

Theory and Applications of Natural Language Processing
Monographs

Verena Rieser
Oliver Lemon

Reinforcement Learning for Adaptive Dialogue Systems

A Data-driven Methodology for
Dialogue Management and Natural
Language Generation

 Springer

Theory and Applications of Natural Language Processing

Series Editors:

Graeme Hirst (Textbooks)

Eduard Hovy (Edited volumes)

Mark Johnson (Monographs)

Aims and Scope

The field of Natural Language Processing (NLP) has expanded explosively over the past decade: growing bodies of available data, novel fields of applications, emerging areas and new connections to neighboring fields have all led to increasing output and to diversification of research.

“Theory and Applications of Natural Language Processing” is a series of volumes dedicated to selected topics in NLP and Language Technology. It focuses on the most recent advances in all areas of the computational modeling and processing of speech and text across languages and domains. Due to the rapid pace of development, the diversity of approaches and application scenarios are scattered in an ever-growing mass of conference proceedings, making entry into the field difficult for both students and potential users. Volumes in the series facilitate this first step and can be used as a teaching aid, advanced-level information resource or a point of reference.

The series encourages the submission of research monographs, contributed volumes and surveys, lecture notes and textbooks covering research frontiers on all relevant topics, offering a platform for the rapid publication of cutting-edge research as well as for comprehensive monographs that cover the full range of research on specific problem areas.

The topics include applications of NLP techniques to gain insights into the use and functioning of language, as well as the use of language technology in applications that enable communication, knowledge management and discovery such as natural language generation, information retrieval, question-answering, machine translation, localization and related fields.

The books are available in printed and electronic (e-book) form:

- * Downloadable on your PC, e-reader or iPad
- * Enhanced by Electronic Supplementary Material, such as algorithms, demonstrations, software, images and videos
- * Available online within an extensive network of academic and corporate R&D libraries worldwide
- * Never out of print thanks to innovative print-on-demand services
- * Competitively priced print editions for eBook customers thanks to MyCopy service
<http://www.springer.com/librarians/e-content/mycopy>

For further volumes:

<http://www.springer.com/series/8899>

Verena Rieser • Oliver Lemon

Reinforcement Learning for Adaptive Dialogue Systems

A Data-driven Methodology for
Dialogue Management and Natural
Language Generation

 Springer

Verena Rieser
School of Mathematical
and Computer Sciences
Heriot-Watt University
Edinburgh EH14 4AS
United Kingdom
v.t.rieser@hw.ac.uk

Oliver Lemon
School of Mathematical
and Computer Sciences
Heriot-Watt University
Edinburgh EH14 4AS
United Kingdom
o.lemon@hw.ac.uk

ISSN 2192-032X e-ISSN 2192-0338
ISBN 978-3-642-24941-9 e-ISBN 978-3-642-24942-6
DOI 10.1007/978-3-642-24942-6
Springer Heidelberg Dordrecht London New York

Library of Congress Control Number: 2011942326

Mathematics Subject Classification (2010): 68-XX, 68Txx, 68T05, 68T50, 68T37, 68T42

© Springer-Verlag Berlin Heidelberg 2011

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Preface

The past decade has seen a revolution in the field of spoken dialogue systems. As in other areas of Computer Science and Artificial Intelligence, data-driven methods are now being used to drive new methodologies for system development and evaluation. These methods are proving to be more robust, flexible, and adaptive than the largely rule-based approaches which preceded them.

We hope that this book is a contribution to that ongoing change. It describes, in detail, a new methodology for developing spoken dialogue systems – in particular the Dialogue Management and Natural Language Generation components – which starts with human data, and culminates in evaluation with real users. The journey therefore starts and ends with human behaviour in interaction, and explores methods for learning from the data, for building simulation environments for training and testing systems, and for evaluating the results.

The detailed material covers: Spoken and Multimodal dialogue systems, Wizard-of-Oz data collection, User Simulation methods, Reinforcement Learning, and Evaluation methodologies.

