Springer Theses
Recognizing Outstanding Ph.D. Research

Tim Baarslag

Exploring the Strategy Space of Negotiating Agents

A Framework for Bidding, Learning and Accepting in Automated Negotiation

Springer

Springer Theses

Recognizing Outstanding Ph.D. Research

Aims and Scope

The series "Springer Theses" brings together a selection of the very best Ph.D. theses from around the world and across the physical sciences. Nominated and endorsed by two recognized specialists, each published volume has been selected for its scientific excellence and the high impact of its contents for the pertinent field of research. For greater accessibility to non-specialists, the published versions include an extended introduction, as well as a foreword by the student's supervisor explaining the special relevance of the work for the field. As a whole, the series will provide a valuable resource both for newcomers to the research fields described, and for other scientists seeking detailed background information on special questions. Finally, it provides an accredited documentation of the valuable contributions made by today's younger generation of scientists.

Theses are accepted into the series by invited nomination only and must fulfill all of the following criteria

- They must be written in good English.
- The topic should fall within the confines of Chemistry, Physics, Earth Sciences, Engineering and related interdisciplinary fields such as Materials, Nanoscience, Chemical Engineering, Complex Systems and Biophysics.
- The work reported in the thesis must represent a significant scientific advance.
- If the thesis includes previously published material, permission to reproduce this must be gained from the respective copyright holder.
- They must have been examined and passed during the 12 months prior to nomination.
- Each thesis should include a foreword by the supervisor outlining the significance of its content.
- The theses should have a clearly defined structure including an introduction accessible to scientists not expert in that particular field.

More information about this series at http://www.springer.com/series/8790

Tim Baarslag

Exploring the Strategy Space of Negotiating Agents

A Framework for Bidding, Learning and Accepting in Automated Negotiation

Nominated as an outstanding PhD thesis by
Delft University of Technology, The Netherlands

Springer

Author
Dr. Tim Baarslag
Interactive Intelligence Group, Department of Intelligent Systems, Faculty of Electrical Engineering, Mathematics and Computer Science
Delft University of Technology
Delft
The Netherlands

Supervisors
Prof. Catholijn M. Jonker
Interactive Intelligence Group, Department of Intelligent Systems, Faculty of Electrical Engineering, Mathematics and Computer Science
Delft University of Technology
Delft
The Netherlands

Dr. Koen V. Hindriks
Interactive Intelligence Group, Department of Intelligent Systems, Faculty of Electrical Engineering, Mathematics and Computer Science
Delft University of Technology
Delft
The Netherlands

ISSN 2190-5053 ISSN 2190-5061 (electronic)
Springer Theses
ISBN 978-3-319-28242-8 ISBN 978-3-319-28243-5 (eBook)
DOI 10.1007/978-3-319-28243-5

Library of Congress Control Number: 2015958900

Printed on acid-free paper

This Springer imprint is published by SpringerNature
The registered company is Springer International Publishing AG Switzerland

Parts of this thesis have been published in the following articles:

The chapters in this thesis are based on publications in scientific journals and/or peer-reviewed conference proceedings, including the appendices. At the beginning of every chapter, we specify where parts have already been published. The full list is given below.

Journal Papers

Tim Baarslag, Mark J.C. Hendrikx, Koen V. Hindriks, and Catholijn M. Jonker. Learning about the opponent in automated bilateral negotiation: a comprehensive survey of opponent modeling techniques. *Autonomous Agents and Multi-Agent Systems*, pages 1–50, 2015

Raz Lin, Sarit Kraus, Tim Baarslag, Dmytro Tykhonov, Koen V. Hindriks, and Catholijn M. Jonker. Genius: An integrated environment for supporting the design of generic automated negotiators. *Computational Intelligence*, 30(1):48–70, 2014

Tim Baarslag, Katsuhide Fujita, Enrico H. Gerding, Koen V. Hindriks, Takayuki Ito, Nicholas R. Jennings, Catholijn M. Jonker, Sarit Kraus, Raz Lin, Valentin Robu, and Colin R. Williams. Evaluating practical negotiating agents: Results and analysis of the 2011 international competition. *Artificial Intelligence*, 198:73–103, May 2013

Tim Baarslag, Koen V. Hindriks, and Catholijn M. Jonker. Effective acceptance conditions in real-time automated negotiation. *Decision Support Systems*, 60:68–77, Apr 2014

Conference Papers

Tim Baarslag, Alexander S.Y. Dirkzwager, Koen V. Hindriks, and Catholijn M. Jonker. The significance of bidding, accepting and opponent modeling in automated negotiation. In *21st European Conference on Artificial Intelligence, volume 263 of Frontiers in Artificial Intelligence and Applications*, pages 27–32, 2014

