



American Planning Association
Planning Advisory Service
Creating Great Communities for All



person 100%

person 100%

person 100%

person 100%

bicycle 99%

bicycle 100%

car 91%

car 99%

PAS REPORT 604

PLANNING WITH ARTIFICIAL INTELLIGENCE

Thomas W. Sanchez, PhD

The American Planning Association will lead the way to equitable, thriving communities by creating unique insights, as well as innovative and practical approaches that enable the planning community to anticipate and successfully adapt to the needs of a rapidly changing world.

Since 1949, the American Planning Association's Planning Advisory Service has provided planners with expert research, authoritative information on best practices and innovative solutions, and practical tools to help them manage on-the-job challenges. PAS Reports are produced in the Research Department of APA. Joel Albizo, FASAE, CAE, Chief Executive Officer; Petra Hurtado, PHD, Director of Research and Foresight; Ann F. Dilemuth, AICP, PAS Editor. APA membership includes access to all PAS publications, including PAS Reports, *PAS Memo*, and *PAS QuickNotes*. Learn more at planning.org/pas. Email: pasreports@planning.org

©June 2023 American Planning Association, 205 N. Michigan Ave., Suite 1200, Chicago, IL 60601-5927; planning.org. All rights reserved. No part of this publication may be reproduced or utilized in any form or by any means without permission in writing from APA.

ISBN: 978-1-61190-212-9

ABOUT THE AUTHOR

Thomas W. Sanchez, PhD, is currently a Professor of Urban Affairs and Planning at Virginia Tech. In 2024, he will join the faculty of the Department of Landscape Architecture and Urban Planning at Texas A&M University. He earned his PhD in city planning from Georgia Tech and a masters degree in city and regional planning from Cal Poly, San Luis Obispo. His research and teaching focuses on urban planning, planning methods, technology, and scholarly impact. His books include *Networks in the Knowledge Economy* (with Denise Bedford), published in 2021, and *Planning Knowledge and Research*, published by Routledge in 2018. Sanchez is an active member of the American Planning Association, where he has served as the chair of the Education Committee and as a member of APA's Artificial Intelligence (AI) Foresight Community.

ON THE COVER

Image by recep-bg/Getty Images; bounding boxes generated by YOLO v3 Object Detector (<https://iashin.ai/detector>)



TABLE OF CONTENTS

EXECUTIVE SUMMARY 4

CHAPTER 1 INTRODUCTION 7

- The Future Is Digital 8
- Why AI, and Why Now 9
- About This Report 11

CHAPTER 2 THE BASICS OF AI 12

- A Brief History of AI 13
- Machine Learning 15
- Neural Networks 20
- Natural Language Processing 22
- Computer Vision 23
- Conclusion 24

CHAPTER 3 AI IN CURRENT PRACTICE 26

- Computer Vision for Neighborhood Conditions Assessment 27
- Machine Learning for Boundary Analysis 31
- Natural Language Processing for Public Involvement 35
- Neural Networks 38
- Conclusion 40

CHAPTER 4 CHALLENGES TO ADOPTING AI IN PLANNING 42

- Fear and Uncertainty 43
- Lack of Skills 44
- Unclear Goals 45
- Data Needs 46
- Transparency and Explainability 46
- Bias 47
- Other Ethical Issues 48
- Conclusion 50

CHAPTER 5 GETTING STARTED 51

- Connecting AI With Planning Practice 53
- Strategic Planning and Visioning For AI 55
- Now Is the Time 56

APPENDIX: GLOSSARY 57

REFERENCES 62

ACKNOWLEDGMENTS 66

EXECUTIVE SUMMARY

Today, artificial intelligence (AI) is altering the human experience, causing profound social, economic, and technical changes that touch all of us. The convergence of very large datasets and the capabilities to analyze them represent a potentially transformative period for the planning profession.

Now is the time for planners to examine the responsible deployment of AI to maximize its benefits and reduce its drawbacks, particularly in the face of its existing uncertainties. Having general knowledge of what AI is and how it operates, as well as understanding the most important practical and ethical aspects surrounding its use, is essential to implementing AI in ways that best serve planning practice.

AI holds the potential to automate repetitive, labor-intensive tasks in urban planning, allowing for more efficient decision-making in the planning process and freeing up planners' time to focus on higher-value work. However, not all facets of planning are suitable for AI applications. To identify opportunities for improvement, planners will need to assess their procedures against AI capabilities and determine the necessary steps to implement changes. As AI continues to advance, its use in urban planning is expected to expand, creating a need for awareness and knowledge among planning practitioners, researchers, and educators.

This PAS Report introduces planners to basic AI methods and concepts, and it provides current examples of how AI is already being used to augment planning activities. It seeks to provide the information planners need to begin recognizing how AI methods might fit planning tasks that they are working on. Planners may not be responsible for building AI solutions, but if they understand the general framework and the language of AI, they can better communicate with application developers and data scientists.

AI TOOLS IN PRACTICE

This PAS Report focuses on four primary application areas of AI. In **machine learning** (ML), machines (i.e., computers)

detect and “learn” patterns, and then use this information as part of decision-making processes. ML is used to enable functions such as *sentiment analysis*, which examines text to determine the attitude that it conveys (e.g., positive, negative, or neutral); *spam detection*, in which algorithms learn to categorize email messages as valid or spam and filter out the latter; *recommendation systems*, which use a consumer's prior purchases, actions, or characteristics to anticipate additional products the consumer “may also like”; and *predictive maintenance* and *anomaly detection*, which predicts the likelihood that machinery will break down, determines maintenance scheduling, and monitors and identifies equipment problems in real time.

Artificial neural networks (NNs) are computerized systems based on how the brain processes information that can form connections between two items and learn to correlate them with one another to make predictions. NNs are used in *speech recognition*, which translates spoken words into text for virtual assistants and live captioning, and general *pattern recognition*, used extensively in fields such as biology, psychology, health, and marketing.

Natural language processing (NLP) comprises computer methods that analyze human language, text, or verbal communication to derive meaning. NLP systems analyze the ideas conveyed by words, and they can also correct grammar, convert recorded speech to text, and translate between languages. NLP combined with ML is used to operate AI *virtual assistants*, such as Apple's Siri and Amazon's Alexa. NLP also drives *text analysis* and *content creation* by large language models such as OpenAI's ChatGPT.

Computer vision (CV) focuses on enabling computers to interpret and understand visual information, such as images or video. CV applications include *raw image data pro-*

cessing, such as the conversion of satellite imagery to output maps, and *detailed image analysis*, used in facial recognition systems and other applications. CV is increasingly being used in conjunction with NNs in the transportation sector for systems that can classify and track vehicle types and roadway usage, monitor pavement and infrastructure conditions, and allow self-driving vehicles to “see” and respond to their surroundings.

ML, NNs, NLP, and CV can be combined to create AI-powered analysis tools that can help planners process large amounts of data, identify patterns and trends, and make better-informed decisions. These application areas are also being used to develop smart city technologies, which can automate and improve city services such as transportation, energy and water management, public health, safety, and security.

This PAS Report offers specific examples of how these AI applications are currently being used in the field, including the following:

- In Tuscaloosa, Alabama, CV is being used to assess neighborhood conditions throughout the city. Smart cameras mounted on municipal garbage trucks collect GPS-located image data from neighborhoods, which is processed using an ML model to generate blight scores. This has allowed for efficient identification of blighted areas and targeting of activities and resources for blight intervention and remediation planning.
- In Ector County, Texas, a school district used AI methods and spatial analysis to improve social equity and incorporate school districts as significant partners in community planning. By analyzing the socioeconomic characteristics of the district’s physical environment, planners were able to visualize the district’s standing and identify areas in need of improvement.
- In Long Beach, California, NLP is being used to capture and evaluate public input. More than 14 departments use the city’s sentiment analysis platform, which collects data from social media, news sites, and city hotlines to analyze residents’ conversations and categorize comments by topic and sentiment (positive, negative, or neutral). This approach helps quantify attention to specific issues and evaluate community perceptions. NLP is also being used in New Zealand to engage and communicate with Indigenous people in university and municipal settings.

The report also explores how researchers are using NNs to analyze transportation data to detect travel impacts and study street network design, disaster and evacuation route

planning, health and safety, and social equity; analyze and model land use and land-use change; and inform urban design and development.

CHALLENGES TO AI ADOPTION

The adoption of AI in urban planning presents both challenges and opportunities. This PAS Report examines seven key challenges to AI adoption within the planning profession: fear and uncertainty, the need for new skills, changing data needs, unclear goals, transparency and explainability, bias, and ethical considerations. By proactively addressing these challenges, planners can harness the power of AI to improve planning processes, drive innovation, and ultimately create more sustainable and equitable urban environments.

The report highlights setting a clear vision, organizational goals, and performance measures as key components for effective AI implementation in urban planning, as well as the critical role of high-quality data and data management skills. It also emphasizes the importance of transparent and explainable AI systems for public understanding and acceptance, and it reiterates the vital need to address bias in AI systems and develop ethical standards for responsible AI use that center fairness, transparency, privacy, and accountability.

LOOKING AHEAD TO AN AI FUTURE

The use of AI is gaining momentum and offers a wide range of benefits for creating vibrant and sustainable communities. AI will not replace planners; instead, it will provide a useful set of tools for a subset of planning tasks and enhance planning practice by allowing planners to focus on the human components of planning.

How can planners begin to integrate AI into their planning work? This report suggests that planners consider five areas of strategic intervention in planning—community visioning; plan making; standards, policies, and incentives; development work; and public investments—as a starting framework to identify areas where AI can be applied.

The report also recommends developing a strategic organizational plan for the adoption of AI, which should include a thorough analysis of the current state of the organization, identification of potential areas for AI adoption, and a clear vision for the implementation of AI, including addressing potential challenges and ethical considerations.

In late 2022, conversational chatbots with advanced language features gained significant attention. These chatbots, such as OpenAI's ChatGPT, use natural language input and can produce various output styles such as blog posts, essays, and recipes. As AI algorithms and ML techniques continue to advance, these may be some of the first AI tools widely used to support data-driven and efficient planning. The adoption of AI tools in urban planning, however, will be largely incremental as skill development occurs through professional training and university programs.

AI is an emerging technology that is here to stay, and it will play a significant role in the future operation and planning of cities. Despite the challenges, AI applications have the potential to improve the understanding and development of more habitable and equitable urban spaces. As more and more data is being generated all around us, urban planners will increasingly use AI to analyze and process these data to detect patterns, make predictions, and have a better understanding of urban dynamics.

This PAS Report explores the substantial opportunities and challenges presented by AI technologies, with more being revealed each day. Responsible and ethical AI use can help planners in their work, enhance current planning procedures, increase efficiency, and allow planners to refocus their work on the human components of planning. In reading this report, planners will begin the process of increasing their knowledge and skills with appropriate AI methods to enhance planning practice in the service of creating more livable, resilient, and sustainable communities.

CHAPTER 1

INTRODUCTION

Today, artificial intelligence (AI) is altering the human experience, causing profound social, economic, and technical changes that touch all of us. The convergence of very large datasets and the capabilities to analyze them represent a potentially transformative period for the planning profession.

At the cusp of significant innovation in analytical processes, it is prudent for planners to examine the responsible deployment of AI to maximize its benefits and reduce its drawbacks, particularly in the face of its existing uncertainties. In addition, having general knowledge of what AI is and how it operates, as well as understanding the most important practical and ethical aspects surrounding its use, is essential to implementing AI in ways that best serve planning practice.

Claims—both positive and negative—are frequently exaggerated to the point of being deceptive, and the term “AI” is frequently used as a marketing phrase rather than a detailed description of the employed methodologies. Regarding scope, definitions, and priorities, there might be conflicting views or disagreement even among those working in the field. To facilitate better policy, identify both opportunities and risks, and evaluate claims made by others, planners should seek clarification on what AI means; what components can make up a system; the wide variety of ways performance and accuracy, fairness, accountability, privacy, and security can come into play; and the complexity of weighing these factors against each other while planning for high-quality places to live.

This PAS Report introduces planners to basic AI methods and concepts, and it provides current examples of how AI is already being used to augment planning activities. While there are potential applications to improve office efficiencies (see Daniel and Petit 2021), the focus of this report is on planning tools and activities. The intent is to provide enough background so that planners can start to recognize how AI methods might fit planning tasks that they are working on. Planners may not be responsible for building AI solutions, but they should understand the general frame-

work and the language of AI so they can communicate with application developers and data scientists.

THE FUTURE IS DIGITAL

In 2022, the American Planning Association (APA) and the Lincoln Institute of Land Policy collaborated to identify approximately 100 trends pertinent to urban planning (Hurtado et al. 2022). The resulting [*2022 Trend Report for Planners*](#), the first of its kind, identified several trends related to AI and urban planning:

- **Digital transformation of cities.** Today, digital technologies touch nearly every aspect of life, including how individuals live, work, learn, shop, and move around the city, as well as how businesses engage with their customers and how we communicate with one another.
- **AI in everyday life.** Because AI is already transforming the local landscape, it is crucial for planners to comprehend how to apply AI fairly and effectively.
- **Urban infrastructure and AI.** AI systems are increasingly being used for urban technology and infrastructure, both in the public and private sectors. This includes automating systems and their maintenance, along with optimizing these processes.
- **AI-based planning tools and upskilling needs.** Aside from the expected effects on cities and communities and the planners who work in them, AI is also likely to change the planning profession itself.
- **AI ethics.** Concerns about human rights, civil liberties, privacy, and social equity are important when AI systems are used in communities or AI is used as a planning tool.

The [2023 Trend Report for Planners](#) (Hurtado et al. 2023) updates and expands on the potential role of AI in urban planning. One factor with significant implications is the accelerating rate of development and availability of open-source, consumer-grade AI. This means that these tools are no longer limited to researchers but can be applied by anyone without specialized training. And as planners begin to explore these tools, city managers and other public administrators will also need to increase their familiarity with AI concepts to manage and support their staff in handling AI-related tasks (Yigitcanlar, Agdas, and Degirmenci 2022).

Digitalization is a factor contributing to the rise of AI in the public realm because it enables the large-scale data capture that is needed to support AI (Andrews et al. 2022; Gomez and DeAngelis 2022). The adoption of new AI-based planning tools will require planners to learn new data management skills, acquire an understanding of infrastructure innovations using AI (Hurtado, Hitchings, and Rouse 2021), and understand ethical concerns about AI (Andrews et al. 2022). This will require upskilling, or learning new skills, through professional development and planning education (Hurtado et al. 2022).

While a large proportion of the American public has concerns about AI and computer applications to enhance human intelligence, they also express a degree of optimism about its potential benefits. A [2021 survey by the Pew Research Center](#) suggests that these perceptions are frequently based on worries about autonomy, unintended consequences, and the degree of change that these developments might entail for people and society (Rainie et al. 2022). Respondents expressed the greatest concern for “loss of humans’ jobs,” “surveillance, hacking, and digital privacy,” and “lack of human connection, qualities,” while at the same time identifying “makes life, society better,” “saves time, more efficient,” and “inevitable progress, is the future” as potential positive outcomes (Ranie et al. 2022, 21). The survey also shows that most people associate AI with video surveillance and self-driving cars and less so with its many other capabilities, such as data analysis and decision support.

Also in 2021, a [survey of APA members](#) asked about planners’ knowledge of AI, perceived levels of appropriateness for planning, and the likelihood of adopting AI tools in the near future. The survey results suggest that planners currently have low levels of familiarity and experience with AI applications, while having a generally favorable outlook on AI adoption for the planning profession (Sanchez et al. 2022). Respondents indicated that they expected AI to play an important role in future planning practice, though only

Plan to Adopt AI




Figure 1.1. Survey respondents reporting that they or their organizations plan to adopt AI (Thomas Sanchez)


seven percent of respondents planned to start using AI tools in the near future (Figure 1.1). But if general trends are any indication, opportunities for AI use in urban planning will expand as the technology continues to advance.

WHY AI, AND WHY NOW

This PAS Report defines artificial intelligence (AI) as the use of computers to mimic or improve upon human intelligence processes, such as reasoning and experience-based learning. For planners, AI represents a valuable tool that can remove the repetitive and labor-intensive aspects of analyses performed to support decision-making. It also represents new opportunities for knowledge access and management. With AI, planners can process large amounts of data and identify patterns that may not be immediately apparent, which can inform more effective decision-making.

While AI refers to a very broad range of analytical methods, the primary elements involve a sequence of inputs, analysis, and outputs (Figure 1.2). An AI program takes inputs of a large quantity of data, analyzes it according to a

Figure 1.2. The basic analytical process for AI
(Thomas Sanchez)



set of rules, and produces outputs, which can take a variety of forms depending on the types of inputs and types of questions being addressed. What makes AI “intelligent” is its ability to use results from the analysis and outputs to update or improve the overall process: to learn. Planners might note that this logic is not unlike the basic process of planning practice—gathering information about an issue or a challenge, analyzing it to answer important questions, and producing plans or policies that improve the situation or address that challenge.

Decades of AI research by computer scientists have focused on understanding the complexities of human thought processes and logic. The objective has been not only to understand how humans think, but also to develop the means to improve upon it in a way that could be applied to nearly anything that involves intelligence and decision-making (Batty 2022). AI is now associated with a wide range of consumer goods and services that we use every day, from product design to marketing to customer feedback.

AI is also starting to have impacts on government operations and policy making. Governments all over the world are using AI technology to influence their responses to some of the most pressing problems of our time, such as the COVID-19 pandemic (OECD 2020), climate change (Dimock 2022), and the implementation of new data laws and governance structures (Rahnama and Pentland 2022). In the United States, agencies at different levels of government are now using AI-related tools to address the many challenges they face. One such example is the growing use of chatbots by state governments to respond to information requests about unemployment cases and COVID-19 (NASCIO 2020).

Research into AI for urban planning began in the 1960s, influenced by the development of computer technolo-

gies with a broad range of scientific and industrial applications (Brail 1987). The development and implementation of planning applications such as transportation and land-use forecasting systems, however, were significantly limited by the lack of large-scale datasets and computing capabilities. Over time, academic and industrial research using advanced quantitative and spatial analysis commonly associated with AI has increased steadily; however, few of these approaches have been adopted by the planning profession. Much of the current planning-related discussion about AI is related to smart city technologies used to capture and analyze data for optimization processes, such as traffic management and energy production. Less attention is being paid to the use of AI in urban planning and associated decision-making activities, including scenario planning and generative design.