This book is therefore intended as research guide which navigates through a detailed case study in data-driven methods for development and evaluation of spoken dialogue systems. Common challenges associated with this approach are discussed and example solutions provided, for example, how to learn from limited amounts of data. As such, we hope it will provide insights, lessons, and inspiration for future research and development – not only for spoken dialogue systems in particular, but for data-driven approaches to human-machine interaction in general.

Edinburgh,
September 2011

Verena Rieser
Oliver Lemon

Acknowledgements

The research presented here would not have been possible without interaction with our research community, and we owe special thanks to the participants of the TALK and CLASSiC projects for discussing the ideas presented in this book.

There is a lively international research community in Spoken Dialogue Systems whose interest, criticism, and encouragement have helped to shape this work. We are especially indebted to Steve Young and his colleagues at Cambridge University, who have had a deep influence on our ideas. Members of the Interaction Lab¹ at Heriot-Watt University's School of Mathematics and Computer Science have also helped to discuss and develop these ideas over several years.

In particular we would also like to thank Dr. Xingkun Liu, Dr. Helen Hastie, Dr. Ivana Kruijff-Korbayová, Dr. Tilman Becker, and other colleagues from Saarland and Edinburgh Universities for helping with the data collection and evaluation involved in this work. We especially thank Professor Manfred Pinkal for his guidance.

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7, 2007-2013) under grant agreement number 216594 ("Computational Learning in Adaptive Systems for Spoken Conversation", CLASSiC project: www.classic-project.org), the EC FP6 project "TALK: Talk and Look, Tools for Ambient Linguistic Knowledge" (IST 507802, www.talk-project.org), from EPSRC grant numbers EP/E019501/1 and EP/G069840/1, and from the International Research Training Group "Language Technology and Cognitive Systems", Saarland University, funded by the German Science Foundation (DFG).

Finally, the first author would like to thank her parents Franz and Tatjana Rieser for their support and encouragement. The second author thanks his family for the decades of rewarding learning experiences which have made this book possible.

¹ <http://www.macs.hw.ac.uk/InteractionLab/>

Contents

1	Introduction	1
1.1	The Design Problem for Spoken Dialogue Systems	1
1.2	Overview	2
1.3	Structure of the Book	4
 Part I Fundamental Concepts		
2	Background	9
2.1	Human-Computer Interaction	10
2.2	Dialogue Strategy Development	11
2.2.1	Conventional Development Lifecycle	12
2.2.2	Evaluation and Strategy Quality Control	13
2.2.3	Strategy Implementation	17
2.2.4	Challenges for Strategy Development	19
2.3	Literature review: Learning Dialogue Strategies	21
2.3.1	Machine Learning Paradigms	21
2.3.2	Supervised Learning for Dialogue Strategies	22
2.3.3	Dialogue as Decision Making under Uncertainty	23
2.3.4	Reinforcement Learning for Dialogue Strategies	24
2.4	Summary	26
3	Reinforcement Learning	29
3.1	The Nature of Dialogue Interaction	30
3.1.1	Dialogue is Temporal	30
3.1.2	Dialogue is Dynamic	31
3.2	Reinforcement Learning-based Dialogue Strategy Learning	32
3.2.1	Dialogue as a Markov Decision Process	32
3.2.2	The Reinforcement Learning Problem	36
3.2.3	Model-based vs. Simulation-based Strategy Learning	42
3.3	Dialogue Simulation	45
3.3.1	Wizard-of-Oz Studies	45

3.3.2	Computer-based Simulations	46
3.3.3	Discussion	47
3.4	Application Domains	48
3.4.1	Information-Seeking Dialogue Systems.....	48
3.4.2	Multimodal Output Planning and Information Presentation...	49
3.4.3	Multimodal Dialogue Systems for In-Car Digital Music Players	52
3.5	Summary	52
4	Proof-of-Concept: Information Seeking Strategies	53
4.1	Introduction	53
4.1.1	A Proof-of-Concept Study.....	54
4.2	Simulated Learning Environments	55
4.2.1	Problem Representation	55
4.2.2	Database Retrieval Simulations	56
4.2.3	Noise Model	57
4.2.4	User Simulations	58
4.2.5	Objective and Reward Function	59
4.2.6	Application Scenarios	60
4.3	Threshold-based Baseline.....	61
4.4	Reinforcement Learning Method.....	63
4.4.1	Training the Policies	63
4.5	Results	65
4.6	Summary	69

