

Neural Approaches to Conversational AI

**Question Answering, Task-oriented
Dialogues and Social Chatbots**

Other titles in Foundations and Trends® in Information Retrieval

Web Forum Retrieval and Text Analytics: A Survey

Doris Hoogeveen, Li Wang, Timothy Baldwin and Karin M. Verspoor

ISBN: 978-1-68083-350-8

Display Advertising with Real-Time Bidding (RTB) and Behavioural Targeting

Jun Wang, Weinan Zhang and Shuai Yuan

ISBN: 978-1-68083-310-2

Applications of Topic Models

Jordan Boyd-Graber, Yuening Hu and David Mimno

ISBN: 978-1-68083-308-9

Searching the Enterprise

Udo Kruschwitz and Charlie Hull

ISBN: 978-1-68083-304-1

Neural Approaches to Conversational AI

Question Answering, Task-oriented Dialogues and Social Chatbots

Jianfeng Gao

Microsoft Research
jfgao@microsoft.com

Michel Galley

Microsoft Research
mgalley@microsoft.com

Lihong Li

Google Brain
lihong@google.com

now

the essence of knowledge

Boston — Delft

Foundations and Trends[®] in Information Retrieval

Published, sold and distributed by:

now Publishers Inc.
PO Box 1024
Hanover, MA 02339
United States
Tel. +1-781-985-4510
www.nowpublishers.com
sales@nowpublishers.com

Outside North America:

now Publishers Inc.
PO Box 179
2600 AD Delft
The Netherlands
Tel. +31-6-51115274

The preferred citation for this publication is

J. Gao, M. Galley and L. Li. *Neural Approaches to Conversational AI*. Foundations and Trends[®] in Information Retrieval, vol. 13, no. 2-3, pp. 127–298, 2019.

ISBN: 978-1-68083-553-3

© 2019 J. Gao, M. Galley and L. Li

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, mechanical, photocopying, recording or otherwise, without prior written permission of the publishers.

Photocopying. In the USA: This journal is registered at the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923. Authorization to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by now Publishers Inc for users registered with the Copyright Clearance Center (CCC). The 'services' for users can be found on the internet at: www.copyright.com

For those organizations that have been granted a photocopy license, a separate system of payment has been arranged. Authorization does not extend to other kinds of copying, such as that for general distribution, for advertising or promotional purposes, for creating new collective works, or for resale. In the rest of the world: Permission to photocopy must be obtained from the copyright owner. Please apply to now Publishers Inc., PO Box 1024, Hanover, MA 02339, USA; Tel. +1 781 871 0245; www.nowpublishers.com; sales@nowpublishers.com

now Publishers Inc. has an exclusive license to publish this material worldwide. Permission to use this content must be obtained from the copyright license holder. Please apply to now Publishers, PO Box 179, 2600 AD Delft, The Netherlands, www.nowpublishers.com; e-mail: sales@nowpublishers.com

Foundations and Trends[®] in Information Retrieval

Volume 13, Issue 2-3, 2019

Editorial Board

Editors-in-Chief

Maarten de Rijke
University of Amsterdam
The Netherlands

Yiqun Liu
Tsinghua University
China

Editors

Ben Carterette
University of Delaware

Charles L.A. Clarke
University of Waterloo

Claudia Hauff
Delft University of Technology

Diane Kelly
University of Tennessee

Doug Oard
University of Maryland

Ellen M. Voorhees
*National Institute of Standards and
Technology*

Fabrizio Sebastiani
Consiglio Nazionale delle Ricerche, Italy

Hang Li
Bytedance Technology

Ian Ruthven
University of Strathclyde, Glasgow

Jaap Kamps
University of Amsterdam

James Allan
University of Massachusetts, Amherst

Jian-Yun Nie
Université de Montreal

Jimmy Lin
University of Maryland

Leif Azzopardi
University of Glasgow

Lynda Tamine
University of Toulouse

Mark D. Smucker
University of Waterloo

Rodrygo Luis Teodoro Santos
Universidade Federal de Minas Gerais

Ryen White
Microsoft Research

Shane Culpepper
RMIT

Soumen Chakrabarti
Indian Institute of Technology

Tie-Yan Liu
Microsoft Research

Editorial Scope

Topics

Foundations and Trends[®] in Information Retrieval publishes survey and tutorial articles in the following topics:

- Applications of IR
- Architectures for IR
- Collaborative filtering and recommender systems
- Cross-lingual and multilingual IR
- Distributed IR and federated search
- Evaluation issues and test collections for IR
- Formal models and language models for IR
- IR on mobile platforms
- Indexing and retrieval of structured documents
- Information categorization and clustering
- Information extraction
- Information filtering and routing
- Metasearch, rank aggregation and data fusion
- Natural language processing for IR
- Performance issues for IR systems, including algorithms, data structures, optimization techniques, and scalability
- Question answering
- Summarization of single documents, multiple documents, and corpora
- Text mining
- Topic detection and tracking
- Usability, interactivity, and visualization issues in IR
- User modelling and user studies for IR
- Web search

Information for Librarians

Foundations and Trends[®] in Information Retrieval, 2019, Volume 13, 5 issues. ISSN paper version 1554-0669. ISSN online version 1554-0677. Also available as a combined paper and online subscription.

Contents

1	Introduction	2
1.1	Who Should Read this Paper?	4
1.2	Dialogue: What Kinds of Problems?	5
1.3	A Unified View: Dialogue as Optimal Decision Making	7
1.4	The Transition of NLP to Neural Approaches	9
2	Machine Learning Background	13
2.1	Machine Learning Basics	13
2.2	Deep Learning	17
2.3	Reinforcement Learning	23
3	Question Answering and Machine Reading Comprehension	29
3.1	Knowledge Base	30
3.2	Semantic Parsing for KB-QA	31
3.3	Embedding-based Methods	32
3.4	Multi-Step Reasoning on KB	34
3.5	Conversational KB-QA Agents	40
3.6	Machine Reading for Text-QA	44
3.7	Neural MRC Models	47
3.8	Conversational Text-QA Agents	53

4	Task-oriented Dialogue Systems	56
4.1	Overview	57
4.2	Evaluation and User Simulation	62
4.3	Natural Language Understanding and Dialogue State Tracking	68
4.4	Dialogue Policy Learning	73
4.5	Natural Language Generation	83
4.6	End-to-end Learning	86
4.7	Further Remarks	88
5	Fully Data-Driven Conversation Models and Social Bots	90
5.1	End-to-End Conversation Models	91
5.2	Challenges and Remedies	95
5.3	Grounded Conversation Models	103
5.4	Beyond Supervised Learning	105
5.5	Data	107
5.6	Evaluation	109
5.7	Open Benchmarks	112
6	Conversational AI in Industry	114
6.1	Question Answering Systems	114
6.2	Task-oriented Dialogue Systems (Virtual Assistants)	118
6.3	Chatbots	121
7	Conclusions and Research Trends	124
	Acknowledgements	128
	References	129

Neural Approaches to Conversational AI

Jianfeng Gao¹, Michel Galley² and Lihong Li³

¹*Microsoft Research; jfgao@microsoft.com*

²*Microsoft Research; mgalley@microsoft.com*

³*Google Brain; lihong@google.com*

ABSTRACT

The present paper surveys neural approaches to conversational AI that have been developed in the last few years. We group conversational systems into three categories: (1) question answering agents, (2) task-oriented dialogue agents, and (3) chatbots. For each category, we present a review of state-of-the-art neural approaches, draw the connection between them and traditional approaches, and discuss the progress that has been made and challenges still being faced, using specific systems and models as case studies.

1

Introduction

Developing an intelligent dialogue system¹ that not only emulates human conversation, but also answers questions on topics ranging from latest news about a movie star to Einstein’s theory of relativity, and fulfills complex tasks such as travel planning, has been one of the longest running goals in AI. The goal has remained elusive until recently. We are now observing promising results both in academia and industry, as large amounts of conversational data become available for training, and the breakthroughs in deep learning (DL) and reinforcement learning (RL) are applied to conversational AI.

Conversational AI is fundamental to natural user interfaces. It is a rapidly growing field, attracting many researchers in the Natural Language Processing (NLP), Information Retrieval (IR) and Machine Learning (ML) communities. For example, SIGIR 2018 has created a new track of Artificial Intelligence, Semantics, and Dialog to bridge research in AI and IR, especially targeting Question Answering (QA), deep semantics and dialogue with intelligent agents.

¹“Dialogue systems” and “conversational AI” are often used interchangeably in the scientific literature. The difference is reflective of different traditions. The former term is more general in that a dialogue system might be purely rule-based rather than AI-based.

Recent years have seen the rise of a small industry of tutorials and survey papers on deep learning and dialogue systems. Yih *et al.* (2015b), Yih *et al.* (2016), and Gao (2017) reviewed deep learning approaches for a wide range of IR and NLP tasks, including dialogues. Chen *et al.* (2017e) presented a tutorial on dialogues, with a focus on task-oriented agents. Serban *et al.* (2015; 2018) surveyed public dialogue datasets that can be used to develop conversational agents. Chen *et al.* (2017b) reviewed popular deep neural network models for dialogues, focusing on supervised learning approaches. The present work substantially expands the scope of Chen *et al.* (2017b) and Serban *et al.* (2015) by going beyond data and supervised learning to provide what we believe is the first survey of neural approaches to conversational AI, targeting NLP and IR audiences.² Its contributions are:

- We provide a comprehensive survey of the neural approaches to conversational AI that have been developed in the last few years, covering QA, task-oriented and social bots with a unified view of optimal decision making.
- We draw connections between modern neural approaches and traditional approaches, allowing us to better understand why and how the research has evolved and to shed light on how we can move forward.
- We present state-of-the-art approaches to training dialogue agents using both supervised and reinforcement learning.
- We sketch out the landscape of conversational systems developed in the research community and released in industry, demonstrating via case studies the progress that has been made and the challenges that we are still facing.

²One important topic of conversational AI that we do not cover is Spoken Language Understanding (SLU). SLU systems are designed to extract the meaning from speech utterances and their application are vast, ranging from voice search in mobile devices to meeting summarization. The present work does encompass many Spoken Dialogue Systems – for example Young *et al.* (2013) – but does not focus on components related to speech. We refer readers to Tur and De Mori (2011) for a survey of SLU.

1.1 Who Should Read this Paper?

This paper is based on tutorials given at the SIGIR and ACL conferences in 2018 (Gao *et al.*, 2018a; Gao *et al.*, 2018b), with the IR and NLP communities as the primary target audience. However, audiences with other backgrounds (such as machine learning) will also find it an accessible introduction to conversational AI with numerous pointers, especially to recently developed neural approaches.

We hope that this paper will prove a valuable resource for students, researchers, and software developers. It provides a unified view, as well as a detailed presentation of the important ideas and insights needed to understand and create modern dialogue agents that will be instrumental to making world knowledge and services accessible to millions of users in ways that seem natural and intuitive.

This survey is structured as follows:

- The rest of this chapter introduces dialogue tasks and presents a unified view in which open-domain dialogue is formulated as an optimal decision making process.
- Chapter 2 introduces basic mathematical tools and machine learning concepts, and reviews recent progress in the deep learning and reinforcement learning techniques that are fundamental to developing neural dialogue agents.
- Chapter 3 describes question answering (QA) agents, focusing on neural models for knowledge-base QA and machine reading comprehension (MRC).
- Chapter 4 describes task-oriented dialogue agents, focusing on applying deep reinforcement learning to dialogue management.
- Chapter 5 describes social chatbots, focusing on fully data-driven neural approaches to end-to-end generation of conversational responses.
- Chapter 6 gives a brief review of several conversational systems in industry.
- Chapter 7 concludes the paper with a discussion of research trends.

Turn	Dialogue
1	<i>usr</i> : Good morning!
	<i>agt</i> : Good morning, I hope you had a great weekend.
2	<i>usr</i> : not really.
	<i>agt</i> : You sound depressed. What can I do for you?
3	<i>usr</i> : I'd like to know where sales are lagging behind our forecast?
	<i>agt</i> : The worst region is [country], where sales are 15% below projections.
4	<i>usr</i> : Do you know why?
	<i>agt</i> : The forecast for [product] growth was overly optimistic.
5	<i>usr</i> : How can we turn this around?
	<i>agt</i> : Here are the 10 customers in [country] with the most growth potential, per our CRM model.
6	<i>usr</i> : Can you set up a meeting with the CTO of [company]?
	<i>agt</i> : Yes, I've set up a meeting with [person name] for next month when you are in [location].
7	<i>usr</i> : Thanks!

Figure 1.1: A human-agent dialogue during the process of making a business decision. (*usr*: user, *agt*: agent) The dialogue consists of multiple segments of different types. Turns 1 and 2 are a social chat segment. Turns 3 to 5 are a QA segment. Turns 6 and 7 are a task-completion segment.

1.2 Dialogue: What Kinds of Problems?

Fig. 1.1 shows a human-agent dialogue during the process of making a business decision. The example illustrates the kinds of problems a dialogue system is expected to solve:

- **question answering:** the agent needs to provide concise, direct answers to user queries based on rich knowledge drawn from various data sources including text collections such as Web documents and pre-compiled knowledge bases such as sales and marketing datasets, as the example shown in Turns 3 to 5 in Fig. 1.1.
- **task completion:** the agent needs to accomplish user tasks ranging from restaurant reservation to meeting scheduling (e.g., Turns 6 to 7 in Fig. 1.1), and to business trip planning.
- **social chat:** the agent needs to converse seamlessly and appropriately with users — like a human as in the Turing test — and provide useful recommendations (e.g., Turns 1 to 2 in Fig. 1.1).

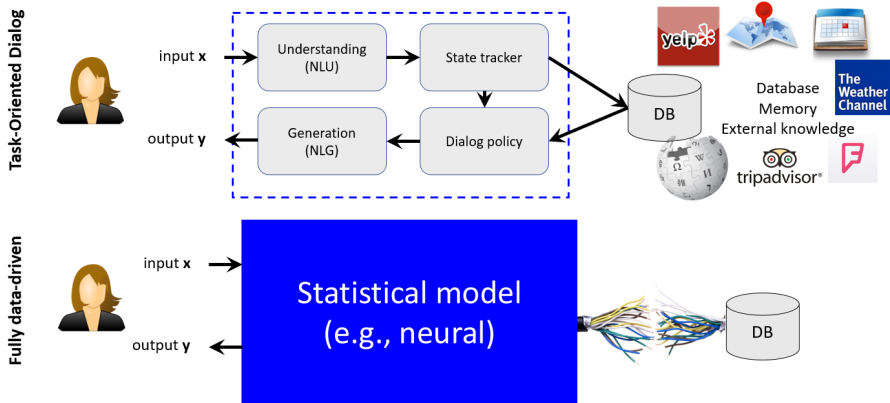


Figure 1.2: Two architectures of dialogue systems for (Top) traditional task-oriented dialogue and (Bottom) fully data-driven dialogue.