Tim Baarslag, Mark J.C. Hendrikx, Koen V. Hindriks, and Catholijn M. Jonker. Predicting the performance of opponent models in automated negotiation. In *International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT), 2013 IEEE/WIC/ACM*, volume 2, pages 59–66, Nov 2013

Tim Baarslag and Koen V. Hindriks. Accepting optimally in automated negotiation with incomplete information. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '13, pages 715–722, Richland, SC, 2013. International Foundation for Autonomous Agents and Multiagent Systems

Tim Baarslag. Designing an automated negotiator: Learning what to bid and when to stop. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '13, pages 1419–1420, Richland, SC, 2013. International Foundation for Autonomous Agents and Multiagent Systems

Book Chapters

Tim Baarslag. *What to Bid and When to Stop*. Dissertation, Delft University of Technology, Sep 2014

Tim Baarslag, Koen V. Hindriks, Mark J.C. Hendrikx, Alex S.Y. Dirkzwager, and Catholijn M. Jonker. Decoupling negotiating agents to explore the space of negotiation strategies. In Ivan Marsa-Maestre, Miguel A. Lopez-Carmona, Takayuki Ito, Minjie Zhang, Quan Bai, and Katsuhide Fujita, editors, *Novel Insights in Agent-based Complex Automated Negotiation, volume 535 of Studies in Computational Intelligence*, pages 61–83. Springer, Japan, 2014

Tim Baarslag, Koen V. Hindriks, and Catholijn M. Jonker. Acceptance conditions in automated negotiation. In Takayuki Ito, Minjie Zhang, Valentin Robu, and Tokuro Matsuo, editors, *Complex Automated Negotiations: Theories, Models, and Software Competitions*, volume 435 of *Studies in Computational Intelligence*, pages 95–111. Springer Berlin Heidelberg, 2013

Tim Baarslag, Koen V. Hindriks, and Catholijn M. Jonker. A tit for tat negotiation strategy for real-time bilateral negotiations. In Takayuki Ito, Minjie Zhang, Valentin Robu, and Tokuro Matsuo, editors, *Complex Automated Negotiations: Theories, Models, and Software Competitions*, volume 435 of *Studies in Computational Intelligence*, pages 229–233. Springer Berlin Heidelberg, 2013

Katsuhide Fujita, Takayuki Ito, Tim Baarslag, Koen V. Hindriks, Catholijn M. Jonker, Sarit Kraus, and Raz Lin. The second automated negotiating agents competition (ANAC 2011). In Takayuki Ito, Minjie Zhang, Valentin Robu, and Tokuro Matsuo, editors, *Complex Automated Negotiations: Theories, Models, and Software Competitions*, volume 435 of *Studies in Computational Intelligence*, pages 183–197. Springer Berlin Heidelberg, 2013

Reyhan Aydoğan, Tim Baarslag, Koen V. Hindriks, Catholijn M. Jonker, and Pınar Yolum. Heuristic-based approaches for CP-nets in negotiation. In Takayuki Ito, Minjie Zhang, Valentin Robu, and Tokuro Matsuo, editors, *Complex Automated Negotiations: Theories, Models, and Software Competitions*, volume 435 of *Studies in Computational Intelligence*, pages 113–123. Springer Berlin Heidelberg, 2013

Tim Baarslag, Koen V. Hindriks, Catholijn M. Jonker, Sarit Kraus, and Raz Lin. The first automated negotiating agents competition (ANAC 2010). In Takayuki Ito, Minjie Zhang, Valentin Robu, Shaheen Fatima, and Tokuro Matsuo, editors, *New Trends in Agent-based Complex Automated Negotiations*, volume 383 of *Studies in Computational Intelligence*, pages 113–135, Berlin, Heidelberg, 2012. Springer-Verlag

Tim Baarslag, Mark J.C. Hendrikx, Koen V. Hindriks, and Catholijn M. Jonker. Measuring the performance of online opponent models in automated bilateral negotiation. In Michael Thielscher and Dongmo Zhang, editors, *AI 2012: Advances in Artificial Intelligence*, volume 7691 of *Lecture Notes in Computer Science*, pages 1–14. Springer Berlin Heidelberg, 2012

Tim Baarslag, Koen V. Hindriks, and Catholijn M. Jonker. Towards a quantitative concession-based classification method of negotiation strategies. In David Kinny, Jane Yung-jen Hsu, Guido Governatori, and Aditya K. Ghose, editors, *Agents in Principle, Agents in Practice*, volume 7047 of *Lecture Notes in Computer Science*, pages 143–158, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg

Workshop Papers

Tim Baarslag, Rafik Hadfi, Koen V. Hindriks, Takayuki Ito, and Catholijn M. Jonker. Optimal non-adaptive concession strategies with incomplete information. In *Proceedings of The Seventh International Workshop on Agent-based Complex Automated Negotiations (ACAN 2014)*, 2014

