1

Introduction

Fei Fang, Milind Tambe, Bistra Dilkina, and Andrew J. Plumptre

Artificial intelligence (AI) has undoubtedly been integrated into our lives. It is hard to enumerate all the aspects of our personal lives that are influenced by AI. Nowadays, the most common way of using the internet, web search, is powered by AI. Almost all the recommender systems are powered by sophisticated AI algorithms that decide what advertisement we see when browsing a website, what items are suggested to us when shopping on Amazon or eBay, what videos, news, and posts are brought to our attention on YouTube, CNN, Facebook, and Twitter. As another example, facial recognition supported by AI techniques is being used in more and more situations to verify the identity of a person, including unlocking of smartphones, security screening at the airport, and verifying customers at the gate during the boarding of a flight. Similarly, speech recognition and synthesis techniques have been embedded into AI-powered home assistants, real-time speech-to-text translation or language translation tools, and automated customer service systems with which we frequently interact. Self-driving cars are no longer a dream.

In short, we are already in an era of AI with all these AI-powered tools and applications. Indeed, with the breakthroughs made in the past few decades, AI has taken over many tasks that have traditionally been completed by humans. This has led to transformation and evolution in many domains and industries, ranging from the IT industry to traditional industries such as finance (e.g., whether to approve a consumer loan, how much one should pay for insurance) and health care (e.g., analyzing medical images automatically).

With all of these rapid advances in the field of AI and the pervasive use of AI, many in both the scientific community and the lay public are increasingly interested in the impact that AI may have on society. Together with these changes brought about by AI, there is a series of concerns. One set of concerns is related to the significant potential for job loss. Another is related to the potential risk of using AI to make critical decisions and, in fact, ensuring the safety of humans from a future powerful AI in charge of such decisions (perhaps similar to what has been depicted in Hollywood movies like *The Terminator*). Yet a third set of concerns focuses on the ethical, moral, and legal challenges that arise from the use of AI in our daily lives – life-and-death questions that come to the forefront when AI must interact with society, as in self-driving cars making choices in life-critical situations.

While these are no doubt legitimate concerns, this book continues with the emphasis in the Artificial Intelligence for Social Good – to push AI research to bring more beneficial results for society. The thrust of the current book series is that AI can be harnessed to address a wide range of social problems and make a significant positive impact. Indeed, the theme of AI for Social Good has been highlighted at different venues where researchers have presented a set of promising emerging applications to promote social good. However, a comprehensive book series that organizes these applications into different topics has so far been missing. Accordingly, each book in our series focuses on how AI could work with a particular field and the social good that could be achieved in that field. The first book in this series, *AI and Social Work*, focused on the fields of AI and social work to address societal problems requiring intervention in social interactions, particularly with low-resource communities (e.g., homeless youth). Topics included HIV prevention, substance abuse prevention, suicide prevention, and others.

This book is the second in this series, taking on a very different domain of impact. In this book we focus on how AI researchers can work with conservation scientists and others in the field of conservation to assist in handling challenges related to the conservation of wildlife, forests, fisheries, rivers, and our environment. The book provides an introduction to this rich and yet new area of research, and outlines key challenges and opportunities. There is tremendous and growing interest in this area, and it is important for us to understand how the significant progress recently achieved in AI could be used to benefit conservation.

The book showcases the interdisciplinary collaboration among many different disciplines in allowing the fields of AI and conservation to partner. We, the editors of this book, are indeed part of this interdisciplinary collaboration, as AI researchers who have spent significant research effort in applying AI to address challenges in conservation (Fang, Tambe, Dilkina) collaborating with a conservation scientist and practitioner who has guided and sought AI efforts for conservation (Plumptre). The book provides an in-depth view of a key thrust of our joint research over the past five years – in this joint work, we have focused on the use of AI in the service of wildlife conservation, specifically how AI can be used to better understand patterns in wildlife poaching and enhance security to combat poaching. The joint effort in this key thrust has led to an important application tool named Protection Assistant for Wildlife Security, or PAWS (Yang et al., 2014; Fang et al., 2016), which has been tested and deployed in countries around the world, including Malaysia, Uganda, and Botswana. The book also features some chapters from key teams of AI researchers partnering with conservation scientists, covering other directions and showcasing various challenges in environmental conservation and how they could be tackled by AI. In these chapters, the book presents work resulting from interdisciplinary collaborations among researchers in computer science, ecology, economics, and psychology.

