MORGAN & CLAYPOOL PUBLISHERS

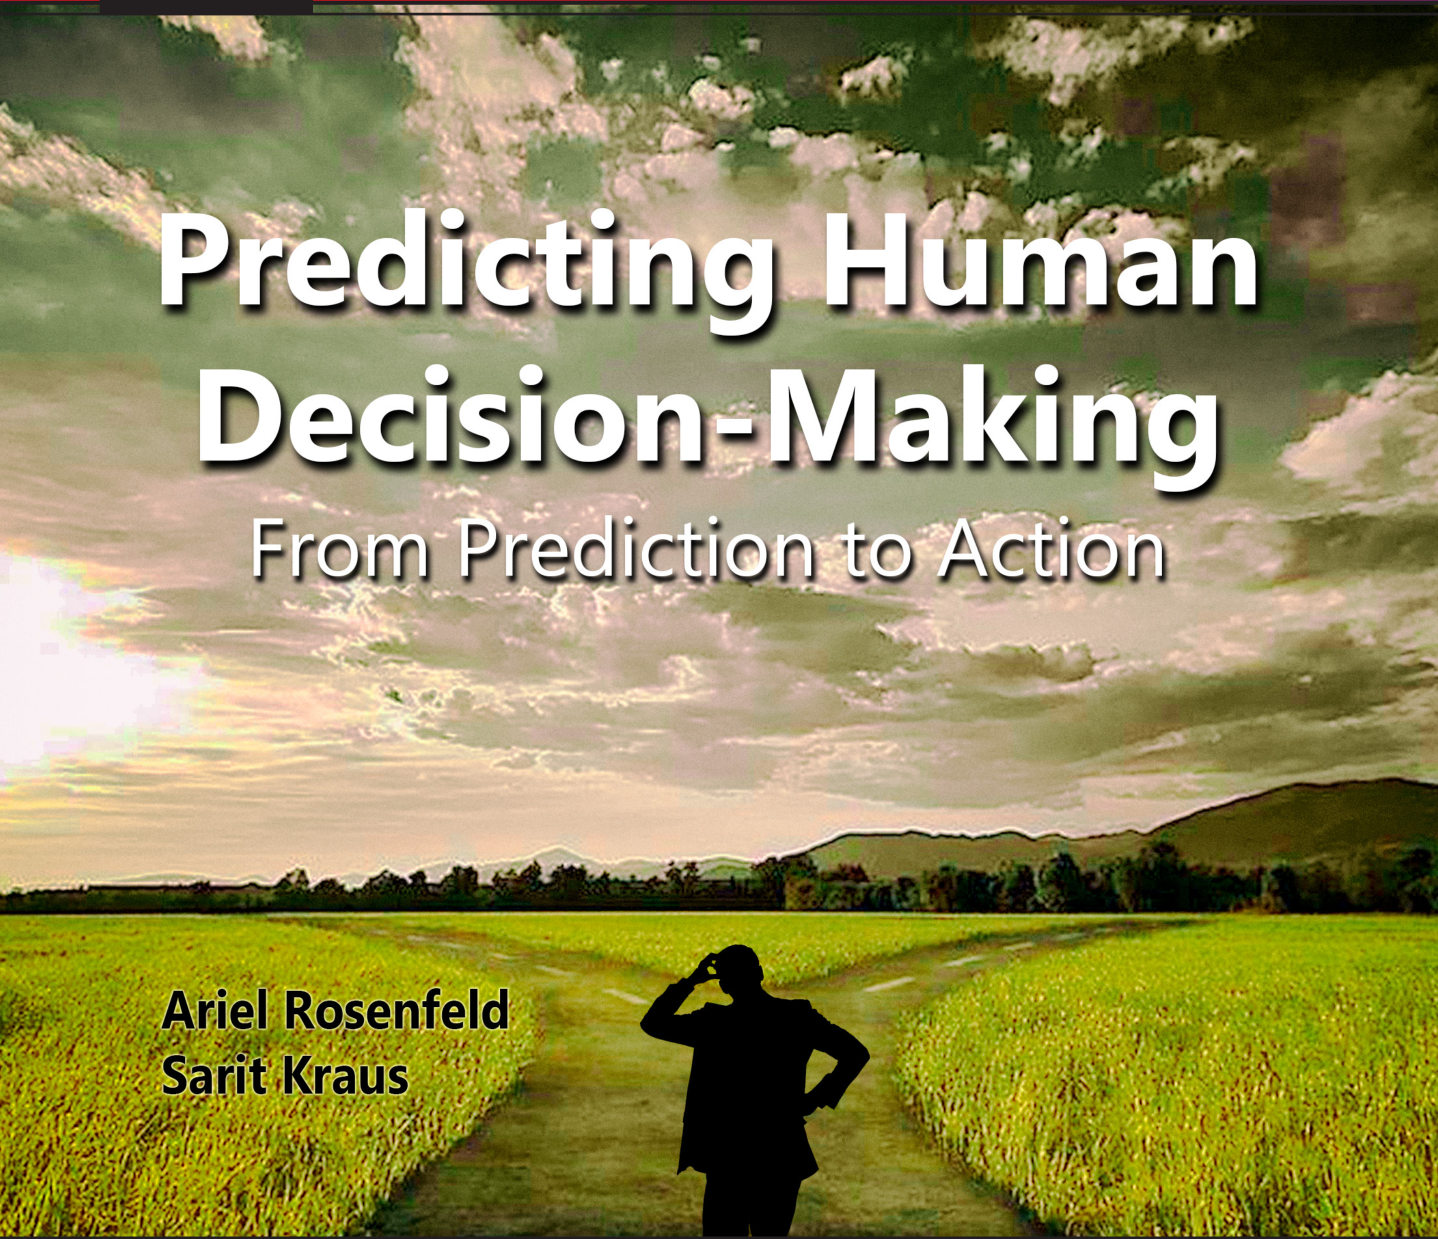

Predicting Human Decision-Making

From Prediction to Action

Ariel Rosenfeld
Sarit Kraus

SYNTHESIS LECTURES ON ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING

Ronald J. Brachman and Peter Stone, *Series Editors*

Predicting Human Decision-Making

From Prediction to Action

Synthesis Lectures on Artificial Intelligence and Machine Learning

Editors

Ronald J. Brachman, ***Jacobs Technion-Cornell Institute at Cornell Tech***
Peter Stone, ***University of Texas at Austin***

Predicting Human Decision-Making: From Prediction to Action
Ariel Rosenfeld and Sarit Kraus
2018

Game Theory for Data Science: Eliciting Truthful Information
Boi Faltings and Goran Radanovic
2017

Multi-Objective Decision Making
Diederik M. Roijers and Shimon Whiteson
2017

Lifelong Machine Learning
Zhiyuan Chen and Bing Liu
2016

Statistical Relational Artificial Intelligence: Logic, Probability, and Computation
Luc De Raedt, Kristian Kersting, Sriraam Natarajan, and David Poole
2016

Representing and Reasoning with Qualitative Preferences: Tools and Applications
Ganesh Ram Santhanam, Samik Basu, and Vasant Honavar
2016

Metric Learning
Aurélien Bellet, Amaury Habrard, and Marc Sebban
2015

Graph-Based Semi-Supervised Learning
Amarnag Subramanya and Partha Pratim Talukdar
2014

Robot Learning from Human Teachers
Sonia Chernova and Andrea L. Thomaz
2014

General Game Playing
Michael Genesereth and Michael Thielscher
2014

Judgment Aggregation: A Primer
Davide Grossi and Gabriella Pigozzi
2014

An Introduction to Constraint-Based Temporal Reasoning
Roman Barták, Robert A. Morris, and K. Brent Venable
2014

Reasoning with Probabilistic and Deterministic Graphical Models: Exact Algorithms
Rina Dechter
2013

Introduction to Intelligent Systems in Traffic and Transportation
Ana L.C. Bazzan and Franziska Klügl
2013

A Concise Introduction to Models and Methods for Automated Planning
Hector Geffner and Blai Bonet
2013

Essential Principles for Autonomous Robotics
Henry Hexmoor
2013

Case-Based Reasoning: A Concise Introduction
Beatriz López
2013

Answer Set Solving in Practice
Martin Gebser, Roland Kaminski, Benjamin Kaufmann, and Torsten Schaub
2012

Planning with Markov Decision Processes: An AI Perspective
Mausam and Andrey Kolobov
2012

Active Learning
Burr Settles
2012

Computational Aspects of Cooperative Game Theory
Georgios Chalkiadakis, Edith Elkind, and Michael Wooldridge
2011

Representations and Techniques for 3D Object Recognition and Scene Interpretation
Derek Hoiem and Silvio Savarese
2011

A Short Introduction to Preferences: Between Artificial Intelligence and Social Choice
Francesca Rossi, Kristen Brent Venable, and Toby Walsh
2011

Human Computation
Edith Law and Luis von Ahn
2011

Trading Agents
Michael P. Wellman
2011

Visual Object Recognition
Kristen Grauman and Bastian Leibe
2011

Learning with Support Vector Machines
Colin Campbell and Yiming Ying
2011

Algorithms for Reinforcement Learning
Csaba Szepesvári
2010

Data Integration: The Relational Logic Approach
Michael Genesereth
2010

Markov Logic: An Interface Layer for Artificial Intelligence
Pedro Domingos and Daniel Lowd
2009

Introduction to Semi-Supervised Learning
XiaojinZhu and Andrew B.Goldberg
2009

Action Programming Languages
Michael Thielscher
2008

Representation Discovery using Harmonic Analysis
Sridhar Mahadevan
2008

Essentials of Game Theory: A Concise Multidisciplinary Introduction
Kevin Leyton-Brown and Yoav Shoham
2008

A Concise Introduction to Multiagent Systems and Distributed Artificial Intelligence
Nikos Vlassis
2007

Intelligent Autonomous Robotics: A Robot Soccer Case Study
Peter Stone
2007

Predicting Human Decision-Making: From Prediction to Action

Ariel Rosenfeld and Sarit Kraus

www.morganclaypool.com

ISBN: 9781681732749 paperback
ISBN: 9781681732756 ebook
ISBN: 9781681732763 hardcover

DOI 10.2200/S00820ED1V01Y201712AIM036

A Publication in the Morgan & Claypool Publishers series
SYNTHESIS LECTURES ON ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING

Lecture #36
Series Editors: Ronald J. Brachman, *Jacobs Technion-Cornell Institute at Cornell Tech*
Peter Stone, *University of Texas at Austin*
Series ISSN
Print 1939-4608 Electronic 1939-4616

Predicting Human Decision-Making

From Prediction to Action

Ariel Rosenfeld
Weizmann Institute of Science, Israel

Sarit Kraus
Bar-Ilan University, Israel

SYNTHESIS LECTURES ON ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING #36

MORGAN & CLAYPOOL PUBLISHERS

ABSTRACT

Human decision-making often transcends our formal models of "rationality." Designing intelligent agents that interact proficiently with people necessitates the modeling of human behavior and the prediction of their decisions. In this book, we explore the task of automatically predicting human decision-making and its use in designing intelligent human-aware automated computer systems of varying natures—from purely conflicting interaction settings (e.g., security and games) to fully cooperative interaction settings (e.g., autonomous driving and personal robotic assistants). We explore the techniques, algorithms, and empirical methodologies for meeting the challenges that arise from the above tasks and illustrate major benefits from the use of these computational solutions in real-world application domains such as **security, negotiations, argumentative interactions, voting systems, autonomous driving, and games**. The book presents both the traditional and classical methods as well as the most recent and cutting-edge advances, providing the reader with a panorama of the challenges and solutions in predicting human decision-making.

KEYWORDS

intelligent agents, human decision-making, prediction models, human-agent interaction, decision theory, game theory, machine learning, human factors, applications

Contents

Preface xi

Acknowledgments xv

1 Introduction 1

1.1 The Premise 1
1.2 Prediction Tasks Taxonomy 3
1.3 Exercises 5

2 Utility Maximization Paradigm 7

2.1 Single Decision-Maker–Decision Theory 7
2.1.1 Decision-Making Under Certainty 8
2.1.2 Decision-Making Under Uncertainty 9
2.2 Multiple Decision-Makers–Game Theory 10
2.2.1 Normal Form Games 11
2.2.2 Extensive Form Games 14
2.3 Are People Rational? A Short Note 16
2.4 Exercises 17

3 Predicting Human Decision-Making 21

3.1 Expert-Driven Paradigm 21
3.1.1 Utility Maximization 21
3.1.2 Quantal Response 23
3.1.3 Level-*k* 24
3.1.4 Cognitive Hierarchy 26
3.1.5 Behavioral Sciences 28
3.1.6 Prospect Theory 33
3.1.7 Utilizing Expert-Driven Models 35
3.2 Data-Driven Paradigm 36
3.2.1 Machine Learning: A Human Prediction Perspective 36
3.2.2 Deep Learning—The Great Redeemer? 38
3.2.3 Data—The Great Barrier? 40

3.2.4 Additional Aspects in Data Collection 45
3.2.5 The Data Frontier 46
3.2.6 Imbalanced Datasets 47
3.2.7 Levels of Specialization: Who and What to Model 48
3.2.8 Transfer Learning 51
3.3 Hybrid Approach 54
3.3.1 Expert-Driven Features in Machine Learning 54
3.3.2 Additional Techniques For Combining Expert-Driven and Data-Driven Models 55
3.4 Exercises 56

4 From Human Prediction to Intelligent Agents 61
4.1 Prediction Models in Agent Design 61
4.2 Security Games 63
4.3 Negotiations 67
4.4 Argumentation 71
4.5 Voting 74
4.6 Automotive Industry 77
4.7 Games That People Play 79
4.8 Exercises 82

5 Which Model Should I Use? 87
5.1 Is This a Good Prediction Model? 87
5.2 The Predicting Human Decision-making (PHD) Flow Graph 88
5.3 Ethical Considerations 90
5.4 Exercises 93

6 Concluding Remarks 95

Bibliography 97

Authors' Biographies 129

Index 131

Preface

Human decision-making often transcends our formal models of "rationality." Designing intelligent agents that interact proficiently with people necessitates the modeling of human behavior and the prediction of their decisions. In this book, we explore the task of automatically predicting human decision-making and its use in designing intelligent human-aware automated computer systems of varying natures—from purely conflicting interaction settings (e.g., security and games) to fully cooperative interaction settings (e.g., autonomous driving and personal robotic assistants). We explore the techniques, algorithms, and empirical methodologies for meeting the challenges that arise from the above tasks and illustrate major benefits from the use of these computational solutions in real world application domains such as **security, negotiations, argumentative interactions, voting systems, autonomous driving and games**. The book presents both the traditional and classical methods as well as the most recent and cutting-edge advances, providing the reader with a panorama of the challenges and solutions in predicting human decision-making.

The book is written with two intentions in mind. First and foremost, **it is intended for students, researchers, and the general population who seek to broaden their knowledge** and become familiar with the task of predicting human decision-making and the development of intelligent agents based on such predictions. Second, **it is intended to serve as a textbook**. This textbook, which includes more than 60 exercises (including programming exercises) and is accompanied by PowerPoint presentations, could fit as part of an introductory or advanced AI course or stand as a short course/seminar/workshop in its own right. Visit our website for further information: `https://sites.google.com/view/predicting-human-dm`.

This book follows a tutorial given by both authors titled "Predicting Human Decision-Making: Tools of the Trade" given at the thirty-first AAAI Conference on Artificial Intelligence on February 5, 2017, in San Francisco, California. An extended version of the tutorial was given by the first author at the 19th European Agent Systems Summer School (EASSS) from August 8–9, 2017, in Gdańsk, Poland.

PREREQUISITES

In the course of this book, we occasionally refer to techniques and concepts common in decision-making, machine learning, game theory, and artificial intelligence. The *basis* for these concepts is covered as part of the book and does not require (substantial) further reading. Some concepts are only *mentioned* in the book and those are followed by appropriate references in order to keep the book's focus intact.

All the same, basic familiarity with decision-making, machine learning, game theory, and artificial intelligence concepts is encouraged, but not mandatory.

The book is filled with real-world examples, varying across different domains, accompanying, illustrating, and explaining the introduced notions. We hope that this variety will allow readers of different backgrounds to find the book engaging.

ROADMAP

The book is organized as follows.

- **Chapter 1** sets the scene by discussing the prediction of human decision-making from several perspectives. It further presents a classification of human decision-making prediction tasks.
- **Chapter 2** presents the basic notions of *rational decision-making*. The utility maximization paradigm is discussed, and the basic notions of decision theory and game theory are defined and exemplified. The chapter ends with a discussion about people's (ir)rationality.
- **Chapter 3** discusses the common techniques deployed for predicting human decision-making. We review the three major approaches: the *expert-driven, data-driven, and hybrid approaches*, as well as their advantages and limitations. This chapter further illustrates the effectiveness of the different approaches in real-world applications.
- **Chapter 4** takes a more practical approach, investigating the challenges of utilizing prediction models in the real world. *Six* highly popular domains are discussed and compared (*Security Games, Negotiations, Argumentation, Games, Autonomous Cars, and Voting*). Each domain is surveyed briefly, followed by a comprehensive discussion of the prediction of human decision-making in the domain at hand. Domain specific and general insights such as applicable assumptions and beneficially deployed techniques are highlighted.
- **Chapter 5** elevates the level of insights derived from Chapters 2, 3, and 4 and provides a thorough discussion of what makes a good prediction model, followed by the **Predicting Human Decision-making (PHD) flow graph**, aimed at providing directions for tackling new prediction tasks and environments. The chapter ends with a discussion of the ethical considerations surrounding the topic.
- **Chapter 6** concludes the book and provides future directions for the field.

TEACHING WITH THIS BOOK

We have written this book with clear intention that it be used as (part of) a course text. The book is primarily intended for advanced undergraduates or beginning-to-middle graduate students in computer science, engineering, or management.

Although different instructors may wish to spend more or less time covering the different chapters of the book, we provide the following time estimations along with practical advice.

- The introduction (**Chapter 1**) sets the scene and introduces terminology used for the remainder of the book, do not skip it. *It should take about 1–2 hours.*

- If students have taken a *Game Theory* course, or come with a strong background in economics, **Chapter 2** may be significantly reduced. Otherwise, *we would recommend spending 2–4 hours on the topic.* Note that the discussion over human rationality can easily be turned into an intriguing class discussion.

- **Chapter 3** is by far the longest chapter of the book. It consists of "two-and-a-half" main sections introducing the expert-driven paradigm (*1–2 hours*), the data-driven paradigm (*1–3 hours*), and the hybrid approach (*half an hour*). The time estimations here relate to computer science students as discussed above. Students from other backgrounds may require significantly more time to understand the data-driven and hybrid approaches. Note that if students have not encountered any machine learning models in the past the instructor should devote at least half an hour to illustrate a few basic learning algorithms in practice (e.g., using Weka [331]).

- Six domains are discussed in **Chapter 4**. From our experience, the last two (*autonomous driving and games*) are the most popular among students. *Allow between 1–4 hours to review some or all of the application domains*. Note that reading and even presenting the left-out domains may be given to students as an exercise.

- **Chapters 5** and **6** summarize the book and revisit the introduced techniques and many of book's notions and concepts. It could easily be turned into a class discussion, for the most part. *Allow for at least 1–2 hours*. The topic of "ethical considerations" may be expanded or reduced per the students' interest or time constraints.

Overall, if one would want to use the book as a *self-contained course*, we would expect such a course to take between 12–18 hours.

Another option, perhaps more relevant for most computer science instructors, is to integrate this book as part of a basic or advanced AI course. Given the students' (expected) prior knowledge on game theory and machine learning, the book can be stripped to its bare essentials of 4–6 hours of teaching. Our recommendation is as follows.

1. Quick overview of the Introduction (*half an hour*).

2. Fundamental notions of Chapter 3 (*2–3 hours*).

3. A few examples from Chapter 4 (*1–1.5 hours*).

4. The PHD graph-flow and conclusions (*1 hour*).

Note that each chapter ends with exercises. Consider using those as a basis for course homework assignments or class discussions.

Additional materials such as PowerPoint presentations and additional references are available at the book's webpage: `https://sites.google.com/view/predicting-human-dm`.

READING THIS BOOK

The book deliberately sits on the fence between artificial intelligence and other fields. Thus, readers from varying disciplines such as computer science, engineering, cognition, and social sciences can find the book, for the most part, accessible and engaging.

We would recommend readers with a strong decision-theoretic and game-theoretic background to lightly go over **Chapter 2** or skip it altogether. Readers with a strong **background** in machine learning are still encouraged to read **Section 3.2** (which focuses on machine learning), as it provides emphasis on the human perspectives on machine learning, commonly not highlighted in machine learning courses.

Ariel Rosenfeld and Sarit Kraus
January 2018

Acknowledgments

To our families, students, collaborators, and friends.

Predicting your decisions, sometimes against all better judgment, has always been hard…but hey, at least we got to publish a book about it.

Ariel Rosenfeld and Sarit Kraus
January 2018

CHAPTER 1

Introduction

> "Trying to understand the behavior of some people is like trying to smell the color 9"
>
> *Unknown*

People make millions of decisions everyday. Just now—**you**, the reader, decided to start reading this book (and we are grateful for that). People's ability to predict each other's decisions in a *fast and accurate way* constitutes an imperative factor in what we consider to be **intelligent behavior**. This ability is part of what enables us, as a species, to effectively interact with, cooperate, and influence one another on a daily basis. While the canonical question of how people *actually* make decisions is of great interest to mankind, in the scope of this book we seek to address two more modest questions.

1. *How can a **machine** predict human decision-making?*
2. *How can a **machine** leverage the prediction of human decision-making in order to perform in an **intelligent** manner?*

1.1 THE PREMISE

*Would you consider an autonomous car that is incapable of predicting what people (pedestrians, drivers, or its own passengers) are about to do next **intelligent** or even **safe**? Most probably not.*

From an *Artificial Intelligence* (AI) perspective, we are interested in endowing our machines and software with so-called "intelligence." While the question of what constitutes as "intelligence" remains debatable, it usually relies on human-associated qualities such as learning, reasoning and adaptation to changes. One of these qualities is our ability to predict (rather successfully in many cases) the behavior and decisions of others, an ability which we humans share with other primates [190]. In his seminal paper [317], *Alan Turing* (1912–1954) proposed to replace the philosophical question "Can machines think?" with an operational challenge of constructing an automated agent that is able to carry on a dialogue with a person well enough to be indistinguishable from a person. We know this challenge today as the *Turing Test*. *Barbara Grosz* [127] hypothesizes that given the transformations in computer use as well as the significant advances in computer science and the cognitive and brain sciences, Turing might pose a slightly different challenge: Can an automated agent team member behave, in the long term and in uncertain, dynamic environments, in such a way that people on the team will not notice that it is not human? Thus, if an automated agent is embodied in an environment shared with people, it is self evident that the prediction of human decision-making is crucial for an automated agent to be considered intelligent.

An *automated agent* is any automated tool that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators [278, Chapter 2]. These automated agents can be either physical (e.g., robots) or virtual (e.g., software). In this relatively broad definition one can find both highly complex algorithmic trading agents that operate in high-risk fast-paced markets as well as a simplistic automated door operation agent that responds to possible movement in the door's vicinity, deciding whether or not to open the door.[1]

As technology progresses, we find ourselves working with automated agents with increased frequency, from intelligent virtual assistants like *Siri*, *Cortana*, *Google Assistant*, and *Alexa* to make our lives more convenient [126], to doctors collaborating with *IBM's Watson* to make better diagnostic and treatment decisions [158]. Consider a *personal robotic assistant* in a home environment [31]. Such a robot would help with (almost) everything—from cooking and cleaning to health care and companionship. In a sense, these robots should be analogous to a sharp assistant, friend, or relative that knows the user's beliefs and interests and can proficiently interact with her. Consequently, predicting "what would a/this human do in a given context?" will enable a personal robotic assistant to figure out for itself, *with minimal to no human intervention*, what it is it needs to do, be it by mimicking, reacting to, reasoning with, or improving upon the predicted human decision(s).

Human decision-making is a highly complex *cognitive process* resulting in the selection of a belief or a course of action among several alternative possibilities. *Understanding and predicting* human decision-making are chief concerns in multiple fields, from social sciences such as psychology and economy through neurobiology and cognitive science. Unfortunately, similar to other cognitive processes such as creativity and emotions, only a small portion of the multitude of factors which effect human decision-making is currently understood by scientists. From what scientists have already established, we know that human decision-making is influenced by a large set of factors that vary across different individuals and groups of individuals [51], including past experience [167], decision complexity [99], emotions [29], and many cognitive biases.[2] The interested reader may refer to Dietrich [80] and Baumeister [28] for concise *psychological* surveys of human decision-making theories. Researchers are also currently just beginning to decipher what exactly happens in our brains when we make decisions [204, 264]. This apparent gap in the scientific understanding of human decision-making makes the prediction of human decisions very complex. Specifically, relying solely on past evidence and hypotheses from different fields may prove to be insufficient and highly inaccurate. We believe that the investigation of human decision-making through an AI perspective can shed a much desired light on the less-explored, yet highly applicable, facets of this challenge.

[1] We refer the interested reader to Franklin and Graesser [102] for a thorough discussion of what turns "simple software" into an automated agent.

[2] A comprehensive list is available at http://rationalwiki.org/wiki/List_of_cognitive_biases

The challenge depicted here is not an easy one to say the least. Capturing the subtle details inherent in human decision-making, generalizing findings across people, settings, and domains, and acting in an intelligent manner based on a given prediction have proven to be highly complex issues [319]. Nevertheless, the enormous potential of successful development and deployment of human-aware agents, capable of predicting human decision-making, can bring about a much desired leap in the way agents are designed and deployed in the real world.

Designing intelligent agents, capable of predicting human decisions and making intelligent decisions appropriately, is what this book is all about.

Throughout the book we decompose the above challenge into two separate, yet extremely related, tasks:

1. **Prediction** of human decision-making and behavior.

2. **Acting** appropriately in the environment based on the prediction.

Both tasks, as well as the interrelated challenges and the integration of solutions thereof, are highlighted throughout the book.

1.2 PREDICTION TASKS TAXONOMY

In order to make sense out of the multitude of agents and domains that (may) use the predicting of human decision-making, we offer the following classification. Each prediction task can be classified according to the different aspects or dimensions underlying it and the agent's use of the prediction model. We highlight the prominent classification criteria:

Human-Agent Interaction Setting

It is common to consider three human-agent interaction settings based on the alignment (or lack thereof) between the agent's goals and its user(s).

1. *Cooperative Settings*. Agents that interact with people as teammates or represent people on different platforms usually share the same goal(s) with their users. If an agent's goal is completely aligned with the human's goal then the agent is considered to be a *cooperative agent*. *Personal robotic assistants*, such as the ones described before, belong to this category.

2. *Adversarial Settings*. Many of AI's greatest achievements have focused on settings where the agent's success results in worse outcomes for the human. These settings include competitive games such as Poker [222], Chess [57], or Go [295]. If an agent's goal is directly in contrast to that of the human then the agent is said to be an *adversarial agent*. Carmel and Markovitch [58] were the first to show that an agent can significantly improve its performance by adequately modeling its opponent(s).

3. *Partially Conflicting Settings*. Any agent that does not fit the fully cooperative or fully adversarial cases is considered to have partially conflicting interests to those of the human (or *partially conflicting interests agent*). In this category one may find agents that are designed to influence people to adopt healthier lifestyle choices. The agent cannot merely consider its own benefit (e.g., minimizing the user's calorie consumption) but also needs to consider the user's preferences and needs in order to build a successful relationship with the user (e.g., [105]).

Platform:

Embodied agents are agents whose execution is paired with a physical body (e.g., a robot). Conversely, *software agents*, such as a chatbot mounted on a smartphone (e.g., Siri), are nor paired with a physical body. The platform on which an agent is mounted can have a vast impact on its perception, actions and human interaction capabilities. In turn, these may have a significant impact on the type and quality of prediction the agent can perform. For instance, for embodied agents, resources are usually limited. This poses a significant challenge when implementing a sophisticated prediction and intelligent decision-making mechanisms.

Activity:

Agent are situated in an environment (physical or virtual). In that environment an agent may be either an **active actor** or **observer**, depending on its capabilities to influence. For example, an observer agent in an economical setting may record and analyze competitors' prices, watch stock manipulation by insider trading and rumors, etc. The agent does not take an active action to affect its environment other than logging its findings. Conversely, active actors are designed to take active actions on their environment which are aimed at influencing it and/or other agents situated in it. For example, an economical agent as discussed above can be amended with trading capabilities, making it an active actor in the market.

Prediction Circle:

Prediction tasks can be broadly classified into three types based on their agent's interaction needs in practice.

1. **"Humanless" circle.** In many settings the prediction of human decision-making is used by the agent despite having no human in its environment. For instance, a Super Mario playing agent may learn from human demonstrations and deploy the learned policy in the fully automated environment of the Super Mario game [178].

2. **Single human circle.** If an agent seeks to predict the decision of a single human which is situated in its environment, such as the case in the classical games of Chess and Go, the prediction task falls into this category. Personal assistant agents such as Siri fall into this category as well.

3. **Multi-human circle** Perhaps the most complex prediction settings involve multiple people. Consider an agent engaging in an election setting with multiple human voters. Pre-

dicting what each voter would do and acting accordingly may be highly different than predicting what a single voter would do and voting on her behalf (acting as a proxy).

Fixed Strategy vs. Learning:

Some prediction tasks may require learning and adaptation to change, especially when people are involved. Predicting what video a person would want to watch next on Youtube[3] naturally requires the agent to adapt to changing preferences and trends. On the other hand, many of the most successfully deployed agents rely on a fixed prediction policy. For instance, the ARMOR, currently developed for intelligently randomizing resource allocation at the Los-Angeles international airport (LAX) [251], predicts the likelihood of attackers to target different assets using a fixed prediction policy.

Prediction tasks and agents may be further classified with respect to additional aspects and dimensions. In the scope of this book, the above categorization would suffice. The articulation of additional classification dimensions is left as Exercise 1.2.

1.3 EXERCISES

1.1. (Level 1) Provide other examples of intelligent agents (or real-world opportunities) where the prediction of human decision-making is (can be) deployed. Characterize the agents using the above classifications.

1.2. (Level 1) What other dimensions divide human decision-making prediction tasks? Provide an additional dimension by which prediction tasks can be classified along with real-world agents that can be separated based on the proposed dimension.

1.3. (Level 2) What practical insights to the design of intelligent agents that interact with people can you derive from Alan Turing's claim that "The processes of inference used by the machine need not be such as would satisfy the most exacting logicians" [317, page 457].

1.4. (Level 1) Define the following terms: *Decision-maker*, *Intelligent decision-maker*, *Automated decision-maker*, *Decision-making*, and *Prediction model*.

1.5. (Level 2) What are the major differences between the challenge of predicting human decision-making and plan, activity, and intent recognition research [303]?

[3] www.youtube.com

CHAPTER 2

Utility Maximization Paradigm

That's your game theory? Rock Paper Scissors with statistics?

Peter Watts, Blindsight

The study of decision-making, human or automated, is an interdisciplinary effort, studied by mathematicians, computer scientists, economists, statisticians, psychologists, biologists, political and social scientists, philosophers and others [138, Chapter 1.2]. Decision-making is centered around a **decision-maker**, an agent, human, or otherwise (e.g., automated), who selects a choice from available options. A decision-maker will be called *intelligent* if there is an underlying reasoning mechanism for its choices.

Economists typically assume that an agent's behavior is motivated primarily by material incentives, and that decisions are governed mainly by self-interest and rationality. In this context, **rationality** means that decision-makers use all available information in a logical and systematic way, so as to make the best choices they can given the alternatives at hand and the objective to be reached [171].

Utility theory provides a well-established starting point for modeling and analyzing decisions. It also implies that decisions are made in a forward-looking way, by fully taking into account future consequences of current decisions. In other words, so-called extrinsic incentives are assumed to shape economic behavior. **Utility** is a measure representing the satisfaction experienced by an agent from a service, good or state. Though a utility cannot always be directly measured (how satisfied am I from a slice of pizza?), it is reasonable to assume that a rational agent would want to maximize its obtained utility.

It turns out that the notion of utility is a very powerful tool for representing and analyzing an agent's decision-making. We will differ between two environmental settings: *single decision-maker* and *multiple decision-makers*.

2.1 SINGLE DECISION-MAKER–DECISION THEORY

A **single decision-maker environment** is represented as follows:

Let Σ be a finite set of possible choices for the decision-maker and X be a finite set of possible outcomes $(x_1, x_2, \ldots, x_N)$. We assume that the decision-maker has a **preference relation** $\preceq$ over X where $x_i \preceq x_j$ is interpreted as "x_i is preferred at least as x_j". A preference relation needs to fulfill the following properties.

1. **Reflexive**: $\forall x_i . x_i \preceq x_i$.
2. **Complete:** $\forall x_i, x_j . \ x_i \preceq x_j \vee x_j \preceq x_i$.
3. **Transitive:** $\forall x_i, x_j, x_k . \ x_i \preceq x_j \wedge x_j \preceq x_k \Rightarrow x_i \preceq x_k$.

A **utility function** $u : X \mapsto \mathbb{R}$ is a real-valued function defined over the relevant set X of outcomes.

Lemma 2.1 *Every preference relation $\preceq$ over a finite or countable set X can be **represented** as a utility function u such that $x_i \preceq x_j \iff u(x_i) \leq u(x_j)$.*

The proof of Lemma 2.1 is left as Exercise 2.1 at the end of this chapter.

Note that Lemma 2.1 refers to finite or countable outcome sets. Is it hopeless to expect the theory to work with uncountable sets? Not at all. We need the following definition to complete our transformation from preferences to utilities.

1. The **upper contor set** of $x \in X$ is defined to be $\preceq (x) = \{y \in X : y \preceq x\}$.
2. The **lower contor set** of $x \in X$ is defined to be $\succeq (x) = \{y \in X : x \preceq y\}$.
3. A preference relation $\preceq$ over X is said to be **continuous** if $\preceq (x)$ and $\succeq (x)$ are closed sets.

Perhaps the main result of this section is due to *Gerard Debreu* (1921–2004). Lemma 2.2 is also known as **Debreu's Representation Theorem**.

Lemma 2.2 *Suppose that $X \subset \mathbb{R}^n$. A preference relation $\preceq$ is reflexive, complete, transitive and continuous if and only if there exists a utility function $u : X \mapsto \mathbb{R}$ that represents it.*

The proof for Lemma 2.2 is left as Exercise 2.3 at the end of this chapter.

As a result, in Lemmas 2.1 and 2.2, we know that when one's preference relation is reflexive, complete, transitive, and either continuous or over a countable set of outcomes, her preference can be entirely described by a utility function. From this point onward, we will use real-valued utilities instead of preference relations, as those will allow us to use numeric techniques.

2.1.1 DECISION-MAKING UNDER CERTAINTY

In **decision-making under certainty**, we assume an **outcome function** $g : \Sigma \rightarrow X$ exists, mapping each possible choice to the *known* outcome of that decision. Namely, in settings where the decision-maker knows exactly what the consequence of each choice will be (g), the **rationally optimal decision** would be:

$$\sigma^* \in \arg\max_{\sigma \in \Sigma} u(g(\sigma))$$

where u is a utility function representing the decision-maker's preferences.

2.1.2 DECISION-MAKING UNDER UNCERTAINTY

In **decision-making under uncertainty**, for every possible choice, there are multiple possible consequences. Therefore, each possible choice introduces a probability distribution over X. Formally, let $a \in \Sigma$ be a possible choice. By selecting a, an outcome from X is sampled according to the probability distribution P_a associated with a. We denote $P_a(i)$ as the probability of outcome x_i given action a. Let P be the set of all probability distributions—$P = \{P_a | a \in \Sigma\}$.

A utility function as defined in Lemma 2.1 captures the link between *preference over outcomes* and utilities. However, in decision-making under uncertainty *preference over uncertain outcomes* (which are usually represented as lotteries) is needed. We would need a theory that constructs a decision-maker's preferences on the lotteries from his preferences on the outcomes. The most well-known such theory is the **Von Neumann and Morgenstern Representation Theorem** (Theorem 2.3) [320]. Before introducing it, we define two (rather technical) axioms:

Independence Axiom: for any $P_a, P_b, P_c \in P$, and any $\alpha \in (0, 1]$.

$$\alpha P_a + (1-\alpha) P_c \preceq \alpha P_b + (1-\alpha) P_c \iff P_a \preceq P_b.$$

Continuity Axiom: for any $P_a, P_b, P_c \in P$.

$$P_a \preceq P_b \Rightarrow \exists \alpha, \beta \in (0,1) \cdot \alpha P_a + (1-\alpha) P_c \preceq P_b \preceq \beta P_a + (1-\beta) P_c.$$

Theorem 2.3 Von Neumann and Morgenstern's Representation Theorem [320].
*A preference relation $\preceq$ on P which satisfies the **Independence** and **Continuity** axioms can be represented by a utility function u such that*

$$p_a \preceq p_b \iff U(P_a) = \sum_{x_i \in X} u(x_i) p_a(x_i) \le \sum_{x_i \in X} u(x_i) p_b(x_i) = U(P_b).$$

Given Theorem 2.3, decision-making is reduced to optimization, where a rational choice would follow the **expected-utility maximization** approach:

$$\sigma^* \in \arg\max_{\sigma \in \Sigma} EU(\sigma) = \sum_{x_i \in X} u(x_i) p_\sigma(x_i).$$

Intuitively, the standard rational decision-making theory (founded in von-Neumann and Morgenstern's work) follows *expected-utility maximization*, which means that choice σ **should be** preferred to choice σ' if and only if the expected utility from selecting σ is higher than the expected utility from selecting σ'.

Von Neumann and Morgenstern's theorem is a very important result for measuring the strength of a rational agent's preferences over sure options. Some of the most celebrated results in rational decision theory address the associated *challenges*, namely in showing what conditions on preferences suffice for the existence of a pair of utility and probability functions relative to which the agent can be represented as maximizing expected utility. We refer the interested reader to Gilboa [120] for a comprehensive discussion over the theoretical, philosophical, and mathematical properties of decision-making under uncertainty.

To illustrate the presented ideas consider Example 2.4.

Example 2.4 Let us consider a *personal robotic assistant*, as described in Chapter 1, which is delegated by its user with the responsibility of ordering Chinese food. There are two restaurants from which this order can be made. If no uncertainty exists, the robot should place an order to the restaurant which is more preferred by its user. However, if outcomes are stochastic, for example, there is some distribution over the possible outcomes (e.g, "good food", "bad food") for each restaurant, then a prediction over the user's preferred decision could rely on a (subjective) expected utility approach.

Thus far, we implicitly assumed that if uncertainty is present in the environment it is induced by so-called "nature." Namely, a stochastic process which does not reason about the world and does not strive to achieve any goal (i.e., it does not seek to maximize utility). This "nature" process simply samples an outcome given the decision-maker's choice. For instance, nature "samples" the weather for the afternoon (i.e., rain or not) unconditionally on the decision-maker's decision in the morning (i.e., taking an umbrella or not).[1] However, "nature" can also model the behavior of *other agents* who occupy the environment in many settings. If an agent assumes/knows that the other agents in the environment *do not reason about its actions at all*, they may be modeled as part of "nature." Namely, not every agent in the environment should be considered as a decision-maker in order for a specific agent to make a rational decision. Specifically, a *single decision-maker environment* does **not** mean a single *agent* is suited in the environment. If other agents within the environment are expected to perform reasoning on others, then one may resort to modeling the environment as a **multiple decision-makers environment** as described below.

Note that when it comes to making **real decisions**, objective probabilities as assumed above may not be available. Instead, **Subjective Expected Utility** (SEU), as developed by Savage [283], Jeffery [161] and others, was developed to model individual differences between decision-makers. SEU is built on the premise that decision-makers associate beliefs evaluations for each outcome.