Planning the cities of the future will require appropriate expertise and proper planning methods. There is general agreement that urbanization will continue to increase as a function of economic opportunity, while at the same time equity, environmental pressures, and infrastructure needs will continue to be major challenges. In addition, technology, connectedness, and digitalization are poised to accelerate at a rapid pace (Frem, Rajadhyaksha, and Woetzel 2018; National League of Cities 2022; Antunes, Barroca, and de Oliveira 2021; European Commission 2019; UN Habitat 2022; World Economic Forum 2022). AI is expected to play an important role in future development activities as well as in the planning and management of cities for sustainability, resilience, and equity in both the short and long terms. Planners must be prepared for the changes this means in how the cities of the future are planned, designed, and managed.

Like most other fields, the planning profession has greatly benefited from adopting new technologies to aug-

ment planning practices. Just as urban planning gradually adopted computer technologies during the 1980s and GIS during the 1990s, the profession is now on an inevitable progression toward adopting AI methods. Unlike in earlier decades, today data is widely available, computer technology is pervasive, and AI is being used all around us.

But not all aspects of planning are suitable or appropriate for advanced analytical methods and other capabilities of AI. AI will not take a planner's job anytime soon (if ever) because of the complexity of planning and the processes involved with plan making, but it can be a valuable tool, especially for tasks that are repetitive and time consuming. As with other innovative tools and technologies, planners will need to assess the procedures they use against the capabilities of AI to identify opportunities for improvement and determine the steps that will be required to make those changes.

ABOUT THIS REPORT

This PAS Report provides an introduction to AI for urban planners. It highlights the basics of AI that can replace or help rethink planning analyses and processes. At this time, little AI is being used in everyday planning practice, despite a substantial and growing body of scholarship that applies AI methods to planning questions. Much of this research performs specific types of planning analyses, but these methods are yet to be used in the field for preparing or evaluating plans or policies. Closing this gap has its challenges, and will start with planning practitioners, researchers, and educators becoming aware of the opportunities and potential benefits of adopting new approaches. This report seeks to make those connections.

This chapter has introduced the concept of AI and explained why it is a technology whose time is coming. Chapter 2, **The Basics of AI**, offers a primer on AI methods through a general discussion of four main AI application areas that have potential planning relevance: machine learning, neural networks, natural language processing, and computer vision. The main objectives of the chapter are to introduce basic AI terminology and concepts and encourage readers to think about how their planning activities share similarities with the methods described.

Chapter 3, **AI in Current Practice**, provides examples of AI applications that are already being used for local government and planning purposes. It highlights the current usage of AI by planners and includes additional discussion to broaden understanding of how that use may extend to

other planning activities. While AI includes a broad range of applications, the focus here is on tools with the potential to augment planning analyses.

Chapter 4, **Challenges to Adopting AI in Planning**, discusses some of the concerns that have been expressed about the use of AI. Some of these concerns are based on known or suspected problems such as bias, error, and misapplication leading to poor outcomes. Stakeholder concerns include those from the public, government officials, and planners themselves. Much of this concern relates to privacy and the confidentiality of data collection. Another concern is that AI methods may perpetuate discriminatory historical analog methods that have had adverse impacts on neighborhoods and communities. Careful consideration of AI use and examination of its logic will be vital in avoiding those pitfalls.

The report concludes with Chapter 5, **Getting Started**. It offers guidance and a strategic framework for how planners should consider future AI methods today and over the long term and forecasts an ever-expanding scale of AI applications across planning and related fields.

A **Glossary**, provided as an appendix, provides definitions of AI-related terminology used in the report. Where terms are highlighted in the text, readers can click through to find the definitions.

CHAPTER 2

THE BASICS OF AI

The term artificial intelligence (AI) applies to many different technologies and methods. In the 1950s, AI pioneer [John McCarthy](#) came up with the term “artificial intelligence,” which he and his colleagues defined as “making a machine behave in ways that would be called intelligent if a human were so behaving” (McCarthy et al. 1955, 11). More than 60 years later, AI researchers Stuart Russell and Peter Norvig defined AI as “as the study of agents that receive percepts from the environment and perform actions” (Russell and Norvig 2016, viii). Though AI may be defined in different ways based on uses and objectives, its key elements—as in both of these definitions—include *intelligence* and *behavior*, or action based on that intelligence.

This chapter provides a general background on AI, a rapidly evolving field that is affecting nearly every aspect of our lives. It discusses some of the basic techniques and ap-

plications of AI so that readers can gain an appreciation for the field through a sample of AI technologies. The chapter covers four main areas of AI applications: **machine learning** (ML), **neural networks** (NN), **natural language processing** (NLP), and **computer vision** (CV). Within each of these, it briefly describes subareas, with some example applications. Some subareas are related to or overlap with others, in some cases sharing data needs and analytical processes.

To innovate, planners should understand the general capabilities of AI methods and technologies so they can recognize opportunities for change and improvement (Batty 2022). Planners will also be involved in discussions about analyses that will increasingly incorporate AI methods, and to do so, they must understand these terms and concepts. This chapter seeks to provide this information to increase planners’ understanding of these approaches and the opportunities that they represent.

As readers consider the following characteristics and applications of AI, they may begin to recognize methods that could apply to individual planning tasks as well as organizational processes. Chapter 3 will take the discussion further and explore some current applications of AI to planning practice.

A BRIEF HISTORY OF AI

The notion of nonhuman intelligent machines goes back to ancient times. The mythical Greek inventor [Daedalus](#) was reputed to have made statues that appeared lifelike and could walk. In ancient China, there are accounts of automatons



Figure 2.1. “Maillardet’s automaton,” created in the early 1800s (Daderot/Wikimedia)

AI Milestones




Figure 2.2. Timeline of significant events in the development of AI (Thomas Sanchez)

from as early as the 3rd century BCE, when a mechanical orchestra was assembled for the emperor (Mayor 2018). And statues coming to life were a common occurrence in myths and legends, such as that of [Pygmalion](#) creating [Galatea](#).

Much later, during the Renaissance, machinists created life-sized automata resembling humans that could write, draw, or play music (Russell 2017). In the early 1800s, "[Maillardet's automaton](#)" (Figure 2.1, p. 13), an automaton created by Swiss mechanist Henri Maillardet with the ability to write three poems and draw four illustrations, toured Europe and the United States. By the late 19th century, these types of devices had become very popular, with automata created by French dollmaker Gustave Vichy being displayed as toys in upscale Parisian department stores.

These efforts focused on simulating the physical form and actions of the human body, but not the abilities of the human mind. Though during this time philosophers from all over the world had been creating formal reasoning techniques to understand the human mind, there were no scientific methods and capabilities to build mechanisms that could simulate the logic processes they were theorizing (Crawford 2021). It would not be until the 20th century, with the development of computers, that "artificial intelligence" in mechanized form was conceptualized.

By some accounts, modern computing began in the 1940s when Princeton mathematician [John von Neumann](#) envisioned a computer-like process that could store pro-

grams (i.e., digital computer instructions). The computer program could be executed using stored memory, more closely achieving the functionality of a human brain. This aligned with the work of [Alan Turing](#)'s World War II code-breaking device that successfully decrypted German military messages (featured in the movie [Imitation Game](#)). The device, known as a "[bombe](#)," far surpassed human, manual processes for testing combinations and problem solving, but was an analog machine rather than a digital computer (Dyson 2012).

An important early AI milestone was a 1956 conference held at Dartmouth University sponsored by the Defense Advanced Research Projects Agency (DARPA). Participants included AI researchers [Marvin Minsky](#), [Oliver Selfridge](#), [Allen Newell](#), [Herbert Simon](#), and [John McCarthy](#). As noted above, the term "artificial intelligence" is commonly attributed to McCarthy, while Newell and Simon have been credited with one of the first AI programs, the [Logic Theorist](#).

The Dartmouth University conference generated substantial enthusiasm about the potential for a "thinking machine" with the capacity to "learn," and significant financial support began to flow from government and businesses for this research (Buchanan 2005). Substantial gains in AI during the next 20-year period included the [General Problem Solver](#) (GPS) [algorithm](#), developed in the late 1950s, which was fundamental to more advanced cognitive structures, and the AI programming language [LISP](#).

Despite decades of well-funded international research, however, scientists were unable to successfully build their envisioned “intelligent machines” and there was a relative lull in scientific advancement. Researchers struggled due to a lack of funding in an “AI winter” from the mid-1970s until the 1990s. But by the end of the 1990s, American corporations had gained renewed interest in AI, partly as a result of considerable advancements in computer hardware. This also included large investments by countries such as Japan to create “[fifth-generation computing](#)” that enhanced computing speeds and AI.

Over the last 20 years, corporations such as IBM, Amazon, Google, and Baidu have been successfully applying AI techniques in specialized fields to significant commercial advantage. Many of the digital services we use daily (the internet, personal computing, communications, navigation, and others) contain AI. As a result, these technologies are increasingly influencing many aspects of our lives.

The four main areas of AI in use today with which planners should be familiar—machine learning (ML), neural networks (NN), natural language processing (NLP), and computer vision (CV) (Figure 2.3)—are all being tested on planning-related activities. Most of the current applications of AI can be considered “weak AI,” which simply mechanizes limited human thought and behavior patterns. The growing pervasiveness of AI technologies and the unforesee-


able potential consequences of “strong AI” or even artificial “superintelligence,” however, as described in the sidebar on p. 16, are a concern to many.

MACHINE LEARNING

In [data science](#), the term [machine learning](#) (ML) is frequently used interchangeably with AI, though ML is a branch of AI. The basic principle of ML is that with appropriate data, computers can “learn.” The goal of ML is to enable machines (i.e., computers) to detect and “memorize” patterns more quickly and accurately than humans, and then use this information as part of decision-making processes.

An important distinction to make here about different types of AI is that of “rule-based” versus “learning-based” systems. A [rule-based AI system](#) relies on predefined rules, [heuristics](#), and [algorithms](#) that are created by human experts. These rules are used to make decisions or perform tasks based on specific input data. The main advantage of rule-based AI is that it is easy to understand and explain the decision-making process, and it works well for well-defined problems with limited variability. It can also produce results quickly, as long as the rules are clear and comprehensive (U.S. GAO 2021). However, the main disadvantage of rule-based AI is its limited flexibility, as the system can

Figure 2.3. Four main areas of AI in use today: machine learning, neural networks, natural language processing, and computer vision (Thomas Sanchez)



LEVELS OF AI

When discussing specific examples of AI methods, it is useful to first understand the evolution of approaches and applications. AI is a broad term encompassing several layers of sophistication, resource needs, and objectives. There are three general levels of AI referred to today: artificial narrow intelligence (ANI), artificial general intelligence (AGI), and artificial super intelligence (ASI) (Figure 2.4) (Bostrom 2014).

Artificial narrow intelligence (ANI), which performs simple tasks, represents most of today's AI technology. Also known as "weak AI," ANI simulates basic human logic and behavior as defined by a limited set of rules and circumstances; it does not reproduce human "intelligence" in the larger sense. Examples include internet search tools, facial recognition, and text-to-speech conversion—systems that can be taught to carry out particular tasks. Online commerce and entertainment services such as Netflix, Spotify, and Amazon use narrow AI to classify user preferences resulting in recommendations. In the last decade, advances in machine learning and deep learning (discussed later in this chapter) have enabled important advancements in ANI.

ANI typically only has a small amount of memory and is considered to be "reactive." Reactive AI mimics the human thinking involved with handling various stimuli without accessing prior knowledge or deep reasoning; it usually lacks memory and data storage capabilities. A computer playing chess, a game with well-defined rules, is one example of ANI. More sophisticated AI can collect and store data, allowing computers to access previous facts to guide future decisions. ANI can be seen as a building block for larger systems that can be integrated with other AI tools and processes.


Artificial general intelligence (AGI), also called "strong AI," is focused on performing a full range of tasks. It represents the proposition that machines can be trained to learn human behaviors, think like humans, and use the fundamental aspects of human consciousness. To perform with near-human levels of intelligence and carry out every task that a human being is capable of, AGI will need to master critical capabilities related to [sensory perception](#), [motor skills](#), [natural language skills](#), human-like creativity, and comprehension of human social and emotional characteristics. This will produce powerful AI with extensive knowledge and cognitive computing skills that can handle data efficiently and at incredible speeds, giving it advantages over human cognitive skills.

AGI is expected to understand the context of a task, which leads to self-evaluation, learning, error correction, and improved performance on an unsupervised basis. A [Level 5 autonomous vehicle](#) (AV) without the need for human intervention is one example of an AGI application. Being able to handle any or all scenarios that arise while driving illustrates the level of complexity that would be expected from an AGI system.

Going beyond AGI, the concept of **artificial super intelligence** (ASI) has been described as "an intellect that is much smarter than the best human brains in practically every field, including scientific creativity, general wisdom, and social skills" (Bostrom 1998, 1). ASI represents the point at which robots (or other intelligent agents) will develop sufficient self-awareness and self-vigilance to outperform human intelligence and behavioral abilities.

Machines that possess "superintelligence" will have neural capacities exceeding that of the human brain, making

Figure 2.4. The three levels of AI (Thomas Sanchez)



them capable of conceiving abstractions and interpretations beyond that of humans. As conceptualized, ASI will replicate human behavioral intelligence, using better memory and processes superior to human intelligence to solve problems and digest information, events, and stimuli more quickly. Not only will ASI be able to comprehend human expressions, but it is also expected to communicate emotions, beliefs, and desires of its own. Such superintelligent robots will have decision-making and problem-solving abilities that are significantly more reliable and superior to those of humans.

As noted above, today's AI technologies fall on the narrow end of the AI spectrum. Scientists and engineers are still working to develop computers with cognitive abilities comparable to humans, though they have made some significant advancements, such as IBM's [Watson](#) supercomputer. A significant number of scientists and professionals contend that AGI will not be viable for at least several decades, with predictions ranging from 2030 to

2060 before the arrival of such a system (Goertzel 2014). Some experts predict that given recent developments, ASI may exist by the end of the 21st century (Kelly 2017).

Although having such powerful technologies at our disposal may sound enticing, the idea itself is fraught with unintended and unknown consequences. We can only speculate as to what effect this may have on the lives of human beings. Dystopian views of an ASI future—as depicted by science fiction in which humans are dominated by robots—assume that there is an innate evil within these systems and that the superintelligent agents will turn against the people that created them. English theoretical physicist [Stephen Hawking](#) has cautioned, “The development of full artificial intelligence could spell the end of the human race” (Cellan-Jones 2014). A significant amount of concern has been expressed by experts in the field, alerting us to possible dangers and encouraging us to be thoughtful about the potential repercussions of this new technology to maximize the benefits and avoid harm.

only handle situations defined by its rules. Additionally, it is difficult to maintain and update when the problem domain changes or expands, and it cannot learn or adapt to new situations.

In contrast, a [learning-based AI system](#) learns and adapts to new situations by analyzing data and adjusting its algorithms accordingly. This is typically achieved through techniques such as ML, [deep learning](#) (discussed further below), and [reinforcement learning](#). Instead of relying on predefined rules, [learning-based AI](#) systems develop their own understanding of the problem domain through experience and data analysis. This allows them to become more flexible and adaptive to new and changing situations. The main advantage of learning-based AI is its ability to handle complex, real-world problems that require adaptation and generalization. It can also improve its performance over time as it learns from new data. However, the main disadvantage of learning-based AI is that its decision-making process can be more difficult to understand and explain, and it may require large amounts of data and computational power to train effectively.

[Supervised learning](#) and [unsupervised learning](#) are the two main aspects of ML data frameworks. In supervised

learning, a collection of accurately [labeled data](#) is critical to the process. The algorithm picks up the proper correlations from these instances and uses them to associate newly encountered [unlabeled data](#). In unsupervised learning, unlabeled data is used as the first input, and the algorithm does not consider outputs. Unsupervised learning can be used to identify commonalities by clustering together comparable data points while not having the actual label or [classification](#) of the observation. While supervised learning involves the computer being “taught” by labeled data, in unsupervised learning, the computer “teaches itself.”

ML algorithms are frequently employed in prediction models; thus, ML and predictive analytics go hand in hand. Analyzing upcoming events in light of prior results from related events is one example (similar to time-series statistical analyses using labeled data). Four primary ML applications are sentiment analysis, spam detection, recommendation systems, and predictive maintenance.

Sentiment Analysis

Sentiment analysis examines text using ML and [natural language processing](#) (NLP) technologies (discussed further below) to determine the perspective or attitude that it




Figure 2.5. Unsupervised clustering for sentiment analysis (Thomas Sanchez)

conveys. The process usually classifies or predicts where a text’s sentiment falls based on a continuous scale from positive to negative.

Sentiment analyses are increasingly used by businesses to summarize large sets of customer reviews in text form. Sentiment analysis algorithms can be applied to text samples of customer comments after gathering input through various channels, such as Twitter, Facebook, Yelp, and customer service platforms, to learn how customers felt about a product, service, or experience. Automating this process is especially useful in cases where there are many responses in an [open-ended text format](#), such as free-form comments.

A scored word list is one approach to sentiment analysis. The processed list comprises words that are assigned scores ranging from -3 to +3 (very dissatisfied, dissatisfied, slightly dissatisfied, neutral, slightly satisfied, satisfied, and very satisfied). To determine a final sentiment index, the text is [parsed](#) into individual words that are compared to the scored word list. In a simple example, if a customer said, “I love the product, but it was delivered late,” sentiment analysis would identify two words on the processed word list, “love” and “late” (with scores of +3 and -2, respectively), and disregard the other words. These two words produce a score of +1; therefore, this statement would be categorized as “slightly satisfied.”

Industry can use very sophisticated implementations of sentiment analysis, including precise scores from algorithms for lengthy texts. In addition, [reinforcement learning](#) models can continue to add intelligence to the process—by updating the scored word list in the previous case, for example. NLP and ML methods can be combined to create complicated models with different algorithms for sentiment analysis

depending on the complexity of the data and the level of precision desired.

There are three primary models for sentiment analysis. **Automated systems** rely on clustering ML models (i.e., unsupervised labeling) or [network analysis](#). Comments can be clustered based on the similarity and co-occurrence of words or terms with no prior word scoring available (Figure 2.5). An analyst then reviews the clusters and assigns an appropriate label (e.g., positive, neutral, negative, or other). No score along a scale or index is used in this case.

Rule-based techniques, as described above, rely on predefined rules to categorize data. [Tokenization](#), [parsing](#), and [stemming](#) are a few methods needed to preprocess the text. The earlier example of scoring “I love the product, but it was delivered late” is a rule-based strategy. Rule-based systems can be customized or tuned to improve performance over time. By creating more intelligent rules, these algorithms can be modified depending on the circumstances, such as changes in the topic or terminology involved. To guarantee consistent and enhanced results, rule-based models need to be routinely updated. This can be done by adding additional correctly labeled data, which takes more time and resources.