Tim Baarslag. Accepting optimally in automated negotiation with incomplete information. In *Proceedings of the 25th Benelux Conference on Artificial Intelligence*, 2013

Tim Baarslag, Koen V. Hindriks, Mark J.C. Hendrikx, Alex S.Y. Dirkzwager, and Catholijn M. Jonker. Decoupling negotiating agents to explore the space of negotiation strategies. In *Proceedings of The Fifth International Workshop on Agent-based Complex Automated Negotiations (ACAN 2012)*, 2012

Tim Baarslag, Koen V. Hindriks, Mark J.C. Hendrikx, Alex S.Y. Dirkzwager, and Catholijn M. Jonker. Decoupling negotiating agents to explore the space of negotiation strategies. In *Proceedings of the 24th Benelux Conference on Artificial Intelligence*, 2012

Tim Baarslag, Koen V. Hindriks, and Catholijn M. Jonker. Acceptance conditions in automated negotiation. In Patrick De Causmaecker, Joris Maervoet, Tommy Messelis, Katja Verbeeck, and Tim Vermeulen, editors, *Proceedings of the 23rd Benelux Conference on Artificial Intelligence*, pages 363–365, 2011

Tim Baarslag, Koen V. Hindriks, and Catholijn M. Jonker. Acceptance conditions in automated negotiation. In *Proceedings of ICT.Open 2011*, 2011

For my parents

Supervisor's Foreword

I first met Tim in 2010, when he started as a Ph.D. student in the Interactive Intelligence Group of the Faculty of Electrical Engineering, Mathematics and Computer Science at Delft University of Technology. As head of the group, and as Tim's promoter, we interacted in weekly face-to-face meetings during the writing of his thesis and at various conferences and events. In these four years, I came to know Tim as a hardworking individual who is full of initiative. Tim has shown a great talent for formulating a research agenda and for recognizing the potential use of mathematical tools. His mathematical insights enabled him to independently come up with a number of tested theories on optimal strategies for negotiation, which outperform all existing mechanisms. As a result, his Ph.D. dissertation has yielded important practical and theoretical insights into the design of automated negotiators.

The dissertation was assessed by an international committee and was unanimously awarded with honors (*cum laude*), which is conferred to less than 5 % of Ph.D. candidates of Delft University of Technology. The dissertation has also won the *2014 Victor Lesser Distinguished Dissertation Runner-up Award* in recognition of an exceptional and highly impressive Ph.D. dissertation in the area of autonomous agents and multiagent systems. I am proud to now see it appear in *Springer Theses*, recognizing its excellence and high impact on research.

During his Ph.D., Tim Baarslag has successfully cooperated with leading researchers in the field, including Prof. Nicholas R. Jennings and Prof. Sarit Kraus. A prime example is the joint publication in *artificial intelligence* about the international negotiation competition he coorganized. The journal editor, G. Lakemeyer, commented as follows on this paper: "*(…) the paper is an excellent example of the kinds of reports we would like to see published (…), providing deep insights into the state of the art of a particular kind of AI systems, from which lessons for future developments can be learned*."

All research results presented in this dissertation are published in top journals in the field; e.g., *artificial intelligence, decision support systems, autonomous agents and multiagent systems, and computational intelligence*. His work on the significance of automated negotiators provides a new perspective on the challenges in the field and resulted in a publications in leading conferences on AI, such as the *12th International Conference on Autonomous Agents and Multi-agent Systems (AAMAS 2013)* and the *21st European Conference on Artificial Intelligence (ECAI 2014)*. The success of his ideas is demonstrated throughout this book; for example, his results on negotiation learning techniques won the Best Paper Award in the *IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT)*.

To top it off, an agent that successfully applied his dissertation framework won first prize in *The 2013 International Automated Negotiating Agents Competition (ANAC)*.

Delft
December 2015

Prof. Catholijn M. Jonker
Dr. Koen V. Hindriks

Acknowledgements

I have thoroughly enjoyed my life as a Ph.D., and I could not have done it without the help and support—and the love and warmth—of the people around me.

First and foremost, I would like to thank my copromotor Koen Hindriks and my promotor Catholijn Jonker for all their suggestions and useful comments. Koen, thank you for being so dedicated and involved in my work, while at the same time giving me the freedom to pursue my own research ideas. Thank you for all the help, for having a supervision style based on trust, and for putting your faith in my proposals.

Catholijn, thank you for believing in me and for supporting me when it mattered most. Thank you for the intense discussions, for the brainstorm sessions, and for the opportunity to go to *AAMAS 2010*, just a few months into my Ph.D. It gave me a lot of inspiration, and it really kick-started my research. Thank you for the nice words and encouragements; these little things matter a lot.