2.2 MULTIPLE DECISION-MAKERS–GAME THEORY

While decision theory is concerned with the choices of individual agents, **game theory** is concerned with interactions of agents whose decisions affect each other. If all players have same

[1] Contrary to the notion of Murphy's Laws.

preferences, then the problem is coordination: ensuring that all players "pull in the same direction."

Game theory is a mathematical theory of interaction between self-interested agents, which are usually referred to as **players**. In this book we only consider **non-cooperative** games, where players are in competition and cooperative behavior is virtually impossible due to the absence of external means to enforce such cooperative behavior (e.g., contracts). We will discuss two popular game forms: **normal form** (also know as strategic form) [320] and **extensive form** games [239].

2.2.1 NORMAL FORM GAMES

Let $N = \{1, \ldots, n\}$ be a set of players. Each player i has a finite set of possible **pure strategies** Σ_i that he can deploy in the game. A strategy $\sigma_i \in \Sigma_i$ provides a complete definition of how a player will play a game given any situation the player may face. In normal form games, the players choose their strategies simultaneously; each player must choose a strategy σ_i without knowing what strategies the other players have chosen.

The combination of players' strategies results in a **strategy profile**:

$$\vec{\sigma} = \langle \sigma_1, \ldots, \sigma_n \rangle \in \Sigma_1 \times \cdots \times \Sigma_n.$$

Let X be a finite set of possible outcomes for the game $(x_1, x_2, \ldots, x_N)$. An **outcome function** $g : \Sigma_1 \times \ldots \times \Sigma_n \to X$ maps each strategy profile to an outcome. We assume that each player can be represented using a utility function as defined above for the single decision-maker setting. For annotation simplicity, we will drop the outcome function g and use utility functions of the form:

$$u_i : \Sigma_1 \times \ldots \times \Sigma_n \to \mathbb{R}.$$

It is customary to represent normal form games as a **payoff matrix**. A payoff matrix summarizes the pure strategies available to each player and the associated utilities for each player resulting from each possible strategy profile. Namely, each cell in the payoff matrix represents one strategy profile and consists of all players' utilities from the corresponding strategy profile.

Consider the following classical example, the **Prisoners' Dilemma**:

Example 2.5 The classic Prisoners' Dilemma is defined as follows.

- A set of players $N = \{1, 2\}$, where 1 stands for "Player 1" and 2 stands for "Player 2."
- For each $i \in N$, a set of strategies $\Sigma_i = \{C, D\}$, where C stands for "Cooperate" and D stands for "Defect."

- Utility functions u_1 for player 1 and u_2 for player 2 are expressed as follows.

  For Player 1:
  $u_1(D, C) = 3 > u_1(C, C) = 2 > u_1(D, D) = 1 > u_1(C, D) = 0$

  For Player 2:
  $u_2(C, D) = 3 > u_2(C, C) = 2 > u_2(D, D) = 1 > u_2(D, C) = 0$

- The game can be represented as the following *payoff matrix*:

		Player 2	
		C	D
Player 1	C	(2, 2)	(0, 3)
	D	(3, 0)	(1, 1)

In order for player i to find the best action for her (recall that she strives to maximize her expected utility), reasoning about what other player(s) will do is usually essential. However, unlike single decision-maker environments, other players, in turn, perform a similar reasoning process as well.

Recall that the utility a player receives as a result of a game depends on the combination of strategies that all players choose—the strategy profile. Therefore, given that all other players act **rationally**—which in our context means maximizing their expected utility—different **solution concepts** can be used to analytically derive the optimal strategy for a single player. In this chapter, we will discuss the two most prominent solution concepts: **dominant strategies** and **Nash equilibrium**.

For convenience, we denote the strategy profile obtained by replacing the i^{th} component of strategy profile $\vec{\sigma}$ with σ_i' as

$$(\vec{\sigma}_{-i}, \sigma_i') = \langle \sigma_1, \ldots, \sigma_i', \ldots, \sigma_n \rangle.$$

A strategy σ_i is considered a **dominant strategy** for player i if, no matter what strategies other players chooses, i will do at least as well playing σ_i as it would doing anything else. Formally, σ_i is a dominant strategy if:

$$\forall\, \vec{\sigma} \in \Sigma_1 \times \cdots \times \Sigma_n, \forall \sigma_i' \in \Sigma_i \cdot u_i(\vec{\sigma}_{-i}, \sigma_i) \geq u_i(\vec{\sigma}_{-i}, \sigma_i').$$

Note that if a rational player has a dominant strategy, then she has no incentive to choose any other strategy. Therefore, using the concept of dominant strategies, if a dominant strategy does exist, it should be played. In Example 2.5, for both players, the strategy "Defect" is a dominant strategy. Therefore, it is reasonable to predict that rational players would result in the strategy profile $\langle D, D\rangle$. However, there are many cases where no dominant strategy exists, as is the case in the classic **Battle of the Sexes** in Example 2.6.

Example 2.6 The classic Battle of the Sexes game consists of two players (one male and one female) where the two players seek to coordinate their choice of whether to go shopping or watch a football game. Unfortunately, the decision has to be simultaneous and no prior coordination is possible. The payoff matrix is represented as follows.

		Player 2	
		Shopping	*Football*
Player 1	*Shopping*	(1, 2)	(0, 0)
	Football	(0, 0)	(2, 1)

Without a doubt, the most influential and celebrated solution concept in game theory to date is **Nash equilibrium**. A strategy profile $\vec{\sigma}$ is a **pure Nash equilibrium** if no player could benefit from deviating from his strategy, *assuming that all other players keep their strategies*. Formally, $\vec{\sigma}$ is a Nash equilibrium if:

$$\forall i \in N \; \forall \sigma_i' \in \Sigma_i \cdot u_i(\vec{\sigma}_{-i}, \sigma_i) \geq u_i(\vec{\sigma}_{-i}, \sigma_i')$$

or, equivalently,

$$\forall i \in N \cdot \sigma_i = \arg\max_{\sigma_i} u_i(\vec{\sigma}_{-i}, \sigma_i).$$

Intuitively, a pure Nash equilibrium means that each player is using a **best response** strategy to other players' choices. Therefore, nobody can benefit by deviating from a Nash equilibrium alone (without coordinating with others).

In Example 2.5, the strategy profile $\langle D, D\rangle$ is a pure Nash equilibrium.

Lemma 2.7 *Not every game has a pure Nash equilibrium and some games have more than a single Nash equilibrium.*

The proof of Lemma 2.7 is left as Exercise 2.10.

Thus far, we assumed that player i chooses a pure strategy $\sigma_i \in \Sigma_i$. However, a player can also choose to use a *mixed strategy*. A mixed strategy for player i is a probability distribution over the pure strategies Σ_i, denoted ms_i. Let $M\Sigma_i = \{ms_i\}$ be the set of all possible mixed strategies for player i. Intuitively, by using a mixed strategy, a player randomly selects a pure strategy according to a chosen probability distribution. Since probabilities are continuous, there are an infinite number of mixed strategies available to each player if $|\Sigma_i| > 1$. One can regard a pure strategy $\sigma_i \in \Sigma_i$ as a degenerate case of a mixed strategy $ms_i \in M\Sigma_i$, in which σ_i is assigned a probability of 1 and all other pure strategies are assigned a probability of 0.

As a natural extension of our definition of a pure Nash equilibrium, a mixed strategy profile $\vec{ms} = \langle ms_1, \ldots, ms_i, \ldots, ms_N \rangle$ is considered a mixed Nash equilibrium if:

$$\forall i \in N \ \forall ms'_i \in M\Sigma_i \cdot u_i(\vec{ms}_{-i}, ms_i) \geq u_i(\vec{ms}_{-i}, ms'_i)$$

where $\vec{ms}_{-i}$ takes the usual meaning.

Intuitively, in a mixed Nash equilibrium, each of the selected distributions should have the property that it is a *best response* to the other distributions; similarly to the pure Nash equilibrium. This means that each action which is assigned with positive probability is among the actions that are best responses, in expectation of the distribution(s) chosen by the other player(s).

Given a mixed strategy ms_i, we define the support for ms_i as:

$$Support(ms_i) = \{\sigma_i \in \Sigma_i | \sigma_i \text{ is assigned a positive probability under } ms_i\}.$$

It turns out that, unlike pure Nash equilibrium, at least one mixed Nash equilibrium exists in every finite game.

Theorem 2.8 Nash's Existence Theorem [225, 226]. *Every game with a finite number of players and a finite number of pure strategies for each player has at least one Nash equilibrium.*

This insight earned John Forbes Nash the Nobel Memorial Prize in Economics in 1994 (which he shared with Reinhard Selten and John Harsanyi). The proof of Nash's Theorem (Theorem 2.8) is rather technical. A detailed proof along with intuitive interpretations is available in [162].

It is easy to verify that in the two-player rock-paper-scissors game, with a utility of 1 for the winning player, -1 for the losing player, and 0 for both if there is a draw, the mixed strategy profile where each player randomly selects an action uniformly is a mixed Nash equilibrium. No pure Nash equilibrium exists in the game and none of the players have a dominant strategy.

The analysis of the Battle of the Sexes game (Example 2.6) is left as Exercise 2.9.

2.2.2 EXTENSIVE FORM GAMES

Normal form games capture key aspects of strategical decision-making in environments with multiple decision-makers. However, in normal form games, players choose their actions simultaneously with no knowledge of the choices of their counterparts. In many environments, players choose their actions in turns, adding a temporal aspect to the decision-making process. Such games are usually represented as *extensive form games*.

In the scope of this chapter, we will discuss *complete* and perfect information Extensive form games. Namely, the structure of the game and the payoffs of the players are commonly known (complete) and that player observes other players' moves which were played before theirs. Note that, in general, a game with complete information may or may not have perfect information, and vice versa.

An extensive form game of *complete* and *perfect* information consists of a set of players $N = \{1, \ldots, n\}$ and a rooted tree T, called the *game tree*. Each non-terminal node is assigned to one of the players and denoted a *decision node*. Player i must choose an action for each of her decision nodes according to the available actions, represented as edges. *Terminal node* (leaf) consists of a *payoff vector* $\langle r_1, \ldots, r_n \rangle$ where r_i is the payoff of player i from the resulting realization of the game. The game is realized by following a path from the root node to a leaf according to the players' decisions.

A pure strategy is thus a player's selection of precisely one outgoing edge (action) for each of the player's decision nodes. A mixed strategy is defined similarly by selecting a probability distribution at each decision node.

It is customary to visually represent extensive form games in the form of a tree, see Example 2.9.

Example 2.9 Consider a sequential version of the prisoner's dilemma presented in Example 2.5. Let player 1 (P_1) move first. Player 2 (P_2), given player 1's move, moves second. For illustrating new ideas to come, we now assume that the utilities of both players in the $\langle D, D \rangle$ path are -1. The game can be represented as follows.

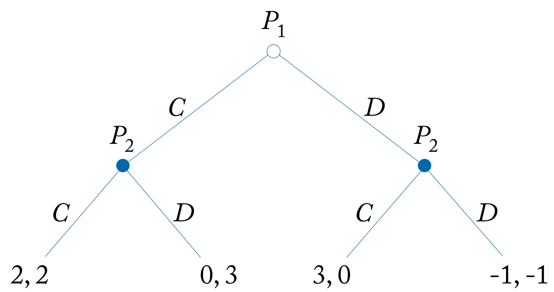

Lemma 2.10 *In extensive form games with complete and perfect information, any mixed strategy for player i will result in a lower or equal utility for player i compared to some pure strategy available to player i.*

Following Lemma 2.10, which is left as Exercise 2.16 at the end of the chapter, we assume players of an extensive form game will use pure strategies.

Similar to normal form games, the combination of players' strategies results in a strategy profile of the form:

$$\vec{\sigma} = \langle \sigma_1, \ldots, \sigma_n \rangle \in \Sigma_1 \times \cdots \times \Sigma_n.$$

A Nash equilibrium is a strategy profile $\vec{\sigma}$ which satisfies the following condition:

$$\forall i \in N \ \forall \sigma_i' \in \Sigma_i \cdot u_i(\vec{\sigma}_{-i}, \sigma_i) \geq u_i(\vec{\sigma}_{-i}, \sigma_i'),$$

or, equivalently,

$$\forall i \in N \cdot \sigma_i = \arg\max_{\sigma_i} u_i(\vec{\sigma}_{-i}, \sigma_i)$$

where $\vec{\sigma}_{-i}$ is a strategy profile without player i's strategy.

Getting back to Example 2.9, one can easily verify that two strategy profiles satisfy the above definition: First, P_1 chooses D, while P_2 selects D in the left node and C in the right. Second, P_1 chooses C while P_2 chooses D in both her nodes. The second equilibrium is supported by P_2's *threat* of playing D in the right subtree. However, this threat is *not credible* since if P_1 plays D, P_2 has a dominant strategy, which is to play C. Thus, although there are two Nash equilibria, only one is reasonable. In order to rule out non-plausible Nash equilibria we use the *Subgame Perfect Nash Equilibrium* (SPNE).

Each decision node in a game tree induces a *subgame*. Specifically, the tree rooted in each decision node constitutes as a well-defined game. Thus, an SPNE is a strategy profile which is a Nash equilibrium at every subtree of the game.

Intuitively, SPNE is a stronger generalization of a Nash equilibrium, where the former requires the strategy profile to be a Nash equilibrium at every subgame including the entire game tree, which is the only condition for the latter.

Theorem 2.11 SPNE Existence Theorem [209]. *Every finite extensive form game of perfect information has an SPNE.*

The proof of the above theorem is by providing an algorithm for deriving an SPNE for every perfect extensive form game. The algorithm commonly used is the backward induction algorithm, also known as *Zermelo's algorithm* which was developed by Ernst Zermelo (1871–1953).

The process of Zermelo's algorithm is rather simple. Start by determining optimal behavior at the lowest (final) decision node at each path of the game tree. This can be achieved by solving a single decision-maker's problem. Note that no reasoning regarding the other player(s) is needed for deriving an optimal move at any decision node at a final decision node. Then replace the decision node by a terminal node resulting from the optimal choice calculated in the previous step. The process is repeated until we reach the root node.

In finite extensive form games with perfect information, Zermelo's algorithm is guaranteed to terminate in time polynomial in the tree size and is guaranteed to find an SPNE.

The page is too short to hold all basic notions of Game Theory. The interested reader is thus encouraged to read [32, 210] for further details and examples.

2.3 ARE PEOPLE RATIONAL? A SHORT NOTE

Rationality is widely used as an assumption of the behavior of individuals in microeconomic models and analyses and appears in almost all economics textbook treatments of human decision-making. The Nobel Memorial Prize in Economic Science of 2002 was awarded to two experimental economists: *Daniel Kahneman* and *Vernon Smith*. Both scientists shared the award for testing the limits of the standard economic theory of choice in predicting the actions of real people [159]. Interestingly enough, the two received the award for seemingly contradictory reasons.

Kahneman received the award for his joint work with *Amos Tversky* (1937–1996) on human judgment and decision-making. Their results demonstrated irrational wrinkles and deviations in the normally assumed rational behavior as described above.

Smith received the award for setting laboratory experimental methodologies by which he investigated economical mechanisms. He found that people were acting (approximately) rationally in most tested settings, verifying economical rational predictions in markets.

Yisrael Aumann, the recipient of the Nobel Memorial Prize in Economic Science of 2005, explains this conflict by hypothesizing that people are rational in their decision-making *procedures* but not in their every decision [18]. Namely, the *rule* that people apply to solve a particular problem is *usually* the right one, but not *necessarily* the right one, as pure rationality suggests. According to this paradigm, when predicting people's decisions in the real world, one should examine the rules and heuristics people deploy, which in many cases can in fact lead to optimal decisions yet in quite a few settings lead to sub-optimal decision-making.

There is also considerable literature on understanding whether and why people deviate from formal models of deductive reasoning, e.g., [39, 165]. Interestingly enough, non-human primates exhibit similar decision-making tendencies to those of humans, suggesting a strong biological base for the presented human decision-making behavior [280].

2.4 EXERCISES

2.1. (Level 1) Prove Lemma 2.1.

2.2. (Level 2) Alice is a decision-maker who needs to choose a single item out of a *finite set* of items (X). Each item is represented as a vector of size n where the first element indicates the monetary value of the item, the second indicates the usefulness of the item to Alice and so on. All features are represented as **integer values** in the $[0, 9]$ range. Alice uses the following preference relation: $x = \langle x_1, \ldots, x_n \rangle \preceq \langle y_1, \ldots, y_n \rangle = y$ if and only if $x_1 > y_1$ or, if $x_1 = y_1$ and $x_2 > y_2$, or $x_1 = y_1, x_2 = y_2$ and $x_3 > y_3$ and so on. (A) (Level 1) Is Alice's preference relation reflexive, complete and transitive? (B) (Level 1) Provide a utility function $u : X \mapsto \mathbb{R}$ which represents Alice's preference. (C) (Level 2) Now assume each feature is represented as a **real number** in the $[0, 9]$ range. Prove that Alice's preference relation is not continuous. (D) (Level 3) Assuming each feature is represented as a **real number** in the $[0, 9]$ range, prove that Alice's preference relation cannot be represented using a utility function (hint: without loss of generality, assume $n = 2$).

2.3. (Level 3) Following Lemma 2.2: (A) Suppose that $X \subset \mathbb{R}^n$. Prove that if a preference relation $\preceq$ is reflexive, complete, transitive, and continuous then there exists a utility function $u : X \mapsto \mathbb{R}$ that represents it. (B) Suppose that $X = \mathbb{R}^n$. Prove that if a preference relation $\preceq$ can be represented as a utility function $u : X \mapsto \mathbb{R}$ then it is reflexive, complete, transitive, and continuous.

2.4. (Level 2) Construct a preference relation on $\mathbb{R}$ that is not continuous, but admits a utility representation.

2.5. (Level 2) Bob considers weather to invest in stock A or stock B. Given the economic uncertainty in his country, there are two "nature" states, *stability* and *unrest*, which will be revealed after he makes his choice. Bob's revenues are represented in the following table:

		Stability	*Unrest*
Bob	*StockA*	3	9
	StockB	16	0

Bob is assumed to have preferences which monotonically increase with revenues. (A) If the preferences are represented by a utility function, what are the arguments of the function? (B) Suppose that Bob is maximizing expected utility and believes that the probability of *stability* is $1/3$ and the probability of *unrest* is $2/3$. Express Bob's objective function. What is a rational decision for Bob? (C) Bob chose to invest in stock A. Should he be considered irrational? Can "conservatism" or "pessimism" explain his decision in a way that would still define him as rational? Explain.

2.6. (Level 1) How can decision theory explain why many people both purchase lottery tickets (implying risk-loving preferences) and insure against losses (implying risk aversion)?

2.7. (Level 2) Consider the following experiment conducted by *Maurice Allais*. Imagine yourself choosing between the following two alternatives: (A) win 1 million dollars for sure; or (B) a 10% chance of Winning 5 million dollars, a 89% chance of winning 1 million dollars, and a 1% chance of winning nothing. Which one would you choose? Now consider the following two alternatives: (C) an 11% chance of winning 1 million dollars and an 89% chance of winning nothing; or (D) a 10% chance of winning 5 million dollars and a 90% chance of winning nothing. Which one would you choose? In many surveys, subjects who were offered these alternatives chose A over B and D over C. This is often refereed to as the "Allais Paradox" [5]. Explain why this is considered a paradox and what significance this paradox has on the expected utility maximization paradigm.

2.8. (Level 1) Explain the following claim: "a *single decision-maker environment* does **not** mean a single *agent* is occupying the environment."

2.9. (Level 1) Analyze the *Battle of the Sexes* (Example 2.6). Who are the players? What are their pure strategies? Does any player have a dominant strategy? Find all Nash equilibria.

2.10. (Level 1) Prove Lemma 2.7.

2.11. (Level 1) Alice and Bob are playing a card-game (say Poker). Each player gets two cards from the deck which the other player cannot see and then the game plays out. During the game players make decisions (e.g., raise or fold in Poker). Bob is well known for playing a fixed-policy where during every turn his action can be anticipated perfectly *if one were to know his cards*. Playing the part of Alice, how would you model your decision-making environment, as a single decision-maker or multiple decision-makers? Explain.

2.12. (Level 2) Present an example of a two-player normal form game in which one of the players has a dominant strategy that is **not** a pure strategy.

2.13. (Level 1) Analyze the two-player rock-paper-scissors-lizard-Spock game `http://www.samkass.com/theories/RPSSL.html`. Who are the players? What are the possible pure strategies? Find all Nash equilibria.

2.14. (Level 2) Prove that, if a player has two dominant strategies, then for every strategic choice made by her opponents, the two strategies yield her equal payoffs.

2.15. (Level 2) A strategy s is called a *strictly* dominant strategy if, no matter what the other players choose, the player is *strictly* better off playing s than any other strategy. (A) Prove that, if a player has a *strictly* dominant strategy s, then that player plays s in all Nash equilibria. (B) (Level 3) Prove that strictly dominant strategies are always pure.

2.16. (Level 2) Prove Lemma 2.10.

2.17. (Level 1) Show Zermelo's algorithm execution in Example 2.9.

2.18. (Level 1) Consider a modified version of the Battle of the Sexes (Example 2.6) where Player 1 chooses first and Player 2 sees Player 1's choice before making her own. Use the same utilities as defined for the normal form version. (A) Describe the modified game as a game tree. (B) Find all Nash equilibria.

2.19. (Level 1) Following Auman's view on human rationality (Section 2.3), describe a setting in which people (often) deploy the correct decision rule yet the realization of that rule may produce sub-optimal decisions.

CHAPTER 3

Predicting Human Decision-Making

"All models are wrong, but some are useful"

Gorge Box

Designing intelligent agents that interact proficiently with people necessitates the prediction of human decision-making. We present and discuss three prediction paradigms for predicting human decision-making which are common in many real world application fields. These paradigms are illustrated and compared across a wide variety of domains and applications in Chapter 4.

3.1 EXPERT-DRIVEN PARADIGM

An **expert-driven model** for predicting human decision-making is a mathematical formulation, *articulated by an expert*, which is assumed to adequately predict people's choices in a given setting or across different settings.

Experts of different disciplines have provided a variety of models to predict people's decisions based on their experience, theoretical assumptions and domain knowledge. In this section, we review and exemplify the most prominent quantitative models which are commonly used for the design of automated agents.

3.1.1 UTILITY MAXIMIZATION

In Chapter 2, we presented the *expected utility maximization* paradigm. Despite its appealing theoretical properties, assuming that (most) people follow the presented principles is inherently flawed. Specifically, people have limited cognitive abilities [219] which prevent them from considering and fully evaluating all options to their fullest extent even if they seek to maximize expected utility. The seemingly "irrational" behavior that people present in an overwhelming number of behavioral experiments (e.g., [54, 168, 205] to name a few) makes the use of utility maximization principles seem irrelevant.

It turns out that, despite its limitations, predicting people's decisions using utility maximization principles can be quite effective in many settings. In his book [32], Binmore provides several guidelines for when he concludes **Game Theory** principles (which prescribe rational decision-making in multiple decision-maker settings—see Section 2.2) should be able to adequately predict people's decisions: (1) players' incentives are adequate and sufficiently large such that they can override social norms (and in particular, norms of cooperation); (2) the game is sufficiently simple; and (3) players are given enough opportunities for trial-and-error learning. These guidelines, which we refer to as **Binmore's guidelines**, which are also supported in additional works such as [89], can be easily generalized to a single decision-maker setting as well.

Perhaps the most renowned success of this approach in recent years is in the security field. To exemplify the success of this approach and explain the general criteria that made this approach successful, consider the task of placing randomized road checkpoints in order to protect an airport against potential attacks by terrorists, drug smugglers etc. The game-theory-based application named **ARMOR** was developed precisely for this task and has been in use for almost ten years in order to protect the Los Angeles international airport (LAX) [251]. The system prescribes where and when security officers should set up checkpoints to check cars driving into LAX. The authors use a game-theoretic model to represent the interaction between the security officers and the potential attacks and *assume* that the possible human attackers are *rational*, and therefore maximize expected utility. A thorough investigation of the prediction of human decision-making in security settings is available in Section 4.2.

The highly successful deployment of the *ARMOR* system is not surprising given Binmore's guidelines. Security experts claim that possible human attackers usually have sufficient time and resources to conduct surveillance, attempt "mock attacks" and reason about the decision problem, which is where, if anywhere, to attack the airport. Thus, the rationality assumption leads to beneficial recommendations by the *ARMOR* system. Note that, naturally, social norms would not play a role in the attackers' decision-making in this case.

However, in practice, it is rarely the case that utility maximization models and reality are aligned. The evaluation of the utility maximization notion with real world data reveals significant discrepancies in many domains (recall the discussion on human rationality in Section 2.3). The obvious way to cope with these discrepancies is to assume that people follow utility maximization *with noise*. Namely, people have a "trembling hand" which with some probability performs unintended, non-optimal strategies. This "slip of the hand" or tremble notion is attributed to *Reinhard Selten* (1930–2016) [288]. This trembling hand is usually assumed to be a random variable with a zero mean. Namely, on expectancy, this "tremble," has no effect on the decision-making process.

When significant deviations from utility maximization predictions are observed, or when seemingly non-utility-maximizing decision-making patterns arise, different prediction models may be needed. The following models, which we discuss next, adopt the **Bounded Rationality** approach, first introduced by the 1978 Nobel-laureate and the 1975 Turing Award recipient *Herbert Simon* (1916–2001) [296]. Bounded rationality suggests that people have a limited ability to compute the expected utility of every single decision alternative and thus might not be able to choose the optimum strategy as suggested in classical decision-making and game theory literature. Instead, bounded rationality suggests that people use a more heuristic approach that may lead to *suboptimal decision-making* in some settings. We discuss several bounded rationality models, common in agent design, in the following.

3.1.2 QUANTAL RESPONSE

Quantal response assumes that humans are expected utility maximizers who noisily estimate each strategy's expected utility [214]. Formally, a person is assumed to maximize the expected value of her utility function which is perpetuated by a zero-mean random noise. Specifically, let us assume that for a strategy profile $\vec{s}$ the "true" expected utility of person i is $u_i(\vec{s})$. Then, according to quantal response, the human estimates the utility as $\hat{u_i}(s) = u(s) + \epsilon_{s_i}$ where $\epsilon(s_i)$ is the noise term associated with the person's strategy. As a result, the person may select any strategy with positive probability, as assigned by her quantal response function.

The most common quantal response function is the **logit quantal response** (often referred to as simply "quantal response function"), where each possible strategy s_i is assigned a non-negative probability as follows:

$$\mathbb{P}(s_i) = \frac{exp[\lambda \cdot u(\vec{s}_{-i}, s_i)]}{\Sigma_{s'_i} exp[\lambda \cdot u(\vec{s}_{-i}, s'_i)]}$$

where $\lambda \in [0, \infty)$ is the **rationality parameter**, indicating how rational the decision-maker is assumed to be. Namely, when $\lambda = 0$ the decision-maker uses a uniform random choice over the possible strategies regardless of their expected utility. Conversely, when $\lambda \to \infty$, quantal response converges to a "pure" utility maximization as discussed before. The proof of these two observations is left as Exercise 3.2.

Simply put, quantal response suggests that instead of strictly maximizing expected utility, as suggested in Section 3.1.1, individuals respond *stochastically* in games: the chance of selecting a non-optimal strategy increases as the cost of such an error decreases. The difference between utility maximization and quantal response is illustrated in Example 3.1.

Example 3.1 Consider a decision-maker with three strategies: s_1, s_2, and s_3. The decision-maker calculated the expected utility of each strategy denoted $EU[s_1] = 1.01$, $EU[s_2] = 1$, and $EU[s_3] = 0.25$. By assuming that the decision-maker follows the utility maximization paradigm, the decision-maker is expected to choose strategy s_1 *with certainty* as it has the highest expected utility. Using the quantal response, high-utility actions are assigned higher probability, depending on the rationality parameter λ. Figure 3.1 depicts this example with varying λ values.

The quantal response notion was found useful in predicting people's decisions in a wide variety of economic settings. For example, in all-pay auctions [8], first-price auctions [122], and in alternating-offer bargaining [122].

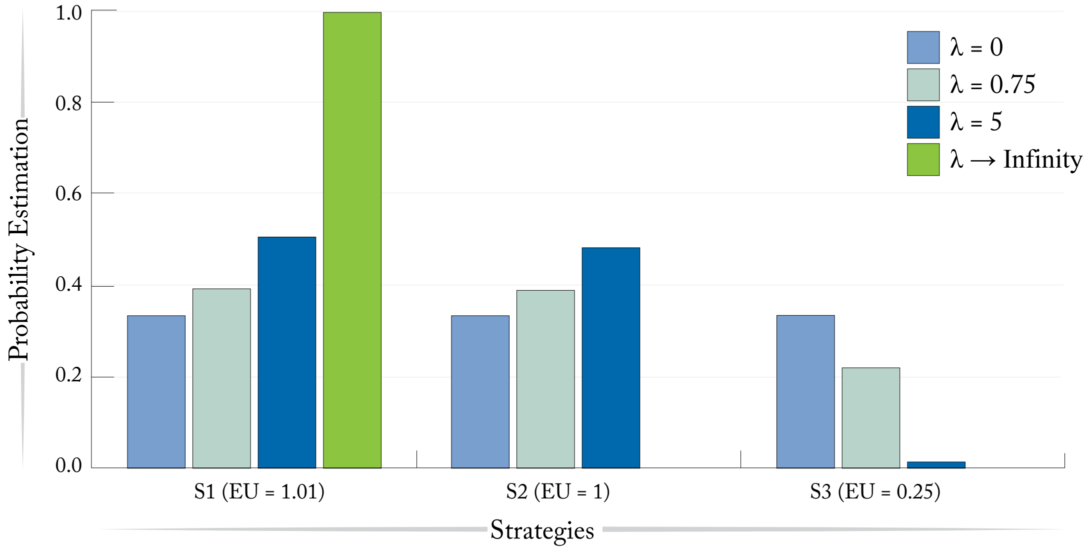

Figure 3.1: Probability estimation for a quantal response decision-maker over three strategies per Example 3.1.

From an intelligent agent design perspective, quantal response has been used to predict people's decision-making in an attempt to influence their decisions. In [20], the authors found that by using a quantal response prediction model of human decision-making, their information revealing agents performed better at influencing human decisions compared to expected utility maximization prediction methods. Another relevant example comes from the field of negotiations. Using a quantal response function, several works have been shown to be able to negotiate effectively with people (e.g., [96]). We provide a detailed discussion of the negotiation domain in Section 4.3.

The definition of **quantal response equilibrium** is left as Exercise 3.4.

The interested reader is encouraged to read [214] for a thorough discussion over the relation between quantal response and other game-theoretical concepts.

3.1.3 LEVEL-k

Level-k model assumes that humans can perform only a bounded number of **iterations of strategic reasoning**. Level-k models were introduced by Stahl and Wilson [300, 301] and Nagel [224]. Formally, a person who follows the level-k model is associated with a level $k \in \mathbb{N} \cup \{0\}$, corresponding to the number of iterations of reasoning the person is able to perform. A $level - 0$ person is nonstrategic and follows a simple decision rule (most commonly, playing randomly or playing an intuitive default strategy). A level-k person, for $k \geq 1$, best responds to the strategy played by level-$(k-1)$ agents. If a level-k agent has more than one best response, it is common to assume that she mixes uniformly over them. This iterative strategical thinking process is a *step-by-step reasoning process* rather than *circular concepts* such as a Nash equilibrium. Level-k models have been shown to adequately predict people's decisions in a variety of experimental studies such as [71] and consequent works.

Similar to the quantal response model (Section 3.1.2), where the λ parameter had to be set, in order to predict a person's decisions using a level-k model one would need to set k appropriately. While k can be estimated from collected data, more troubling is the need to determine what the level-0 player would do, as it recursively defines what a level k player would do. Wright and Layton-Brown [334] provide several level-0 meta-models which they show to work well across a variety of settings. It is important to note, however, that level-0 reasoners may not be "dumb." For example, if a person starts to analyze the game carefully but gets confused or makes an error, she might make a choice that appears random, much like how a small calculation error in a long proof can lead to odd results.

Arad and Rubinstein [9] designed a simple two-player game that naturally triggers level-k reasoning. The game works as follows: each player requests an integer amount of money between 11 and 20 NIS. Each player will receive the amount s/he requests. A player will receive an additional amount of 20 NIS if s/he asks for exactly one NIS less than the other player. The experiment is aimed at triggering level-k reasoning where a level-0 reasoner would request 20 NIS, a level-1 reasoner would request 19 and so on. The game's level-k models are illustrated in Figure 3.2. An experiment was conducted with 108 students from Tel Aviv University (the first author of this book happens to be one of those students) which showed that level-k model subjects did not use more than $k = 3$ of reasoning in choosing their actions. Past studies of level-k reasoning reached the same conclusion. Crawford et al. [73] reviews a large body of evidence suggesting that level-k models can out-predict other models in various other environments as well.

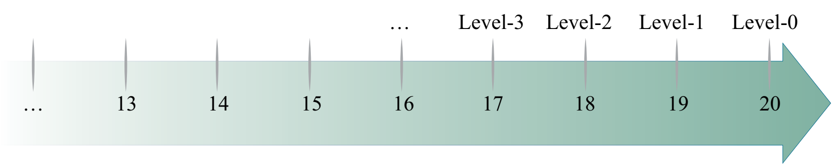

Figure 3.2: Arad and Rubinstein's 11-20 Money Request Game analyses through level-k modeling.

Note that level-k models are agnostic over whether individuals stop the iterated reasoning because of their own cognitive constraints, or because of their beliefs over the cognitive constraints of their opponents. Nevertheless, players with a high IQ have been shown to use more steps of thinking than others [121]. Specifically, in *Chess* stronger players' moves have been shown to be best explained by larger values of k [33].

There are several extensions of the level-k model. Most prominent is the inclusion of noisy actions. Namely, a level-k decision-maker may decide on a non-best-response action with an ϵ probability. An additional extension in that spirit is the integration of quantal response within the level-k modeling, denoted $Qlevel-k$. In a $Qlevel-k$ model, the reasoner responds using a quantal response to $Qlevel-k-1$ reasoners. Specifically, a $Qlevel-k$ decision-maker responds quantally to level-$k-1$ decision-makers (which in turn respond quantally to level-$k-2$ and so forth, unlike level-k decision-makers which use best response to level-$k-1$ (which in turn best respond to level-$k-2$ and so forth). Note that in order to deploy the $Qlevel-k$ model, one would need to set the λ parameter for each level.

3.1.4 COGNITIVE HIERARCHY

Closely related to the level-k model (Section 3.1.3), the **cognitive hierarchy model** was proposed by [55] to account for people's assumed iterative reasoning process. In the cognitive hierarchy model, a decision-maker is assumed to be associated with a level $k \in \mathbb{N} \cup \{0\}$, denoted *cognitive hierarchy level*, corresponding to the number of iterations of reasoning the person is able to perform. However, unlike the level-k model, a reasoner of cognitive hierarchy level k best responds to a *distribution* of cognitive hierarchal reasoners of lower levels and not just to reasoners of level $k-1$. Specifically, while a level-k reasoner best responds to the level-$k-1$ reasoners, cognitive hierarchy level k best responds to an hypothesized distribution of lower levels from 0 to $k-1$. We use the notation $f_k(j)$ to denote the *belief* that a decision-maker of cognitive hierarchy level k has regarding the proportion of cognitive hierarchy level j decision-makers in the population.

Note that under both level-k and cognitive hierarchy, a decision-maker of any level ignores the possibility that other decision-makers may be making as many strategical iterations as she does or even more.

Similar to the level-k model, the iterative process begins with decision-makers of cognitive hierarchy level 0, who are assumed not to perform any strategical reasoning and merely choose according to some rule of thumb (commonly, assumed to be uniform distribution).

In order to predict a person's decisions using a cognitive hierarchy model, one would need to set k appropriately. Furthermore, the level 0 rule would need to be set. However, in addition to the above (which are also needed in level-k model), when using a cognitive hierarchy model, the decision-maker's hypothesized distribution of cognitive hierarchy levels in the population is needed (as she is assumed to best respond to that distribution). Camerer et al. [55] advocate a single-parameter distribution in which the cognitive hierarchy levels of decision-makers in the population are distributed according to a *Poisson* distribution. Namely, the proportion of cognitive hierarchy level j in the population is assumed to be

$$f(j) = \frac{\tau^j e^{-\tau}}{j!}$$

where τ is a positive real number equal to the expected cognitive level in the population.

The Poisson distribution assumes that the frequency of very high cognitive levels drops off quickly for higher values of k. For example, if the average number of thinking steps in the population is 1.5 ($\tau = 1.5$), then less than 2% of players are expected to have five or more steps of thinking. Camerer et al. [55] found that τ values of between 1 and 2 explain empirical results of human decision-making in about 100 games, suggesting that assuming a value of 1.5 could give reliable predictions for many other games as well.

The authors assume that each reasoner of cognitive hierarchy level k reasons according to the above distribution but consider only reasoners of cognitive hierarchy levels 0 to $k-1$. As a result, the above distribution does not necessarily sum to 1 for the reasoner. The authors suggest bypassing this issue by normalizing the reasoner's distribution; they divide each applicable $f(j)$ by the sum of all applicable $f(j)$-s. Simply put, for a reasoner of cognitive hierarchy level k, she assumes that the population is distributed according to

$$f_k(j) = \begin{cases} \frac{\tau^j e^{-\tau}}{j! \sum_{l=0}^{k-1} f(l)} & \text{if } j < k \\ 0 & \text{otherwise.} \end{cases}$$

The above "normalized" (truncated) Poisson distribution is the most popular assumption over the population distribution, yet it is not the only one. Other forms include a normalized uniform distribution and exponential decay distributions (see [134]). The formulation of these notions is left as an exercise.

A simple setting which illustrates apparent cognitive hierarchy behavior is the **beauty contest game** proposed by *John Maynard Keynes* (1883–1946) [182]. Keynes likens the stock market to a newspaper contest in which people guess which faces *others* will consider to be the most beautiful. Keynes notes that "It is not the case of choosing those which, to the best of one's judgment, are really the prettiest, nor even those which average opinion genuinely thinks the prettiest. We have reached the third degree, where we devote our intelligences to *anticipating what average opinion expects the average opinion to be*. And there are some, I believe, who practice the fourth, fifth, and higher degrees" [182, p. 156]. The essence of Keynes's observation is captured in Example 3.2.

Example 3.2 Each player is asked to pick a number between 0 and 100. The player whose number is closest to $\frac{2}{3}$ of the average wins a prize. The rest get nothing. A cognitive hierarchy level 0 player will select a number non-strategically by sampling a number at random or choosing a number which may have special significance to the player (in either case, the player's choice is indistinguishable from a randomly generated number by other players).

The cognitive hierarchy model seems to accurately explain people's behavior in several empirical investigations of Example 3.2 [41]. Specifically, past studies have shown that players' choices spiked at $50 \cdot \frac{2}{3}^k$ for $k = 1, 2, 3$, as suggested by the level-k model. Nevertheless, most players choose different numbers, mainly in the 1-33 interval, suggesting a possible reasoning over the distribution of players. Additional examples for the use of cognitive hierarchy models for two-player games are available in [125].

Note that unlike level-k models, it is uncommon to associate error rates with cognitive hierarchy models. Namely, a person is assumed to respond perfectly to her belief over the level distribution. However, the incorporation of quantal response was proposed by Wright and Leyton-Brown in [333]. The authors show the benefit of the integrated model (quantal response with cognitive hierarchy) in several publicly accessible databases.

3.1.5 BEHAVIORAL SCIENCES

Behavioral sciences such as psychology and cognitive science provide a wide gamut of observations, theories, empirical studies, and general criteria to explain and predict how people make decisions. However, as noted at the beginning of this section, only a few of the proposed models provide a mathematical model which is needed for automated agent's design.