Part II Policy Learning in Simulated Environments

5	The Bootstrapping Approach to Developing Reinforcement Learning-based Strategies.....	73
5.1	Motivation	74
5.1.1	Term Definition	75
5.1.2	Related Work	76
5.2	Advantages for Learning from WOZ Data	77
5.2.1	Challenges for Learning from WOZ Data	78
5.3	The Bootstrapping Method.....	79
5.3.1	Step 1: Data Collection in a Wizard-of-Oz Experiment	79
5.3.2	Step 2: Build a Simulated Learning Environment.....	81
5.3.3	Step 3: Train and test a strategy in simulation.....	81
5.3.4	Step 4: Test with Real Users	82
5.3.5	Step 5: Post-Evaluation	82
5.4	Summary	82

- 6 Data Collection in a Wizard-of-Oz Experiment 85**
 - 6.1 Experimental Setup 86
 - 6.1.1 Recruited Subjects: Wizards and Users 89
 - 6.1.2 Experimental Procedure and Task Design 90
 - 6.2 Noise Simulation 90
 - 6.2.1 Related Work 90
 - 6.2.2 Method 91
 - 6.2.3 Results and Discussion 91
 - 6.3 Corpus Description 92
 - 6.4 Analysis 94
 - 6.4.1 Qualitative Measures 94
 - 6.4.2 Subjective Ratings from the User Questionnaires 95
 - 6.5 Summary and Discussion 98

- 7 Building Simulation Environments from Wizard-of-Oz Data 101**
 - 7.1 Dialogue Strategy Learning with Simulated Environments 101
 - 7.1.1 Method and Related Work 103
 - 7.1.2 Outline 106
 - 7.2 Database Description 107
 - 7.3 Action Set Selection 108
 - 7.3.1 Method and Related Work 108
 - 7.3.2 Annotation Scheme 108
 - 7.3.3 Manual Annotation 110
 - 7.3.4 Action Set for Learning 111
 - 7.4 State Space Selection 112
 - 7.4.1 Method and Related Work 112
 - 7.4.2 Task-based State Space Features 113
 - 7.4.3 Feature Selection Techniques for Domain-specific State
Space Features 114
 - 7.5 MDP and Strategy Design 118
 - 7.5.1 Motivation 118
 - 7.5.2 Implementation 118
 - 7.5.3 Hierarchical Reinforcement Learning in the ISU Approach . . . 119
 - 7.5.4 Further System Behaviour 120
 - 7.6 Wizard Behaviour 122
 - 7.6.1 Method and Related Work 122
 - 7.6.2 Supervised Learning: Rule-based Classification 124
 - 7.7 Noise Simulation: Modelling the Effects of Mis-Communication . . . 125
 - 7.7.1 Method and Related Work 125
 - 7.7.2 Simulating the Effects of Non- and Mis-Understandings . . . 127
 - 7.8 User Simulation 128
 - 7.8.1 Method and Related Work 129
 - 7.8.2 User Actions 132
 - 7.8.3 A Simple Bi-gram Model 133
 - 7.8.4 Cluster-based User Simulation 134

- 7.8.5 Smoothed Bi-gram User Simulation 136
- 7.8.6 Evaluation of User Simulations 138
- 7.8.7 Speech Act Realisation Dependent on the User Goal 139
- 7.9 Reward and Objective Functions 142
 - 7.9.1 Method and Related Work 142
 - 7.9.2 Linear Regression for Information Acquisition 146
 - 7.9.3 Non-linear Rewards for Information Presentation 148
 - 7.9.4 Final Reward 150
- 7.10 State-Space Discretisation 151
- 7.11 Learning Experiments 152
 - 7.11.1 Training with SHARSHA 152
 - 7.11.2 Results for Testing in Simulation 154
 - 7.11.3 Qualitative Strategy Description 155
 - 7.11.4 Strategy Implementation 157
 - 7.11.5 Discussion and Error Analysis 158
- 7.12 Summary 162

