

Deep Learning for Conversational AI

NAACL 2018 Tutorial

Pei-Hao Su¹, Nikola Mrkšić¹, Iñigo Casanueva¹ Ivan Vulić^{1,2},

¹ PolyAI

² University of Cambridge

{eddy-su, nikola, inigo, ivan}@poly-ai.com

1 Objectives

Spoken Dialogue Systems (SDS) have great commercial potential as they promise to revolutionise the way in which humans interact with machines. The advent of deep learning led to substantial developments in this area of NLP research, and the goal of this tutorial is to familiarise the research community with the recent advances in what some call the most difficult problem in NLP.

From a research perspective, the design of spoken dialogue systems provides a number of significant challenges, as these systems depend on: **a)** solving several difficult NLP and decision-making tasks; and **b)** combining these into a functional dialogue system pipeline. A key long-term goal of dialogue system research is to enable open-domain systems that can converse about arbitrary topics and assist humans with completing a wide range of tasks. Furthermore, such systems need to autonomously learn on-line to improve their performance and recover from errors using both signals from their environment and from implicit and explicit user feedback. While the design of such systems has traditionally been modular, domain and language-specific, advances in deep learning have alleviated many of the design problems.

The main purpose of this tutorial is to encourage dialogue research in the NLP community by providing the research background, a survey of available resources, and giving key insights to application of state-of-the-art SDS methodology into industry-scale systems. We plan to introduce researchers to the pipeline framework for modelling goal-oriented dialogue systems, which includes three key components: **1)** Language Understanding; **2)** Dialogue Management; and **3)** Language Generation. The differences between goal-oriented dialogue systems and chat-bot style conversational agents will be explained in order to show the motivation behind the design of both, with the main focus on the pipeline SDS framework. For each key component, we will define the research problem, provide a brief literature review and introduce the current state-of-the-art approaches. Complementary resources (e.g. available datasets and toolkits) will also be discussed. Finally, future work, outstanding challenges, and current industry practices will be presented. All of the presented material will be made available online for future reference.

2 Tutorial Overview

2.1 Part I: Introduction to Statistical Dialogue Systems

The modular architecture of a goal-oriented spoken dialogue system will be introduced and the range of approaches available for each component, from rule-based to (increasingly) statistical methods will be discussed. The key architectural requirements of goal-oriented spoken dialogue systems will be emphasised and the differences to chat-bot style systems will be explained. Based on this introduction, the key challenges for machine learning will be identified and the options available for moving from the current generation of limited domain systems to fully open-domain conversational agents will be presented. A particular focus will be on learning techniques which enable a system to incrementally increase its naturalness, robustness and coverage over time by interaction on-line with real users.

2.2 Part II: Language Understanding and Dialogue State Tracking

In this part, we will present the *language understanding* module, which is the first component of the SDS pipeline. This module takes as input the users' spoken/written utterances and converts them to an abstract

representation that the downstream dialogue management component can (learn to) operate and reason with. We plan to give an overview of: **a)** rule-based systems; **b)** conventional approaches which split the language understanding problem into Spoken Language Understanding (SLU) and Belief Tracking (BT); and **c)** the most recent models which learn to perform the two tasks jointly. In presenting these approaches, we will focus on two key challenges: **1)** mitigating the effect of automatic speech recognition (ASR) errors; and **2)** dealing with the ambiguity introduced by the linguistic variations available to users in expressing their intentions in various dialogue contexts. Finally, the impact that recent advances in representation learning have had on language understanding will be discussed: these very recent fully statistical approaches hold promise to drive progress in domain adaptation for dialogue systems, both across *different dialogue domains* and *across different languages*.

2.3 Part III: Dialogue Management and Reinforcement Learning

This part will focus on how the turn-taking process is managed in an SDS. The role of the dialogue manager is to map the inferred belief state into a meaningful system action, accounting for the uncertainty propagated from the upstream components. The basics of reinforcement learning (RL) will first be introduced, followed by its practical application to the dialogue management task. We will cover: **a)** tabular-based RL, which is only tractable for simplified problems; **b)** Gaussian process-based RL, which enables fast policy learning; and **c)** deep (neural network-based) RL which has the potential to eliminate the explicit need for hand-crafted feature engineering. We will also show how a dialogue policy can be trained off-line on corpora via supervised learning, and on-line with a user simulator or through direct interaction with human users using RL. When learning with human users, task success can be hard to measure and user feedback is often unreliable and difficult to obtain. To deal with this, a literature review will be covered, and especially Gaussian Process estimators will be presented which minimise the burden on users of providing explicit feedback and mitigate the problems of noisy user feedback.