In his popular book, *Dan Ariely* advocates that people, to great extent, are **"Irrationally Predictable"** [11]. Namely, Ariely notes that simple observations and manipulations can have a far-reaching effect on people's decisions and can assist in predicting them. While these behavioral observations are not mathematically formulated for the most part, utilizing them in automated agent design can be very useful. For example, Hajaj et al. [137] recently showed that comparison shopping agents, such as the ones listed in `Shoppingbots.info`, can increase their revenues substantially by displaying prices sequentially in a way that will leverage known decision-makers' cognitive biases. Specifically, they instantiate the anchoring-and-adjustment insight from Kahneman and Tversky [170] along with *ad-hoc, domain-specific heuristics* to derive a computational method which they show to increase a comparison shopping agent's revenue.

Unfortunately, articulating a *mathematical* behavioral science-based model can be very complex and may require extensive domain knowledge and the incorporation of expert-defined heuristics. Let us consider The Ultimatum Game in Example 3.3.

Example 3.3 The Ultimatum Game The ultimatum game consists of two players. The first player (the proposer) conditionally receives a sum of money (say, $10) and proposes how to divide the sum between the proposer and the other player. The second player (the responder) chooses to either accept or reject this proposal. If the second player accepts, the money is split according to the proposal. If the second player rejects, neither player receives any money. The game is typically played only once so that reciprocation should not pose a significant issue.

Güth et al. [130] was the first experimental study of this game. The mean offer that human proposers gave in this study was a split where the proposer receives 63% of the money and the responder receives 37% of the money. Subsequent studies have revalidated these findings and found that offers below 20–30% are very likely to be rejected whereas offers which propose at least 50% to the responder are almost always accepted. The results vary in different parameters of the game such as cultures and different levels of familiarity between the proposer and the

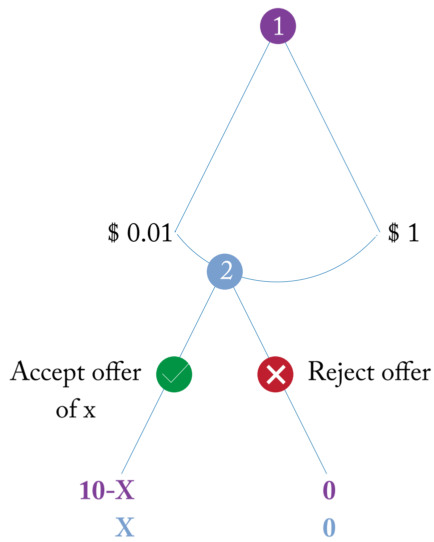

Figure 3.3: The Ultimatum Game (Example 3.3) represented as an extensive form game.

responder (see [150, 151] for discussion on these issues). This behavioral evidence suggests a general guideline as to how one should predict how a human proposer and a human responder would act in an ultimatum game setting. However, how can one translate the above insights into a *mathematically viable model* that would work in a new domain or in a new setting?

Consider an intelligent agent that negotiates a possible trade deal with a human interlocutor. The agent wants to offer a "take-it-or-leave-it" contract to the human. Such an offer can be easily viewed as a proposal as to how to split the profit from the trade deal. The situation can be modeled as an Ultimatum Game (Example 3.3). Therefore, the agent may take the experimental evidence discussed above and use it as a guideline, for instance by assuming that the likelihood that the human would accept an offer below 50% diminishes according to some expert-based functional form (which parameters may be learned from appropriate data, if such exists). We provide a thorough discussion over the negotiation domain in Section 4.3.

The above examples are part of a long list of well-known and well-studied **cognitive biases** of human decision-makers. Cognitive biases are tendencies to think and act in certain ways that can lead to systematic deviations from a standard of rationality or good judgment [141]. These biases are often studied in psychology and behavioral economics and are usually confirmed by replicable research, testing the hypothesized bias in different experimental settings including different cultures, genders, etc. A significant effort is also made in investigating the evolutionary background of these biases (e.g., [64]). We are interested in these cognitive biases as they have significant effects on humans' decision-making. Thus, identifying and appropriately leveraging them in intelligent agent design may be very effective. We divide these biases into two groups:

general biases, which have been identified across many different domains and environmental settings, and **domain specific biases**, which are (suspected to be) limited to a specific domain or to a narrow group of environmental settings.

Note that while the list of known decision-making biases gets longer by the day, only a small fraction of these biases are currently in use by automated agents for predicting human decision-making. We will discuss this issue in Chapter 4.

General Biases

We will discuss only a few prominent cognitive biases which seem to be presented across a wide range of settings. We will exemplify their use in agent design as well. Ideally, each of the mentioned biases should have a chapter of its own. Such a chapter would discuss the evolutionary background of the bias and the multitude of experimental results that validate and question it. However, due to space limitations, we will restrict ourselves to a short and concise definition of each of the prominent biases, along with a practical example of where and when this bias can be of use in agent design. The interested reader is encouraged to refer to one of the many books on the topic—the authors of this book would recommend the following easy-access books [11, 81, 169, 253, 324, 343].

- **Anchoring** is the tendency of people to rely too heavily, or "anchor," on a past reference or on one trait or piece of information when making decisions. Kahneman and Tversky were the first to document the anchoring bias in an experiment involving a roulette wheel marked with integers ranging from 0–100. Each participant witnessed a spin of the roulette wheel. They were then asked whether they thought that the percentage of United Nations member countries that was from Africa was greater or smaller than the number spun on the wheel. Next, they were asked to estimate the true percentage. Participants who saw the wheel stop on the number 10 guessed, on average, that the actual percentage of African countries belonging to the United Nations was 25%. In contrast, those who saw the wheel stop on the number 65 guessed, on average, that the percentage from Africa was 45%. A common example for the use of Anchoring in agent design stems from automated negotiating agents. For example, the NegoChat-A agent [276] leverages this bias at the beginning of the negotiation process by presenting an offer which it does not expect the human to accept. Yet, the offer serves as an anchor for further negotiation offers. We provide a detailed overview of the methods used in the design of human-interacting automated negotiation agents in Section 4.3.

- **Bandwagon effect**, also known as **herding**, refers to the tendency of people to do (or believe) things because many other people do (or believe) the same things. The phrase "jump on the bandwagon" first appeared in American politics in 1848 when Dan Rice (1823–1900), a famous and popular circus clown of the time, used his bandwagon and its music to gain attention for his political campaign appearances. The idea underlying the effect is that decisions or beliefs spread among people with the probability of any individual adopting it increase as more people "jump on the bandwagon." This often happens regardless of the initial belief or preference of the decision-maker. The effect is often leveraged in recommendation settings, for example, by offering items to people accompanied with explanations such as "best selling" or "people like you also liked…," thus, to some extent, aiming to trigger the effect. From a multi-agent perspective, this effect can facilitate the design of agents that support people or act autonomously in voting systems, as people show clear biases toward voting for the leader of a poll or to their second-most preferred candidate as long as it receives many votes in the polls [305]. Voting settings are discussed in Section 4.5 in more detail.

- **Loss Aversion** refers to people's tendency to prefer avoiding losses over acquiring equivalent gains. Loss aversion is the reason we see phrases like "last chance" or "Hurry! Only 2 left!" in marketing campaigns so often. Loss aversion can be considered as "Playing Not to Lose" as opposed to "Playing to Win." Loss aversion can explain many everyday phenomena such as the unwillingness of many people to sell their house or stocks for less money than they paid for them. In [345], the authors show that automated agents can use debiasing techniques to partially decrease the impact of loss aversion, and thus, potentially, improve the agent's performance. From the opposing perspective, Gunaratne et al. [129] show how loss aversion can be triggered in decision-makers by manipulating the information presentation by an automated agent. Note that loss aversion is also part of the **prospect theory** described in Section 3.1.6.

- **Default bias**, which is also known as the **status quo bias**, is the tendency of people to opt for the default supplied option. A default usually refers to that option which a decision-maker will receive if she does not make an active choice or if she avoids taking the time to consider or adopt an alternative state to the status quo. The default bias can be seen in religion. More than 90% of religious people belong to the religion of their birth, namely, the default religion offered to them. From a design perspective, leveraging the default bias can have a far reaching effect on people's decision. For instance, many countries add to their driver's license application form an organ donation check box which applicants are asked to mark if they wish to be added to the organ donors database. Interestingly enough, countries which use a slightly different check box where applicants can mark the box if whey wish **not** to be added to the organ donors database have significantly higher organ donation rates: organ donation rates can vary from approximately 15% to approximately 95% with this small change [164]. Similar results are shown in web users agreeing to receive e-mail advertisements [163]. From an agent design perspective, Lee et al. [198] show that the default bias can be a powerful tool for automated persuasive agents (in their experiment, a persuasive robot) to propel people toward self-beneficial behavior. By placing a healthy snack in the "default" location on the robot's tray (compared to an unhealthy snack placed in a non-default location), human subjects significantly chose the healthy snack more often than under controlled conditions, thereby promoting a healthier diet.

- **Framing** is the phenomenon where people may make different decisions based on the same information, depending on how that information is presented. Different types of framing effects have been identified: including risky choice framing (e.g., the risk of losing 10 out of 100 lives vs. the opportunity to save 90 out of 100 lives), attribute framing (e.g., beef that is 95% lean vs. 5% fat), and goal framing (e.g., motivating people by offering a $5 reward vs. imposing a $5 penalty) [199]. From a design perspective, an agent can adjust the framing of information without changing their content, in order to influence decision-makers in a desired way. For example, in the field of human-robot interaction, Souza et al. [299] show that by framing the robots' information, human operators can significantly change their decisions and commands to the robots.

Note that it is not completely and immediately clear how one should apply the above general biases, separately and together, to agent technologies in a given setting. Therefore, in many cases, extensive domain knowledge and the incorporation of expert-defined heuristics may be required. Nevertheless, a few more holistic theories have emerged by combining several of these (and other) biases. Most notable is the *prospect theory*, as described in Section 3.1.6. Before describing the theory, we will complete our discussion on the use of behavioral biases by discussing domain-specific biases.

Domain-Specific Biases

In the previous section, we devoted our attention to general biases which have been found across different environments and people. However, in different fields, experts have identified many more biases and tendencies of people which stem from their domain of interest. To exemplify the use of this approach, consider the Virtual Suspect system [36]. The system is intended for inexperienced law enforcement personnel and is used to train them in interrogation strategies. As part of the system, an automated agent plays the part of the suspect and therefore has to simulate a human interrogatee. Ample research in criminology shows that suspects often display various biases during the interrogation which leads law enforcement personnel to develop different techniques for driving the interrogatee to tell the truth. These techniques are often depicted in popular TV shows such as *Law and Order*, *CSI*, *Criminal Minds*, and others. The developers of the Virtual Suspect system were able to adequately predict human suspects' decisions (e.g., deny, come up with an alibi, etc.) based on expert-based modeling of their emotional state and leveraging expert-knowledge in criminology.

Another example is the modeling of human geographical movement. Experts in the field, inspired by Newton's law of gravity, have hypothesized that the flow of individuals between two locations decreases exponentially with the physical distance thereof. These so-called "gravity models" and their subsequent models have a long tradition in urban planning and have been used to model and intelligently adapt to a wide variety of social phenomena such as human migration, inter-city communication, and traffic flows [231]. Agents may utilize these models in different ways, for example by exploiting the predicted movement of people in order to improve caching of multimedia files in distributed systems [284].

Both of the above examples, as well as many others, are *analytical models* which are crafted by experts for a *specific domain and setting*. In many cases, the very existence of these models remains within the domain of interest (e.g., criminology or geography) and thus requires extensive domain knowledge on the part of the agent designer.

3.1.6 PROSPECT THEORY

Prospect Theory [172], proposed by the 2002 Nobel laureates *Daniel Kahneman* and *Amos Tversky* (1937–1996), offers a descriptive theory of how people actually make decisions. This theory is one of the most widely cited theories in economic literature.

Prospect Theory consists of two main novel notions compared to traditional methods: First, the theory assumes that people derive utilities from *gains and losses* which are measured relative to some *reference point*, rather than from the resulting outcome of the decision. Namely, a utility function should receive as an argument the relative change as compared to the decision-maker's reference point, denoted as $\Delta(x_i)$ where x_i denotes a possible outcome of her decision-making. We assume $\Delta(x_i) \in \mathbb{R}$ where positive numbers indicate gains and negative numbers indicate losses. The utility function is assumed to follow the loss-aversion bias, namely, the "pain" of losing α dollars should outweigh the "pleasure" of gaining α dollars. Furthermore, the utility function is assumed to be concave in the region of gains but convex in the region of losses. This element of prospect theory is known as diminishing sensitivity. This means that by replacing a gain (or loss) of α dollars with $\alpha + 1$ dollars the marginal change in utilities is smaller than the replacement of β dollars with $\beta + 1$ dollars if $\alpha < \beta$. This notion is illustrated in Figure 3.4. Second, the theory assumes that a human decision-maker is using a non-linear probability weighing. Namely, people do not weigh outcomes by their objective probabilities but rather by transformed probabilities, called weighted probabilities or decision probabilities, often denoted using π instead of p. The weighing is done such that people overweigh low probabilities and underweigh high probabilities.

The combination of both notions together suggests that people reason over potential *prospects* instead of potential *expected utilities*. The prospect of a decision option σ is defined, using our notation from Chapter 2, as

$$Prospect(\sigma) = \sum_{x_i \in X} \pi_\sigma(x_i) \cdot u(\Delta(x_i))$$

as opposed to the expected utility of σ which is defined as $EU(\sigma) = \sum_{x_i \in X} p_\sigma(x_i) \cdot u(x_i)$.

It is common to use a utility function of the form

$$u(\Delta(x_i)) = \begin{cases} \Delta(x_i)^a & \text{if } \Delta(x_i) \geq 0 \\ -\lambda(-\Delta(x_i))^b & \text{otherwise} \end{cases}$$

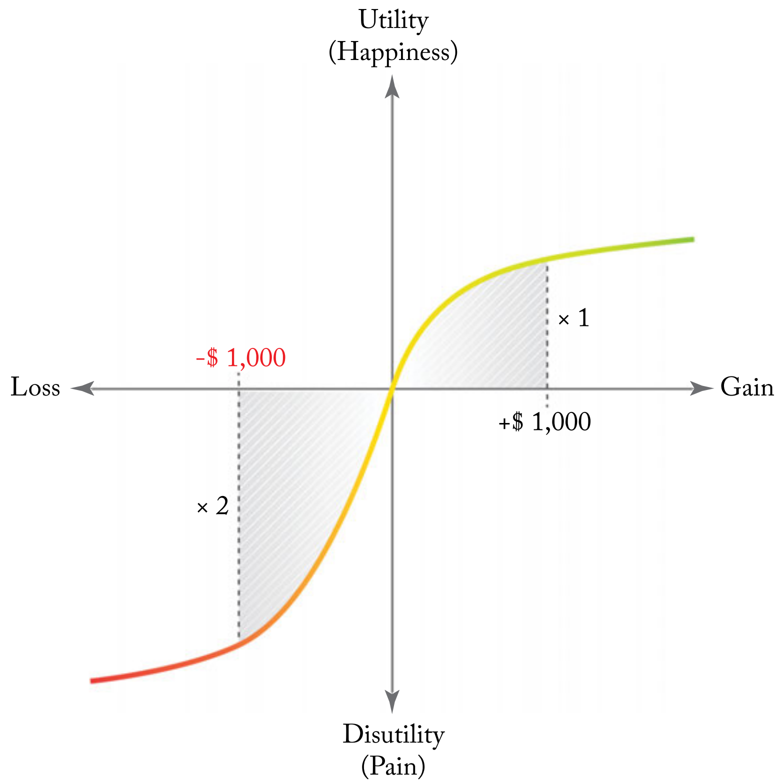

Figure 3.4: Illustration of the non-linear weighing of gains and loses underlying the Prospect Theory.

as proposed in the original paper introducing the theory.

The above utility function consists of the coefficient of loss aversion (λ), the gain satiation coefficient (a), and the loss satiation coefficient (b). It is generally assumed that the a and b parameters are in the $(0, 1)$ range and that $2 < \lambda < 4$. It is further common to assume that $a = b$. Tversky and Kahneman suggested that $\lambda = 2.25$, $a = 0.88$, and $b = 0.88$.

The transformation of probabilities to decision weights is commonly assumed to follow

$$\pi_i = \frac{p_i^c}{(p_i^c + (1 - p_i)^c)^{\frac{1}{c}}}$$

as done in the original paper.

The nonlinear transformation of objective probabilities to decision weights is expressed via the c parameter. It is common to use *two* probability weighting functions, one for the probabilities associated with "gain" outcomes (with parameter c_+) and one associated with "lose" outcomes (with parameter c_-). Tversky and Kahneman found that $c_+ = 0.61$ and $c_- = 0.69$. It is common to find c parameters estimated in the $[0.5,1)$ range.

The prospect theory was later extended into a Cumulative Prospect Theory [318]. The main modification to Prospect Theory is that *cumulative probabilities* are transformed, rather than the probabilities themselves. We refer the interested reader to [95] for a discussion of the empirical difference between the two.

From agent design perspective, it has been shown that prospect theory-based agents can better simulate market behavior than expected utility maximizing agents [77]. This naturally fits empirical evidence showing that the financial professionals follow prospect theory principles more often than expected utility maximization ones [2]. Using Prospect Theory, an agent can order presented alternatives to human decision-makers in order to drive them to make a desired investment [21].

Note that the above expert-driven models are the most well-known and most deployed methods in agent design. Naturally, many other expert-driven models exist. The interested reader may consider other prominent models such as **regret minimization** [157], **strategy elimination** [69], **Stackelberg Reasoning** [68], and **Team Reasoning** [26].

3.1.7 UTILIZING EXPERT-DRIVEN MODELS

In many domains, expert knowledge is usually **qualitative**. Namely, expert knowledge is often formulated as general criteria or rules (as exemplified in Section 3.1.5) and is not available in a mathematical formulation, needed for most agent designs. Extracting domain knowledge from human experts (also known as **knowledge acquisition** from human experts) in a way that can be easily used by an automated agent may be very complex [236]. It is rarely the case that a domain expert can "pin-point the true decision-making model" and provide exact rules, formula and parameters for people's decisions. As a result, two knowledge acquisition techniques have been in use.

1. **Interviews and Questionnaires**, by which a designer elicits domain knowledge in the form of decision-makers' utilities, features, heuristics, or decision-making characteristics with which she can deploy an expert-driven model. For instance, to develop the *ARMOR* [251] agent, the designers conducted intensive interviews with security experts, trying to reason about the utilities of potential attackers from attacking different airport terminals. Furthermore, they were able to conclude that attackers are (or can be assumed to be) expected utility maximizers. As a result, a game-theoretical approach was successfully deployed by the designers.

2. **The designer becomes a domain expert herself**, relying on introspection to articulate the domain knowledge in a mathematical way. For instance, in order to predict human argumentative decisions (see Section 4.4), the first author had to become an expert in the computational argumentation field.

The two methods often suffer from a significant gap between the qualitative expert knowledge and the resulting mathematical formulation thereof. As a result, the process of articulating expert knowledge in an expert-driven model is usually iterative: an expert-driven model is constructed based on qualitative expert-knowledge and then tested with a few real-world or synthetic examples. The qualitative knowledge is then reformulated to account for non-intuitive results and the process is repeated.

Tuning expert-driven models' parameters may be done using expert knowledge or a *limited* amount of data. As a result, expert-driven models mostly require little to no contextual data, making them especially favored in domains with established behavioral assumptions and in domains where obtaining data imposes major costs or when relevant data is scarce.

Note that use of expert knowledge can also be utilized using a **hybrid approach** as we will discuss in Section 3.3.

3.2 DATA-DRIVEN PARADIGM

More than 60 years ago, the psychologist *Paul Meehl* (1920–2003) put forward an "outrageous claim"—that mechanical, *data-driven algorithms* could better predict human behavior and decisions than a trained clinical psychologist, and with much simpler criteria [216]. Paraphrasing on Meehl's claim, he claimed that expert-based models, normally deployed by human experts, are expected to be outperformed by data-based models. To a certain extent, he was right.

A **data-driven model** for predicting human decision-making is a machine learning-based model which is trained using contextual data about people's decisions, and realized to a given setting. Today, most computer scientists predict people's decisions through machine learning techniques.

In the following, we assume that the reader is familiar with the basic terminology in machine learning. We recommend the following texts for a complete and through overview of the machine learning field [218, 220, 293]. For a more practically oriented introduction we recommend [82, 349]. For an easy-access evaluation of different machine learning algorithms, which requires no programming knowledge, we recommend using the **Waikato Environment for Knowledge Analysis (Weka)** [331]. A recent extension thereof is the Auto-Weka 2 [187] which automatically searches through the space of WEKA's learning algorithms and their hyperparameters.

3.2.1 MACHINE LEARNING: A HUMAN PREDICTION PERSPECTIVE

Generally, the basic premise of machine learning is to build algorithms that leverage statistically based methods to produce a prediction model. In our context, the resulting prediction model should produce a prediction as to a human's decision in a (most commonly, unseen) decision-making setting.

A decision-making setting is commonly represented as a vector of features $\vec{x}$ that describe the decision environment and, possibly, the decision-maker. For example, if we wish to predict the time at which a worker would leave the workplace today, $\vec{x}$ should consist of features that describe the decision-making environment such as the day of the week, the company type (e.g., high tech, public office), etc. Additional features which describe the decision-maker may also be included, such as the time the worker started the work day, the time at which the worker left the workplace the previous day, whether the worker uses public transportation, etc. The process of describing a decision-making setting using a set of features is often called *feature extraction* (also known as *feature engineering* or *feature construction*). Naturally, features are intended to be as informative as possible in order to facilitate the subsequent learning and generalization of contextual examples. Recent advances in *deep learning* may relieve some of the feature extraction burden in many cases, yet it imposes other limitations as we further discuss in Section 3.2.2.

Machine learning algorithms are typically classified into three broad categories.

1. **Supervised learning:** Supervised learning algorithms require a set of **training examples**, each consisting of a past decision-making setting extended with the actual choice made by a human decision-maker (i.e, the examples are **labeled** with the human decision). Continuing with our previous example, a supervised learning algorithm will require a set of training examples consisting of past decisions made by a specific worker or a group of different workers, each with the correct *label*—the decision that the decision-maker has made in the setting in question. Given the labeled training set, a wide variety supervised learning algorithm such as **decision trees** [259], **deep neural networks** [124], or **Support Vector Machine** (SVM) [70] may be applied. The goal of these models is to approximate the mapping of decision settings to actual decisions using different underlying assumptions. Note that the goal is not to "reverse engineer" the human decision-making process but rather to approximate the decision outcome. An example for the use of this approach in agent design is predicting which arguments a person is likely to put forward during an argumentative dialog in order to provide beneficial arguments for that person to use or for persuading the person to change her mind. We discuss both examples as part of the Argumentation domain in Section 4.4.

2. **Unsupervised learning:** Unsupervised learning algorithms do not require the training data to include labels. Unsupervised learning algorithms such as **K-means** and **Hierarchical clustering** [337] are used for modeling the underlying structure or distribution in the data in order to learn more about the data (e.g., discovering patterns) or as a means toward constructing a prediction model (e.g., feature construction). Continuing with our previous example, if the company did not keep track of when different employees left the workplace, the encountered decision-making settings can reveal interesting insights that may be used later for providing a prediction. For example, *clustering* the employees into types (most commonly known as clusters) or deriving *association rules* such as "employees with a large workload tend to come earlier to work" may provide additional strength to a future predic-

tion model. A practical example comes from the automotive industry, which we discuss in Section 4.6. By dividing drivers into clusters according to age, gender, etc. automotive functionalities such as adaptive cruise control can be better configured to the driver's satisfaction by automated agents. We discuss this and other automotive applications which use the prediction of human decision-making in Section 4.6.

3. **Reinforcement learning:** Reinforcement learning is concerned with how agents ought to take actions in an unknown environment so as to maximize some cumulative reward over time. Namely, popular reinforcement learning algorithms such as **Q-learning** and **Rmax** [304] are developed and deployed by agents as to learn an optimal mapping from states in the environment to desired actions. The use of reinforcement learning for the prediction of human decision-making takes one of two forms. (1) Approximating and predicting how human decisions evolve in reinforcement learning environments (e.g., repeated play of an unknown game). Erev and Roth [87] found that reinforcement learning models outperform baseline models in modeling and predicting how human decisions change and adapt in a broad range of repeated economical environments, (2) By using a variant of reinforcement learning called **inverse reinforcement learning**, an agent can model human decision-making from a set of human-generated demonstrations (i.e., training examples) [227]. Ziebart et al. [348] showed that inverse reinforcement learning can be very effective in inferring destinations and routes of human drivers based on partial driving trajectories. Inverse reinforcement learning is strongly related to the notion of **learning from demonstrations** [65], in which a person seeks to teach an agent (usually a robotic agent) how to operate in a complex environment. Using our terminology, in most cases, the person shows the robot what she would do in a given decision-making setting and expects the robot to adequately predict and carry out the action she would take in future decision-making settings.

3.2.2 DEEP LEARNING—THE GREAT REDEEMER?

The recent successes of **deep learning** [197] have demonstrated that predictive accuracy of data-driven models can often be considerably enhanced and that feature extraction can often be relieved. These successes have been mainly demonstrated in computer vision, speech recognition and natural language processing. The major advantage of deep learning compared to standard supervised learning lies in deep learning's ability to automatically learn internal representations of the necessary processing steps such as detecting useful features with only "raw" input as the training signals. Thereby, deep learning can significantly relieve many of the feature construction concerns.

The most impressive and world-renowned use of deep learning for predicting human decision-making is found in the game Go, as described in Example 3.4. The use of this pre-

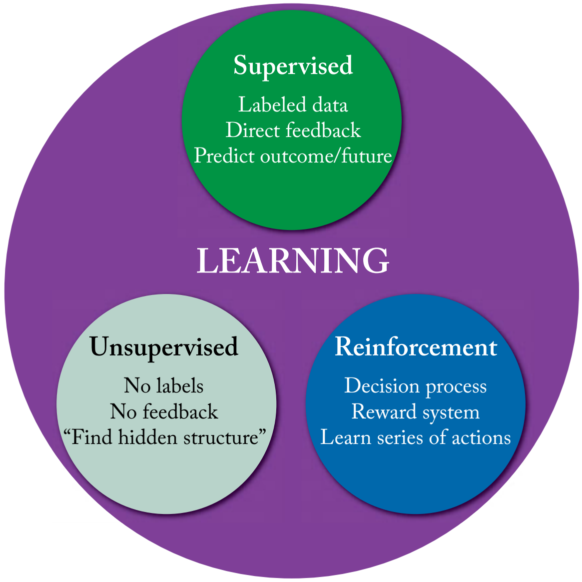

Figure 3.5: High-level view of the three main strands of Machine Learning.

diction model allowed the Go playing agent **AlphaGo** [295] to beat expert human Go players, which marked a significant milestone in AI history [208].

Example 3.4 Due to its enormous state space, the game of Go cannot be solved using exhaustive search. Instead, Silver et al. [295] presented a novel approach to mitigate this problem which *heavily* relies on the prediction of expert human Go players' game decisions (i.e., game moves). First, the authors trained a deep network prediction model on a set of 30 million game moves made by strong human players to predict what would be an expert's next move given a board position. The prediction model was extremely accurate, achieving an impressive 57% accuracy.[1] Using this prediction model, the authors then use a reinforcement learning approach to adjust the learned model toward the goal of winning games rather than maximizing prediction accuracy. In addition, another model is trained for estimating the chance of each player winning the game given a board position. These models are integrated within a *Monte Carlo Tree Search* (MCTS) method which uses the prediction of expert moves in its expansion and rollout phases.

Note that the recent achievements in the game of **Heads-Up Texas Hold'em Poker**, most notably the **DeepStack** [221] and **Libratus** [49] agents which conclusively defeated human professional players [117], *do not* explicitly use the prediction of human moves. We believe that this is a research direction worth exploring.

[1]Why should 57% accuracy be considered impressive in the game of Go? This question is left as Exercise 3.17.

Similar to Example 3.4, the prediction of human decision-making was also successfully deployed for the task of predicting human decisions in unrepeated two-player normal form games [140]. The authors show that a deep learning architecture can outperform many of the expert-based and hybrid approaches (such as the ones we discuss in Section 3.1) in a very convincing manner.

Impressive successes were also demonstrated for the prediction of human decisions from a visual perspective. For example, deep learning was successfully deployed to predict whether two individuals in a video will hug, kiss, shake hands, or slap five based on more than 600 hours of unlabeled videos from YouTube[2] [321]. In a similar spirit, human drivers' steering decisions were successfully predicted using a deep neural network that maps raw pixels from a single front-facing camera directly to steering commands based on dozens of hours of unlabeled steering behavior [38].

To date, deep learning seems to be most powerful when a lot of training examples are available and especially when there is some feature locality in space and time such as evident in videos, images, voice and many games such as Go and Atari. Recall that in addition to enhanced prediction qualities, as discussed before, an important advantage of deep learning is that one does not have to worry about the feature engineering. However, all that glitters is not gold. *At least not yet*. To date, large human decisions corpora are not widely available and collecting such big datasets can be highly complex (see Section 3.2.3). However, we speculate that this will change in the near future given the technological advancements in wearable devices, smart phones and online platforms. Furthermore, a deep learning infrastructure needs to be designed such that it could be trained in reasonable time and also, potentially, leverage domain knowledge on human decision-making. Thus, deep learning algorithms usually require much more experience on the designer's part as compared to "old fashioned/off-the-shelf" methods such as decision trees and SVMs, which can be found in *Weka*. Also, in many cases, it is the lack of self-explainability of the deep learning models which makes them disadvantageous compared to simplistic, yet possibly less accurate, models which are easier to understand, explain and maintain (e.g., Bayesian networks, K-nearest-neighbors, etc.). Law and medicine are the first of many fields that require explainable AI developments.

3.2.3 DATA—THE GREAT BARRIER?

Google's Research Director *Peter Norvig* is often quoted for saying "We (Google) don't have better algorithms, we just have more data," emphasizing the great importance of data in today's world. Machine learning is naturally based on data.

For predicting human decision-making, contextual training data on past decision-making settings could be valuable. However, obtaining contextual data raises two main issues: (1) obtaining **enough** data; and (2) obtaining the **right** data.

[2] https://www.youtube.com/

Obtaining data on human decision-making is very complex compared to standard prediction tasks. For example, in order to predict the weather, historical weather data can be easily obtained from most national meteorological services, thus obtaining quality data should not pose a great concern. However, data on human decision-making is usually not widely available and unfortunately, even if past data on human decisions is available, it does not necessarily fit the specific setting and assumptions required for a new prediction task. Altman et al. [6] showed that machine learning methods can be easily deployed to predict a person's decision in game-theoretic settings based on her decisions in earlier games. Naturally, the proposed model requires data that identifies a person across different games. Identified repeated-interaction datasets are usually not available.

Before one turns to collecting data independently, a highly complex and resource consuming task as we discuss next, an extensive search for available data should be performed. Several open source data repositories are available, such as Harvard's Dataverse,[3] Zenodo,[4] Open ICPSR,[5] Kaggle,[6] and UCI Machine Learning Repository,[7] to name a few. Several government entities, such as the U.S Government[8] and public organizations, are following suit. These and others enable the sharing, preserving, citing, exploring, and analyzing of previously collected (and usually labeled) data. In general, there is an ever-growing multidisciplinary recognition of the benefits resulting from publishing data [188, 195, 242]. Nevertheless, many researchers do not publish their collected data on open source repositories. In many cases it is worthwhile to visit a researcher web-page to search for specific data or (kindly) ask a researcher for the data by e-mail. Note that avoiding making data public can be explained, in some cases, with ethic concerns such as anonymity and privacy. In such cases, the above arguments do not apply.

To date, many researchers who investigate human decision-making turn to collect their own data. Given the discussion above, that is unfortunate yet understandable.

Data collection is the process of gathering and measuring information on targeted variables. Ideally, it would be highly valuable to use an **observational approach** where human decisions are gathered in the *real world*, namely, in decision-making settings which were neither triggered nor artificially manipulated by the researcher. It is self-evident that the observational approach is the most ecologically valid approach to model and predict human decisions "in the wild." With the advances in different technologies such as on-line shopping platforms, smart phones, wearable devices, and others, it seems that observational data on human decision-making is becoming increasingly prevalent. For example, smart phone navigational applications such as WAZE[9] or Google Maps[10] can observe driving decisions (e.g., take route *a* or route *b*

[3] https://dataverse.org/
[4] https://zenodo.org/
[5] https://www.openicpsr.org/
[6] https://www.kaggle.com/datasets
[7] http://archive.ics.uci.edu/ml/
[8] https://www.data.gov/
[9] http://www.waze.com
[10] http://maps.google.com

given the time, congestions, etc.) with and without the system's recommendations. Note that once the navigational application provides a suggested route, it is most likely to influence the user's decision, violating the observational premise. Hence, the major limitation of the observational approach.

When a researcher manipulates or artificially creates decision-making settings for data collection she departs from the observational approach and turns to an **experimental approach**. Using an experimental approach, data is gathered from people who willingly agree to participate in the data collection, who are fully aware that they are subject to artificial decision-making settings and, most importantly, *the decision-making settings faced by the experiment participants are governed by strict data collection protocols and standards*. **The correct design of data collection from people is a vastly under-appreciated art.** We review this process below.

Phase 1 – Recruiting Human Participants

Data collection normally starts by recruiting human participants, right? *Wrong*. Before recruiting human participants, an **Institutional Review Board** (IRB) must approve the proposed recruitment protocol. The purpose of the IRB is to assure that appropriate measures are taken to protect the rights and welfare of human participants. This includes how their anonymity is kept, how they will be compensated for their participation, etc.

Armed with an IRB-approved recruitment protocol, researchers may start recruiting participants. A very common practice in academic research is to recruit students and faculty as study participants. While this practice is acceptable, in many cases its worthwhile to go to the extra effort of recruiting non-academic participants. This may be done using crowd-sourcing platforms such as Amazon Mechanical Turk (AMT)[11] or by posting ads using social media, etc.

Since there is an infinite number of possible participant groups one can recruit, it is customary to recruit an (approximately) equal number of men and women and strive to achieve a diverse group in terms of age, education level, etc. The number of participants in data collection varies significantly in the literature and there does not seem to be a "magic number." The number of participants is usually determined in an ad-hoc fashion, as we will describe soon. As a general rule of thumb, it is customary to recruit at least 15 participants for each examined condition in the data collection. **Note that sampling the decisions of a small and similar group of people will likely result in poor prediction accuracy in real-world deployment.**

It is of the utmost importance to get the participants' written informed consent before continuing with the data collection.

Phase 2 – Pre-questionnaire

In some cases, a researcher may want to collect additional information using psychological or behavioral questionnaires. For example, the *Minnesota Multiphasic Personality Inventory* [142] may provide significant insights into adult personality and psychopathology. While such ques-

[11] `https://www.mturk.com/`

tionnaires may provide very valuable information *on the participants of the data collection*, one should consider if this information would be available in the future when the intended agent is deployed. For example, consider an autonomous car which seeks to predict a driver's driving style (e.g., aggressive, safety-first, etc.). It is reasonable to speculate that obtaining a driver's personality traits in data collection will assist in predicting her intended driving style. Nevertheless, one may wonder if such psychological information would be available for a new driver who will use the autonomous car in the future. Psychological questionnaires are generally very resource consuming, and are rarely administered in AI.

Phase 3 – Presenting Decision-Making Settings

In this phase, the participant is presented with decision-making settings and is required to make choices. But, **how do we choose which decision-making settings to present to a participant?** The choice will bear a significant effect on the quality and diversity of the collected data which in turn will effect the trained prediction model and the intended agent that will use the prediction model. Consider the following **true** experience.

Example 3.5 Can we build an automated negotiator, with no expert designed rules, that would do well *across cultures?* This is the main question of [136]. Participants from three countries were recruited—the U.S., Israel and Lebanon. Due to logistical constraints (time differences, access to participants, etc.), asking participants of one group to negotiate with participants of another group was too complex. Therefore, an *observational* approach was adopted where participants were randomly coupled within their group and were asked to negotiate over a given topic as they would do in "the real world" (i.e., with minimal influence from the experiment designer). Note that in the examined settings people negotiated repeatedly with their partners, *but agreements were not enforceable*, as with many verbal agreements in the real world. Using the collected data, three machine learning models were trained, one for each culture. Thus far, everything went according to plan. However, once an agent was designed based on the data collected in *Lebanon*, the agent turned out to be very "nasty"—it promised one thing and did another. Namely, it did not keep its commitments throughout the negotiation process. The agent performed very poorly in preliminary testing with human participants who, in turn, started acting "nasty" as well.

So, what went wrong in Example 3.5? It turns out that collectivist societies such as Lebanon are more homogeneous and display less variance in the extent to which they fulfill commitments [112]. Specifically, in data collected from Lebanon, virtually **everybody** kept all of their commitments. So, what would an expected utility maximizing agent do if it assumes that its human counterpart would always keep her commitments? You guessed it, it would not keep its own commitments.

The above example highlights the problem of generalizing *observed* human decisions which may be *biased*. Specifically, the collected data from Lebanon did not reflect all realistic settings and thus the resulting prediction model was highly biased. A common strategy for mitigating this problem is acquiring additional data on "hard-to-reach" or *off-the-path* parts of the decision-making space via experiments. However, this may not be as simple as it may sounds.

Consider the following task faced by judges everyday: deciding whether a defendant will await trial at home (bail) or at jail (remand). The judge's task hinges on the prediction of what a defendant would do if he were to be released until trial (e.g., commit another crime or not), a prediction task which can be significantly improved using data-based methods [184]. However, available data on defendants' decisions suffers from the **selective labels problem** [192]. Namely, one would know what a defendant would have done if sent to await trial at home only if she is *actually* sent to await trail at home. Naturally, judges' decisions are made during the *data collection phase* and those are a *consequence of our existing predictions*. This is very likely to lead to a self-fulfilling prophecy.

In order to ensure that the intended model provides adequate predictions for all settings (also for those which are unlikely, yet possible), it is important to collect data on a wide range of decision-making settings, not only the most common and obvious ones. Ideally, we would want to observe people's decisions for every possible decision-making setting and use the observed data with minimal (if any) generalization. However, this is usually infeasible. Think of a judge who would release *all defendants* only to learn which defendants would commit a crime, or a behavioral economist who asks experiment participants to make decisions as though they were under all possible two-player normal form game conditions in a questionnaire.

The common technique to overcome the above limitations is somewhat similar to the Expectation-Maximization (EM) technique in statistics [78]. In an iterative process, the designer first speculates a possible division of the decision settings space into different **choice zones**. It is advisable to rely on domain knowledge in this phase. Namely, we divide the possible decision-making settings into clusters or types based on our *assumed* human behavior in these clusters. We then articulate *representative* decision-making settings for each cluster *as well as randomly selected settings when applicable*[12] and obtain participants' responses thereof. Then, using post-hoc analysis of the learned model, the designer seeks apparent defects in the learned model and revises its assumed *choice zones*, usually using hold-out and cross-validation analysis. Specifically, the designer should examine the model using a held-out subset of decisions which the trained model did not encounter in training. Another common, yet non-cost-efficient, way to examine the adequacy of the learned model is to deploy it using an automated agent in preliminary testing (usually with very few human participants, if any). Then, the designer should carefully examine the agent's behavior and revise its assumptions. The iterative process then continues for "as many times as needed." **Remember, a prediction model is only as good as its agent performance.** Namely, a prediction model is considered useful if and only if the agent that uses it performs well according to a pre-defined measure (usually measured in terms of received utility). Specifically, there is no "gold standard" for what constitutes a good prediction model, yet it is common to expect more accurate prediction models to bring about better agent

[12] It is debatable whether or not releasing a defendant at random is ethically acceptable.

performance. As a general rule of thumb, for a reasonable prediction task, two or three iterations would suffice. Furthermore, from our experience, extremely high prediction accuracy of a model ($\approx$ 90) on human decision-making usually indicates some error or **overfitting**. Perhaps the most common error for novice researchers and practitioners is using the evaluation set as part of the training data.