Of course, I would also like to thank my colleagues of the *Interactive Intelligence Group* for making these years so much more enjoyable. Many of you have undoubtedly skipped straight to this section, but I like to think you are just saving the rest for later. In particular, I thank Hani Alers for his helpful advice and for organizing all kinds of cool stuff; Reyhan Aydogan for her cheerfulness and for being my surrogate mom in Delft; Joost Broekens for being so passionate and positive; Christian Detweiler for the music, the stage-diving, and the useless trivia; Nike Gunawan for her good mood and organizing the awesome Antwerp trip; Maaike Harbers for the "voeten op tafel" conversations about life and career; Alex Kayal for bringing a much-needed new impulse to the group, for the conversation slipups, and for making Minnesota awesome; Iris van de Kieft for sharing the experience of embarking on our Ph.D. adventure; Thomas King for his wits and for his innovative guitar techniques; Iulia Lefter for patiently answering all my questions about defense regulations; Hantao Liu for all the postdoc advice; Arman Noroozian for the math and puzzles and for his strength; Alina Pommeranz for the hallway laughs and for introducing me to Frisbee; Judith Redi for a healthy dose of realism; Dmytro Tykhonov for his wise tips and advice; Wietske Visser for her kindness and helpfulness; and Chang Wang for the history lessons and the impeccable impersonations.

Furthermore, I thank Anita Hoogmoed, Ruud de Jong, and Bart Vastenhouw for their unwavering support and cheerfulness and for always being willing to lend a hand.

I also received great support from outside the *TU Delft*. I thank my coauthors and collaborators from *The University of Southampton*, in particular the *Agents, Interaction and Complexity Group*, with special thanks to Nicholas Jennings, Enrico Gerding, Valentin Robu, and Colin Williams.

Great thanks also go to my colleagues at *Bar-Ilan University* and *Ben-Gurion University of the Negev*, especially Sarit Kraus and Raz Lin for all the collaborative and organizational effort and Kobi Gal for all the support. I thank Pınar Yolum from *Boğaziçi University* for our collaborations together, and a special thanks goes to Ivan Marsa-Maestre for helping me turn the page in Taipei.

I thank my sensei away from home, Takayuki Ito, for his help and for generously hosting me at *Nagoya Institute of Technology*. I also enjoyed collaborating with Katsuhide Fujita, now at *Tokyo University of Agriculture and Technology*. I thank Shantanu Chakraborty, Rafik Hadfi, and Shun Okuhara for making my stay at Nagoya such a blast. I am also grateful for a great and inspiring time in *Shizuoka University*, thanks to the warm welcome of Naoki Fukuta and Yoshinori Tsuruhashi.

A big thank you goes to my master students, Alex Dirkzwager and Mark Hendrikx, who I supervised with pleasure for the last two years. Without you, I would not have been able to do all the work I wanted to do. You both did an amazing job.

I also like to thank my housemates of the Jesseplaats for dropping by on all sorts of occasions, preferably uninvited, and for making sure I never came home to an empty house after a conference. I thank my band members and good friends Ivo Esseveld and Evgeny Rezunenko for making my life more musical and for occasionally showing up for band practice. I would like to thank and remember Mississippi John Hurt, John Fahey, and Jack Rose for making me want to pick up the guitar every day.

Derk van Veen and Tamara Vreeburg, thank you for being there for me through the good and the bad times. Thank you Elwin Man, for more than 15 years of beers, music, and countless chess blitz matches. I thank my trustworthy paranymphs Andreas Goetze and Hugo Looijestijn for more than a decade of yearly holidays, mathematics, and endless bar talks.

I thank my parents for everything. Thank you Ed, for giving me love for science. It makes me sad that you will not see how helpful and inspiring our talks have been for me, but you will always be there in my heart. Thank you so much Willy, for your love and encouragement, and for being such an awesome mom. I love you both.

Lastly, I thank my sweet Christina for her love, and for accepting me; for broadening my perspective on life; and for adding a welcome dash of Mediterranean spirit to my life in Delft.

Contents

About the Author

Tim Baarslag was born on May 27, 1983, in Utrecht, The Netherlands. He studied mathematics at the Faculty of Mathematics and Computer Science of Utrecht University. Tim received his propaedeutic degree in mathematics and computer science in 2012 (with honors) on the topic of the foundations of mathematics. Tim obtained his master's degree in mathematics in 2007 (with honors), with a focus on the complexity analysis of recursive algorithms. Tim obtained an additional bachelor's degree in computer science at Utrecht University in 2008 (with honors). He completed his master of science teaching mathematics in the Graduate School of Teaching and Learning, University of Amsterdam in 2009.

In 2014, Tim obtained his Ph.D. (with honors) at the Faculty of Electrical Engineering, Mathematics and Computer Science of Delft University of Technology on the topic of intelligent decision support systems for automated negotiation. During this period, Tim has twice been a visiting researcher at the Nagoya Institute of Technology of Japan.