This book aims to inspire new AI research that can help address conservation challenges. These advances in AI research can bring benefits to the conservation of key natural resources, and help us advance toward the goals listed in the United Nations's 2030 Agenda for Sustainable Development. Simultaneously, the book is intended for conservation scientists who may be interested in using AI approaches and techniques; the variety of examples in the book may yield clues as to how AI could be used, and the strengths and limitations of current AI tools. Accordingly, the intended audience for the book is: (1) researchers in computer science who are interested in AI and social good and in particular applications of AI in conservation; (2) researchers and practitioners in conservation and related scientific fields who are interested in how AI techniques can be used for conservation; and (3) other researchers interested in both computer science and conservation who may be interested in starting research in this area.

Since this book aims to speak to multiple different audiences, including conservation and AI researchers, a brief overview of these two disciplines is warranted. Unfortunately, a few paragraphs on these topics can hardly hope to do each area much justice. Moreover, such a short treatment will inevitably be biased – reflecting our own personal biases within our respective disciplines. That said, we hope these next few paragraphs help you, our readers, to be grounded to some extent in the work of these two fields.

What Is Conservation Science?

Conservation Biology

Conservation biology is about the management of nature, the maintenance, loss, and restoration of biodiversity on the earth with the aim of protecting species, their habitats, and ecosystems from excessive rates of extinction and the erosion of biotic interactions (Soule, 1980, 1985). The extinction of species, the loss of habitat, and the decline of biodiversity and biological systems around the world raises significant concerns about the sustainable well-being of human society; together they may contribute to poverty, starvation, and will even reset the course of evolution on this planet. Conservation biologists study these causes, trends, and their impacts, as well as action plans to combat undesired change (Groom et al., 2006).

Conservation biology started as a scientific discipline in the mid-1970s during the First International Conference on Research in Conservation Biology, held at the University of California, San Diego Douglas, 1978. The field of conservation biology and the concept of biodiversity emerged together, helping to crystallize the modern era of conservation science and policy, and bridging the gap between theory in ecology and evolutionary genetics and conservation policy and practice. It has grown rapidly to include many subdisciplines, such as conservation genetics, conservation planning, conservation social science, conservation physiology, ex-situ and in-situ conservation, restoration science, and many others. It is an interdisciplinary subject drawing on natural and social sciences and the practice of natural resource management. Conservation biology is closely connected to ecology, biology, and economy, among other related disciplines (Van Dyke, 2008).

Conservation of Natural Resources

The conservation of natural resources is a topic for both research and practice. The practice of natural resource conservation exists in history in different cultures (Kala, 2005; Murphree, 2009). Within conservation biology, the research on conservation of natural resources can be further classified based on the type of natural resources being considered, such as water conservation, energy conservation, marine conservation, wildlife conservation, habitat conservation, soil conservation, forest conservation, wetland conservation, etc. In this book, the work we discuss will cover several of these topics – including wildlife conservation, forest conservation, water conservation, and marine conservation – but our emphasis will be on wildlife conservation.

Managing Conservation Areas

Since the establishment of the first national park, Yellowstone National Park in the USA, in 1872 there has been a need to manage protected areas both for their wildlife and also for the people living around them. Initially, these areas were established to conserve wilderness and areas for recreation, but over time there has been an increasing focus on the conservation of biodiversity and the richness of life on earth. Over the years, the percentage of the planet covered by protected areas has increased to about 15 percent on land, and about 10 percent of coastal areas together with 4 percent of global ocean areas. The Aichi targets of the Convention on Biological Diversity (CBD) set a goal of 17 percent of terrestrial and 10 percent of marine areas for conservation by 2020, but these targets are already being reassessed as the area required to conserve all species on the planet will need to be larger than this.