2.4 Part IV: Response Generation and End-to-End Dialogue Modelling

In this part of the tutorial, methods of statistical language generation will be presented, which map abstract system dialogue acts back into natural language. We will first explain how Recurrent Neural Network language models can be used to generate sentences, and how a structured meaning representation such as a dialogue act can be used to condition the generation process. We will also show that attention and gating mechanisms can be used to better model internal content selection and prevent semantic repetitions, which leads to more coherent and natural responses. Next, we will frame the response generation task in a broader context by treating end-to-end dialogue modelling as a conditional response generation task. We will draw connections between this approach and other chat-bot style conversational agents, showing that explicit language grounding is crucial for goal-oriented dialogue response generation. Finally, we will address the difficulty of collecting corpora for training the SDS systems in general and the generation module in particular. We will also discuss how a pipelined Wizard-of-Oz data collection framework can be used to collect significant amounts of data at acceptably low cost.

2.5 Part V: SDS Systems in Conversational AI Applications and Current Challenges

The conversational interfaces hold promise to construct a fully natural way of communication between the human and the machine. In the final part of the tutorial, we will frame modern dialogue research sub-problems in the context of broader NLP research: we will outline once more recent trends in the development of modular dialogue systems, explaining how these complement the long-term goals of broader AI research. We will also discuss the current status of deploying SDS systems beyond the core academic research: we will analyse their impact and usefulness in industry-scale applications and their potential for conversational AI. We will place special emphasis on the key challenges and open questions in our pursuit of creating open-domain statistical dialogue systems across different languages.

We will conclude by listing publicly available software packages and implementations, available training datasets and evaluation protocols, and sketching future research avenues in this domain.

3 Structure and Summary

Part I: Introduction to Statistical Dialogue Systems (30 minutes)

- Overview of statistical dialogue systems: related work, current trends.
- Pipeline approaches vs. chat-bot style conversational agents.
- Long-term SDS goals and its relation to conversational AI.

Part II: Language Understanding and Dialogue State Tracking (40 minutes)

- Survey of approaches for performing language understanding in spoken dialogue systems.
- The impact of advances in semantic representation learning on understanding in dialogue systems.
- Fully statistical language understanding: towards open-domain SDS systems across languages.

Part III: Dialogue Management and Reinforcement Learning (40 minutes)

- Reinforcement learning approaches for managing the turn-taking dialogue task.
- Dialogue evaluation and reward estimation for practical policy learning.

Part IV: Response Generation and End-to-End Dialogue Modelling (40 minutes)

- Response generation from structured meaning representations.
- End-to-End dialogue modelling: Models, evaluations, and data collection.

Part V: SDS Systems in Conversational AI Applications and Current Challenges (30 minutes)

- Publicly available software packages and implementations, available training datasets and evaluation protocols.
- SDS systems and conversational interfaces: research vs. industry demands.
- Key challenges and open questions in the pursuit of creating open-domain statistical dialogue systems across different languages.

4 About the Speakers

Pei-Hao Su <https://eddy0613.github.io/>; email: eddysu@poly-ai.com

Pei-Hao (Eddy) Su is a co-founder and CTO of PolyAI, a London-based startup looking to use the latest developments in NLP to create a general machine learning platform for deploying spoken dialogue systems. He holds a PhD from the Dialogue Systems group, University of Cambridge, where he worked under the supervision of Professor Steve Young. His research interests centre on applying deep learning, reinforcement learning and Bayesian approaches to dialogue management and reward estimation, with the aim of building systems that can learn directly from human interaction. He has given several invited talks at academia and industry such as Apple, Microsoft, General Motor and DeepHack.Turing. He received the best student paper award at ACL 2016.