The most advanced, effective, and popular method for sentiment analysis is a **hybrid methodology**, which combines the advantages of both automatic and rule-based systems. Hybrid models benefit from both the flexibility of customization and the strength of machine learning. An example is a self-updating word list built using a standardized database. ML can be used to monitor and update the word list following an organization’s requirements, but outcomes must be periodically validated by humans. Such methods are examples of [“human-in-the-loop”](#) systems.

Spam Detection

Everyone is familiar with spam, or junk email—unsolicited email sent in bulk to a large list of addresses, often with messages that are meant to deceive or trick recipients into sending money or revealing confidential information. As [worms](#), [viruses](#), [Trojan horses](#), and other [malware](#) are often sent via spam, this is increasingly being recognized as a serious security issue. [Knowledge engineering](#) (or rule creation) and ML are two popular methods for detecting spam emails.

In knowledge engineering, a process similar to sentiment analysis, email messages are categorized as spam or valid based on a set of rules. These rules can be constructed by an individual using a filter, or the software an email provider uses can learn over time as a user classifies emails as spam. Since rules need to be updated often, this strategy may not ensure effective results depending on the volume and variety of language being used.

For spam detection, ML is more effective than knowledge engineering. Based on a set of labeled training samples—emails that have already been categorized as valid or spam—an ML algorithm “learns” the categorization criteria to cluster or classify text-based data (i.e., emails) and filter spam. Leading email providers such as Gmail, Yahoo Mail, and Outlook use combinations of ML techniques in their spam filters to address the dangers posed by spam. These approaches go beyond simply scanning spam emails using pre-existing rules because ML can adapt to changing strategies and create new rules (i.e., learn) using what it has analyzed.

Recommendation Systems

Another ML subclass known as recommendation systems is used to rank or rate products or services for users or customers. These systems anticipate the ratings a user would give to a certain product based on previous purchases, actions, or characteristics.

Recommendation systems (Figure 2.6) are frequently employed by well-known corporations, including Google, Instagram, Spotify, Amazon, Reddit, and Netflix, to boost user and platform engagement. To encourage the use of its platform to listen to music, Spotify, for instance, will suggest tracks that are similar to those a user has “liked” or listened to frequently. Based on the information it has gathered about similar users, Amazon uses recommendations to suggest certain products to different consumers.

Recommendation systems can be built in a variety of ways. Some employ formulaic and algorithmic methods, such as [PageRank](#), while others use modeling-based methods, such as [collaborative filtering](#), [content-based methods](#),




Figure 2.6. Example of a conceptual recommendation engine (Thomas Sanchez)


and [link prediction](#). Complexity does not necessarily equate to better performance. Just as with other ML methods, simple strategies and tactics can improve performance over time. Major corporations including Reddit, CNN, and Google have promoted content on their platforms using straightforward, formulaic implementations of recommendation engines. The more customer or user information available (e.g., prior purchases, items searched, demographic information, and other interests), the better the potential recommendations.

Predictive Maintenance and Anomaly Detection

Predictive maintenance (PdM) uses ML algorithms trained on industrial sensor data to predict future events, such as determining the likelihood that machinery or equipment will break down. PdM can determine maintenance scheduling by combining data, statistics, ML, and modeling (Mobley 2002).

To monitor and identify equipment problems in real time, PdM considers expected service intervals in addition to data-driven insights based on measurements of operational conditions. The data collected over time will eventually produce patterns and [anomalous activities](#) that can be classified accordingly (i.e., as malfunctions), allowing for the detection of irregularities in automated activities before they develop into significant problems that could harm the operation. For instance, a spike in energy usage that falls outside of a specified error range can be detected and researched further (Figure 2.7, p. 20). By using this predictive analysis, PdM helps maintain the operating efficiency of equipment in service while avoiding expensive repairs. This approach can be applied to processes that involve a mix of equipment.

Figure 2.7. Example of anomaly detection for a value out of expected range (Thomas Sanchez)



Because it can identify potential faults and can reduce downtime and waste, PdM is becoming valuable for system monitoring. Sensors can track machine and operating conditions, and this information can be integrated with manually collected data and cloud storage systems. A rich dataset for PdM involves the widest variety of diagnostic and maintenance information over long periods. Organizations are better able to implement maintenance plans when they are needed—and anticipate the next equipment breakdown or failure—when they can identify patterns and signals in sensor data.

NEURAL NETWORKS

Artificial neural networks, also called neural networks or neural nets (NN), are computerized systems based on how the brain processes information. They are a common approach used in ML to identify patterns, determine the likelihood of a specific outcome, and learn by using feedback loops (Swingler 1996). They can form connections between two items and learn to correlate them with one another, which is the basis for making predictions.

NNs are more practical today given the increasing availability of very large datasets (also known as “[big data](#)”). NN variants can combine numerous NNs to attain a greater level of complexity. Similar problems can be analyzed with [multiple regression analysis](#) and NN; however, NN by nature can discover nonlinear and interactive relationships, while these structures need to be predetermined for regression models.

The basic architecture for NNs was modeled after the structure of the human brain. Because of this, many terms used in NN are drawn from the field of neurology. As the neuron is the core component of the human brain, the “perceptron” is the analogous computational element of NN structures (Figure 2.8, p. 21). Similar to the neuron, the perceptron takes in a set of input signals and converts them into output signals. The perceptron then attempts to analyze the data by stacking numerous layers, or groups of computing units, that can recognize repeatedly occurring values. The job of each layer is to interpret a particular pattern within the data.


The term “[deep learning](#)” (DL) refers to a neural network with a depth of multiple layers. In contrast to the basic AI functionalities available through ML, DL allows for systems that can carry out more complicated tasks and learn a wider variety of functionalities. The main uses of DL today include speech recognition and more general pattern recognition.

Speech Recognition

From voice search to virtual assistants, AI-powered speech recognition systems are increasingly becoming part of daily life. DL advances have produced great gains in this area.

Automatic speech recognition (ASR) involves translating spoken words into text, often in the form of a command. The most cutting-edge software of today can recognize different accents and dialects. ASR is frequently used with virtual assistants, live captioning, and clinical transcription

Figure 2.8. Neural structure in the brain (neuron) and an artificial neural network (perceptron) (Thomas Sanchez)



in a wide range of fields. For many applications, accuracy is critical. These tools will likely become even more widely used over time, replacing the keyboard as a primary input device (Mitchell 2019). Speech recognition systems are distinguished from voice recognition because they extract words from spoken language, where voice recognition identifies a person through their speech, similar to the uniqueness of a fingerprint.

[Natural language processing](#) (NLP) models, discussed further below, are important to ASR. NLP can be used to add capitalization and punctuation and check grammar in transcripts at the end of the ASR process, in addition to its use in [large language models](#). After NLP post-processing, transcripts can be used for tasks involving language modeling, such as sentiment analysis, text analytics, chatbots, and personal assistants.

Both traditional statistical algorithms and DL can be used to develop speech recognition algorithms to translate audio into text. However, DL for speech recognition over the past few years has exceeded the performance of statistical models (Sarker 2021). DL algorithms are more accurate even with noisy data and are better at understanding accents, dialects, context, and numerous languages. DL speech recognition models can be developed using many software tools and services from Google, Amazon, and Microsoft, and voice recognition models can be created with the aid of open-source toolkits. Like other AI tools, speech and voice recognition are becoming easily accessible through popular programming language modules in [Python](#) and [R](#), as well as web-based platforms.

Pattern Recognition

The capacity to recognize patterns is a fundamental aspect of intelligent systems. Humans naturally look for patterns within whatever type of information they are consuming. *Descriptive* pattern recognition classifies the identified pat-

terns; *explorative* pattern recognition tries to identify data patterns in general. AI pattern recognition addresses both approaches and uses a variety of techniques, depending on the application and type of data. As a result, pattern recognition comprises a large body of information and skills that are often unrelated.

In general, [pattern recognition](#) is an automated way to detect regularities or predictable sequences and orders of values within a sample of data. Computers monitor their surroundings, learn to discern different patterns of interest from their backdrop, and then reason logically about the categories of the patterns (Rich, Knight, and Nair 2009). The specified objects are categorized during the recognition process. Definitions of pattern recognition vary because of the broad range of data types that can be analyzed. Words or text, images, and audio recordings can all be used as data inputs for pattern recognition. As a result, pattern recognition is more general than [computer vision](#) (discussed below), which focuses on image identification. Biology, psychology, health, marketing, and computer vision are just a few of the engineering and scientific fields where automatic and machine-based recognition, description, categorization, and grouping of patterns are used extensively.

A pattern can be any object of interest that appears over time or space and has some level of significance, regardless of medium. Common illustrations of patterns are fingerprints, letters, numbers, faces, or speech signals. An observation of a pattern can be made quantitatively by using statistical techniques, or it can be made visually, such as in pictures and films. The classification or categorization task is typically used to represent the recognition problem. Classifications are either created by the system (supervised categorization) or are discovered through pattern similarity (unsupervised classification). The use of these methods depends on the types and volume of data.

NATURAL LANGUAGE PROCESSING

Natural language processing (NLP) is a subfield of AI in which computer methods analyze human language, text, or verbal communication to derive meaning. NLP can perform tasks such as automatic summarization, sentiment analysis, speech recognition, and *topic modeling*. For text, NLP is used to analyze the ideas conveyed by terms, sentences, and phrases (Obermeier 1989). NLP systems can also correct grammar, convert recorded speech to text, and translate between languages.

NLP is complicated by the fact that human language usage is imprecise, with great variety in how it is used. To understand human language is to also understand concepts and context and how they interact to create meaning. Despite language being relatively easy for the human mind to learn, ambiguities of language make it difficult for NLP to be precise and consistent.

NLP algorithms are commonly based on ML algorithms. Instead of hand-coded rules, NLP can rely on ML to automatically learn these rules by analyzing documents and samples of text. Like other areas of AI, more training data usually leads to more accurate models and performance. Common applications of NLP include virtual assistants, text analysis, and content creation.

Virtual Assistants

Chatbots, also known as virtual assistants, can communicate with users through a variety of applications and platforms. They can understand voice commands and initiate requested tasks. They can also extract information from conversations for analysis or input to other processes or tasks (Brandtzaeg and Folstad 2017). These technologies are being integrated into automation processes to help increase productivity.

Task-oriented chatbots are the simplest type of chatbots, serving well-defined purposes to execute specific tasks. Using rule-based NLP combined with ML, chatbots can respond to user queries in conversational ways and are often used to support customer service functions. Task-oriented chatbots can handle traditional and general inquiries, such as providing hours of operation, business locations, or other information that does not involve variables or require decision-making.

Predictive AI chatbots are more advanced and personalized than the task-oriented variety. Conversational virtual assistants, such as Apple's Siri and Amazon's Alexa, are contextually aware, using *natural language understanding* (NLU), NLP, and ML to acquire information during opera-

tion. They use predictive intelligence and other AI capabilities to personalize conversations. These chatbots store a user's preferences while offering responses and recommendations, similar to that of a recommendation system. *Large language models* (LLMs), such as OpenAI's *ChatGPT*, also serve as virtual assistants by interpreting user prompts and generating information and content about those prompts. The sidebar on p. 23 delves into LLMs in more detail.

Text Analysis

Text mining is an AI method that uses NLP to transform or extract text from documents and databases into structured data for analysis with ML. Text mining can identify facts, relationships, and assertions within documents or other textual data.

Knowledge-based businesses examine large sets of documents with text mining to discover new insights, extract important topics or themes, or answer research questions. Structured data generated from the mined text can be stored in a database or used in a business intelligence dashboard and used for descriptive analysis. A simple example of this is word or term frequency analysis. Frequently used or related terms in a document often indicate the overall topic of a document.

Content Creation

Content creation can be expensive and time consuming for organizations that rely on frequent communications with users or consumers. AI and NLP can be used to create content from available data resources. In this process, a computer uses rules of communication (e.g., sentence structure, grammar, spelling) and stored data to compose material or information. Specified parameters that direct the creative process can be adjusted to fit the content.

AI can perform content creation in a variety of ways, depending on the specific type of AI and its capabilities. Some AI systems, such as NLP models and LLMs, can be trained to generate text that is similar to human language. These models can take a given prompt or input and generate text that follows certain rules or patterns, such as a specific writing style or topic. Other AI systems, such as ML models, can be trained on large collections of existing content, such as articles or videos, and then use that training to generate new content that is similar to the existing content. The specific approach used by an AI system for content creation will depend on its design and intended use.

In late 2022, conversational chatbots emerged and gained significant attention for having advanced language features. These include the input or prompt format, which

LARGE LANGUAGE MODELS

Large language models (LLMs), such as OpenAI's [GPT-4](#), are AI systems that interpret and generate human language. These models use a type of NN known as a [transformer](#), which is composed of interconnected nodes or "neurons" arranged in layers. These layers work together to process language and learn from the patterns in data. This results in the ability to predict the likelihood of word or term sequences.

To create an LLM, developers assemble a huge amount of text data from the internet and other digital sources. The resulting dataset teaches the model to recognize language patterns and structures. When data is fed into the model, it is divided into smaller units called [tokens](#), which are then processed by the NN. As the model has been trained on a wide variety of texts, it can produce coherent and contextually appropriate responses, making its conversational output easily understandable for users.

For someone interested in building their own LLM, the process begins with acquiring computational resources, as these models require substantial computational power and memory. Next, a diverse and comprehensive text dataset must be compiled from various sources, such as websites, books, and articles. This data must then be cleaned and preprocessed, converting the text into a suitable format for training. An appropriate NN architecture, like the [transformer](#),

should be chosen, followed by feeding the preprocessed text data into the model for training. The model's performance should be continuously evaluated and improved using a separate dataset. Once the model achieves a viable performance level, it can be deployed for various applications based on the content of the model.

Building an LLM for a specific profession or discipline is feasible, but it requires significant resources, expertise, and access to relevant, high-quality data. For disciplines with a smaller scope or less available data, it may be more practical to customize pretrained language models, such as GPT-4, to cater to the specific needs of that field. By doing so, professionals can leverage the power of LLMs without needing to build one from scratch, ultimately saving time and resources while still benefiting from the capabilities of these AI systems.

For the planning profession, LLMs trained with local zoning codes, building regulations, and land-use plans could assist with development review and generate ideas for future plans. Additionally, LLMs could facilitate communication between planners, policy makers, and community members by generating user-friendly visualizations and explanations of planning concepts and local policies and regulations, ensuring that diverse perspectives are considered and fostering a more inclusive and democratic approach to urban development.

uses natural language rather than a search term made up of keywords commonly used for web searches. The other notable feature is the range of output styles that chatbots can produce, such as blog posts, outlines, recipes, essays, and other formats generated from the language model. This has significant implications for any activities involving text-based communications. One example of an AI content-creating chatbot is OpenAI's [ChatGPT](#), which, given the prompt "How does AI perform content creation?", generated the prior paragraph.


In the future, these tools are likely to become more sophisticated, incorporating advanced AI algorithms and ML techniques that include text, analytics, graphics, and video. For example, they may be able to process vast amounts of data in real time, improve predictions, review development proposals, and provide planners with more detailed

and interactive visualizations. Already, ChatGPT can be used to write computer code to create simple websites and draft workflows for software programs such as ArcGIS, SketchUp, Photoshop, TwinMotion, Unity, Unreal Engine, or CAD software (Plass 2023). LLMs such as ChatGPT hold great promise as knowledge management tools that can be used by planning organizations, planning consultants, and others involved in the development process to support more data-driven and efficient planning while helping to incorporate the needs of residents and promote sustainable growth.

COMPUTER VISION

[Computer vision](#) (CV) is a subfield of AI and computer science that focuses on enabling computers to interpret and

Figure 2.9. Edge detection to identify objects
(Photo by Nenad Radojčić/Unsplash;
edge detection by
PineTools Online Edge
Detection)



understand visual information, such as images or video. The aim is to encompass both visual and cognitive functions to capture humans' ability to make sense of what they see (Wasserman 2020). Both human vision and CV systems include sensors and interpreters. In a CV system, a camera is the equivalent of a human eye. Other hardware used to collect information include [proximity sensors](#), [laser scanners](#), and [radar](#), which can be used in different combinations depending on the context and objective.

CV emerged in the late 1960s and developed in close parallel to AI. Tasks such as [color segmentation](#) and [edge detection](#) (Figure 2.9) were developed early and formed the basis of CV applications, but computational reasoning for CV has been a significant challenge to the scientific community. Humans can easily distinguish between foreground and background objects in 2D images, but computers currently lack reliable depth perception. More importantly, they cannot extract higher-level information from images unless related training data has been provided.

Along with CV, image processing involves the analysis of raw imagery to extract or transform the information, usually to prepare images as input for further analysis. For example, operations to reduce noise, augment contrast, and rotate data are image-processing tasks that can be performed at the pixel level and do not require a complex understanding of the objects in the image or the overall context. Image processing is also critical for land use and land cover detection derived from satellite imagery (Figure 2.10, p. 25).

CV informed by rule-based intelligence is commonly used for industrial purposes. These uses include inspection tasks in production or manufacturing, such as determining whether containers are clean, empty, and undamaged and verifying that a product is properly packaged. CV can also address more sophisticated problems, such as facial recognition and detailed image analysis. Facial recognition attempts

to identify a person based on facial structure and features associated with images contained in a database. Biometrics including facial recognition, voice recognition, fingerprint matching, and retinal scans all rely on pattern recognition drawing on DL methods. These technologies are mainly used for security and law enforcement, which can have controversial applications and social justice concerns.


There are an increasing number of CV applications related to transportation analysis. The use of this technology in vehicle classification, for example, has improved significantly with advancements in AI (Figure 2.11, p. 25). Roadways and waterways can be monitored to detect, track, and classify vehicles using a range of sensors including closed-circuit television (CCTV) cameras, [Lidar](#), and [thermal imaging](#) devices (Dimitrakopoulos, Uden, and Varlamis 2020). Within vehicles, CV technologies can be used to segment different traffic configurations, including DL algorithms that recognize traffic signs. Driver assistance systems in autonomous vehicles use these along with detection algorithms for pedestrians, lane striping, and other vehicles. Deep NNs have recently also been used to develop AV collision avoidance systems (Everett, Chen, and How 2018).

CV can also be used to monitor the condition of asphalt pavement and concrete infrastructure. A CV-based vulnerability detection system can assess pavement condition and provides data that can be tracked over time. Similarly, CV can help perform more efficient and consistent assessments of neighborhood and property conditions using large amounts of training data (imagery) from supervised learning sources.