With the growth of conservation areas, there has also been incredible growth in the human population, from about 1.2 billion in 1872 to 7.4 billion today. This has changed the nature of conservation, as the demand for land has greatly increased, causing fragmentation and loss of habitat, together with the increasingly unsustainable use of natural resources. The challenges a protected area manager has to tackle today are very different to those of 145 years ago, and are increasingly complex and often involve tradeoffs between various management options. Many of the management practices in conservation improve with the experience of the manager as they learn how to negotiate the issues they face. There is a need to improve the ability of managers to make wise decisions, and conservation science aims to provide answers to improve management and conservation. Conservation science aims to help identify questions, such as where to conserve, when to conserve, and how to conserve.

What Is AI?

From the Turing Test to National Strategies

Many people have heard about the Turing Test, proposed by Alan Turing in 1950 (Turing, 1950). The Turing Test is designed to answer the question “Can the machine think?” Briefly, the test consists of a human judge conversing with a human and a machine via text. The judge needs to identify the machine reliably, while the machine’s objective is to mimic a human as closely as possible in an attempt to fool the judge. This test is highly influential yet widely criticized. Even if a machine can do well in this game, can we claim the machine has intelligence? This relates to the question of how AI should be defined. Indeed, there is no accepted definition of AI, but generally speaking we will consider AI to include computer systems (including robots) that perform tasks that require intelligence.

Since the creation of the Turing Test, research in AI has continued to advance with the design of new algorithms and tools. AI progressed from being an academic discipline that few had heard about, to a discipline that is often in the news, to a discipline that is now mentioned in national research strategies in multiple countries, including the USA and UK.

History of AI

The discipline of AI was founded in the 1950s. Four scientists are generally considered to be founding fathers: Allen Newell and Herbert Simon of Carnegie Mellon University, John McCarthy of Stanford University, and Marvin Minsky of Massachusetts Institute of Technology. They had the optimistic vision that "aspects of intelligence can be so precisely described that a machine can be made to simulate it" (McCarthy et al., 1955). Since then, AI has experienced optimism and disappointment. The early years of optimism in AI were based on computer programs performing tasks that could be considered difficult for humans, such as playing chess, proving theorems, and solving puzzles. In these tasks, interacting with the outside world was not a priority and uncertainty was not a major concern. For example, statements are either true or false; chess moves are deterministic. Furthermore, without the need to rapidly interact with the outside world, a program could reason about a situation in depth and then react. Indeed, early progress in AI was dominated by the use of logic and symbolic systems, heuristic reasoning, and search algorithms, which did not involve probabilistic reasoning or uncertainty.

In the 1980s, this optimism about AI research led to integrated AI systems that could combine machine learning to autonomously learn from interactions of the AI system with the world, automated reasoning and planning to act in the world, possibly speech and dialogue in natural language to interact with the world, social interactions with humans and other entities, and so on. Such integrated intelligent agents included robots, but also avatars in virtual worlds. This vision of AI involved gathering data from the world and using machine learning to build a predictive model of how the world would evolve. However, it did not stop only at making predictions but also included planning and reasoning from first principles, that is, providing prescriptions and taking actions in the world.

This early optimism was followed by the AI winter around 1990, when outsiders criticized AI research for not living up to its promises, and this led to a general reduction in funding and interest in AI. Researchers realized that what was easy for humans (e.g., recognizing objects in images, understanding human speech as English words) turned out to be difficult for AI systems.