One may envision that the above dialogue can be collectively accomplished by a set of agents, also known as *bots*, each of which is designed for solving a particular type of task, e.g., QA bots, task-completion bots, social chatbots. These bots can be grouped into two categories, *task-oriented* and *chitchat*, depending on whether the dialogue is conducted to assist users to achieve specific tasks, e.g., obtain an answer to a query or have a meeting scheduled.

Most of the popular personal assistants in today's market, such as Amazon Alexa, Apple Siri, Google Home, and Microsoft Cortana, are task-oriented bots. These can only handle relatively simple tasks, such as reporting weather and requesting songs. An example of a chitchat dialogue bot is Microsoft XiaoIce. Building a dialogue agent to fulfill complex tasks as in Fig. 1.1 remains one of the most fundamental challenges for the IR and NLP communities, and AI in general.

A typical task-oriented dialogue agent is composed of four modules, as illustrated in Fig. 1.2 (Top): (1) a Natural Language Understanding (NLU) module for identifying user intents and extracting associated information; (2) a state tracker for tracking the dialogue state that captures all essential information in the conversation so far; (3) a

dialogue policy that selects the next action based on the current state; and (4) a Natural Language Generation (NLG) module for converting agent actions to natural language responses. In recent years there has been a trend towards developing fully data-driven systems by unifying these modules using a deep neural network that maps the user input to the agent output directly, as shown in Fig. 1.2 (Bottom).

Most task-oriented bots are implemented using a modular system, where the bot often has access to an external database on which to inquire about information to accomplish the task (Young *et al.*, 2013; Tur and De Mori, 2011). Social chatbots, on the other hand, are often implemented using a unitary (non-modular) system. Since the primary goal of social chatbots is to be AI companions to humans with an emotional connection rather than completing specific tasks, they are often developed to mimic human conversations by training DNN-based response generation models on large amounts of human-human conversational data (Ritter *et al.*, 2011; Sordoni *et al.*, 2015b; Vinyals and Le, 2015; Shang *et al.*, 2015). Only recently have researchers begun to explore how to ground the chitchat in world knowledge (Ghazvininejad *et al.*, 2018) and images (Mostafazadeh *et al.*, 2017) so as to make the conversation more contentful and interesting.

1.3 A Unified View: Dialogue as Optimal Decision Making

The example dialogue in Fig. 1.1 can be formulated as a decision making process. It has a natural hierarchy: a top-level process selects what agent to activate for a particular subtask (e.g., answering a question, scheduling a meeting, providing a recommendation or just having a casual chat), and a low-level process, controlled by the selected agent, chooses primitive actions to complete the subtask.

Such hierarchical decision making processes can be cast in the mathematical framework of *options* over Markov Decision Processes (MDPs) (Sutton *et al.*, 1999b), where options generalize primitive actions to higher-level actions. In a traditional MDP setting, an agent chooses a primitive action at each time step. With options, the agent can choose a “multi-step” action which for example could be a sequence of primitive actions for completing a subtask.

If we view each option as an action, both top- and low-level processes can be naturally captured by the reinforcement learning framework. The dialogue agent navigates in a MDP, interacting with its environment over a sequence of discrete steps. At each step, the agent observes the current state, and chooses an action according to a policy. The agent then receives a reward and observes a new state, continuing the cycle until the episode terminates. The goal of dialogue learning is to find optimal policies to maximize expected rewards. Table 1.1 formulates an sample of dialogue agents using this unified view of RL, where the state-action spaces characterize the complexity of the problems, and the rewards are the objective functions to be optimized.

The unified view of hierarchical MDPs has already been applied to guide the development of some large-scale open-domain dialogue systems. Recent examples include Sounding Board ³, a social chatbot that won the 2017 Amazon Alexa Prize, and Microsoft XiaoIce ⁴, arguably the most popular social chatbot that has attracted more than 660 million users worldwide since its release in 2014. Both systems use a hierarchical dialogue manager: a master (top-level) that manages the overall conversation process, and a collection of skills (low-level) that handle different types of conversation segments (subtasks).

The reward functions in Table 1.1, which seem contradictory in CPS (e.g., we need to minimize CPS for efficient task completion but maximize CPS for improving user engagement), suggest that we have to balance the long-term and short-term gains when developing a dialogue system. For example, XiaoIce is a social chatbot optimized for user engagement, but is also equipped with more than 230 skills, most of which are QA and task-oriented. XiaoIce is optimized for *expected* CPS which corresponds a long-term, rather than a short-term, engagement. Although incorporating many task-oriented and QA skills can reduce CPS in the short term since these skills help users accomplish tasks *more efficiently* by minimizing CPS, these new skills establish XiaoIce as an efficient and trustworthy personal assistant, thus strengthening the emotional bond with human users in the long run.

³<https://sounding-board.github.io/>

⁴<https://www.msxiaobing.com/>

Table 1.1: Reinforcement Learning for Dialogue. CPS stands for Conversation-turns Per Session, and is defined as the average number of conversation-turns between the bot and the user in a conversational session.

dialogue	state	action	reward
QA	understanding of user query intent	clarification questions or answers	relevance of answer, (min) CPS
task-oriented	understanding of user goal	dialogue-act and slot/value	task success rate, (min) CPS
chitchat	conversation history and user intent	responses	user engagement, measured in CPS
top-level bot	understanding of user top-level intent	options	user engagement, measured in CPS

Although RL provides a unified ML framework for building dialogue agents, applying RL requires training the agents by interacting with real users, which can be expensive in many domains. Hence, in practice, we often use a hybrid approach that combines the strengths of different ML methods. For example, we might use imitation and/or supervised learning methods (if there is a large amount of human-human conversational corpus) to obtain a reasonably good agent before applying RL to continue improving it. In the paper, we will survey these ML approaches and their use for training dialogue systems.

1.4 The Transition of NLP to Neural Approaches

Neural approaches are now transforming the field of NLP and IR, where symbolic approaches have been dominating for decades.

NLP applications differ from other data processing systems in their use of language knowledge of various levels, including phonology, morphology, syntax, semantics and discourse (Jurafsky and Martin, 2009). Historically, much of the NLP field has organized itself around the architecture of Fig. 1.3, with researchers aligning their work with one component task, such as morphological analysis or parsing. These tasks can be viewed as resolving (or realizing) natural language ambiguity (or diversity) at different levels by mapping (or generating) a natural language sentence to (or from) a series of human-defined, unambiguous, symbolic representations, such as Part-Of-Speech (POS) tags, context

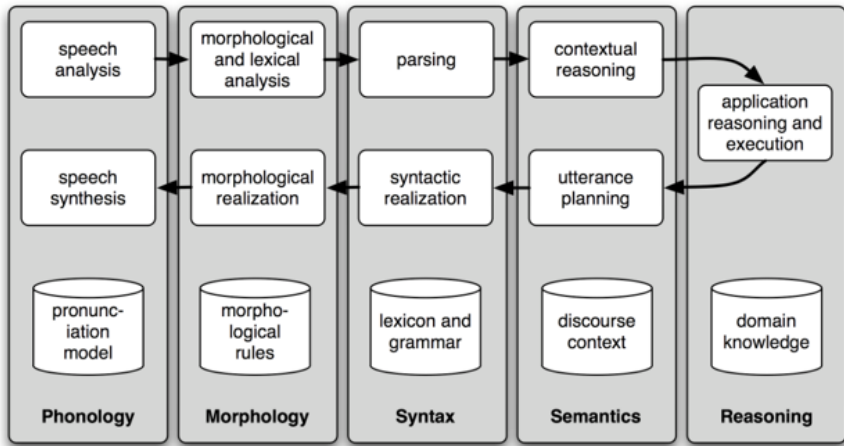


Figure 1.3: Traditional NLP Component Stack. Figure credit: Bird *et al.* (2009).

free grammar, first-order predicate calculus. With the rise of data-driven and statistical approaches, these components have remained and have been adapted as a rich source of engineered features to be fed into a variety of machine learning models (Manning *et al.*, 2014).

Neural approaches do not rely on any human-defined symbolic representations but learn in a *task-specific* neural space where task-specific knowledge is *implicitly* represented as semantic concepts using low-dimensional continuous vectors. As Fig. 1.4 illustrates, neural methods in NLP tasks (e.g., machine reading comprehension and dialogue) often consist of three steps: (1) *encoding* symbolic user input and knowledge into their neural semantic representations, where semantically related or similar concepts are represented as vectors that are close to each other; (2) *reasoning* in the neural space to generate a system response based on input and system state; and (3) *decoding* the system response into a natural language output in a symbolic space. Encoding, reasoning and decoding are implemented using neural networks of different architectures, all of which may be stacked into a deep neural network trained in an end-to-end fashion via back propagation.

End-to-end training results in tighter coupling between the end application and the neural network architecture, lessening the need for traditional NLP component boundaries like morphological analysis and

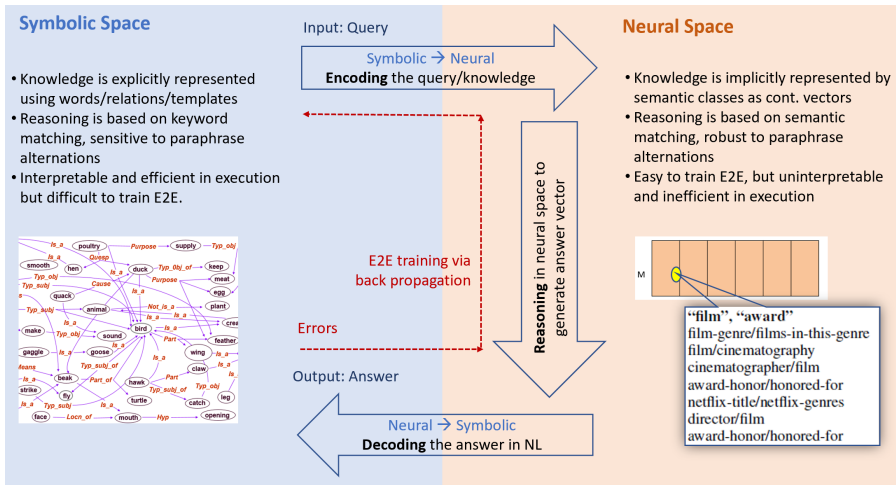


Figure 1.4: Symbolic and Neural Computation.

parsing. This drastically flattens the technology stack of Fig. 1.3, and substantially reduces the need for feature engineering. Instead, the focus has shifted to carefully tailoring the increasingly complex architecture of neural networks to the end application.

Although neural approaches have already been widely adopted in many AI tasks, including image processing, speech recognition and machine translation (e.g., Goodfellow *et al.*, 2016), their impact on conversational AI has come somewhat more slowly. Only recently have we begun to observe neural approaches establish state-of-the-art results on an array of conversation benchmarks for both component tasks and end applications and, in the process, sweep aside the traditional component-based boundaries that have defined research areas for decades. This symbolic-to-neural shift is also reshaping the conversational AI landscape by opening up new tasks and user experiences that were not possible with older techniques. One reason for this is that neural approaches provide a consistent representation for many modalities, capturing linguistic and non-linguistic (e.g., image and video (Mostafazadeh *et al.*, 2017)) features in the same modeling framework.

There are also works on hybrid methods that combine the strengths of both neural and symbolic approaches e.g., (Mou *et al.*, 2016; Liang *et*

al., 2016). As summarized in Fig. 1.4, neural approaches can be trained in an end-to-end fashion and are robust to paraphrase alternations, but are weak in terms of execution efficiency and explicit interpretability. Symbolic approaches, on the other hand, are difficult to train and sensitive to paraphrase alternations, but are more interpretable and efficient in execution.

References

- Agichtein, E., D. Carmel, D. Pelleg, Y. Pinter, and D. Harman. 2015. “Overview of the TREC 2015 LiveQA Track”. In: *TREC*.
- Ai, H. and D. J. Litman. 2008. “Assessing Dialog System User Simulation Evaluation Measures Using Human Judges”. In: *Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics (ACL)*. 622–629.
- Ai, H., A. Raux, D. Bohus, M. Eskenazi, and D. Litman. 2007. “Comparing Spoken Dialog Corpora Collected with Recruited Subjects versus Real Users”. In: *Proceedings of the 8th SIGdial Workshop on Discourse and Dialogue*. 124–131.
- Al-Rfou, R., M. Pickett, J. Snider, Y. Sung, B. Strope, and R. Kurzweil. 2016. “Conversational Contextual Cues: The Case of Personalization and History for Response Ranking”. *CoRR*. abs/1606.00372.
- Albrecht, J. and R. Hwa. 2007. “A Re-examination of Machine Learning Approaches for Sentence-Level MT Evaluation”. In: *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*. Prague, Czech Republic. 880–887.
- Allen, J. F., D. K. Byron, M. O. Dzikovska, G. Ferguson, L. Galescu, and A. Stent. 2001. “Toward Conversational Human-Computer Interaction”. *AI Magazine*. 22(4): 27–38.

- Angeli, G., P. Liang, and D. Klein. 2010. “A Simple Domain-Independent Probabilistic Approach to Generation”. *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing (EMNLP)*: 502–512.
- Auer, S., C. Bizer, G. Kobilarov, J. Lehmann, R. Cyganiak, and Z. Ives. 2007. “DBpedia: A nucleus for a web of open data”. In: *The semantic web*. Springer. 722–735.
- Bahdanau, D., K. Cho, and Y. Bengio. 2015. “Neural machine translation by jointly learning to align and translate”. In: *Proc. of ICLR*.
- Banerjee, S. and A. Lavie. 2005. “METEOR: An Automatic Metric for MT Evaluation with Improved Correlation with Human Judgments”. In: *ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*. 65–72.
- Bao, J., N. Duan, M. Zhou, and T. Zhao. 2014. “Knowledge-based question answering as machine translation”. In: *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Vol. 1. 967–976.
- Bapna, A., G. Tür, D. Hakkani-Tür, and L. P. Heck. 2017. “Towards Zero-Shot Frame Semantic Parsing for Domain Scaling”. In: *Proceedings of the 18th Annual Conference of the International Speech Communication Association (INTERSPEECH)*. 2476–2480.
- Barlier, M., J. Pérolat, R. Laroche, and O. Pietquin. 2015. “Human-Machine Dialogue as a Stochastic Game”. In: *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*. 2–11.
- Baxter, J. and P. Bartlett. 2001. “Infinite-Horizon Policy-Gradient Estimation”. *Journal of Artificial Intelligence Research*. 15: 319–350.
- Baxter, J., P. Bartlett, and L. Weaver. 2001. “Experiments with Infinite-Horizon, Policy-Gradient Estimation”. *Journal of Artificial Intelligence Research*. 15: 351–381.
- Bell, J. 1999. “Pragmatic reasoning: Inferring contexts”. In: *International and Interdisciplinary Conference on Modeling and Using Context*. Springer. 42–53.