Nikola Mrkšić mi.eng.cam.ac.uk/~nm480; email: nikola@poly-ai.com

Nikola Mrkšić is a co-founder and CEO of PolyAI, a London-based startup looking to use the latest developments in NLP to create a general machine learning platform for deploying spoken dialogue systems. He holds a PhD from the Dialogue Systems group, University of Cambridge, where he worked under

the supervision of Professor Steve Young. His research is focused on belief tracking in human-machine dialogue, specifically in moving towards building open-domain, cross-lingual language understanding models that are fully data-driven. He is also interested in deep learning, semantics, Bayesian nonparametrics, unsupervised and semi-supervised learning. He previously gave a tutorial on word vector space specialisation at EACL 2017, and will teach a course on the same topic at ESSLLI 2018. He also gave invited talks at the REWORK AI Personal Assistant summit and the Chatbot Summit.

Iñigo Casanueva <http://mi.eng.cam.ac.uk/~ic340/>; email: inigo@poly-ai.com

Iñigo Casanueva is a Machine Learning engineer at PolyAI, a London-based startup looking to use the latest developments in NLP to create a general machine learning platform for deploying spoken dialogue systems. He got his PhD from the University of Sheffield and later he worked as Research Assistant in the Dialogue Systems group, University of Cambridge. His main research interest focuses on increasing the scalability of machine learning based dialogue management, looking for methods to make deep learning and/or reinforcement learning applicable to real world dialogue management tasks. He has published several papers on the topic, two of them nominated to best paper award.

Ivan Vulić <https://sites.google.com/site/ivanvulic/>; email: iv250@cam.ac.uk

Ivan Vulić is a Senior Research Associate in the Language Technology Lab at the University of Cambridge. He holds a PhD from KU Leuven, obtained summa cum laude. Ivan is interested in representation learning, human language understanding, distributional, lexical, and multi-modal semantics in monolingual and multilingual contexts, and transfer learning for enabling cross-lingual NLP applications. He co-lectured a tutorial on monolingual and multilingual topic models and applications at ECIR 2013 and WSDM 2014, a tutorial on word vector space specialisation at EACL 2017, and a tutorial on cross-lingual word representations at EMNLP 2017. He will lecture a course on word vector space specialisation at ESSLLI 2018. He has given invited talks at academia and industry such as Apple Inc., University of Cambridge, UCL, University of Copenhagen, Paris-Saclay, and Bar-Ilan University.

5 Other Information

- Previous tutorial editions: N/A
- Audience size (estimate): 100-120
- Special requirements: None
- Acceptable venues: 1. EMNLP, 2. NAACL-HLT, 3. ACL (sorted starting with the most preferable venue)

Selected Bibliography

- Antoine Bordes and Jason Weston. 2016. Learning end-to-end goal-oriented dialog. In *Proceedings of ICLR (Conference Track)*, volume abs/1605.07683.
- Paweł Budzianowski, Stefan Ultes, Pei-Hao Su, Nikola Mrkšić, Tsung-Hsien Wen, Iñigo Casanueva, Lina M. Rojas Barahona, and Milica Gasic. 2017. Sub-domain modelling for dialogue management with hierarchical reinforcement learning. In *Proceedings of SIGDIAL*, pages 86–92, August.
- Nina Dethlefs, Helen F. Hastie, Heriberto Cuayáhuítl, Yanchao Yu, Verena Rieser, and Oliver Lemon. 2016. Information density and overlap in spoken dialogue. *Computer Speech & Language*, 37:82–97.
- Ondřej Dušek and Filip Jurcicek. 2015. Training a natural language generator from unaligned data. In *Proceedings of ACL*, pages 451–461.
- Milica Gašić, Dilek Hakkani-Tür, and Asli Çelikyilmaz. 2017a. Spoken language understanding and interaction: Machine learning for human-like conversational systems. *Computer Speech & Language*, 46:249–251.

