Interactional Dynamics and the Emergence of Language Games

Arash Eshghi
Interaction Lab
Heriot-Watt University
eshghi.a@gmail.com

Igor Shalyminov
Interaction Lab
Heriot-Watt University
o.lemon@hw.ac.uk

Oliver Lemon
Interaction Lab
Heriot-Watt University
o.lemon@hw.ac.uk

Abstract

Meaning is highly activity-specific, in that the action that a particular sequence of words is taken to perform is severely underdetermined in the absence of an overarching activity, or a 'language-game'. In this paper, we combine a formal, incremental model of interactional dynamics and contextual update - Dynamic Syntax and Type Theory with Records (DS-TTR) - with Reinforcement Learning for word selection. We show, using an implemented system, that trial and error generation with a DS-TTR lexicon – a process we have dubbed `babbling` – leads to particular domain-specific dialogue acts to be learned and routinised over time; and thus that higher level dialogue structures - or how actions fit together to form a coherent whole - can be learned in this fashion. This method therefore allows incremental dialogue systems to be automatically bootstrapped from small amounts of unannotated dialogue transcripts, yet capturing a combinatorially large number of interactional variations. Even when the system is trained from only a single dialogue, we show that it supports over 8000 new dialogues in the same domain. This generalisation property results from the structural knowledge and constraints present within the grammar, and highlights limitations of recent state-of-the-art systems that are built using machine learning techniques only.

1 Introduction

Meaning is highly activity-specific, in that the action that a particular sequence of words is taken to perform, together with any perlocutionary effect that action might give rise to, is severely underdetermined in the absence of a particular overarching activity, or a 'language-game'. Wittgenstein famously argued that the structure of a language-game, or how actions fit together to form a coherent whole, is irreducible. Arguably, this is the most unyielding obstacle facing not only theoretical approaches to pragmatics, but also dialogue system developers today. This suggests that particular dialogue structures are emergent, learned, and very frequently adjusted during interaction (Mills and Gregoromichelaki, 2010; Mills, 2011; Healey, 2008; Larsson and Cooper, 2008).

Despite this, recent and ongoing work in formal dialogue modelling suggests that not only language processing mechanisms, but also certain basic principles of contextual dynamics in dialogue do generalise across domains (Ginzburg, 2012; Kempson et al., 2016; Eshghi et al., 2015; Kempson et al., 2015; Purver et al., 2010). Even in a simple domain, there's a lot of *interactional variation* that does not ultimately affect the overall communicative goal of a dialogue. For example, the dialogues in Fig. 1 (specifically the top two rows, where the lexicon is held constant) all lead to a context in which the user wants to buy a phone by LG. These dialogues can be said to be *pragmatically synonymous* for this domain. Arguably, a good model of interactional dynamics should be able to capture this synonymy.

In this paper, we show, using an implemented system (Kalatzis et al., 2016), that given Dynamic Syntax and Type Theory with Records (DS-TTR) (Kempson et al., 2001; Eshghi et al., 2012; Eshghi et al., 2015) as a low-level, incremental model of interactional and contextual dynamics, one can see dialogue acts, together with their associated local dialogue structures and procedural conventions as emergent and learned from interaction; and thus

Interactional	USR: I would like an LG lap- err, phone SYS: okay.	USR: I would like a phone by LG. SYS: sorry a what? USR: a phone by LG. SYS: okay.	SYS: what would you like? USR: an LG phone SYS: okay.
	SYS: what would you like? USR: a phone SYS: by which brand? USR: LG SYS: okay	SYS: you'd like a ...? USR: a phone SYS: by what brand? USR: LG. SYS: okay	SYS: so would you like a computer? USR: no, a phone. SYS: okay. by which brand? USR: LG. SYS: okay.
Lexical	USR: I want an LG phone. SYS: okay.	USR: Can I have an LG phone? SYS: Sure.	SYS: What do you want to buy? USR: a phone SYS: by which make? USR: LG SYS: Okay.