- Bellemare, M. G., S. Srinivasan, G. Ostrovski, T. Schaul, D. Saxton, and R. Munos. 2016. “Unifying Count-Based Exploration and Intrinsic Motivation”. In: *Advances in Neural Information Processing Systems (NIPS)*. 1471–1479.
- Berant, J., A. Chou, R. Frostig, and P. Liang. 2013. “Semantic parsing on Freebase from question-answer pairs”. In: *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*. 1533–1544.
- Bertsekas, D. P. and J. N. Tsitsiklis. 1996. *Neuro-Dynamic Programming*. Athena Scientific.
- Bhatia, P., M. Gavalda, and A. Einolghozati. 2017. “soc2seq: Social Embedding meets Conversation Model”. *CoRR*. abs/1702.05512.
- Bird, S., E. Klein, and E. Loper. 2009. *Natural Language Processing with Python*. O’Reilly.
- Black, A. W., S. Burger, A. Conkie, H. W. Hastie, S. Keizer, O. Lemon, N. Merigaud, G. Parent, G. Schubiner, B. Thomson, J. D. Williams, K. Yu, S. J. Young, and M. Eskénazi. 2011. “Spoken Dialog Challenge 2010: Comparison of Live and Control Test Results”. In: *Proceedings of the 12th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*. 2–7.
- Bohus, D. and E. Horvitz. 2009. “Models for Multiparty Engagement in Open-World Dialog”. In: *Proceedings of the 10th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*. 225–234.
- Bohus, D. and E. Horvitz. 2011. “Multiparty Turn Taking in Situated Dialog: Study, Lessons, and Directions”. In: *Proceedings of the 12th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*. 98–109.
- Bohus, D. and A. I. Rudnicky. 2009. “The RavenClaw Dialog Management Framework: Architecture and Systems”. *Computer Speech & Language*. 23(3): 332–361.
- Bohus, D., C. W. Saw, and E. Horvitz. 2014. “Directions Robot: In-the-wild Experiences and Lessons Learned”. In: *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. 637–644.

- Bollacker, K., C. Evans, P. Paritosh, T. Sturge, and J. Taylor. 2008. "Freebase: a collaboratively created graph database for structuring human knowledge". In: *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*. ACM. 1247–1250.
- Bordes, A., Y.-L. Boureau, and J. Weston. 2017. "Learning End-to-End Goal-Oriented Dialog". In: *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Bordes, A., N. Usunier, A. Garcia-Duran, J. Weston, and O. Yakhnenko. 2013. "Translating embeddings for modeling multi-relational data". In: *Advances in neural information processing systems*. 2787–2795.
- Bos, J., E. Klein, O. Lemon, and T. Oka. 2003. "DIPPER: Description and Formalisation of an Information-State Update Dialogue System Architecture". In: *Proceedings of the 4th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*. 115–124.
- Braunschweiler, N. and A. Papangelis. 2018. "Comparison of an End-to-end Trainable Dialogue System with a Modular Statistical Dialogue System". In: *Proceedings of the 19th Annual Conference of the International Speech Communication Association (INTERSPEECH)*. 576–580.
- Budzianowski, P., S. Ultes, P.-H. Su, N. Mrkšić, T.-H. Wen, I. Casanueva, L. M. Rojas-Barahona, and M. Gašić. 2017. "Sub-domain Modelling for Dialogue Management with Hierarchical Reinforcement Learning". In: *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue (SIGDIAL)*. 86–92.
- Callison-Burch, C., P. Koehn, C. Monz, and J. Schroeder. 2009. "Findings of the 2009 Workshop on Statistical Machine Translation". In: *Proceedings of the Fourth Workshop on Statistical Machine Translation*. Athens, Greece. 1–28.
- Caruana, R. 1998. "Multitask learning". In: *Learning to learn*. Springer. 95–133.

- Casanueva, I., P. Budzianowski, P.-H. Su, S. Ultes, L. M. Rojas-Barahona, B.-H. Tseng, and M. Gašić. 2018. “Feudal Reinforcement Learning for Dialogue Management in Large Domains”. In: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*. 714–719.
- Cer, D., C. D. Manning, and D. Jurafsky. 2010. “The Best Lexical Metric for Phrase-based Statistical MT System Optimization”. In: *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics. HLT '10*. Los Angeles, California. 555–563.
- Chandramohan, S., M. Geist, F. Lefèvre, and O. Pietquin. 2011. “User Simulation in Dialogue Systems Using Inverse Reinforcement Learning”. In: *Proceedings of the 12th Annual Conference of the International Speech Communication Association (INTERSPEECH)*. 1025–1028.
- Chapelle, O. and L. Li. 2012. “An Empirical Evaluation of Thompson Sampling”. In: *Advances in Neural Information Processing Systems 24 (NIPS)*. 2249–2257.
- Chen, D., J. Bolton, and C. D. Manning. 2016a. “A thorough examination of the CNN/Daily Mail reading comprehension task”. *arXiv preprint arXiv:1606.02858*.
- Chen, D., A. Fisch, J. Weston, and A. Bordes. 2017a. “Reading Wikipedia to answer open-domain questions”. *arXiv 1704.00051*.
- Chen, H., X. Liu, D. Yin, and J. Tang. 2017b. “A Survey on Dialogue Systems: Recent Advances and New Frontiers”. *arXiv preprint arXiv:1711.01731*.
- Chen, J., C. Wang, L. Xiao, J. He, L. Li, and L. Deng. 2017c. “Q-LDA: Uncovering Latent Patterns in Text-based Sequential Decision Processes”. In: *Advances in Neural Information Processing Systems 30*. 4984–4993.
- Chen, L., B. Tan, S. Long, and K. Yu. 2018. “Structured Dialogue Policy with Graph Neural Networks”. In: *Proceedings of the 27th International Conference on Computational Linguistics (COLING)*. 1257–1268.

- Chen, L., X. Zhou, C. Chang, R. Yang, and K. Yu. 2017d. “Agent-Aware Dropout DQN for Safe and Efficient On-line Dialogue Policy Learning”. In: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2454–2464.
- Chen, Y.-N., A. Celikyilmaz, and D. Hakkani-Tür. 2017e. “Deep Learning for Dialogue Systems”. In: *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Tutorial Abstracts)*. 8–14.
- Chen, Y.-N. and J. Gao. 2017. “Open-Domain Neural Dialogue Systems”. In: *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Tutorial Abstracts)*. 6–10.
- Chen, Y.-N., D. Hakkani-Tür, G. Tur, J. Gao, and L. Deng. 2016b. “End-to-End Memory Networks with Knowledge Carryover for Multi-Turn Spoken Language Understanding”. In: *Proceedings of The 17th Annual Meeting of the International Speech Communication Association*. 3245–3249.
- Cho, K., B. van Merriënboer, D. Bahdanau, and Y. Bengio. 2014a. “On the Properties of Neural Machine Translation: Encoder–Decoder Approaches”. In: *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*. Doha, Qatar. 103–111.
- Cho, K., B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. 2014b. “Learning Phrase Representations using RNN Encoder–Decoder for Statistical Machine Translation”. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Doha, Qatar. 1724–1734.
- Choi, E., H. He, M. Iyyer, M. Yatskar, W.-t. Yih, Y. Choi, P. Liang, and L. Zettlemoyer. 2018. “QuAC: Question Answering in Context”. *arXiv preprint arXiv:1808.07036*.
- Clark, P., I. Cowhey, O. Etzioni, T. Khot, A. Sabharwal, C. Schoenick, and O. Tafjord. 2018. “Think you have Solved Question Answering? Try ARC, the AI2 Reasoning Challenge”. *arXiv preprint arXiv:1803.05457*.
- Colby, K. M. 1975. *Artificial Paranoia: A Computer Simulation of Paranoid Processes*. New York, NY, USA: Elsevier Science Inc.

- Cole, R. A. 1999. "Tools for Research and Education in Speech Science". In: *Proceedings of International Conference of Phonetic Sciences*. 1277–1280.
- Core, M. G. and J. F. Allen. 1997. "Coding Dialogs with the DAMSL Annotation Scheme". In: *Proceedings of AAAI Fall Symposium on Communicative Action in Humans and Machines*. 28–35.
- Corston-Oliver, S., M. Gamon, and C. Brockett. 2001. "A Machine Learning Approach to the Automatic Evaluation of Machine Translation". In: *Proceedings of 39th Annual Meeting of the Association for Computational Linguistics*. Toulouse, France. 148–155.
- Côté, M.-A., Á. Kádár, X. Yuan, B. Kybartas, T. Barnes, E. Fine, J. Moore, M. Hausknecht, L. El Asri, M. Adada, W. Tay, and A. Trischler. 2018. "TextWorld: A Learning Environment for Text-based Games". arXiv:1806.11532.
- Crook, P. A., A. Marin, V. Agarwal, K. Aggarwal, T. Anastasakos, R. Bikkula, D. Boies, A. Celikyilmaz, S. Chandramohan, Z. Feizollahi, R. Holenstein, M. Jeong, O. Z. Khan, Y.-B. Kim, E. Krawczyk, X. Liu, D. Panic, V. Radostev, N. Ramesh, J.-P. Robichaud, A. Rochette, L. Stromberg, and R. Sarikaya. 2016. "Task Completion Platform: A Self-serve Multi-domain Goal Oriented Dialogue Platform". In: *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (HLT-NAACL): Demonstrations Session*. 47–51.
- Cuayáhuitl, H., S. Renals, O. Lemon, and H. Shimodaira. 2010. "Evaluation of a Hierarchical Reinforcement Learning Spoken Dialogue System". *Computer Speech and Language*. 24(2): 395–429.
- Cuayáhuitl, H., S. Yu, A. Williamson, and J. Carse. 2016. "Deep Reinforcement Learning for Multi-Domain Dialogue Systems". *arXiv preprint arXiv:1611.08675*.
- Dai, B., A. Shaw, N. He, L. Li, and L. Song. 2018a. "Boosting the Actor with Dual Critic". In: *Proceedings of the Sixth International Conference on Learning Representations (ICLR)*.

- Dai, B., A. Shaw, L. Li, L. Xiao, N. He, Z. Liu, J. Chen, and L. Song. 2018b. “SBEEED: Convergent Reinforcement Learning with Nonlinear Function Approximation”. In: *Proceedings of the Thirty-Fifth International Conference on Machine Learning (ICML-18)*. 1133–1142.
- Dang, H. T., D. Kelly, and J. J. Lin. 2007. “Overview of the TREC 2007 Question Answering Track”. In: *TREC*. Vol. 7. 63.
- Dann, C., T. Lattimore, and E. Brunskill. 2017. “Unifying PAC and Regret: Uniform PAC Bounds for Episodic Reinforcement Learning”. In: *Advances in Neural Information Processing Systems 30 (NIPS)*. 5717–5727.
- Das, A., S. Kottur, K. Gupta, A. Singh, D. Yadav, J. M. Moura, D. Parikh, and D. Batra. 2017a. “Visual Dialog”. In: *CVPR*.
- Das, R., S. Dhuliawala, M. Zaheer, L. Vilnis, I. Durugkar, A. Krishnamurthy, A. Smola, and A. McCallum. 2017b. “Go for a walk and arrive at the answer: Reasoning over paths in knowledge bases using reinforcement learning”. *arXiv preprint arXiv:1711.05851*.
- Daubigney, L., M. Gašić, S. Chandramohan, M. Geist, O. Pietquin, and S. J. Young. 2011. “Uncertainty Management for On-Line Optimisation of a POMDP-Based Large-Scale Spoken Dialogue System”. In: *Proceedings of the 12th Annual Conference of the International Speech Communication Association (INTERSPEECH)*. 1301–1304.
- Daubigney, L., M. Geist, S. Chandramohan, and O. Pietquin. 2012. “A Comprehensive Reinforcement Learning Framework for Dialogue Management Optimization”. *IEEE Journal of Selected Topics in Signal Processing*. 6(8): 891–902.
- Dayan, P. and G. E. Hinton. 1993. “Feudal Reinforcement Learning”. In: *Advances in Neural Information Processing Systems 5 (NIPS)*. 271–278.

- DeVault, D., R. Artstein, G. Benn, T. Dey, E. Fast, A. Gainer, K. Georgila, J. Gratch, A. Hartholt, M. Lhommet, G. M. Lucas, S. Marsella, F. Morbini, A. Nazarian, S. Scherer, G. Stratou, A. Suri, D. R. Traum, R. Wood, Y. Xu, A. A. Rizzo, and L.-P. Morency. 2014. "SimSensei Kiosk: A Virtual Human Interviewer for Healthcare Decision Support". In: *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. 1061–1068.
- Devlin, J., M.-W. Chang, K. Lee, and K. Toutanova. 2018. "BERT: Pre-training of deep bidirectional transformers for language understanding". *arXiv preprint arXiv:1810.04805*.
- Dhingra, B., L. Li, X. Li, J. Gao, Y.-N. Chen, F. Ahmed, and L. Deng. 2017. "Towards End-to-End Reinforcement Learning of Dialogue Agents for Information Access". In: *ACL (1)*. 484–495.
- Dhingra, B., H. Liu, Z. Yang, W. W. Cohen, and R. Salakhutdinov. 2016. "Gated-attention readers for text comprehension". *arXiv preprint arXiv:1606.01549*.
- Dietterich, T. G. 2000. "Hierarchical Reinforcement Learning with the MAXQ Value Function Decomposition". *Journal of Artificial Intelligence Research*. 13: 227–303.
- Ding, Y., Y. Liu, H. Luan, and M. Sun. 2017. "Visualizing and Understanding Neural Machine Translation". In: *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Vancouver, Canada. 1150–1159.
- Dodge, J., A. Gane, X. Zhang, A. Bordes, S. Chopra, A. Miller, A. Szlam, and J. Weston. 2016. "Evaluating prerequisite qualities for learning end-to-end dialog systems". In: *ICLR*.
- Dunn, M., L. Sagun, M. Higgins, U. Guney, V. Cirik, and K. Cho. 2017. "SearchQA: A new Q&A dataset augmented with context from a search engine". *arXiv preprint arXiv:1704.05179*.
- Eckert, W., E. Levin, and R. Pieraccini. 1997. "User Modeling for Spoken Dialogue System Evaluation". In: *Proceedings of the 1997 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*. 80–87.