As of June 2014, Tim is a research fellow in the Agents, Interaction and Complexity Group at the University of Southampton, where he works on devising intelligent negotiation systems that help users obtain better protection of their privacy-sensitive data.

Chapter 1
Introduction

Leo Hendrik Baekeland, born in 1863, was always at the head of his class. He graduated at sixteen and received his doctor's degree maxima cum laude when he was still only twenty-one. By 1891, he had opened an office in the US as an independent consultant and invented a type of photographic paper that could be developed under artificial light. In 1899, Leo Baekeland was still struggling with his Velox photosensitive manufacturing business. One day he received an invitation letter from George Eastman-Kodak, who had established the Eastman Kodak Company in Rochester, New York. George suggested that if Baekeland was willing to sell his Velox manufacturing company, he was welcome to visit him for a talk. During the long carriage ride up to Rochester, Baekeland planned to ask for $50,000, but kept wondering if he would be able to get even $25,000 for his manufacturing process. George Eastman invited Leo Baekeland into his office, and fortunately for Baekeland, Eastman spoke first and right away offered him one million dollars. Baekeland immediately took the offer. He could now afford to do his research in a well-equipped laboratory and went on to invent the first plastic, Bakelite.

1.1 Negotiation

Negotiation is a core activity in human society to form alliances, to reach trade agreements, and to resolve conflicts. One cannot overstate the importance of negotiation and the centrality it has taken in our everyday lives. People negotiate everywhere, in business as well as their personal lives [1], mostly without realizing they do so [2]. Negotiation not only occurs in obvious instances, such as job negotiation, politics, acquiring a house, or haggling at the marketplace. We also use it in various everyday situations, such as setting a calendar date with a friend, asking for a refund, or agreeing on a deadline.

The field of negotiation is an important topic of research in economics [3, 4], artificial intelligence [5–10], game theory [3, 5, 6, 9, 11–13], and social psychology [14]. The last two decades have seen a growing interest in the *automation* of negotiation and e-negotiation systems [6, 8, 15–17], for example in the setting of e-commerce [18–21]. This interest is fueled by the promise of automated agents

T. Baarslag, *Exploring the Strategy Space of Negotiating Agents*,
Springer Theses, DOI 10.1007/978-3-319-28243-5_1

being able to negotiate on behalf of human negotiators, and to find better outcomes than human negotiators [18, 21–25].

Negotiation agents can alleviate some of the efforts required of people during negotiations and make negotiation problems more manageable and comprehensible for negotiators [26]. The potential benefits of automation include the reduced time and negotiation costs resulting from automation [21, 27–29], the potential increase in negotiation usage since the user can avoid social confrontation [21, 30], the ability to improve the negotiation skills of the user [24, 31, 32], and the possibility of finding more interesting deals by exploration of the outcome space [21, 31]. There are also many unexpected uses of automated negotiation; for example controlling the load in an electrical grid [33], locating available parking spaces [34], playing *Civilization IV* [35], routing telephone calls [2], or Mars rovers coordinating autonomously who is better equipped for a given task [36]. Thus, success in developing an automated agent with negotiation capabilities has great advantages and implications.

1.2 An Automated Negotiator

Automated negotiation research deals with two main topics [6, 37], which both have received their fair share of attention in the field.

From a system design or mechanism design point of view, devising an effective *negotiation protocol* is the most important concern (e.g. [2, 38–40]). Negotiation protocols are the set of rules that govern the way the negotiation takes place. This covers the number of participants and the valid actions of the participants in every particular negotiation state (e.g., which messages can be sent by whom, to whom, and at what stage). It also specifies the structure of the possible agreements, and what operations are allowed to change the contents of proposed offers.

In other cases, such as in this thesis, the *agent's decision making model* is the dominant concern (e.g. [8, 41, 42]). The main focus here is on the reasoning modules and strategies that the negotiating agents employ to make their decision in order to achieve their objectives. When the protocol is such that it leaves room for strategic reasoning, the success of a self-interested agent is determined by the effectiveness of its decision making model.

In order to be successful, a negotiating agent needs to be able to perform a variety of tasks. First of all, the agent needs to be able to interact with the others in a given *negotiation setting* that defines the different parameters of the negotiation (see Fig. 1.1). During negotiation, the agent exchanges proposals with the other participants in order to reach an acceptable agreement, which is a contract that all negotiating parties agree upon. The range of contracts being negotiated over (i.e., the set of all possible negotiation outcomes) is called the *negotiation domain*. Of course, the proposals must be submitted according to certain rules and be valid according to constraints set by the negotiation protocol. Every agent has *preferences* over the negotiation domain, which define the particular *negotiation scenario*.