Figure 1: Some Interactional and Lexical Variations in a Shopping Domain

that fully incremental dialogue systems can be bootstrapped from raw, unannotated example successful dialogues within a particular domain.

The model we present below combines DS-TTR with Reinforcement Learning for incremental word selection, where dialogue management and language generation are treated as one and the same decision/optimisation problem, and where *the corresponding Markov Decision Process is automatically constructed*. Using our implemented system, we demonstrate that using this system one can generalise from very small amounts of raw dialogue data, to a combinatorially large space of interactional variations, including phenomena such as question-answer pairs, over-answering, self- and other-corrections, split-utterances, and clarification interaction, when most of these are not even observed in the original data (see section 4.1).

1.1 Dimensions of Pragmatic Synonymy

There are two important dimensions along which dialogues can vary, but nevertheless, lead to very similar final contexts: interactional, and lexical. Interactional synonymy is analogous to syntactic synonymy - when two distinct sentences are parsed to identical logical forms - except that it occurs not only at the level of a single sentence, but at the dialogue or discourse level - Fig. 1 shows examples. Importantly as we shall show, this type of synonymy can be captured by grammars/models of dialogue context.

Lexical synonymy relations, on the other hand, hold among utterances, or dialogues, when different words (or sequences of words) express meanings that are sufficiently similar in a particular domain or activity - see Fig 1. Unlike syntactic/interactional synonymy relations, lexical ones can often break down when one moves to another domain: lexical synonymy relations are domain specific. Here we do not focus on these, but merely note that lexical synonymy relations can be captured using Distributional Methods (see e.g. Lewis & Steedman (2013)), or methods akin to Eshghi & Lemon (2014) by grounding domain-general semantics into the non-linguistic actions within a domain.

2 Dynamic Syntax (DS) and Type Theory with Records (TTR)

Dynamic Syntax (DS) a is a word-by-word incremental semantic parser/generator, based around the Dynamic Syntax (DS) grammar framework (Cann et al., 2005) especially suited to the fragmentary and highly contextual nature of dialogue. In DS, words are conditional actions - semantic updates; and dialogue is modelled as the interactive and incremental construction of contextual and semantic representations (Eshghi et al., 2015) - see Fig. 2. The contextual representations afforded by DS are of the fine-grained semantic content that is jointly negotiated/agreed upon by the interlocutors, as a result of processing questions and answers, clarification requests, acceptances, self-/other-corrections etc. The upshot of this is that using DS, we can not only track the semantic content of some current turn as it is being constructed (parsed or generated) word by word, but also the context of the conversation as whole, with the latter also encoding the grounded/agreed content of the conversation (see e.g. Fig. 2, and see Eshghi et al. (2015); Purver et al. (2010) for details of the model). Crucially for our model below, the inherent incrementality of DS together with the word-level, as well as cross-turn, parsing constraints it provides, enables the word-by-word exploration of the space of grammatical dialogues, and the semantic and contextual representations that result from them.

These representations are Record Types (RT, see Fig. 2) of Type Theory with Records (TTR, (Cooper, 2005)), useful for incremental specification of utterance content, underspecification, as well as richer representations of the dialogue context (Purver et al., 2010; Purver et al., 2011; Eshghi et al., 2012). For reasons of lack of space, we only note that the TTR calculus provides, in addition to other operations, the *subtype checking operation*, $\sqsubseteq$, among Record Types (RT), and that of the Maximally specific Common Supertype (MCS) of two RTs, which both turn out to be crucial for the automatic construction of our MDP model, and feature checking (for more detail on the DS-TTR Hybrid, see (Eshghi et al., 2012; Hough and Purver, 2014)).

3 The overall BABBLE method

We start with two resources: a) a DS-TTR parser *DS* (either learned from data (Eshghi et al., 2013), or constructed by hand), for incremental language processing, but also, more generally, for tracking the context of the dialogue using Eshghi et al.'s model of feedback (Eshghi et al., 2015; Eshghi, 2015); b) a set D of transcribed successful dialogues in the target domain.