- El Asri, L., J. He, and K. Suleman. 2016. “A Sequence-to-Sequence Model for User Simulation in Spoken Dialogue Systems”. In: *Proceedings of the 17th Annual Conference of the International Speech Communication Association (INTERSPEECH)*. 1151–1155.
- El Asri, L., R. Laroche, and O. Pietquin. 2012. “Reward Function Learning for Dialogue Management”. In: *Proceedings of the Sixth Starting AI Researchers’ Symposium (STAIRS)*. 95–106.
- Engel, Y., S. Mannor, and R. Meir. 2005. “Reinforcement Learning with Gaussian Processes”. In: *Proceedings of the 22nd International Conference on Machine Learning (ICML)*. 201–208.
- Eric, M., L. Krishnan, F. Charette, and C. D. Manning. 2017. “Key-Value Retrieval Networks for Task-Oriented Dialogue”. In: *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue (SIGDIAL)*. 37–49.
- Ernst, D., P. Geurts, and L. Wehenkel. 2005. “Tree-Based Batch Mode Reinforcement Learning”. *Journal of Machine Learning Research*. 6: 503–556.
- Fang, H., S. Gupta, F. Iandola, R. K. Srivastava, L. Deng, P. Dollár, J. Gao, X. He, M. Mitchell, J. C. Platt, L. Zitnick, and G. Zweig. 2015. “From captions to visual concepts and back”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1473–1482.
- Fatemi, M., L. El Asri, H. Schulz, J. He, and K. Suleman. 2016. “Policy Networks with Two-stage Training for Dialogue Systems”. In: *Proceedings of the 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*. 101–110.
- Fedorenko, D. G., N. Smetanin, and A. Rodichev. 2017. “Avoiding Echo-Responses in a Retrieval-Based Conversation System”. *CoRR*. abs/1712.05626.
- Fung, P., D. Bertero, Y. Wan, A. Dey, R. H. Y. Chan, F. B. Siddique, Y. Yang, C. Wu, and R. Lin. 2016. “Towards Empathetic Human-Robot Interactions”. *CoRR*. abs/1605.04072.
- Galley, M., C. Brockett, A. Sordoni, Y. Ji, M. Auli, C. Quirk, M. Mitchell, J. Gao, and B. Dolan. 2015. “deltaBLEU: A Discriminative Metric for Generation Tasks with Intrinsically Diverse Targets”. In: *ACL-IJCNLP*. 445–450.

- Gao, J. 2017. “An Introduction to Deep Learning for Natural Language Processing”. In: *International Summer School on Deep Learning, Bilbao*.
- Gao, J., M. Galley, and L. Li. 2018a. “Neural Approaches to Conversational AI”. In: *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. ACM. 1371–1374.
- Gao, J., M. Galley, and L. Li. 2018b. “Neural Approaches to Conversational AI”. *Proc. of ACL 2018, Tutorial Abstracts: 2–7*.
- Gao, J., X. He, W.-t. Yih, and L. Deng. 2014a. “Learning continuous phrase representations for translation modeling”. In: *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Vol. 1. 699–709.
- Gao, J., P. Pantel, M. Gamon, X. He, and L. Deng. 2014b. “Modeling interestingness with deep neural networks”. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2–13.
- Gardner, M., P. Talukdar, J. Krishnamurthy, and T. Mitchell. 2014. “Incorporating vector space similarity in random walk inference over knowledge bases”. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 397–406.
- Gašić, M., C. Breslin, M. Henderson, D. Kim, M. Szummer, B. Thomson, P. Tsiakoulis, and S. J. Young. 2013. “On-line Policy Optimisation of Bayesian Spoken Dialogue Systems via Human Interaction”. In: *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 8367–8371.
- Gašić, M., D. Kim, P. Tsiakoulis, C. Breslin, M. Henderson, M. Szummer, B. Thomson, and S. Young. 2014. “Incremental On-line Adaptation of POMDP-based Dialogue Managers to Extended Domains”. In: *Proceedings of the 15th Annual Conference of the International Speech Communication Association (INTERSPEECH)*. 140–144.
- Gašić, M., N. Mrkšić, P.-h. Su, D. Vandyke, T.-H. Wen, and S. J. Young. 2015. “Policy Committee for Adaptation in Multi-domain Spoken Dialogue Systems”. In: *Proceedings of the 2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*. 806–812.

- Gašić, M. and S. J. Young. 2014. “Gaussian Processes for POMDP-based Dialogue Manager Optimization”. *IEEE Trans. Audio, Speech & Language Processing*. 22(1): 28–40.
- Georgila, K., J. Henderson, and O. Lemon. 2006. “User Simulation for Spoken Dialogue Systems: Learning and Evaluation”. In: *Proceedings of the 9th International Conference on Spoken Language Processing (INTERSPEECH)*. 1065–1068.
- Ghazvininejad, M., C. Brockett, M.-W. Chang, B. Dolan, J. Gao, W.-t. Yih, and M. Galley. 2018. “A Knowledge-Grounded Neural Conversation Model”. In: *5110–5117*.
- Giménez, J. and L. Màrquez. 2008. “A Smorgasbord of Features for Automatic MT Evaluation”. In: *Proceedings of the Third Workshop on Statistical Machine Translation*. 195–198.
- Goodfellow, I., Y. Bengio, and A. Courville. 2016. *Deep Learning*. MIT Press.
- Goodfellow, I., J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. 2014. “Generative adversarial nets”. In: *NIPS*. 2672–2680.
- Graham, Y. and T. Baldwin. 2014. “Testing for Significance of Increased Correlation with Human Judgment”. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Doha, Qatar. 172–176.
- Graham, Y., T. Baldwin, and N. Mathur. 2015. “Accurate Evaluation of Segment-level Machine Translation Metrics”. In: *NAACL-HLT*.
- Graves, A. and J. Schmidhuber. 2005. “Framewise Phoneme Classification with Bidirectional LSTM and Other Neural Network Architectures”. *Neural Networks*. 18: 602–610.
- Gu, J., Z. Lu, H. Li, and V. O. Li. 2016. “Incorporating Copying Mechanism in Sequence-to-Sequence Learning”. In: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 1631–1640.
- Gu, S., T. Lillicrap, Z. Ghahramani, R. E. Turner, and S. Levine. 2017. “Q-Prop: Sample-Efficient Policy Gradient with An Off-Policy Critic”. In: *Proceedings of the 5th International Conference on Learning Representations (ICLR)*.

- Guu, K., J. Miller, and P. Liang. 2015. "Traversing knowledge graphs in vector space". *arXiv preprint arXiv:1506.01094*.
- Hakkani-Tür, D., G. Tür, A. Celikyilmaz, Y.-N. Chen, J. Gao, L. Deng, and Y.-Y. Wang. 2016. "Multi-Domain Joint Semantic Frame Parsing using Bi-directional RNN-LSTM". In: *Proceedings of the 17th Annual Conference of the International Speech Communication Association (INTERSPEECH)*. 715–719.
- Hakkani-Tür, D., G. Tur, L. Heck, A. Fidler, and A. Celikyilmaz. 2012. "A Discriminative Classification-Based Approach to Information State Updates for a Multi-Domain Dialog System". In: *Proceedings of the 13th Annual Conference of the International Speech Communication Association (INTERSPEECH)*. 330–333.
- Hartikainen, M., E.-P. Salonen, and M. Turunen. 2004. "Subjective Evaluation of Spoken Dialogue Systems using SERVQUAL Method". In: *Proceedings of the 8th International Conference on Spoken Language Processing (INTERSPEECH)*. 2273–2276.
- Hausknecht, M. and P. Stone. 2015. "Deep Recurrent Q-Learning for Partially Observable MDPs". In: *Proceedings of the AAAI Fall Symposium on Sequential Decision Making for Intelligent Agents*. 29–37.
- He, J., J. Chen, X. He, J. Gao, L. Li, L. Deng, and M. Ostendorf. 2016. "Deep Reinforcement Learning with a Natural Language Action Space". In: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL)*. 1621–1630.
- He, S., C. Liu, K. Liu, and J. Zhao. 2017a. "Generating Natural Answers by Incorporating Copying and Retrieving Mechanisms in Sequence-to-Sequence Learning". In: *ACL*. Vol. 1. 199–208.
- He, W., K. Liu, Y. Lyu, S. Zhao, X. Xiao, Y. Liu, Y. Wang, H. Wu, Q. She, X. Liu, T. Wu, and H. Wang. 2017b. "DuReader: a Chinese Machine Reading Comprehension Dataset from Real-world Applications". *arXiv preprint arXiv:1711.05073*.
- Henderson, J., O. Lemon, and K. Georgila. 2008. "Hybrid Reinforcement/Supervised Learning of Dialogue Policies from Fixed Data Sets". *Computational Linguistics*. 34(4): 487–511.

- Henderson, M. 2015. "Machine Learning for Dialog State Tracking: A Review". In: *Proceedings of The First International Workshop on Machine Learning in Spoken Language Processing*.
- Henderson, M., B. Thomson, and J. D. Williams. 2014a. "The 3rd Dialog State Tracking Challenge". In: *Proceedings of the 2014 IEEE Spoken Language Technology Workshop (SLT)*. 324–329.
- Henderson, M., B. Thomson, and J. D. Williams. 2014b. "The Second Dialog State Tracking Challenge". In: *Proceedings of the 15th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*. 263–272.
- Henderson, M., B. Thomson, and S. J. Young. 2013. "Deep Neural Network Approach for the Dialog State Tracking Challenge". In: *Proceedings of the 14th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*. 467–471.
- Hermann, K. M., T. Kocisky, E. Grefenstette, L. Espeholt, W. Kay, M. Suleyman, and P. Blunsom. 2015. "Teaching machines to read and comprehend". In: *Advances in Neural Information Processing Systems*. 1693–1701.
- Hewlett, D., A. Lacoste, L. Jones, I. Polosukhin, A. Fandrianto, J. Han, M. Kelcey, and D. Berthelot. 2016. "WikiReading: A novel large-scale language understanding task over wikipedia". *arXiv preprint arXiv:1608.03542*.
- Hill, F., A. Bordes, S. Chopra, and J. Weston. 2015. "The Goldilocks principle: Reading children's books with explicit memory representations". *arXiv preprint arXiv:1511.02301*.
- Hinton, G. E. and D. Van Camp. 1993. "Keeping the neural networks simple by minimizing the description length of the weights". In: *Proceedings of the sixth annual conference on Computational learning theory*. ACM. 5–13.
- Hinton, G., L. Deng, D. Yu, G. E. Dahl, A. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, and B. Kingsbury. 2012. "Deep Neural Networks for Acoustic Modeling in Speech Recognition: The Shared Views of Four Research Groups". *IEEE Signal Processing Magazine*. 29(6): 82–97.

- Hochreiter, S. 1991. "Untersuchungen zu dynamischen neuronalen Netzen. Diploma thesis, Institut für Informatik, Lehrstuhl Prof. Brauer, Technische Universität München".
- Hochreiter, S. and J. Schmidhuber. 1997. "Long short-term memory". *Neural computation*. 9(8): 1735–1780.
- Hofmann, K., L. Li, and F. Radlinski. 2016. "Online Evaluation for Information Retrieval". *Foundations and Trends in Information Retrieval*. 10(1): 1–117.
- Holtzman, A., J. Buys, M. Forbes, A. Bosselut, D. Golub, and Y. Choi. 2018. "Learning to Write with Cooperative Discriminators". In: *ACL*. Melbourne, Australia. 1638–1649.
- Hori, C. and T. Hori. 2017. "End-to-end Conversation Modeling Track in DSTC6". *CoRR*. abs/1706.07440.
- Hori, C., T. Hori, S. Watanabe, and J. R. Hershey. 2015. "Context Sensitive Spoken Language Understanding using Role Dependent LSTM Layers". *Tech. rep.* No. TR2015-134. Mitsubishi Electric Research Laboratories.
- Hori, C., J. Perez, K. Yoshino, and S. Kim. 2017. "The Sixth Dialog State Tracking Challenge". <http://workshop.colips.org/dstc6>.
- Horvitz, E. 1999. "Principles of Mixed-Initiative User Interfaces". In: *Proceeding of the Conference on Human Factors in Computing Systems (CHI)*. 159–166.
- Houthooft, R., X. Chen, Y. Duan, J. Schulman, F. D. Turck, and P. Abbeel. 2016. "VIME: Variational Information Maximizing Exploration". In: *Advances in Neural Information Processing Systems 29 (NIPS)*. 1109–1117.
- Hu, M., Y. Peng, and X. Qiu. 2017. "Mnemonic reader for machine comprehension". *arXiv preprint arXiv:1705.02798*.
- Huang, H.-Y., C. Zhu, Y. Shen, and W. Chen. 2017. "FusionNet: Fusing via Fully-Aware Attention with Application to Machine Comprehension". *arXiv preprint arXiv:1711.07341*.
- Huang, Q., Z. Gan, A. Celikyilmaz, D. O. Wu, J. Wang, and X. He. 2018. "Hierarchically Structured Reinforcement Learning for Topically Coherent Visual Story Generation". *CoRR*. abs/1805.08191.

- Huang, P.-S., X. He, J. Gao, L. Deng, A. Acero, and L. Heck. 2013. “Learning deep structured semantic models for web search using clickthrough data”. In: *Proceedings of the 22nd ACM international conference on Conference on information & knowledge management*. 2333–2338.
- Huang, X., A. Acero, and H.-W. Hon. 2001. *Spoken language processing: A guide to theory, algorithm, and system development*. Prentice Hall.
- Huber, B., D. McDuff, C. Brockett, M. Galley, and B. Dolan. 2018. “Emotional Dialogue Generation using Image-Grounded Language Models”. In: *CHI*. 277:1–277:12.
- Inaba, M. and K. Takahashi. 2016. “Neural Utterance Ranking Model for Conversational Dialogue Systems”. In: *SIGDIAL*. 393–403.
- Iyyer, M., W.-t. Yih, and M.-W. Chang. 2017. “Search-based neural structured learning for sequential question answering”. In: *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Vol. 1. 1821–1831.
- Jafarpour, S., C. J. Burges, and A. Ritter. 2010. “Filter, rank, and transfer the knowledge: Learning to chat”. *Advances in Ranking*. 10: 2329–9290.
- Jaksch, T., R. Ortner, and P. Auer. 2010. “Near-optimal Regret Bounds for Reinforcement Learning”. *Journal of Machine Learning Research*. 11: 1563–1600.
- Jia, R. and P. Liang. 2017. “Adversarial examples for evaluating reading comprehension systems”. *arXiv preprint arXiv:1707.07328*.
- Jiang, N., A. Krishnamurthy, A. Agarwal, J. Langford, and R. E. Schapire. 2017. “Contextual Decision Processes with Low Bellman Rank are PAC-Learnable”. In: *Proceedings of the 34th International Conference on Machine Learning (ICML)*. 1704–1713.
- Jiang, N. and L. Li. 2016. “Doubly Robust Off-policy Evaluation for Reinforcement Learning”. In: *Proceedings of the 33rd International Conference on Machine Learning (ICML)*. 652–661.
- Joshi, M., E. Choi, D. S. Weld, and L. Zettlemoyer. 2017. “TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension”. *arXiv preprint arXiv:1705.03551*.