This phase of the data collection is the most resource consuming. As a result, it is advisable to take the time and effort to organize the experiment protocol properly in order to ensure that the right data, and enough of it, are obtained. This phase is often called **experimental design** or **design of experiments** in the literature. This phase is also related to the active learning paradigm of machine learning [290], although active learning is rarely applied for human decision-making.

To complete the story of Example 3.5, in the second iteration of Phase 3 participants from Lebanon were asked to negotiate with an automated agent that was programmed to be significantly less reliable when fulfilling its agreements compared to human negotiators in Lebanon. The obtained data enabled more diverse negotiation behavior and decisions and thus resulted in a better prediction model that, in turn, translated into better automated negotiators.

Phase 4 – Post-questionnaire

Similar to Phase 2, in most cases it is useful to administer a questionnaire following Phase 3. In a **post-questionnaire**, participants are asked to rate various aspects of Phase 3, focusing on key evaluation questions to boost the reliability of the data collection. These may include questions as to how clearly the decision-making settings were presented to them in the data collection. Post-questionnaires tend to be short in order to reduce the amount of time participants need to complete them, and therefore increase the response rate and quality. Guidelines for efficient post-questionnaire design are available in [311].

In some cases, a researcher may want to collect additional information for future reference. For example, the **Nasa Task Load indeX** (TLX) [139] is commonly applied to assess the perceived workload of the participant. An example for the use of Nasa TLX post-questionaires in data collection is given in [265] where, in order to construct a model of human operators' skills in multirobot tasks, participants were asked to operate a group of robots and their decisions were recorded. An automated advising agent was later constructed to provide advice for human operators and relied on the trained model. Participants' TLX scores, obtained from the data collection process, were later used as a baseline for evaluating the benefit of the advising agent.

3.2.4 ADDITIONAL ASPECTS IN DATA COLLECTION

Additional important aspects of data collection focus on maintaining the validity of the data and boosting its reliability and replicability. These aspects are not specific or unique to the data collection of human decision-making, hence for further reading we recommend referring to one of the many sources on data collection standards and best practices such as [44, 118, 234].

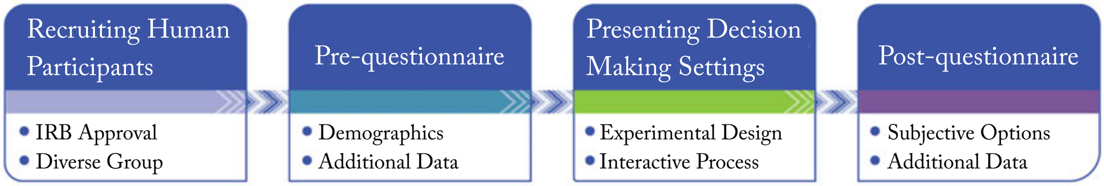

Figure 3.6: High-level view of the data collection process.

An under-represented method for data collection in agent design is the **think-aloud method**. The think-aloud method requires participants to say whatever comes into their mind as they make decisions during the experiment. This might include, but is not restricted to, what they are looking at, thinking and feeling. All verbalizations are transcribed and then analyzed. This may allow the researcher to better evaluate the adequacy of her experimental design. This may also provide a unique view into the decision-making *process* rather than only the *final decision*. The method usually requires manual annotation of the provided verbalization which may be very resource demanding.

Note that the data you, the reader, collect in your research can be of much use to other researchers in your field and in other fields as well. Unless the data is confidential, please share it using one of the open data-sharing platforms or your personal webpage where it can accessible to all.

3.2.5 THE DATA FRONTIER

Developing a data-driven (or, as we will soon discuss, hybrid) model necessitates the collection of contextual data. The traditional observational and experimental data collection approaches were discussed in Section 3.2.3. However, an important dimension to data collection is the use of novel data sources through which decision-makers or decision-making settings may be better characterized and understood.

Consider the following prediction task: predicting whether an individual is going to commit suicide. Today, suicide is a leading cause of death worldwide [330], making this a task of utmost social importance. Naturally, the experimental data collection approach is inadequate. However, how should one use an observational approach in this setting? Characterizing an individual using demographics or even historical mental health data may not suffice. Specifically, many of the factors that are expected to effect one's mental state are hidden from researchers. The use of new data collection techniques and the use of new data sources for the prediction of human decision-making is what we refer to as the **data frontier**.

Social media platforms such as Twitter can be used as effective data sources. Getting back to our prediction task from before, social media has been shown to be a platform for individuals to express suicidal thoughts, behaviors, and intent [232]. The use of hybrid techniques for detecting people at risk was proposed in the past (e.g., [1]), mostly relying on expert-articulated lexicons of "risky vocabulary" or changes in social profiles. The task is often referred to as suicidal screening. However, existing prediction models seem to lack real-world impact, mainly due to their poor ability to distinguish suicidal thoughts from other psychological conditions such as depression and their lack of effective intervention means [67]. To the best of our knowledge, to date, there exists no controlled study testing whether prediction tools of this sort can effectively reduce suicidal behaviors.

The use of social media to predict people's decision-making was also successful for the task of HIV prevention [338], intercepting terrorist acts [84] and predicting antisocial behavior [191], to name a few.

Additional data sources provide unique opportunities. For example, detecting non-verbal cues such as facial expressions can be used to enhance the prediction of human players' strategic decisions in games [246], psychological tests and neuroimaging can be effectively utilized to better predict criminal behavior, substance abuse and educational decisions (among other lifestyle decisions) [106], and the integration of multiple sensors can be used for predicting when one person would interrupt another [98]. Generally speaking, the combination and fusion of multiple data sources and sensors such as video, sound, bodily vital signs, and others can significantly enhance a prediction model's quality. Consider the task of predicting when a driver is about to perform a dangerous maneuver. The state-of-the-art solution provided in [160], capable of predicting drivers' maneuvers several seconds before they take place, uses an integration of video of the driver inside the car and the road in front, the vehicle's dynamics, global position coordinates (GPS), street maps, and others from which multi-modal data from both inside and outside the vehicle can be generated.

The challenge of using multiple resources lies in identifying relevant resources and fusing the different sources correctly. With the advancements in wearable devices, Internet of Things (IoT) and other technologies, this data frontier is expected to keep expanding in the future.

3.2.6 IMBALANCED DATASETS

The prediction of *rare* decisions poses a great challenge. For concreteness, let us consider the following problem instance; many companies invest significant efforts and resources in predicting which "trusted individual", also known as "insider", will choose to maliciously break security policy—commonly, by stealing or selling company secrets. This is often referred to as the **insider threat** problem [22, 233].

Most individuals in an organization are benign and would not engage in malicious activities. Only a small fraction of individuals choose to act maliciously. Due to the lack of extensive examples for the latter case, developing a prediction model, and especially a data-driven model, is extremely challenging. Let us consider a data-driven approach where each person is labeled "inclined toward acting maliciously" or not. A great deal of data sources are available at one's disposal, including the employee's e-mails, demographics, uploads, and downloads onto memory sticks, and much more. However, the set of training examples for which the label is "inclined toward acting maliciously" is usually **very** small, usually by several orders of magnitude, compared to the complementary set. This problem is often known in machine learning as the **imbalanced dataset problem**.

To understand the problem better, consider the following prediction model: for any individual, regardless of any information, predict that she will not engage in malicious behavior. If non-malicious individuals consist of 99.9% of the employee population, the proposed model would have an expected accuracy of 99.9% just by saying that no one is malicious. This is clearly a problem because many machine learning algorithms are designed to maximize accuracy.

To overcome the challenges associated with the imbalanced datasets, several machine learning techniques have been proposed as surveyed in [143]. The most popular techniques include collecting more data (which in many cases is infeasible or very expensive), changing the performance metric (e.g., by using the **receiver operating characteristic (ROC) curve** [94]), resampling the training set using **over-sampling** of the minority set or **under-sampling** the majority set and statistical techniques such as **SMOTE** [62]. Another common solution, which seems to be the dominant approach for mitigating the insider threat problem, is adopting an expert-driven or hybrid approach, thus relying less on data (e.g., [173]). We discuss the use of hybrid models in Section 3.3.

3.2.7 LEVELS OF SPECIALIZATION: WHO AND WHAT TO MODEL

In tandem to obtaining data, one must decide the level of specialization for the intended prediction model. Intuitively, if the prediction model is intended to address a modest task of predicting specific human decisions for a specific decision-making setting, the collected data process should reflect that. For example, if we wish to predict what $Dr.House$ will choose to eat for lunch it makes sense to collect data on $Dr.House$ at lunchtime over some period of time.

We discuss the two main dimensions of prediction models' specialization: *personalization* and *situationalization*.

Personalization

Personalization refers to the decision-maker's "weight" in the prediction model. Every human decision-making prediction model can be classified as one of three types.

1. **Generalized models** do not model the decision-maker within the decision setting at all, hence these models are generalized across all decision-makers in the training data. These models rely solely on the decision-making environment, namely, the decision options and circumstances. This modeling approach is useful when low variation in humans decisions is expected and when data is scarce.

2. **Semi-personalized models** consider both the decision-maker's characteristics as well as the decision-making environment's characteristics. The decision-maker's features such as gender, age and others are incorporated within the decision setting feature vector $\vec{x}$. This

allows for partial personalization of the model, as different decision-makers will potentially be associated with different predictions based on their features. *This is the most popular personalization level today*.

3. **Fully-personalized models** are constructed for each specific human of interest. Namely, a separate prediction model is constructed for each human and is trained based only on that specific human's decisions. This modeling approach is useful when very high variation between humans is expected and when substantial data on a decision-maker is available.

To exemplify the difference between the above personalization levels, consider Example 3.6.

Example 3.6 Automotive climate control systems are heavy energy consumers. As a result, an automated agent may be used to persuade drivers to save energy by suggesting more economically beneficial settings that would keep the driver comfortable. To that end, a prediction model of which climate control settings a driver would be willing to adopt if suggested to her is very valuable. In [23], a prediction model for which climate control settings a driver would consider acceptable was constructed as a *generalized model*. Specifically, all drivers would receive the same advice, which depends solely on environmental factors (e.g., outside temperatures). This was later extended in [266] into a *semi-personalized model*, where drivers' actions influence future prediction. Namely, each driver is represented using her past climate control settings and reactions to the agent's advised settings and thus the prediction of the driver's future decisions (to accept or reject a proposed climate control setting) is modified accordingly. Lastly, in [267], a *fully personalized model* was suggested such that each driver is represented using a unique parameterized model which is updated over time using the driver's interactions with the system. The resulting prediction model is fully personalized such that prediction and actions of driver x are unrelated and bare no effect on the predictions for driver y once the models are deployed.

Generalized models are especially useful for mitigating the **cold start problem**. The cold start problem is related to the sparsity of information regarding a decision-maker or a decision-making environment. For example, when predicting which item(s) a *new* human user, on whom we have no prior knowledge, may seek to purchase in an e-commerce platform, a generalized model based on tending items may prove valuable. An *extremely* successful deployment of a generalized model is *AlphaGo*, which we discussed in Example 3.4. Note that the prediction model used by AlphaGo does not seek to predict **a specific person's decisions** but rather the likelihood that an expert Go player will make a move. In principle, generalized models are very useful when the decision-maker's characteristics are assumed to have little to no effect over the decision. Consider a **clinical decision support system** [223] which is intended to recommend which tests to administer to a patient given her symptoms and medical records. A possible approach may include the prediction of what an experienced human doctor would do. Given the clinical decisions of several highly experienced human doctors, a generalized model of what an expert would decide should be more useful than a personalized model of what a specific *Dr. House* would decide.

Semi-personalized prediction models are very effective across a wide variety of prediction settings. These models leverage data from (possibly many) different people, yet use the decision-maker's characteristics to personalize the prediction, thus "enjoy the best of both worlds". For example, Last.fm[13] and YouTube[14] predict what songs or videos a user is likely to see (or would want to see) next by observing which songs or videos the user has listened to on a regular basis (the user's decisions) and comparing those against the decisions of other users. Common techniques deployed in this realm include **collaborative filtering** [46]. Semi-personalized models are also common in **targeted advertising** [254] which use semi-personalized prediction models based on user demographics and personal information as well as other (similar) people's past actions for deciding on which advertisements a human user would be most likely to click.

Fully personalized models have the benefit of modeling a specific decision-maker. As one would expect, adequate fully personalized models require extensive data on the decision-maker in question. Virtual assistants such as **Google Now**[15] strive to provide fully personalized prediction models. For example, say a human user drives to her gym everyday at 7pm. We would expect the virtual assistant to predict that and react accordingly (e.g., notify me on possible traffic congestions). Indeed, personalized virtual assistants do just that. Another example is the **DARKO** prediction model for a specific human's decisions in a home environment using a camera [263]. By following a video stream of a person at home, a personal robotic assistant can utilize personalized prediction models to allow better reactivity to the user's activity in shared environments [186]. In order to train a fully personalized prediction model one would need to collect data generated from a specific person, *a highly difficult task in most settings*. Today, only a few large cooperations such as Google and Facebook, as well as government authorities (presumably), hold such personalized data on people.

Situationalization

Situationalization refers to the decision environments (or situations) in which we seek to evaluate the intended prediction model. There exist three broad levels of situationalization which are defined with respect to the entire decision environment space of the task in question.

1. **Narrow models** are trained and evaluated on a single decision-making environment in the domain. Specifically, the prediction model is not intended for generalization across different environments, however it may be intended for generalization across different decision-makers. For instance, in an **ultimatum game** (Example 3.3) where an agent plays the proposer and a human decision-maker acts as the receiver, a prediction model may be useful to predict the probability that a human decision-maker would accept a specific offer,

[13] https://www.last.fm/
[14] https://www.youtube.com/
[15] https://www.google.com/landing/now/

for example the Nash equilibrium strategy offer. Several studies examined this question exactly, *while focusing on a specific offer or a few specific offers without the intention of generalizing across them*. An interesting implication of these investigations is through the prism of cultural differences; namely, by learning how individuals in different cultures may react to an offer. A survey of these efforts is available in [237].

2. **Broad models** generalize data across an abundance of decision-making circumstances yet do not cover the entire decision-making space. These models are intended for addressing a subset of decision-making environments which do not necessarily represent the entire space of possible decision-making environments. For example, consider the models listed in Example 3.6 which are designed to predict human decisions in climate control environments which are *only considered summer conditions*. Naturally, the trained model is not expected to perform well in winter conditions but it was not designed with that goal in mind.

3. **Holistic models** are intended for evaluation under *all* possible decision-making environments in the context of the task in question. Holistic models are expected to train on a very varied dataset which can reflect the possible heterogeneity in decision-making environments. For example, a model for the prediction of an expert **Go** player's moves (Example 3.4) is not limited to the confines of specific game positions.

See Figure 3.7 for a graphic illustration of the examples discussed in this section on the **personalization-situationalization space**.

3.2.8 TRANSFER LEARNING

In a traditional learning environment, such as supervised learning, if we intend to deploy a model for some prediction task in domain A (e.g., predicting one's investment decisions), we assume that we are provided with training data from the same domain, A. Namely, according to the model's situationalization, the model is expected to generalize the training examples for the same domain A alone. For instance, given an investor's history of investment, we could predict the likelihood that she will invest in a new stock. **Transfer learning** allows us to (re-)use data from domain A to a new, related domain B (e.g., predicting one's insurance decisions). Specifically, using transfer learning the designer seeks to transfer as much knowledge as she can from the source domain (domain A), from which training examples have been gathered, to our target domain (domain B). Using our example above, a designer may seek to transfer knowledge about the decision-maker and decision-making features from one's investment decisions to her insurance decisions, assuming the two domains are related.

Andrew Ng, chief scientist at Baidu and professor at Stanford, is often quoted as claiming that "transfer learning will be the next driver of machine learning commercial success" (*NIPS 2016 tutorial*).

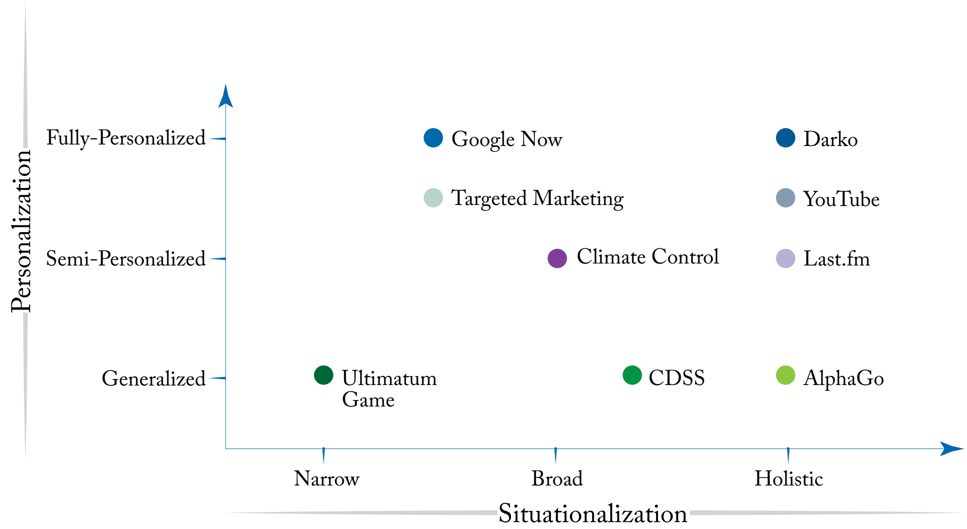

Figure 3.7: Examples presented in Section 3.2.7 depicted on the personalization-situationalization space.

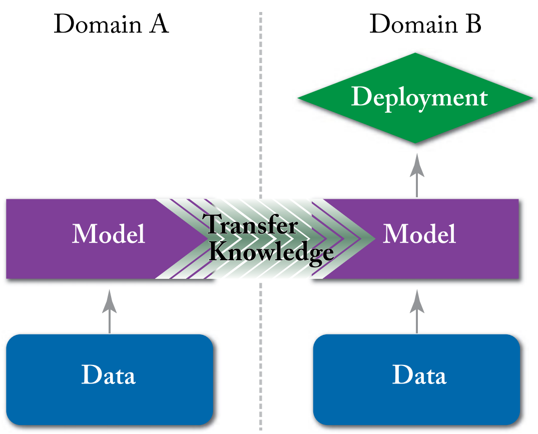

Figure 3.8: Transfer Learning illustration.

The ability of a prediction model to apply knowledge learned in previous domains to novel domains, which share some commonality, is inspired by a human being's ability to do so. Successful use of transfer learning can significantly reduce time and costs and can bring about a more holistic prediction perspective, spanning across various domains in one's life. This in turn can translate to better adoption of intelligent agents that provide a wide variety of services (e.g., *personal robotic assistants*).

Different techniques have been proposed in the literature to perform transfer learning (see [241] for a survey). These techniques are mainly motivated by having a robot transfer knowledge on how to accomplish one task to accomplishing a similar task [310]. However, to the best of our knowledge, **to date, no work has successfully transfered human decision examples from one domain to another in order to train a prediction model**. And not for lack of trying.

Consider the following example based on [269].

Example 3.7 The prediction of human decisions in argumentative dialogs (i.e., which arguments to present and when) can be very useful for argumentative agents (see a thorough discussion in Section 4.4). In order to predict what arguments a person is likely to present next in a dialog, *training data* of past dialogs *on the same topic* have been shown to be highly effective (more effective than using an expert-driven approach). However, in practice, obtaining past dialogs on every possible topic is infeasible. The common setting is one in which previous argumentative discussions on *different topics* are available. In this case, transfer learning could potentially be used to transfer the observed argumentative decisions from one set of domains to a target domain. Namely, given a partial discussion in the target domain, for which the agent has no prior dialogs from which to learn, a prediction is generated according to users' argumentative selections in *other domains*. Unfortunately, as discussed in detail in [269], the attempt was highly unsuccessful. Apparently, given conversations on the target topic, conversations on different topics (even from the same person) do not enhance the prediction accuracy. Moreover, when no conversations over the desired domain are available, it is better to use expert-driven heuristics rather than deploying transfer learning.

So what went wrong in Example 3.7? Seemingly, the authors did everything right: their feature-based transfer method, where each decision in one domain is mapped as a decision in the target domain, seems adequate. In addition, they tried to transfer argumentative decisions by using **the same person**, having each participant engage in three argumentative dialogs and transferring the person's decisions from the first two domains to the third. The authors provide two explanations for the unsuccessful use of transfer learning. These explanations also highlight the main difficulties in applying transfer learning when predicting human decision-making in general.

1. Different domains may drive significantly different decision-making behaviors. In Example 3.7 people may not use a cross-domain deliberation style—people might deliberate differently over different topics, depending on varying factors such as their knowledge of the topic, their attitude toward the discussed issue, etc. In general, given the multitude of factors that influence human decision-making, it is very hard to know in advance which domains should be considered similar with respect to human decision-making.

2. Even if a person follows a similar decision-making process across different domains, the factors that influence the decision-making in both domains may amount to completely

different decisions which are hard to transfer. In Example 3.7, the tested domains may have triggered different outcomes despite people following fixed argumentative patterns. Namely, in topics of personal interest (e.g., politics), people's decisions may have been different as a result of the emotional factors involved. Transferring the person's decisions to a less emotionally charged domain may prove problematic.

Due to the reasons above, the use of transfer learning in the prediction of human decision-making is mostly considered an open challenge.

3.3 HYBRID APPROACH

Hybrid prediction models combine methods from both the expert-driven and the data-driven paradigms. These models explore how these two paradigms can be combined to bring about better predictions than the two paradigms alone. The basic premise of the hybrid approach is the assumption that theoretical models can benefit from empirical evidence and data which conversely can be better leveraged using theoretical models. This premise in aligned with the call for combining theoretical game theory and experimental game theory for better predicting people,s decisions raised by the 1970 Nobel-laureate *Paul Samuelson* (1915–2009) [279].

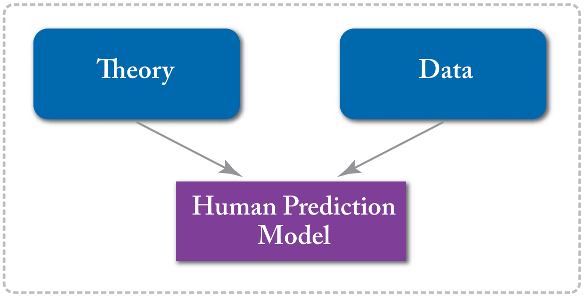

Figure 3.9: Hybrid models integrate expert-driven techniques (i.e., theory) with data-driven techniques (i.e., data).

The most common and straightforward technique to combine the two paradigms is by injecting expert-driven theories and models as *features* within a machine learning approach. We discuss this technique as well as others in the following.

3.3.1 EXPERT-DRIVEN FEATURES IN MACHINE LEARNING

Manual feature extraction is, by definition, an expert-driven process. However, beyond the basic representation of a decision-making setting, the use of well-known expert-driven theories and models as features within a machine learning approach is worth exploring. We discuss a few recent successful deployments of this approach.

Recently, such hybrid models were proposed for predicting human decision-making in decision theoretic settings under risk and ambiguity [230, 252]. These models use expert-driven models and behavioral biases such as **utility maximization** and **loss aversion** as additional features for representing the decision setting. This approach was shown to outperform expert-driven and data-driven approaches in a recent human choice prediction competition [86].

Another recent example comes from the negotiation domain (discussed in detail in Section 4.3). The prediction of what will be the next bid by a human interlocutor can be very useful in planning the negotiation course. **NegoChat-A** [276], the state-of-the-art human-agent negotiation agent, uses a hybrid approach by combining the expert-driven **Aspiration Adaptation Theory** [289] as additional features to the machine learning prediction model. These additional features have been shown to enhance the agent's prediction accuracy compared to a basic machine learning method (which does not include expert-driven features) or a pure Aspirational Adaptation Theory prediction approach. The prediction of human reciprocity in bilateral negotiation was also shown to benefit from using a hybrid model [108].

In predicting argumentative decisions of people in argumentative dialogs, the well-established **Argumentation Theory** (see [322] for a summary) can provide expert-based prediction, similar to the way Game Theory can provide predictions in appropriate settings. However, the predictive accuracy of Argumentation Theory models was shown to be very low in real settings [268]. Nevertheless, the state-of-the-art prediction model [269] has shown that argumentation theory (as well as other behavioral biases such as **confirmation bias** [229]—the tendency to search for, interpret, favor, and recall information in a way that confirms one's pre-existing beliefs) can be very effective in enhancing the prediction accuracy of machine learning models.

3.3.2 ADDITIONAL TECHNIQUES FOR COMBINING EXPERT-DRIVEN AND DATA-DRIVEN MODELS

Other techniques for combining expert-driven and data-driven models are used sporadically in the literature. For example, several models (both expert-driven and data-driven) may be combined into a single model using **ensemble methods**. In its simplest form one can consider combining several models using averaging (for continuous predictions), or *majority voting* (for classification of discrete data). Erev et al. [88] provide a real-world example of combining both expert-driven and data-expert models in such a fashion.

Other possible techniques include reducing the size of the decision-making space using expert-driven models. Specifically, the use of **focal points** [286] was shown to significantly reduce the size of the decision-making space. Focal points are a concept introduced by 2005 Nobel-laureate *Thomas Schelling* (1921–2016), identifying decisions that people tend to make in the absence of communication or prior knowledge just because they seem natural or special. Consider the following example: two people unable to communicate with each other are each shown a panel of four squares and asked to select one; if and only if they both select the same one, they will each receive a prize. Three of the squares are blue and one is red. Assuming they each know nothing about the other player, but that they each want to win the prize, then they will, reasonably, both choose the red square. The red square is hence a focal point. Focal points may play an important role in shaping the expectations of human decision-makers and hence can be leveraged in predicting them. In domains such as the job candidate selection [352] and "Pick the Pile" games [274], by assuming human decisions are heavily biased toward focal points, the number of the possible predicted decisions can be significantly reduced and, conversely, prediction accuracy is shown to increase.

It is important to note that, in many cases, expert-driven models are trained using data. For example, a **quantal response** model uses a *rationality parameter* which is usually learned from available data. While the line separating expert-driven and hybrid models is not crisp, it is our opinion that the mere introduction of data does not turn a quantal response model (or any of the models described in Section 3.1 for that matter) into a hybrid model. Expert-driven models are developed primarily through decision-making assumptions and expert-based theory. As a result, the use of data in these models does not change the assumptions and theory behind the intended model and takes a marginal role at best.

3.4 EXERCISES

3.1. (Level 1) What real-world application domains, other than security, can potentially benefit from assuming that people are fully rational?

3.2. (Level 1) Prove the following properties of the logit quantal response model: A) When the rationality parameter λ is set to 0, then a quantal response decision-maker will follow a uniform random choice over the possible strategies regardless of their expected utility. B) When $\lambda \to \infty$, quantal response converges to a "pure" utility maximization as discussed before.

3.3. (Programming) Implement a script (or excel/access) sheet which receives the number of possible strategies for a decision-maker and its rationality parameter (λ) and returns the logit quantal response probability distribution over the possible strategies.

3.4. (Level 2) Define the quantal response equilibrium based on Section 3.1.2. Prove that it is a generalization of a Nash equilibrium.

3.5. (Level 1) Analyze the Prisoner's Dilemma (Example 2.5) under the following conditions: (A) both prisoners use quantal response with $\lambda = 0.5$; (B) both prisoners are level-k reasoners with $k = 1, 2$ (mention which level-0 behavior you assumed and why it is reasonable); and (C) both prisoners have a cognitive hierarchy level k where both prisoners believe that the prisoner population is half of level 1 and the other half of level 2.

3.6. (Level 1) Find all pure Nash equilibrium in the *beauty contest game* (Example 3.2).

3.7. (Level 1) The silk market in Beijing, China is famous for its unorthodox selling technique. Each item is *significantly* overpriced to begin with. Once a costumer enters the store the seller negotiates a price with the costumer which in some cases can reach as little as 10% of the original pricing. Which cognitive biases are leveraged by the sellers?

3.8. (Level 1) Two restaurants differ in their tip systems. Restaurant A adds the common tip (10%) to the bill, clearly mentioning that the tip could be removed if a box is checked on the receipt. Restaurant B clearly mentions on its receipts that the common tip (10%) could be added to the bill if the appropriate box on the receipt is checked. Which restaurant would you consider to receive more tips? Why?

3.9. (Level 2) Analyze Rubinstein's e-mail game [277] using the expert-based approaches presented in Section 3.1. Whenever parameters need tuning, set them as you see fit.

3.10. (Level 2) Prove that the Tversky and Kahneman probability weighting function (presented in Section 3.1.6) infinitely overweights infinitesimal probabilities and infinitely underweights near-one probabilities. Namely, show that

$$\forall x_i \cdot \lim_{p_\sigma(x_i)\to 0} \frac{\pi_\sigma(x_i)}{p_\sigma(x_i)} = \lim_{p_\sigma(x_i)\to 1} \frac{1-\pi_\sigma(x_i)}{1-p_\sigma(x_i)} = \infty.$$

3.11. (Level 1) Consider a lottery ticket L that pays 10^6 with probability 10^{-6}. How much would a rational decision-maker be willing to pay to buy the lottery ticket? Now suppose that there is insurance that can save the decision-maker a loss of 10^6 that will occur with probability 10^{-6} (e.g., her house can burn down). How much would a rational decision-maker be willing to pay for full insurance against this risk? Repeat these two questions under the assumption that the decision-maker follows the prospect theory with a zero-wealth as the reference point, $w(p) = p$ and

$$v(y) = \begin{cases} y & \text{if } y \geq 0 \\ 2y & \text{otherwise.} \end{cases}$$

3.12. (Level 1) What is the **Focal Points Theory** [286]?

3.13. (Level 2) What is the **Aspiration Adaptation Theory** [289]?

3.14. (Programing) Try your data-driven modeling skills on one of the many data-sets on Kaggle `https://www.kaggle.com/`, some of which are aimed at predicting human decision-making.

3.15. (Level 1) Explain in your own words what is the **selective labels problem**. Identify another domain where the problem may manifest itself.

3.16. (Level 1) What is deep learning, and how does it contrast with other supervised learning algorithms?

3.17. (Level 2) The prediction model used by AlphaGo (Example 3.4 is considered to be highly accurate, yet it achieves "only" 57% accuracy. On the other hand, a model for predicting whether one will brush her teeth tomorrow evening with 99% accuracy should be considered rather poor in most cases. What is the reason? Can you think of easy ways to avoid this phenomena?

3.18. (Level 1) What are the different methods for collecting data on human decision-making? Elaborate on the main advantages and disadvantages of the within-subject and between-subjects experimental designs.

3.19. (Level 2) How can one evaluate a prediction model using the Receiver operating characteristic (ROC) curve [94]?

3.20. (Level 2) Provide a few new real-world examples, not mentioned in this chapter, for human decision-making prediction tasks and situate them on the personalization-situationalization space (Figure 3.7). Explain your choices.

3.21. (Level 2) Which prediction approach would you take in order to predict what stocks people are likely to invest in more today? Discuss the different options. What is the main barrier for the using such prediction models to make money in practice?

3.22. (Level 1) A large company noticed that their best and most experienced employees are leaving prematurely. Define a human decision-making prediction problem that, once solved, can be used by the company to mitigate the problem. How would you go about modeling the problem? What data sources would you use, if any? What are the problems you are likely to tackle and how do you plan to approach them?

3.23. (Level 2) A doctor finds out that almost 30% of her scheduled patients do not show up. As a result she has been "over-booking" patients. Some days, her solution works perfectly with the "right" number of patients arriving while on other days there is significant overcrowding in the clinic. Discuss how would you go about helping the doctor schedule her patients better. Which prediction approach would you adopt and why?

3.24. (Level 2) The Global Terrorism Database (`https://www.start.umd.edu/gtd/`) is an open-source database including information on terrorist attacks around the world from 1970 with annual updates. Can we expect a fully data-driven model, based of the above data set, to be of human expert prediction quality? Why? How would you go about predicting next year's terrorist attacks?

3.25. (Level 2) An autonomous car seeks to adapt its driving style (e.g., speed, distance from other vehicles) to its passengers. To that end, the car tries to predict what decisions its passenger would have made if she was the driver. However, the car cannot ask the driver to demonstrate her driving behavior as it is fully autonomous with no wheel or gas pedal. How would you predict one's decisions if you cannot observe their decisions? Can you think of a "work-around"? How would you mitigate the cold start problem where a new passenger enters the car? What would you do if there is more than a single passenger?

CHAPTER 4

From Human Prediction to Intelligent Agents

"Action speaks louder than words but not nearly as often."

Mark Twain

The prediction of human decision-making can, in many cases, enhance our understanding of how people make decisions and endow automated agents with the ability to predict human decisions, an ability we associate with **intelligence**. However, paraphrasing on Mark Twain, "**intelligent action** speaks louder than *predictions* but not nearly as often."

The maxim of this book can be phrased as follows: **a prediction model is as good as its agent performance.** In this context, the challenge of using a prediction model in an intelligent agent is of utmost importance. This point of view was also highlighted in a recent *Science Magazine* essay [302] by Subrahmanian and Kumar which claimed that "A predictive model must also provide one or more prescriptions for potential future actions…" Note that a prediction model may employ various assumptions regarding people's decision-making. However, it is important to remember that the usefulness or accuracy of a prediction model is **not** attributed to the correctness of its underlying assumptions but rather to the fact that people make decisions **as if the assumptions are correct**.

We first present the generic framework in which prediction models are normally integrated within the design of intelligent agents. Next, we systematically discuss and compare several real world domains in which the prediction of human decision-making plays a role in the design of intelligent agents. Through this comparison we aim to identify which assumptions and domain characteristics make a specific model or approach more suitable than others. We discuss our insights, as well as our own personal experience, providing best practices in Chapter 5.

4.1 PREDICTION MODELS IN AGENT DESIGN

Intelligent agents may use human prediction models in many different ways. The most common way to leverage such prediction models is through the integration thereof within the agent's reasoning and/or optimization process, as depicted in Figure 4.1.

In Chapter 3, we focused on the dashed component of Figure 4.1, denoting the prediction of human decision-making.

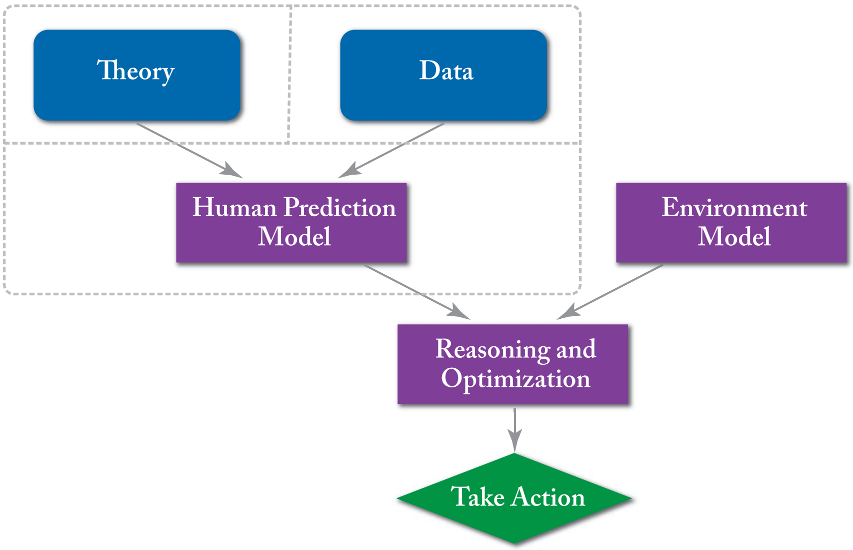

Figure 4.1: High-level illustration of the the integration of human prediction models in intelligent agent design. The dashed area is discussed in Chapter 3.

Note that in order to proficiently implement the above methodology, components outside the dashed area should be accounted for. Specifically, designing an optimization problem or a reasoning mechanism which contemplates a prediction model of human decisions along with environmental factors may not be straightforward. In some cases, the designed problem along with the associated prediction model could be readily mapped into a well-known optimization or reasoning category such as **Linear Programing** (LP) [76], **Markov Decision Process** (MDP), [258] or **regret minimization** [30], while in other cases new formulations need to be developed.

The Optimizer's Curse

If the prediction model perfectly reflects human decision-making, then an agent can select its actions through the expected utility maximization paradigm using optimization and reasoning techniques such as the ones stated above. Theoretically, on average, such an approach will receive the expected utility if the whole process is repeated many times. In reality, however, human prediction models are far from perfect. These models usually oversimplify the real decision-making setting and the representation of the decision-maker. They tend to generalize across people and provide predictions across different decision-makers. Therefore, if an agent consistently takes the prediction at face value and selects a **best response**, it should expect the value of the chosen action to be less than what was estimated, even if the value estimates are unbiased(!). This phenomenon is often referred to as the **curse of the optimizer** [298].

To illustrate the curse of the optimizer in our context, consider a *personal robotic assistant* that needs to take one of three actions (e.g., cleaning the floor, doing the dishes, or ordering food, denoted a_1, a_2, a_3). Let us assume that the true expected value of the three actions is exactly zero ($\mu_1 = \mu_2 = \mu_3 = 0$). The robot uses an estimation of the expected value of each action based on the prediction of the human's decisions (e.g., will the user choose to eat soon). We assume that the estimations are normally distributed with the mean equal to the true value, which is zero, with a standard deviation of one. We assume that the expected utilities are conditionally unbiased. If an agent selects the highest estimation of the three, the value of that selection is 0.85 on expectancy (the calculation is left to the reader as Exercise 4.2). Namely, if the agent evaluates the estimators at every decision setting (which we assume to be unconditional), then the agent is likely to face significant deviations from the expected utilities it considered (0.85 on expectation in the above example). Intuitively, by selecting the action with the highest utility estimate, the agent favors the overly optimistic estimates and that is the source of the bias.

Theoretically speaking, the optimizer's curse is very troubling. However, in practice, its impact may be limited. First, the greater the difference between the best alternative and the other alternatives (on expectancy), the more likely it is that the agent will select the best alternative. Namely, if one alternative is significantly better than others on expectancy, it is also significantly more likely to be selected by the agent. Furthermore, the estimates of expected utilities are usually (highly) correlated as they consider common (or even the very same) features in the prediction model. The extent to which the curse hinders human interacting agents' performance is yet to be investigated in real world domains.

Directly overcoming the optimizer's curse can be accomplished using a Bayesian approach to interpret the expected utility estimations. Specifically, one may define a prior on the true expected utility of every action. Then, rather than ranking alternatives based on the calculated expected utilities (i.e., evidence), one may use Bayes' rule to determine the posterior distribution of the expected utility of each action and rank the alternatives accordingly. The potential benefit of using this approach in human interacting agent design has yet to be investigated in practice. See more details in [298].

Note that the use of a prediction model which does not adequately reflect human decisions may significantly hinder an agent's performance [196, 212]. Guidelines and techniques for the design of intelligent agents, irrespective to the use of prediction models, are presented in popular AI textbooks such as [278, 332].

4.2 SECURITY GAMES

Today, security is a critical concern around the world. A security setting usually consists of a **defender** (e.g., police) which has a limited number of resources (e.g., officers) that seeks to protect a large set of **targets** (e.g., airport terminals) from an **adversary** (e.g., terrorists) [306]. Naturally, the prediction of an adversary's decisions (e.g., if, when, and where to attack) is critical for the success allocation of defender resources by the automated agent.

Automated agents have been shown to outperform human security experts by providing security predictions and efficient resource allocations that are significantly better than the current practices. These allocations are not only better, but they are also derived much faster than manual ones.