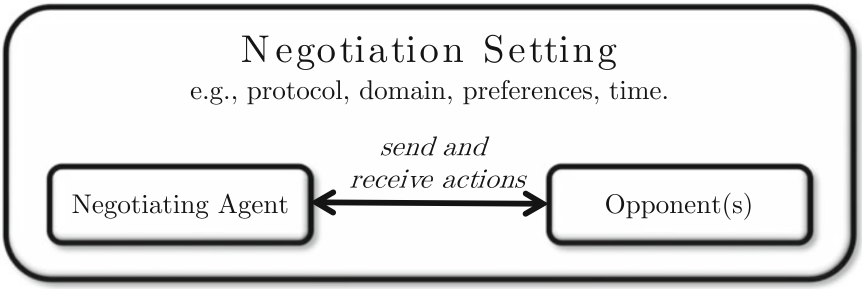

Fig. 1.1 The setting for an automated negotiator

The agent designer can select a number of *performance measures* to assess the success of a negotiating agent. The most popular way is to assign a certain *utility* to the outcomes that are reached by the agent. Other measures that the agent designer might choose are the duration of the negotiation (i.e., how fast the agent is able to reach agreements), or fairness of the outcome (i.e., whether the agreement satisfies all negotiation parties).

1.2.1 Generic Negotiation Strategies

With the constant introduction of new negotiation domains, negotiating agents may encounter different types of opponents with different characteristics. Therefore, an important research topic in automated negotiation is the design of agents that can perform well in a *variety of circumstances*. Such generic automated negotiation agents should be capable of negotiating proficiently within arbitrary negotiation scenarios, with opponents that are diverse in their behavior.

A number of automated negotiation strategies have been proposed that are designed to operate in specific and relatively simple scenarios and are often based on simplifying assumptions (e.g., [32, 39, 43–46]). A typical example of such an assumption is that the opponent strategies and preferences are known or partially known. This is generally unrealistic, as negotiators tend to avoid revealing their private information [47], because the shared information may be used to the revealer's disadvantage [48].

Examples of more general agent negotiators are increasingly available in the literature. Every year, automated negotiation agents are improving in various ways and have proven to be successful in many regards (for an exposition, see Chap. 2 and Appendix B). They all have their unique strengths and weaknesses and are based on a variety of techniques, such as game-trees [49], generic trade-off algorithms [43, 45], concession curves [50, 51], statistical analysis [52, 53], wavelet decomposition [54–56], and Gaussian process regression [57–59].

Each technique is used for various aspects of the negotiation process, such as preference learning, strategy prediction, making concessions, or choosing when to accept. However, when testing the performance of the resulting agents, varying performance measures for the negotiation outcome are used, but their inner components are not inspected. This makes it very difficult to meaningfully compare the agents, let alone their underlying techniques. As a result, we lack a reliable way to pinpoint the most effective constituents of a negotiating agent. This makes it virtually impossible to determine the reasons for an agent's success or to provide incremental improvement over existing agent designs.

To put it succinctly:

Problem We lack a fundamental approach to build comparably effective, general automated negotiators in an incremental fashion that enables us to understand how their underlying techniques influence their performance.

1.3 Bidding, Learning, and Accepting

There is a wide variety of currently existing sophisticated agent strategies and architectures, but we show in this thesis that there is some common structure to their overall design. For example, every agent decides whether the opponent's offer is acceptable, and if not, what offer should be proposed instead. In addition, when the agent decides on the counter-offer, it considers its own utility, but it usually also takes the opponent's utility into account. We elaborate on this topic in Chap. 3, but for now it suffices to say that we distinguish three distinct *components* of a negotiating agent strategy, each of which we analyze separately in this thesis:

- **Bidding strategy**. Given the current negotiation state, what are the appropriate bids to be made?
- **Opponent model**. How can we learn what the opponent wants, and how do we take this into account?
- **Acceptance strategy**. Should we accept the opponent's bid, reject it, or walk away from the negotiation altogether?

There are two major advantages of distinguishing between the different components of a negotiating agent's strategy: first, given performance measures for the individual components, it allows the study of the behavior and performance of the components in isolation. For example, it becomes possible to compare the accuracy of the opponent modeling components of a set of agents, and to pinpoint the best opponent model among them. Second, we can assemble, from already existing components, new negotiating agents in a plug and play fashion (see Fig. 1.2), e.g.: replacing the opponent model of an agent and then examining whether this makes a difference in performance. Such a procedure enables us to combine the individual components to systematically explore the space of possible negotiation strategies. Finding a good negotiating strategy then boils down to deciding *what to bid*, *how to learn*, and *when to accept*.