- Jung, S., C. Lee, K. Kim, M. Jeong, and G. G. Lee. 2009. “Data-driven User Simulation for Automated Evaluation of Spoken Dialog Systems”. *Computer Speech and Language*. 23: 479–509.
- Jurafsky, D. and J. H. Martin. 2009. *Speech & language processing*. Prentice Hall.
- Jurafsky, D. and J. H. Martin. 2018. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. Draft of August 12th, 2018. Website: <https://web.stanford.edu/~jurafsky/slp3>. Prentice-Hall.
- Kaelbling, L. P., M. L. Littman, and A. P. Moore. 1996. “Reinforcement Learning: A Survey”. *Journal of Artificial Intelligence Research*. 4: 237–285.
- Kakade, S. 2001. “A Natural Policy Gradient”. In: *Advances in Neural Information Processing Systems 13 (NIPS)*. 1531–1538.
- Kalchbrenner, N. and P. Blunsom. 2013. “Recurrent Continuous Translation Models”. In: *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*. Seattle, Washington, USA. 1700–1709.
- Kannan, A. and O. Vinyals. 2016. “Adversarial Evaluation of Dialogue Models”. In: *NIPS Workshop on Adversarial Training*.
- Khandelwal, U., H. He, P. Qi, and D. Jurafsky. 2018. “Sharp Nearby, Fuzzy Far Away: How Neural Language Models Use Context”. In: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Melbourne, Australia. 284–294.
- Kim, S., L. F. D’Haro, R. E. Banchs, J. D. Williams, and M. Henderson. 2016a. “The Fourth Dialog State Tracking Challenge”. In: *Proceedings of the 7th International Workshop on Spoken Dialogue Systems (IWSDS)*. 435–449.
- Kim, S., L. F. D’Haro, R. E. Banchs, J. D. Williams, M. Henderson, and K. Yoshino. 2016b. “The Fifth Dialog State Tracking Challenge”. In: *Proceedings of the 2016 IEEE Spoken Language Technology Workshop (SLT-16)*. 511–517.
- Kočisky, T., J. Schwarz, P. Blunsom, C. Dyer, K. M. Hermann, G. Melis, and E. Grefenstette. 2017. “The NarrativeQA Reading Comprehension Challenge”. *arXiv preprint arXiv:1712.07040*.

- Koehn, P. 2004. "Statistical Significance Tests for Machine Translation Evaluation". In: *Proceedings of EMNLP 2004*. Barcelona, Spain. 388–395.
- Koehn, P., F. J. Och, and D. Marcu. 2003. "Statistical Phrase-based Translation". In: *Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology - Volume 1. NAACL '03*. Edmonton, Canada. 48–54.
- Koller, A. and M. Stone. 2007. "Sentence Generation as a Planning Problem". In: *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics (ACL-07)*. 336–343.
- Komatani, K., N. Kanda, M. Nakano, K. Nakadai, H. Tsujino, T. Ogata, and H. G. Okuno. 2006. "Multi-Domain Spoken Dialogue System with Extensibility and Robustness against Speech Recognition Errors". In: *Proceedings of the SIGDIAL 2006 Workshop*. 9–17.
- Konda, V. R. and J. N. Tsitsiklis. 1999. "Actor-Critic Algorithms". In: *Advances in Neural Information Processing Systems 12 (NIPS)*. 1008–1014.
- Kondadadi, R., B. Howald, and F. Schilder. 2013. "A Statistical NLG Framework for Aggregated Planning and Realization". In: *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (ACL)*. 1406–1415.
- Kotti, M., V. Diakouloukas, A. Papangelis, M. Lagoudakis, and Y. Stylianou. 2018. "A Case Study on the Importance of Belief State Representation for Dialogue Policy Management". In: *Proceedings of the 19th Annual Conference of the International Speech Communication Association (INTERSPEECH)*. 986–990.
- Kulesza, A. and S. M. Shieber. 2004. "A Learning Approach to Improving Sentence-Level MT Evaluation". In: *Proceedings of the 10th International Conference on Theoretical and Methodological Issues in Machine Translation*. Baltimore, MD.
- Kumar, A., O. Irsoy, P. Ondruska, M. Iyyer, J. Bradbury, I. Gulrajani, V. Zhong, R. Paulus, and R. Socher. 2016. "Ask me anything: Dynamic memory networks for natural language processing". In: *International Conference on Machine Learning*. 1378–1387.

- Lai, G., Q. Xie, H. Liu, Y. Yang, and E. Hovy. 2017. “RACE: Large-scale reading comprehension dataset from examinations”. *arXiv preprint arXiv:1704.04683*.
- Langkilde, I. and K. Knight. 1998. “Generation that Exploits Corpus-Based Statistical Knowledge”. In: *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics (COLING-ACL)*. 704–710.
- Lao, N. and W. W. Cohen. 2010. “Relational retrieval using a combination of path-constrained random walks”. *Machine learning*. 81(1): 53–67.
- Lao, N., T. Mitchell, and W. W. Cohen. 2011. “Random walk inference and learning in a large scale knowledge base”. In: *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. 529–539.
- Larsson, S. and D. R. Traum. 2000. “Information State and Dialogue Management in the TRINDI Dialogue Move Engine Toolkit”. *Natural Language Engineering*. 6(3–4): 323–340.
- Lee, J. Y. and F. Deroncourt. 2016. “Sequential Short-Text Classification with Recurrent and Convolutional Neural Networks”. In: *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*. 515–520.
- Lee, S. and R. Jha. 2019. “Zero-Shot Adaptive Transfer for Conversational Language Understanding”. In: *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI-19)*, to appear.
- Lei, W., X. Jin, Z. Ren, X. He, M.-Y. Kan, and D. Yin. 2018. “Sequicity: Simplifying Task-oriented Dialogue Systems with Single Sequence-to-Sequence Architectures”. In: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*. 1437–1447.
- Lemon, O. 2011. “Learning What to Say and How to Say It: Joint Optimisation of Spoken Dialogue Management and Natural Language Generation”. *Computer Speech & Language*. 25(2): 210–221.
- Levin, E., R. Pieraccini, and W. Eckert. 2000. “A Stochastic Model of Human-Machine Interaction for Learning Dialog Strategies”. *IEEE Transactions on Speech and Audio Processing*. 8(1): 11–23.

- Lewis, M., D. Yarats, Y. Dauphin, D. Parikh, and D. Batra. 2017. “Deal or No Deal? End-to-End Learning of Negotiation Dialogues”. In: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP-17)*. 2443–2453.
- Li, J., M. Galley, C. Brockett, J. Gao, and B. Dolan. 2016a. “A diversity-promoting objective function for neural conversation models”. In: *NAACL-HLT*. 110–119.
- Li, J., M. Galley, C. Brockett, J. Gao, and B. Dolan. 2016b. “A Persona-Based Neural Conversation Model”. In: *ACL*. 994–1003.
- Li, J., A. H. Miller, S. Chopra, M. Ranzato, and J. Weston. 2017a. “Dialogue learning with human-in-the-loop”. In: *ICLR*.
- Li, J., A. H. Miller, S. Chopra, M. Ranzato, and J. Weston. 2017b. “Learning Through Dialogue Interactions by Asking Questions”. In: *ICLR*.
- Li, J., W. Monroe, A. Ritter, D. Jurafsky, M. Galley, and J. Gao. 2016c. “Deep Reinforcement Learning for Dialogue Generation”. In: *EMNLP*. 1192–1202.
- Li, J., W. Monroe, T. Shi, S. Jean, A. Ritter, and D. Jurafsky. 2017c. “Adversarial Learning for Neural Dialogue Generation”. In: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Copenhagen, Denmark. 2157–2169.
- Li, L., H. He, and J. D. Williams. 2014. “Temporal Supervised Learning for Inferring a Dialog Policy from Example Conversations”. In: *Proceedings of the 2014 IEEE Spoken Language Technology Workshop (SLT)*. 312–317.
- Li, L., J. D. Williams, and S. Balakrishnan. 2009. “Reinforcement Learning for Spoken Dialog Management using Least-Squares Policy Iteration and Fast Feature Selection”. In: *Proceedings of the 10th Annual Conference of the International Speech Communication Association (INTERSPEECH)*. 2475–2478.
- Li, X., Y.-N. Chen, L. Li, J. Gao, and A. Celikyilmaz. 2017d. “End-to-End Task-Completion Neural Dialogue Systems”. In: *Proceedings of the 8th International Joint Conference on Natural Language Processing (IJCNLP)*. 733–743.

- Li, X., Y.-N. Chen, L. Li, J. Gao, and A. Celikyilmaz. 2017e. “Investigation of Language Understanding Impact for Reinforcement Learning Based Dialogue Systems”. CoRR abs/1703.07055.
- Li, X., L. Li, J. Gao, X. He, J. Chen, L. Deng, and J. He. 2015. “Recurrent Reinforcement Learning: A Hybrid Approach”. arXiv:1509.03044.
- Li, X., Z. C. Lipton, B. Dhingra, L. Li, J. Gao, and Y.-N. Chen. 2016d. “A User Simulator for Task-Completion Dialogues”. CoRR abs/1612.05688.
- Li, X., S. Panda, J. Liu, and J. Gao. 2018. “Microsoft Dialogue Challenge: Building End-to-End Task-Completion Dialogue Systems”. *arXiv preprint arXiv:1807.11125*.
- Li, Y. 2018. “Deep Reinforcement Learning”. arXiv 1810.06339.
- Liang, C., J. Berant, Q. Le, K. D. Forbus, and N. Lao. 2016. “Neural symbolic machines: Learning semantic parsers on freebase with weak supervision”. *arXiv preprint arXiv:1611.00020*.
- Lin, C.-Y. 2004. “ROUGE: A package for automatic evaluation of summaries”. In: *ACL workshop*. 74–81.
- Lin, L.-J. 1992. “Self-Improving Reactive Agents Based on Reinforcement Learning, Planning and Teaching”. *Machine Learning*. 8(3–4): 293–321.
- Lipton, Z. C., J. Gao, L. Li, X. Li, F. Ahmed, and L. Deng. 2018. “BBQ-Networks: Efficient Exploration in Deep Reinforcement Learning for Task-Oriented Dialogue Systems”. In: *AAAI*. 5237–5244.
- Lita, L. V., M. Rogati, and A. Lavie. 2005. “BLANC: Learning Evaluation Metrics for MT”. In: *Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing. HLT '05*. Vancouver, British Columbia, Canada. 740–747.
- Litman, D. J. and J. F. Allen. 1987. “A Plan Recognition Model for Subdialogues in Conversations”. *Cognitive Science*. 11(163–200).
- Liu, B. and I. Lane. 2016. “Attention-based Recurrent Neural Network Models for Joint Intent Detection and Slot Filling”. In: *Proceedings of the 17th Annual Conference of the International Speech Communication Association (INTERSPEECH)*. 685–689.

- Liu, B. and I. Lane. 2017. “Iterative Policy Learning in End-to-end Trainable Task-oriented Neural Dialog Models”. In: *Proceedings of the 2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. 482–489.
- Liu, B. and I. Lane. 2018. “Adversarial Learning of Task-Oriented Neural Dialog Models”. In: *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue (SIGDIAL)*. 350–359.
- Liu, B., G. Tür, D. Hakkani-Tür, P. Shah, and L. P. Heck. 2018a. “Dialogue Learning with Human Teaching and Feedback in End-to-End Trainable Task-Oriented Dialogue Systems”. In: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*. 2060–2069.
- Liu, C.-W., R. Lowe, I. Serban, M. Noseworthy, L. Charlin, and J. Pineau. 2016. “How NOT To Evaluate Your Dialogue System: An Empirical Study of Unsupervised Evaluation Metrics for Dialogue Response Generation”. In: *EMNLP*. 2122–2132.
- Liu, H., Y. Feng, Y. Mao, D. Zhou, J. Peng, and Q. Liu. 2018b. “Action-dependent Control Variates for Policy Optimization via Stein’s Identity”. In: *Proceedings of the 6th International Conference on Learning Representations (ICLR)*.
- Liu, Q., L. Li, Z. Tang, and D. Zhou. 2018c. “Breaking the Curse of Horizon: Infinite-horizon Off-policy Estimation”. In: *Advances in Neural Information Processing Systems 31 (NIPS-18)*. 5361–5371.
- Liu, R., W. Wei, W. Mao, and M. Chikina. 2017. “Phase Conductor on Multi-layered Attentions for Machine Comprehension”. *arXiv preprint arXiv:1710.10504*.
- Liu, X., Y. Shen, K. Duh, and J. Gao. 2018d. “Stochastic answer networks for machine reading comprehension”. In: *ACL*. 1694–1704.
- Lowe, R., M. Noseworthy, I. V. Serban, N. Angelard-Gontier, Y. Bengio, and J. Pineau. 2017. “Towards an Automatic Turing Test: Learning to Evaluate Dialogue Responses”. In: *ACL*. 1116–1126.
- Lowe, R., N. Pow, I. Serban, and J. Pineau. 2015. “The Ubuntu Dialogue Corpus: A Large Dataset for Research in Unstructured Multi-Turn Dialogue Systems”. In: *SIGDIAL*. 285–294.