Security domains are unique in that the predictive agent (used by the defender), in most cases, cannot consider the adversary in isolation. Namely, the adversary may react to the prediction model used by the defender and may leverage it to its benefit. Therefore, an appropriate prediction by the defender should also consider the adversary's reasoning about the defender. As a result, the main question in predicting an adversary's decisions is—**How strategic does one consider the adversary to be?**

Given the possible strategic interaction between the defender and the adversary, the **Game Theoretic** approach seems to be the most adequate starting-point. Indeed, most commonly, security games are modeled as a **Stackelberg Security Game** (SSG) where the defender "commits" to a (mixed) strategy that the adversary can first observe and then (best) respond to. The underlying motivation for this model is that adversaries can, in many security settings, perform careful surveillance of the defender's actions and thus elicit the defender's policy. Consider Example 4.1 for an illustration.

Example 4.1 Consider a security guard who needs to protect two assets at different locations. Naturally, the security guard cannot be in two places at the same time, therefore complete coverage is impossible. As a result, the security guard needs to decide which of the resources to guard at each time step. A potential attacker can decide which asset to attack. Domain experts have provided you with the following payoff matrix:

		Attacker	
		*Asset*1	*Asset*2
Defender	*Asset*1	(4, -3)	(-1, 1)
	*Asset*1	(-5, 5)	(2, -1)

If players have no prior knowledge of their counterpart, then a Nash equilibrium may be calculated (left as part of Exercise 4.4). However, in practice, the attacker may perform careful surveillance and obtain a distribution of the guard's decisions (e.g., "on Sundays the guard secures Asset 1 in 75% of the cases"). In a sense, the defender "moves first," as the attacker is expected to learn the guard's mixed strategy and then attack. The calculation of the guard's optimal strategy completes Exercise 4.4.

The **ARMOR** agent is a prominent example of the practical benefits of this approach. ARMOR was developed for intelligently randomizing resource allocation at the Los-Angeles international airport (LAX) [251]. The agent prescribes where and when security officers should set up checkpoints to check cars driving into LAX, *assuming they face a perfectly rational adversary* *capable of observing their past actions and which maximizes expected utility*. The highly successful deployment of ARMOR gave rise to additional agents that follow the SSG modeling with the assumption of an adversary's perfect rationality: **IRIS** [316], a scheduler for randomized deployment of U.S. federal air marshals to flights, **TRUSTS** [341], a patrol scheduler for fare-checking in massive transit systems, and others.

Nevertheless, in many security domains the assumption of a *perfectly rational adversary* brings about poor predictions and in turn poor agent performance. As a result, the use of *bounded rational models* such as quantal response and prospect theory were adapted and extended to the security context in order to relax the perfect rationality assumption. These models have been shown to significantly improve prediction accuracy and agent performance in human studies [92, 228, 340]. A recent comparison of bounded rational models in SSG is available at [176]. The parameters of these bounded rational models are normally trained using (usually limited) data available from security agencies or by explicit experimentation in simulated settings. Newly deployed agents such as **PAWS** [91], which is designed to combat illegal poaching through the optimization of human patrol resources, use bounded rational prediction models.

Despite significant improvement in prediction accuracy by bounded rational models, in some environments, adversaries are assumed to use even less strategic reasoning and planning against the defender. Specifically, in some environments the adversary is assumed to be *reactive* to defender's decisions rather than strategic. Namely, non-strategic adversaries do not consider their decisions' effect on the defender's (future) decisions. Non-strategic adversaries in security settings have recently been modeled as **opportunistic adversaries** which choose where and when to attack in real-time based on defender presence and the attractiveness of the potential targets [344]. Predicting the decisions of non-strategic adversaries tends to take a more data-driven approach, such as the **INTERCEPT** model [177], intended at predicting poaching activity as reactive to defender actions based on significant amounts of past data. When data is insufficient for deploying a data-driven approach, an expert-driven or hybrid approach is adopted. For example, The **Traffic Enforcement Allocation Problem** (TEAP) [271] is used for scheduling efficient traffic enforcement in order to mitigate traffic accidents. In order to model drivers' decisions (e.g., drive recklessly or safely given the traffic police's current and past actions), a hybrid approach is adopted by combining both numerical estimations and heuristics found in traffic enforcement literature. The TEAP model is deployed by the Israeli Traffic police [272].

Finally, in some environments adversaries are assumed to follow *minimal to no strategic behavior at all*. Namely, the decisions of adversaries are assumed to be influenced mostly (if not only) by environmental factors. The idea is similar to the 2002 hit movie *Minority Report*, where Tom Cruise plays a police officer in the LAPD "pre-crime unit." In *Minority Report*, the LAPD uses a super-natural predicting model of when and where someone is going to commit a crime and prevent it before anything actually happens. Possible offenders have no way to know whether the police is "on to them" or not and have no way to influence the police's predictions (other than to avoid *any intention* of committing a crime to begin with). In a more realistic manifestation of the *Minority Report* notion, police departments worldwide are now implementing the **predictive policing** approach [249], where historical crime data is used to produce predictions for future crimes. By using data-driven models, police departments can predict which people and locations are at an increased risk of crime, thus creating so-called "heat maps" of criminal activity. For example, Chicago's police department is reported as using such a prediction model to help locate who might be more likely to commit violent crimes [156]. Another example is the **Anti-Poaching Engine** [243], which is currently in deployment in South Africa. The system was developed with the aim of efficiently allocating and coordinating between drones and ground rangers for mitigating illegal poaching. The underlying poachers' prediction model assumes that poachers are *solely influenced by environment factors* such as animal density, accessibility to roads and other terrain-based characteristics. The prediction of non-strategical adversaries' decisions is most naturally done using data-driven models, as long as appropriate data is available.

Security prediction models tend to take a *generalized* approach where a single model is designed for all possible attackers. A few *semi-personalized* models exist, for example in predicting people at higher risk of performing violent crimes, yet to a far lesser extent. Security prediction models strive to be *holistic* yet usually provide only *broad* prediction model situationalization. Specifically, the prediction is usually restricted to a realistic subset of adversary decision-making settings due to the large decision-making setting space. Furthermore, note that most security settings are repeated in nature. As a result, successful prediction of the adversary's decisions and effective counter-action by the defender might lead to changes in the adversary's decisions in the future. These changes may be poorly described by a prediction model, limiting its applicability. Therefore, by relying solely on historical data, a highly biased model may be learned due to the underlying distribution of defender's strategies by which the data was obtained, which cannot be regarded as holistic. In many cases, an **online learning** approach may be used to mitigate the above issue. Namely, the adversary model may be updated as more real world data becomes available. A few theoretically justified strategies for updating the adversary model are available in [133, 297].

The key factor that distinguishes between the above security settings in terms of adversary prediction models is how strategic does one consider the adversary to be. Note that there is an inverse relation between the assumed strategic capabilities of the adversary and the amount of needed data for efficient deployment. Specifically, by assuming pure rationality, as done in *ARMOR*, no data of adversary decision-making is needed for setting parameters or configuration of the prediction model. However, the less rational the adversary is assumed to be, the greater the increase in the number of parameters and thus more data is required for adequate prediction. Unfortunately, lack of appropriate data may drive agent designers toward adopting unrealistic assumptions. The systems discussed in this section are depicted in Figure 4.2.

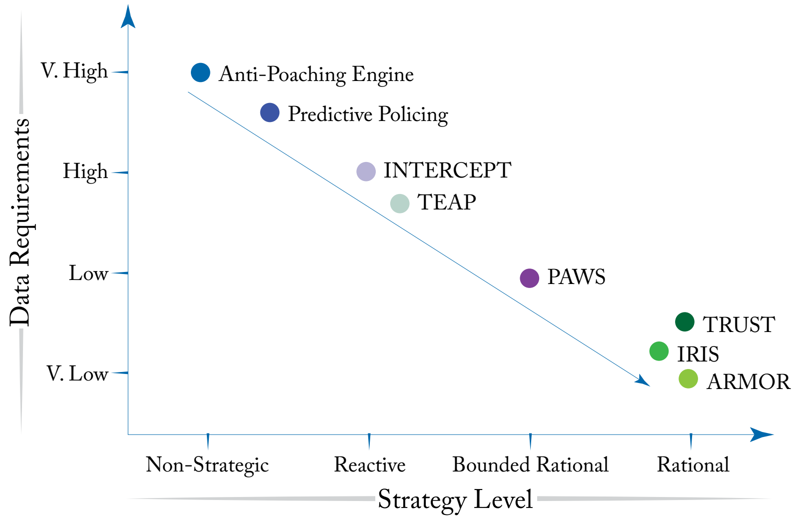

Figure 4.2: The characterization of agents in state-of-the-art security games with respect to the level of strategic reasoning assumed on the adversary's part and the amount of required data for adequate prediction. A trade-off between the two criteria is apparent.

4.3 NEGOTIATIONS

Automated negotiators can be used with humans in the loop or without them. When used with humans, they can alleviate some of the effort required of people during negotiations and can assist people that are less qualified in the negotiation process [180]. Also, there may be situations in which automated negotiators can even replace human negotiators [25, 83]. Another possibility is for people embarking on important negotiation tasks to use these agents as a training tool, prior to actually performing the task [181, 202, 315]. The success of automated negotiators depends on efficiently modeling the negotiation counter-part. Thus, the key difficulty *when negotiating with people* is to cope with human decision-making.

In contrast to the (assumed) fully rational attackers in Security Games as discussed above (e.g., in *ARMOR*), empirical studies have shown that people's negotiation decisions often fall short of what one would consider rational [85, 213]. As a result, automated negotiators that assume their human counterpart to be fully rational often perform badly with people [152, 244].

Could data be used to mitigate the above challenge? Potentially yes. However, due to the extremely wide variety of negotiation settings and protocols it seems that obtaining adequate data for a specific negotiation setting is highly complex. Consider a negotiation setting where one wants to sell a used car to a stranger. This setting is remarkably different than a negotiation setting where two executives negotiate a trading deal after several successful trading deals in the past. Note that the use of transfer learning in this realm is also highly difficult for the same reasons.

Given the complexity of obtaining contextual data on the one hand, and the well developed social science literature on what drives human negotiation decisions on the other, we discuss the different approaches by considering the amount of data needed for their successful deployment. We start with the expert driven approach, relying on expert-articulated heuristics, going through hybrid approaches that combine collected data with heuristics and concluding with data-driven models that cope with the complexities associated with obtaining data in this domain.

One of the most prominent and long-standing techniques for mitigating a human negotiator's irrationality is the use of expert-based models, based on expert knowledge and heuristics. For instance, **Curhan's Subjective Value Inventory** [75] identifies four key factors that can be leveraged to predict which offers people will accept: (1) *material outcome* (e.g., "the extent to which the terms of the agreement benefit you"); (2) *feelings about the self* (e.g., "did you lose face?"); (3) *feelings about the process* (e.g., "did the counterpart listen to your concerns?"); and (4) *feelings about the relationship* (e.g., "did the negotiation build a good foundation for a future relationship"). Following these and other "rules-of-thumb," automated agents that negotiate with people have been developed. The first software agent that negotiated with people, **Diplomat** [189], was designed for the game of Diplomacy [294].[1] A crucial component of the game involves *simultaneous, repeated negotiation*—players are expected to negotiate complex deals in incomplete and imperfect information settings where other agents' goals are usually unknown. Misleading information can be exchanged between the different agents and complex, non-committing alliances and agreements between the agents may be formed. The *Diplomat* agent used a heuristic approach where different *personality traits* of the human player, such as aggressiveness, willingness to take chances and loyalty were estimated for predicting her decisions. The prediction rate of this agent with respect to whether its human opponent will keep the agreements they signed was 92%. Most human players were not able to guess which of the players was played by the automated agent. In a similar spirit, other works such as [107, 166, 235, 281] relied on heuristics to try and predict a human negotiator's decisions. Many real-world trading bots, such as high frequency financial exchanges, advertising exchanges, or sniping agents used in eBay, employ expert-driven rule-based functions which have been hard-coded in advance by human experts [153]. The major limitation of the above expert-driven approach is obvious—it heavily relies on expert-based knowledge and heuristics which may not be available or may not bring about an adequate agent performance.

A competing approach to mitigate a human negotiatior's irrationality is by adopting bounded rationality models. These models were deployed successfully with human negotiators in several systems. For example, the **QOAgent** [203] adopts a bounded rational approach where a finite set of possible negotiator types is assumed based on expert knowledge, each with an

[1] This game is considered to be the ancestor of many popular computer strategy games, such as Sid Meier's Civilization.

additive utility function (e.g., one type might have a long-term orientation regarding the final agreement, while the other type might have a more constrained orientation). After each observed decision, the agent tries to infer which type best suits the opponent based on the assumed decision-making models. The agent considers the utility of both sides when calculating a proposal using an expert-articulated model that leverages domain-specific biases (e.g., Luce numbers [207] which represent probabilistic beliefs regarding the tendency of a negotiator to choose one offer over another). Other works, such as [244], adopt classical bounded rationality models such as the quantal response approach. Ficici and Pfeffer [96] compared the prediction capabilities of several bounded rationality models with human-human negotiations. The authors show that beyond a certain level of sophistication, more complex bounded rational models yield diminishing returns. They further show that simple bounded rational models (such as the quantal response) can be very useful for constructing automated negotiators that negotiate better than people. More recently, Haim et al. [135] showed that adopting a bounded rationality approach can surpass human performance level at three-player market settings. These agents are best classified as *expert-driven* models due to their substantial reliance on expert-articulated heuristics and behavioral models.

Note that the above agents do not assume access to a significant amount of past experiences or negotiation data. However, an agent can possibly improve its negotiation strategies based on such data, if available. An initial attempt in that direction is the **KBAgent** [240] which extended the *QOAgent* by using a generic opponent modeling mechanism, which allows the agents to model their counterpart's population and adapt their behavior to that population. The agent negotiates with each person only once, and uses past negotiation sessions of other people as a knowledge base for generic opponent modeling. The database containing a relatively small number of past negotiation sessions is used to extract the likelihood of acceptance of proposals and which proposals may be offered by the opposite side. The KBAgent still heavily relies on expert-driven components (as does the underlying QOAgent), thus it is best classified as part of the *hybrid approach*. A similar attempt was made by Byde et al. [53] who developed **AutONA**, an automated negotiation agent for negotiations between buyers and sellers over the price and quantity of a given product. While the model can be viewed as one-shot negotiations, for each experiment, AutONA was provided with data from previous experiments, allowing it to adapt its predictions. In order to model the opponent, AutONA attaches a belief function to each player that tries to estimate the probability of a price for a given seller and a given quantity. This belief function, based on expert-based heuristics, is updated based on observed prices in prior negotiations and rounds. Unlike the KBAgent who was shown to outperform human negotiators, AutONA was shown to achieve a human-like negotiation level, but no better. More recently, the **NegoChat-A** and **NegoChat** agents [275, 276] presented a novel hybrid approach which combines machine learning models for the prediction of the likelihood that an offer would be accepted with well-known human biases, such as *Anchoring* and *Aspiration Adaptation Theory* [289] for outperforming the KBAgent. The unique feature of the NegoChat and NegoChat-

A agents is that they present natural text understanding capabilities that demonstrate the need for partial agreements and issue-by-issue interactions, currently not accounted for in most prediction models. Recently, Zick et al. [347] studied a unique negotiation setting in which agents, both human and automated, need to form coalitions and agree on how to share resources. The authors show that hybrid models which rely on basic machine learning techniques and features based on game theory can provide adequate predictions of people's decisions of whether or not to join a coalition in an online study.

For a more data-driven perspective, some works have adopted a machine learning approach where negotiation decisions (offers and responses to offers) are collected and generalized. Most commonly, machine learning models take a *generalized* or *semi-personalized* approach due to the complexities of collecting many negotiation decisions of specific users. For example, Peled et al. [245] presented an agent design for repeated negotiation in incomplete-information settings that learns to reveal information strategically during the negotiation process. The agent used classical machine learning techniques to predict how people make and respond to offers during the negotiation, how they reveal information and their response to potential revelation actions by the agent using a generalized model. A recent survey by Baarslag et al. [24] demonstrates the variety of machine learning methods successfully used to predict human decisions in a wide range of negotiation settings and protocols. The survey shows the wide popularity of Bayesian learning techniques in negotiation prediction tasks. Note that the underlying notion behind this line of works is to develop agents without (or at least, to a minimal extent) relying on human expert knowledge. However, a major concern in using this approach hinges on feature construction, which may bear a significant effect on the prediction model's accuracy [274].

Figure 4.3 depicts the negotiation agents discussed in this section with respect to their prediction approach.

Due to the high number of possible negotiation settings which may vary dramatically (number of players, negotiation protocol, reward structure, temporal aspects, etc.) it is hard to conclusively say that a certain prediction method is preferable to another. However, the above discussion illustrates three dominant phenomena.

1. It seems that the expert-driven approach, and specifically the heuristic-based approach, stands out in their successful deployment with people. This outcome may be attributed to the complexities of obtaining data combined with a large number of factors that influence human decision-making in negotiation as discussed earlier in this section. These factors can be leveraged by domain experts using hand-coded heuristics and rules which can be effectively deployed by automated agents.

2. In *repeated negotiation settings*, data on past interactions is valuable. Namely, as a result one can see hybrid agents, which combine expert-driven heuristics with collected data from past interactions, which perform adequately.

3. Due to the complexity of obtaining human negotiation decisions in real-world settings, most data-driven models take a generalized or a semi-personalized approach. To date, these models seem to be outperformed by expert-driven and hybrid models.

Recent work in the field has also focused on creating adequate platforms for the design and evaluation of negotiation agents and the research of their ability to compete with as well as train human partners. One such platform, the *IAGO platform* [215], was recently proposed. By creating testable agents, direct performance comparisons can potentially be made, allowing the quantification and improvement of human-agent interaction in the negotiation realm.

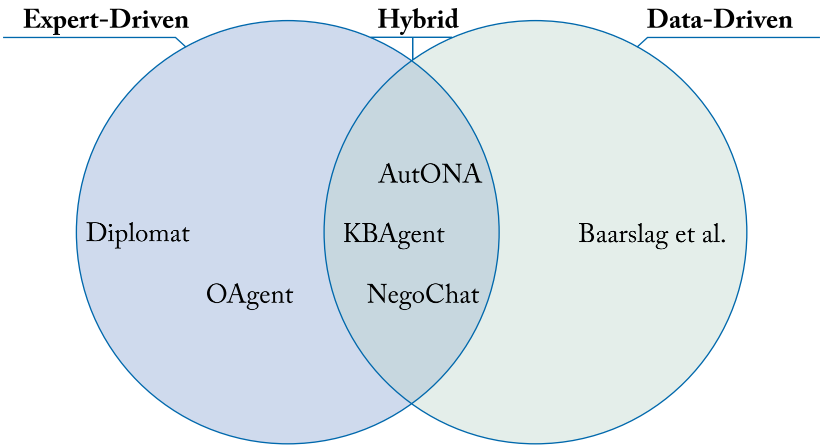

Figure 4.3: Categorization of automated negotiation agents discussed in Section 4.3.

We speculate that future developments will bring about a greater emphasis on automated agents which negotiate on one's behalf (as also speculated in [25]). Such agents will be expected to consider a fully personalized approach, be it as part of the expert-driven, hybrid, or data-driven models, for predicting their user's decisions. Note that humans often rely on costless, non-binding signals (often known as *cheap talk*) to establish cooperative relationships in repeated settings. Existing negotiation agents, of all prediction approaches, still fall short on this criterion.

4.4 ARGUMENTATION

A key human skill, utilized across many domains and activities, is the ability to *argue*. Politicians argue for their election manifestos, colleagues argue about the best way of solving a task, and we even argue with ourselves before making an important decision. Intelligent agents may benefit from the prediction of human argumentative decisions and behavior in many ways. For instance, by enhancing the argumentative capabilities of their users [287], being more persuasive while engaging with human interlocutors [155], etc. Here we discuss two argumentative prediction tasks: (1) predicting how a person would evaluate an argument (e.g., consider it persuasive or not); and (2) predicting which argument a person would present next in the course of an argumentative dialog.

Similar to negotiation settings (Section 4.3), expert-based theories and models have been developed and proposed to account for argumentative decision-making in different settings. The accumulation of these theories and models is often called the **Argumentation Theory** [17, 322]. In many respects, the Argumentation Theory is the most widely accepted expert-driven approach for predicting agents' argumentative decision-making (human or automated). Note that argumentative decision-making is unique in that the decision-maker is expected to use some form of **non-monotonic logic**. In other words, the introduction of new information or arguments can cause the conclusions or decisions of an arguer to change significantly. Unfortunately, despite its appealing theoretical properties, Argumentation Theory has been shown to provide poor prediction capabilities for **human** argumentative decisions across a wide variety of argumentative settings ranging from formal argument analysis and reasoning [59], through well-accepted argumentative principles [260, 268] to the analysis and reasoning of non-formal arguments on the Web [19]. The most comprehensive study to demonstrate the poor predictive capabilities of Argumentation Theory is that of Rosenfeld and Kraus [269]. The authors examined the most well-established argumentative principles and models with more than $1,000$ human participants across a wide range of argumentative settings and assumptions. The authors also examined the formulation of bounded rationality and heuristic prediction models which have also fallen short. Altogether, unlike negotiation settings where the expert-driven approach was shown to provide valuable predictions and in turn allowed for the design of beneficial agents, here, the use of expert-driven approaches seems inadequate for prediction and agent design. This lack of applicability as to agent design is supported by [269] who examined the use of Argumentation Theory and bounded rationality models in the design of argumentative agents aimed at supporting human argumentation. These agents performed rather poorly, resulting in a low subjective benefit for their human users. Note that other expert-based computational techniques for modeling human reasoners such as [37, 132, 154] were recently proposed. These methods have thus far not been tested with human subjects.

It is important to note that in **specific** application domains such as legal [14], medical [101], and critical reasoning [308], Argumentation Theory is able to provide adequate argumentation-support systems. These domains are unique in their emphasis on *strict protocols and reasoning rules*, making argumentation theory very applicable. This is not the case in most argumentative settings.

The good news is that people love to argue, especially online. The emergence of web discourse and argumentation platforms such as ChangeMyView, an active community on Reddit,[2] and social media debates allows the Natural Language Processing (NLP) and specifically the argumentation mining community to make significant achievements in predicting which arguments a person would find more compelling. For instance, automated classifiers have been

[2] https://reddit.com/r/changemyview

proposed to automatically score the persuasiveness of argumentative essays [93, 250]. Recently, Habernal and Gurevych [131] provided a deep learning model for comparing the persuasiveness of arguments in head-to-head comparisons using a large corpus of 16,000 arguments. A similar approach was proposed by Wei et al. [326], relying on a hybrid approach based on expert-articulated features. Surprisingly, studies show that specific non-trivial interaction patterns between users as well as linguistic factors can significantly enhance the prediction accuracy of people's argumentative decisions [307].

For the task of predicting which arguments a person is likely to put forward in an argumentative dialog, obtaining a large amount of data is impractical. Namely, while people post arguments and comments on the web debating politics, sports, and other issues, bilateral argumentative settings like a debate between two spouses on which vehicle they should purchase, or a debate between two researchers on whether to send a joint paper to Conference *A* or *B* is very hard to find online or to collect experimentally. As a result, for this task, a **hybrid approach**, leveraging data together with expert-articulated features, was shown to provide adequate predictions and beneficial agents. For instance, Rosenfeld and Kraus [269] presented a hybrid approach for the prediction of which arguments a person is likely to present in a dialog. The authors present novel argument provision agents that assist their human users by providing arguments for them to use while engaging in an argumentative dialog using a variety of argument provision policies based on the prediction model. A prediction model may also be used to derive a persuasive policy as proposed in [270] where, unlike the former case, the prediction model is integrated within an optimization technique aimed at maximizing the likelihood of a successful persuasive interaction with human interlocutors (rather than providing arguments that a person would find useful to use on her own in a dialog). The integration of expert-articulated features, which on their own pose little predictive power, is used in order to mitigate the challenge of generalizing the few argumentative decision examples that the authors were able to obtain.

It is important to note that unlike negotiation agents that can "fool" people into thinking they are human under certain settings, there is still a long way to go before the argumentation field reaches that level of performance. The use of human experts that manually classify arguments and relations between arguments in order to train data-driven and hybrid is still a very common practice. We believe that future developments in NLP, such as the automatic mapping of natural language statements into formatted arguments, the automatic extraction of arguments from texts and the automatic identification of relations between natural language arguments, as well as other open problems of great importance, will enable additional developments in the argumentation field.

To summarize, unlike Security Games (Section 4.2) and Negotiations (Section 4.3), in predicting human argumentative decisions the **data-driven and hybrid approaches seem to be more suitable** than the expert-driven approach. We can point to two possible reasons. First, in security and negotiation settings, it is (fairly) clear what the engaging parties are striving to achieve (e.g., the adversary in security games seeks to maximize the likelihood of a successful and beneficial attack, a negotiator seeks to maximize expected value, etc.) and substantial expert-knowledge has been accumulated to support different decision-making rules and hypotheses. However, in many argumentative settings it is completely unclear what the arguing parties are trying to achieve. People argue for a wide variety of reasons. These reasons may vary from wanting to establish one's beliefs as correct, as advocated in [119, p. 10]—"By spotting weaknesses, mistakes, and falsehoods in your own and other people's arguments, you stand a much better chance of holding to and acting on true beliefs."—to arguing to fulfill socio-evolutionary needs as advocated in [217, p. 284]—"Our ancestors were neither living in harmony with one another nor waging constant war... argumentation may have played at least as important a role in their social lives as in ours"—to arguing just for the sake of arguing as many people (too often) do. In any case, a multitude of factors such as incomplete knowledge over a reasoner's values, prior knowledge and reasoning capabilities seem to make the expert-driven approach unsuitable for the task of predicting human decision-making in most argumentative settings. Exception may be found in application domains such as law and medicine, where expert-articulated rules are very common and accepted by the engaging parties (e.g., precedents and legislation in Law, Hippocrates' oath in medicine). Second, even when the argumentative goal is fairly clear, for instance in persuasive settings where a professor seeks to persuade her students to enroll in graduate studies, people tend to resist changing their decisions [174]. Therefore, the prediction of which argument(s) would change one's mind is highly complex (as change is scarce). Despite the disheartening results above, due to the recent development in online debating platforms, social media, and advancements in NLP, prediction tasks such as predicting which argument(s) a person would find more compelling based on linguistics and interaction dynamics, or which arguments a person is likely to put forward in a dialog based on the dialog dynamics, can now be addressed appropriate l y.

Note that the expert-driven research on argumentation can be interpreted as a **normative approach** to how one should argue from a logical perspective as opposed to the practical perspective. Therefore, it is important to note that the scientific value of such works is well-appreciated and is not in any way questioned here except for its applicability to **human** arguers.

4.5 VOTING

Voting is a key component in group decision-making and has been since its use in ancient Greece in the 6th century BC. Similar to argumentative settings, voting is used as a social mechanism for conflict resolution between agents (in this context, referred to as **voters**), be they human or automated. The conflict is over which alternative(s) to adopt out of a set of alternatives (i.e., candidates) standing for election [72].[3] It is important to note that voting is not restricted to political decision-making alone and in fact takes a significant role in the business world (e.g.,

[3]The case of "multiwinner voting" where more than a single alternative is chosen is inherently more challenging than the "single winner voting" case [90].

board of directors' decision-making) and in other social decision-making settings (e.g., decision-making on a University's requirement committees).

Similar to Argumentation Theory in predicting argumentative decisions, different voting procedures and rules have been proposed and investigated over the years (see [43] for a survey) which are grouped together as part of the **Social Choice Theory**. While there are many voting procedures and rules, none of them is "perfect," as shown by the 1972 Nobel laureate *Kenneth Arrow* (1921–2017) in his famous impossibility theorem [13] (left as Exercise 4.11). Perhaps the most commonly applied voting procedure in practice is *plurality*, where each voter selects one candidate (or none if the voter can abstain), and the candidate with the most votes wins the election.[4] Assuming that a person would vote **sincerely**, namely reporting her true preferences, seems to be natural in two-alternative elections. Specifically, given two alternatives, the choice to vote for one's non-favorite alternative is dominated by voting for the preferred alternative, meaning that regardless of other voters' choices, voting for the preferred alternative is the rational thing to do. However, in many real-world voting settings, and specifically when more than two alternatives are available, a voter may be better off misrepresenting her preference and acting strategically. This phenomenon is call a **voting manipulation** or **strategic voting** [309]. Consider Example 4.2 for an illustration of voting manipulation.

Example 4.2 One of the most well discussed examples of voting manipulation comes from the 2000 U.S. elections where many voters who ranked third-party candidate Ralph Nader first, voted for their second choice (typically Al Gore) assuming that casting a vote for Nader provided no practical chance of making him the winner (e.g., given pre-election opinion polls) [146].

Unfortunately, as shown by the **Gibbard-Satterthwaite Theorem** [116, 282], the development of a voting procedure that guarantees that voters cannot manipulate their voting successfully (e.g., improve the chance of their preferred alternative to be the winner) when there are three or more alternatives is *impossible* (under basic assumptions). As a result, every voting mechanism can be theoretically manipulated. Combining this insight with the occurrence of strategic voting in many real world elections [100, 256] suggests that some people may engage in strategic voting under different circumstances [27], making the assumption that all people vote sincerely impractical.

In negotiation settings, expert knowledge was used to mitigate the prediction challenge. Therefore, one may wonder under what circumstances a person would try to manipulate her vote and how? Several expert-driven models were proposed, mostly to **explain strategic voting in hindsight** (e.g., [194, 262, 342, 351] to name a few). These models capture and resort to different behavioral biases such as the **bandwagon effect** (Section 3.1.5). These models are often **highly generalized in terms of personalization level and very narrow in their situationalization level**. Namely, proposed models do not consider any personalization nor do they consider the decision-making of voters across different voting settings and sets of alternatives. Part of the difficulty in

[4] Other voting procedures and illustrations of people's voting behavior are available in [291].

modeling *human* voters lies in the fact that voting is a result of subjective **preferences over the alternatives**. These preferences are unknown to an outside observer in natural experiments and may be hard for one to articulate and represent in controlled experiments (unlike negotiation settings where incentives may be set more easily). Consider Example 4.2—while the decision outcome of voters was observed (many voting for Gore), the preferences (e.g., do they prefer Gore over Nader?) and strategical considerations (i.e., were they aware of the polls?) that led to the decision outcome are unobservable in most cases. It is hard to conclusively say that people acted strategically based on the observed data alone. Moreover, even if one's own preferences can be represented and understood (say, in an extensive interview), it is hard to speculate what that person thinks that other voters will choose. This is similar to the Level-k expert-driven model (Section 3.1.3) where reasoners reason about others which in turn are expected to reason about them recursively. This is especially complex since one's incentives to engage in strategic voting may also radically depend on context and emotional factors (e.g., how strongly the person feels about the alternatives, such as the case in many political voting settings).

The challenge stated above also applies to data-driven methods. Namely, while observed argumentative choices were useful in argumentative prediction settings, here, data has to be collected while controlling for the underlying preference of the voters. To the best of our knowledge, only two papers [305, 351] have examined this question with human voters in a computational context. Tal et al. [305] examined people's voting behavior in various online settings under the plurality rule. The authors studies two popular settings. First, a single voter who is asked to cast a single vote after seeing a large pre-election poll (e.g., U.S. presidential elections). Second, a repeated voting setting until a decision is made (e.g., hiring committee). The results demonstrate that most people select the natural "default" vote, namely *voting for the most-preferred candidate in the single-shot setting, or keep voting for the same candidate in the iterative setting* when there is no "clear and obvious" strategic behavior. When a "simple" strategic decision can be made that improves their short-term payoff (e.g., voting for their second most preferred alternative would make it the winner), people usually choose the strategic alternative. This is often called a **myopic best-response**. Similar results were observed in Zou et al. [351] who investigated people's strategic voting in collaborative social event scheduling using the online scheduling website *Doodle* (`www.doodle.com`). Strategic behavior in Doodle can take different forms. For instance, when scheduling a meeting, the invited participants are asked to vote for all time slots for which they are available. However, while many time slots may be feasible for a participant, she may decide to "hide" some of her feasible time slots hoping that a more convenient (or preferred) time slot would be selected. Interestingly, most participants seem to engage in strategic voting. Furthermore, the authors speculate that there might be social pressure for voters to vote for as many slots as possible in order to appear flexible. Interestingly, strategic voting in the form of voting for unpopular slots in addition to one's preferred slots is also detected, presumably in order to appear flexible.

Automated agents may participate in voting systems *with people* for two main reasons: acting autonomously seeking to pursue their own design goal [34] or acting as proxies for individual people who delegated the voting task to the agents [110]. In either case, an agent may significantly benefit from accurately predicting people's voting decisions and manipulating its own vote to improve its well-being.

Automated agents which engage in strategic voting **with other automated agents** have been widely studied and discussed within the computational social choice literature [45].

4.6 AUTOMOTIVE INDUSTRY

The seeds of autonomous cars where planted somewhere in the 1500's, centuries before the invention of the automobile. Back then, Leonardo da Vinci designed a cart that could move without being pushed or pulled, a "self-propelled cart" if you will. Since then, impressive milestones have been achieved in this realm, such as Whitehead Torpedo (1868), which could efficiently maintain depth and changed the face of navel warfare, Mechanical Mike aircraft autopilot (1933), which was used in an around-the-world flight in 1933, reaching the autonomous cars we begin to see today [325].

Automated agents have a number of advantages over human beings that make them better drivers: they have more sensors, they are not emotional, they never get distracted or tired, and so on. However, in order for such agents to operate safely and intelligently in the real world, the prediction of what people—be they pedestrians, other drivers or the car's own passengers - are about to do next is essential.

Unlike previous sections, here, **data is usually abundant**. Millions of miles of driving data are being collected using front-facing and driver-facing cameras in different vehicles throughout the world [103]. The availability of large data sets (such as KITTI [111]) along with the inherent spatio-temporal nature of driving decisions make deep learning architectures especially attractive. Indeed, the use of deep learning is showing much promise: Bojarski et al. [38] presented a deep learning architecture that is shown to predict a human driver's choice of steering angle such that it can drive autonomously for about 98% of the time without human intervention. Fridman et al. [104] were able to predict driver's decisions using a deep learning technique based on drivers' glances. Predicting drivers' *dangerous decisions*, such as performing dangerous maneuvers, in advance was also demonstrated using deep learning in [160]. Similar deep learning methods were deployed in other studies as well (e.g., [323, 336]). These predictions may be used in a number of ways, from learning a baseline policy upon which an agent can improve (similar to the AlphaGo) or for providing driving assistance such as alerting a driver before she performs a dangerous maneuver.

Chen et al. [63] proposed a less expensive deep learning approach by transferring data from simulated driving to real-world driving. The learned affordances need to be manually associated with car actions, which is expensive, as was the case with older rule-based systems (e.g., expert-systems [115]). Interestingly, despite having their prediction model train on drivers' decisions from *a simulated environment* (the TORCS car simulation game [335]), the prediction model performs adequately on real data as well.[5]

Prediction models in this realm are usually **generalized** providing predictions for an average or good driver and do not consider driver characteristics. Despite recent attempts to personalize an autonomous car's driving behavior to individuals [179],[6] we are unaware of a personalized or semi-personalized prediction model successfully tested with human data. The state-of-the-art models often take **broad situationalization** aimed at providing a prediction across a wide range of decision-making settings. Note, however, that current models strive to provide an holistic approach, yet due to the wide variety of decision-making settings (e.g., fog, snow, gravel, construction zone, driving norms, cultures, etc.) current models are not expected to provide adequate predictions across every remotely possible decision-making setting.

Note that the automotive industry's drive to predict drivers' decisions is not restricted to the quest toward fully autonomous driving. An important field of study in this realm is **driver experience and engagement** with vehicular systems. Recent evidence suggests that drivers' current user experience often does not meet drivers' wishes, making many drivers desire more natural car systems that can "understand and predict their desires and actions" [200, 261]. A recent example is the effort to personalize and improve the thermal comfort system in modern cars. Note that *this effort is completely complementary to the drive for developing self-driving cars*. Unfortunately, unlike with driving decisions, collecting data on people's decisions in other automotive systems may be more complex.

The thermal comfort of humans has been exhaustively For that purpose, extensive investigated over the last four decades from the expert-driven perspective, resulting in the **ISO 7730 standard**[7] [15]. This standard, which was also found to be applicable in car cabins, is aimed at predicting the degree of thermal comfort of an *average person* exposed to a certain *steady* environment (see [74] for a recent survey). It relies on the assumption that user-specific parameters are available, such as thermal sensitivity, clothing and activity level. These parameters are hard to obtain automatically. The standard is in fact a **hybrid model** as it is the result of human experiments and expert knowledge. It is a broad, semi-personalized method for predicting what cabin temperature a driver would find comfortable.

While the ISO 7730 standard provides a prediction model for what temperature a driver would find comfortable, it does not prescribe predictions if one does not set the parameters online according to the person in question. Furthermore, it does not predict what actions a driver would take to bring about the desired cabin temperature. For example, Alice's desired interior cabin temperature is 21°C and she wishes to reach it as fast as possible (she does not mind enduring extreme settings in the process). Bob, on the other hand, wants to reach 19°C but re-

[5] A similar concept was also successfully deployed in predicting human operators' decisions while controlling a set of robots [265].

[6] In racing video games a number of studies have trained cars to race in a human-like style using supervised learning (e.g., [60, 314]).

[7] Also known as **Fanger's Predicted Mean Vote** (PMV) criteria.

frains from settings in which the fan speed is higher than 3 and thus prefers milder adjustments. The task of predicting what actions a driver would take to adjust her thermal comfort level is considered with different goals in mind, from reducing energy consumption [23, 266] to improving drivers' satisfaction [123, 267]. The proposed models rely on a **data-driven approach**, leveraging collected data. Note that these models train on **significantly less data than the models discussed before**, as obtaining data on drivers' thermal comfort decisions is much harder and more expensive to obtain. To date, drivers' "non-driving" decisions (e.g., tuning the radio, changing the climate control settings, etc.) are not recorded automatically. The above models take a semi-personalized to fully personalized approach and consider a broad situationalization. These prediction techniques and subsequent agents are being considered for implementation in future General Motors (GM) vehicles.

The prediction of human decision-making in driving also raises important *ethical* questions, such as what an "ethical autonomous car" should do in accordance to human drivers' ethical decision-making [201]. Unfortunately, people do not necessarily agree on what constitutes as a correct ethical choice, nor do they apply the same ethical standards to others as they do themselves. For example, when an autonomous car has to choose between sacrificing its passenger(s) or other road users, most people agree that a utilitarian approach should be taken, namely, saving as many people as possible. However, despite this ethical standpoint, most people would like **others** to buy a car that employs this ethical reasoning, but they would themselves prefer to ride in an autonomous car that protects its passengers at all costs [40]. An extended discussion about ethical consideration in predicting human decision-making is given in Section 5.3.

4.7 GAMES THAT PEOPLE PLAY

Game playing is an area of research in AI from its inception. Ever since, a major driving force behind much of the technical and theoretical progress in AI has been attributed to the task of achieving human-like or super-human performance in different game settings. This drive has led to many of the most celebrated achievements in AI, such as the successes in **Chess** [57], **Checkers** [285], **Poker** [42], and **Go** [295]. In this section we focus on designing **good** game playing agents using the prediction of human player's decisions.

An automated game playing agent may benefit from the prediction of human decision-making in two major ways. First, by predicting what moves a human player is likely to make, an agent may adapt its actions to the human player and better plan its future actions. In a common adversarial setting, such as in Chess or Go, this use is often known as **opponent modeling** [145]. Second, through the prediction of adequate move selection by *human players* (hopefully, good human players), an agent may use different AI techniques to improve its own playing policy by relying on human anticipated moves as a baseline. This technique is often associated with the terms **behavior cloning** and **imitation learning** [10]. Note that the difference between the two approaches described above is situated in how the prediction model is used by the agent. Namely, the same prediction model can potentially be used for both purposes.