Overall, we will demonstrate the following steps (see (Kalatzis et al., 2016) for more details):

1. Automatically induce the Markov Decision Process (MDP) state space, S, and the dialogue goal, G_D, from D;
2. Automatically define the state encoding function $F : C \rightarrow S$; where $s \in S$ is a (binary) state vector, designed to extract from the current context of the dialogue, the semantic features observed in the example dialogues D; and $c \in C$ is a DS context, viz. a pair of TTR Record Types: $\langle c_p, c_g \rangle$, where c_p is the content of the current, *PENDING* clause as it is being constructed, but not necessarily fully grounded yet; and c_g is the content already jointly built and *GROUNDED* by the interlocutors (loosely following the DGB model of (Ginzburg, 2012)).
3. Define the MDP action set as the *DS* lexicon L (i.e. actions are words);
4. Define the reward function R as reaching G_D, while minimising dialogue length.

We then solve the generated MDP using Reinforcement Learning, with a standard Q-learning method, implemented using BURLAP (MacGlashan, 2015): train a policy $\pi : S \rightarrow L$, where L is the DS Lexicon, and S the state space induced using F. The system is trained in interaction with a (semantic) simulated user, also automatically built from the dialogue data (see (Kalatzis et al., 2016) for details).

The state encoding function F , as shown in Figure 2 the MDP state is a binary vector of size $2 \times |\Phi|$, i.e. twice the number of the RT features. The first half of the state vector contains the grounded features (i.e. agreed by the participants) ϕ_i, while the second half contains the current semantics being incrementally built in the current dialogue utterance. Formally:
$s = \langle F_1(c_p), \ldots, F_m(c_p), F_1(c_g), \ldots, F_m(c_g) \rangle$; where $F_i(c) = 1$ if $c \sqsubseteq \phi_i$, and 0 otherwise. (Recall that $\sqsubseteq$ is the RT subtype relation).

4 Discussion

We have so far induced two prototype dialogue systems, one in an 'electronic shopping' domain (as exemplified by the dialogues in Fig. 1) and another in a 'restaurant-search' domain showing that incremental dialogue systems can be automatically created from small amounts of dialogue transcripts - in this case both systems were induced from a single successful example dialogue.

From a theoretical point of view, this shows that DS-TTR as an incremental model of interactional dynamics, with a domain-specific reward signal/goal is sufficient for certain word sequences becoming routinised and learned as ways of performing specific kinds of speech act within the domain, without any prior, procedural specifications of such actions. Thus, a dialogue system learns not only *what* it needs to do, but also *how* and *when* to do it (e.g. in a 'restaurant-booking' task, it learns to ask "What kind of cuisine would you like?", in a situation where the user says she wants to book a table, but does not provide information about restaurant type): higher-, discourse-level dialogue structure is emergent from interaction in such a setting.

From the practical point of view of dialogue system development, the major benefits of this approach are in (1) more naturally interactive dialogue systems as the resulting systems are incremental and are thus able to handle inherently in-

Grounded Semantics (c_g)	Pending Semantics (c_p)	Dialogue so far
$\left[\begin{array}{ll} x2 & : e \\ e2_{=like} & : es \\ x1_{=USR} & : e \\ p2_{=pres(e2)} & : t \\ p5_{=subj(e2,x1)} & : t \\ p4_{=obj(e2,x2)} & : t \\ p11_{=phone(x2)} & : t \end{array}\right]$	$\left[\begin{array}{ll} x2 & : e \\ e2_{=like} & : es \\ x1_{=USR} & : e \\ p2_{=pres(e2)} & : t \\ p5_{=subj(e2,x1)} & : t \\ p4_{=obj(e2,x2)} & : t \\ p11_{=phone(x2)} & : t \\ x3 & : e \\ p10_{=by(x2,x3)} & : t \\ p9_{=brand(x3)} & : t \\ p10_{=question(x3)} & : t \end{array}\right]$	SYS: What would you like? USR: a phone SYS: by which brand?

