
Reinforcement Learning For Adaptive Dialogue Systems A Data Driven Methodology For Dialogue Management And Natural Language Generation Theory And Applications Of Natural Language Processing

ECAI 2014

Reinforcement Learning

Natural Human-Interface Technologies

Adaptive Technologies for Training and Education

Empirical Methods in Natural Language Generation

Computational Learning for Conversational Interfaces

Hierarchical Reinforcement Learning for Spoken Dialogue Systems

Artificial Paranoia

Reinforcement Learning for Adaptive Dialogue Systems

Data-oriented Methods and Empirical Evaluation

25th International Conference, ICONIP 2018, Siem Reap, Cambodia, December 13-16, 2018, Proceedings, Part III

Data-oriented Methods and Empirical Evaluation

Enablements, Analyses, and Evaluation

22nd European Conference on Artificial Intelligence, 29 August - 2 September 2016, The Hague, The Netherlands - Including

Prestigious Applications of Artificial Intelligence (PAIS 2016)

Universal Access in Human-Computer Interaction. Design for All and eInclusion

Language Processing, Software, Commercialization, and Emerging Directions

Handbook of Technical Communication

Future and Emerging Trends in Language Technology. Machine Learning and Big Data

Theory and Applications of Models of Computation
Man-Machine Dialogue
Semantics in Adaptive and Personalised Systems
Empirical Methods in Natural Language Generation
5th International Conference, ICAART 2013, Barcelona, Spain, February 15-18, 2013. Revised Selected Papers
Computational Learning for Conversational Interfaces
ECAI 2016
Spoken Dialogue Systems for Ambient Environments
Neural Information Processing
7th Hellenic Conference on AI, SETN 2012, Lamia, Greece, May 28-31, 2012, Proceedings
Social Robotics
Computational Linguistics and Intelligent Text Processing
8th International Conference, ICSR 2016, Kansas City, MO, USA, November 1-3, 2016 Proceedings
State-of-the-Art
6th Annual Conference, TAMC 2009, Changsha, China, May 18-22, 2009. Proceedings
Learning Optimal Discourse Strategies in a Spoken Dialogue System
Modeling the Listener's Perspective in Passive and Interactive Scenarios
15th International Conference, CICLing 2014, Kathmandu, Nepal, April 6-12, 2014, Proceedings, Part I
Data-Driven Methods for Adaptive Spoken Dialogue Systems
Automated Machine Learning and Meta-Learning for Multimedia
13th International Conference, IVA 2013, Edinburgh, UK, August 29-31, 2013, Proceedings
Second International Workshop, IWSDS 2010, Gotemba, Shizuoka, Japan, October 1-2, 2010. Proceedings

*Reinforcement Learning For Adaptive
Dialogue Systems A Data Driven
Methodology For Dialogue
Management And Natural Language
Generation Theory And Applications Of
Natural Language Processing*

Downloaded from <ftp.wtvq.com> by guest

JANIAH KADE

ECAI 2014 Springer

Adaptive Multimodal Interactive Systems introduces a general framework for adapting multimodal interactive systems and comprises a detailed discussion of each of the steps required for

adaptation. This book also investigates how interactive systems may be improved in terms of usability and user friendliness while describing the exhaustive user tests employed to evaluate the presented approaches. After introducing general theory, a generic approach for user modeling in interactive systems is presented, ranging from an observation of basic events to a description of higher-level user behavior. Adaptations are presented as a set of patterns similar to those known from software or usability engineering. These patterns describe recurring problems and present proven solutions. The authors include a discussion on when and how to employ patterns and provide guidance to the system designer who wants to add adaptivity to interactive systems. In addition to these patterns, the book introduces an adaptation framework, which exhibits an abstraction layer using Semantic Web technology. Adaptations are implemented on top of this abstraction layer by creating a semantic representation of the adaptation patterns. The patterns cover both graphical interfaces as well as speech-based and multimodal interactive systems.