CONCLUSION

The rapid evolution of AI in concert with computer technologies is expected to increase with each successive im-

Figure 2.10. Computer vision used to categorize land-use characteristics in satellite image analysis (USGS.gov)



provement or innovation in our understanding of cognitive processes (Kurzweil 2013). This includes the ability to build machines capable of matching or exceeding the capabilities and speeds of the human brain. *Quantum computing* is expected to revolutionize computing and AI in this respect, primarily due to the theoretical computing speeds that can be achieved (Kaku 2015). In addition, the popularity and growth of scientific exploration and applications are experiencing growing momentum that will lead to further advancement.

The selected application areas of AI briefly described in this chapter are examples of narrow or weak AI that are broadly available and potentially useful to most planning organizations. In particular, the rapid emergence and continued development of LLMs could play a significant role in revolutionizing urban planning. These AI-powered models can analyze vast quantities of data, from city demographics to traffic patterns and environmental concerns, helping planners identify optimal solutions for sustainable growth and livability. This includes generating data-driven insights, optimizing designs, and facilitating collaboration among stakeholders. As with the other AI methods discussed, there are concerns about the sensitivity of these models to bias and past errors of judgment that will require trained human experts to supervise these applications and review their outputs.

Just as urban planning has gradually adopted new technologies over time, we can expect that AI will be used to handle tasks that are currently time consuming and in need of improvement. But for this to happen means that urban planners must be knowledgeable about the associated costs and benefits, both of AI use and for the organizations they work for. The objective is to use better information with improved analysis techniques to inform planning processes. The next chapter illustrates how several of these AI methods are already being used in current planning practice.



Figure 2.11. “**Bounding boxes**” applied by computer vision to identify different types of vehicles on a roadway (Photo by Keszthelyi Timi/Unsplash; bounding boxes created using [Online Object Detector](#))

CHAPTER 3

AI IN CURRENT PRACTICE

Artificial intelligence (AI) offers possibilities to improve the planning and management of cities in many ways. As discussed in Chapter 2, four important application areas planners should know about are machine learning (ML), artificial neural networks (NNs), natural language processing (NLP), and computer vision (CV).

These four areas have the potential to solve a wide range of challenging problems with a variety of data types and techniques. ML allows computers to learn from data and improve their performance over time, and CV allows computers to “see,” enabling them to recognize objects and scenes in images and videos. NLP enables computers to interpret and generate human language, making it possible for them to interact with people more naturally and intuitively. This is a crucial technology for applications such as voice-controlled assistants, chatbots, and text analysis. Finally, NNs are capable of learning from data and making decisions based on that data, which is key in applications such as image and speech recognition.

These four application areas are driving significant advances in AI technology and can be combined to create AI-powered analysis tools, which can help planners process large amounts of data and identify patterns and trends. Planners can use these approaches to classify data, make predictions, and solve complex problems—all relevant to planning practice. This can help planners to make better-informed decisions about the development and management of urban systems.

Another way that urban planners can use AI is through the development of smart city technologies (Hurtado, Hitchings, and Rouse 2021). These technologies can use AI and other advanced technologies to automate and improve city services, such as transportation, energy and water management, public health, safety, and security. The goal of using AI in these ways is to make better decisions and plans that lead to more livable and sustainable cities.

This chapter provides examples of AI currently being used by public agencies for planning purposes. Given the newness of AI in planning, these cases represent early

adopters who are pioneering these approaches and making the first attempts to integrate AI into planning processes. The first is an example of using CV to observe neighborhood conditions, the second uses ML methods for evaluating district characteristics and boundaries, and the third is an example of NLP to assess public input. The chapter also reviews a mix of research approaches using NN. Each contributes to planning practice in different ways.

These examples are all applications of weak AI (see the sidebar in Chapter 2) in which AI was used to identify and explain patterns useful for planning activities. For each of the examples, the application area is briefly discussed, drawing from both planning and AI concepts. Readers are encouraged to consider each example in the context of the problems being addressed, the stakeholders involved, the data and methods used, and the application of results. A specific planning problem can be approached differently by different types of organizations depending on the scope involved, the resources available, and how it fits into the institutional mission. The examples also highlight the context of the problem, the partnerships involved, and other considerations that may have influenced the selected method.

COMPUTER VISION FOR NEIGHBORHOOD CONDITIONS ASSESSMENT

Urban planners use neighborhood conditions assessments as part of the planning process because it allows them to gather important information about the current state of a neighborhood. This information can include the condition of the housing, property conditions, road conditions, and environmental conditions.

By conducting a thorough assessment of neighborhood conditions, urban planners can better understand the needs and concerns of the community and use this information to develop plans that will improve the area and enhance the lives of the people who live there (Zheng et al. 2017). Additionally, neighborhood conditions assessments can help urban planners identify potential challenges and opportunities in the area, which can inform the development of more effective and targeted plans.

One example of how urban planners use neighborhood conditions assessment is to identify areas of a neighborhood in need of repair or improvement. For instance, an assessment might reveal that a high proportion of the housing in a particular area is in poor condition, or that there are few public parks or other amenities available for residents to enjoy. Based on this information, urban planners can develop plans to invest in housing rehabilitation programs or create new public spaces and amenities in the area.

Another example of how urban planners might use a neighborhood conditions assessment is to prioritize the allocation of resources. For instance, an assessment might reveal that a particular neighborhood has a higher-than-average concentration of low-income households or a higher-than-average crime rate. Based on this information, urban planners can prioritize the allocation of resources to address these issues, such as providing financial assistance to help low-income families afford housing and investing in crime prevention programs.

Planners should pay special attention to the ramifications of neighborhood conditions assessments that can lead to discriminatory decisions when such assessments may unfairly target specific neighborhoods. For example, these assessments may unfairly target specific neighborhoods, leading to the allocation of resources, services, or opportunities in a discriminatory manner. In the 1930s, the U.S. government-sponsored Home Owners' Loan Corporation (HOLC) used neighborhood assessments to create residential security maps, commonly referred to as redlining. These maps assigned ratings to neighborhoods based on factors such as race, ethnicity, and socioeconomic status, with minority and low-income communities receiving the lowest ratings (Tate et al. 1997). The redlining practice resulted in the denial of loans and insurance to these communities, reinforcing existing racial and economic segregation (Massey and Denton 1993).

Even though redlining is now illegal, some argue that modern neighborhood assessments can still perpetuate the same discriminatory outcomes. For example, cities

may use neighborhood assessments to target specific areas for “improvements,” leading to the displacement of low-income residents and a further concentration of resources in wealthier neighborhoods (Newman and Wyly 2006). In this case, neighborhood assessments can be seen as a tool for perpetuating existing inequalities. As with any planning technique, planners must be mindful of considering both intended and unintended consequences in ensuring that AI is used in ethical and equitable ways.

Urban Blight Detection

Urban blight is a term that is used to describe areas of a city or town that are in a state of decline or decay. These areas may have a high concentration of abandoned or dilapidated buildings (Figure 3.1), including obsolete industrial buildings as widely experienced in Detroit and other older Rust Belt cities (Schilling and Pinzon 2016); high rates of crime and poverty; and limited access to basic services and amenities.

Urban blight can have a range of negative impacts on a community, including reducing property values, lowering the quality of life for residents, and contributing to social



Figure 3.1. Physical indicators of blight include abandoned, boarded up, and graffitied buildings (Library of Congress, Carol M. Highsmith Archive)

and economic problems. In many cases, urban blight is the result of decades of lack of investment and neglect and can be difficult and costly to address.

Neglect and disrepair also signal deeper problems that can grow worse over time. Residents who have the means eventually move away from these areas, further impacting the vitality of these areas. A lack of investment by local government and property owners, along with residential shifts, accelerates and prolongs blight (Hosseini et al. 2017).

Identifying blighted urban areas can highlight a range of negative impacts that blight imposes on low-income, disadvantaged residents who live in these areas. Some of the most significant impacts include:

- **Reduced access to basic services and amenities.** Blighted urban areas often have limited access to basic services and amenities, such as clean drinking water, reliable electricity, and public transportation. This can make it difficult for residents to meet their basic needs and can harm their overall quality of life.
- **Increased exposure to crime and violence.** Blighted urban areas often have higher rates of crime and violence compared to other parts of the city. This can create a sense of fear and insecurity among residents and can make it difficult for them to feel safe in their neighborhoods.
- **Lower property values.** Blighted urban areas often have lower property values compared to other parts of the city. This can make it difficult for residents to sell their homes or to access affordable housing options and can contribute to a cycle of poverty and economic distress.
- **Social and economic isolation.** Blighted urban areas can be socially and economically isolated from the rest of the city. This can make it difficult for residents to access job opportunities and other economic resources and can contribute to a sense of isolation and disconnection from the larger community.

By bringing attention to these issues, urban planners and policymakers can work to develop strategies and programs to address the associated causes of poor neighborhood conditions to improve the lives of residents in these communities. However, planners should be aware that because there is often a connection between social and economic disadvantage and poor physical conditions, the act of labeling an area as “blighted” can potentially disenfranchise and destroy struggling, underresourced neighborhoods. This is because the label “blight” carries a negative connotation and can stigmatize a community,

which can lead to a lack of investment and support from policymakers and other stakeholders. This lack of investment and support can then exacerbate the problems that are already present in the community, such as inadequate housing, limited access to basic services and amenities, and high rates of crime and poverty.

Furthermore, labeling an area as “blighted” can also have legal and financial implications. For instance, particularly in the urban renewal programs of the 1960s and 1970s, the designation of “blight” was used as a justification for the use of eminent domain, which allowed the government to seize private property for public use. This resulted in the displacement of residents and the destruction of neighborhoods, which had devastating impacts on those communities. Planners must acknowledge such past harms and ensure that blight detection activities do not perpetuate such legacies but rather are used to support and improve neighborhoods for the benefit of existing residents.

Blighted conditions are commonly reported based on complaints from owners of surrounding properties. In addition to their poor aesthetic conditions, blighted properties also have negative impacts on surrounding property values, which are a concern for property owners and impact the city’s tax revenue. Relying on complaints of blighted conditions, however, results in uneven reporting that does not represent the entire jurisdiction and can exhibit spatial bias.

A better approach is to use a systematic and uniform sampling method. Ideally, this would involve observing properties on all accessible streets either on foot or by vehicle. But neighborhood conditions surveys can be costly due to the time and labor involved, and relying on multiple individuals who may apply criteria in different ways may lead to biased or unreliable data. To confront these problems, jurisdictions such as Memphis, Tennessee, and Tuscaloosa, Alabama, have used CV technology mounted on vehicles to capture neighborhood condition data and ML to identify symptoms of blight. A similar process can be used for other urban elements such as street pavement conditions, public parking lot use, and urban flooding (National League of Cities 2020; Ibrahim, Haworth, and Cheng 2020).

Blight Detection in Tuscaloosa

In 2017, Tuscaloosa mayor Walt Maddox and the city council directed staff to oversee a community conditions survey, particularly looking at blight in West Tuscaloosa (City of Tuscaloosa 2018). The survey resulted in over 270 properties being identified as having not met property maintenance or condemnation thresholds. Based on the results of the analy-



Figure 3.2. Cameras used on Tuscaloosa garbage trucks to collect conditions imagery (City of Tuscaloosa)

sis, an executive order was issued to address blight within the study area. Nearly 30 percent of these property issues were resolved quickly due to the relatively minor issues involved. However, obstacles such as a lack of property maintenance inspectors, nonconsolidated code enforcement, inefficient workflows, and COVID-19 slowed progress in the area.

Not long after, then-executive director of the Tuscaloosa Office of Urban Development Brendan Moore, AICP, and University of Alabama economics professor Erik Johnson collaborated on an idea to address issues related to “curb appeal” and blight in the city. Their idea was to use existing city data, CV, and an AI model to generate blight scores. Working together, Moore and Johnson designed a “smart camera” with GPS-based image collection that could be

mounted on municipal garbage trucks to collect image data from neighborhoods in the city (Figure 3.2).

As images are collected on a garbage truck’s route, they are processed using an ML model. After the garbage trucks complete their routes, the images are automatically uploaded via wifi in the garbage vehicle parking lot. The processing of the video assigns a score, indicating the degree of blight or code violations present. A human reviewer then inspects the results for accuracy and adherence to municipal code.

Tuscaloosa’s model was trained with over 4,000 images from building inspections, planning, environmental services, and police documentation to identify indicators of blight and code violations. Most of these images were already on file with the city due to established documentation systems (via case files, follow-ups, violations, etc.), but some new images were hand-collected of properties exhibiting no blight. As with other ML analyses, training data sets are very important for more accurate classifications. The substantial amount of training data that AI models require can pose a challenge to organizations that do not have a robust photo inventory. As noted, Tuscaloosa used images of properties both with and without visible blight, permitting a representative sample of housing conditions across the city (Figure 3.3).

The garbage truck–mounted cameras allowed for properties to be inspected more efficiently and effectively, covering a greater amount of the city during routine maintenance checks. Garbage trucks use predefined routes on a regular schedule, while other municipal vehicles may only be used sporadically or as needed. Using vehicles that do not have established routes poses the challenge of inconsistent and skewed data collection, which goes against the model’s intention.




Figure 3.3. Images from the dataset used to train Tuscaloosa’s blight analysis AI model: left, an example of a blighted structure with boarded-up windows and overgrown vegetation; right, an example of a nonblighted structure (City of Tuscaloosa)




Figure 3.4. Map of Tuscaloosa blight analysis with both weighted (original) and local blight scores (City of Tuscaloosa)

Scores can be used to generate reports for both internal and external uses. Internally, reports can describe existing violations, prioritize properties based on blight scores, or indicate how a property's condition trends, among others. Externally, reports can be trained to send violation letters to property owners, list typical repair fees, or provide contacts and information for assistance programs.

Though weighted scores consider all properties independently, relative (local) scores can also be generated (Figure 3.4). This method allows for a property's conditions to be weighed against its neighbors, accounting for minor conditions that may be normal in one area, but abnormal locally. This approach improves early intervention in remediation, proactively addressing minor issues before they become unmanageable and cost burdensome. Using a local score also permits a more equitable distribution of inspections across the city, reducing geographical and socioeconomic bias that may exist through traditional enforcement methods.

The City of Tuscaloosa plans to further involve stakeholder groups in blight intervention and remediation planning activities. Such groups will include public and private partnerships, with an emphasis on those who can directly assist in providing funding opportunities or educational resources related to property improvements. The City has a particular interest in public-private partnerships related to affordable housing development, including partnerships with Habitat for Humanity, Community Services Programs

of West Alabama, and West Alabama's Chamber of Commerce. Internally, the City's Department of Community & Neighborhood Services (DCNS) provides resources related to opportunities for home-improvement programs and financial assistance and directly assists constituents with program applications. Blight abatement programs and community intervention teams closely integrate with code enforcement officers' and inspectors' workflows, directly connecting these beneficial resources via blight reports and face-to-face interactions with constituents.

Once the model's dataset is further established, the City will also incorporate predictive analyses of blight and code violations. Such patterns may be layered with other historic datasets to understand temporal trends of violation severity, allowing officials to anticipate workloads and prioritize early intervention. Future improvements to the application's model will include a smart score to evaluate additional variables, such as the presence of school-aged children in a neighborhood, property and violent crimes, streetlights, and historic code enforcement data. This type of multivariate analysis can transition toward further proactive intervention, prioritizing property issues and socioeconomic factors more strategically and equitably. A multivariate approach can assist municipalities in making informed decisions when allocating resources to make the most impact on an affected area.

Securing "buy-in" is a significant hurdle to those seeking new technology; challenges include a lack of understanding of future benefits, upfront implementation costs, and the need to train staff to maintain and optimize the application or technology. Appropriate technical staff must be involved in technology projects to optimize the data's use and outcomes and must be able to effectively communicate results with stakeholders and decision-makers.

(Special thanks to Brendan Moore, AICP, and Jayla Blanke from the City of Tuscaloosa for providing information about this project.)

MACHINE LEARNING FOR BOUNDARY ANALYSIS

Zones or subareas are often used for geographic analysis and planning purposes because they provide a way to group similar areas or regions within a larger geographic area. This can be useful for a variety of purposes, such as analyzing and comparing data across different areas, identifying patterns and trends within a particular region, and developing strategies for addressing issues or challenges within a

specific area. For example, a city might use zones or subareas to identify locations that are experiencing high levels of environmental degradation, identify the sources of environmental impact, and then develop strategies for conserving and protecting those areas. Using zones or subareas for geographic analysis and planning allows for a more targeted and tailored approach to addressing issues and challenges within a specific region, rather than applying a one-size-fits-all approach across the entire area.

There are many different types of zones and subareas that can be found within municipalities and regions. These can include zoning districts, such as residential, commercial, and industrial districts; destination areas, such as cultural or tourism districts; and limited-purpose local government special districts, such as school or water districts. Because the composition of zones and subareas changes over time, modifying boundaries can often be necessary to maintain their intended purpose and appropriateness for analysis. Modification triggers can include changes in population or demographics, shifts in economic or social conditions, or the need to address issues related to equity or access. Changing district boundaries can be a complex process that involves several steps and considerations.

ML can be used for boundary analysis in several ways, including mapping and visualization, boundary classification, predictive modeling, and clustering. In addition, ML can be used to evaluate the need for modifications related to changes in composition. Overall, ML has the potential to be a powerful tool for geographic boundary analysis, providing insights into the boundaries of a particular area and helping to inform decision-making related to land use, resource allocation, and other issues.

School District Boundaries

School districts are responsible for providing educational services to the students who live within the district. These districts are typically defined by political boundaries, such as city, county, or state lines, and are usually administered by a board of education or a similar governing body, which is responsible for setting policies, making budget decisions, and overseeing the operation of the schools within the district.

School districts are an important part of the education system in the United States, as they provide a framework for organizing and delivering public education services to students. In addition to operating traditional public schools, many school districts also offer alternative educational programs, such as magnet schools, charter schools, and online learning programs. School districts are usually funded

through a combination of local, state, and federal sources, and may also receive funding from private sources, such as foundations and philanthropic organizations.

School district boundaries influence the quantity and quality of educational resources that are available to schools, as well as the mix of social, economic, and cultural backgrounds of students within those districts (Jargowsky and El Komi 2011). When school district boundaries result in separate and unequal systems, this can have particularly negative effects on a variety of noncognitive outcomes, as the fundamental argument in *Brown v. Board of Education* reminds us. District boundaries classify students into distinct “market areas” rather than separate schools within the same system, which is the main source of today’s public schools’ segregation (Owens 2017). Recent years have seen an increase in the number of new school district boundaries as smaller communities strive to break away from larger school districts.