As AI systems started interacting with the real world, an entirely new set of tools was needed to handle reasoning with uncertainty and for quick reactions to unanticipated events. These difficulties led to new quantitative tools from decision theory, optimization, and game theory being brought into AI. This period also coincided with an emphasis on fast, reactive reasoning (rather than purely in-depth reasoning) in AI, particularly for robots interacting with the world. Despite the criticism and difficulties, AI researchers continued working hard and making progress in various aspects of AI. Encouraging news started flowing again to slowly lift research out of the AI winter, including the famous victory of IBM supercomputer Deep Blue over the world chess champion, Garry Kasparov.

New Era of AI

We are now in a new era of AI. The increasing computing power and the ability to gather and store large quantities of data has led to promising new tools and approaches. Machine learning using deep neural networks has led to significant new successes in tasks such as recognizing objects in images and natural language understanding. The computer program AlphaGo defeating world champions at the game Go was viewed as another milestone of AI – the game of Go, a classic strategy board game for two players, was considered to be a very complex game that required a high level of intelligence to play.

This new era has brought significant new interest and investment from industry due to the use of AI in commercial applications, from ad auctions to the search, recommendation, verification, and translation tools that are used by millions of people. It has also led to some misunderstanding in the general public that AI is only about machine learning and prediction, and that unless there are vast quantities of data, AI cannot function. Often not well understood is AI's ability to provide prescriptions or to plan intervention in the world by reasoning through millions of possibilities – for example, using millions of simulations from first principles of how entities may interact in the world, and, more importantly, how such interventions can be integrated with machine learning to accomplish more complex tasks. In our work described in Part I of this book, and the other chapters outlined in Part II, we will often focus on this ability to plan interventions and their integration.

Subfields of AI

As can be seen from the development of AI, AI research has been divided into subfields, including machine learning, multi-agent systems, natural language processing, computer vision, planning and scheduling, reasoning under uncertainty, search and constraint satisfaction, heuristic search and optimization, knowledge representation and reasoning, and others. Nonetheless, AI systems in the real world may require integration of research from several subfields.

Partnership of AI and Conservation Science

Despite the obvious differences in AI and conservation science, they also complement each other in important ways, which makes the partnership of AI and conservation science a productive one. The shared interest of AI and conservation science are centered around a few major topics. First, measuring and monitoring the ecosystem often requires statistical inference, for which machine learning models and algorithms can be developed and applied. For example, the problem of estimating the distribution of wildlife species or poaching activities based on occurrence data can be viewed as predicting the probability distribution given a set of training data, which is naturally addressed in the AI subfield of machine learning.

Second, the conservation and management of natural resources often involves planning, scheduling, and optimization, which are the focus of study in several subfields of AI. For example, to manage invasive species, a land manager must decide how and where to fight the invasion – for example, she may choose a location to eradicate invasive plants or plant more native plants. Finding the optimal location and treatment is an optimization problem that may exceed the capability of a human manager, but can be addressed through developing Markov decision process-based models and algorithms which are heavily studied in AI. In fact, the partnership of AI and conservation science dates back to 1989, when a conservation planning problem in South Australia was modeled as an optimization problem, and mathematical programming was first used to solve the problem (Cocks and Baird, 1989). This solution approach was followed by Marxan, a decision support software for conservation planning developed in 2005 that is widely used across the globe.

Third, conservation of natural resources often involves understanding how multiple agents interact. For example, conservation agencies send rangers to protected areas to detect and deter poaching activities conducted by poachers, and both parties behave in ways that are best for themselves, leading to strategic interplay. Similarly, conservation may also require understanding the conflict between humans and wildlife. Such interactions are modeled via a multi-agent system, the focus of a key subfield in AI. Therefore, frameworks and approaches in multi-agent systems become a natural fit to understand and analyze the interactions as well as to optimize the actions of some agents (e.g., planning the patrol routes for the rangers) to achieve the conservation goal.