Many games have an optimal strategy for playing against *fully rational opponents*, however, empirical studies suggest that people rarely converge to fully rational play [87]. Therefore, similar to the way human players take advantage of the peculiarities of their (human) opponents, predicting a human opponent's decisions may yield better agent performance, especially in games involving high degrees of uncertainty, deceptiveness and stochasticity.

In predicting human game decision-making, **expert and domain knowledge are usually highly available**. Usually, after playing a game for long enough, human players gain some sense of what constitutes a good playing policy and what one may expect from other human players. On the other hand, **data availability may vary** between games. Namely, when data is abundant (such as the case in the automotive industry), data-driven models seem to perform best in game settings as well. The most recent illustration of the above claim is the *AlphaGo* agent (Example 3.4), where the use of deep learning for predicting human game decision-making was shown to outperform past expert-driven agents (as well as human experts) which rely on expert-articulated rules and heuristics. On the other hand, when the intended game is sufficiently complex (i.e., multiplayer, real-time, etc.), or when obtaining data may be expensive, expert-driven, and hybrid approaches seem to perform well and relieve some of the data collection requirements. *The latter is the most common case in game playing agents and, in fact, it is the case in most modern commercialized computer games* [206].

Traditionally, game playing agents that are designed to play highly complex games, such as real-time video games, have been designed using *finite state machines* or other *handcrafted behavioral scripts*. As a result, the prediction of an opponent's decisions, however sometimes done implicitly, is traditionally done in an *expert-driven fashion*. These agents heavily relied on expert-driven heuristics articulated by the agent's designers based on their knowledge, intuitions and experience. The advantage of this approach is clear: no human decision-making data is needed and, given sufficient domain knowledge, the automated agent is expected to play well. Note that other expert-driven prediction methods have also been proposed, yet to a **significantly lesser extent**. These include game theoretical models (e.g., [109]) and bounded rationality models (e.g., [353]). Non-heuristic models are harder to instantiate in highly complex games and, as such, the focus of the expert-driven approach was, and still is, placed on handcoded heuristics by domain experts.

As noted before, domain knowledge is usually highly available. As a result, the combination of domain knowledge with human-generated data has been shown to be very beneficial in a variety of games. There are several ways to integrate domain knowledge within a game playing agent. First, the agent's designer may represent the agent's decision-making space in a sophisticated manner such that it will encompass, to a large extent, what drives human decision-making. For example, Thurau et al. [313] created game playing agents for the computer game Quake II. Different algorithms are presented that learn from human-generated data. Yet, the authors leverage game-specific expert knowledge and devise hierarchal "levels of learning," inspired by how people make decisions in the game. Namely, the agent learns strategic, tactical and reactive behavior separately per the designers' domain experience. Similarly, in the representation of decision-making settings in the game of **Mario** [178], different learning methods have been investigated which rely heavily on domain specific abstraction [238], similarities between different states and actions [273], and others. All of the above are expected to be devised by a human expert. Another related line of research investigates the provision of direct biasing to the agent's decision-making by relying on (non-expert) humans to provided feedback [185, 247] or demonstrations [52]. For example, Aler et al. [4] show how human demonstrations for what a robotic soccer player should do in different game settings can be generalized using machine learning techniques and used later to control a robotic soccer agent efficiently. Priesterjahn et al. [257] proposed a different approach, where the behavior of artificial opponents in a game is created and adjusted through automatically learned rules derived from human players' actions. Note that human input is highly domain specific. Hence, the resulting models are best classified as **hybrid models**.

From a data-driven perspective, human game decision-making may be leveraged without relying on domain specific knowledge at all, often known as **general game playing** [114]. The **DRON** technique, which stands for Deep Reinforcement Opponent Network, was recently proposed in this spirit [144]. *DRON* predicts an opponent's decisions based on past observations through a deep learning architecture. The technique then uses the prediction to compute an adaptive response, exploiting the human opponent's predicted idiosyncrasies. In many senses, this approach is similar to that used in AlphaGo (see Example 3.4) [295], where a deep learning prediction model was trained to predict what an expert's next move would be given a board position. Facial expressions were also used for predicting people's strategic decisions in a data-driven model in the *Centipede* game [246]. Key facial points were extracted from video snippets of the players' faces, with no domain-specific characteristics, in order to train a classifier to predict participants' game decisions. Note that, due to the ample expert knowledge collected in most game settings, the use of pure data-driven models is relatively scarce.

In many cases, agents that play games well do not play in a **human-like style**. Therefore, some game playing agents are designed with a different goal in mind: playing in the style of human players, either a particular human or humans in general [149]. The notion behind this approach is that AI should be utilized to make a game more *enjoyable and realistic*. The task of imitating human game play is a variation of the famous **Turing test** [148], and has led to the establishment of various competitions such as the **2k BotPrize** [147] and the **Mario AI Championship** [292]. The results of these competitions suggest that in many cases it is harder to create a human-like game playing agent than to create a high-performing agent that achieves superhuman performance. Nevertheless, significant successes were developed in this realm (see a recent survey in [312]). The common methodology deployed for this "imitation" task is naturally a **data-driven or hybrid** one. Namely, the prediction of human decisions is based on (sometimes limited amounts of) data, which in many cases is combined with different heuristics and game-specific knowledge. Unlike before, the success of such agents is usually evaluated through a human user study.

Most of the best-performing agents for general video games are based on tree search methods, specifically on the **Monte Carlo Tree Search** (MCTS) [248]. MCTS [50] is used for stochastic planning and game playing, and offers a more focused search of the game tree as well as the incorporation of opponent modeling. This was the method of choice in AlphaGo [295] (Example 3.4). There are different ways to incorporate the prediction of human decision-making within the MCTS framework. For example, Ponsen et al. [255] propose estimating the opponent's cards and actions in the game of **Poker** and integrate the two models in the MCTS procedure. The proposed model is generalized at first, providing a prior to the MCTS based on the entire obtained training data, yet as more experienced is gained from a specific human player the model becomes personalized. A similar approach was also deployed for the popular card game "Cheat" (also known as "I Doubt It" and "Bullshit") [35]. Other recent examples include [79, 183], where the standard MCTS decision-making is biased toward decisions made more often by human players, thus emulating human playing style to a certain extent.

To summarize, agents may benefit from the prediction of human decision-making in games in various ways; from using the prediction for achieving super-human performance to adjusting game difficulty for the human player's enjoyment. Two major insights may be derived from the discussion above.

1. **Hybrid models** perform very well across a wide variety of games. These approaches level the existing game-specific knowledge and expertise along with human-generated data to bring about adequate prediction models and good game playing agents.

2. **Data-driven models** perform well when sufficient data is available. In such cases, it seems that the use of game-specific knowledge is not needed (such as the case in Go).

There is a growing interest in creating agents for **general game playing**. Namely, addressing the question of how one would create a single agent that is capable of playing any game it is given without prior knowledge. In such settings, relying on domain specific knowledge will take a lesser role and conversely, generalized, data-driven models are likely to take a stronger role. We further speculate that future games will strive to provide a more personalized experience, especially for the task of making games more enjoyable. As a result, *personalized prediction models* that are able to transfer user's decisions across different gaming platforms will be more widely adopted.

4.8 EXERCISES

4.1. (Level 1) Consider a prediction model capable of adequately predicting what a person would do next in a game of Chess. The model is given a board position and returns a probability distribution over the legal moves. How would you use the prediction model for the design of a good Chess playing agent? Provide three options which differ in their time and memory demands.

4.2. (Level 2) An agent's disappointment (or surprise) from performing action a is defined as the difference between the expected value from performing action a and the actual value received as a result of that action. Consider k possible actions where each is estimated with an unbiased normally distributed estimator with a standard deviation of one. (A) Prove that if $k = 3$ and the true expected value of each action is 0 then an agent's disappointment is expected to be 0.85. (B) Prove that the expected disappointment of the agent increases with k.

4.3. (Level 1) An agent gained access to several unconditional prediction models. To make its decision, the agent runs the prediction models and chooses the action which received the highest confidence value for any of the models. Namely, the action for which one of the models returned the highest confidence regadless of other predictions. What is the main problem one can expect to see from the agent? Provide a numerical example to justify your claim.

4.4. (Level 2) Consider Example 4.1. (A) Find all Nash equilibrium of the game given that the attacker cannot perform surveillance. (B) Solve the game given that the attacker performs perfect surveillance. [Hint—assume that the guard secures Asset 1 with probability p and Asset 2 with probability $1 - p$. What would the attacker do in such a case? What would the guard want to do given this analysis?]

4.5. (Level 1) What is the unique notion introduced by the subjective quantal response prediction model [228]? Can this notion be adapted to non-security game settings? If so, provide a reasonable example.

4.6. (Level 1) In a repeated security setting where the adversary is assumed to be playing a best response to the defender's last action, how should the agent conduct its moves? Explain and provide an example to illustrate your claim.

4.7. (Level 2) In a maximum security prison, security guards need to plan their patrols to intercept potential breakouts. How would you recommend approaching the problem? Discuss the necessary steps as well as the advantages and limitations for a few reasonable approaches.

4.8. (Level 2) Articulate a heuristic-based negotiation policy for an automated seller in a physical shop. Assume the seller knows the true cost of each item and can send and receive purchase offers from costumers in the shop. Make sure you address the negotiation protocol (e.g., if an offer is rejected, how does the agent generate new offers, if any). If your policy were leaked to the human buyer, can she take advantage of it to improve her negotiation strategy? If so, how?

4.9. (Level 1) In order for an automated agent to negotiate on a human's behalf the agent must predict the preference of its users. This issue is considered to be one of the major issues preventing wider deployment of automated negotiation agents. Explain why.

4.10. (Level 1) The role of non-verbal signals such as facial expressions and arm gestures is extremely important in human interaction. What are the major challenges in predicting and generating non-verbal cues in human-agent argumentative interactions? Relate to [66] in your answer.

4.11. (Level 2) Explain, in your own words, what is "*Arrow's (impossibility) theorem*" (also known as the "General Possibility Theorem" or "Arrow's paradox") [13].

4.12. (Programming) The **Lemonade Stand Game** is a game for three agents (vendors) which interact by choosing a position (usually from a finite set of options) on an "island" in order to sell lemonade to the island's population. The rewards depend on the actions of all of the agents such that the further away an agent is from its competitors, the higher the reward it receives. Program an agent that plays the Lemonade Stand Game in its competition version [350]. What approach did you take? Why? Say the agent is supposed to compete *against people*, how would you modify your agent? Provide full details.

4.13. (Programming) Try your deep learning skills in the recent MIT course on *Deep Learning for Self-Driving Cars* [103]. Use the browser-based deep learning simulations (e.g., `http://selfdrivingcars.mit.edu/deeptrafficjs/`) to design a self-driving car. How would you integrate a human prediction model in your model?

4.14. (Level 1) In the classical Chess match between IBM's **Deep Blue** and *Gery Kasparov* in 1997, Deep Blue's operators were reported to have "tweaked" Deep Blue's responses specifically to Kasparov's playing style [56]. How would you characterize Deep Blue's opponent modeling technique?

4.15. (Level 2) Healthcare agents are generally considered to address a partially-conflicting interaction setting. In many application settings these agents heavily rely on the prediction of human decision-making. (A) In [339], the authors investigate the problem of finding an optimal incentive structure for encouraging people to become vaccinated against influenza. What assumptions were made about the human decision-makers? How were these assumptions justified? (B) Provide 2–3 additional real-world settings and identify which assumptions and characteristics make particular prediction approach more suitable than others.

4.16. (Level 1) Discuss the following claim: *Using a prediction model that achieves higher accuracy than chance-level is always beneficial for an automated agent.*

4.17. (Level 1) Is it possible that an agent that responds quantly to the predictions of a constructed prediction model performs better than an agent than maximizes expected utility? Explain.

CHAPTER 5

Which Model Should I Use?

"Sometimes Science is more Art than Science, Morty. A lot of people don't get that"

Rick Sanchez, Rick and Morty (S01E06)

The prediction of human decision-making can take a wide variety of different approaches and techniques as we discuss in Chapter 3 and as illustrated in Chapter 4. However, the answer to the question **"What prediction method should I use for task** X**?"** is always **"It depends".** As demonstrated in Chapter 4, the decision of which prediction approach to adopt depends on many factors, specifically **data availability**, **expert knowledge**, and **domain characteristics**.

Even the most experienced professionals cannot conclusively say which approach will perform best ex-ante. From a data-driven approach, this is often called the **"No Free Lunch Theorem".** Intuitively, it states that no single algorithm works best for every problem. However, some guidelines and best practices may assist one in finding a suitable approach and reduce the costs and efforts of developing unsuitable prediction models. We start our discussion by addressing the question of *what makes a good prediction model*. Following are the *guidelines for how one should go about developing a human decision-making prediction model*. We finish the chapter with a disucssion about the ethical considerations.

5.1 IS THIS A GOOD PREDICTION MODEL?

How can one claim that one prediction method is better than another? A common technique to evaluate the quality of a prediction model, be it of human decision-making or not, is the use of statistical analysis. It is common in statistics to assess the performance of a model using common measures such as the area under the **receiver operating characteristic (ROC) curve**, accuracy, recall, etc. (see [211, 218] for introductory books discussing the statistical analysis of prediction models). We refer to this approach as **statistical testing**. Naturally, every statistical measure has to be considered carefully. For example, the most common measure one looks for is *accuracy*. However, accurate predictions may significantly suffer from **overfitting** such that accurate predictions on training data generalize poorly to test data (a problem also associated with the **Accuracy paradox** [346, p. 118]). Remember that in virtually all prediction environments, let alone the prediction of human decision-making which is inherently noisy, a tradeoff occurs between the complexity of the learned model (e.g., in terms of the number of parameters) and its successful generalization of observed data. Namely, complex models (i.e., many parameters) are likely to fit training data very well yet may overfit it as well.

Recall the maxim of this book: **a prediction model is only as good as its agent performance.** Through this viewpoint, one should deploy the developed prediction models in a controlled experimental design and "put the agents (and thereby, the prediction models) to the test." We refer to this approach as **end-to-end testing**. Note that end-to-end testing may be highly expensive in both time and cost (recruiting participants, developing the needed code for the agent(s), etc.). Nevertheless, end-to-end testing is imperative for obtaining a complete, real world testing of the **actual** benefits of a prediction model. When comparing more than a single prediction model in an end-to-end testing, it is vital to maintain all other components of the agents constant (as much as possible) and vary only the prediction model. Otherwise, poor agent performance may be attributed to one of other components of the agent, not necessarily the prediction component (see Figure 4.1). Surprisingly, in some cases, prediction models with lower accuracy levels may perform better with people compared to more accurate ones. The reader is encouraged to reflect on this point in Exercise 5.1.

In tandem to the above considerations, one may be interested in the **interpretability** of a prediction model. For instance, a decision tree prediction model is usually considered to be much more natural for most people to understand and to a certain extent, perhaps, *trust*. Consider a magic 8-ball that **seems** to be capable of providing perfect predictions of whether a criminal up for parole will commit another crime if released. A parole officer could potentially benefit from such a prediction tool. However, the lack of understanding of what factors are accounted for in the model and the lack of interpretability of the prediction process (which human judges, jurors, or correction officers are expected to provide in human parole decision-making) may make such a magical tool be rejected despite its appealing properties. Note that without an interpretable model, relevant authorities may not be willing to test prediction models and associated agents in field trials, thus completely defeating the spirit of innovative AI applications (e.g., in security settings [177]).

5.2 THE PREDICTING HUMAN DECISION-MAKING (PHD) FLOW GRAPH

Before one starts collecting data, interviewing domain experts and implementing machine learning algorithms and intelligent agents, it is best to consider what prediction methods worked best in similar prediction tasks. Following extensive deliberations between the authors, based on our experience in the field and the domains surveyed in this book we present and illustrate the **Predicting Human Decision-making flow graph**, abbreviated as the **PHD flow graph**. The *PHD flow graph* is depicted in Figure 5.1.

The *PHD flow graph* works as follows. Given a new prediction task, the first question a designer should ask herself is whether there is **available data**. This corresponds to our recommendation in Section 3.2.3 to perform an extensive search for available data before one turns to collect her own data. If no data is available one should consider **collecting data** explicitly (as discussed in Section 3.2.3). Recall the complexities and costs associated with data collection before

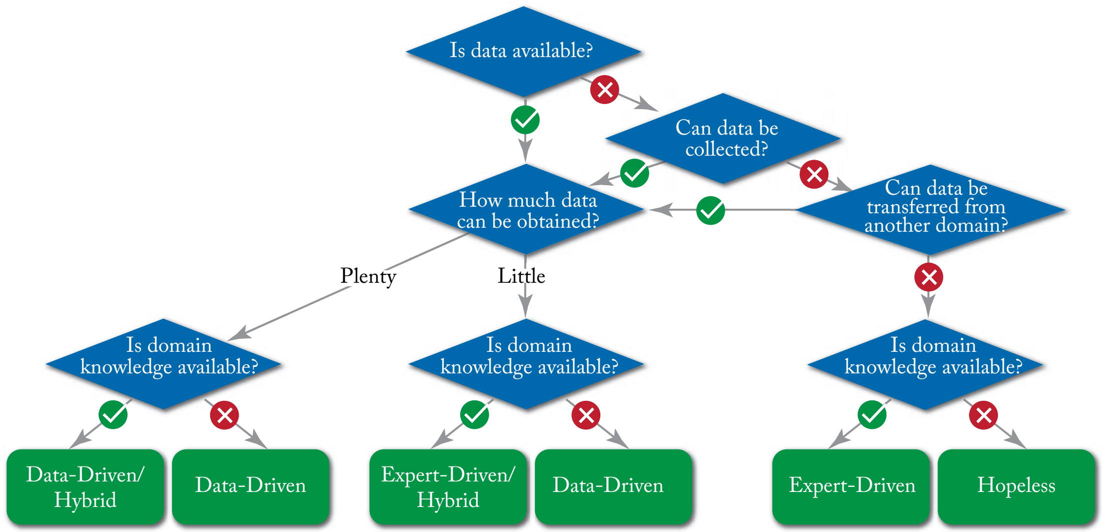

Figure 5.1: The Predicting Human Decision-making (PHD) flow graph. By following the questions in the graph's internal nodes (starting from it root (Is data available?), a terminal node is reached—pointing to the most well-suited technique(s) to address a given prediction task.

answering the question to the affirmative. If the answer is negative, then the last data option is **transferring related data** from other fields. Unfortunately, as discussed in Section 3.2.8, the use of transfer learning in the prediction of human decision-making is mostly considered an open challenge.

If data is unavailable and cannot be collected or transferred from other fields then one may resort to **expert knowledge**. If such knowledge exists, then the use of **expert-driven** prediction is the main way to go. Otherwise, the prediction task is somewhat "hopeless." To exemplify this part of the PHD flow let us revisit the **ARMOR** agent from Section 4.2. The prediction task underlying the *ARMOR* agent is predicting where an attacker would choose to attack an airport. Naturally, data is not available and collecting such data is virtually impossible (i.e., how many examples does one have from the past?). Nevertheless, expert knowledge does exist. Namely, experts assume attackers to be rational since they can perform careful surveillance over the defender's actions and decide when and where to attack accordingly. As such, the task of predicting where and when an attack would occur on an airport fits the expert-driven approach (specifically, game-theoretical models). A similar argument applies to the **IRIS** and **TRUSTS** agents from the security domain, also discussed in Section 4.2. Additional examples come from the negotiation domain (Section 4.3 where the negotiation agents **Diplomat** and **QOAgent** (among others) take an expert-driven approach based on the significant expert knowledge accumulated in the field. In certain complex games (Section 4.7), for which obtaining data may be expensive and domain knowledge is available, this approach is used as well. A slightly different setting is encountered when one tackles the prediction task of identifying **who is likely to choose to attempt to assassinate his prime minister or president?**. In Israel, only a single attempt (unfortunately successful) was made to assassinate the Prime Minister *Yitzhak Rabin* (1922–1995) in 1995. Of course, data is unavailable and collecting or transferring appropriate data is impractical. Unfortunately, expert knowledge of what drives one to assassinate his prime minister have not matured (unlike expert knowledge on potential terrorists which sadly has accumulated around the world by intelligence agencies as we shall discuss next). As a result, it seems that the task of adequately predicting who is likely to attempt to assassinate his prime minister is virtually hopeless at the moment.

If a significant amount of quality data can be obtained, data-driven models seem to be the natural choice. However, if expert knowledge is available, such knowledge may be used to enhance the prediction quality. For instance, consider the task of predicting the persuasiveness of arguments on the web (Section 4.4). Plenty of argument examples are available online, however, the injection of argumentative and linguistic knowledge such as NLP features and non-trivial interaction patterns between users is shown to boost prediction quality. Note that in some cases domain knowledge exists but is not used; this is the case for the **Anti-Poaching Engine** for battling illegal poaching (Section 4.2). It takes a data-driven approach, avoiding considering expert-based knowledge on the possible strategical interaction between the defenders and the poachers. In the case where no domain knowledge is available, the use of data-driven models seems most natural. This is the case in predicting drivers' decisions for **autonomous driving** directly from visual input.

To better understand the *PHD flow graph*, the reader is encouraged to refer to the exercises at the end of this chapter.

5.3 ETHICAL CONSIDERATIONS

Joseph Weizenbaum (1923–2008), the creator of the *ELIZA* program [327] (among many other important achievements), was an advocate for the restrictive usage of AI. He claimed that certain tasks should never be done by machines which he argued had no compassion and intuition, qualities which are required for many tasks such as nursing [328]. Weizenbaum's main claim was that there is more to the human mind than just the brain, and that a simple replication (or approximation) of the brain will never achieve realistic human behavior. While some readers may agree with the sentiment of Weizenbaum's claim, with the enhanced capabilities of predicting human decision-making, and potentially human *ethical* decision-making, one may wonder if and where would the line should be drawn.

The study of how ethical considerations can be injected into an automated agent is studied within the emerging field of *machines ethics* [97]. Many different facets play a role in endowing an agent with moral capabilities ([61] provides an up-to-date comprehensive survey). In this

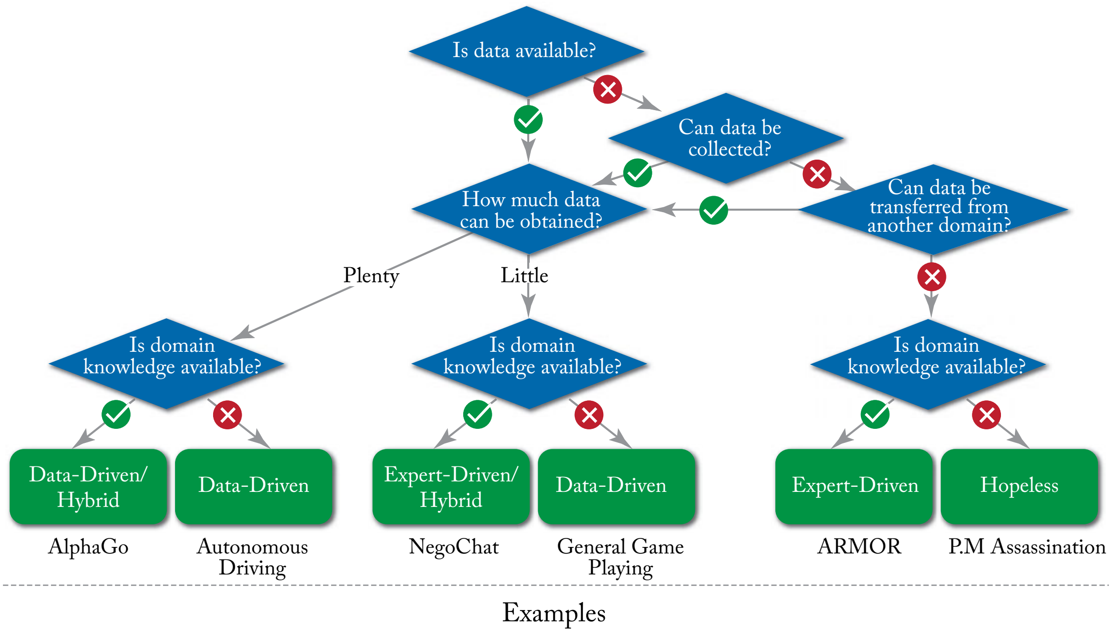

Figure 5.2: Representative agents and prediction tasks using the *PHD flow graph*.

section, we discuss a few directions where the prediction of human ethical decision-making may be used to advance this important field.

The first issue we discuss is how an agent can leverage the prediction of human *ethical* decision-making in order to make ethical decisions on its own. The most straightforward way to predict what a human's ethical decisions is by using **expert-articulated rules** that are expected to apply to the intended encounters between the designed agent and humans [47]. One may consider *Asimov's rules* for robotics [16] as an example of expert-articulated hierarchal rules. Expert rules suffer from two major limitations. First, it is hard for an expert to ensure that the intended agent(s) behave ethically under all possible decision-making settings it/they may face. *Underspecification* of the expert-driven model may result in unethical behavior of the agent. Conversely, *overspecification* of ethical rules may prove detrimental, as reasoning over a large set of rules may be very expensive. Second, no personalization is taken into account. Specifically, *who is to guarantee that a set of ethical rules and their relations are applicable and fixed across all people?* See [175] for a discussion on how different cultures consider machine hazards and ethics.

Consider the MedEthEx agent which acts as an "ethical advisor" for evaluating medical decisions [7]. MedEthEx takes as input a list of *ethical duties* for a doctor provided by an expert where duties can be overruled by a stronger duty/rule, in a similar spirit to Brook's famous subsumption agent architecture [48]. The goal of the MedEthEx agent is to learn the preference order of the duties and recommend ethical courses of action. To that end, the agent uses training examples consisting of ethical dilemmas and the decision made by the doctor, thus providing a data-driven dimension to the expert-articulated rules. A similar approach was proposed by Armstrong [12] where a set of expert-articulated ethical utility functions is given and a linear weighting between the functions is learned to capture the probability of each of them being the "correct" one according to contextual training examples.

From a more data-driven approach, Guarini [128] explores the use of an artificial neural network which is trained using examples of ethical dilemmas that are labeled as ethically permissible or not. A reinforcement learning approach may also be applicable where an agent seeks to find the right policy given an unknown "ethical utility function" [3].

In its current state, there seems to be little agreement about what makes one prediction model better than another for predicting ethical decision-making. Nevertheless, most models seem to follow a broad generalized approach, assuming that (most) people follow the same ethical standard across different decision-making settings. We speculate that the investigation of human ethical decision-making in different domains may reveal important insights for driving this research forward. Specifically, ethical decision-making in medicine, autonomous driving, and law may present different phenomena that can be levered by hybrid approaches that combine domain knowledge and data. The study of *transfer learning* [241] in this realm proposes an interesting opportunity to study the nature of ethical decision-making across cultures and settings.

We now turn to discuss the ethical considerations of the use of prediction models for human decision-making in different domains.

Consider the prediction of what (potential) criminals would do next [249]. Powerful prediction can allow police to focus on those most likely to commit crimes, thus reduce criminal activity, promote social values, and allow for better relationships between police and normative citizens. Sounds good, doesn't it? However, such prediction models may use **racial profiling** and infringe on civil liberties with little accountability. Those who advocate the use of automated prediction models usually claim that when a police officer is allocating resources or selects individuals for inspection on the street or at the airport, it still deploys some prediction model, though sometimes implicitly, thus computerized prediction models can in fact *reduce* unethical or socially unacceptable considerations. For instance, the New York Police Department's infamous stop-and-frisk program was conclusively shown to target African-American and Latino men significantly more often than Caucasian men [113]. Another debatable question is whether the use of different information sources is ethically acceptable. For example, should one's medical record or bank details be used to predict if she is going to commit treason? On some topics the debate seems to rage on while in other cases a large consensus seems to be more feasible. For example, a recent survey shows that online screening for suicidal thoughts and behaviors (Section 3.2.5) appears to be socially acceptable by most people [329].

The use of prediction models and the information that feeds into them, be these models part of a *decision-support tool* for decision-makers or an *autonomous decision tool*, is subject to debate in popular media and among legislative authorities. Balancing the potential benefits with the potential consequences is a key challenge in this realm. We speculate that if more parts of the data collection and prediction model development process were to be transparent, it would increase the acceptance of prediction models in practice. In any case, it is important to discuss ethically charged issues with one's university ethics advisory committee, most commonly the **Institutional Review Board** (IRB), or other acceptable ethical function at one's institute. The authors of this book adopted a practice of consulting the university's ethical advisor at least once a year.

5.4 EXERCISES

5.1. (Level 1) Discuss the main differences between **statistical testing** and **end-to-end testing**. Why should end-to-end testing be considered more reliable?

5.2. (Level 1) Legal reasoning in argumentation (Section 4.4) has very strict, well-accepted rules and protocols that are followed most of the time. On the other hand, data is rather abundant as legal proceedings are recorded and transcribed. Which method would you use?

5.3. (Level 2) In order to reduce costs, supermarket chains use a new technology where the customers themselves scan the products instead of human cashiers. A human supervisor is in charge of picking people for screening—namely, the human supervisor seeks to predict which customer is more likely to misrepresent her shopping. How would you go about modeling the prediction task? What data or expert knowledge would you seek and how? [Hint: *selective labeling problem*.]

5.4. (Level 1) How would you model the prediction task of how a member of Parliament would vote on a new legislative act? How would you handle the *cold-start* problem? Explain.

5.5. (Level 1) Some patients do not comply with doctors' medical advice such as medication dosage, diet, therapy, etc. How would you model the task of which patients are likely to disregard doctors' advice?

5.6. (Level 2) The selection of papers for publication in refereed conferences and journals goes through the evaluations of individual **Program Committee** (PC) members. Each PC member is asked to read a submitted paper and provide a recommendation (usually on a Likert-scale ranging from "strong reject" through "borderline" to "strong accept." Given a PC member's past recommendations on different papers, could one develop an agent to predict a reviewer's recommendation on a new paper? Consider the case of *blind review* (the paper's authors are unknown during reviewing). [Mention the "NIPS Experiment" in your answer [193].]

5.7. (Programming) **Agar.io** (`http://agar.io/`) is a massive multiplayer online action game. Human players control a cell on a map representing a petri dish. The goal is to gain as much mass as possible by eating agar (food pellets) and cells smaller than the player's cell, while avoiding larger ones which can eat the player's cell(s). Build an automated agent aimed at playing against human players.

5.8. (Level 1) Consider *Asimov's rules* as the ethical decision rules for a robot in their original ordering of importance [16]. (A) Can you think of a setting where the ethical decision derived from Asimov's rules would not coincide with your idea of an ethical decision? (B) What are the implications of changing the ordering of the three rules? Illustrate.

5.9. (Level 2) What are the major social, cultural, and ethical tensions that emerge due to the use of human decision-making prediction? Provide two examples that were not mentioned in this chapter.

CHAPTER 6

Concluding Remarks

"Success consists of going from failure to failure without loss of enthusiasm"

Winston Churchill

Predicting human decision-making is both widely beneficial and deeply problematic. In this book, we reviewed and illustrated the main challenges, techniques, algorithms, and empirical methodologies for predicting human decision-making and their use in intelligent agent design.

We explored the three major prediction approaches, thoroughly discussed and exemplified in Chapter 3:

1. Expert-driven.
2. Data-driven.
3. Hybrid.

But this book is not just about predicting human decision-making, it is also about designing intelligent agents based on such prediction models. In this vein, we examined some of the most popular and intriguing domains to-date that have shown to benefit from the prediction of human decision-making. These domains are discussed, and both domain-specific and general insights are highlighted in Chapter 4. For tackling new prediction tasks and developing novel agents we provide the **Predicting Human Decision-making (PHD) flow graph** in Chapter 5, concluding our exploration of the field in the context of this book.

As advancements in technology are made, specifically human interacting technology (e.g., smart phones), the need for efficient prediction models of human decisions becomes more prevalent. In recent years, we have seen an ever-growing emphasis on data-driven models which is attributed to two main factors: (1) the large quantities of human-generated data we can now easily collect and store; and (2) the major advancements in machine learning methods such as deep learning.

We expect to see the following major trends in the near future.

1. **Deep learning** is expected to lead to new and exciting applications that rely on the prediction of human decision-making. We believe that AlphaGo (Example 3.4) is just one of many possible AI achievements that can leverage deep learning and human decision-making prediction. We see new graduate students and many engineering teams in industry that have already abandoned many of the classical data-driven approaches and shifted to

deep learning. We expect to see this trend manifest itself in the prediction of human behavior and decision-making more extensively in the near future.

2. As significant amounts of data are gathered on each of us everyday by different technological instruments, we expect to see more **observational data** (which is "in the wild" as opposed to explicit experimentation). This data will most probably be rich in terms of recorded modalities (video, sound, etc.) and will allow better **personalization** of prediction models that will slowly become a standard. This will assist in building a relationship between human users and automated systems which deploy the prediction.

3. **Hybrid models** are expected to shed new light on the possible mutual development of both the expert-driven and data-driven approaches. This will possibly allow for new breakthroughs in both paradigms. Hybrid models are also expected to be applied to new domains and settings, previously considered too complex for expert and data-driven models.

4. **Transfer learning** in human decision-making is expected to be a catalyst for adopting human decision-making prediction models across new domains. Furthermore, the use of transfer learning is expected to put a stronger emphasis on *data-driven and hybrid models* capable of providing adequate predictions across a wide range of domains and settings.

5. **New application domains** are likely to trigger new research questions and drive the development of new prediction approaches and methods.

6. **End-to-end testing**, where prediction models are integrated within an agent design and tested with actual people, will become a standard tool in a researcher's toolbox. Recall the maxim of this book: **A prediction model is as good as its agent's performance**.

We hope that this book encouraged you, the reader, to keep pushing forward the knowledge and prediction abilities of human decision-making. **There is more to this challenge than meets the eye**.

Bibliography

[1] Amayas Abboute, Yasser Boudjeriou, Gilles Entringer, Jérôme Azé, Sandra Bringay, and Pascal Poncelet. Mining twitter for suicide prevention. In *International Conference on Applications of Natural Language to Data Bases/Information Systems*, pages 250–253, Springer, 2014. DOI: 10.1007/978-3-319-07983-7_36. 46

[2] Mohammed Abdellaoui, Han Bleichrodt, and Hilda Kammoun. Do financial professionals behave according to prospect theory? An experimental study. *Theory and Decision*, 74(3):411–429, 2013. DOI: 10.1007/s11238-011-9282-3. 35

[3] David Abel, James MacGlashan, and Michael L. Littman. Reinforcement learning as a framework for ethical decision making. In *AAAI Workshop: AI, Ethics, and Society*, 2016. 92

[4] Ricardo Aler, Jose M. Valls, David Camacho, and Alberto Lopez. Programming robosoccer agents by modeling human behavior. *Expert Systems with Applications*, 36(2):1850–1859, 2009. DOI: 10.1016/j.eswa.2007.12.033. 81

[5] Maurice Allais. Allais paradox. In *Utility and Probability*, pages 3–9, Springer, 1990. DOI: 10.1007/978-1-349-20568-4_2. 18

[6] Alon Altman, Avivit Bercovici-Boden, and Moshe Tennenholtz. Learning in one-shot strategic form games. In *European Conference on Machine Learning*, pages 6–17, Springer, 2006. DOI: 10.1007/11871842_6. 41

[7] Michael Anderson, Susan Leigh Anderson, and Chris Armen. Medethex: A prototype medical ethics advisor. In *Proc. of the National Conference on Artificial Intelligence*, 21, pages 1759–1765, 2006. 91

[8] Simon P. Anderson, Jacob K. Goeree, and Charles A. Holt. Rent seeking with bounded rationality: An analysis of the all-pay auction. *Journal of Political Economy*, 106(4):828–853, 1998. DOI: 10.1007/978-3-540-79182-9_13. 23

[9] Ayala Arad and Ariel Rubinstein. The 11–20 money request game: A level-k reasoning study. *The American Economic Review*, 102(7):3561–3573, 2012. DOI: 10.1257/aer.102.7.3561. 25

[10] Brenna D. Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57(5):469–483, 2009. DOI: 10.1016/j.robot.2008.10.024. 79

[11] Dan Ariely. *Predictably Irrational*, HarperCollins, New York, 2008. DOI: 10.1037/e653632011-003. 28, 30

[12] Stuart Armstrong. Motivated value selection for artificial agents. In *AAAI Workshop: AI and Ethics*, 2015. 92

[13] Kenneth J. Arrow. *Social Choice and Individual Values*, 12, Yale University Press, 2012. 75, 84

[14] Kevin D. Ashley. *Modeling Legal Arguments: Reasoning with Cases and Hypotheticals*, MIT Press, 1991. 72

[15] ASHRAE. Standard 55-2013. Thermal environmental conditions for human occupancy. *American Society of Heating, Refrigerating and Air-Conditioning Engineering*, 2013. 78

[16] Isaac Asimov. Runaround. *Astounding Science Fiction*, 29(1):94–103, 1942. 91, 94

[17] Katie Atkinson, Pietro Baroni, Massimiliano Giacomin, Anthony Hunter, Henry Prakken, Chris Reed, Guillermo Simari, Matthias Thimm, and Serena Villata. Towards artificial argumentation. *AI Magazine*, 2017. DOI: 10.1609/aimag.v38i3.2704. 72

[18] Yisrael Aumann. Keynote talk at *IAAI Day*, Ariel University, Israel, October 2014. 17

[19] Edmond Awad, Jean-François Bonnefon, Martin Caminada, Thomas W. Malone, and Iyad Rahwan. Experimental assessment of aggregation principles in argumentation-enabled collective intelligence. *ACM Transactions on Internet Technology (TOIT)*, 17(3):29, 2017. DOI: 10.1145/3053371. 72

[20] Amos Azaria, Zinovi Rabinovich, Claudia V. Goldman, and Sarit Kraus. Strategic information disclosure to people with multiple alternatives. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 5(4):64, 2015. DOI: 10.1145/2558397. 24

[21] Amos Azaria, Ariella Richardson, and Sarit Kraus. An agent for the prospect presentation problem. In *Proc. of the International Conference on Autonomous Agents and Multi-agent Systems*, pages 989–996, 2014. 35

[22] Amos Azaria, Ariella Richardson, Sarit Kraus, and V. S. Subrahmanian. Behavioral analysis of insider threat: A survey and bootstrapped prediction in imbalanced data. *IEEE Transactions on Computational Social Systems*, 1(2):135–155, 2014. DOI: 10.1109/tcss.2014.2377811. 47

[23] Amos Azaria, Ariel Rosenfeld, Sarit Kraus, Claudia V. Goldman, and Omer Tsimhoni. Advice provision for energy saving in automobile climate-control system. *AI Magazine*, 36(3):61–72, 2015. DOI: 10.1609/aimag.v36i3.2603. 49, 79

[24] 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*, 30(5):849–898, 2016. DOI: 10.1007/s10458-015-9309-1. 70

[25] Tim Baarslag, Michael Kaisers, Enrico H. Gerding, Catholijn M. Jonker, and Jonathan Gratch. When will negotiation agents be able to represent us? The challenges and opportunities for autonomous negotiators. In *Proc. of the 26th International Joint Conference on Artificial Intelligence, (IJCAI)*, pages 4684–4690, 2017. DOI: 10.24963/ijcai.2017/653. 67, 71

[26] Michael Bacharach. Interactive team reasoning: A contribution to the theory of co-operation. *Research in Economics*, 53(2):117–147, 1999. DOI: 10.1006/reec.1999.0188. 35

[27] Anna Bassi. Voting systems and strategic manipulation: An experimental study. *Journal of Theoretical Politics*, 27(1):58–85, 2015. DOI: 10.1177/0951629813514300. 75

[28] Roy Baumeister. The psychology of irrationality: Why people make foolish, self-defeating choices. *The Psychology of Economics Decisions*, 1:3–16, 2003. 2