Creating districts with geographic boundaries is an analytical and political challenge because of the spatial context of residential and demographic patterns. Boundaries can be manipulated to segregate or integrate different groups, such as in the gerrymandering of electoral voting districts in which boundaries are used to separate groups. When creating or modifying district boundaries, planners should consider the characteristics of a population or development through a spatial lens, which adds to the complexity of the task. With school districts, it is important to consider how to better balance underlying community characteristics to not further divide socioeconomic groups. From a planning and policy design perspective, this represents a rule-based system that can facilitate ongoing evaluation and updates. The following example highlights an analytical strategy used by a Texas school district to address this challenge.

Equity Analysis in Ector County Independent School District

As new development increases population levels and changes demographic characteristics in a community, shifts in school supply and demand mean that district boundaries need to be adjusted to accommodate changes in the number of students.

School districts maintain data about their student populations to monitor school performance. As might be expected, the data shows that student populations that are economically disadvantaged exhibit lower levels of academic success compared to more affluent student populations. This suggests that there may be inequities in school resources

and performance. To address this challenge, planners from consulting firm Parkhill and the Ector County Independent School District (ECISD) in Texas used AI methods and spatial analysis to simplify strategies aimed at improving social equity, along with incorporating school districts as significant partners in community planning.

Part of the population in ECISD resides in areas where essential services—including water and sewage access, internet connectivity, and other important infrastructure—have been neglected. Due to the increasing obstacles that families in these locations must overcome, these locations put student populations at a disadvantage. Even if these obstacles are not considered in the academic performance scoring matrix, they will still have an impact on school districts' abilities to meet Texas Education Agency (TEA) performance criteria.

A spatial assessment of equity is a useful way to visualize where a school district stands in terms of socioeconomic characteristics relative to locational factors. This example highlights the link between a school district's socioeconomic status and its physical environment. In Ector County, the results of the spatial analysis showed a high degree of correlation between the physical environment and the socioeconomic burdens of a community.




Figure 3.5. An example of a decision tree; random forest analysis makes predictions based on large numbers of these (Thomas Sanchez)

The Parkhill Planning and Analytics team used two algorithms, random forests (RF) and K-means clustering (KM), to confirm the relationship between the socioeconomic strain on pupils in the district and their physical environment:

- **Random forests (RF)** is an ML technique that is used to build predictive models by combining the predictions of multiple decision trees (Figure 3.5). In an RF analysis, a large number of decision trees are trained on a dataset and used to make predictions. Each decision tree is trained on a randomly selected subset of the data, and the final prediction is made by averaging the predictions of all the decision trees. RF analysis is often used for classification tasks in which the goal is to predict the class or category that a particular data point belongs to. It is considered to be a very effective and reliable method for making predictions and is widely used in a variety of fields.
- **K-means clustering (KM)** is an iterative algorithm that tries to divide the dataset into K pre-defined subgroups (clusters) that do not overlap and to which each data point belongs only once (Figure 3.6). It tries to make the data points inside each cluster as similar as possible while keeping the clusters as far apart as possible. It puts data points into clusters so that the total distance between the data points and the cluster's centroid (the arithmetic mean of all the data points in that cluster) is the smallest. The result assigns membership to each cluster or group.





Figure 3.6. A K-means cluster analysis, in which AI seeks to make data points inside each cluster as similar as possible while keeping the clusters as far apart as possible (Thomas Sanchez)

Figure 3.7. Spatial variables used to identify locational disadvantages in the ECISD analysis (Parkhill)



In this case, the application of ML involved training data to predict the classification of economically disadvantaged (ECODIS) or non-ECODIS categories for the student population of this district, solely based on their physical location. Geospatial aspects such as proximity to floodplains, parks, schools, urban heat zones, or noise zones were used in the ML model, as well as a combination of all these (Figure 3.7).

The resulting map shows three clusters defined by the K-means algorithm, which grouped similar areas based on the school district population characteristics and their geospatial burden of accessibility to assets and resources (Figure 3.8, p. 36). The model effectively predicted the ECODIS category of the student population and identified which combination of variables had the greatest explanatory value. It also predicted the school district's classification with an accuracy of 74 percent, meaning that physical environmental factors alone could explain nearly 75 percent of the variation in school district conditions.

In most school districts, families are asked to self-report their socioeconomic status. Over time, this self-reporting has been increasingly difficult due to low response rates; hence, the use of this methodology can help school districts


target areas where enrollment and self-reporting are a priority. This is a household-level analysis that allows an examination of the proportion of predicted ECODIS students for the whole district. This information can then be used to balance ECODIS and non-ECODIS students by elementary-school, middle-school, and high-school attendance zones.

The most valuable information the school district gains from this prediction of ECODIS is which areas they need to focus on for requesting a self-reporting status. Care must be taken in applying ML in cases like this, however, because the model was trained using data that already contains socioeconomic inequities (i.e., bias). The clustering algorithm enabled the analysis to delve deeper into the population's inequalities that are heavily influenced by geospatial factors.

Once an assessment is conducted, school districts (in collaboration with cities or communities) can prepare strategic plans to modify and improve boundaries based on fairness and equity standards. The data and results generated by the analysis also help to increase the transparency of the process being used by the district.

This example of ML applications provides a strong analytic approach to assessing district boundaries that can

Figure 3.8. Assessment of spatially burdened population in ECISD, defined in three clusters; blue represents highly burdened portions of the population (Parkhill)



be conducted continuously. The first aspect focuses on data collection and data accuracy. It is challenging to obtain or maintain complete demographic datasets over time due to natural changes in household tenure. The second aspect is tracking district composition changes over time that can influence boundary configurations. The approach discussed addresses both of these issues with analytical tools.

(Special thanks to Julio Carrillo, AICP, from Parkhill's Planning and Analytics team for providing information about this project.)

NATURAL LANGUAGE PROCESSING FOR PUBLIC INVOLVEMENT

It is a fundamental pillar of planning that those who live in a community have the right to participate by providing input on how their future will be planned. Because these individuals are the ones who live with the impact of plans, they should be included in every phase of the project, from conceptualization to implementation. Planners understand the close relationship between public participation and participatory democracy.

Democratic systems rely on representative modes, as it would be cumbersome and ineffective for every member of society to participate fully and on an individual basis (San-

chez and Brenman 2013). Information and communications technologies (ICTs) can be used to lower barriers to information sharing, resulting in stronger linkages and more rapid partnership and connection building. It is crucial to distinguish between the actions of the public (i.e., community members) and governmental institutions to comprehend various points of view. Although this is a simplified approach, it highlights how those who are governed and those who are governing, as well as their intermediaries, use technology differently.

The idea that cities can become more intelligent, connected, and responsive to resident needs is one hope of using new ICTs as part of the public participation process (Afzalan, Sanchez, and Evans-Cowley 2017). The emphasis on community members highlights the potential of online participatory technology to enable residents to take an active role in determining the future of their cities.

The transition from traditional planning approaches to the deployment of digital services marks a paradigm shift, not only for participatory methods but for the entire urban planning field. This new planning paradigm was described by Silva (2010) as a mix of two aspects: the broad use of ICTs and the involvement of numerous urban stakeholders. Under this planning paradigm, technology, particularly e-participation tools, replaces in-person, verbal, or handwrit-

ten public comments during community meetings with on-line platforms, allowing for participation at any time, from any place, by anyone. This represents a potential revolution in capturing public input about planning matters because, in theory, anyone with access to the internet could participate.

As we all know, it is more complicated than that, and expanding participation and gaining new voices are still challenges for planning organizations. Planners must continue to promote equitable participation and better incorporate the needs and perspectives of underrepresented community members in urban development plans. While the use of online platforms does open up public input opportunities for those who might not have had the ability to attend a public meeting in person, not all community members have access to the internet or digital devices, and planners must keep this in mind when working to achieve truly inclusive and accessible public participation and engagement processes.

This envisioned shift to online participatory technology has been accelerated by the COVID-19 pandemic, which required replacing in-person public meetings with online platforms to accommodate the need for social distancing. Both virtual and hybrid meeting formats appear to be becoming the new normal, as is the shift from telephone and mail-in surveys to website feedback forms and social media polls. This growth in online public input opportunities is creating larger amounts of public input, and the question—and challenge—is how to read, digest, and summarize those comments, especially when they number in the thousands or millions.

As discussed in Chapter 2, AI tools such as NLP, sentiment analysis, and topic analysis are being used for these specific purposes. Comments, questions, and suggestions from the public can be drawn from emails, comment boxes on websites, social media, or other electronic forms. The two main questions of interest are (1) what are the topics or themes of input? and (2) is the sentiment expressed positive, negative, or neutral—or somewhere along that continuum?

Depending on the level of recruitment involved for public input, coordinated and focused comments can start to take the form of “crowdsourcing,” a concept popularized by Jeff Howe from *Wired* magazine (Howe 2006). The notion of the “wisdom of the crowds” or collection intelligence is very applicable to urban planning problems (Afzalan and Sanchez 2017), which the City of Long Beach, California, is discovering, as described in the case study example below.

NLP can also be used to support inclusive participation of underrepresented populations, which may have

cultural differences in language such as slang and dialect. Each dialect has its own unique set of words and phrases, as well as syntax and grammar. Dialects can affect the accuracy of NLP algorithms, as algorithms may struggle to recognize words, phrases, and grammar from different dialects. But with proper training, NLP algorithms can learn to recognize, understand, and communicate in different dialects, opening up a wide range of engagement opportunities. The example provided in the sidebar on p. 37 highlights work using NLP being done in New Zealand with Indigenous populations.

Sentiment Analysis in Long Beach

In late 2019, the City of Long Beach, California’s Technology and Innovation Department (TID) issued a Start Up in Residence (STIR) challenge (now called the [Smart City Challenge](#)) to crowdsource resident feedback. Community engagement software provider Zencity submitted a successful proposal for the pilot opportunity, which started in February 2020. The TID then received approval from the city council to enter into a contract. It continues to manage the program, which is currently used by more than 14 City departments.

In its crowdsourced resident feedback program, the City uses sentiment analysis on input drawn from a broad range of sources, including not only social media (such as Twitter) but also news sites and Long Beach 311 (Pignataro 2021). As a result, it reflects the interests and the sentiments of residents who contribute using these modes. Additionally, to focus on the conversations of residents, tweets are filtered with selected keywords to exclude those not originating from the Long Beach area that may reflect regional or national dialogue. The sentiment analysis only includes comments expressed on fully public platforms.

Data is collected from all social media interactions (posts, tweets, likes, comments, shares, retweets, etc.) as well as communications from city hotline sources. The total number of interactions shows the level of interest in a topic as expressed by community members who provide their opinions. The platform categorizes comments and interactions by topic and sentiment—positive, negative, or neutral—which reflect levels of satisfaction or unhappiness with specific aspects of the city. Using this methodology quantifies the attention or interest that particular situations garner and evaluates how community members viewed these issues.

In 2021, nearly 5,360,000 interactions were analyzed, with 98 percent obtained from social media (e.g., Facebook, Twitter, and Instagram), one percent from local media websites, and one percent from city hotlines (311). Although

COMMUNICATING WITH MINORITY POPULATIONS IN NEW ZEALAND

Public engagement of Indigenous people takes an additional level of effort in translating and understanding both language and culture. The quality of these translations is very important to communicate respect and goodwill. A focus of the AI platform [FranklyAI](#) is the revitalization of Indigenous languages. Users of the platform take a lot of pride in their language and culture.

Translations must be of high quality for users, especially for individuals wanting to learn these languages. This can involve translating technical terms that do not exist in Indigenous languages. In these instances, close work is required to make sure that any terminology adequately captures the concepts while remaining as transparent to the users as possible. Further, conversation style and flow are unique to each language, so a particular effort is needed to ensure conversations are as natural as possible. Engaging and accurate conversation is only achieved through cultural co-design.

In New Zealand, the FranklyAI team worked with the University of Auckland to develop multilingual and multicultural equivalents of their English student services virtual assistant in te reo Māori, Gagana Samoa, Tongan, Vietnamese, and Mandarin. The team developed multiple AI skills to complement existing virtual assistants and cater to a wider range of the university community through multilingual capabilities. The non-English virtual assistants engage with users in culturally appropriate ways while maintaining the same capability as the English assistants.

There are important concepts in te reo Māori and other Indigenous languages that do not have direct English translations. Engaging in native dialects allows Indigenous speakers to engage confidently, be heard, and improve their experience navigating large amounts of university-specific

information. As a public institution, the University of Auckland is required to provide accessible resources in te reo Māori, as well as actively promote and incorporate te reo and Tikanga Māori through all of its activities.

Another FranklyAI project focused on developing a public consultation conversational agent tailored to Sāmoans and the wider Sāmoan community. This agent, named Tala, features full support for conversations in the Sāmoan language through the use of real-time language translation. This allows clients to continue developing Tala in English, while end users can fluently converse with the agent in Sāmoan.

The Sāmoan language has two ways in which people may converse—traditional Sāmoan, which is used in more formal situations, such as speaking to elders, and more common modern-day conversational Sāmoan. For this conversational agent, it was important to maintain a respectful and appropriate tone so as not to offend anyone, but also to keep the language as informal as possible so that the conversation feels real and natural.

A particularly innovative aspect of this approach is the emphasis on a conversational tone and style that mimics person-to-person conversation, rather than a rigid and impersonal survey form. Survey questionnaires are the traditional approach to public consultation, but it can be challenging to recruit participants. When a combination of fluency, translation, and a human touch can be used, experience has shown that gathering personal opinions and insights can be much improved compared to impersonal survey formats.

(Special thanks to Matthew Ensor from FranklyAI and John Duffy from Beca for providing information about this project.)

one percent may seem like a small proportion, in this case, it represents nearly 54,000 comments or items of information included in the analysis. Figure 3.9 (p. 38) shows the leading topics for 2021, and as might be expected, public health was at the top. Figure 3.10 (p. 38) shows a summary of the sentiment analysis of public comments. Topics can also be tracked over time to gauge changes in public perceptions.

The City used the results of the sentiment analysis to support multiple phases of FY 2023 virtual budget town

halls by providing a synthesis of the qualitative data (e.g., open-ended comments, public remarks, etc.) received from the engagement process. The City also crowdsourced resident feedback on COVID-19 vaccination messaging around safety and effectiveness by analyzing public social media posts. The insights from these analyses proved significantly valuable and were forwarded to various city stakeholders to inform decisions regarding the rollout and distribution of the vaccine, hesitancy surrounding the booster, and inocula-




Figure 3.9. Discourse summary of topics from the 2021 City of Long Beach sentiment analysis (City of Long Beach)

tion for children. Using this information, the City of Long Beach Health Department was able to adjust its messaging when it determined that messaging was not resonating with community members. The City was also able to communicate the safety of the vaccine to residents, clarify information about the booster and its efficacy, and address concerns about the safety and necessity of inoculating children.

This approach to determining public sentiment is not without challenges. AI applications can be used to generate content that mimics public input but instead is intentionally biased to skew results that are based on frequency and sentiment (such as those shown in Figures 3.9 and 3.10). Such schemes could involve bad actors flooding public input channels with many AI-generated messages supporting or opposing a development proposal to influence planning commissions or city councils as they decide the fate of development proposals. Planners will need to be aware of the potential for the use of AI to spread misinformation and stay abreast of developments in measures to counteract this threat through verification of sources, use of AI-generated content detectors, and ground-truthing through face-to-face engagement of community members. These issues are discussed further in Chapter 4.

(Special thanks to Ryan Kurtzman, Smart Cities Manager, and Lea Eriksen, Director of Technology and Innovation and Chief Information Officer, City of Long Beach, for providing information about this project.)

In the **overall online discourse**, positive sentiment was twice as high as negative sentiment (30% vs. 15%, respectively).




Figure 3.10. Percentages of positive (green), negative (red), and neutral (gray) sentiment from the 2021 City of Long Beach sentiment analysis (City of Long Beach)

NEURAL NETWORKS

As discussed in Chapter 2, artificial and deep neural networks (NNs) are methods for identifying patterns and making predictions based on those patterns. Similar to multiple regression analysis, the inputs are a set of characteristics (or variables) and the output is a set of predictions.

The most commonly applied regression technique in planning analysis, ordinary least squares regression (OLS), bases predictions on linear correlations. Because the relationships between complex urban systems are rarely linear and usually based on the interaction of several variables at a time, OLS may not be the best analytical approach. However, NNs are complex, nonlinear models that can be difficult to interpret and explain. While NNs can often achieve high accuracy on a given task, understanding the specific relationships between the input variables and the output is more difficult compared to OLS.

Though few planning departments seem to have yet integrated NN applications into their everyday workflows, there is currently a significant amount of academic research using NN methods for urban planning that may hold promise for better predictions and modeling in areas such as transportation, land use, and urban design. The following examples are drawn primarily from academic

research and are expected to be adopted by practitioners as planners acquire the necessary knowledge and skills about these methods.

Transportation

Many types of questions involved with urban transportation planning have been analyzed with NN methods. One reason is that transportation activities generate a lot of data that is easy to capture. This includes vehicle counts, passenger counts, travel distances, trip origins, and trip destinations.

Transportation infrastructure is mostly made up of linear features such as roadways and transit routes, but for analysis purposes, these are often converted into grids, with each cell representing locational characteristics and spatial relationships (Figure 3.11). For example, a cell may indicate that it has a transit route or stop located in it, and one attribute is the number of jobs within a specified distance of that cell. This is done by querying adjacent cells for this information and associating it with the cell having the transport feature of interest.

Not only can planners use NN to analyze transportation data to detect travel impacts, but the richness of data that NN applications can evaluate can also be used to inform the design or modification of street networks or public transit routes to better accommodate these travel patterns. Examples of recent transportation-related research using NN include the following:

- **Street network design.** Researchers have used a variety of NNs, including deep learning NNs, [convolutional NNs](#) (CNNs), and [recurrent NNs](#) (RNNs), using data on street performance along with associated land-use information (transport features, land-use types, density, topography, etc.) to design, evaluate, and modify street alignments and corridors (see Barreda-Luna et al. 2022; Fang et al. 2022; Tu and Zhang 2022).
- **Disaster and evacuation route planning.** Researchers are accessing large datasets with both historical and current features and applying them to all aspects of emergency preparedness and response. Both timely and accurate information and quick and effective coordination across the responding entities are essential (see Aqib et al. 2020; Chen et al. 2022; Ortiz et al. 2020).
- **Health and safety.** Transportation analysis needs extend beyond designing and evaluating physical infrastructure improvements and mobility outcomes. Researchers are using NNs to assess environmental impacts, such as air quality, as well as safety concerns, such as collision risk,




Figure 3.11. Example of a gridded model for road network design (Thomas Sanchez)

resulting from urban transportation investments (see Abduljabbar et al. 2019; Cai et al. 2019; Iyer 2021).