$$\text{RT Feature } (\phi_i): \left[\begin{array}{ll} x10 & : e \\ p15_{=brand(x10)} & : t \end{array}\right] \left[\begin{array}{ll} e3_{=like} & : es \\ p2_{=pres(e3)} & : t \end{array}\right] \left[\begin{array}{ll} x10 & : e \\ x8 & : e \\ p14_{=by(x8,x10)} & : t \end{array}\right] \left[\begin{array}{ll} e3_{=like} & : es \\ x5_{=usr} & : e \\ p7_{=subj(e3,x5)} & : t \end{array}\right] \left[\begin{array}{ll} x8 & : e \\ e3_{=like} & : es \\ p6_{=obj(e3,x8)} & : t \end{array}\right]$$

$$\text{State: } \left\langle \begin{array}{lccccc} & F_1\downarrow & F_2\downarrow & F_3\downarrow & F_4\downarrow & F_5\downarrow \\ \text{Pending:} & 1, & 1, & 1, & 1, & 1, \\ \text{Grounded:} & 0, & 1, & 0, & 1, & 1 \end{array} \right\rangle$$

Figure 2: Semantics to MDP state encoding with RT features

cremental dialogue phenomena such as continuations, interruptions, and self-repair (see (Hough, 2015) for the DS-TTR model of self-repair); and (2) reduced development time and cost. To evaluate (2), below we consider the number of different dialogues that can be processed based on only 1 example training dialogue.

4.1 Number of interactional variations captured

Here we establish, as an example of the power of the method implemented, a lower-bound on the number of dialogue variants that can be processed based on training from *only 1 example dialogue*. Consider the training dialogue (which has only 2 'slots' and 4 turns) below:

SYS: What would you like?
USR: a phone
SYS: by which brand?
USR: by Apple

Parsing this dialogue establishes (as described above) a dialogue context that is required for success. The DS grammar is able to parse and generate many variants of the above turns, which lead to the same dialogue contexts being created, and thus also result in successful dialogues. To quantify this, we count the number of interactional variants on the above dialogue which can be parsed/generated by DS, and are thus automatically supported after training the system on this dialogue. Note that we do not take into account possible syntactic and lexical variations here, which would again lead to a large number of variants that the system can handle.

The *DS* grammar can parse several variants of the first turn, including overanswering ("I want an Apple laptop"), self-repair ("I want an Apple laptop, err, no, an LG laptop"), and ellipsis ("a laptop"), whose combinatorics give rise to 16 different ways the user can respond (not counting lexical and syntactic variations). These variations can also happen in the second user turn. If we consider the user turns alone, there are at least 256 variants on the above dialogue which we demonstrate that the trained system can handle. If we also consider similar variations in the two system turns (ellipsis, questions vs. statement, utterance completions, continuation, etc), then we arrive at a lower bound for the number of variations on the training dialogue of 8,192.

This remarkable generative power is due to the generalisation power of the DS grammar, combined with the system's DM/NLG policy which is created by searching through the space of possible (successful) dialogue variants.

5 Conclusion and ongoing work

We show how incremental dialogue systems can be automatically learned from example successful dialogues in a domain, with Dialogue Acts and discourse structure emergent rather specified in advance. This method allows rapid domain transfer – simply collect some example (successful) dialogues in a 'slot-filling' domain, and retrain. At present this is fully automated, and only requires checking that the DS lexicon covers the input data. We are currently applying this method to the problem of learning (visual) word meanings (groundings) from interaction.

References

Ronnie Cann, Ruth Kempson, and Lutz Marten. 2005. *The Dynamics of Language*. Elsevier, Oxford.

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

Arash Eshghi and Oliver Lemon. 2014. How domain-general can we be? Learning incremental dialogue systems without dialogue acts. In *Proceedings of Semdial 2014 (DialWatt)*.