[29] Antoine Bechara, Hanna Damasio, and Antonio R. Damasio. Emotion, decision making and the orbitofrontal cortex. *Cerebral Cortex*, 10(3):295–307, 2000. DOI: 10.1093/cercor/10.3.295. 2

[30] David E. Bell. Regret in decision making under uncertainty. *Operations Research*, 30(5):961–981, 1982. DOI: 10.1287/opre.30.5.961. 62

[31] Aubrianna Bilyea, Nitin Seth, and Hussein A. Abdullah. Robotic assistants in personal care: A scoping review. *Medical Engineering and Physics*, 2017. DOI: 10.1016/j.medengphy.2017.06.038. 2

[32] Ken Binmore. *Game Theory: A Very Short Introduction*, 173, Oxford University Press, 2007. DOI: 10.1093/actrade/9780199218462.001.0001. 16, 21

[33] Tamal Biswas and Kenneth Regan. Measuring level-k reasoning, satisfying, and human error in game-play data. In *IEEE 14th International Conference on Machine Learning and Applications (ICMLA)*, pages 941–947, 2015. DOI: 10.1109/icmla.2015.233. 25

[34] Moshe Bitan, Ya'akov Gal, Sarit Kraus, Elad Dokow, and Amos Azaria. Social rankings in human-computer committees. In *Proc. of the 27th AAAI Conference on Artificial Intelligence (AAAI)*, 2013. 77

[35] Moshe Bitan and Sarit Kraus. Combining prediction of human decisions with ISMCTS in imperfect information games. *arXiv Preprint arXiv:1709.09451*, 2017. 82

[36] Moshe Bitan, Galit Nahari, Zvi Nisin, Ariel Roth, and Sarit Kraus. Psychologically based virtual-suspect for interrogative interview training. In *Proc. of the 31st AAAI Conference on Artificial Intelligence*, pages 993–1000, 2017. DOI: 10.1007/978-3-319-47665-0_42. 32

[37] Elizabeth Black, Amanda J. Coles, and Christopher Hampson. Planning for persuasion. In *Proc. of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 933–942, International Foundation for Autonomous Agents and Multiagent Systems, 2017. 72

[38] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Prasoon Goyal, Lawrence D. Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, et al. End to end learning for self-driving cars. *arXiv Preprint arXiv:1604.07316*, 2016. 40, 77

[39] Jean-François Bonnefon. A theory of utility conditionals: Paralogical reasoning from decision-theoretic leakage. *Psychological Review*, 116(4):888, 2009. DOI: 10.1037/a0017186. 17

[40] Jean-François Bonnefon, Azim Shariff, and Iyad Rahwan. The social dilemma of autonomous vehicles. *Science*, 352(6293):1573–1576, 2016. DOI: 10.1126/science.aaf2654. 79

[41] Antoni Bosch-Domenech, Jose G. Montalvo, Rosemarie Nagel, and Albert Satorra. One, two, (three), infinity,...: Newspaper and lab beauty-contest experiments. *The American Economic Review*, 92(5):1687–1701, 2002. DOI: 10.1257/000282802762024737. 27

[42] Michael Bowling, Neil Burch, Michael Johanson, and Oskari Tammelin. Heads-up limit hold'em poker is solved. *Science*, 347(6218):145–149, 2015. DOI: 10.1145/3131284. 79

[43] Steven J. Brams and Peter C. Fishburn. Voting procedures. *Handbook of Social Choice and Welfare*, 1:173–236, 2002. 75

[44] Eurostat, Handbook of recommended practices for questionnaire development and testing in the european statistical system. *Technical Report*, 2006. 45

[45] Felix Brandt, Vincent Conitzer, and Ulle Endriss. Computational social choice. *Multiagent Systems*, pages 213–283, 2012. 77

[46] John S. Breese, David Heckerman, and Carl Kadie. Empirical analysis of predictive algorithms for collaborative filtering. In *Proc. of the 14th Conference on Uncertainty in Artificial Intelligence*, pages 43–52, Morgan Kaufmann Publishers, Inc., 1998. 50

[47] Gordon Briggs and Matthias Scheutz. Sorry, I can't do that": Developing mechanisms to appropriately reject directives in human-robot interactions. In *AAAI Fall Symposium Series*, 2015. 91

[48] Rodney Brooks. A robust layered control system for a mobile robot. *IEEE Journal on Robotics and Automation*, 2(1):14–23, 1986. DOI: 10.1109/jra.1986.1087032. 91

[49] Noam Brown and Tuomas Sandholm. Safe and nested subgame solving for imperfect-information games. *arXiv Preprint arXiv:1705.02955*, 2017. 39

[50] Cameron B. Browne, Edward Powley, Daniel Whitehouse, Simon M. Lucas, Peter I. Cowling, Philipp Rohlfshagen, Stephen Tavener, Diego Perez, Spyridon Samothrakis, and Simon Colton. A survey of Monte Carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in Games*, 4(1):1–43, 2012. DOI: 10.1109/tciaig.2012.2186810. 82

[51] Wändi Bruine de Bruin, Andrew M. Parker, and Baruch Fischhoff. Individual differences in adult decision-making competence. *Journal of Personality and Social Psychology*, 92(5):938, 2007. DOI: 10.1037/0022-3514.92.5.938. 2

[52] Tim Brys, Anna Harutyunyan, Halit Bener Suay, Sonia Chernova, Matthew E. Taylor, and Ann Nowé. Reinforcement learning from demonstration through shaping. In *IJCAI*, pages 3352–3358, 2015. 81

[53] Andrew Byde, Michael Yearworth, Kay-Yut Chen, and Claudio Bartolini. Autona: A system for automated multiple 1-1 negotiation. In *E-Commerce, CEC. IEEE International Conference on*, pages 59–67, 2003. DOI: 10.1109/coec.2003.1210234. 69

[54] Colin F. Camerer. *Behavioral Game Theory: Experiments in Strategic Interaction*, Princeton University Press, 2003. 21

[55] Colin F. Camerer, Teck-Hua Ho, and Juin-Kuan Chong. A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, 119(3):861–898, 2004. DOI: 10.1162/0033553041502225. 26, 27

[56] Murray Campbell. Knowledge discovery in deep blue. *Communications of the ACM*, 42(11):65–67, 1999. DOI: 10.1145/319382.319396. 84

[57] Murray Campbell, A. Joseph Hoane, and Feng-hsiung Hsu. Deep blue. *Artificial Intelligence*, 134(1-2):57–83, 2002. DOI: 10.1016/s0004-3702(01)00129-1. 3, 79

[58] David Carmel and Shaul Markovitch. Opponent modeling in multi-agent systems. In *International Joint Conference on Artificial Intelligence*, pages 40–52, Springer, 1995. DOI: 10.1007/3-540-60923-7_18. 3

[59] Federico Cerutti, Nava Tintarev, and Nir Oren. Formal arguments, preferences, and natural language interfaces to humans: An empirical evaluation. In *Proc. of the 21st European Conference on Artificial Intelligence*, pages 207–212, IOS Press, 2014. DOI: 10.3233/978-1-61499-419-0-207. 72

[60] Benoit Chaperot and Colin Fyfe. Improving artificial intelligence in a motocross game. In *IEEE Symposium on Computational Intelligence and Games*, pages 181–186, IEEE, 2006. DOI: 10.1109/cig.2006.311698. 78

[61] Vicky Charisi, Louise Dennis, Michael Fisher Robert Lieck, Andreas Matthias, Marija Slavkovik Janina Sombetzki, Alan F. T. Winfield, and Roman Yampolskiy. Towards moral autonomous systems. *arXiv Preprint arXiv:1703.04741*, 2017. 90

[62] Nitesh V. Chawla, Kevin W. Bowyer, Lawrence O. Hall, and W. Philip Kegelmeyer. Smote: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16:321–357, 2002. 48

[63] Chenyi Chen, Ari Seff, Alain Kornhauser, and Jianxiong Xiao. Deepdriving: Learning affordance for direct perception in autonomous driving. In *Proc. of the IEEE International Conference on Computer Vision*, pages 2722–2730, 2015. DOI: 10.1109/iccv.2015.312. 77

[64] Li Chen, Marco de Gemmis, Alexander Felfernig, Pasquale Lops, Francesco Ricci, and Giovanni Semeraro. Human decision making and recommender systems. *ACM Transactions on Interactive Intelligent Systems (TIIS)*, 3(3), 2013. DOI: 10.1145/2533670.2533675. 29

[65] Sonia Chernova and Andrea L. Thomaz. *Robot Learning from Human Teachers*, Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers, 2014. DOI: 10.2200/s00568ed1v01y201402aim028. 38

[66] Vijay Chidambaram, Yueh-Hsuan Chiang, and Bilge Mutlu. Designing persuasive robots: How robots might persuade people using vocal and nonverbal cues. In *Proc. of the 7th Annual ACM/IEEE International Conference on Human-robot Interaction*, pages 293–300, 2012. DOI: 10.1145/2157689.2157798. 84

[67] Helen Christensen, Philip J. Batterham, and Bridianne O'Dea. E-health interventions for suicide prevention. *International Journal of Environmental Research and Public Health*, 11(8):8193–8212, 2014. DOI: 10.3390/ijerph110808193. 47

[68] Andrew M. Colman and Michael Bacharach. Payoff dominance and the stackelberg heuristic. *Theory and Decision*, 43(1):1–19, 1997. DOI: 10.1023/A:1004911723951. 35

[69] Vincent Conitzer and Tuomas Sandholm. A generalized strategy eliminability criterion and computational methods for applying it. In *AAAI*, pages 483–488, 2005. 35

[70] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine Learning*, 20(3):273–297, 1995. DOI: 10.1007/bf00994018. 37

[71] Miguel Costa-Gomes, Vincent P. Crawford, and Bruno Broseta. Cognition and behavior in normal-form games: An experimental study. *Econometrica*, 69(5):1193–1235, 2001. DOI: 10.1111/1468-0262.00239. 25

[72] Gary W. Cox. *Making Votes Count: Strategic Coordination in the World's Electoral Systems*, Cambridge University Press, 1997. DOI: 10.1017/cbo9781139174954. 74

[73] Vincent P. Crawford, Miguel A. Costa-Gomes, and Nagore Iriberri. Structural models of non-equilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature*, 51(1):5–62, 2013. DOI: 10.1257/jel.51.1.5. 25

[74] Cristiana Croitoru, Ilinca Nastase, Florin Bode, Amina Meslem, and Angel Dogeanu. Thermal comfort models for indoor spaces and vehicles—current capabilities and future perspectives. *Renewable and Sustainable Energy Reviews*, 44:304–318, 2015. DOI: 10.1016/j.rser.2014.10.105. 78

[75] Jared R. Curhan, Hillary Anger Elfenbein, and Heng Xu. What do people value when they negotiate? Mapping the domain of subjective value in negotiation. *Journal of Personality and Social Psychology*, 91(3):493, 2006. DOI: 10.1037/0022-3514.91.3.493. 68

[76] George Dantzig. *Linear Programming and Extensions*, Princeton University Press, 2016. DOI: 10.1515/9781400884179. 62

[77] Paulo André Lima de Castro, Anderson Rodrigo Barreto Teodoro, Luciano Irineu de Castro, and Simon Parsons. Expected utility or prospect theory: Which better fits agent-based modeling of markets? *Journal of Computational Science*, 17:97–102, 2016. DOI: 10.1016/j.jocs.2016.10.002. 35

[78] Arthur P. Dempster, Nan M. Laird, and Donald B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 1–38, 1977. 44

[79] Sam Devlin, Anastasija Anspoka, Nick Sephton, Peter I. Cowling, and Jeff Rollason. Combining gameplay data with Monte Carlo tree search to emulate human play. In *12th Artificial Intelligence and Interactive Digital Entertainment Conference*, 2016. 82

[80] Cindy Dietrich. Decision making: Factors that influence decision making, heuristics used, and decision outcomes. *Inquiries Journal*, 2(02), 2010. 2

[81] Rolf Dobelli. *The Art of Thinking Clearly: Better Thinking, Better Decisions*, Hachette UK, 2013. 30

[82] Pedro Domingos. A few useful things to know about machine learning. *Communications of the ACM*, 55(10):78–87, 2012. DOI: 10.1145/2347736.2347755. 36

[83] Eugene A. Durenard. *Professional Automated Trading: Theory and Practice*, John Wiley & Sons, 2013. DOI: 10.1002/9781118755846. 67

[84] The Economist. Israel is using social media to prevent terrorist attacks, 2016. https://www.economist.com/news/middle-east-and-africa/21697083-new-paradigm-intelligence-israel-using-social-media-prevent-terrorist 47

[85] Ido Erev and Alvin E. Roth. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review*, 88(4):848–881, 1998. 67

[86] Ido Erev, Eyal Ert, Ori Plonsky, Doron Cohen, and Oded Cohen. From anomalies to forecasts: Toward a descriptive model of decisions under risk, under ambiguity, and from experience. *Psychological Review*, pages 1–56, 2017. DOI: 10.1037/rev0000062. 55

[87] Ido Erev and Alvin E. Roth. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review*, pages 848–881, 1998. 38, 80

[88] Ido Erev, Alvin E. Roth, Robert L. Slonim, and Greg Barron. Combining a theoretical prediction with experimental evidence, 2002. DOI: 10.2139/ssrn.1111712. 55

[89] Ido Erev, Alvin E. Roth, Robert L. Slonim, and Greg Barron. Predictive value and the usefulness of game theoretic models. *International Journal of Forecasting*, 18(3):359–368, 2002. DOI: 10.1016/s0169-2070(02)00023-7. 22

[90] Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. Multiwinner voting: A new challenge for social choice theory. *Trends in Computational Social Choice*, 2017. 74

[91] Fei Fang, Thanh Hong Nguyen, Rob Pickles, Wai Y. Lam, Gopalasamy R. Clements, Bo An, Amandeep Singh, Brian C. Schwedock, Milind Tambe, and Andrew Lemieux. Paws-a deployed game-theoretic application to combat poaching. *AI Magazine*, 38(1):23–36, 2017. DOI: 10.1609/aimag.v38i1.2710. 65

[92] Fei Fang, Peter Stone, and Milind Tambe. When security games go green: Designing defender strategies to prevent poaching and illegal fishing. In *IJCAI*, pages 2589–2595, 2015. 65

[93] Noura Farra, Swapna Somasundaran, and Jill Burstein. Scoring persuasive essays using opinions and their targets. In *Proc. of the 10th Workshop on Innovative Use of NLP for Building Educational Applications*, pages 64–74, 2015. DOI: 10.3115/v1/w15-0608. 73

[94] Tom Fawcett. An introduction to ROC analysis. *Pattern Recognition Letters*, 27(8):861–874, 2006. DOI: 10.1016/j.patrec.2005.10.010. 48, 58

[95] Hein Fennema and Peter Wakker. Original and cumulative prospect theory: A discussion of empirical differences. *Journal of Behavioral Decision Making*, 10:53–64, 1997. DOI: 10.1002/(sici)1099-0771(199703)10:1%3C53::aid-bdm245%3E3.0.co;2-1. 35

[96] Sevan G. Ficici and Avi Pfeffer. Modeling how humans reason about others with partial information. In *Proc. of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems*, 1, pages 315–322, International Foundation for Autonomous Agents and Multiagent Systems, 2008. 24, 69

[97] Michael Fisher, Christian List, Marija Slavkovik, and Alan Winfield. Engineering moral agents—from human morality to artificial morality (Dagstuhl Seminar 16222). *Dagstuhl Reports*, 6(5):114–137, 2016. 90

[98] James Fogarty, Scott E. Hudson, Christopher G. Atkeson, Daniel Avrahami, Jodi Forlizzi, Sara Kiesler, Johnny C. Lee, and Jie Yang. Predicting human interruptibility with sensors. *ACM Transactions on Computer-human Interaction (TOCHI)*, 12(1):119–146, 2005. DOI: 10.1145/1057237.1057243. 47

[99] J. Kevin Ford, Neal Schmitt, Susan L. Schechtman, Brian M. Hults, and Mary L. Doherty. Process tracing methods: Contributions, problems, and neglected research questions. *Organizational Behavior and Human Decision Processes*, 43(1):75–117, 1989. DOI: 10.1016/0749-5978(89)90059-9. 2

[100] Robert Forsythe, Thomas Rietz, Roger B. Myerson, and Robert Weber. An experimental study of voting rules and polls in three-candidate elections. *International Journal of Game Theory*, 25(3):355–383, 1996. DOI: 10.1007/bf02425262. 75

[101] John Fox, David Glasspool, Dan Grecu, Sanjay Modgil, Matthew South, and Vivek Patkar. Argumentation-based inference and decision making—a medical perspective. *IEEE Intelligent Systems*, 22(6), 2007. DOI: 10.1109/mis.2007.102. 72

[102] Stan Franklin and Art Graesser. Is it an agent, or just a program?: A taxonomy for autonomous agents. *Intelligent Agents III Agent Theories, Architectures, and Languages*, pages 21–35, 1997. DOI: 10.1007/bfb0013570. 2

[103] Lex Fridman. MIT 6.s094: Deep learning for human-centered semi-autonomous vehicles, 2017. 77, 84

[104] Lex Fridman, Heishiro Toyoda, Sean Seaman, Bobbie Seppelt, Linda Angell, Joonbum Lee, Bruce Mehler, and Bryan Reimer. What can be predicted from six seconds of driver glances? In *Proc. of the CHI Conference on Human Factors in Computing Systems*, pages 2805–2813, ACM, 2017. DOI: 10.1145/3025453.3025929. 77

[105] Jon Froehlich, Leah Findlater, and James Landay. The design of eco-feedback technology. In *Proc. of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1999–2008, ACM, 2010. DOI: 10.1145/1753326.1753629. 4

[106] John D. E. Gabrieli, Satrajit S. Ghosh, and Susan Whitfield-Gabrieli. Prediction as a humanitarian and pragmatic contribution from human cognitive neuroscience. *Neuron*, 85(1):11–26, 2015. DOI: 10.1016/j.neuron.2014.10.047. 47

[107] Ya'akov Gal, Sarit Kraus, Michele Gelfand, Hilal Khashan, and Elizabeth Salmon. An adaptive agent for negotiating with people in different cultures. *ACM Transactions on Intelligent Systems and Technology*, 3(1):8, 2011. DOI: 10.1145/2036264.2036272. 68

[108] Ya'akov Gal and Avi Pfeffer. Modeling reciprocal behavior in human bilateral negotiation. In *Proc. of the National Conference on Artificial Intelligence*, 22, pages 815–820, 2007. 55

[109] Sam Ganzfried and Tuomas Sandholm. Game theory-based opponent modeling in large imperfect-information games. In *The 10th International Conference on Autonomous Agents and Multiagent Systems*, 2, pages 533–540, International Foundation for Autonomous Agents and Multiagent Systems, 2011. 80

[110] Leonardo Garrido and Katia Sycara. Multi-agent meeting scheduling: Preliminary experimental results. In *Proc. of the 2nd International Conference on Multiagent Systems*, pages 95–102, 1996. 77

[111] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The KITTI dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013. DOI: 10.1177/0278364913491297. 77

[112] Michele Gelfand, Garriy Shteynberg, Tiane Lee, Janetta Lun, Sarah Lyons, Chris Bell, Joan Y. Chiao, C. Bayan Bruss, May Al Dabbagh, Zeynep Aycan, et al. The cultural contagion of conflict. *Philosophical Transactions on Royal Society B*, 367(1589):692–703, 2012. DOI: 10.1098/rstb.2011.0304. 43

[113] Andrew Gelman, Jeffrey Fagan, and Alex Kiss. An analysis of the New York city police department's "stop-and-frisk" policy in the context of claims of racial bias. *Journal of the American Statistical Association*, 102(479):813–823, 2007. DOI: 10.1198/016214506000001040. 92

[114] Michael Genesereth, Nathaniel Love, and Barney Pell. General game playing: Overview of the AAAI competition. *AI Magazine*, 26(2):62, 2005. DOI: 10.1609/aimag.v26i2.1813. 81

[115] Joseph C. Giarratano and Gary Riley. *Expert Systems*, PWS Publishing Co., 1998. 77

[116] Allan Gibbard. Manipulation of voting schemes: A general result. *Econometrica: Journal of the Econometric Society*, pages 587–601, 1973. DOI: 10.2307/1914083. 75

[117] Elizabeth Gibney. How rival bots battled their way to poker supremacy. *Nature*, 543(7644):160–161, 2017. DOI: 10.1038/nature.2017.21580. 39

[118] Deidre Giesen, Vivian Meertens, Rachel Vis-Visschers, and Dirkjan Beukenhorst. Questionnaire development. *Technical Report*, 2012. 45

[119] Michael A. Gilbert. *How to Win an Argument: Surefire Strategies for Getting your Point Across*, John Wiley, New York, 1996. 74

[120] Itzhak Gilboa. *Theory of Decision Under Uncertainty*, Cambridge University Press Cambridge, 2009. DOI: 10.1017/cbo9780511840203. 10

[121] David Gill and Victoria L. Prowse. Cognitive ability and learning to play equilibrium: A level-k analysis. *Analysis*, 2012. DOI: 10.2139/ssrn.2043336. 25

[122] Jacob K. Goeree, Charles A. Holt, and Thomas R. Palfrey. Quantal response equilibrium and overbidding in private-value auctions. *Journal of Economic Theory*, 104(1):247–272, 2002. DOI: 10.1006/jeth.2001.2914. 23

[123] Claudia V. Goldman, Asaf Degani, Omer Deutsch, and Omer Tsimhoni. On holding and containment in HMI: Theoretical foundations, framework, and initial trials. *Technical Report*, GM, 2015. 79

[124] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*, MIT Press, 2016. `http://www.deeplearningbook.org` 37

[125] Carlos Gracia-Lázaro, Luis Mario Floría, and Yamir Moreno. Cognitive hierarchy theory and two-person games. *Games*, 8(1):1, 2017. DOI: 10.3390/g8010001. 28

[126] Phil Gray. The rise of intelligent virtual assistants. `https://www.interactions.com/rise-intelligent-virtual-assistants/` 2

[127] Barbara Grosz. What question would turing pose today? *AI Magazine*, 33(4):73, 2012. DOI: 10.1609/aimag.v33i4.2441. 1

[128] Marcello Guarini. Particularism and the classification and reclassification of moral cases. *IEEE Intelligent Systems*, 21(4):22–28, 2006. DOI: 10.1109/mis.2006.76. 92

[129] Junius Gunaratne and Oded Nov. Informing and improving retirement saving performance using behavioral economics theory-driven user interfaces. In *Proc. of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 917–920, 2015. DOI: 10.1145/2702123.2702408. 31

[130] Werner Güth, Rolf Schmittberger, and Bernd Schwarze. An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior and Organization*, 3(4):367–388, 1982. DOI: 10.1016/0167-2681(82)90011-7. 28

[131] Ivan Habernal and Iryna Gurevych. What makes a convincing argument? Empirical analysis and detecting attributes of convincingness in web argumentation. In *EMNLP*, pages 1214–1223, 2016. DOI: 10.18653/v1/d16-1129. 73

[132] Emmanuel Hadoux, Ariel Rosenfeld, and Jean-Baptiste Corrégé. From psychological persuasion to abstract argumentation: A step forward. In *Proc. of the 5th Workshop on Artificial Intelligence and Cognition (AIC)*, 2017. 72

[133] Nika Haghtalab, Fei Fang, Thanh Hong Nguyen, Arunesh Sinha, Ariel D. Procaccia, and Milind Tambe. Three strategies to success: Learning adversary models in security games. In *IJCAI*, pages 308–314, 2016. 66

[134] P. Richard Hahn, Kristian Lum, and Carl Mela. Testing cognitive hierarchy theories of beauty contest games. *Technical Report*, Duke University, Working paper, 2010. 27

[135] Galit Haim, Bo An and Sarit Kraus, Human—computer negotiation in a three player market setting. *Artificial Intelligence*, 246:34–52, 2017. DOI: 10.1016/j.artint.2017.01.003. 69

[136] Galit Haim, Ya'akov Gal, Michele Gelfand, and Sarit Kraus. A cultural sensitive agent for human-computer negotiation. In *Proc. of the 11th International Conference on Autonomous Agents and Multiagent Systems*, pages 451–458, 2012. 43

[137] Chen Hajaj, Noam Hazon, and David Sarne. Enhancing comparison shopping agents through ordering and gradual information disclosure. *Autonomous Agents and Multi-Agent Systems*, 31(3):696–714, 2017. DOI: 10.1007/s10458-016-9342-8. 28

[138] Sven Ove Hansson. Decision theory: A brief introduction, 2005. `https://www.goodreads.com/book/show/17214227-decision-theory` 7

[139] Sandra G. Hart and Lowell E. Staveland. Development of NASA-TLX (task load index): Results of empirical and theoretical research. *Advances in Psychology*, 52:139–183, 1988. DOI: 10.1016/s0166-4115(08)62386-9. 45

[140] Jason S. Hartford, James R. Wright, and Kevin Leyton-Brown. Deep learning for predicting human strategic behavior. In *Advances in Neural Information Processing Systems*, pages 2424–2432, 2016. 40

[141] Martie G. Haselton, Daniel Nettle, and Damian R. Murray. The evolution of cognitive bias. *The Handbook of Evolutionary Psychology*, 2005. DOI: 10.1002/9780470939376.ch25. 29

[142] Starke Rosecrans Hathaway and John Charnley McKinley. *Minnesota Multiphasic Personality Inventory; Manual*, Psychological Corporation, 1951. 42

[143] Haibo He and Edwardo A. Garcia. Learning from imbalanced data. *IEEE Transactions on Knowledge and Data Engineering*, 21(9):1263–1284, 2009. DOI: 10.1109/tkde.2008.239. 48

[144] He He, Jordan Boyd-Graber, Kevin Kwok, and Hal Daumé III. Opponent modeling in deep reinforcement learning. In *International Conference on Machine Learning*, pages 1804–1813, 2016. 81

[145] Pablo Hernandez-Leal, Michael Kaisers, Tim Baarslag, and Enrique Munoz de Cote. A survey of learning in multiagent environments: Dealing with non-stationarity. *arXiv Preprint arXiv:1707.09183*, 2017. 79

[146] Michael Herron and Jeffrey Lewis. Did Ralph Nader spoil a Gore presidency? A ballot-level study of green and reform party voters in the 2000 presidential election. *Quarterly Journal of Political Science*, 2(3):205–226, 2007. DOI: 10.1561/100.00005039. 75

[147] Philip Hingston. The 2k BotPrize. In *Computational Intelligence and Games, (CIG). IEEE Symposium on*, page 1, 2009. DOI: 10.1109/cig.2009.5286505. 81

[148] Philip Hingston. A turing test for computer game bots. *IEEE Transactions on Computational Intelligence and AI in Games*, 1(3):169–186, 2009. DOI: 10.1109/tciaig.2009.2032534. 81

[149] Philip Hingston. *Believable Bots: Can Computers Play Like People?*, Springer, 2012. DOI: 10.1007/978-3-642-32323-2. 81

[150] Arend Hintze and Ralph Hertwig. The evolution of generosity in the ultimatum game. *Scientific Reports*, 6:34102, 2016. DOI: 10.1038/srep34102. 29

[151] Elizabeth Hoffman, Kevin McCabe, and Vernon Smith. Reciprocity in ultimatum and dictator games: An introduction. *Handbook of Experimental Economics Results*, 1:411–416, 2008. DOI: 10.1016/s1574-0722(07)00046-7. 29

[152] Penina Hoz-Weiss, Sarit Kraus, Jonathan Wilkenfeld, David R. Andersen, and Amy Pate. Resolving crises through automated bilateral negotiations. 172(1):1–18, 2008. DOI: 10.1016/j.artint.2007.05.007. 67

[153] Wenyan Hu and Alvaro Bolivar. Online auctions efficiency: A survey of ebay auctions. In *Proc. of the 17th International Conference on World Wide Web*, pages 925–934, ACM, 2008. DOI: 10.1145/1367497.1367621. 68

[154] Anthony Hunter. Modelling the persuade in asymmetric argumentation dialogues for persuasion. In *IJCAI*, pages 3055–3061, 2015. 72

[155] Anthony Hunter. Computational persuasion with applications in behaviour change. In *COMMA*, pages 5–18, 2016. 71

[156] Mara Hvistendahl. Can "predictive policing" prevent crime before it happens. *Science Magazine*, 2016. `www.sciencemag.org/news/2016/09/can-predictive-policing-prevent-crime-it-happens` DOI: 10.1126/science.aah7353. 66

[157] Nathanael Hyafil and Craig Boutilier. Regret minimizing equilibria and mechanisms for games with strict type uncertainty. In *Proc. of the 20th Conference on Uncertainty in Artificial Intelligence*, pages 268–277, AUAI Press, 2004. 35

[158] IBM. Watson in healthcare. `https://www.ibm.com/watson/health/` 2

[159] The Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel 2002. 17

[160] Ashesh Jain, Avi Singh, Hema S. Koppula, Shane Soh, and Ashutosh Saxena. Recurrent neural networks for driver activity anticipation via sensory-fusion architecture. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 3118–3125, 2016. DOI: 10.1109/icra.2016.7487478. 47, 77

[161] Richard C. Jeffrey. *The Logic of Decision*, McGraw-Hill, 1965. 10

[162] Albert Xin Jiang and Kevin Leyton-Brown. A tutorial on the proof of the existence of nash equilibria. *University of British Columbia Technical Report TR-2007-25.pdf*, 2009. 14

[163] Eric J. Johnson, Steven Bellman, and Gerald L. Lohse. Defaults, framing and privacy: Why opting in-opting out. *Marketing Letters*, 13(1):5–15, 2002. DOI: 10.1023/A:1015044207315. 31

[164] Eric J. Johnson and Daniel Goldstein. Medicine. Do defaults save lives? *Science*, 302(5649):1338–1339, 2003. DOI: 10.1126/science.1091721. 31

[165] Philip N. Johnson-Laird and Ruth M. J. Byrne. Conditionals: A theory of meaning, pragmatics, and inference. *Psychological Review*, 109(4):646, 2002. DOI: 10.1037/0033-295x.109.4.646. 17

[166] Catholijn M. Jonker, Valentin Robu, and Jan Treur. An agent architecture for multi-attribute negotiation using incomplete preference information. 15(2):221–252, 2007. DOI: 10.1007/s10458-006-9009-y. 68

[167] E Ásgeir Juliusson, Niklas Karlsson, and Tommy Gärling. Weighing the past and the future in decision making. *European Journal of Cognitive Psychology*, 17(4):561–575, 2005. DOI: 10.1080/09541440440000159. 2

[168] Daniel Kahneman. A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist*, 58(9):697, 2003. DOI: 10.1037/0003-066x.58.9.697. 21

[169] Daniel Kahneman. *Thinking, Fast and Slow*, Macmillan, 2011. 30

[170] Daniel Kahneman, Paul Slovic, and Amos Tversky. *Judgment Under Uncertainty: Heuristics and Biases*, Cambridge University Press, 1982. DOI: 10.1017/cbo9780511809477. 28

[171] Daniel Kahneman and Vernon Smith. Foundations of behavioral and experimental economics. *Nobel Prize in Economics Documents*, 1, 2002. 7

[172] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society*, pages 263–291, 1979. DOI: 10.2307/1914185. 33

[173] Miltiadis Kandias, Alexios Mylonas, Nikos Virvilis, Marianthi Theoharidou, and Dimitris Gritzalis. An insider threat prediction model. In *International Conference on Trust, Privacy and Security in Digital Business*, pages 26–37, Springer, 2010. DOI: 10.1007/978-3-642-15152-1_3. 48

[174] Rosabeth Moss Kanter. Ten reasons people resist change. *Harvard Business Review*, 2012. 74

[175] Frédéric Kaplan. Who is afraid of the humanoid? Investigating cultural differences in the acceptance of robots. *International Journal of Humanoid Robotics*, 1(03):465–480, 2004. DOI: 10.1142/s0219843604000289. 91

[176] Debarun Kar, Fei Fang, Francesco M. Delle Fave, Nicole Sintov, Milind Tambe, and Arnaud Lyet. Comparing human behavior models in repeated stackelberg security games: An extended study. *Artificial Intelligence*, 240:65–103, 2016. DOI: 10.1016/j.artint.2016.08.002. 65

[177] Debarun Kar, Benjamin Ford, Shahrzad Gholami, Fei Fang, Andrew Plumptre, Milind Tambe, Margaret Driciru, Fred Wanyama, Aggrey Rwetsiba, Mustapha Nsubaga, et al. Cloudy with a chance of poaching: Adversary behavior modeling and forecasting with real-world poaching data. In *Proc. of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 159–167, 2017. 65, 88

[178] Sergey Karakovskiy and Julian Togelius. The Mario AI benchmark and competitions. *IEEE Transactions on Computational Intelligence and AI in Games*, 4(1):55–67, 2012. DOI: 10.1109/tciaig.2012.2188528. 4, 81

[179] Leo Kent. The humanized driving project, 2015. `https://www.citylab.com/life/2015/02/teaching-driverless-cars-to-fit-your-human-driving-style/386369/` 78

[180] Gregory E. Kersten and Hsiangchu Lai. Negotiation support and e-negotiation systems: An overview. *Group Decision and Negotiation*, 16(6):553–586, 2007. DOI: 10.1007/s10726-007-9095-5. 67

[181] Gregory E. Kersten and Gordon Lo. Aspire: An integrated negotiation support system and software agents for e-business negotiation. *International Journal of Internet and Enterprise Management*, 1(3):293–315, 2003. DOI: 10.1504/ijiem.2003.003822. 67

[182] John Maynard Keynes. *General Theory of Employment, Interest and Money*, 1936. 27

[183] Ahmed Khalifa, Aaron Isaksen, Julian Togelius, and Andy Nealen. Modifying MCTS for human-like general video game playing. In *IJCAI*, pages 2514–2520, 2016. 82

[184] Jon Kleinberg, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan. Human decisions and machine predictions. *Technical Report*, National Bureau of Economic Research, 2017. DOI: 10.3386/w23180. 44

[185] W. Bradley Knox and Peter Stone. Combining manual feedback with subsequent MDP reward signals for reinforcement learning. In *Proc. of AAMAS*, May 2010. 81

[186] Hema S. Koppula and Ashutosh Saxena. Anticipating human activities using object affordances for reactive robotic response. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(1):14–29, 2016. DOI: 10.15607/rss.2013.ix.006. 50

[187] Lars Kotthoff, Chris Thornton, Holger H. Hoos, Frank Hutter, and Kevin Leyton-Brown. Auto-WEKA 2.0: Automatic model selection and hyperparameter optimization in WEKA. *Journal of Machine Learning Research*, 18:1–5, 2017. 36

[188] John Kratz and Carly Strasser. Data publication consensus and controversies. *F1000 Research*, 3, 2014. DOI: 10.12688/f1000research.3979.1. 41

[189] Sarit Kraus and Daniel Lehmann. Designing and building a negotiating automated agent. *Computational Intelligence*, 11(1):132–171, 1995. DOI: 10.1111/j.1467-8640.1995.tb00026.x. 68

[190] Christopher Krupenye, Fumihiro Kano, Satoshi Hirata, Josep Call, and Michael Tomasello. Great apes anticipate that other individuals will act according to false beliefs. *Science*, 354(6308):110–114, 2016. DOI: 10.1126/science.aaf8110. 1

[191] Srijan Kumar, Justin Cheng, and Jure Leskovec. Antisocial behavior on the Web: Characterization and detection. In *Proc. of the 26th International Conference on World Wide Web Companion*, pages 947–950, 2017. 47

[192] Himabindu Lakkaraju, Jon Kleinberg, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan. The selective labels problem: Evaluating algorithmic predictions in the presence of unobservables. In *Proc. of the International Conference on Knowledge Discovery and Data Mining (KDD)*, 2017. DOI: 10.1145/3097983.3098066. 44

[193] John Langford. The NIPS experiment. *Communications of the ACM*, 2015. 94

[194] Jean-François Laslier. The leader rule: A model of strategic approval voting in a large electorate. *Journal of Theoretical Politics*, 21(1):113–136, 2009. DOI: 10.1177/0951629808097286. 75

[195] Bryan Lawrence, Catherine Jones, Brian Matthews, Sam Pepler, and Sarah Callaghan. Citation and peer review of data: Moving towards formal data publication. *International Journal of Digital Curation*, 6(2):4–37, 2011. DOI: 10.2218/ijdc.v6i2.205. 41

[196] Alessandro Lazaric, Mario Quaresimale, and Marcello Restelli. On the usefulness of opponent modeling: The Kuhn poker case study. In *Proc. of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 1345–1348, 2008. 63

[197] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015. DOI: 10.1038/nature14539. 38

[198] Min Kyung Lee, Sara Kiesler, and Jodi Forlizzi. Mining behavioral economics to design persuasive technology for healthy choices. In *Proc. of the SIGCHI Conference on Human Factors in Computing Systems*, pages 325–334, ACM, 2011. DOI: 10.1145/1978942.1978989. 31

[199] Irwin P. Levin, Sandra L. Schneider, and Gary J. Gaeth. All frames are not created equal: A typology and critical analysis of framing effects. *Organizational Behavior and Human Decision Processes*, 76(2):149–188, 1998. DOI: 10.1006/obhd.1998.2804. 32

[200] Li Li, Ding Wen, Nan-Ning Zheng, and Lin-Cheng Shen. Cognitive cars: A new frontier for ADAS research. *IEEE Transactions on Intelligent Transportation Systems*, 13(1):395–407, 2012. DOI: 10.1109/tits.2011.2159493. 78

[201] Patrick Lin. Why ethics matters for autonomous cars. In *Autonomous Driving*, pages 69–85, Springer, 2016. DOI: 10.1007/978-3-662-48847-8_4. 79

[202] Raz Lin, Ya'akov Gal, Sarit Kraus, and Yaniv Mazliah. Training with automated agents improves people's behavior in negotiation and coordination tasks. *Decision Support Systems*, (to appear), 2012. DOI: 10.1016/j.dss.2013.05.015. 67

[203] Raz Lin, Sarit Kraus, Jonathan Wilkenfeld, and James Barry. Negotiating with bounded rational agents in environments with incomplete information using an automated agent. 172(6-7):823–851, 2008. DOI: 10.1016/j.artint.2007.09.007. 68

[204] John Lisman. The challenge of understanding the brain: Where we stand in 2015. *Neuron*, 86(4):864–882, 2015. DOI: 10.1016/j.neuron.2015.03.032. 2

[205] George Loewenstein. Emotions in economic theory and economic behavior. *American Economic Review*, 90(2):426–432, 2000. DOI: 10.1257/aer.90.2.426. 21

[206] Simon M. Lucas, Michael Mateas, Mike Preuss, Pieter Spronck, and Julian Togelius. Artificial and computational intelligence in games: Integration (Dagstuhl seminar 15051). In *Dagstuhl Reports*, 5, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2015. DOI: 10.4230/DagRep.5.1.207. 80

[207] R. Duncan Luce. *Individual Choice Behavior: A Theoretical Analysis*, Courier Corporation, 2005. DOI: 10.1037/14396-000. 69

[208] *Science Magazine*. From AI to protein folding: Our breakthrough runners-up, 2016. 39

[209] Andreu Mas-Colell, Michael Dennis Whinston, Jerry R. Green, et al. *Microeconomic Theory*, 1, Oxford University Press, New York, 1995. 16

[210] M. Maschler, Eilon Solan, and Shmuel Zamir. *Game Theory*, Cambridge University Press, 2013. DOI: 10.1017/cbo9780511794216. 16

[211] James T. Sincich McClave, Terry James T. McClave, and Terry Sincich. *Statistics*, Number QA 276.12. M33 2006, 2006. 87

[212] Peter McCracken and Michael Bowling. Safe strategies for agent modelling in games. In *AAAI Fall Symposium on Artificial Multi-agent Learning*, pages 103–110, 2004. 63