- **Social equity.** There are significant ethical issues with ML. Numerous studies have discovered that ML algorithms can give worse prediction results for underprivileged groups like low-income populations, women, and racial minorities, which can result in discriminatory treatment of these groups (see Khan, Adnan, and Iqbal 2022; Newell 2022; Zheng, Wang, and Zhao 2021).

As evidenced, transportation is a multifaceted area highly interconnected with other planning topics and concerns. These examples highlight the many ways planners could use ML and NN methods to perform assessments and make predictions to benefit planning processes and improve transportation options and operations in communities.

Land Use

Land use is another area of planning that has used NN and ML approaches for many years, as very large amounts of land-use data can be collected for analysis and modeling purposes.

Land use and land-use change are typically studied through earth observations, remote sensing, and field measurements, all of which allow for monitoring and quantifying changes in land-use patterns at local, regional, and global levels (MohanRajan, Loganathan, and Manoharan 2020). Besides earth observation and remote sensing techniques, which require advanced satellite equipment and expert knowledge for data interpretation, researchers are now using land-use change modeling approaches. The United States

Geological Survey (USGS) is an excellent source for this data. In addition to raw imagery data, USGS also provides processed data, such as its [Land Use/Land Cover dataset](#).

Land-use change models provide representations and strategies that can help analyze, understand, and assist in the planning and management of land and natural resources. Many such models are based on inductive approaches, which start with studying the observations and then developing explanations. These models are often based on a suite of explanatory variables that contribute to the change process. However, the main factors of land-use change are directly related to human interactions and decision-making processes that are often difficult to accurately model and predict.

Figure 3.12 shows an example of a land-use change model in which each grid cell on a map contains information about the input layer variables of slope, density, elevation, and accessibility, as well as the type of land use that occurred between 2003 and 2015 (Gharaibeh et al. 2020). These independent variable inputs (X_1 – X_5) predicted whether a change occurred or not (1, 0) or what type of change occurred (land-use type) as the dependent variable output.

Because most of this data is derived from satellite imagery, the amount of data may be limited by the resolution of the imagery or the unit of analysis for reporting purposes. The results of the NN are the weights (similar to regression coefficients) among the input variables and the resulting land-use change type. The trained model can then be used to predict the future state of each grid cell based on trends over a specified time period.




Figure 3.12. A neural network model for land-use change (Thomas Sanchez)

Urban Design

Economic, ecological, social, and mobility factors are just a few of the numerous criteria that influence urban design solutions. The need to account for this wide variety of elements that interact with one another presents a challenge for urban designers.

Because the outcome of an urban development project affects many aspects of daily life, the objective is to arrive at the best possible design. The challenge is specifying the parameters that will ultimately determine the best solution for a certain type of design project. There are many kinds of rules, and they can be based on anything from local ordinances to common knowledge to a person's appreciation of beauty.

The design of an entire city block is an example of a challenge that includes many variables. This includes the composition of open space and the interactions of buildings with the open space, the urban block itself, and the buildings of the neighboring blocks. Understanding the city's block typologies, as well as the surrounding area (including streets, squares, and any nearby railways), is essential for block design. Further, the block should fit in naturally with the surrounding urban environment, which necessitates thinking about the layout and building styles of neighboring blocks (Zheng and Yuan 2021). However, defining all of these requirements for block design is challenging.

[Generative adversarial networks](#) (GANs) allow NNs to harness their ability to learn visual information independently to generate new content. This is not too different from content creation with a language model using NLP. In this case, the "text" is images and designs that can be fit together based on common characteristics such as the physical dimensions, the types of surrounding land uses, and the visual appearance of building types. The fundamental benefit of using this method for generative design is that it provides examples or alternatives from a large collection of examples. The algorithm learns the visual aspects that determine the existing city's structure and then can suggest designs such as city block patterns emulating this structure. Figure 3.13 (p. 41) shows how OpenAI's [DALL-E](#) can create urban design scenarios using such methods. In the future, 3D models such as those in "digital twins" could be applied to design generation in the early design steps (Batty 2018).

CONCLUSION

During the early development of AI applications in urban planning, researchers sought to create powerful optimiza-

Figure 3.13. New York City's traffic-clogged Brooklyn–Queens Expressway as reimagined by DALL-E (created by Zach Katz using DALL-E)



tion tools that could help to synthesize human approaches with expert systems. For a variety of reasons, many of these innovations were never used widely in planning practice. There remains a substantial need in developing innovative approaches to the understanding and development of more sustainable and equitable urban places.

The examples discussed in this chapter illustrate a variety of new approaches that are being used by planners and developed by planning researchers. The breadth and depth of AI applications for urban planning will continue to increase with time as these applications are tested, improved, and become more widely used. Like the adoption of geographic information systems (GIS) in urban planning, the adoption of AI tools will be incremental as skill development occurs through professional training and university urban planning programs.

Also like GIS, more data will become available and more platforms will be developed to facilitate usage. Forward-looking public agencies will also focus on better integrating departments and functions of government to increase efficiency and enhance the intelligence of their operations.

Planners can use the examples in this chapter to think about how AI applications might be useful to them in their planning work, but they will likely face challenges in implementing new AI approaches in their day-to-day workflows. The next chapter outlines the main challenges to AI implementation and offers guidance on how to overcome those challenges.

CHAPTER 4

CHALLENGES TO ADOPTING AI IN PLANNING

The potential of artificial intelligence (AI) methods for planning was first recognized in the late 1950s and early 1960s, though further adoption progress was constrained by the absence of relevant data about urban places and processes. This situation is rapidly changing due to the digitalization of formerly analog transactions as well as the implementation of sensors that monitor movement patterns, land-use changes, energy usage, and other types of information relevant to planning and urban dynamics. Given the enormous growth in data collection, accelerated processing, and growing acceptance of AI techniques, now is the time for planners to learn about these innovative approaches.

The adoption of new technologies is often less disruptive than might be anticipated because it is usually a gradual process that occurs incrementally. Consider geographic information systems (GIS), possibly the most significant technological advancement for planners of the past few decades. The use of GIS has progressed from the creation and updating of maps by specialized GIS technicians to the relatively sophisticated spatial analysis now available to anyone with access to a smartphone and GIS applications. This has taken place gradually over many years and has positively transformed the practice of planning with minimal negative disruptive impacts. The adoption of AI by planners could follow a similar path.

All new technologies have their challenges and potential barriers to adoption. But by proactively identifying those challenges and finding solutions to potential negative disruptive effects, planners can help usher in the benefits that new technology can bring to planning practice, the planning profession, and their communities.

As discussed in this report, AI is a set of many tools, not a single platform or application. It is already integrated into many types of software, including word processing, email, spreadsheets, and GIS. While there will likely be new stand-alone AI applications or platforms developed for planners, most will be layered onto existing software or easy-to-access web-based platforms. However, as with statistical methods used for making predictions, planners will need to know the underlying assumptions and be able to communicate the reliability of results. In some cases, new approaches to planning problems may be very different from those used

in the past. These methods will involve new training and educational efforts.

This chapter addresses the challenges associated with the adoption of AI technology by planners. Seven challenges, in particular, will impact how quickly and efficiently urban planners will implement AI:

- Fear and uncertainty
- The need for new skills
- Changing data needs and considerations
- Unclear goals
- Incorporating transparency and explainability
- Encountering bias
- Addressing ethical issues that may arise with new methods and data

The chapter explores each challenge and provides suggestions to help planners and planning organizations understand each of these.

FEAR AND UNCERTAINTY

Fear and uncertainty are inevitable during the transition to new technologies along the road to digital transformation. Though fear can be an effective decision-making motivator, it can also deter us from making big changes or from taking actions that might be advantageous in the long run. Anxiety can be a warning indication of a potentially poor choice, signaling that the suggested course of action has a sizable

risk. If a planning department is considering adopting AI technologies, it is therefore important to make sure that staff are fully informed about AI methods and involved in assessing perceived risks to adoption. At the same time, planners and administrators should be prepared to address apprehension on the part of the public as stakeholders in the planning process (Yigitcanlar et al. 2022).

How can we tell when our anxieties are well-founded or whether they are being fueled by unrelated factors? The answer to this question lies in analyzing and understanding the source of the fear. Adoption anxiety is generally shaped by three related types of fear: fear of making the decision, fear of the unknown, and fear of failure (Hindle 2021; Rogers, Singhal, and Quinlan 2014).

When faced with important decisions, our fear of the consequences of making the wrong choice can cause us to take no action or minimize the advantages of innovation. Making such decisions can be particularly challenging when we are forced to step outside of our comfort zones. Choosing from a list of consultants and solutions can be daunting. Understanding an organization's features, the dangers of doing nothing, and the risks of various potential solutions is crucial as we move through the decision-making process.

It is critical to clearly outline the project's goals and the benefits of applying AI in that situation and to evaluate the risk of not achieving those goals. This enables an organization or individual to weigh prospective technology risks against the initial prospects or rewards that led them to pursue the project. After that, cooperation with knowledgeable staff or consultants to methodically address and reduce any dangers that the technology may introduce can help to take advantage of the opportunity. This includes deconstructing the broader decision into a process that is analytical, fact driven, and focuses on the opportunity balanced with the risk.

When something is unfamiliar to us, we may focus on the dangers involved with the change rather than the hazards associated with our inaction. To solve this, we should fully understand the innovation or technology under consideration. It is important to comprehend the big picture in addition to providing technical answers to specific concerns. The challenging process of digital transformation necessitates the support of administrators, managers, and staff, as well as considerable IT and operator retraining. Therefore, it is crucial to consider not just the technological component but also the operational and human components. Even when we have a high-level understanding of how the solution is supposed to function, embracing new technology can be demanding. Fortunately, competent technology providers are

aware of this and should offer thorough support for communication, transition, collaboration, and training.

The fear that a new technology will not succeed may be further divided into three categories: concerns that (1) the technology will not be established, (2) the technology will not work for a particular use case, or (3) the transition to or adoption of the technology will not proceed smoothly. As described in the prior scenario, it is again critical to quantitatively assess the technology in light of the advantages it can bring to a planning organization or project. It is also vital to understand the implementation goals as they relate to organizational value and determine why the change is needed for the optimum strategy. Innovation should be assessed based on the quantifiable rewards as well as the softer benefits.

Also, when we lack the necessary expertise, adopting new technology can require further consideration. Finding out where each component is in the adoption or maturity cycle, both within and outside of the planning application or sector, can help to allay worries about the reliability and effectiveness of a possible solution. Testimonials and live demonstrations can also help. Speaking with others who have been through this process is also beneficial to learn what has and has not worked for them. There is a reason we might be reluctant to be early adopters, but thankfully most significant automation technologies have already seen their first, second, and third adopters.

For all stakeholders who will be impacted, effective communication is essential for the diffusion of innovations. Managers should put themselves in the position of their staff and consider how the adoption of new technology will affect their day-to-day activities, both positively and negatively. Leaders should also identify any supporters and detractors of the new technology. Internal champions are essential for promoting new technologies, and they are frequently those who are closest to the day-to-day activities of individual planners or staff. The change will go more smoothly if the main advantages and justification for new technology are communicated, especially to the doubters. Buy-in does not happen when technology is pursued simply for technology's sake.

LACK OF SKILLS

The digitalization of urban planning is requiring planners to learn new skills. But AI is not the only driver of this change. Today's workplace is in a constant state of change with new technologies, changes in workplace culture, and evolving business practices. For many professions, the COVID-19

pandemic required the rapid adoption of video conferencing technologies and protocols to accommodate remote work, and electronic alternatives quickly replaced in-person contact. By necessity, we managed to learn new technologies quickly, some of which replaced long-standing previous practices. Updating skill sets to keep up with such changes is becoming a given.

The adoption of AI will differ among planners depending on their roles and responsibilities. It may include learning new AI-related terminology, software packages, programming languages, advanced statistical methods, and other AI-related concepts as discussed in this report. Upskilling to learn and maintain these skills will be an important element of adopting these new methods (Andrews et al. 2022).

In addition to learning some of these AI-specific skills, planners will increasingly need to be conversant in other urban technology applications, such as those being used to create smart cities. Technology investments in infrastructure, the Internet of Things (IoT), and other technologies will transform how urban places function and therefore will be factors in how cities are planned and managed. To stay relevant in a changing world, planners will need to know about these technologies and be prepared to engage with other disciplines involved in shaping the urban landscape, such as civil engineers, architects, landscape architects, and public administrators.

Though most planners will not be directly involved with developing or working with new AI applications, they may be engaging with technologists, such as computer scientists and application developers, to help them do so. As new urban technologies are being implemented, planners should have seats at the table to represent the interests of community stakeholders as well as promote connections to better understand urban futures.

The readiness of planners to adapt and best use these new technologies is a crucial factor at this point. There are at least three factors involved in this: (1) the ability of AI now and in the future to have a significant impact on urban social and constructed environments, (2) the ability of AI now and in the future to meet planners' needs, and (3) the readiness and ability of planners to use AI technologies. The first of these is supported by research, and the second will emerge as we see planners embracing new tools. Last, but not least, whether planners employ AI in their work will be greatly influenced by their existing degree of awareness and expertise in the field.

The successful implementation of these new methodologies will require new expertise in data analytic techniques

and information systems, which has immediate implications for planning practice. Both internal and external elements, such as the knowledge and skill sets of new employees (i.e., young planners), will play a role in the adoption process for planning organizations. Upskilling might take place through training exercises for experienced planners, perhaps led by newly hired employees who bring these abilities into an organization.

The adoption of new digital technology also entails investments in new computer equipment and systems, as well as human resources. It is important to evaluate the potential benefit, which will also depend on how many current procedures will be disrupted or altered.

UNCLEAR GOALS

When an organization makes investments in employee training, new data, and perhaps new hardware to implement AI technologies, it must also practice effective performance management. As previously mentioned, the primary purpose of adopting AI is to increase efficiency and analytical productivity. New approaches that involve investments of organizational resources (i.e., time and money) should be carefully planned for and therefore have an expected return on investment.

Fairfax County, Virginia, offers an excellent example of [IT planning and governance](#) (Fairfax County 2022). The county has a “planning and land use system modernization” initiative, which “supports County plans to advance economic development and competitiveness, enhance business processes, provide better customer service, and achieve increased reliability in plan review, approval, permitting, and inspections” (Fairfax County 2022, 18). In this case, the County has strategically linked the performance of system enhancements to overall competitiveness.

One of the main performance concerns that prevents effective management is a lack of clarity regarding what has to be accomplished. Planners have extensive knowledge and experience with this subject through their experiences in the comprehensive planning process. An ambiguous vision and unclear objectives lead to projects with no expectations and leadership for which it is impossible to create performance measures. Without specified objectives, it is impossible to define what constitutes effective performance.

It is recommended to start with a distinct vision for the use of AI (as discussed in Chapter 5), as well as organizational goals to optimize its use by staff. Together with long-

and short-term objectives, the vision and goals should be communicated to staff clearly and frequently. Coordinated oversight and evaluation of AI use within the organization is also important to ensure goals are being met, aims are being achieved, and budget resources are appropriate. Technology should never be implemented simply for technology's sake, and AI is no exception to this principle.

DATA NEEDS

Urban planning has a tradition of being data-hungry for the many types of analysis that are used in the plan-making process (Boeing et al. 2021) and data is critical for effective AI implementation. Three primary factors are data volume, quality, and management.

One of the biggest obstacles to AI implementation is the lack of sufficient high-quality data (Roh, Heo, and Whang 2019). Algorithms can often require thousands or millions of observations. A model risks poor performance when additional data is added if it is initially trained on inadequate amounts of data. Planning relies on information about people and households, as well as their behaviors, which usually comes from a combination of federal, state, and local data. In some cases, the data will need to be anonymized, which is standard practice for data collected by the U.S. Census. This includes approaches such as “data swapping,” which transfers data points from one observation to another, increasing security but potentially decreasing the value of the data for analysis to draw conclusions about real-world situations (U.S. Census Bureau 2021).

Potentially using more finely grained disaggregated data for AI models means running the risk of infringing on personal privacy rules that will constrain the availability of data. In the United States, the General Data Protection Regulation (GDPR), Payment Card Industry Data Security Standard (PCI DSS), Health Insurance Portability and Accountability Act (HIPAA), Federal Information Security Management Act of 2002 (FISMA), Family Educational Rights and Privacy Act (FERPA), and Gramm-Leach-Bliley Act are just a few of the laws that forbid the use of confidential or sensitive information (i.e., personally identifiable information, or PII) as input for machine learning (ML) models (Moallem 2018). While urban planners have not often used credit card, health care, or educational data in the past, the increasing ubiquity of data collection and its commercialization will require ethical decisions about the appropriateness of certain data that may reveal personal identities.

Data quality and integrity can be a challenge as well. Any data that an ML model reads will be used to inform its assumptions. If the data is inaccurate or incomplete, the rules produced will be fundamentally flawed. Training data must accurately reflect the population by including data sets that represent every demographic category. According to several studies, dark-skinned females are mistakenly identified by facial recognition 40 percent more frequently than white males (U.S. GAO 2020). Nonrepresentative training data sets are mostly to blame for this. Even if the data is inclusive of all demographic groups, it may still suffer from being incomplete or out of date.

Compounding this challenge is that many planners are not effectively trained in the basic elements of data management (Batty 2021). Without effective IT management practices, planning organizations can find it difficult to answer basic questions, such as how many databases an organization has, which database contains a particular piece of information, or how data was originally collected. Within organizations, siloed functional groups and poor communication create challenges for sharing data resources with coworkers, particularly for policy makers and administrators. Given that AI techniques rely heavily on data, these are serious issues. Organizations that lack the tools to understand and manage their data will find it more difficult to benefit from AI.

Public agencies struggle to deploy and manage AI systems because their staff lack fundamental data management skills (Sarker, Wu, and Hossin 2018). Data governance and management practices should extend beyond IT departments to planning staff so they can better use and protect data assets. Though the learning curve for data administration is manageable, it is far more difficult to acquire the necessary abilities to create AI solutions. Planning organizations may need to prioritize the recruitment of planners with AI training and expertise.