Arash Eshghi, Julian Hough, Matthew Purver, Ruth Kempson, and Eleni Gregoromichelaki. 2012. Conversational interactions: Capturing dialogue dynamics. In S. Larsson and L. Borin, editors, *From Quantification to Conversation: Festschrift for Robin Cooper on the occasion of his 65th birthday*, volume 19 of *Tributes*, pages 325–349. College Publications, London.

Arash Eshghi, Julian Hough, and Matthew Purver. 2013. Incremental grammar induction from child-directed dialogue utterances. In *Proceedings of the 4th Annual Workshop on Cognitive Modeling and Computational Linguistics (CMCL)*, pages 94–103, Sofia, Bulgaria, August. Association for Computational Linguistics.

A. Eshghi, C. Howes, E. Gregoromichelaki, J. Hough, and M. Purver. 2015. Feedback in conversation as incremental semantic update. In *Proceedings of the 11th International Conference on Computational Semantics (IWCS 2015)*, London, UK. Association for Computational Linguisitics.

Arash Eshghi. 2015. DS-TTR: An incremental, semantic, contextual parser for dialogue. In *Proceedings of Semdial 2015 (goDial), the 19th workshop on the semantics and pragmatics of dialogue*.

Jonathan Ginzburg. 2012. *The Interactive Stance: Meaning for Conversation*. Oxford University Press.

Patrick G. T. Healey. 2008. Interactive misalignment: The role of repair in the development of group sub-languages. In R. Cooper and R. Kempson, editors, *Language in Flux*. College Publications.

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. 2015. *Modelling Incremental Self-Repair Processing in Dialogue*. Ph.D. thesis, Queen Mary University of London.

Dimitrios Kalatzis, Arash Eshghi, and Oliver Lemon. 2016. Bootstrapping incremental dialogue systems: using linguistic knowledge to learn from minimal data. In *Proceedings of the NIPS 2016 workshop on Learning Methods for Dialogue*, Barcelona.

Ruth Kempson, Wilfried Meyer-Viol, and Dov Gabbay. 2001. *Dynamic Syntax: The Flow of Language Understanding*. Blackwell.

Ruth Kempson, Ronnie Cann, Arash Eshghi, Eleni Gregoromichelaki, and Matthew Purver. 2015. Ellipsis. In Shalom Lappin and Chris Fox, editors, *The Handbook of Contemporary Semantics*. Wiley-Blackwell.

Ruth Kempson, Ronnie Cann, Eleni Gregoromichelaki, and Stergios Chatzikiriakidis. 2016. Language as mechanisms for interaction. *Theoretical Linguistics*, 42(3-4):203–275.

Staffan Larsson and Robin Cooper. 2008. Corrective feedback and concept updates. In *Proceedings of The 2nd Swedish Language Technology Conference (SLTC-08)*.

Mike Lewis and Mark Steedman. 2013. Combined distributional and logical semantics. *Transactions of the Association for Computational Linguistics*, 1:179–192.

James MacGlashan. 2015. Burlap: Brown-umbc reinforcement learning and planning. In *http://burlap.cs.brown.edu/*.

G. Mills and E. Gregoromichelaki. 2010. Establishing coherence in dialogue: sequentiality, intentions and negotiation. In *Proceedings of SemDial (PozDial)*.

Gregory Mills. 2011. The emergence of procedural conventions in dialogue. In *IProceedings of the 33rd Annual Conference of the Cognitive Science Society*.

Matthew Purver, Eleni Gregoromichelaki, Wilfried Meyer-Viol, and Ronnie Cann. 2010. Splitting the 'I's and crossing the 'You's: Context, speech acts and grammar. In P. Łupkowski and M. Purver, editors, *Aspects of Semantics and Pragmatics of Dialogue. SemDial 2010, 14th Workshop on the Semantics and Pragmatics of Dialogue*, pages 43–50, Poznań, June. Polish Society for Cognitive Science.

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.