Part III Evaluation and Application

- 8 Comparing Reinforcement and Supervised Learning of Dialogue Policies with Real Users 167**
 - 8.1 Policy Integration into a Dialogue System 168
 - 8.1.1 The DUDE Rapid Dialogue Development Tools 168
 - 8.1.2 Extensions to DUDE 170
 - 8.2 Experimental Setup 174
 - 8.2.1 Technical Setup 174
 - 8.2.2 Primary Driving Task 174
 - 8.2.3 Subjects and Procedure 175
 - 8.2.4 Task Types 176
 - 8.2.5 User Questionnaires 176
 - 8.3 Results 177
 - 8.3.1 Subjective User Ratings 178
 - 8.3.2 Objective Dialogue Performance 181
 - 8.4 Discussion of Real User Evaluation Results 182
 - 8.5 Meta-Evaluation 183
 - 8.5.1 Transfer Between Simulated and Real Environments 183
 - 8.5.2 Evaluation of the Learned Reward Function 184
 - 8.6 Summary 188
- 9 Adaptive Natural Language Generation 189**
 - 9.1 Introduction 190
 - 9.1.1 Previous Work on Information Presentation in SDS 190
 - 9.2 NLG as Planning Under Uncertainty 192
 - 9.3 Wizard-of-Oz Data Collection 192
 - 9.3.1 Experimental Setup and Data Collection 193

- 9.3.2 Surface Realiser 193
- 9.3.3 Human “Wizard” Baseline Strategy 194
- 9.4 The Simulation / Learning Environment 195
 - 9.4.1 User Simulations 195
 - 9.4.2 Database Matches and “Focus of Attention” 197
 - 9.4.3 Data-driven Reward Function 197
- 9.5 Reinforcement Learning Experiments 198
 - 9.5.1 Experimental Set-up 199
 - 9.5.2 Results 199
- 9.6 Evaluation with real users 202
- 9.7 Conclusion 203

- 10 Conclusion 205**
 - 10.1 Contributions 206
 - 10.2 Discussion 207
 - 10.2.1 Lessons Learned 208
 - 10.2.2 RL for Commercial Dialogue Strategy Development 209
 - 10.3 Outlook: challenges for future statistical dialogue systems 210

- Example Dialogues 213**
 - A.1 Wizard-of-Oz Example Dialogues 213
 - A.2 Example Dialogues from Simulated Interaction 216
 - A.3 Example Dialogues from User Testing 218

- Learned State-Action Mappings 223**
 - References 229

- About the Authors 253**

Acronyms

ASR	Automatic Speech Recognition
DA	Dialogue Act
DB	Database
DM	Dialogue Management
GUI	Graphical User Interface
HCI	Human Computer Interaction
IP	Information Presentation
ISU	Information State Update
MDP	Markov Decision Process
ML	Machine Learning
NLG	Natural Language Generation
NLP	Natural Language Processing
PARADISE	PARAdigm for DIAlogue System Evaluation
POMDP	Partially Observable Markov Decision Process
RL	Reinforcement Learning
SA	Speech Act
SASSI	Subjective Assessment of Speech System Interfaces
SDS	Spoken Dialogue System
SL	Supervised Learning
SLU	Spoken Language Understanding
TTS	Text-to-Speech
VOIP	Voice-Over-Internet Protocol
WER	Word-Error Rate
WOZ	Wizard-of-Oz

Chapter 1

Introduction

The past decade has seen something of a revolution in the field of spoken dialogue systems. As in other areas of Computer Science and Artificial Intelligence, data-driven methods are being used to drive new methodologies for system development and evaluation. These methods are proving to be more robust, flexible, and adaptive than the rule-based approaches which preceded them.

We hope that this book makes a contribution to that revolution. It describes, in detail, a new methodology for developing spoken dialogue systems – in particular the Dialogue Management component – which starts with human data, and ends with evaluation with real users. Related methods are now being developed further by a number of researchers worldwide. The journey begins and ends with human behaviour in interaction, and en route we explore methods for learning from such data, for building simulation environments for training and testing our systems, and methods for evaluating the results.