- Lu, Z. and H. Li. 2014. “A Deep Architecture for Matching Short Texts”. In: *Advances in Neural Information Processing Systems 27*. Curran Associates, Inc. 1368–1375.
- Luan, Y., C. Brockett, B. Dolan, J. Gao, and M. Galley. 2017. “Multi-Task Learning for Speaker-Role Adaptation in Neural Conversation Models”. In: *IJCNLP*. 605–614.
- Luong, T., H. Pham, and C. D. Manning. 2015. “Effective Approaches to Attention-based Neural Machine Translation”. In: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Lisbon, Portugal. 1412–1421.
- Madotto, A., C.-S. Wu, and P. Fung. 2018. “Mem2Seq: Effectively Incorporating Knowledge Bases into End-to-End Task-Oriented Dialog Systems”. In: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*. 1468–1478.
- Mairesse, F. and S. Young. 2014. “Stochastic Language Generation in Dialogue using Factored Language Models”. *Computational Linguistics*. 40(4): 763–799.
- Manning, C., M. Surdeanu, J. Bauer, J. Finkel, S. Bethard, and D. McClosky. 2014. “The Stanford CoreNLP natural language processing toolkit”. In: *Proceedings of 52nd annual meeting of the association for computational linguistics: system demonstrations*. 55–60.
- Mao, J., W. Xu, Y. Yang, J. Wang, Z. Huang, and A. Yuille. 2015. “Deep Captioning with Multimodal Recurrent Neural Networks (m-RNN)”. In: *ICLR*.
- McCann, B., J. Bradbury, C. Xiong, and R. Socher. 2017. “Learned in translation: Contextualized word vectors”. In: *Advances in Neural Information Processing Systems*. 6297–6308.
- McCann, B., N. S. Keskar, C. Xiong, and R. Socher. 2018. “The Natural Language Decathlon: Multitask Learning as Question Answering”. *arXiv preprint arXiv:1806.08730*.
- McTear, M. F. 2002. “Spoken Dialogue Technology: Enabling the Conversational User Interface”. *ACM Computing Surveys*. 34(1): 90–169.
- Mei, H., M. Bansal, and M. R. Walter. 2017. “Coherent Dialogue with Attention-based Language Models”. In: *AAAI*. 3252–3258.

- Melamud, O., J. Goldberger, and I. Dagan. 2016. “context2vec: Learning generic context embedding with bidirectional lstm”. In: *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*. 51–61.
- Mesnil, G., Y. Dauphin, K. Yao, Y. Bengio, L. Deng, D. Z. Hakkani-Tür, X. He, L. P. Heck, G. Tür, D. Yu, and G. Zweig. 2015. “Using Recurrent Neural Networks for Slot Filling in Spoken Language Understanding”. *IEEE/ACM Transactions on Audio, Speech & Language Processing*. 23(3): 530–539.
- Mesnil, G., X. He, L. Deng, and Y. Bengio. 2013. “Investigation of Recurrent-Neural-network Architectures and Learning Methods for Spoken Language Understanding”. In: *Proceedings of the 14th Annual Conference of the International Speech Communication Association (INTERSPEECH)*. 3771–3775.
- Mikolov, T., I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. 2013. “Distributed representations of words and phrases and their compositionality”. In: *Advances in neural information processing systems*. 3111–3119.
- Minsky, M. and S. Papert. 1969. *Perceptrons: An Introduction to Computational Geometry*. MIT Press.
- Mitchell, T. 1997. *Machine Learning*. New York: McGraw-Hill.
- Mnih, V., Adrià, P. Badia, M. Mirza, A. Graves, T. P. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu. 2016. “Asynchronous Methods for Deep Reinforcement Learning”. In: *Proceedings of the 33rd International Conference on Machine Learning (ICML)*. 1928–1937.
- Mnih, V., K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Belle-mare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis. 2015. “Human-level Control through Deep Reinforcement Learning”. *Nature*. 518: 529–533.
- Morgenstern, L. and C. L. Ortiz. 2015. “The Winograd Schema Challenge: Evaluating Progress in Commonsense Reasoning”. In: *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence. AAAI’15*. Austin, Texas. 4024–4025.

- Mostafazadeh, N., C. Brockett, B. Dolan, M. Galley, J. Gao, G. Spithourakis, and L. Vanderwende. 2017. "Image-Grounded Conversations: Multimodal Context for Natural Question and Response Generation". In: *IJCNLP*. 462–472.
- Mou, L., Z. Lu, H. Li, and Z. Jin. 2016. "Coupling distributed and symbolic execution for natural language queries". *arXiv preprint arXiv:1612.02741*.
- Mrkšić, N., D. Ó. Séaghdha, B. Thomson, M. Gašić, P.-H. Su, D. Vandyke, T.-H. Wen, and S. J. Young. 2015. "Multi-domain Dialog State Tracking using Recurrent Neural Networks". In: *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing (ACL)*. 794–799.
- Mrkšić, N., D. Ó. Séaghdha, T.-H. Wen, B. Thomson, and S. J. Young. 2017. "Neural Belief Tracker: Data-Driven Dialogue State Tracking". In: *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL)*. 1777–1788.
- Munos, R. and C. Szepesvári. 2008. "Finite-Time Bounds for Sampling-Based Fitted Value Iteration". *Journal of Machine Learning Research*. 9: 815–857.
- Narasimhan, K., T. D. Kulkarni, and R. Barzilay. 2015. "Language Understanding for Text-based Games using Deep Reinforcement Learning". In: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 1–11.
- Neelakantan, A., B. Roth, and A. McCallum. 2015. "Compositional vector space models for knowledge base completion". *arXiv preprint arXiv:1504.06662*.
- Nguyen, D. Q. 2017. "An overview of embedding models of entities and relationships for knowledge base completion". *arXiv preprint arXiv:1703.08098*.
- Nguyen, T., M. Rosenberg, X. Song, J. Gao, S. Tiwary, R. Majumder, and L. Deng. 2016. "MS MARCO: A human generated machine reading comprehension dataset". *arXiv preprint arXiv:1611.09268*.

- Och, F. J. 2003. "Minimum Error Rate Training in Statistical Machine Translation". In: *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics*. Sapporo, Japan. 160–167.
- Och, F. J. and H. Ney. 2003. "A Systematic Comparison of Various Statistical Alignment Models". *Computational Linguistics*. 29(1): 19–51.
- Och, F. J. and H. Ney. 2004. "The Alignment Template Approach to Statistical Machine Translation". *Comput. Linguist.* 30(4): 417–449.
- Oh, A. H. and A. I. Rudnicky. 2002. "Stochastic Natural Language Generation for Spoken Dialog Systems". *Computer Speech & Language*. 16(3–4): 387–407.
- Osband, I., C. Blundell, A. Pritzel, and B. V. Roy. 2016. "Deep Exploration via Bootstrapped DQN". In: *Advances in Neural Information Processing Systems 29 (NIPS-16)*. 4026–4034.
- Osband, I. and B. V. Roy. 2017. "Why is Posterior Sampling Better than Optimism for Reinforcement Learning?" In: *Proceedings of the 34th International Conference on Machine Learning (ICML)*. 2701–2710.
- Pado, S., D. Cer, M. Galley, D. Jurafsky, and C. D. Manning. 2009. "Measuring Machine Translation Quality as Semantic Equivalence: A Metric Based on Entailment Features". *Machine Translation*: 181–193.
- Paek, T. 2001. "Empirical Methods for Evaluating Dialog Systems". In: *Proceedings of the 2nd Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*. 1–9.
- Paek, T. and R. Pieraccini. 2008. "Automating Spoken Dialogue Management Design using Machine Learning: An Industry Perspective". *Speech Communication*. 50(8–9): 716–729.
- Papangelis, A., M. Kotti, and Y. Stylianou. 2018a. "Towards Scalable Information-Seeking Multi-Domain Dialogue". In: *Proceedings of the 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 6064–6068.

- Papangelis, A., P. Papadakos, Y. Stylianou, and Y. Tzitzikas. 2018b. “Spoken Dialogue for Information Navigation”. In: *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue (SIGDIAL)*. 229–234.
- Papineni, K., S. Roukos, T. Ward, and W.-J. Zhu. 2002. “BLEU: a method for automatic evaluation of machine translation”. In: *ACL*. 311–318.
- Parr, R. and S. J. Russell. 1998. “Reinforcement Learning with Hierarchies of Machines”. In: *Advances of Neural Information Processing Systems 10 (NIPS)*. 1043–1049.
- Pasupat, P. and P. Liang. 2015. “Compositional semantic parsing on semi-structured tables”. *arXiv preprint arXiv:1508.00305*.
- Peng, B., X. Li, J. Gao, J. Liu, and K.-F. Wong. 2018. “Integrating planning for task-completion dialogue policy learning”. CoRR abs/1801.06176.
- Peng, B., X. Li, L. Li, J. Gao, A. Celikyilmaz, S. Lee, and K.-F. Wong. 2017. “Composite Task-Completion Dialogue Policy Learning via Hierarchical Deep Reinforcement Learning”. In: *EMNLP*. 2231–2240.
- Pennington, J., R. Socher, and C. Manning. 2014. “GloVe: Global vectors for word representation”. In: *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*. 1532–1543.
- Peters, J., S. Vijayakumar, and S. Schaal. 2005. “Natural Actor-Critic”. In: *Proceedings of the 16th European Conference on Machine Learning (ECML)*. 280–291.
- Peters, M. E., M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer. 2018. “Deep contextualized word representations”. *arXiv preprint arXiv:1802.05365*.
- Pietquin, O. and T. Dutoit. 2006. “A Probabilistic Framework for Dialog Simulation and Optimal Strategy Learning”. *IEEE Transactions on Audio, Speech & Language Processing*. 14(2): 589–599.
- Pietquin, O., M. Geist, S. Chandramohan, and H. Frezza-Buet. 2011. “Sample-efficient Batch Reinforcement Learning for Dialogue Management Optimization”. *ACM Transactions on Speech and Language Processing*. 7(3): 7:1–7:21.

- Pietquin, O. and H. Hastie. 2013. “A Survey on Metrics for the Evaluation of User Simulations”. *The Knowledge Engineering Review*. 28(1): 59–73.
- Precup, D., R. S. Sutton, and S. P. Singh. 2000. “Eligibility Traces for Off-Policy Policy Evaluation”. In: *Proceedings of the 17th International Conference on Machine Learning (ICML)*. 759–766.
- Przybocki, M., K. Peterson, S. Bronsart, and G. Sanders. 2009. “The NIST 2008 Metrics for machine translation challenge—overview, methodology, metrics, and results”. *Machine Translation*. 23(2): 71–103.
- Puterman, M. L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. New York: Wiley-Interscience.
- Rajpurkar, P., R. Jia, and P. Liang. 2018. “Know What You Don’t Know: Unanswerable Questions for SQuAD”. *arXiv 1806.03822*.
- Rajpurkar, P., J. Zhang, K. Lopyrev, and P. Liang. 2016. “SQuAD: 100,000+ questions for machine comprehension of text”. *arXiv preprint arXiv:1606.05250*.
- Ram, A., R. Prasad, C. Khatri, A. Venkatesh, R. Gabriel, Q. Liu, J. Nunn, B. Hedayatnia, M. Cheng, A. Nagar, E. King, K. Bland, A. Wartick, Y. Pan, H. Song, S. Jayadevan, G. Hwang, and A. Pettigru. 2018. “Conversational AI: The Science Behind the Alexa Prize”. *CoRR*. abs/1801.03604.
- Ramshaw, L. A. and M. Marcus. 1995. “Text Chunking using Transformation based Learning”. In: *Third Workshop on Very Large Corpora (VLC at ACL)*. 82–94.
- Ranzato, M., S. Chopra, M. Auli, and W. Zaremba. 2015. “Sequence Level Training with Recurrent Neural Networks”. *arXiv 1511.06732*.
- Ravuri, S. V. and A. Stolcke. 2015. “Recurrent Neural Network and LSTM Models for Lexical Utterance Classification”. In: *Proceedings of the 16th Annual Conference of the International Speech Communication Association (INTERSPEECH)*. 135–139.
- Ravuri, S. V. and A. Stolcke. 2016. “A Comparative Study of Recurrent Neural Network Models for Lexical Domain Classification”. In: *Proceedings of the 2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 6075–6079.

- Reddy, S., D. Chen, and C. D. Manning. 2018. "CoQA: A Conversational Question Answering Challenge". *arXiv preprint arXiv:1808.07042*.
- Ren, G., X. Ni, M. Malik, and Q. Ke. 2018a. "Conversational Query Understanding Using Sequence to Sequence Modeling". In: *Proceedings of the 2018 World Wide Web Conference on World Wide Web*. International World Wide Web Conferences Steering Committee. 1715–1724.
- Ren, L., K. Xie, L. Chen, and K. Yu. 2018b. "Towards Universal Dialogue State Tracking". In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2780–2786.
- Rich, C., C. L. Sidner, and N. Lesh. 2001. "COLLAGEN: Applying Collaborative Discourse Theory to Human-Computer Interaction". *AI Magazine*. 22(4): 15–25.
- Richardson, M., C. J. Burges, and E. Renshaw. 2013. "MCTest: A challenge dataset for the open-domain machine comprehension of text". In: *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*. 193–203.
- Richardson, S. D., W. B. Dolan, and L. Vanderwende. 1998. "Mind-Net: acquiring and structuring semantic information from text". In: *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics-Volume 2*. 1098–1102.
- Rieser, V. and O. Lemon. 2008. "Learning Effective Multimodal Dialogue Strategies from Wizard-of-Oz Data: Bootstrapping and Evaluation". In: *Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics (ACL)*. 638–646.
- Rieser, V. and O. Lemon. 2010. "Natural Language Generation as Planning under Uncertainty for Spoken Dialogue Systems". In: *Empirical Methods in Natural Language Generation: Data-oriented Methods and Empirical Evaluation*. Vol. 5790. *Lecture Notes in Computer Science*. Springer. 105–120.
- Rieser, V. and O. 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.

- Rieser, V., O. Lemon, and X. Liu. 2010. "Optimising Information Presentation for Spoken Dialogue Systems". In: *Proceedings of the Forty-Eighth Annual Meeting of the Association for Computational Linguistics (ACL-10)*. 1009–1018.
- Ritter, A., C. Cherry, and W. Dolan. 2011. "Data-driven response generation in social media". In: *EMNLP*. 583–593.
- Roemmele, M., C. Bejan, and A. S. Gordon. 2011. "Choice of Plausible Alternatives: An Evaluation of Commonsense Causal Reasoning". In: *AAAI Spring Symposium - Technical Report*.
- Rosenblatt, F. 1957. "The perceptron: A perceiving and recognizing automaton". *Report No. 85-460-1*. Ithaca, New York: Project PARA, Cornell Aeronautical Laboratory.
- Rosenblatt, F. 1962. *Principles of Neurodynamics*. New York: Spartan Books.
- Ross, S., G. J. Gordon, and D. Bagnell. 2011. "A Reduction of Imitation Learning and Structured Prediction to No-Regret Online Learning". In: *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics (AISTATS)*. 627–635.
- Roy, N., J. Pineau, and S. Thrun. 2000. "Spoken Dialogue Management using Probabilistic Reasoning". In: *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics (ACL)*. 93–100.
- Russo, D. J., B. Van Roy, A. Kazerouni, I. Osband, and Z. Wen. 2018. "A Tutorial on Thompson Sampling". *Foundations and Trends in Machine Learning*. 11(1): 1–96.
- Saha, A., V. Pahuja, M. M. Khapra, K. Sankaranarayanan, and S. Chandar. 2018. "Complex Sequential Question Answering: Towards Learning to Converse Over Linked Question Answer Pairs with a Knowledge Graph". *arXiv preprint arXiv:1801.10314*.
- Salimans, T., I. J. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. 2016. "Improved Techniques for Training GANs". *CoRR*. abs/1606.03498.
- Sarikaya, R., G. E. Hinton, and A. Deoras. 2014. "Application of Deep Belief Networks for Natural Language Understanding". *IEEE/ACM Transactions on Audio, Speech & Language Processing*. 22(4): 778–784.