- Milica Gašić, Nikola Mrkšić, Lina Maria Rojas-Barahona, Pei-Hao Su, Stefan Ultes, David Vandyke, Tsung-Hsien Wen, and Steve J. Young. 2017b. Dialogue manager domain adaptation using gaussian process reinforcement learning. *Computer Speech & Language*, 45:552–569.
- Matthew Henderson, Milica Gašić, Blaise Thomson, Pirros Tsiakoulis, Kai Yu, and Steve J. Young. 2012. Discriminative spoken language understanding using word confusion networks. In *Proceedings of IEEE SLT*, pages 176–181.
- Matthew Henderson, Blaise Thomson, and Steve Young. 2013. Deep neural network approach for the Dialog State Tracking Challenge. In *Proceedings of SIGDIAL*, pages 467–471.
- Matthew Henderson, Blaise Thomson, and Steve Young. 2014. Word-based dialog state tracking with recurrent neural networks. In *Proceedings of SIGDIAL*, pages 292–299.
- Matthew Henderson. 2015. *Discriminative Methods for Statistical Spoken Dialogue Systems*. Ph.D. thesis, University of Cambridge.
- Youngsoo Jang, Jiyeon Ham, Byung-Jun Lee, Youngjae Chang, and Kee-Eung Kim. 2016. Neural dialog state tracker for large ontologies by attention mechanism. In *IEEE SLT*, pages 531–537.
- Seokhwan Kim, Luis Fernando D’Haro, Rafael E. Banchs, Jason D. Williams, Matthew Henderson, and Koichiro Yoshino. 2016. The fifth dialog state tracking challenge. In *Proceedings of IEEE SLT*, pages 511–517.
- Jiwei Li, Will Monroe, Alan Ritter, Dan Jurafsky, Michel Galley, and Jianfeng Gao. 2016. Deep reinforcement learning for dialogue generation. In *Proceedings of EMNLP*, pages 1192–1202, November.
- Jiwei Li, Will Monroe, Tianlin Shi, Sébastien Jean, Alan Ritter, and Dan Jurafsky. 2017. Adversarial learning for neural dialogue generation. In *Proceedings of EMNLP*, pages 2157–2169.
- Bing Liu and Ian Lane. 2016. Joint online spoken language understanding and language modeling with recurrent neural networks. In *Proceedings of SIGDIAL*, pages 22–30.
- Chia-Wei Liu, Ryan Lowe, Iulian Serban, Mike Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. How NOT to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. In *Proceedings of EMNLP*, pages 2122–2132.
- Nikola Mrkšić, Diarmuid Ó Séaghdha, Blaise Thomson, Milica Gašić, Pei-Hao Su, David Vandyke, Tsung-Hsien Wen, and Steve Young. 2015. Multi-domain dialog state tracking using recurrent neural networks. In *Proceedings of ACL*, pages 794–799.
- Nikola Mrkšić, Diarmuid Ó Séaghdha, Blaise Thomson, Milica Gašić, Lina M. Rojas-Barahona, Pei-Hao Su, David Vandyke, Tsung-Hsien Wen, and Steve Young. 2016. Counter-fitting word vectors to linguistic constraints. In *Proceedings of NAACL-HLT*, pages 142–148.
- Nikola Mrkšić, Diarmuid Ó Séaghdha, Tsung-Hsien Wen, Blaise Thomson, and Steve Young. 2017a. Neural Belief Tracker: Data-driven dialogue state tracking. In *Proceedings of ACL*, pages 1777–1788.
- Nikola Mrkšić, Ivan Vulić, Diarmuid Ó Séaghdha, Ira Leviant, Roi Reichart, Milica Gašić, Anna Korhonen, and Steve Young. 2017b. Semantic specialization of distributional word vector spaces using monolingual and cross-lingual constraints. *Transactions of the Association for Computational Linguistics*, 5:309–324.
- Julien Perez and Fei Liu. 2017. Dialog state tracking, a machine reading approach using memory network. In *Proceedings of EACL*, pages 305–314.
- Verena Rieser and Oliver Lemon. 2011. Learning and evaluation of dialogue strategies for new applications: Empirical methods for optimization from small data sets. *Computational Linguistics*, 37(1):153–196.
- Iulian Vlad Serban, Alessandro Sordani, Yoshua Bengio, Aaron C. Courville, and Joelle Pineau. 2016. Building end-to-end dialogue systems using generative hierarchical neural network models. In *Proceedings of AAAI*, pages 3776–3784.
- Pei-Hao Su, David Vandyke, Milica Gašić, Nikola Mrkšić, Tsung-Hsien Wen, and Steve Young. 2015. Reward shaping with recurrent neural networks for speeding up on-line policy learning in spoken dialogue systems. In *Proceedings of SIGDIAL*, pages 417–421, September.
- Pei-Hao Su, Milica Gašić, Nikola Mrkšić, Lina M. Rojas Barahona, Stefan Ultes, David Vandyke, Tsung-Hsien Wen, and Steve Young. 2016. On-line active reward learning for policy optimisation in spoken dialogue systems. In *Proceedings of ACL*, pages 2431–2441.