This book is therefore intended as a guide which navigates through a detailed case study in data-driven methods for development and evaluation of spoken dialogue systems. It focusses on Dialogue Management and Natural Language Generation, rather than speech recognition and spoken language understanding. As such, we hope that it can provide insights and lessons for future research and development – not only for spoken dialogue systems in particular, but for data-driven approaches to building better human-machine interaction in general.

1.1 The Design Problem for Spoken Dialogue Systems

The design of Spoken Dialogue Systems (SDS) is not only concerned with integrating speech and language processing modules such as Automatic Speech Recognition (ASR), Spoken Language Understanding (SLU), Natural Language Generation (NLG), and Text-to-Speech (TTS) synthesis systems. It also requires the development of skills for “what to say next”: *dialogue strategies* which take into account the performance of these components, the nature of the user’s tasks (e.g.

information-seeking, tutoring, or robot control), and other features of the operating environment such as the user's behaviour and preferences. The great variability and unpredictability of these factors makes dialogue strategy design an extremely difficult task for human developers. In conventional, rule-based, dialogue development many expensive iterations of manual design and re-design are necessary in order to produce good strategies. In addition, such hand-coded strategies are not re-usable from task to task, are not scalable, require a substantial amount of human labour and expertise, and are not guaranteed to be optimal.

For these reasons machine learning methods (such as Reinforcement Learning) for dialogue strategy design have been a leading research area for several years. These statistical computational learning approaches offer several key potential advantages over the standard rule-based hand-coding approach to dialogue systems development (Lemon and Pietquin, 2007):

- a data-driven automatic development cycle
- provably optimal action policies
- a principled mathematical model for action selection
- possibilities for generalisation to unseen states
- reduced development and deployment costs.

However, in cases where a system is designed from scratch, there is often no suitable in-domain data to enable such a design. Collecting dialogue data without a working prototype is problematic, leaving the developer with a classic “chicken-and-egg” problem. One of the main issues that this book addresses is how to use a data-driven development methodology when little or no in-domain data exists.

1.2 Overview

In this book we propose to learn dialogue strategies by simulation-based Reinforcement Learning (RL) (Sutton and Barto, 1998), where a simulated environment is learned from small amounts of Wizard-of-Oz (WOZ) data. Using WOZ data rather than data from real Human-Computer Interaction (HCI) allows us to learn optimal strategies for domains where no working dialogue system already exists. Automatic strategy learning has been applied to dialogue systems which have already been deployed in the real world using handcrafted strategies. In such work, strategy learning was performed based on already present extensive online-operation experience, e.g. (Henderson et al, 2005, 2008; Singh et al, 2002). In contrast to this preceding work, our approach enables strategy learning in domains where no prior system is available. Optimised learned strategies are then available from the first moment of online-operation, and labour-intensive handcrafting of dialogue strategies is avoided. This independence from large amounts of in-domain dialogue data allows researchers to apply RL to new application areas beyond the scope of existing dialogue systems. We call this method “bootstrapping”.

This book first provides the general proof-of-concept that RL-based strategies outperform handcrafted strategies which are manually tuned for a wide spectrum of application scenarios. After theoretically motivating our approach, we turn to the practical problem of how to learn optimal strategies for new application domains where no prior system or in-domain data are available. We propose to learn dialogue strategies by simulation-based RL, where the simulated environment is learned from small amounts of WOZ data. We therefore introduce a 5-step procedure:

1. Collect data in a WOZ experiment.
2. Use this data to construct a simulated learning environment using data-driven methods only.
3. Train a RL-based dialogue policy by interacting with the simulated environment. We compare this policy against a supervised baseline. This comparison allows us to measure the relative improvements over the WOZ strategies contained in the training data.
4. Evaluate the learned policy with real users.
5. Show that “bootstrapping” from WOZ data is a valid estimate of real HCI by comparing different aspects of the 3 corpora gathered so far: the WOZ study, the dialogues generated in simulation, and the final user tests.

It should be noted that these steps are not unique to the method introduced in this book, but most of the defined steps are required for any simulation-driven approach to strategy learning (though the last step is specific to our method). A main contribution of this book is that all these steps are now performed starting with a limited WOZ data set, and specific methods are introduced to build and validate the obtained simulations.