- Schapire, R. E. and Y. Singer. 2000. “BoosTexter: A Boosting-based System for Text Categorization”. *Machine Learning*. 39(2/3): 135–168.
- Schatzmann, J., K. Georgila, and S. Young. 2005a. “Quantitative Evaluation of User Simulation Techniques for Spoken Dialogue Systems”. In: *Proceedings of the 6th SIGdial Workshop on Discourse and Dialogue*. 45–54.
- Schatzmann, J., M. N. Stuttle, K. Weilhammer, and S. Young. 2005b. “Effects of the User Model on Simulation-based Learning of Dialogue Strategies”. In: *Proceedings of the IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*. 220–225.
- Schatzmann, J., K. Weilhammer, M. Stuttle, and S. Young. 2006. “A Survey of Statistical User Simulation Techniques for Reinforcement-Learning of Dialogue Management Strategies”. *The Knowledge Engineering Review*. 21(2): 97–126.
- Schatzmann, J. and S. Young. 2009. “The Hidden Agenda User Simulation Model”. *IEEE Transactions on Audio, Speech, and Language Processing*. 17(4): 733–747.
- Schulman, J., S. Levine, P. Abbeel, M. I. Jordan, and P. Moritz. 2015a. “Trust Region Policy Optimization”. In: *Proceedings of the Thirty-Second International Conference on Machine Learning (ICML)*. 1889–1897.
- Schulman, J., P. Moritz, S. Levine, M. Jordan, and P. Abbeel. 2015b. “High-Dimensional Continuous Control Using Generalized Advantage Estimation”. arXiv:1506.02438.
- See, A., P. J. Liu, and C. D. Manning. 2017. “Get to the point: Summarization with pointer-generator networks”. *arXiv preprint arXiv:1704.04368*.
- Seo, M., A. Kembhavi, A. Farhadi, and H. Hajishirzi. 2016. “Bidirectional attention flow for machine comprehension”. *arXiv preprint arXiv:1611.01603*.
- Serban, I. V., R. Lowe, L. Charlin, and J. Pineau. 2015. “A Survey of Available Corpora for Building Data-Driven Dialogue Systems”. *arXiv preprint arXiv:1512.05742*.

- Serban, I. V., R. Lowe, P. Henderson, L. Charlin, and J. Pineau. 2018. "A Survey of Available Corpora for Building Data-Driven Dialogue Systems: The Journal Version". *Dialogue & Discourse*. 9(1): 1–49.
- Serban, I. V., A. Sordoni, Y. Bengio, A. C. Courville, and J. Pineau. 2016. "Building End-To-End Dialogue Systems Using Generative Hierarchical Neural Network Models". In: *AAAI*. 3776–3783.
- Serban, I. V., A. Sordoni, R. Lowe, L. Charlin, J. Pineau, A. Courville, and Y. Bengio. 2017. "A hierarchical latent variable encoder-decoder model for generating dialogues". In: *AAAI*. 3295–3301.
- Shah, P., D. Z. Hakkani-Tür, B. Liu, and G. Tür. 2018. "Bootstrapping a Neural Conversational Agent with Dialogue Self-Play, Crowdsourcing and On-Line Reinforcement Learning". In: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*. 41–51.
- Shang, L., Z. Lu, and H. Li. 2015. "Neural Responding Machine for Short-Text Conversation". In: *ACL-IJCNLP*. 1577–1586.
- Shao, C. C., T. Liu, Y. Lai, Y. Tseng, and S. Tsai. 2018. "DRCD: a Chinese Machine Reading Comprehension Dataset". *arXiv preprint arXiv:1806.00920*.
- Shao, Y., S. Gouws, D. Britz, A. Goldie, B. Strope, and R. Kurzweil. 2017. "Generating High-Quality and Informative Conversation Responses with Sequence-to-Sequence Models". In: *EMNLP*. 2210–2219.
- Shen, Y., J. Chen, P. Huang, Y. Guo, and J. Gao. 2018. "M-Walk: Learning to Walk in Graph with Monte Carlo Tree Search". *CoRR*. abs/1802.04394.
- Shen, Y., X. He, J. Gao, L. Deng, and G. Mesnil. 2014. "A latent semantic model with convolutional-pooling structure for information retrieval". In: *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*. ACM. 101–110.
- Shen, Y., P. Huang, M. Chang, and J. Gao. 2016. "Implicit ReasonNet: Modeling Large-Scale Structured Relationships with Shared Memory". *CoRR*. abs/1611.04642.

- Shen, Y., P.-S. Huang, M.-W. Chang, and J. Gao. 2017a. “Traversing Knowledge Graph in Vector Space without Symbolic Space Guidance”. *arXiv preprint arXiv:1611.04642*.
- Shen, Y., P.-S. Huang, J. Gao, and W. Chen. 2017b. “ReasonNet: Learning to stop reading in machine comprehension”. In: *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM. 1047–1055.
- Shen, Y., X. Liu, K. Duh, and J. Gao. 2017c. “An Empirical Analysis of Multiple-Turn Reasoning Strategies in Reading Comprehension Tasks”. *arXiv preprint arXiv:1711.03230*.
- Shum, H., X. He, and D. Li. 2018. “From Eliza to XiaoIce: Challenges and Opportunities with Social Chatbots”. *CoRR*. abs/1801.01957.
- Singh, S. P., D. Litman, M. J. Kearns, and M. Walker. 2002. “Optimizing Dialogue Management with Reinforcement Learning: Experiments with the NJFun System”. *Journal of Artificial Intelligence Research*. 16: 105–133.
- Socher, R., D. Chen, C. D. Manning, and A. Ng. 2013. “Reasoning with neural tensor networks for knowledge base completion”. In: *Advances in neural information processing systems*. 926–934.
- Sordoni, A., P. Bachman, A. Trischler, and Y. Bengio. 2016. “Iterative alternating neural attention for machine reading”. *arXiv preprint arXiv:1606.02245*.
- Sordoni, A., Y. Bengio, H. Vahabi, C. Lioma, J. Grue Simonsen, and J.-Y. Nie. 2015a. “A Hierarchical Recurrent Encoder-Decoder for Generative Context-Aware Query Suggestion”. In: *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management. CIKM '15*. Melbourne, Australia: ACM. 553–562.
- Sordoni, A., M. Galley, M. Auli, C. Brockett, Y. Ji, M. Mitchell, J.-Y. Nie, J. Gao, and B. Dolan. 2015b. “A neural network approach to context-sensitive generation of conversational responses”. In: *NAACL-HLT*. 196–205.
- Stanojević, M. and K. Sima'an. 2014. “Fitting Sentence Level Translation Evaluation with Many Dense Features”. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Doha, Qatar. 202–206.

- Stent, A., R. Prasad, and M. A. Walker. 2004. “Trainable Sentence Planning for Complex Information Presentations in Spoken Dialog Systems”. In: *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL)*. 79–86.
- Stone, M., C. Doran, B. L. Webber, T. Bleam, and M. Palmer. 2003. “Microplanning with Communicative Intentions: The SPUD System”. *Computational Intelligence*. 19(4): 311–381.
- Strehl, A. L., L. Li, and M. L. Littman. 2009. “Reinforcement Learning in Finite MDPs: PAC Analysis”. *Journal of Machine Learning Research*. 10: 2413–2444.
- Strub, F., H. de Vries, J. Mary, B. Piot, A. C. Courville, and O. Pietquin. 2017. “End-to-end Optimization of Goal-driven and Visually Grounded Dialogue Systems”. In: *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*. 2765–2771.
- Su, P.-H., M. Gasic, N. Mrksic, L. M. Rojas-Barahona, S. Ultes, D. Vandyke, T.-H. Wen, and S. J. Young. 2016a. “On-line Active Reward Learning for Policy Optimisation in Spoken Dialogue Systems”. In: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL)*. Vol. 1. 2431–2441.
- Su, P.-H., M. Gašić, N. Mrkšić, L. Rojas-Barahona, S. Ultes, D. Vandyke, T.-H. Wen, and S. Young. 2016b. “Continuously Learning Neural Dialogue Management”. arXiv preprint: 1606.02689.
- Su, P.-H., M. Gašić, and S. Young. 2018a. “Reward Estimation for Dialogue Policy Optimisation”. *Computer Speech & Language*. 51: 24–43.
- Su, P.-H., D. Vandyke, M. Gašić, D. Kim, N. Mrkšić, T.-H. Wen, and S. Young. 2015. “Learning from Real Users: Rating Dialogue Success with Neural Networks for Reinforcement Learning in Spoken Dialogue Systems”. In: *Proceedings of the 16th Annual Conference of the International Speech Communication Association (INTERSPEECH)*. 2007–2011.
- Su, S.-Y., X. Li, J. Gao, J. Liu, and Y.-N. Chen. 2018b. “Discriminative Deep Dyna-Q: Robust Planning for Dialogue Policy Learning”. In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 3813–3823.

- Su, S.-Y., K.-L. Lo, Y. T. Yeh, and Y.-N. Chen. 2018c. “Natural Language Generation by Hierarchical Decoding with Linguistic Patterns”. In: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*. 61–66.
- Suchanek, F. M., G. Kasneci, and G. Weikum. 2007. “YAGO: a core of semantic knowledge”. In: *Proceedings of the 16th international conference on World Wide Web*. ACM. 697–706.
- Suhr, A., S. Iyer, and Y. Artzi. 2018. “Learning to Map Context-Dependent Sentences to Executable Formal Queries”. *arXiv preprint arXiv:1804.06868*.
- Sutskever, I., O. Vinyals, and Q. Le. 2014. “Sequence to sequence learning with neural networks”. In: *NIPS*. 3104–3112.
- Sutton, R. S. 1990. “Integrated architectures for learning, planning, and reacting based on approximating dynamic programming”. In: *Proceedings of the seventh international conference on machine learning*. 216–224.
- Sutton, R. S. 1988. “Learning to Predict by the Methods of Temporal Differences”. *Machine Learning*. 3(1): 9–44.
- Sutton, R. S. and A. G. Barto. 2018. *Reinforcement Learning: An Introduction*. 2nd edition. MIT Press.
- Sutton, R. S., D. McAllester, S. P. Singh, and Y. Mansour. 1999a. “Policy Gradient Methods for Reinforcement Learning with Function Approximation”. In: *Advances in Neural Information Processing Systems 12 (NIPS)*. 1057–1063.
- Sutton, R. S., D. Precup, and S. P. Singh. 1999b. “Between MDPs and Semi-MDPs: A Framework for Temporal Abstraction in Reinforcement Learning”. *Artificial Intelligence*. 112(1–2): 181–211. An earlier version appeared as Technical Report 98-74, Department of Computer Science, University of Massachusetts, Amherst, MA 01003. April, 1998.
- Szegedy, C., W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. J. Goodfellow, and R. Fergus. 2013. “Intriguing properties of neural networks”. *CoRR*. abs/1312.6199.
- Szepesvári, C. 2010. *Algorithms for Reinforcement Learning*. Morgan & Claypool.

- Talmor, A. and J. Berant. 2018. “The Web as a Knowledge-base for Answering Complex Questions”. *arXiv preprint arXiv:1803.06643*.
- Tang, D., X. Li, J. Gao, C. Wang, L. Li, and T. Jebara. 2018. “Subgoal Discovery for Hierarchical Dialogue Policy Learning”. In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2298–2309.
- Tesauro, G. 1995. “Temporal Difference Learning and TD-Gammon”. *Communications of the ACM*. 38(3): 58–68.
- Thomas, P. S. and E. Brunskill. 2016. “Data-Efficient Off-Policy Policy Evaluation for Reinforcement Learning”. In: *Proceedings of the 33rd International Conference on Machine Learning (ICML)*. 2139–2148.
- Thompson, W. R. 1933. “On the Likelihood that One Unknown Probability Exceeds Another in View of the Evidence of Two Samples”. *Biometrika*. 25(3–4): 285–294.
- Tiedemann, J. 2012. “Parallel Data, Tools and Interfaces in OPUS”. In: *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC’12)*. Istanbul, Turkey: European Language Resources Association (ELRA). 2214–2218.
- Toutanova, K., V. Lin, W.-t. Yih, H. Poon, and C. Quirk. 2016. “Compositional learning of embeddings for relation paths in knowledge base and text”. In: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Vol. 1. 1434–1444.
- Traum, D. R. 1999. “Speech Acts for Dialogue Agents”. In: *Foundations of Rational Agency*. Springer. 169–201.
- Trischler, A., T. Wang, X. Yuan, J. Harris, A. Sordoni, P. Bachman, and K. Suleman. 2016. “NewsQA: A machine comprehension dataset”. *arXiv preprint arXiv:1611.09830*.
- Tromp, J. and G. Farneböck. 2006. “Combinatorics of Go”. In: *Proceedings of the Fifth International Conference on Computers and Games. Lecture Notes in Computer Science*. No. 4630. 84–99.
- Tur, G. and R. De Mori. 2011. *Spoken language understanding: Systems for extracting semantic information from speech*. John Wiley & Sons.