This productive collaboration between AI and conservation is a new intellectual space that empowers new research for both fields. It goes beyond just taking a problem in conservation biology and applying existing AI techniques to solve it. Instead, it creates new challenges for both sides and fosters innovative science in AI and conservation biology, with general applications. Take our work on predicting poaching threat, described later in this book, as an illustrative example. The data collected from previous patrols provide information about where snares (a common poaching tool) were found and how often different regions in the protected area have been patrolled; the question to be answered is where the poachers will place snares in the future. From the perspective of AI, predicting the distribution of poaching activity introduced several new challenges for which existing algorithms could not be directly applied. First, the level of poaching threat varies spatially and temporally, and poachers may adapt to how the rangers patrol by circumventing areas that have been heavily patrolled in the past. Second, unlike many other domains that have a tremendous amount of data, the anti-poaching domain often has a training dataset that is highly imbalanced, biased, and noisy. The imbalances come from the fact that patrollers patrol vast areas, but they often find actual poaching signs or snares in small parts of the areas they patrol. Furthermore, since the patrollers cannot cover all the regions in a protected area and the selection of patrol routes is decided by the site manager, the data collected suffer significantly from sampling bias. Finally, the noisiness comes from imperfect detection. When a patroller does not find any poaching activity in a region, it is possible that poaching activity actually occurred but the patroller failed to detect it because the poaching signs remain well hidden in the forest. Thus, we had to propose new algorithms to address these research challenges, advancing the techniques in machine learning from limited labeled data. From the perspective of conservation, while conservation tools such as SMART (Spatial Monitoring And Reporting Tool) have been used to record and analyze the coverage of patrols and distribution of poaching activities found, there is not much work on predicting future poaching threats using machine learning models. Such predictive models radically shift the view of conservation site management from summarizing historical poaching activities and patrols to predicting the future and proactively planning patrols.

The interdisciplinary space also leads to challenges that have been underexplored in these two disciplines. One set of challenges is how to integrate AI with human insight and knowledge. For example, how to elicit valuable prior knowledge from conservation researchers and practitioners when building an AI tool. Practitioners working on conservation often have a much better understanding of the domain than what is recorded in a dataset. In predicting poaching threat, the dataset only contains where the rangers went and what they found, but the rangers know more than that. They may know which features make a region attractive to poachers, which villages the poachers come from, what external factors make the overall level of poaching threat increase or decrease, and so on. Such prior knowledge, if used properly, may significantly boost the performance of the machine learning algorithm, but there is a lack of research on how to make proper use of it. Also, it is important to build tools so that human experts can provide iterative inputs to improve the performance of an algorithm. Another example is how to interpret solutions to conservation challenges provided by AI tools. For AI researchers, the research questions include how to provide explainable solutions and effectively convey the solutions to conservation researchers and practitioners. For conservation science researchers, the research question is how to interpret the solution to provide conservation insights.

In addition to the challenges in knowledge elicitation and interpretation, another set of challenges center on the validation and evaluation of prediction or planning tools. While evaluating predictions on a readily available test set (dataset in the lab) is a more traditional approach to evaluating predictive models, an important method of evaluating a predictive model in this interdisciplinary space is to test it in the real world. However, it is difficult to conduct tests under perfectly controlled conditions in the real world. These difficulties also extend to the evaluation of planning tools. For example, to evaluate a patrol planning tool, ideally one would have two identical portions of a protected area, with one running the planning tool and one not; but finding such identical areas is difficult. Furthermore, improved patrolling could lead to an increase in the numbers of captures of poaching tools such as snares, but it may also cause poachers to reduce poaching and thus a reduction of the numbers of snares found. Finally, there may also be contamination in the experiment because poachers may shift poaching activities from one area to another precisely because of improved patrolling in a neighboring area. Many of these challenges are still open questions, forming important directions for future research in how to evaluate AI algorithms in the field.

In fact, AI and conservation is one of the key research thrusts in the broader research field of computational sustainability (Gomes, 2009), which focuses on ways that computational models, methods, and tools can help solve some of the key sustainability challenges that our planet and society face. This growing interdisciplinary research area provides a two-way street of benefits, by harnessing the power of computing and specifically AI for real-world