We apply this framework to optimise multimodal¹ Dialogue Management strategies and Natural Language Generation. In the first case we consider Dialogue Management and content selection as two closely interrelated problems for information seeking dialogues: the decision of *when* to present information depends on *how many* pieces of information to present and the available options for *how* to present them, and vice versa. We therefore formulate the problem as a hierarchy of joint learning decisions which are optimised together.

The second study describes a new approach to generating Natural Language in interactive systems. Natural Language Generation (NLG) addresses the problem of “how to say” an utterance, once “what to say” has been determined by the Dialogue Manager. We treat NLG as planning under uncertainty for information-seeking dialogue systems, where the strategy for information presentation and its associated attributes are incrementally selected using hierarchical learning. This hierarchical approach to DM and NLG has recently been explored by other researchers (Dethlefs and Cuayahuitl, 2010; Dethlefs and Cuayahuitl, 2011), and a utility-based approach to NLG is discussed by van Deemter (2009a).

¹ “Multimodal” dialogue systems are those which do not only use speech input and output, but which also use other information modalities, such as graphics (as in our case), or gesture, gaze, facial expressions, and so on.

Our results in both studies show that RL significantly outperforms supervised learning (SL) when interacting in simulation as well as for interactions with real users. For optimising multimodal Dialogue Management, the RL-based policy gains on average 50-times more reward than the SL policy when tested in simulation, and almost 18-times more reward when interacting with real users. Users also subjectively rate the RL-based policy on average 10% higher. For optimising Natural Language Generation, the trained information presentation strategies significantly improve dialogue task completion, with up to a 9.7% increase (30% relative) compared to the deployed dialogue system which uses conventional, hand-coded presentation prompts.

One focus of this book is to optimise dialogue strategies with respect to real user preferences. A major advantage of RL-based dialogue strategy development is that the dialogue strategy can be automatically trained and evaluated using the same objective function (Walker, 2005). Despite its central importance for RL, quality assurance for objective functions has received little attention so far. In fact, the reward function is one of the most hand-coded aspects of RL (Paek, 2006). Clearly, automatic optimisation and evaluation of dialogue policies, as well as quality control of the objective function, are closely inter-related problems: how can we make sure that we optimise a system according to real users' preferences? This book is the first to explore learning with data-driven, non-linear objective functions. We also propose a new method for meta-evaluation of the objective function.

Note that chapters 4 to 8 are significantly revised, updated, and extended versions of material from (Rieser, 2008).

1.3 Structure of the Book

Chapter 2 (Background)

This chapter provides the reader with relevant background knowledge for the research. After introducing some general information about Spoken Dialogue Systems, we contrast different methods applied in research and industry to develop dialogue strategies. We show how these two approaches fail to meet the current challenges for strategy design and argue for the use of statistical methods. In particular, we propose the use of Reinforcement Learning, as it uses trial-and-error exploration with delayed rewards which, we argue, is a natural model for human dialogue.

Chapter 3 (Reinforcement Learning)

This chapter provides technical background on RL for dialogue strategy development and discusses simulation-based learning in particular. We also introduce the application domain.

Chapter 4 (Proof-of-Concept: Information Seeking Strategies)

Here we develop the theoretical proof-of-concept that RL-based strategies outperform hand-coded strategies, which are tuned to the same objective function. We show this for a wide range of application scenarios, e.g. for different user types and noise conditions. This chapter also demonstrates how to apply simulation-based RL to solve a complex and challenging problem for information-seeking dialogue systems: which questions to ask the user, how many database search results to present, and when to present them, given the competing trade-offs between the length of the results list, the length of the interaction, the type of database, and noise in the speech recognition environment. In this chapter the reader will receive a deeper understanding of the principles and advantages of Reinforcement Learning for dialogue strategy learning. We also explain why data-driven simulation approaches are preferred over manually constructed simulated environments (as done here for the theoretical proof-of-concept).

Chapter 5 (A Bootstrapping Approach to Develop Reinforcement Learning-based Strategies)

This chapter introduces a 5-step procedure model to bootstrap optimal RL-based strategies for WOZ data. The resulting strategies are tailored to the application environment, do not require a working prototype system, and are optimised with respect to real user preferences. We also explain how we meet the challenges when learning from WOZ data. In particular, we introduce the problem of how to construct a simulated environment from limited amounts of WOZ data. We also discuss the fact that a WOZ study itself is a simulation of real HCI.