TRANSPARENCY AND EXPLAINABILITY

Transparency in the public realm can be a challenge, and it has several facets. For AI, the main priority for planners should be being open about the use of AI for making predictions, recommendations, or decisions (Gaur and Sahoo 2022). Transparency also involves enabling users to understand the elements of an AI system. The ability to deliver clear, relevant information regarding the outputs of an AI system and the justification for its use is another example of

transparency, as is the facilitation of open, multistakeholder conversations and the creation of specialized organizations, where required, to promote public understanding and acceptance of AI systems.

Because AI and ML models may be too technically challenging to be practical or helpful for understanding an output, transparency generally does not include disclosing source code or sharing proprietary datasets. Source code and datasets could also be considered intellectual property, which has its own types of legal protections.

Enabling those who will be affected by the output from an AI-based decision to understand how it was reached is referred to as *explainability*. This requires giving stakeholders simple-to-understand information that will allow those who are negatively impacted to contest outcomes, and, when possible, the causes and logic that led to those outcomes. For some AI systems, requiring explainability may negatively impact accuracy, performance, privacy, and security, as it may necessitate condensing the solution variables to a set small enough for humans to understand, which may not be optimal in complex, high-dimensional problems. This may not likely be the case for many planning-related analyses, however.

The main factors in a decision; the data, logic, or algorithm behind the specific outcome (Wright 2019); or an explanation of why similar-looking circumstances generated different outcomes may all be included when AI actors explain an outcome in clear and simple terms, as appropriate to the context. If applicable, personal data protection standards should be respected while allowing people to understand and contest the conclusion.

BIAS

Human biases have been extensively studied, from field experiments that show how much these biases may influence results to other tests that show the prejudices we do not realize that we have. The extent to which these human biases can infiltrate AI systems and cause detrimental outcomes has become a topic of discussion in society over the last few years (Roselli, Matthews, and Talagala 2019). These hazards are important considerations for organizations considering the adoption of AI systems.

AI can assist in identifying and minimizing the effects of human biases, but it also has the potential to exacerbate the issue by systematically introducing biases into sensitive application domains. Algorithms are susceptible to bias in several ways, even when sensitive factors like gender,

ethnicity, or sexual orientation are accounted for (Roselli, Matthews, and Talagala 2019). Even large amounts of accumulated training data can be incomplete or inaccurate, reflecting previous poor decisions or biased analyses resulting from historical conditions.

We are all susceptible to and responsible for combating bias. Not only does bias harm those discriminated against, but it also harms everyone else by limiting people's participation in the economy and society. As a result of fostering distrust and delivering skewed results, bias lowers AI's potential for use in government, business, and society in general.

As complex as deep learning models can be, however, it should be noted that the human brain is the ultimate "black box." Humans can exaggerate or not even be aware of the reasons why they made a particular decision or selected a specific alternative. It may be nearly impossible to detect the source or type of bias in a human brain, while the logic and data used by computers can be carefully analyzed.

Solutions to combating bias involve expanding academic study regarding AI prejudice in urban planning and public policy, and carefully considering the many ways that AI may supplement our current methods of making decisions. We can improve AI decision-making by using models that are meticulously constructed to avoid past biases. It is also important to immediately resolve any instances we witness of bias in AI. There may be no easy fixes in these cases, however. Defining and assessing "fairness" is exceedingly difficult. Researchers have devised technical definitions of fairness, such as mandating that models have similar outcome values across socioeconomic groups (Corbett-Davies and Goel 2018). Different fairness criteria typically cannot be satisfied at the same time, which is a considerable challenge.

Researchers have also made strides in strategies that can improve how AI systems match fairness requirements and metrics, whether by preprocessing data, reviewing results, or embedding appropriate and transparent rules as part of the data training process (Corbett-Davies and Goel 2018). "Counterfactual fairness" is a method that ensures a model's conclusions would be true even if sensitive characteristics like race, gender, or sexual orientation were modified (Kusner et al. 2017). This approach can be used in complex situations in which some impacts from sensitive qualities that affect outcomes are viewed as fair, while others are viewed as unfair. The model could be used, for instance, to ensure that an applicant's race had no bearing on whether or not they were approved for a mortgage, while still allowing the lender to include race as demographic information for later reporting.

It is also necessary to determine when a system is deemed fair for use by deciding under what circumstances automated decision-making can be permitted (Araujo et al. 2018). In some cases, “*human-in-the-loop*” algorithms or AI system responses, which include human intervention or review, will be needed to maintain control or oversight, especially in unusual circumstances that may not have significant machine intelligence to draw upon. These issues call for interdisciplinary approaches from planners, engineers, designers, ethicists, and social scientists. This can be problematic in the case of planning, where political forces influence decision-making.

Some basic steps are recommended to address potential ethical concerns relating to bias. Planners should keep up to date on the ethical dimensions of rapidly evolving urban technologies, including AI, which are being continuously uncovered as more and more such applications are being implemented. In addition, planners and planning organizations can create accountable procedures that help reduce prejudice when they use AI, including the use of technical tools or operational techniques such as oversight committees or external evaluations. The sidebar on p. 49 offers additional suggestions, resources, and best practices for the responsible and ethical use of AI.

Ongoing organizational discussions about potential biases in plans or analyses can increase awareness to determine whether particular decisions were fair. We can now hold people to higher standards because we have more sophisticated techniques for checking for bias in machines. This can include comparing the outcomes of algorithm performance to human decisions and using explainability approaches, which, depending on the AI technique used and the level of detail needed, can help identify factors that influenced the model’s results (Belle and Papantonis 2021).

Also, as mentioned earlier, planners should consider how bias can be reduced by combining the efforts of people and algorithms. Human-in-the-loop systems can offer suggestions that staff can verify or select from. Humans can intervene to help interpret results and explain levels of confidence in these algorithms. Planners, engineers, architects, computer scientists, and other experts can confer in the event of contentious or difficult-to-explain algorithmic performance.

AI provides numerous potential advantages for addressing serious societal issues, such as the effects of human prejudice. But it will not be achievable until the public believes that these technologies will provide accurate and useful results. AI can assist people in overcoming bias, but only if collaboration among stakeholders seeks to address bias in AI.

OTHER ETHICAL ISSUES

AI and big data are increasingly associated with many aspects of our day-to-day lives. In the realm of urban planning and development, the data collection sensors of smart city initiatives capture human activities at different scales, collecting large amounts of data about us both with and without our knowledge (Chang 2021). Further, data is used to make decisions, such as medical diagnoses, credit reporting, and consumer recommendations, that we seem to have little control over. Scientists and scholars have expressed significant concerns and warnings about bias embedded in search engine results and algorithms (Noble 2018) and how big data and algorithms harm the poor and reinforce racism and inequality (O’Neil 2016; Eubanks 2018).

Many of the examples we hear about related to the misapplication of AI are associated with law enforcement, surveillance, mistaken identity, and hacking, and involve negative impacts on individuals, based on data or other information at the individual level. It is understandable that the public, as well as planners, might assume that similar negative outcomes may arise while planning communities and neighborhoods. Very few if any of the analyses that urban planners perform, however, use individual-level data for persons, so it is unlikely that the use of AI in planning will have the same potential to target individuals for negative outcomes. More generally, though, planners will need to ensure that we do not build AI systems that replicate detrimental decisions made in the past related to segregation and racial discrimination.

In addition, planners should be aware of potential areas of ethical concern drawn from broader experiences with AI. These general areas of concern relate to bias, privacy, and misinformation.

Bias. As detailed in the prior section, planners are already aware of issues related to data that does not adequately represent populations of concern. Data from public input or characteristics from limited numbers of groups or neighborhoods does not represent the whole of the community and needs to be interpreted accordingly. The degree of misrepresentation can increase as the size of datasets increase. More data does not necessarily translate into better analyses, so planners should use caution. Data needs for statistical validity should be balanced with ethical concerns as well as emphasizing transparency throughout the data collection and analysis process. In addition, the logic built into AI models should be carefully reviewed for biased assumptions and socioeconomic profiling drawn from historical practices.

RESOURCES FOR THE ETHICAL USE OF AI

The ethical use of artificial intelligence (AI) is a complex and multifaceted issue that touches on a wide range of concerns, from fairness and transparency to privacy and accountability. At its core, the ethical use of AI involves ensuring that these systems are designed, developed, and deployed in a way that respects the rights and dignity of individuals.

To promote ethical AI, companies, researchers, and other organizations are taking steps that include developing internal guidelines and policies for the responsible use of AI, investing in research and development to improve the fairness and explainability of AI systems, collaborating with external organizations and experts to address ethical issues related to AI, providing transparency around their AI systems and decision-making processes, and educating employees and stakeholders on the ethical use of AI. The following list offers example resources on this topic.

INDUSTRY APPROACHES

[Responsible AI Practices](#) (Google AI). Provides recommended practices and several example areas to build fairness, interpretability, privacy, and security into AI systems.

[AI Ethics](#) (IBM). Includes reports, white papers, tutorials, videos, and best practices on AI explainability, transparency, and ethics.

PUBLICATIONS

[AI Ethics](#) (Mark Coeckelbergh, 2020, MIT Press). This book summarizes the key ethical issues of AI, such as privacy concerns, accountability, and the delegation of decision-making, openness, and prejudice as it manifests itself throughout the entire data science process. It also takes into account how employment will change in an AI economy.

[A Citizen's Guide to Artificial Intelligence](#) (John Zerilli, 2021, MIT Press). The ramifications of AI for ethical, political, legal, and economic issues are covered in this guide. It includes the fundamentals of machine learning and takes into account concerns such as transparency, bias, liability, privacy, and legislation.

[“The Global Landscape of AI Ethics Guidelines”](#) (Anna Jobin, Marcello Lenca, and Effy Vayena, 2019, *Nature Machine Intelligence* 1(9): 389–99). The authors discuss the global convergence of the five ethical principles of transparency, justice and fairness, nonmaleficence, responsibility, and privacy. They emphasize how crucial it is to coordinate efforts to produce guidelines with thorough ethical analysis and effective implementation plans.

[“The Ethics of AI Ethics: An Evaluation of Guidelines”](#) (Thilo Hagendorff, 2020, *Minds and Machines* 30(1): 99–120). This article provides a thorough summary of AI ethics, which includes looking at how ethical norms and principles are used in the research, design, and use of AI systems—as well as how to make the demands of AI ethics more effective.

LINKEDIN LEARNING COURSES

[Foundations of Responsible AI](#). This course, taught by data scientist and AI ethicist Ayodele Odubela, explains the principles of responsible AI as well as the frameworks necessary to apply responsible AI techniques in AI systems.

[Tech On the Go: Ethics in AI](#). This course, also taught by Ayodele Odubela, explores important ethical issues posed by the rapid growth of AI technology and addresses responsible AI, the nature of algorithmic harm, and the relationship between AI and human rights.

[AI Accountability Essential Training](#). This conceptually oriented course taught by Barton Poulson, associate professor of psychology at Utah Valley University, digs into the social concerns, safety challenges, and ethical issues posed by AI and offers potential solutions to key concerns.

[Learning XAI: Explainable Artificial Intelligence](#). This course, taught by Aki Ohashi, director of business development for PARC, a Xerox company, explains how explainable AI (XAI) works and the value it provides to data science-related businesses and initiatives from legal and commercial perspectives.

Threats to privacy. Privacy is arguably the most legitimate ethical concern about AI. The collection of personal data has become simpler due to technologies such as security cameras, smartphones, and the internet. As noted above, however, disaggregated data with characteristics about individuals represents the most serious threat to privacy, and these types of data are not commonly used for planning purposes.

Misinformation. Another significant ethical problem is the use of AI to spread false information. Bad actors may use ML models to produce and spread on social media channels factually incorrect information about contentious urban development issues (Hollander et al. 2020). The extent to which social media played a role in propagating false information during the 2016 presidential election serves as an example of this on a much larger scale (Shu et al. 2017). The spread of misinformation will likely continue on social media platforms, however, with or without the intervention of AI technologies. Slowing the spread of misinformation on social media platforms is a complex and multifaceted challenge that requires a multipronged approach. Overall, it will require a combination of fact-checking, algorithm adjustments, education, community engagement, and legal measures (Anderson and Rainie 2017; Villasenor 2020).

concerns related to its use as new types of data and analysis emerge for AI applications.

The final chapter in this report discusses key considerations for planners and their organizations as they look to the future adoption of AI methods and applications—in particular, the need for education, training, and upskilling for these new methods to be used effectively—and offers strategic planning recommendations for the exploration and adoption of these emerging technologies for planning practice.

CONCLUSION

AI is an emerging technology that is here to stay. As such, though planners and planning organizations will face challenges in adopting this new technology, the benefits of AI may be worth the risk. As AI becomes more commonly used and useful over time, the challenges of adoption will lessen, just as they have in other cases of innovation adoption.

The promises and potential of AI are both intriguing and absorbing, but organizations must carefully prepare before investing resources in AI. It is important to prepare for change, and some interruption is to be expected. We are accustomed to our routines; therefore, getting staff to accept AI may be challenging. As described in this chapter, it is important to create a clear vision and goals for AI use and to develop a thorough understanding of what AI offers an organization. Once an organization adopts AI, it is also important to track performance and evaluate the value and advantages of AI deployment to document a comprehensive understanding of how AI benefits the organization, staff, and stakeholders. And throughout the process of adopting AI, planners should be aware of the potential for ethical

CHAPTER 5

GETTING STARTED

Urban planning has a long, but not visible, history with artificial intelligence (AI). Researchers have been developing innovative land-use and transportation models since the late 1960s with a spirit of discovery that not only changed how we analyzed urban spaces but also how we thought about them.

This is the role of any tools that we use in planning practice: they are lenses for looking at the past, present, and future. The objective is to use the best tool for the job, and the job is to create sustainable cities that meet the needs of all residents. The AI methods discussed in this report have the potential to be very powerful when used for these purposes.

As should now be clear, though it may not always be evident, we interact with AI solutions all the time in our day-to-day lives when we use search engines, social media, e-commerce sites, mobile devices, and other online platforms. As described in Chapter 2, AI is hard at work in the background, making it easier and more effective to search for and retrieve the best available information. With appropriate AI tools, planners should be using the best possible information to make evidence-based decisions using systems that can learn over time and improve performance.

The following excerpt from PAS Report 599, *Smart Cities: Integrating Technology, Community, and Nature*, sums up the situation we face with AI and urban planning:

The world around us is changing and the planning profession needs to evolve with these changes. This may mean adding new processes, tools, and planning competencies as laid out in this report; it may also mean a reinvention of what planners do and what their roles are supposed to be. While the goal of planning remains the creation of great communities for all, the path to get there can be improved and made more effective and inclusive by using state-of-the-art technology. Improved communications between stakeholders, enhanced service quality, and empowerment of all community members to experiment and co-create can be powerful outcomes of smart

cities that planners should emphasize in their work. (Hurtado, Rouse, and Hitchings 2021)

At this stage, a crucial consideration is whether planners are willing to adopt new technologies and make the greatest use of them. As discussed in Chapter 4, there are at least three aspects to this problem:

- The present and future capacities of AI to significantly influence the social and built environments of urban areas
- The present and future capacities of AI to cater to the requirements of urban planners
- The willingness and capacity of urban planners to embrace and make use of technological advancements

There appears to be support for the first of these in the academic literature, and the second of these will emerge over time as we witness planners adopting new technology. The current level of AI awareness and expertise held by planners will have a substantial bearing on the question of whether the profession will make use of these technologies.

Successful implementation of these new methodologies will require new abilities and knowledge regarding data analytic methods and information systems, which has direct implications for planning practice. The process of adoption incorporates both internal elements, such as attitudes and decisions made by planning organizations, as well as external factors, such as the abilities that new employees (i.e., young planners) will bring to planning organizations. This type of “upskilling” can take place either through training activities for practicing planners or through recruits who already possess these talents and introduce them to their organizations. In addition, these necessitate investments

AI RESOURCES

Planners can consult a vast and ever-growing array of books, reports, websites, training courses, and plans and policies to learn more about AI and its potential applications to planning practice. The following resources represent a starting point.

[Artificial Intelligence in Urban Planning and Design: Technologies, Implementation, and Impacts.](#) Imdat As, Prithwish Basu, and Pratap Talwar, editors. 2022, Elsevier. This book provides a review of several AI-based solutions and explores emerging applications of AI technologies in urban design and urban planning. It provides a framework for talking about the theoretical underpinnings of AI, applications of AI in urban design, and systems for generative design based on AI.

[“AI in Planning: Opportunities and Challenges and How to Prepare.”](#) Clint Andrews, Keith Cooke, Alexandra Gomez, Petra Hurtado, Thomas W. Sanchez, Sagar Shah, and Norman Wright. 2022, American Planning Association. This report from APA’s “AI in Planning” Foresight Community summarizes its findings and offers preliminary suggestions on how planners can get ready for AI and its potential effects, how planners can guarantee that AI-based planning tools are used in an equitable and inclusive manner, and what the planner’s role should be in creating and using AI-based planning tools.

[AI in Urban Planning: 3 Ways It Will Strengthen How We Plan for the Future.](#) Huang Zhongwen. 2021, Urban Redevelopment Authority of Singapore. The Urban Re-

development Authority of Singapore (URA) is a leader in the application of AI to urban planning. This website summarizes a few of its activities in the areas of planning tools and innovation and highlights some of URA’s projects.

[Government Measures for a Municipal Algorithms and Data Strategy for an Ethical Promotion of Artificial Intelligence.](#) 2021, Barcelona City Council. This municipal report provides a useful framework for cities anticipating the application of AI. It identifies particular areas of concern, such as legal and ethical questions concerning their constituencies.

[City of Philadelphia IT Strategic Plan.](#) 2019, City of Philadelphia. Philadelphia’s IT strategic plan highlights existing and upcoming projects and tactics intended to integrate and capitalize on technology, innovation, engagement, access, and efficiency. Additionally, it discusses how the city’s IT office is adapting to new ideas and responsibilities in local government.

[NYC Artificial Intelligence Strategy.](#) 2021, City of New York, Mayor’s Office of the Chief Technology Officer. New York’s AI strategy takes into account the city’s AI ecosystem. It defines a healthy AI ecosystem as one in which a diverse range of people and organizations across sectors and communities work cooperatively to promote well-being, fairness, and opportunity for all. This approach aims to take into account the full range of stakeholders working in, training for, building, buying, using, and being impacted by AI.

in new forms of computer technology as well as human resources. The potential benefit will also involve the amount of disruption or change in the practices that are now being used, and this aspect needs to be evaluated.