Reinforcement Learning IOS Press

This monograph gives a complete overview of the techniques and the methods for semantics-aware content representation and shows how to apply such techniques in various use cases, such as recommender systems, user profiling and social media analysis. Throughout the book, the authors provide an extensive analysis of the techniques currently proposed in the literature and cover all the available tools and libraries to implement and exploit such methodologies in real-world scenarios. The book first introduces the problem of information overload and the reasons why

content-based information needs to be taken into account. Next, the basics of Natural Language Processing are provided, by describing operations such as tokenization, stopword removal, lemmatization, stemming, part-of-speech tagging, along with the main problems and issues. Finally, the book describes the different approaches for semantics-aware content representation: such approaches are split into 'exogenous' and 'endogenous' ones, depending on whether external knowledge sources as DBpedia or geometrical models and distributional semantics are used, respectively. To conclude, several successful use cases and an extensive list of available tools and resources to implement the approaches are shown. Semantics in Adaptive and Personalised Systems definitely fills the gap between the extensive literature on content-based recommender systems, natural language processing, and the different types of semantics-aware representations.

Natural Human-Interface Technologies Springer

Theory and Applications of Models of Computation (TAMC) is an international conference series with an interdisciplinary character, bringing together researchers working in computer science, mathematics (especially logic) and the physical sciences. This crossdisciplinary character, together with its focus on algorithms, complexity and computability theory, gives the conference a special flavor and distinction.

TAMC2009 was the sixth conference in the series. The previous meetings

were held during May 17-19, 2004 in Beijing, May 17-20, 2005 in Kunming, May 15-20, 2006 in Beijing, May 22-25, 2007 in Shanghai, and April 25-29, 2008 in Xi'an. TAMC 2009 was held in Changsha,

during May 18–22, 2009. Next year will see a new departure, namely, the first TAMC conference to be held outside of Asia. TAMC 2010 will be held in Prague, capital of the Czech Republic. At TAMC 2009 we had three plenary speakers, Leslie Valiant (Harvard University, USA), Moshe Vardi (Rice University, USA) and Matthew Hennessy (Trinity College, Ireland), giving one-hour talks each. Professor Valiant spoke on “Neural Computations That Support Long Mixed Sequence of Knowledge Acquisition Tasks,” Professor Vardi on “Constraints, Graphs, Algebra, Logic, and Complexity,” and Professor Hennessy on “Distributed Systems and Their Environments.” Their respective abstracts accompanying the talks are included in these proceedings.

[Adaptive Technologies for Training and Education](#) Springer
This book compiles and presents a synopsis on current global research efforts to push forward the state of the art in dialogue technologies, including advances to language and context understanding, and dialogue management, as well as human–robot interaction, conversational agents, question answering and lifelong learning for dialogue systems.

[Empirical Methods in Natural Language Generation](#) Springer
Nature
This monograph is the first survey of neural approaches to conversational AI that targets Natural Language Processing and Information Retrieval audiences. It provides 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. The authors draw connections between modern neural approaches and traditional approaches, allowing readers to

better understand why and how the research has evolved and to shed light on how they can move forward. They also present state-of-the-art approaches to training dialogue agents using both supervised and reinforcement learning. Finally, the authors 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 are still being faced.

[Neural Approaches to Conversational AI](#) is 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.

[Computational Learning for Conversational Interfaces](#) Springer
This book constitutes the thoroughly refereed post-conference proceedings of the 5th International Conference on Agents and Artificial Intelligence, ICAART 2013, held in Barcelona, Spain, in February 2013. The 20 revised full papers presented together with one invited paper were carefully reviewed and selected from 269 submissions. The papers are organized in two topical sections on artificial intelligence and on agents.

[Hierarchical Reinforcement Learning for Spoken Dialogue Systems](#) IOS Press

Reinforcement learning encompasses both a science of adaptive behavior of rational beings in uncertain environments and a computational methodology for finding optimal behaviors for challenging problems in control, optimization and adaptive

behavior of intelligent agents. As a field, reinforcement learning has progressed tremendously in the past decade. The main goal of this book is to present an up-to-date series of survey articles on the main contemporary sub-fields of reinforcement learning. This includes surveys on partially observable environments, hierarchical task decompositions, relational knowledge representation and predictive state representations. Furthermore, topics such as transfer, evolutionary methods and continuous spaces in reinforcement learning are surveyed. In addition, several chapters review reinforcement learning methods in robotics, in games, and in computational neuroscience. In total seventeen different subfields are presented by mostly young experts in those areas, and together they truly represent a state-of-the-art of current reinforcement learning research. Marco Wiering works at the artificial intelligence department of the University of Groningen in the Netherlands. He has published extensively on various reinforcement learning topics. Martijn van Otterlo works in the cognitive artificial intelligence group at the Radboud University Nijmegen in The Netherlands. He has mainly focused on expressive knowledge representation in reinforcement learning settings.