- Ultes, S., P. Budzianowski, I. Casanueva, N. Mrkšić, L. M. Rojas-Barahona, P.-H. Su, T.-H. Wen, M. Gašić, and S. J. Young. 2017a. “Domain-Independent User Satisfaction Reward Estimation for Dialogue Policy Learning”. In: *Proceedings of the 18th Annual Conference of the International Speech Communication Association (INTERSPEECH)*. 1721–1725.
- Ultes, S., P. Budzianowski, I. Casanueva, N. Mrkšić, L. M. Rojas-Barahona, P.-H. Su, T.-H. Wen, M. Gašić, and S. J. Young. 2017b. “Reward-Balancing for Statistical Spoken Dialogue Systems using Multi-objective Reinforcement Learning”. In: *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue (SIGDIAL)*. 65–70.
- Ultes, S., L. M. Rojas-Barahona, P.-H. Su, D. Vandyke, D. Kim, I. Casanueva, P. Budzianowski, N. Mrkšić, T.-H. Wen, M. Gašić, and S. J. Young. 2017c. “PyDial: A Multi-Domain Statistical Dialogue System Toolkit”. In: *Proceedings of the Fifty-fifth Annual Meeting of the Association for Computational Linguistics (ACL), System Demonstrations*. 73–78.
- van den Oord, A., S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu. 2016. “WaveNet: A Generative Model for Raw Audio”. arXiv:1609.03499.
- van Hasselt, H., A. Guez, and D. Silver. 2016. “Deep Reinforcement Learning with Double Q-learning”. In: *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI-16)*. 2094–2100.
- Vaswani, A., N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. 2017. “Attention is All you Need”. In: *Advances in Neural Information Processing Systems 30*. Curran Associates, Inc. 5998–6008.
- Vinyals, O., M. Fortunato, and N. Jaitly. 2015a. “Pointer Networks”. In: *Advances in Neural Information Processing Systems 28*. Curran Associates, Inc. 2692–2700.
- Vinyals, O. and Q. Le. 2015. “A Neural Conversational Model”. In: *ICML Deep Learning Workshop*.
- Vinyals, O., A. Toshev, S. Bengio, and D. Erhan. 2015b. “Show and tell: A neural image caption generator”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3156–3164.

- Vries, H. de, F. Strub, S. Chandar, O. Pietquin, H. Larochelle, and A. C. Courville. 2017. "GuessWhat?! Visual Object Discovery through Multi-modal Dialogue". In: *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 4466–4475.
- Walker, M. A. 2000. "An Application of Reinforcement Learning to Dialogue Strategy Selection in a Spoken Dialogue System for Email". *Journal of Artificial Intelligence Research*. 12: 387–416.
- Walker, M. A., D. J. Litman, C. A. Kamm, and A. Abella. 1997. "PARADISE: A Framework for Evaluating Spoken Dialogue Agents". In: *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics (ACL)*. 271–280.
- Walker, M. A., D. J. Litman, C. A. Kamm, and A. Abella. 1998. "Evaluating Spoken Dialogue Agents with PARADISE: Two Case Studies". *Computer Speech & Language*. 12(4): 317–347.
- Walker, M. A., A. Stent, F. Mairesse, and R. Prasad. 2007. "Individual and Domain Adaptation in Sentence Planning for Dialogue". *Journal of Artificial Intelligence Research*. 30: 413–456.
- Walker, M., C. Kamm, and D. Litman. 2000. "Towards Developing General Models of Usability with PARADISE". *Natural Language Engineering*. 6(3–4): 363–377.
- Wang, C., Y. Wang, P.-S. Huang, A. Mohamed, D. Zhou, and L. Deng. 2017a. "Sequence Modeling via Segmentations". In: *Proceedings of the 34th International Conference on Machine Learning*. 3674–3683.
- Wang, W., N. Yang, F. Wei, B. Chang, and M. Zhou. 2017b. "Gated self-matching networks for reading comprehension and question answering". In: *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Vol. 1. 189–198.
- Wang, Y.-Y., L. Deng, and A. Acero. 2005. "Spoken Language Understanding: An Introduction to the Statistical Framework". *IEEE Signal Processing Magazine*. 22(5): 16–31.
- Wang, Z., H. Chen, G. Wang, H. Tian, H. Wu, and H. Wang. 2014. "Policy Learning for Domain Selection in an Extensible Multi-domain Spoken Dialogue System". In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 57–67.

- Wang, Z., T. Schaul, M. Hessel, H. van Hasselt, M. Lanctot, and N. de Freitas. 2016. “Dueling Network Architectures for Deep Reinforcement Learning”. In: *Proceedings of the Third International Conference on Machine Learning (ICML-16)*. 1995–2003.
- Watkins, C. J. 1989. “Learning from Delayed Rewards”. *PhD thesis*. UK: King’s College, University of Cambridge.
- Wei, W., Q. V. Le, A. M. Dai, and L.-J. Li. 2018. “AirDialogue: An Environment for Goal-oriented Dialogue Research”. In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 3844–3854.
- Weissenborn, D., G. Wiese, and L. Seiffe. 2017. “FastQA: A simple and efficient neural architecture for question answering”. *arXiv preprint arXiv:1703.04816*.
- Weizenbaum, J. 1966. “ELIZA: a Computer Program for the Study of Natural Language Communication Between Man and Machine”. *Commun. ACM*. 9(1): 36–45.
- Welbl, J., P. Stenetorp, and S. Riedel. 2017. “Constructing Datasets for Multi-hop Reading Comprehension Across Documents”. *arXiv preprint arXiv:1710.06481*.
- Wen, T.-H., M. Gašić, N. Mrkšić, L. M. Rojas-Barahona, P.-H. Su, D. Vandyke, and S. J. Young. 2016. “Multi-domain Neural Network Language Generation for Spoken Dialogue Systems”. In: *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (HLT-NAACL)*. 120–129.
- Wen, T.-H., M. Gašić, N. Mrkšić, P.-H. Su, D. Vandyke, and S. J. Young. 2015. “Semantically Conditioned LSTM-based Natural Language Generation for Spoken Dialogue Systems”. In: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 1711–1721.
- Wen, T.-H., D. Vandyke, N. Mrkšić, M. Gašić, L. M. Rojas-Barahona, P.-H. Su, S. Ultes, and S. J. Young. 2017. “A Network-based End-to-End Trainable Task-oriented Dialogue System”. In: *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*. arXiv reprint arXiv:1604.04562. 438–449.

- Wiering, M. and M. van Otterlo. 2012. *Reinforcement Learning: State of the Art*. Springer.
- Williams, J. D. 2006. “Partially Observable Markov Decision Processes for Spoken Dialogue Management”. *PhD thesis*. Cambridge, UK: Cambridge University.
- Williams, J. D. 2008. “Evaluating User Simulations with the Cramér-von Mises Divergence”. *Speech Communication*. 50(10): 829–846.
- Williams, J. D., K. Asadi, and G. Zweig. 2017. “Hybrid Code Networks: Practical and Efficient End-to-end Dialog Control with Supervised and Reinforcement Learning”. In: *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL)*. Vol. 1. 665–677.
- Williams, J. D., M. Henderson, A. Raux, B. Thomson, A. W. Black, and D. Ramachandran. 2014. “The Dialog State Tracking Challenge Series”. *AI Magazine*. 35(4): 121–124.
- Williams, J. D., A. Raux, D. Ramachandran, and A. W. Black. 2013. “The Dialog State Tracking Challenge”. In: *Proceedings of the 14th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*. 404–413.
- Williams, J. D. and S. J. Young. 2007. “Partially Observable Markov Decision Processes for Spoken Dialog Systems”. *Computer Speech and Language*. 21(2): 393–422.
- Williams, J. D. and G. Zweig. 2016. “End-to-end LSTM-based Dialog Control Optimized with Supervised and Reinforcement Learning”.
- Williams, R. J. 1992. “Simple Statistical Gradient-Following Algorithms for Connectionist Reinforcement Learning”. *Machine Learning*. 8: 229–256.
- Winata, G. I., O. Kampman, Y. Yang, A. Dey, and P. Fung. 2017. “Nora the empathetic psychologist”. In: *Proc. Interspeech*. 3437–3438.
- Wu, C.-S., A. Madotto, G. I. Winata, and P. Fung. 2018. “End-to-End Dynamic Query Memory Network for Entity-Value Independent Task-Oriented Dialog”. In: *Proceedings of the 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 6154–6158.

- Wu, J., M. Li, and C.-H. Lee. 2015. “A probabilistic framework for representing dialog systems and entropy-based dialog management through dynamic stochastic state evolution”. *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)*. 23(11): 2026–2035.
- Wu, Q., C. J. Burges, K. M. Svore, and J. Gao. 2010. “Adapting boosting for information retrieval measures”. *Information Retrieval*. 13(3): 254–270.
- Wu, Y., M. Schuster, Z. Chen, Q. V. Le, M. Norouzi, W. Macherey, M. Krikun, Y. Cao, Q. Gao, K. Macherey, J. Klingner, A. Shah, M. Johnson, X. Liu, L. Kaiser, S. Gouws, Y. Kato, T. Kudo, H. Kazawa, K. Stevens, G. Kurian, N. Patil, W. Wang, C. Young, J. Smith, J. Riesa, A. Rudnick, O. Vinyals, G. Corrado, M. Hughes, and J. Dean. 2016. “Google’s Neural Machine Translation System: Bridging the Gap between Human and Machine Translation”. *CoRR*. abs/1609.08144.
- Xing, C., W. Wu, Y. Wu, M. Zhou, Y. Huang, and W. Ma. 2018. “Hierarchical Recurrent Attention Network for Response Generation”. In: *AAAI*. 5610–5617.
- Xiong, C., V. Zhong, and R. Socher. 2016. “Dynamic coattention networks for question answering”. *arXiv preprint arXiv:1611.01604*.
- Xiong, W., T. Hoang, and W. Y. Wang. 2017. “DeepPath: A reinforcement learning method for knowledge graph reasoning”. *arXiv preprint arXiv:1707.06690*.
- Xu, P., A. Madotto, C.-S. Wu, J. H. Park, and P. Fung. 2018. “Emo2Vec: Learning Generalized Emotion Representation by Multi-task Training”. *arXiv preprint arXiv:1809.04505*.
- Xu, Z., B. Liu, B. Wang, C. Sun, X. Wang, Z. Wang, and C. Qi. 2017. “Neural Response Generation via GAN with an Approximate Embedding Layer”. In: *EMNLP*. 617–626.
- Yaman, S., L. Deng, D. Yu, Y.-Y. Wang, and A. Acero. 2008. “An Integrative and Discriminative Technique for Spoken Utterance Classification”. *IEEE Transactions on Audio, Speech & Language Processing*. 16(6): 1207–1214.

- Yan, R., Y. Song, and H. Wu. 2016. "Learning to Respond with Deep Neural Networks for Retrieval-Based Human-Computer Conversation System". In: *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR*. Pisa, Italy: ACM. 55–64.
- Yang, B., W.-t. Yih, X. He, J. Gao, and L. Deng. 2015. "Embedding entities and relations for learning and inference in knowledge bases". In: *ICLR*.
- Yang, F., Z. Yang, and W. W. Cohen. 2017a. "Differentiable learning of logical rules for knowledge base completion". *CoRR*, *abs/1702.08367*.
- Yang, X., Y.-N. Chen, D. Z. Hakkani-Tür, P. Crook, X. Li, J. Gao, and L. Deng. 2017b. "End-to-end Joint Learning of Natural Language Understanding and Dialogue Manager". In: *Proceedings of the 2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 5690–5694.
- Yao, K., G. Zweig, M.-Y. Hwang, Y. Shi, and D. Yu. 2013. "Recurrent Neural Networks for Language Understanding". In: *Proceedings of the 14th Annual Conference of the International Speech Communication Association (INTERSPEECH)*. 2524–2528.
- Yao, K., G. Zweig, and B. Peng. 2015. "Attention with Intention for a Neural Network Conversation Model". In: *NIPS workshop on Machine Learning for Spoken Language Understanding and Interaction*.
- Yao, X. and B. Van Durme. 2014. "Information extraction over structured data: Question answering with Freebase". In: *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Vol. 1. 956–966.
- Yih, W.-t., M.-W. Chang, X. He, and J. Gao. 2015a. "Semantic parsing via staged query graph generation: Question answering with knowledge base". In: *ACL*. 1321–1331.
- Yih, W.-t., X. He, and J. Gao. 2015b. "Deep Learning and Continuous Representations for Natural Language Processing". In: *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Tutorial*.
- Yih, W.-t., X. He, and J. Gao. 2016. "Deep Learning and Continuous Representations for Natural Language Processing". In: *IJCAI: Tutorial*.

- Young, S. J., M. Gašić, S. Keizer, F. Mairesse, J. Schatzmann, B. Thomson, and K. Yu. 2010. “The Hidden Information State Model: A Practical Framework for POMDP-based Spoken Dialogue Management”. *Computer Speech & Language*. 24(2): 150–174.
- Young, S., C. Breslin, M. Gašić, M. Henderson, D. Kim, M. Szummer, B. Thomson, P. Tsiakoulis, and E. T. Hancock. 2016. “Evaluation of Statistical POMDP-Based Dialogue Systems in Noisy Environments”. In: *Situated Dialog in Speech-Based Human-Computer Interaction. Signals and Communication Technology*. Springer. 3–14.
- Young, S., M. Gašić, B. Thomson, and J. D. Williams. 2013. “POMDP-based statistical spoken dialog systems: A review”. *Proceedings of the IEEE*. 101(5): 1160–1179.
- Yu, A. W., D. Dohan, M.-T. Luong, R. Zhao, K. Chen, M. Norouzi, and Q. V. Le. 2018. “QANet: Combining Local Convolution with Global Self-Attention for Reading Comprehension”. *arXiv preprint arXiv:1804.09541*.
- Zhang, J., T. Zhao, and Z. Yu. 2018a. “Multimodal Hierarchical Reinforcement Learning Policy for Task-Oriented Visual Dialog”. In: *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue (SIGDIAL)*. 140–150.
- Zhang, R., J. Guo, Y. Fan, Y. Lan, J. Xu, and X. Cheng. 2018b. “Learning to Control the Specificity in Neural Response Generation”. In: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Melbourne, Australia. 1108–1117.
- Zhang, S., E. Dinan, J. Urbanek, A. Szlam, D. Kiela, and J. Weston. 2018c. “Personalizing Dialogue Agents: I have a dog, do you have pets too?”. In: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Melbourne, Australia. 2204–2213.
- Zhang, S., X. Liu, J. Liu, J. Gao, K. Duh, and B. Van Durme. 2018d. “ReCoRD: Bridging the Gap between Human and Machine Commonsense Reading Comprehension”. *arXiv preprint arXiv:1810.12885*.

- Zhang, Y., M. Galley, J. Gao, Z. Gan, X. Li, C. Brockett, and B. Dolan. 2018e. “Generating Informative and Diverse Conversational Responses via Adversarial Information Maximization”. In: *NeurIPS*. 1815–1825.
- Zhao, T. and M. Eskénazi. 2016. “Towards End-to-End Learning for Dialog State Tracking and Management using Deep Reinforcement Learning”. In: *Proceedings of the 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*. 1–10.
- Zhao, T., A. Lu, K. Lee, and M. Eskenazi. 2017. “Generative Encoder-Decoder Models for Task-Oriented Spoken Dialog Systems with Chatting Capability”. In: *ACL*. 27–36.
- Zhou, L., J. Gao, D. Li, and H.-Y. Shum. 2018. “The design and implementation of XiaoIce, an empathetic social chatbot”. *arXiv preprint arXiv:1812.08989*.