- Pei-Hao Su, Paweł Budzianowski, Stefan Ultes, Milica Gasic, and Steve Young. 2017. Sample-efficient actor-critic reinforcement learning with supervised data for dialogue management. In *Proceedings of SIGDIAL*, pages 147–157.
- Blaise Thomson. 2009. *Statistical methods for spoken dialogue management*. Ph.D. thesis, University of Cambridge.
- Stefan Ultes, Lina M. Rojas Barahona, Pei-Hao Su, David Vandyke, Dongho Kim, Iñigo Casanueva, Paweł Budzianowski, Nikola Mrkšić, Tsung-Hsien Wen, Milica Gašić, and Steve Young. 2017. PyDial: A multi-domain statistical dialogue system toolkit. In *Proceedings of ACL Demos*, pages 73–78.
- David Vandyke, Pei-hao Su, Milica Gašić, Nikola Mrkšić, Tsung-Hsien Wen, and Steve Young. 2015. Multi-domain dialogue success classifiers for policy training. In *Proceedings of ASRU*, pages 763–770.
- Oriol Vinyals and Quoc V. Le. 2015. A neural conversational model. *CoRR*, abs/1506.05869.
- Miroslav Vodolán, Rudolf Kadlec, and Jan Kleindienst. 2017. Hybrid dialog state tracker with ASR features. In *Proceedings of EACL*, pages 205–210.
- Ivan Vulić, Nikola Mrkšić, Roi Reichart, Diarmuid Ó Séaghdha, Steve Young, and Anna Korhonen. 2017. Morph-fitting: Fine-tuning word vector spaces with simple language-specific rules. In *Proceedings of ACL*, pages 56–68.
- Zhuoran Wang, Tsung-Hsien Wen, Pei-Hao Su, and Yannis Stylianou. 2015. Learning domain-independent dialogue policies via ontology parameterisation. In *Proceedings of SIGDIAL*, pages 412–416.
- Joseph Weizenbaum. 1966. ELIZA - A computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9(1):36–45.
- Tsung-Hsien Wen, Milica Gašić, Dongho Kim, Nikola Mrkšić, Pei-Hao Su, David Vandyke, and Steve Young. 2015a. Stochastic language generation in dialogue using recurrent neural networks with convolutional sentence reranking. In *Proceedings of SIGDIAL*, pages 275–284.
- Tsung-Hsien Wen, Milica Gašić, Nikola Mrkšić, Pei-Hao Su, David Vandyke, and Steve Young. 2015b. Semantically conditioned LSTM-based natural language generation for spoken dialogue systems. In *Proceedings of EMNLP*, pages 1711–1721.
- Tsung-Hsien Wen, Milica Gašić, Nikola Mrkšić, Lina M. Rojas-Barahona, Pei-Hao Su, David Vandyke, and Steve Young. 2016. Multi-domain neural network language generation for spoken dialogue systems. In *Proceedings of NAACL-HLT*, pages 120–129.
- Tsung-Hsien Wen, Yishu Miao, Phil Blunsom, and Steve J. Young. 2017a. Latent intention dialogue models. In *Proceedings of ICML*, pages 3732–3741.
- Tsung-Hsien Wen, David Vandyke, Nikola Mrkšić, Milica Gašić, Lina M. Rojas Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2017b. A network-based end-to-end trainable task-oriented dialogue system. In *Proceedings of EACL*, pages 438–449.
- Jason Williams, Antoine Raux, and Matthew Henderson. 2016. The dialog state tracking challenge series: A review. *Dialogue&Discourse*, 7(3):4–33.
- Jason Williams. 2013. Multi-domain learning and generalization in dialog state tracking. In *Proceedings of SIGDIAL*, pages 433–441.
- Steve Young, Milica Gašić, Blaise Thomson, and Jason D. Williams. 2013. POMDP-Based statistical spoken dialog systems: A review. *Proceedings of the IEEE*, 101(5):1160–1179.
- Steve Young. 2010a. Cognitive User Interfaces. *IEEE Signal Processing Magazine*.
- Steve Young. 2010b. Still talking to machines (cognitively speaking). In *Proceedings of INTERSPEECH*, pages 1–10.