[213] Richard D. McKelvey and Thomas R. Palfrey. An experimental study of the centipede game. *Econometrica*, 60(4):803–836, 1992. DOI: 10.2307/2951567. 67

[214] Richard D. McKelvey and Thomas R. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6–38, 1995. DOI: 10.1006/game.1995.1023. 23, 24

[215] Johnathan Mell and Jonathan Gratch. IAGO: Interactive arbitration guide online. *Proc. of the International Conference on Autonomous Agents and Multiagent Systems*, pages 1510–1512, 2016. 71

[216] Paul E. Meehl. Clinical vs. statistical prediction: A theoretical analysis and a review of the evidence, 1954. `https://www.amazon.com/Clinical-Versus-Statistical-Prediction-Theoretical/dp/0963878492` DOI: 10.1037/11281-000. 36

[217] Hugo Mercier and Dan Sperber. *The Enigma of Reason*, Harvard University Press, 2017. 74

[218] Ryszard S. Michalski, Jaime G. Carbonell, and Tom M. Mitchell. *Machine Learning: An Artificial Intelligence Approach*, Springer Science and Business Media, 2013. DOI: 10.1007/978-3-662-12405-5. 36, 87

[219] George A. Miller. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63(2):81, 1956. DOI: 10.1037//0033-295x.101.2.343. 21

[220] Tom Michael Mitchell. *The Discipline of Machine Learning*, 3, Carnegie Mellon University, School of Computer Science, Machine Learning Department, 2006. 36

[221] Matej Moravčík, Martin Schmid, Neil Burch, Viliam Lisy, Dustin Morrill, Nolan Bard, Trevor Davis, Kevin Waugh, Michael Johanson, and Michael Bowling. Deepstack: Expert-level artificial intelligence in heads-up no-limit poker. *Science*, 2017. DOI: 10.1126/science.aam6960. 39

[222] Matej Moravčík, Martin Schmid, Neil Burch, Viliam Lisy, Dustin Morrill, Nolan Bard, Trevor Davis, Kevin Waugh, Michael Johanson, and Michael Bowling. Deepstack: Expert-level artificial intelligence in no-limit poker. *arXiv Preprint arXiv:1701.01724*, 2017. DOI: 10.1126/science.aam6960. 3

[223] Mark A. Musen, Blackford Middleton, and Robert A. Greenes. Clinical decision-support systems. In *Biomedical Informatics*, pages 643–674, Springer, 2014. DOI: 10.1007/978-1-4471-4474-8_22. 49

[224] Rosemarie Nagel. Unraveling in guessing games: An experimental study. *The American Economic Review*, 85(5):1313–1326, 1995. 24

[225] John Nash. Non-cooperative games. *Annals of Mathematics*, pages 286–295, 1951. DOI: 10.2307/1969529. 14

[226] John F. Nash et al. Equilibrium points in n-person games. *Proc. of the National Academy of Sciences*, 36(1):48–49, 1950. DOI: 10.1073/pnas.36.1.48. 14

[227] Andrew Y. Ng, Stuart J. Russell, et al. Algorithms for inverse reinforcement learning. In *ICML*, pages 663–670, 2000. 38

[228] Thanh Hong Nguyen, Rong Yang, Amos Azaria, Sarit Kraus, and Milind Tambe. Analyzing the effectiveness of adversary modeling in security games. In *AAAI*, 2013. 65, 83

[229] Raymond S. Nickerson. Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2(2):175, 1998. DOI: 10.1037//1089-2680.2.2.175. 55

[230] Gali Noti, Effi Levi, Yoav Kolumbus, and Amit Daniely. Behavior-based machine-learning: A hybrid approach for predicting human decision making. *arXiv Preprint arXiv:1611.10228*, 2016. 55

[231] Anastasios Noulas, Salvatore Scellato, Renaud Lambiotte, Massimiliano Pontil, and Cecilia Mascolo. A tale of many cities: Universal patterns in human urban mobility. *PloS One*, 7(5):e37027, 2012. DOI: 10.1371/journal.pone.0037027. 32

[232] Bridianne O'Dea, Stephen Wan, Philip J. Batterham, Alison L. Calear, Cecile Paris, and Helen Christensen. Detecting suicidality on twitter. *Internet Interventions*, 2(2):183–188, 2015. DOI: 10.1016/j.invent.2015.03.005. 46

[233] The Federal Bureau of Investigation. The insider threat an introduction to detecting and deterring an insider spy, 2017. `https://www.fbi.gov/file-repository/insider_threat_brochure.pdf/view` 47

[234] New-Zealand Department of Statistics. A guide to good survey design (4th ed.), *Technical Report*, 2015. 45

[235] S. Orguz Oguz, Ayse Kucukyilmaz, Tevfik Sezgin, and Cagatay Basdogan. Supporting negotiation behavior with haptics-enabled human-computer interfaces. *IEEE Transactions on Haptics*, 5(3):274–284, 2012. DOI: 10.1109/toh.2012.37. 68

[236] Judith Reitman Olson and Henry H. Rueter. Extracting expertise from experts: Methods for knowledge acquisition. *Expert Systems*, 4(3):152–168, 1987. DOI: 10.1111/j.1468-0394.1987.tb00139.x. 35

[237] Hessel Oosterbeek, Randolph Sloof, and Gijs Van De Kuilen. Cultural differences in ultimatum game experiments: Evidence from a meta-analysis. *Experimental Economics*, 7(2):171–188, 2004. DOI: 10.1023/b:exec.0000026978.14316.74. 51

[238] Juan Ortega, Noor Shaker, Julian Togelius, and Georgios N. Yannakakis. Imitating human playing styles in super Mario Bros. *Entertainment Computing*, 4(2):93–104, 2013. DOI: 10.1016/j.entcom.2012.10.001. 81

[239] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*, MIT Press, 1994. 11

[240] Yinon Oshrat, Raz Lin, and Sarit Kraus. Facing the challenge of human-agent negotiations via effective general opponent modeling. In *Proc. of the 8th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 377–384, Budapest, Hungary, 2009. 69

[241] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359, 2010. DOI: 10.1109/tkde.2009.191. 53, 92

[242] Susan Wells Parham, Jake Carlson, Patricia Hswe, Brian Westra, and Amanda Whitmire. Using data management plans to explore variability in research data management practices across domains. *International Journal of Digital Curation*, 11(1):53–67, 2016. DOI: 10.2218/ijdc.v11i1.423. 41

[243] Noseong Park, Edoardo Serra, Tom Snitch, and V. S. Subrahmanian. APE: A data-driven, behavioral model-based anti-poaching engine. *IEEE Transactions on Computational Social Systems*, 2(2):15–37, 2015. DOI: 10.1109/tcss.2016.2517452. 66

[244] Noam Peled, Ya'akov Gal, and Sarit Kraus. A study of computational and human strategies in revelation games, pages 345–352, 2011. DOI: 10.1007/s10458-014-9253-5. 67, 69

[245] Noam Peled, Ya'akov Gal, and Sarit Kraus. An agent design for repeated negotiation and information revelation with people, *Proc. of the 27th AAAI Conference on Artificial Intelligence*, 2013. 70

[246] Noam Peled, Moshe Bitan, Joseph Keshet, and Sarit Kraus. Predicting human strategic decisions using facial expressions. In *IJCAI*, 2013. 47, 81

[247] Bei Peng, James MacGlashan, Robert Loftin, Michael L. Littman, David L. Roberts, and Matthew E. Taylor. A need for speed: Adapting agent action speed to improve task learning from non-expert humans. In *AAMAS*, pages 957–965, 2016. 81

[248] Diego Perez-Liebana, Spyridon Samothrakis, Julian Togelius, Tom Schaul, Simon M. Lucas, Adrien Couëtoux, Jerry Lee, Chong-U Lim, and Tommy Thompson. The 2014 general video game playing competition. *IEEE Transactions on Computational Intelligence and AI in Games*, 8(3):229–243, 2016. DOI: 10.1109/tciaig.2015.2402393. 82

[249] Walt L. Perry. *Predictive Policing: The Role of Crime Forecasting in Law Enforcement Operations*, Rand Corporation, 2013. 66, 92

[250] Isaac Persing and Vincent Ng. Modeling argument strength in student essays. In *ACL (1)*, pages 543–552, 2015. DOI: 10.3115/v1/p15-1053. 73

[251] James Pita, Manish Jain, Janusz Marecki, Fernando Ordóñez, Christopher Portway, Milind Tambe, Craig Western, Praveen Paruchuri, and Sarit Kraus. Deployed armor protection: The application of a game theoretic model for security at the Los Angeles International Airport. In *Proc. of the International Conference on Autonomous Agents and Multiagent Systems*, pages 125–132, 2008. DOI: 10.1017/cbo9780511973031.004. 5, 22, 35, 64

[252] Ori Plonsky, Ido Erev, Tamir Hazan, and Moshe Tennenholtz. Psychological forest: Predicting human behavior. In *AAAI*, pages 656–662, 2017. DOI: 10.2139/ssrn.2816450. 55

[253] Scott Plous. *The Psychology of Judgment and Decision Making*, McGraw-Hill Book Company, 1993. 30

[254] Joe Plummer, Stephen D. Rappaport, Taddy Hall, and Robert Barocci. *The online advertising playbook: Proven strategies and tested tactics from the advertising research foundation*, John Wiley & Sons, 2007. 50

[255] Marc J. V. Ponsen, Geert Gerritsen, and Guillaume Chaslot. Integrating opponent models with Monte-Carlo tree search in poker. In *Interactive Decision Theory and Game Theory*, 2010. 82

[256] William Poundstone. *Gaming the Vote: Why Elections aren't Fair (and What We Can Do About It)*, Hill and Wang Press, New York, 2008. 75

[257] Steffen Priesterjahn, Oliver Kramer, Alexander Weimer, and Andreas Goebels. Evolution of reactive rules in multi player computer games based on imitation. In *International Conference on Natural Computation*, pages 744–755, Springer, 2005. DOI: 10.1007/11539117_105. 81

[258] Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, John Wiley & Sons, 2014. 62

[259] J. Ross Quinlan. Induction of decision trees. *Machine Learning*, 1(1):81–106, 1986. DOI: 10.1007/bf00116251. 37

[260] Iyad Rahwan, Mohammed I. Madakkatel, Jean-François Bonnefon, Ruqiyabi N. Awan, and Sherief Abdallah. Behavioral experiments for assessing the abstract argumentation semantics of reinstatement. *Cognitive Science*, 34(8):1483–1502, 2010. DOI: 10.1111/j.1551-6709.2010.01123.x. 72

[261] Simon Ramm, Joseph Giacomin, Duncan Robertson, and Alessio Malizia. A first approach to understanding and measuring naturalness in driver-car interaction. In *Proc. of the 6th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, pages 1–10, ACM, 2014. DOI: 10.1145/2667317.2667416. 78

[262] Michel Regenwetter and Ilia Tsetlin. Approval voting and positional voting methods: Inference, relationship, examples. *Social Choice and Welfare*, 22(3):539–566, 2004. DOI: 10.1007/s00355-003-0232-z. 75

[263] Nicholas Rhinehart and Kris M. Kitani. First-person forecasting with online inverse reinforcement learning. In *Proc. of the International Conference on Computer Vision (ICCV)*, 2017. 50

[264] Alan E. Rorie and William T. Newsome. A general mechanism for decision-making in the human brain? *Trends in Cognitive Sciences*, 9(2):41–43, 2005. DOI: 10.1016/j.tics.2004.12.007. 2

[265] Ariel Rosenfeld, Noa Agmon, Oleg Maksimov, and Sarit Kraus. Intelligent agent supporting human-multi-robot team collaboration, *Artificial Intelligence Journal*, 252, pages 211–231, 2017. DOI: 10.1016/j.artint.2017.08.005. 45, 78

[266] Ariel Rosenfeld, Amos Azaria, Sarit Kraus, Claudia V. Goldman, and Omer Tsimhoni. Adaptive advice in automobile climate control systems. In *Proc. of the International Conference on Autonomous Agents and Multiagent Systems, AAMAS*, pages 543–551, 2015. 49, 79

[267] Ariel Rosenfeld, Joseph Keshet, Sarit Goldman, and Claudia V. Kraus. Online prediction of exponential decay time series with human-agent application. In *Proc. of the 22nd European Conference on Artificial Intelligence*, pages 595–603, 2016. 49, 79

[268] Ariel Rosenfeld and Sarit Kraus. Argumentation theory in the field: An empirical study of fundamental notions. In *Proc. of the Workshop on Frontiers and Connections Between Argumentation Theory and Natural Language Processing*, 2014. 55, 72

[269] Ariel Rosenfeld and Sarit Kraus. Providing arguments in discussions on the basis of the prediction of human argumentative behavior. *ACM Transactions on Interactive Intelligent Systems*, 6(4):30:1–30:33, December 2016. DOI: 10.1145/2983925. 53, 55, 72, 73

[270] Ariel Rosenfeld and Sarit Kraus. Strategical argumentative agent for human persuasion. In *Proc. of the 22nd European Conference on Artificial Intelligence 29*, 285, pages 320–328, 2016. 73

[271] Ariel Rosenfeld and Sarit Kraus. When security games hit traffic: Optimal traffic enforcement under one sided uncertainty. In *Proc. of the 26th International Conference on Artificial Intelligence, (IJCAI)*, 2017. DOI: 10.24963/ijcai.2017/533. 65

[272] Ariel Rosenfeld, Oleg Maksimov, and Sarit Kraus. Optimizing traffic enforcement: From the lab to the roads. In *Proc. of the 8th Conference on Decision and Game Theory for Security*, 2017. DOI: 10.1007/978-3-319-68711-7_1. 65

[273] Ariel Rosenfeld, Matthew E. Taylor, and Sarit Kraus. Leveraging human knowledge in tabular reinforcement learning: A study of human subjects. In *Proc. of the 26th International Joint Conference on Artificial Intelligence*, pages 3823–3830, 2017. DOI: 10.24963/ijcai.2017/534. 81

[274] Avi Rosenfeld, Inon Zuckerman, Amos Azaria, and Sarit Kraus. Combining psychological models with machine learning to better predict people's decisions. *Synthese*, 189(1):81–93, 2012. DOI: 10.1007/s11229-012-0182-z. 56, 70

[275] Avi Rosenfeld, Inon Zuckerman, Erel Segal-Halevi, Osnat Drein, and Sarit Kraus. Negochat: A chat-based negotiation agent. In *Proc. of the International Conference on Autonomous Agents and Multi-agent Systems*, pages 525–532, 2014. DOI: 10.13140/2.1.3619.7128. 69

[276] Avi Rosenfeld, Inon Zuckerman, Erel Segal-Halevi, Osnat Drein, and Sarit Kraus. Negochat-a: A chat-based negotiation agent with bounded rationality. *Autonomous Agents and Multi-agent Systems*, 30(1):60–81, 2016. DOI: 10.1007/s10458-015-9281-9. 30, 55, 69

[277] Ariel Rubinstein. The electronic mail game: Strategic behavior under "almost common knowledge." *The American Economic Review*, pages 385–391, 1989. DOI: 10.1017/cbo9780511983474.019. 57

[278] Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*, 3rd ed., Pearson Education, 2010. 2, 63

[279] Larry Samuelson. Economic theory and experimental economics. *Journal of Economic Literature*, 43(1):65–107, 2005. DOI: 10.1257/0022051053737816. 54

[280] Laurie R. Santos and Alexandra G. Rosati. The evolutionary roots of human decision making. *Annual Review of Psychology*, 66:321–347, 2015. DOI: 10.1146/annurev-psych-010814-015310. 17

[281] R. Santos, G. Marreiros, C. Ramos, J. Neves, and J. Bulas-Cruz. Personality, emotion, and mood in agent-based group decision making. *IEEE Intelligent Systems*, 26(6):58–66, 2011. DOI: 10.1109/mis.2011.92. 68

[282] Mark Allen Satterthwaite. Strategy-proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975. DOI: 10.1016/0022-0531(75)90050-2. 75

[283] Leonard J. Savage. *The Foundations of Statistics*, John Wiley & Sons, New York, 1953. 10

[284] Salvatore Scellato, Cecilia Mascolo, Mirco Musolesi, and Jon Crowcroft. Track globally, deliver locally: Improving content delivery networks by tracking geographic social cascades. In *Proc. of the 20th International Conference on World Wide Web*, pages 457–466, 2011. DOI: 10.1145/1963405.1963471. 33

[285] Jonathan Schaeffer, Neil Burch, Yngvi Björnsson, Akihiro Kishimoto, Martin Müller, Robert Lake, Paul Lu, and Steve Sutphen. Checkers is solved. *Science*, 317(5844):1518–1522, 2007. DOI: 10.1126/science.1144079. 79

[286] Thomas C. Schelling. *The Strategy of Conflict*, Harvard University Press, 1963. 55, 57

[287] Oliver Scheuer, Frank Loll, Niels Pinkwart, and Bruce M. McLaren. Computer-supported argumentation: A review of the state of the art. *International Journal of Computer-supported Collaborative Learning*, 5(1):43–102, 2010. DOI: 10.1007/s11412-009-9080-x. 71

[288] Reinhard Selten. Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4(1):25–55, 1975. DOI: 10.1007/bf01766400. 22

[289] Reinhard Selten. Aspiration adaptation theory. *Journal of Mathematical Psychology*, 42(2-3):191–214, 1998. DOI: 10.1006/jmps.1997.1205. 55, 57, 69

[290] Burr Settles. Active learning literature survey. *Technical Report*, 2010. 45

[291] Siamak F. Shahandashti. Electoral systems used around the world. *Real-world Electronic Voting: Design, Analysis and Deployment*, page 77, 2016. 75

[292] Noor Shaker, Julian Togelius, Georgios N. Yannakakis, Likith Poovanna, Vinay S. Ethiraj, Stefan J. Johansson, Robert G. Reynolds, Leonard K. Heether, Tom Schumann, and Marcus Gallagher. The Turing test track of the 2012 Mario AI championship: Entries and evaluation. In *Computational Intelligence in Games (CIG), IEEE Conference on*, pages 1–8, 2013. DOI: 10.1109/cig.2013.6633634. 81

[293] Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*, Cambridge University Press, 2014. DOI: 10.1017/cbo9781107298019. 36

[294] Richard Sharp. *The Game of Diplomacy*, 1978. 68

[295] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016. DOI: 10.1038/nature16961. 3, 39, 79, 81, 82

[296] Herbert A. Simon. Rational choice and the structure of the environment. *Psychological Review*, 63(2):129, 1956. DOI: 10.1037/h0042769. 22

[297] Arunesh Sinha, Debarun Kar, and Milind Tambe. Learning adversary behavior in security games: A PAC model perspective. In *Proc. of the International Conference on Autonomous Agents and Multiagent Systems*, pages 214–222, 2016. 66

[298] James E. Smith and Robert L. Winkler. The optimizer's curse: Skepticism and post-decision surprise in decision analysis. *Management Science*, 52(3):311–322, 2006. DOI: 10.1287/mnsc.1050.0451. 62, 63

[299] Paulo E. U. Souza, Caroline P. Carvalho Chanel, Frédéric Dehais, and Sidney Givigi. Towards human-robot interaction: A framing effect experiment. In *International Conference on Systems, Man, and Cybernetics (SMC)*, pages 1929–1934, IEEE, 2016. DOI: 10.1109/smc.2016.7844521. 32

[300] Dale O. Stahl and Paul W. Wilson. Experimental evidence on players' models of other players. *Journal of Economic Behavior and Organization*, 25(3):309–327, 1994. DOI: 10.1016/0167-2681(94)90103-1. 24

[301] Dale O. Stahl and Paul W. Wilson. On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 10(1):218–254, 1995. DOI: 10.1006/game.1995.1031. 24

[302] V. S. Subrahmanian and Srijan Kumar. Predicting human behavior: The next frontiers. *Science*, 355(6324):489–489, 2017. DOI: 10.1126/science.aam7032. 61

[303] Gita Sukthankar, Christopher Geib, Hung Hai Bui, David Pynadath, and Robert P. Goldman. *Plan, Activity, and Intent Recognition: Theory and Practice*, Newnes, 2014. 5

[304] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*, MIT Press Cambridge, 1998. DOI: 10.1109/tnn.1998.712192. 38

[305] Maor Tal, Reshef Meir, and Ya'akov Gal. A study of human behavior in online voting. In *Proc. of the International Conference on Autonomous Agents and Multiagent Systems*, pages 665–673, 2015. 31, 76

[306] Milind Tambe. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*, Cambridge University Press, 2011. DOI: 10.1017/cbo9780511973031. 63

[307] Chenhao Tan, Vlad Niculae, Cristian Danescu-Niculescu-Mizil, and Lillian Lee. Winning arguments: Interaction dynamics and persuasion strategies in good-faith online discussions. In *Proc. of the 25th International Conference on World Wide Web*, pages 613–624, 2016. DOI: 10.1145/2872427.2883081. 73

[308] Yuqing Tang, Elizabeth Sklar, and Simon Parsons. An argumentation engine: Argtrust. In *9th International Workshop on Argumentation in Multiagent Systems*, 2012. 72

[309] Alan D. Taylor. *Social Choice and the Mathematics of Manipulation*, Cambridge University Press, 2005. DOI: 10.1017/cbo9780511614316. 75

[310] Matthew E. Taylor and Peter Stone. Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research*, 10(Jul):1633–1685, 2009. 53

[311] Ellen Taylor-Powell and Marcus Renner. *Collecting Evaluation Data: End-of-Session Questionnaires*, University of Wisconsin—Extension, Cooperative Extension, Program Development and Evaluation, 2009. 45

[312] Sila Temsiririrkkul, Naoyuki Sato, Kenta Nakagawa, and Kokolo Ikeda. Survey of how human players divert in-game actions for other purposes: Towards human-like computer players. In *International Conference on Entertainment Computing*, pages 243–256, Springer, 2017. DOI: 10.1007/978-3-319-66715-7_27. 81

[313] Christian Thurau, Christian Bauckhage, and Gerhard Sagerer. Imitation learning at all levels of game-AI. In *Proc. of the International Conference on Computer Games, Artificial Intelligence, Design and Education*, 5, 2004. 80

[314] Julian Togelius, Renzo De Nardi, and Simon M. Lucas. Towards automatic personalised content creation for racing games. In *Computational Intelligence and Games, (CIG). IEEE Symposium on*, pages 252–259, 2007. DOI: 10.1109/cig.2007.368106. 78

[315] David Traum, William Swartout, Stacy Marsella, and Jonathan Gratch. Fight, flight, or negotiate: Believable strategies for conversing under crisis. In *Intelligent Virtual Agents*, pages 52–64, 2005. DOI: 10.1007/11550617_5. 67

[316] Jason Tsai, Shyamsunder Rathi, Christopher Kiekintveld, Fernando Ordóñez, Milind Tambe, and Milind Tambe. Iris—a tool for strategic security allocation in transportation networks. In *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*, page 88–106, Cambridge University Press, 2011. DOI: 10.1017/cbo9780511973031.005. 65

[317] Alan M. Turing. Computing machinery and intelligence. *Mind*, 59(236):433–460, 1950. DOI: 10.1093/mind/lix.236.433. 1, 5

[318] Amos Tversky and Daniel Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323, 1992. DOI: 10.1007/978-3-319-20451-2_24. 35

[319] Alessandro Vinciarelli, Anna Esposito, Elisabeth André, Francesca Bonin, Mohamed Chetouani, Jeffrey F. Cohn, Marco Cristani, Ferdinand Fuhrmann, Elmer Gilmartin, Zakia Hammal, et al. Open challenges in modelling, analysis and synthesis of human behaviour in human—human and human—machine interactions. *Cognitive Computation*, 7(4):397–413, 2015. DOI: 10.1007/s12559-015-9326-z. 3

[320] John Von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*, Princeton University Press, 1944. 9, 11

[321] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Anticipating visual representations from unlabeled video. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 98–106, 2016. DOI: 10.1109/cvpr.2016.18. 40

[322] Douglas Walton. Argumentation theory: A very short introduction. *Argumentation in Artificial Intelligence*, pages 1–22, 2009. DOI: 10.1007/978-0-387-98197-0_1. 55, 72

[323] Pin Wang and Ching-Yao Chan. Formulation of deep reinforcement learning architecture toward autonomous driving for on-ramp merge. *arXiv Preprint arXiv:1709.02066*, 2017. 77

[324] Duncan J. Watts. *Everything is Obvious:* Once You Know the Answer*, Crown Business, 2011. 30

[325] Marc Weber. Where to? A history of autonomous vehicles, 2014. 77

[326] Zhongyu Wei, Yang Liu, and Yi Li. Is this post persuasive? Ranking argumentative comments in the online forum. In *The 54th Annual Meeting of the Association for Computational Linguistics*, page 195, 2016. DOI: 10.18653/v1/p16-2032. 73

[327] Joseph Weizenbaum. Eliza—a computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9(1):36–45, 1966. DOI: 10.1145/357980.357991. 90

[328] Joseph Weizenbaum. Computer power and human reason: From judgment to calculation, W. H. Freeman and Co., New York, 1976. DOI: 10.1063/1.3037375. 90

[329] Janis Whitlock, Celeste Pietrusza, and Amanda Purington. Young adult respondent experiences of disclosing self-injury, suicide-related behavior, and psychological distress in a web-based survey. *Archives of Suicide Research*, 17(1):20–32, 2013. DOI: 10.1080/13811118.2013.748405. 92

[330] WHO. Suicide data, 2017. `http://www.who.int/mental_health/prevention/suicide/suicideprevent/en/` 46

[331] Ian H. Witten, Eibe Frank, Mark A. Hall, and Christopher J. Pal. *Data Mining: Practical Machine Learning Tools and Techniques*, Morgan Kaufmann, 2016. DOI: 10.1145/507338.507355. xiii, 36

[332] Michael Wooldridge. *An Introduction to Multiagent Systems*, John Wiley & Sons, 2009. 63

[333] James R. Wright and Kevin Leyton-Brown. Beyond equilibrium: Predicting human behavior in normal-form games. In *AAAI*, 2010. DOI: 10.1145/1807406.1807449. 28

[334] James R. Wright and Kevin Leyton-Brown. Level-0 meta-models for predicting human behavior in games. In *Proc. of the 15th ACM Conference on Economics and Computation*, pages 857–874, 2014. DOI: 10.1145/2600057.2602907. 25

[335] Bernhard Wymann, Eric Espié, Christophe Guionneau, Christos Dimitrakakis, Rémi Coulom, and Andrew Sumner. TORCS, the open racing car simulator, 2017. `https://sourceforge.net/projects/torcs/` 78

[336] Huazhe Xu, Yang Gao, Fisher Yu, and Trevor Darrell. End-to-end learning of driving models from large-scale video datasets. *arXiv Preprint arXiv:1612.01079*, 2016. DOI: 10.1109/cvpr.2017.376. 77

[337] Rui Xu and Donald Wunsch. Survey of clustering algorithms. *IEEE Transactions on Neural Networks*, 16(3):645–678, 2005. DOI: 10.1109/tnn.2005.845141. 37

[338] Amulya Yadav, Bryan Wilder, Eric Rice, Robin Petering, Jaih Craddock, Amanda Yoshioka-Maxwell, Mary Hemler, Laura Onasch-Vera, Milind Tambe, and Darlene Woo. Influence maximization in the field: The arduous journey from emerging to deployed application. In *Proc. of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 150–158, 2017. 47

[339] Dan Yamin and Arieh Gavious. Incentives' effect in influenza vaccination policy. *Management Science*, 59(12):2667–2686, 2013. DOI: 10.1287/mnsc.2013.1725. 84

[340] Rong Yang, Christopher Kiekintveld, Fernando Ordonez, Milind Tambe, and Richard John. Improving resource allocation strategy against human adversaries in security games. In *Proc. of the International Joint Conference on Artificial Intelligence*, 22, pages 458–465, 2011. 65

[341] Zhengyu Yin, Albert Xin Jiang, Matthew Paul Johnson, Christopher Kiekintveld, Kevin Leyton-Brown, Tuomas Sandholm, Milind Tambe, and John P. Sullivan. Trusts: Scheduling randomized patrols for fare inspection in transit systems. In *Proc. of the Innovative Applications in Artificial Intelligence Conference IAAI*, 2012. 65

[342] David Ben Yosef, Lihi Naamani-Dery, Svetlana Obraztsova, Zinovi Rabinovich, and Marina Bannikova. Haste makes waste: A case to favour voting bots. In *Proc. of the International Conference on Web Intelligence*, pages 419–425, ACM, 2017. DOI: 10.1145/3106426.3106532. 75

[343] Eliezer Yudkowsky. Cognitive biases potentially affecting judgment of global risks. In Nick Bostrom and Milan M. Cirkovic, Eds., *Global Catastrophic Risks*, chapter 5, pages 91–119, Oxford University Press, Oxford, 2008. 30

[344] Chao Zhang, Shahrzad Gholami, Debarun Kar, Arunesh Sinha, Manish Jain, Ripple Goyal, and Milind Tambe. Keeping pace with criminals: An extended study of designing patrol allocation against adaptive opportunistic criminals. *Games*, 7(3):15, 2016. DOI: 10.3390/g7030015. 65

[345] Yunfeng Zhang, Rachel K. E. Bellamy, and Wendy A. Kellogg. Designing information for remediating cognitive biases in decision-making. In *Proc. of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 2211–2220, 2015. DOI: 10.1145/2702123.2702239. 31

[346] Xingquan Zhu. *Knowledge Discovery and Data Mining: Challenges and Realities: Challenges and Realities*, Igi Global, 2007. DOI: 10.4018/978-1-59904-252-7. 87

[347] Yair Zick, Kobi Gal, Yoram Bachrach, and Moshe Mash. How to form winning coalitions in mixed human-computer settings. In *Proc. of the 26th International Joint Conference on Artificial Intelligence, (IJCAI)*, pages 465–471, 2017. DOI: 10.24963/ijcai.2017/66. 70

[348] Brian D. Ziebart, Andrew L. Maas, J. Andrew Bagnell, and Anind K. Dey. Maximum entropy inverse reinforcement learning. In *AAAI*, 8, pages 1433–1438, 2008. 38

[349] Martin Zinkevich. Rules of machine learning: Best practices for ML engineering. *Technical Report*, 2017. `http://martin.zinkevich.org/rules_of_ml/rules_of_ml.pdf` 36

[350] Martin A. Zinkevich, Michael Bowling, and Michael Wunder. The lemonade stand game competition: Solving unsolvable games. *ACM SIGecom Exchanges*, 10(1):35–38, 2011. DOI: 10.1145/1978721.1978730. 84

[351] James Zou, Reshef Meir, and David Parkes. Strategic voting behavior in doodle polls. In *Proc. of the 18th ACM Conference on Computer Supported Cooperative Work and Social Computing*, pages 464–472, 2015. DOI: 10.1145/2675133.2675273. 75, 76

[352] Inon Zuckerman, Sarit Kraus, and Jeffrey S. Rosenschein. Using focal point learning to improve human—machine tacit coordination. *Autonomous Agents and Multi-agent Systems*, 22(2):289–316, 2011. DOI: 10.1007/s10458-010-9126-5. 56

[353] Inon Zuckerman, Sarit Kraus, Jeffrey S. Rosenschein, and Gal Kaminka. An adversarial environment model for bounded rational agents in zero-sum interactions. In *Proc. of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems*, page 83, 2007. DOI: 10.1145/1329125.1329226. 80

Authors' Biographies

ARIEL ROSENFELD

Ariel Rosenfeld is a Koshland Postdoctoral Fellow at Weizmann Institute of Science, Israel. He obtained a B.Sc. in Computer Science and Economics, graduating magna cum laude from Tel Aviv University, and a Ph.D. in Computer Science from Bar-Ilan University. Rosenfeld's research focus is Human-Agent Interaction and, specifically, the prediction of human decision-making. He has published over 20 papers on related topics at top venues such as AAAI, IJCAI, AAMAS, and ECAI conferences and leading journals. Rosenfeld has a rich lecturing and teaching background spanning over a decade and is currently acting as a lecturer at Bar-Ilan University.

SARIT KRAUS

Sarit Kraus is a Professor of Computer Science at Bar-Ilan University, Israel, and an Adjunct Professor at the University of Maryland. She has focused her research on intelligent agents and multi-agent systems. In particular, she developed Diplomat, the first automated agent that negotiated proficiently with people. Kraus has received the EMET Prize for her expertise and contributions to artificial intelligence, the IJCAI "Computers and Thought Award," the ACM SIGART Agents Research award, and the prestigious Advanced ERC Grant. She also received a special commendation from the city of Los Angeles, together with Professor Tambe, Professor Ordonez, and their students, for the creation of the ARMOR security scheduling system. Kraus has published over 300 papers in leading journals and major conferences.

Index

2k BotPrize, 81

Accuracy paradox, 87
Adversarial agent, 3
Adversary, 63
Agar.io, 94
Alan Turing, 1
AlphaGo, 39
Amos Tversky, 17, 33
Anchoring, 30, 70
Andrew Ng, 51
Anti-Poaching Engine, 66, 90
Argumentation Theory, 55, 72
ARMOR, 5, 22, 35, 64, 89
Artificial intelligence, 1
Asimov's rules, 91, 94
Aspiration Adaptation Theory, 55, 57, 70
Association rules, 38
Automated agent, 2
AutONA, 69
Autonomous decision tool, 93

Bandwagon effect, 30, 75
Battle of the Sexes, 13, 19
Beauty contest game, 27, 57
Behavior cloning, 79
Behavioral sciences, 28
Best response, 13, 14, 62
Binmore's guidelines, 22
Bounded Rationality, 22
Broad model, 51

Cheat game, 82
Checkers, 79
Chess, 3, 25, 79
Choice zones, 44
Clinical decision support system, 49
Clustering, 38
Cognitive bias, 29
Cognitive hierarchy model, 26
Cold start problem, 49
Collaborative filtering, 50
Complete information, 14
Confirmation bias, 55
Continuity, 9
Continuous preferences, 8
Cooperative agent, 3
Cumulative Prospect Theory, 35
Curhan's Subjective Value Inventory, 68
Curse of the optimizer, 62

Dan Ariely, 28
Dan Rice, 30
Daniel Kahneman, 17, 33
DARKO, 50
Data frontier, 46
Data-driven model, 36
Debreu's Representation Theorem, 8
Decision node, 15
Decision trees, 37
Decision-maker, 7
Decision-making under certainty, 8
Decision-making under uncertainty, 9

Decision-support tool, 93
Deductive reasoning, 17
Deep Blue, 84
Deep learning, 38
Deep neural networks, 37
DeepStack, 39
Default bias, 31
Defender, 63
Design of experiments, 45
Diplomat, 68, 89
Domain specific bias, 32
Domain specific bias, 30
Dominant strategy, 12
DRON, 81

Embodied agents, 4
End-to-end testing, 88
Ensemble methods, 55
Ernst Zermelo, 16
Expected-utility maximization, 9, 21
Experimental approach, 42
Experimental design, 45
Expert-driven model, 21
Extensive form game, 14

Fanger's Predicted Mean Vote, 78
Feature construction, 37
Feature engineering, 37
Feature extraction, 37
Focal points, 55
Focal Points Theory, 57
Framing, 32
Fully-personalized model, 49

Game Theory, 10, 16, 22, 64
Game tree, 15
General bias, 30
General game playing, 81, 82
Generalized model, 48
Gerard Debreu, 8
Gery Kasparov, 84
Gibbard-Satterthwaite Theorem, 75
Go, 3, 39, 51, 79
Gravity model, 33

Herbert Simon, 22
Herding, 30
Hierarchical clustering, 37
Holistic model, 51
Human decision-making, 2
Human-Agent Interaction, 3
Humanless circle, 4
Hybrid model, 54

IAGO, 71
Imbalanced dataset, 48
Imitation learning, 79
Independence, 9
Insider threat, 47
Institutional Review Board, 42, 93
Intelligence, 1, 61
INTERCEPT, 65
Interpretability, 88
Inverse reinforcement learning, 38
IRIS, 65, 89
Irrationally Predictable, 28
ISO 7730 standard, 78
Iterations of strategic reasoning, 24

John Forbes Nash, 14
John Maynard Keynes, 27
Joseph Weizenbaum, 90

K-means, 37
KBAgent, 69
Kenneth Arrow, 75
Knowledge acquisition, 35

Labeled examples, 37

Learning from demonstrations, 38
Lemonade Stand Game, 84
Level-k model, 24
Libratus, 39
Linear Programing, 62
Logit quantal response, 23
Loss aversion, 31, 55
Lower contor set, 8

Mario, 81
Mario AI Championship Turing Test, 81
Markov Decision Process, 62
Maurice Allais, 18
Minnesota Multiphasic Personality Inventory, 42
Mixed strategy, 13
Monte Carlo Tree Search, 39, 82
Multiple decision-makers environment, 10
Myopic best-response, 76

Narrow model, 50
Nasa Task Load indeX (TLX), 45
Nash equilibrium, 13
NegoChat, 70
NegoChat-A, 55
No Free Lunch theorem, 87
Nobel Memorial Prize, 14, 17, 22, 33, 54, 55
Non-cooperative game, 11
Non-monotonic logic, 72
Normal form games, 11

Observational approach, 41
Observer agent, 4
Online learning, 66
Opponent modeling, 79
Opportunistic adversaries, 65
Outcome function, 8, 11
Over-sampling, 48
Overfitting, 45, 87
Overspecification, 91

Partially conflicting interests agent, 4
Paul Meehl, 36
Paul Samuelson, 54
PAWS, 65
Payoff matrix, 11
Payoff vector, 15
Perfect information, 14
Personal robotic assistant, 2, 3, 10, 52, 63
Personalization, 48
Personalization-situationalization space, 51
Peter Norvig, 40
PHD flow graph, 88
Player, 11
Poker, 3, 39, 79, 82
post-questionnaire, 45
Predicting Human Decision-making flow graph, 88
Predictive policing, 66
Preference relation, 7
Prisoner's dilemma, 11, 15
Prospect Theory, 31, 33
Pure Nash equilibrium, 13
Pure strategy, 11

Q-learning, 38
QOAgent, 69, 89
Quantal response, 23, 56
Quantal response equilibrium, 24

Racial profiling, 92
Rationality, 7, 12, 16
Rationality parameter, 23, 56
Rationally optimal decision, 8
Receiver operating characteristic (ROC) curve, 48, 87
Regret minimization, 35, 62
Reinforcement learning, 38
Reinhard Selten, 22

Rmax, 38

Security games, 63
Selective labels problem, 44, 57
Semi-personalized model, 48
Single decision-maker environment, 7
Situationalization, 50
SMOTE, 48
Social Choice Theory, 75
Software agents, 4
Solution concepts, 12
Stackelberg Reasoning, 35
Stackelberg Security Game, 64
Statistical testing, 87
Status quo bias, 31
Strategic voting, 75
Strategy elimination, 35
Strategy profile, 11
Subgame, 16
Subgame Perfect Nash Equilibrium, 16
Subjective Expected Utility, 10
Suboptimal decision-making, 23
Super Mario, 4
Supervised learning, 37
Support Vector Machine, 37

Targeted advertising, 50
Targets, 63
Team Reasoning, 35
Terminal node, 15
Think-aloud method, 46
Thomas Schelling, 55
TLX, 45
Traffic Enforcement Allocation Problem, 65
Training examples, 37
Transfer learning, 51, 92
Trembling hand, 22
TRUST, 65, 89
Turing Award, 22
Turing Test, 2, 81

Ultimatum game, 28, 51
Under-sampling, 48
Underspecification, 91
Unsupervised learning, 37
Upper contor set, 8
Utility, 7
Utility function, 8
Utility maximization, 55
Utility theory, 7

Vernon Smith, 17
Von Neumann and Morgenstern, 9
Voting manipulation, 75

Weka, 36

Yisrael Aumann, 17
Yitzhak Rabin, 90

Zermelo's algorithm, 16