Artificial Paranoia John Wiley & Sons

The seven-volume set of LNCS 11301-11307, constitutes the proceedings of the 25th International Conference on Neural Information Processing, ICONIP 2018, held in Siem Reap, Cambodia, in December 2018. The 401 full papers presented were carefully reviewed and selected from 575 submissions. The papers address the emerging topics of theoretical research, empirical studies, and applications of neural information

processing techniques across different domains. The third volume, LNCS 11303, is organized in topical sections on embedded learning, transfer learning, reinforcement learning, and other learning approaches.

Reinforcement Learning for Adaptive Dialogue Systems Springer

This book constitutes the proceedings of the 13th International Conference on Intelligent Virtual Agents, IVA 2013, held in Edinburgh, UK, in August 2013. There was a total of 94 submissions. The 18 full and 18 short papers presented in this volume were carefully reviewed and selected for inclusion in the book. In addition, the volume lists the 34 posters which were on display during the conference. The papers are organized in topical sections named: cognitive models; applications; dialogue, language, speech; non-verbal behaviour; and social, cultural models and agents.

Data-oriented Methods and Empirical Evaluation Reinforcement Learning for Adaptive Dialogue Systems A Data-driven Methodology for Dialogue Management and Natural Language Generation

The Handbook of Technical Communication brings together a variety of topics which range from the role of technical media in human communication to the linguistic, multimodal enhancement of present-day technologies. It covers the area of computer-mediated text, voice and multimedia communication as well as of technical documentation. In doing so, the handbook takes professional and private communication into account. Special emphasis is put on technical communication by means of web 2.0 technologies and its standardization in system development. In summary, the handbook deals with theoretical issues of technical

communication and its practical impact on the development and usage of text and speech technologies.

25th International Conference, ICONIP 2018, Siem Reap, Cambodia, December 13-16, 2018, Proceedings, Part III
Springer

Natural language generation (NLG) is a subfield of natural language processing (NLP) that is often characterized as the study of automatically converting non-linguistic representations (e.g., from databases or other knowledge sources) into coherent natural language text. In recent years the field has evolved substantially. Perhaps the most important new development is the current emphasis on data-oriented methods and empirical evaluation. Progress in related areas such as machine translation, dialogue system design and automatic text summarization and the resulting awareness of the importance of language generation, the increasing availability of suitable corpora in recent years, and the organization of shared tasks for NLG, where different teams of researchers develop and evaluate their algorithms on a shared, held out data set have had a considerable impact on the field, and this book offers the first comprehensive overview of recent empirically oriented NLG research.

Data-oriented Methods and Empirical Evaluation Springer
Science & Business Media

Artificial Intelligence continues to be one of the most exciting and fast-developing fields of computer science. This book presents the 177 long papers and 123 short papers accepted for ECAI 2016, the latest edition of the biennial European Conference on Artificial Intelligence, Europe's premier venue for presenting

scientific results in AI. The conference was held in The Hague, the Netherlands, from August 29 to September 2, 2016. ECAI 2016 also incorporated the conference on Prestigious Applications of Intelligent Systems (PAIS) 2016, and the Starting AI Researcher Symposium (STAIRS). The papers from PAIS are included in this volume; the papers from STAIRS are published in a separate volume in the Frontiers in Artificial Intelligence and Applications (FAIA) series. Organized by the European Association for Artificial Intelligence (EurAI) and the Benelux Association for Artificial Intelligence (BNVKI), the ECAI conference provides an opportunity for researchers to present and hear about the very best research in contemporary AI. This proceedings will be of interest to all those seeking an overview of the very latest innovations and developments in this field.

