate coverage for simple domains in the `DyLan` dialogue system's interpreter module (Eshghi et al., 2011) in Java. We have prototypes for RT lattice inference using Python's `features` package, which allows easy implementation of lattices described here, but we are integrating this into `DyLan` modules in future work.

Acknowledgements

Thanks to the three anonymous REFNET reviewers for their helpful comments.

Hough is supported by DUEL (Disfluencies, Exclamations and Laughter in Dialogue), a project financially supported by the Agence Nationale de la Research (grant number ANR-13-FRAL-0001) and the Deutsche Forschungsgemainschaft. Much of the work was carried out with support from an EPSRC DTA scholarship at Queen Mary University of London.

Purver is partly supported by ConCreTe: the project ConCreTe acknowledges the financial support of the Future and Emerging Technologies (FET) programme within the Seventh Framework Programme for Research of the European Commission, under FET grant number 611733.

References

S.E. Brennan and M.F. Schober. 2001. How listeners compensate for disfluencies in spontaneous speech. *Journal of Memory and Language*, 44(2):274–296.

Robin Cooper, Simon Dobnik, Shalom Lappin, and Staffan Larsson. 2014. A probabilistic rich type theory for semantic interpretation. In *Proceedings of the EACL Workshop on Type Theory and Natural Language Semantics (TTNLS)*, Gothenburg, Sweden, April. Association for Computational Linguistics.

Robin Cooper. 2005. Records and record types in semantic theory. *Journal of Logic and Computation*, 15(2):99–112.

Robin Cooper. 2012. Type theory and semantics in flux. In Ruth Kempson, Nicholas Asher, and Tim Fernando, editors, *Handbook of the Philosophy of Science*, volume 14: Philosophy of Linguistics, pages 271–323. North Holland.

Robert Dale and Ehud Reiter. 1995. Computational interpretations of the gricean maxims in the generation of referring expressions. *Cognitive Science*, 19(2):233–263.

A. Eshghi, M. Purver, and Julian Hough. 2011. Dylan: Parser for dynamic syntax. Technical report, Queen Mary University of London.

Raquel Fernández. 2013. Rethinking overspecification in terms of incremental processing. In *Proceedings of the PRE-CogSci 2013 Workshop on the Production of Referring Expressions*, Berlin,Germany.

Michael C Frank and Noah D Goodman. 2012. Predicting pragmatic reasoning in language games. *Science*, 336(6084):998–998.

Jonathan Ginzburg, Raquel Fernández, and David Schlangen. 2014. Disfluencies as intra-utterance dialogue moves. *Semantics and Pragmatics*, 7(9):1–64, June.

Markus Guhe. 2007. *Incremental Conceptualization for Language Production*. NJ: Lawrence Erlbaum Associates.

Julian Hough and Matthew Purver. 2014. Probabilistic type theory for incremental dialogue processing. In *Proceedings of the EACL 2014 Workshop on Type Theory and Natural Language Semantics (TTNLS)*, pages 80–88, Gothenburg, Sweden, April. Association for Computational Linguistics.

Julian Hough. 2011. Incremental semantics driven natural language generation with self-repairing capability. In *Recent Advances in Natural Language Processing (RANLP)*, pages 79–84, Hissar, Bulgaria.

Casey Kennington and David Schlangen. 2014. Situated incremental natural language understanding using markov logic networks. *Computer Speech & Language*, 28(1):240–255.

Kevin H Knuth. 2005. Lattice duality: The origin of probability and entropy. *Neurocomputing*, 67:245–274.

Emiel Krahmer and Kees Van Deemter. 2012. Computational generation of referring expressions: A survey. *Computational Linguistics*, 38(1):173–218.

W.J.M. Levelt. 1989. *Speaking: From Intention to Articulation*. MIT Press.

Matthew Purver, Arash Eshghi, and Julian Hough. 2011. Incremental semantic construction in a dialogue system. In J. Bos and S. Pulman, editors, *Proceedings of the 9th International Conference on Computational Semantics*, pages 365–369, Oxford, UK, January.

Dan Sperber and Deirdre Wilson. 1986. *Relevance: Communication and Cognition*. Blackwell.