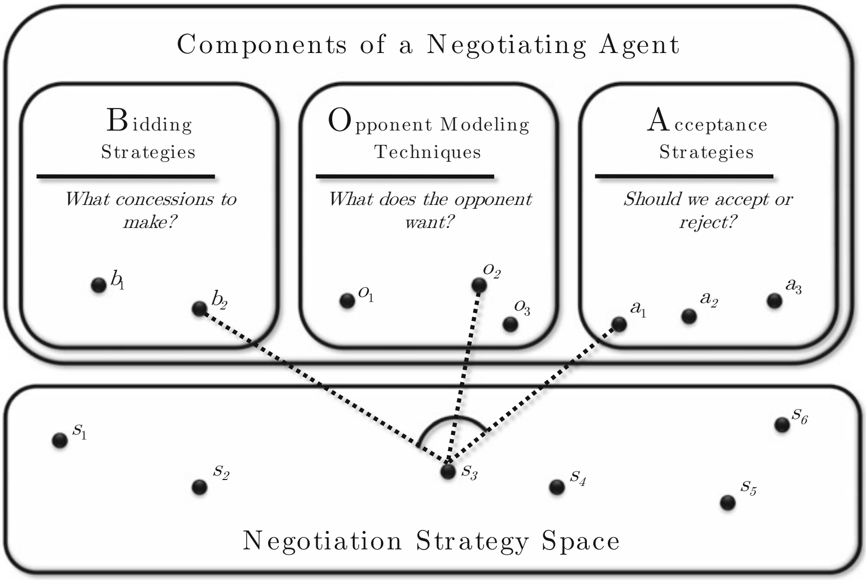

Fig. 1.2 The strategy space of automated negotiators can be explored by combining a bidding strategy with an opponent model and an acceptance strategy

Due to possible dependencies between the components, the agent should be able to combine them in a meaningful way; e.g., purposely selecting 'exploratory offers' to learn more about the opponent's preferences, or considering the opponent's future behavior when deciding whether to accept the opponent's bid. This means that in order to be successful, a negotiating agent should not only have the three components work effectively in an individual manner, but the agent also needs a powerful architecture with which to assemble the components into a negotiation strategy.

1.4 Research Questions

The advantages of a component-based approach for an automated negotiator as outlined above have motivated our concrete research aim as follows:

Thesis Aim
The central aim of this thesis is to research effective ways for a general automated negotiating strategy *to learn*, *to bid*, and *to accept* and to develop a compositional approach for evaluating and combining these components.

Note that our thesis aim consists of two separate aspects: *creating* a component-based approach, and *using* it to analyze and devise the components. That is, our aim involves both a *design* and an *analysis* point of view that together contribute to a more methodological approach for automated negotiation research. We will treat both aspects separately and formulate a set of research questions for each of them.

1.4.1 Designing a Component-Based Automated Negotiation Framework

To develop a compositional approach to evaluate and combine the components, we need to design a negotiation environment that supports negotiation analysis and that implements benchmarks for general automated negotiating strategies and their components. In particular, we need to establish an agent decision making architecture capable component-based negotiation behavior. For this, we need to understand how the bidding, learning and accepting components of a negotiating agent relate to each other and how to combine them in an effective way. Also, given the availability of state of the art negotiation strategies, an important consideration is that existing agent designs can be incorporated into our approach.

Thus, in order to achieve the design aspect of our aim, we address the following questions:

Research Questions I
Designing a component-based automated negotiation framework.

How do we create a negotiation framework that:

1. supports new agent designs and provides insight into the effectiveness of negotiation strategies;
2. facilitates evaluating and combining various negotiation strategy components;
3. enables us to decompose existing, state of the art agent designs into distinct components.

1.4.2 Analyzing the Negotiating Strategy Components

To analyze the components individually, it is necessary to formulate benchmarks and predictors for the performance of the individual components. The performance measures for the bidding strategy, opponent model, and acceptance strategy are likely to be different for each case. With performance measures for every component, we can specify solutions separately in a plug and play fashion. We will consider specific situations (and specific classes of opponents in particular) for which we can find effective solutions, and in some cases, even optimal ones.

Of course, after analyzing the components individually, we need to consider what happens if we assemble them again, and whether combining effective components also improves the overall performance. There could be strong interdependency between the components, and some components can prove to be more important to consider than others.

We formulate three additional questions regarding the analysis aspect of our research aim:

Research Questions II
Analyzing the negotiating strategy components: what to bid, how to learn, and when to accept.

1. What measures can we use to compare and predict the performance of the individual components?
2. Can we pinpoint classes of opponents against which we can find effective components? Can we formulate optimal solutions for any of the components?
3. How does the performance of the components influence the negotiator's performance as a whole, and which components are most important?

1.5 Thesis Scope

Before we describe our research method to answer our research questions, we briefly frame the scope of our work. We will elaborate extensively on our model of negotiation (and on related possibilities) in Chap. 2.

This thesis focuses on bilateral negotiations (i.e., negotiations between two agents), in which the agents exchange offers in turns. While the negotiation domain is known by both agents, the preferences of each player is private information. The agents seek to reach an agreement while aiming to satisfy their own preferences.