Enablements, Analyses, and Evaluation Springer

This book discusses the Partially Observable Markov Decision Process (POMDP) framework applied in dialogue systems. It presents POMDP as a formal framework to represent uncertainty explicitly while supporting automated policy solving. The authors propose and implement an end-to-end learning approach for dialogue POMDP model components. Starting from scratch, they present the state, the transition model, the observation model and then finally the reward model from unannotated and noisy dialogues. These altogether form a significant set of contributions that can potentially inspire substantial further work. This concise manuscript is written in a simple language, full of illustrative examples, figures, and tables.

22nd European Conference on Artificial Intelligence, 29 August - 2 September 2016, The Hague, The Netherlands - Including

Prestigious Applications of Artificial Intelligence (PAIS 2016)

Springer Science & Business Media

The Handbook of Multimodal-Multisensor Interfaces provides the first authoritative resource on what has become the dominant paradigm for new computer interfaces-user input involving new media (speech, multi-touch, hand and body gestures, facial expressions, writing) embedded in multimodal-multisensor interfaces. This three-volume handbook is written by international experts and pioneers in the field. It provides a textbook, reference, and technology roadmap for professionals working in this and related areas. This third volume focuses on state-of-the-art multimodal language and dialogue processing, including semantic integration of modalities. The development of increasingly expressive embodied agents and robots has become an active test bed for coordinating multimodal dialogue input and output, including processing of language and nonverbal communication. In addition, major application areas are featured for commercializing multimodal-multisensor systems, including automotive, robotic, manufacturing, machine translation, banking, communications, and others. These systems rely heavily on software tools, data resources, and international standards to facilitate their development. For insights into the future, emerging multimodal-multisensor technology trends are highlighted in medicine, robotics, interaction with smart spaces, and similar areas. Finally, this volume discusses the societal impact of more widespread adoption of these systems, such as privacy risks and how to mitigate them. The handbook chapters provide a number of walk-through examples of system design and processing, information on practical resources for developing

and evaluating new systems, and terminology and tutorial support for mastering this emerging field. In the final section of this volume, experts exchange views on a timely and controversial challenge topic, and how they believe multimodal-multisensor interfaces need to be equipped to most effectively advance human performance during the next decade.

Universal Access in Human-Computer Interaction. Design for All and eInclusion Springer

The role of artificial intelligence (AI) applications in fields as diverse as medicine, economics, linguistics, logical analysis and industry continues to grow in scope and importance. AI has become integral to the effective functioning of much of the technical infrastructure we all now take for granted as part of our daily lives. This book presents the papers from the 21st biennial European Conference on Artificial Intelligence, ECAI 2014, held in Prague, Czech Republic, in August 2014. The ECAI conference remains Europe's principal opportunity for researchers and practitioners of Artificial Intelligence to gather and to discuss the latest trends and challenges in all subfields of AI, as well as to demonstrate innovative applications and uses of advanced AI technology. Included here are the 158 long papers and 94 short papers selected for presentation at the conference. Many of the papers cover the fields of knowledge representation, reasoning and logic as well as agent-based and multi-agent systems, machine learning, and data mining. The proceedings of PAIS 2014 and the PAIS System Demonstrations are also included in this volume, which will be of interest to all those wishing to keep abreast of the latest developments in the field of AI.

Language Processing, Software, Commercialization, and

Emerging Directions Springer

Artificial Paranoia: A Computer Simulation of Paranoid Processes is a seven-chapter book that begins by explaining the concept, characteristics, and theories of paranoia. Subsequent chapters focus on the explanations, models, and symbol-processing theory of the paranoid mode. Another chapter explores language-recognition processes for understanding dialogues in teletyped psychiatric interviews. The last three chapters explore the central processes of the model, validation, and evaluation.