Chapter 6 (Data Collection in a Wizard-of-Oz Experiment)

Here we describe the experimental setup of the WOZ experiment. We explain which changes to the conventional WOZ method are necessary for strategy learning. In particular, we introduce an utterance distortion method in order to resemble noise conditions for real dialogue systems. In addition, we explore the “intuitive” strategies that were applied by our human wizards.

Chapter 7 (Building a Simulated Learning Environment from Wizard-of-Oz Data)

This chapter uses the WOZ data to construct a simulated learning environment. We therefore introduce methods suited to build and validate simulations from small amounts of data. In particular, we construct the action set and state space by exploring the wizards’ actions. The user and noise simulations are obtained using

frequency-based approaches. The objective function is a predictive model of user ratings obtained by a regression analysis, following the PARADISE framework of (Walker et al, 1997). We then train a RL-based dialogue policy by interacting with the simulated environment, and we compare this strategy against a baseline constructed by Supervised Learning. This comparison allows us to measure the relative improvements over the wizard strategies obtained from the training data.

Chapter 8 (Comparing Reinforcement and Supervised Learning of Dialogue Policies with Real Users)

In this chapter we evaluate the learned strategy with real users. We therefore develop a music-player dialogue system using a rapid development tool, where the learned strategy is implemented using a table look-up between states and learned actions. We report detailed results from the real user tests.

We also post-evaluate our overall “bootstrapping” approach by comparing different aspects of the 3 corpora gathered so far: the WOZ study, the dialogues generated in simulation, and the final user tests. We first evaluate whether strategies learned in simulation do transfer to tests with real users, and we also compare the experimental conditions of the different studies, where we discuss the noise model in particular. Furthermore, we explore whether the objective function used for learning is a realistic estimate of real user preferences.

Chapter 9 (Natural Language Generation)

This chapter further develops the methodology to encompass elements of policy learning for adaptive Natural Language Generation in spoken dialogue systems. This chapter shows that our method can quite easily be applied to new domains and tasks. We show how to develop a data-driven approach to content selection and structuring decisions in NLG, this time in the domain of a restaurant recommendation SDS. We also report results from evaluations with both simulated and real users. The real user evaluation shows that improved NLG can lead to significant improvements in overall dialogue system performance.

Chapter 10 (Conclusion)

Finally, we conclude by summarising the main contributions of this work. We also report on “lessons learned” to provide guidance for future researchers. For example, we discuss the data quantity and quality required for the proposed bootstrapping approach. In the final outlook of this chapter we discuss ongoing challenges for research in this field.

Part I
Fundamental Concepts

Chapter 2

Background




Fig. 2.1 Wolfgang von Kempelen's speaking machine, drawing from *Mechanism of Human Speech* (1791)

Mankind appears to be fascinated by the idea of talking with machines. The first attempts to produce human speech by machine were made in the second half of the 18th century. One of the best known examples is Wolfgang von Kempelen's speak-

ing machine, as described in his book “Mechanism of Human Speech”¹ (1791), see Figure 2. Von Kempelen’s machine was the first that produced not only some speech sounds, but also whole words and short sentences. Clearly, in order to converse with a machine in natural language more than just automatic sound production is necessary. Descartes even thought that it would never be possible to engage in dialogue with machines at all. In his book “Discourse on the Method of Rightly Conducting the Reason, and Searching for Truth in the Sciences”² (1637) he declares:

“... but if there were machines bearing the image of our bodies, and capable of imitating our actions as far as it is morally possible, [...] they could never use words or other signs arranged in such a manner as is competent to us in order to declare our thoughts to others.”

Today, however, it is arguable that such machines now exist, at least for limited application domains, due to major advances in the field of Human-Computer Interaction (HCI) and Spoken Dialogue Systems (SDS). Still, human-machine dialogue is far from resembling the capabilities of human-human dialogue, as we discuss below.