The heart of this thesis consists of the analysis of decision making procedures for a negotiating agent in such a setting. More specifically, if we adhere to the classification used by Lomuscio et al. [21], the focus of this thesis is as follows:

Thesis Scope
This thesis focuses on *one-to-one* negotiations with *alternating offers* on *multiple-issue* domains, using *self-interested* agents with *bounded rationality* and *incomplete information*.

1.6 Dissertation Outline

We give a detailed overview of this thesis in the paragraphs below (see also Fig. 1.3). A summary is available at the end of this section.

Quick Read Guide
For the reader in a hurry, we suggest the following quick read guide for this thesis:

1. Skip the Background chapter entirely, but read the summaries of Appendix A about GENIUS (p. 215) and Appendix B about ANAC (p. 223).
2. Read the full chapter about the BOA framework (Chap. 3, p. 53).
3. Choose one chapter for each component of the BOA framework. We recommend Chap. 5 on optimal acceptance policies (p. 91), Chap. 7 on performance and accuracy of learning methods (p. 129), and Chap. 8 on the classification of bidding strategies (p. 147).
4. Read Sect. 10.4 (p. 186) on how the BOA components fit together.
5. End with our concluding chapter (p. 195).

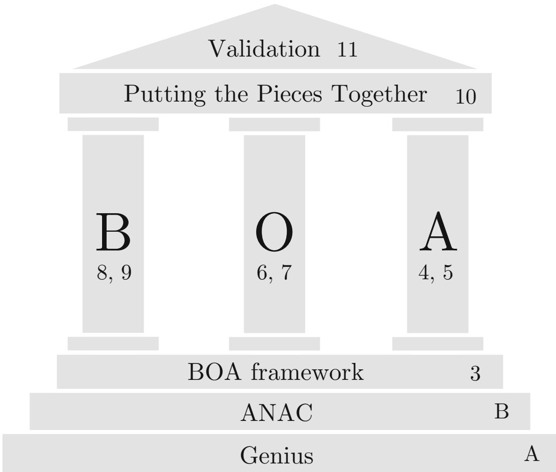

Fig. 1.3 A graphical representation of the outline of this thesis. GENIUS (Appendix A) lays the groundwork for ANAC (Appendix B), and the BOA framework (Chap. 3) builds on top of both of them. In turn, all three support the pillars to component analysis of *Bidding* (Chaps. 8 and 9), *Opponent modeling* (Chaps. 6 and 7), and *Accepting* (Chaps. 4 and 5). We put the pieces together in Chap. 10, culminating in the validation of the BOA framework (Chap. 11)

1.6.1 The Fundamentals

We start Chap. 2 by briefly discussing the background and related work in automated negotiation. We give definitions of the basic terminology used in negotiation literature and we discuss prime examples of existing automated negotiation architectures and strategies. We focus specifically on existing bidding strategies, opponent models and acceptance strategies, and on combining a set of components to explore the negotiation strategy space.

We conclude the background chapter by describing several methodologies for evaluating and comparing negotiation strategies and components. Among our discussed evaluation methods are performance and accuracy measures, agent competitions, and analytical software to assess the outcome of the negotiation. We conclude with a discussion of several evaluation methodologies of negotiation strategies, with an emphasis on *performance* and *accuracy* measures.

1.6.2 The BOA Architecture

Chapter 3 describes the BOA architecture, in which we can develop and integrate the different components of a negotiating agent into one negotiating strategy. We use the BOA architecture to explore the space of possible strategies by studying and recombining different state of the art strategy components.

The BOA architecture is integrated seamlessly into a generic negotiation environment called GENIUS (Appendix A), which is a flexible software environment that facilitates the design, evaluation and analysis of negotiation strategies. GENIUS provides full support for a diversity of different negotiation protocols, scenarios, and agents, which we amend with analytical tools and various existing agents, negotiation scenarios, and protocols from literature. The implementation of the BOA architecture offers the user the ability to create and combine newly developed components using a graphical user interface.

To explore the negotiation strategy space of the negotiation research community, we require a variety of different state of the art negotiating agents, and we need to formulate objective evaluation criteria for them. Appendix B describes the organization and insights gained from four instances of a yearly international negotiation competition (ANAC) held between 2010 and 2013 in conjunction with the International Conference on Autonomous Agents and Multiagent Systems (AAMAS). ANAC acts as an evaluation tool for negotiation strategies, and encourages the design of negotiation strategies and scenarios. Moreover, through ANAC we learn new, improved approaches to effective agent designs, which are accessible as benchmarks for the negotiation research community. We organize the competition, but we also *participate* in it, through which we foster our ties with the automated negotiation community. The agents, domains, and scores of ANAC are used in most chapters of this thesis and are discussed in detail in Appendix C–F.