Handbook of Technical Communication MIT Press

This two-volume set, consisting of LNCS 8403 and LNCS 8404, constitutes the thoroughly refereed proceedings of the 14th International Conference on Intelligent Text Processing and Computational Linguistics, CICLing 2014, held in Kathmandu, Nepal, in April 2014. The 85 revised papers presented together with 4 invited papers were carefully reviewed and selected from 300 submissions. The papers are organized in the following topical sections: lexical resources; document representation; morphology, POS-tagging, and named entity recognition; syntax and parsing; anaphora resolution; recognizing textual entailment; semantics and discourse; natural language generation; sentiment analysis and emotion recognition; opinion mining and social networks; machine translation and multilingualism; information retrieval; text classification and clustering; text summarization; plagiarism detection; style and spelling checking; speech processing; and applications.

Future and Emerging Trends in Language Technology. Machine Learning and Big Data Springer Science & Business Media

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. This book is a unique contribution to that ongoing change. A new methodology for developing spoken dialogue systems is described in detail. The journey 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. The book is a research guide for students and researchers with a background in Computer Science, AI, or Machine Learning. It 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 are provided. This work provides 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.

Theory and Applications of Models of Computation Walter de Gruyter

This book summarizes the main problems posed by the design of human-machine dialogue system and offers ideas on how to continue along the path towards efficient, realistic and fluid communication between humans and machines. A culmination of ten years of research, it is based on the author's development, investigation and experimentation covering a

multitude of fields, including artificial intelligence, automated language processing, man-machine interfaces and notably multimodal or multimedia interfaces. Contents Part 1. Historical and Methodological Landmarks 1. An Assessment of the Evolution of Research and Systems. 2. Man-Machine Dialogue Fields. 3. The Development Stages of a Dialogue System. 4. Reusable System Architectures. Part 2. Inputs Processing 5. Semantic Analyses and Representations. 6. Reference Resolution. 7. Dialogue Acts Recognition. Part 3. System Behavior and Evaluation 8. A Few Dialogue Strategies. 9. Multimodal Output Management. 10. Multimodal Dialogue System Assessment. About the Authors Frédéric Landragin is a computer science engineer and has a PhD from the University of Lorraine, France. He is currently in charge of linguistics research for the French National Center for Scientific Research (CNRS). His studies focus on the analysis and modeling of language interpretation. Man-machine dialogue is one of the applications of this research.

Man-Machine Dialogue Springer

This thesis focuses on the problem of scalable optimization of dialogue behaviour in speech-based conversational systems using reinforcement learning. Most previous investigations in dialogue strategy learning have proposed flat reinforcement learning methods, which are more suitable for small-scale spoken dialogue systems. This research formulates the problem in terms of Semi-Markov Decision Processes (SMDPs), and proposes two hierarchical reinforcement learning methods to optimize sub-dialogues rather than full dialogues. The first method uses a hierarchy of SMDPs, where every SMDP ignores irrelevant state variables and actions in order to optimize a sub-dialogue. The

second method extends the first one by constraining every SMDP in the hierarchy with prior expert knowledge. The latter method proposes a learning algorithm called 'HAM+HSMQ-Learning', which combines two existing algorithms in the literature of hierarchical reinforcement learning. Whilst the first method generates fully-learned behaviour, the second one generates semi-learned behaviour. In addition, this research proposes a heuristic dialogue simulation environment for automatic dialogue strategy learning. Experiments were performed on simulated and real environments based on a travel planning spoken dialogue system. Experimental results provided evidence to support the following claims: First, both methods scale well at the cost of near-optimal solutions, resulting in slightly longer dialogues than the optimal solutions. Second, dialogue strategies learned with coherent user behaviour and conservative recognition error rates can outperform a reasonable hand-coded strategy. Third, semi-learned dialogue behaviours are a better alternative (because of their higher overall performance) than hand-coded or fully-learned dialogue behaviours. Last, hierarchical reinforcement learning dialogue agents are feasible and promising for the (semi) automatic design of adaptive behaviours in larger-scale spoken dialogue systems. This research makes the following contributions to spoken dialogue systems which learn their dialogue behaviour. First, the Semi-Markov Decision Process (SMDP) model was proposed to learn spoken dialogue strategies in a scalable way. Second, the concept of 'partially specified dialogue strategies' was proposed for integrating simultaneously hand-coded and learned spoken dialogue behaviours into a single learning framework. Third, an evaluation with real users of

hierarchical reinforcement learning dialogue agents was essential to validate their effectiveness in a realistic environment.