2.1 Human-Computer Interaction

For computers, holding a conversation is difficult. Engaging in a conversation requires more than just technical language proficiency. When people engage in dialogue, they carry out a purposeful *activity* (Austin, 1962), a *joint action* (Clark, 1996), or a *language game* (Wittgenstein, 1953), which they know how to perform using their communicative skills. Dialogue behaviour is often formally described as a sequence of Speech Acts (SAs) (Searle, 1969). These SAs are organised into structural patterns in the field of Conversation Analysis, e.g. (Levinson, 1983; Sacks et al, 1974). Some of this behaviour follows standardised cultural conventions. For example, the fact that people greet each other is described as the standardised SA “adjacency pair” *greeting-greeting* (Levinson, 1983). Other behaviour is highly context-dependent. For example, in previous work we show that the way people ask for clarification is influenced by various contextual and environmental factors, such as dialogue type, modality, and channel quality (Rieser and Moore, 2005). In addition, people often engage in dialogue to solve a task together. Their behaviour is then driven by the goal of the task as well as by their “mental model” (Johnson-Laird, 1983) of the other dialogue participant.

Humans acquire these communicative skills over time, but for a dialogue system, they need to be developed by a dialogue designer. This usually is an expert who defines a *dialogue strategy*, which “tells” the system what to do in specific situations. The “dialogue strategy” is part of the Dialogue Manager (DM) which controls the

¹ Original title: *Mechanismus der menschlichen Sprache nebst Beschreibung einer sprechenden Maschine*

² Original title: *Discours de la méthode pour bien conduire sa raison, et chercher la vérité dans les sciences*

behaviour of the system. Broadly speaking, a dialogue system has three modules, one each for input, output, and control, as shown in [Figure 2.2](#). The input module commonly comprises Automatic Speech Recognition (ASR) and Spoken Language Understanding (SLU). The control module corresponds to the Dialogue Manager, which executes a dialogue strategy. The output module consists of a Natural Language Generation (NLG) system and a Text-To-Speech (TTS) engine. Usually, these modules are placed in a pipeline model (see [Figure 2.2](#)). The ASR converts the user's speech input (1) into text (2). SLU parses the text into a string of meaningful concepts, intentions, or Speech Acts (3). The Dialogue Manager maintains an internal state and decides what SA action to take next (4). This is what we call a dialogue strategy. For most applications the DM is also connected to a back-end database. In the output module, NLG renders the communicative acts (4) as text (5), and the TTS engine converts text to audio (6) for the user. Interested readers are referred to introductory texts such as (Bernsen et al, 1998; Huang et al, 2001; Jurafsky and Martin, 2000; McTear, 2004).

Human-Computer Interaction (HCI) is the study of interaction between people (users) and computers (such as dialogue systems). Human-machine dialogue differs from human-human dialogue in various ways. The most prominent features are the lack of deep language understanding and the lack of pragmatic competence (communicative skills) of the system. The lack of language understanding is due to errors introduced by less-than-perfect input processing (ASR and NLU), and the common use of shallow semantic representations. The lack of pragmatic competence is mainly due to the limited capabilities (often hand-coded heuristics) of the control module.

A substantial amount of recent work targets the problem of limited language understanding capabilities with so-called “error handling”, e.g. (Bohus, 2007; Frampton, 2008; Skantze, 2007a), or “uncertainty handling” mechanisms, e.g. (Thomson and Young, 2010; Williams, 2006; Williams and Young, 2007a). This book addresses the problem of pragmatic competence: how to improve the communicative skills of a system by providing effective mechanisms to develop better dialogue strategies. In particular, this book explains how to automatically learn these skills from experience using simulated interactions, starting with a small experimental study of human behaviour. We now explain the conventional methods for strategy development.

2.2 Dialogue Strategy Development

There is a wide range of techniques to develop dialogue strategies, and techniques applied in industry are very different from the ones applied in research (Griol et al, 2010; Pieraccini and Huerta, 2005; Williams, 2008). These differences reflect the fact that academia and industry often pursue different objectives. Academic systems often aim to emulate human behaviour in order to generate ‘natural’ behaviour, whereas commercial systems are required to be robust interfaces in order to solve