The discussion of AI tools in this report is meant to be an introduction. The topics included have been simplified so that the reader who is not well-versed in AI or related analytical methods will still be able to grasp the general themes involved. Readers should consult additional sources, such as the references in this report and the sidebar above, the virtually infinite number of websites on these topics, and the

many online courses available, to gain more familiarity. As mentioned in the report, the language of AI is an important place to start so that planners play a role in the design of new technologies as opportunities emerge; the [glossary](#) provided as an appendix to this report offers readers a helpful resource.

CONNECTING AI WITH PLANNING PRACTICE

The intent of this PAS Report is to prepare planners for exploring AI’s potential to be the best tool for a wide range

of planning jobs that can help create vibrant, sustainable communities. Chapter 2 of this report gives planners a foundation in understanding the four primary AI functions that underlie its application to real-world use, Chapter 3 offers examples of how AI is being applied to planning practice along with several research-based examples, and Chapter 4 details associated challenges that may be encountered along the way to adding AI to planning workflows. But planners may still ask: how can I apply this information to my community and integrate AI into my planning work?

A useful way to think about where and how AI might fit into everyday planning practice is to consider five strategic points of intervention connected to the work that planners undertake (Klein 2011): community visioning; plan making; standards, policies, and incentives; development work; and public investments. Planners can think about what tasks comprise the work of each of these points and from there explore how they might be able to use AI to automate any of those tasks. And at the same time, planners can consider how data that is generated through their planning work can be optimally collected and configured to power AI analyses for better data-driven and evidence-based decision-making.

- **Community visioning** is an important element of the planning process, particularly in defining and prioritizing goals. The visioning process can involve extensive data about the community and its residents to help inform scenario design as well as to provide feedback about alternatives. As mentioned earlier, natural language processing (NLP) can be used to analyze community feedback on visions and goals. Generative methods for urban design options (using tools like DALL-E) can potentially increase engagement and interactions with the public. At the same time, community conversations should involve planners introducing the benefits of AI and being transparent about its use in the visioning process.
- **Plan making**, whether for comprehensive or general plans, functional sector-specific plans, or subarea plans, involves a significant amount of information and analysis. These analyses examine a broad spectrum of interrelated systems that have both spatial and temporal dynamics. Many planning analyses rely on detecting trends and patterns to make predictions; as mentioned throughout the report, this is what AI does best, assuming the availability of sufficient data. For instance, planners can use machine learning (ML) and neural network (NN) methods for making predictions about land-use

needs that feed into the comprehensive planning process. Tracking these predictions over time can help to evaluate model performance concerning development impacts and outcomes.

- **Standards, policies, and incentives** developed by planners are related to important evidence for making decisions or recommendations about how land will be used as well as the types of services needed to support those uses. Policy analysis with AI to ensure that regulations make the most sense and result in desired outcomes is analytically very challenging, and further research is needed in this area of planning practice to develop and test appropriate methods.
- **With development work**, planners play an important role in translating visions, goals, and requirements into changes in the built environment. They may be taking the lead in designing and managing development projects, or they may be responsible for reviewing development proposals and working with developers to ensure projects meet regulations and community goals. AI methods have the potential for augmenting these rule-based activities. Examples include using chatbots to manage development application submissions and answer current planning questions or developing algorithms and decision models to evaluate development proposals (Wright 2019). Tying development proposals to the land-use and regulation performance indicators mentioned above can be a source of historical data needed by AI methods.
- **Public investments** that support development projects in the form of services and infrastructure have the potential to generate data about demand, usage, and resident satisfaction. Examples include public parks, streets, and transit, bicycle, and pedestrian facilities. As discussed in Chapter 3, natural language processing (NLP) methods can be used to identify public investment needs and priorities from public feedback (e.g., 311, tweets, and public comments).

The five points above can be used as a starting framework to help planners systematically think through planning tasks and determine where AI might be a better tool for existing processes or might enable new types of analysis previously beyond departmental scope, time, or resources. It can also help planners think about how to best collect and use data to inform AI processes for optimal planning outcomes. This analysis can be an important component of a planning agency's strategic planning and visioning process for AI integration.

STRATEGIC PLANNING AND VISIONING FOR AI

A planning organization may choose to adopt AI for a variety of reasons, such as to improve efficiency, reduce costs, or gain new or better insights not available through current analysis techniques. Preparing a strategic plan for the adoption of AI is a good way to get started and involves several steps, including conducting a thorough analysis of the organization's current state, identifying the possible areas where AI can be used, and developing a clear vision for the implementation of AI within the organization. AI expertise, either internal or external, is critical to the following process.

Step 1: Conduct a current conditions analysis. First, a planning organization should conduct a thorough analysis of its current analysis needs and practices, including its strengths, weaknesses, opportunities, and threats. This analysis can help the organization understand its current capabilities and limitations, as well as identify areas where AI can potentially be used to improve operations.

To conduct this analysis, the organization should first assess its strengths, which may include its budget resources, staff knowledge and skills, stakeholder support, potential partners, and intended work products. The organization should also identify its weaknesses, such as outdated technology, limited expertise in AI, or a lack of data to support AI implementation.

Next, the organization should consider external factors that may present opportunities for the adoption of AI. These may include potential partners at local universities, potential private-sector partners, and the availability of AI platforms and tools. The organization should also be aware of potential threats, such as a lack of stakeholder buy-in for adopting AI, and conduct research on peer experiences about unforeseen or unexpected negative outcomes of AI application use.

To gather the necessary information for this analysis, the organization may need to conduct research, consult with experts, or conduct surveys or interviews with employees and stakeholders. The organization may also need to review its internal data and documents, such as planning products, customer feedback, and performance metrics.

Step 2. Identify implementation opportunities. Next, the organization should build on the information generated in the prior step to identify potential opportunities to adopt AI. It should consider its current operations, processes, and systems to identify specific areas where AI could potentially be used to improve efficiency, reduce costs, or enhance per-

formance. For example, an organization may be able to use AI for tasks such as data collection and analysis, decision-making, or automation of repetitive tasks.

This step may involve undertaking research to understand the latest developments in AI, consulting with experts in the field, and conducting interviews or surveys with staff and stakeholders to gather their perspectives and insights. The organization should be sure to consider its own needs and goals when identifying potential areas for AI adoption.

Step 3: Establish a vision and plan. Once areas for AI adoption have been identified, the organization can then develop a vision for the implementation of AI within the organization. This may involve setting specific goals and objectives for the use of AI, as well as outlining the steps and resources necessary to achieve those goals.

The organization should also consider potential challenges and obstacles that may arise during the implementation of AI, which may include concerns about the impact of AI on jobs, ethical concerns about the use of AI, or compliance with relevant regulations. The organization should develop strategies for addressing these challenges, such as providing employees with training and support, establishing protocols for responsible AI use, or partnering with other organizations to share knowledge and expertise. The organization should also outline the steps and resources necessary to achieve its goals, such as implementing new technologies, training employees, or acquiring data.

Overall, developing a clear vision for the implementation of AI involves identifying potential areas for AI adoption, setting goals and objectives, outlining the steps and resources necessary to achieve those goals, and considering potential challenges and obstacles. The sidebar on p. 56 summarizes this step-by-step approach. By carefully planning for the implementation of AI, an organization can position itself to take advantage of the many benefits of AI while mitigating potential risks and challenges.

Ethics and AI

Throughout the process of preparing a strategic plan for the adoption of AI, the organization should also consider the ethical implications of AI and develop guidelines for the responsible use of AI. This may include establishing protocols for data privacy, transparency, and accountability.

As discussed in Chapter 4, the ethical implications of AI are complex and wide ranging, and they involve considerations about the impacts of AI on society, the economy, and individual rights. Some of the ethical implications of

AI STRATEGIC PLANNING: A SUMMARY

Planners can apply the following step-by-step approach as they begin to explore the integration of AI in everyday planning practice.

1. Conduct a thorough analysis of the organization's current state, including its strengths, weaknesses, opportunities, and threats.
2. Identify potential areas where AI can be used to improve operations or update traditional analysis approaches.
3. Research the latest developments in AI and consult with experts in the field to understand the capabilities and limitations of AI.
4. Develop clear goals and objectives for the use of AI within the organization.
5. Outline the steps and resources necessary to achieve the goals and objectives for AI adoption.
6. Consider potential challenges and obstacles that may arise during the implementation of AI and develop strategies for addressing them.
7. Establish protocols for the collection and use of data as well as for decision-making by AI systems.
8. Develop guidelines for the responsible use of AI, including principles such as transparency, accountability, fairness, and respect for privacy.
9. Consider the potential impact of AI on society and develop strategies for ensuring that the benefits of AI are distributed equitably.
10. Collaborate with other organizations and stakeholders to address the ethical implications of AI and promote responsible use.

AI include potential job losses due to automation, unequal access to the benefits of AI, and concerns about the use of AI for surveillance or other potentially nefarious purposes.

To address these ethical implications, organizations should develop guidelines for the responsible use of AI. These guidelines may include principles such as transparency, accountability, fairness, and respect for privacy. Organizations should also establish protocols for the collection and use of data, as well as for decision-making by AI systems. This may involve collaborating with other organizations and stakeholders, such as government agencies, nonprofits, and academics, to address the ethical implications of AI and promote responsible use.

Overall, the ethical implications of AI are complex and require careful consideration and attention. By developing guidelines for the responsible use of AI and considering the potential impact of AI on society, organizations can ensure that they are using AI in a way that is ethical and beneficial to all stakeholders.

NOW IS THE TIME

AI will play a significant role in the future operation and planning of cities. As more and more data is being generated all around us, urban planners will increasingly use AI to analyze and process this data to detect patterns, make predictions, and have a better understanding of urban dynamics.

But AI will not take over all aspects of planning, and neither will it replace planners. Rather, it will comprise a useful set of tools for a subset of planning tasks. The most suitable tasks for AI will be those that are repetitive, have quantifiable dimensions, and involve data that can be captured and maintained. While there may be concerns about the loss of entry-level positions that typically handle these tasks, shifting needs will be accommodated by changes in skills that young planners have already acquired.

The American Planning Association has acknowledged the accelerating pace of AI application development, particularly as it relates to the planning profession. There are both substantial opportunities and challenges presented by these technologies, with more being revealed each day. AI can help planners in their work, enhance current planning procedures, increase efficiency, and allow planners to refocus their work on the human components of planning. The task at hand is to increase our knowledge and skills with appropriate AI methods to enhance planning practice in the service of creating more livable, resilient, and sustainable communities.

APPENDIX: GLOSSARY

algorithm: a step-by-step procedure or a set of instructions to solve a particular problem or perform a specific task. It can be thought of as a recipe that consists of a sequence of steps, each of which is clearly defined, to achieve the desired outcome.

anomalous activities: any actions or events that deviate from normal or expected behavior. These unusual or unexpected patterns or activities may indicate a problem or threat.

artificial intelligence (AI): As defined by this PAS Report, the use of computers to mimic or improve upon human intelligence processes, such as reasoning and experience-based learning.

big data: extremely large and complex data sets that traditional data processing methods are unable to handle efficiently. Big data can come from various sources such as social media, IoT devices, log files, and sensor data. The data can be analyzed to uncover hidden patterns, unknown correlations, and other useful information that can be used for business and organizational decision-making. See also PAS Report 585, [Big Data and Planning](#).

bounding boxes: rectangular boxes used in computer vision and image processing to define the location and boundaries of objects within an image. They are typically represented by a set of coordinates that define the top-left and bottom-right corners of the rectangle.

classification: a supervised learning task in which an algorithm is trained to categorize or label input data into one of several predefined classes or categories based on the features of the data.

collaborative filtering: a method of making recommendations by identifying patterns and similarities in the preferences and behavior of users. It is based on the idea that people who have similar preferences or behaviors in the past will have similar preferences in the future.

color segmentation: the process of dividing an image into multiple segments or regions, each of which corresponds to a different color or range of colors. The goal is to partition an image into regions that have similar colors or color characteristics and to separate the regions that have different colors or color characteristics. This can be used for various applications such as object detection, object tracking, and image editing.

computer vision (CV): a subfield of AI and computer science that focuses on enabling computers to interpret and understand visual information, such as images or video. The aim is to encompass both visual and cognitive functions to capture humans' ability to make sense of what they see.

content-based methods: a type of recommendation system that makes recommendations based on the characteristics or features of the items being recommended. It works by analyzing the content or attributes of an item and recommending similar items based on those attributes.

convolutional neural network (CNN): a type of artificial neural network designed specifically for processing and analyzing grid-like data, such as images, by detecting patterns and features within the data. CNNs have proven to be very effective in various computer vision tasks, such as image classification, object detection, and image recognition.

data science: an interdisciplinary field that uses scientific methods, processes, algorithms, and systems to extract meaningful knowledge and insights from structured and unstructured data to make informed decisions. It is a combination of several fields, including statistics, data analysis, machine learning, and computer science, that are used to collect, process, and analyze large amounts of data.

deep learning (DL): a neural network with several layers. DL allows for systems that can carry out complicated tasks and learn a wider variety of functionalities; its main uses today include speech recognition and general pattern recognition.

edge detection: a technique used in image processing to identify and locate sharp discontinuities in an image, such as boundaries of objects or changes in texture. It is used to find edges in an image, which can be useful in a variety of applications, such as object recognition, image segmentation, and image compression.

explainability: the ability to understand and interpret the decision-making process of an AI system, particularly machine learning and deep learning models. It is about making the inner workings of these complex models more transparent and comprehensible to humans.

generative adversarial networks (GNNs): a type of AI model that can generate new data that resembles a given dataset. Since GNNs were introduced in 2014 they have become a popular area of research in AI, particularly for tasks such as image synthesis and text generation.

heuristics: rules of thumb that are used to simplify analysis processes. While they may help to assist in decision-making, they may lead to biases or inaccuracies because they are not based on empirical results.

human-in-the-loop: a term used in AI and other computational systems to describe a process in which humans are actively involved in the decision-making or problem-solving workflow of the system. The idea is to combine the strengths of both human intelligence and AI algorithms to achieve better performance and results.

knowledge engineering: a subfield of AI that deals with the acquisition, representation, and maintenance of knowledge in a computer system. It involves acquiring, representing, and reasoning with knowledge in a formal, structured way.

labeled data: data that has a predefined or assigned output or target variable. The output or target variable is also known as the label. Labeled data is also known as supervised data, annotated data, or labeled data. It is often more difficult and expensive to obtain than unlabeled data, but it is necessary for supervised learning and many other tasks. See also *unlabeled data*.

large language model (LLM): a type of machine learning model that is trained to predict the likelihood of a sequence of words. It is used to understand and generate natural language, which is the language spoken and written by humans. LLMs are typically trained on large sets of text data and can then be used for a variety of natural language processing tasks, such as text classification, machine translation, speech recognition, and text generation.

laser scanner: an instrument that uses a laser beam to scan an object or a scene and create a 3D representation of it. The scanner works by emitting a laser beam and measuring the time it takes for the beam to reflect off an object and return to the scanner. By repeating this process multiple times, the scanner can build up a 3D point cloud of the object or scene.

learning-based AI system: an approach used in AI that learns and improves performance through experience. Instead of relying on predetermined rules (i.e., rule-based systems), algorithms and data are used to adapt, generalize, and predict, allowing them to operate more efficiently. See also *rule-based AI system*.

Level 5 autonomous vehicle: a vehicle that is capable of driving itself in any situation, without any human input or intervention. Also known as fully autonomous vehicles, they do not have steering wheels, pedals, or any other manual controls for human drivers. This level of autonomy means that the vehicle can drive itself in any conditions, on any road, and in any environment. It can handle all aspects of driving, including navigating, braking, accelerating, and responding to traffic signals and other road conditions.

Lidar: a remote sensing technology that uses laser light to measure distances between objects. Lidar stands for “light detection and ranging”; it works by emitting a laser beam and measuring the time it takes for the light to bounce back to a sensor. This information is used to create a detailed 3D map of the environment. Lidar is often used in self-driving cars, robotics, and surveying.

link prediction: the task of predicting links that will form in a network in the future. It is a common problem in network analysis and graph mining, which has many applications in fields such as social network analysis, recommendation systems, and bioinformatics.

machine learning (ML): a branch of AI, based on the principle that computers can “learn,” that seeks to enable machines (i.e., computers) to detect and “memorize” patterns and then use this information as part of decision-making processes.

malware: any software specifically designed to harm or exploit computer systems. Malware (short for “malicious software”) can come in many forms and can cause a variety of problems, from stealing personal information to causing system crashes.

motor skills: the ability to control and coordinate the body’s movements, including both gross and fine motor movements. Gross motor skills involve the larger muscle groups and are used for movements such as crawling, walking, and running, while fine motor skills involve the smaller muscle groups and are used for movements such as writing, buttoning a shirt, and using utensils.

multiple regression analysis: a statistical technique used to model the relationship between one or more independent variables (also known as predictors or explanatory variables) and a dependent variable (also known as the response or outcome variable). It is an extension of simple linear regression, which is used to model the relationship between a single independent variable and a dependent variable.

natural language processing (NLP): a subfield of AI in which computer methods analyze human language, text, or verbal communication to derive meaning. NLP is used to analyze the ideas conveyed by terms, sentences, and phrases and can also correct grammar, convert recorded speech to text, and translate between languages.

natural language understanding (NLU): a subfield of AI and natural language processing (NLP) that focuses on interpreting and extracting meaning from human language in a way that is both contextually and semantically accurate.

natural language skills: the ability to understand and use language for communication. This includes the ability to comprehend spoken and written language, as well as the ability to express oneself effectively through speech and writing. Also known as “linguistic skills,” they are essential for communication and social interaction. They are developed early in life and are shaped by a combination of genetic and environmental factors.

network analysis: the process of studying the relationships and interactions between entities in a network. These entities can be anything that can be represented as nodes in a graph, such as people, organizations, web pages, or computer systems. The relationships between these entities are represented as edges in the graph.

neural network (NN) (also artificial neural network): a computerized system, based on how the brain processes information, in which a “perceptron” takes in a set of input signals and converts them into output signals, then attempts to analyze the data by stacking numerous layers, or groups of computing units, that can recognize values that occur repeatedly and interpret patterns within the data.

open-ended text format: a type of data format that allows for free-form, unstructured text input. It is characterized by the fact that there is no predefined set of possible responses or outputs, and the answer can take many different forms.

PageRank: a website’s importance or relevance determined by the number and quality of other websites that link to it. The algorithm assigns a “rank” or “score” to each website based on the number and quality of links pointing to it. The higher the score, the more important the website is considered to be.

parsing: the process of analyzing and breaking down text data into its individual components for further processing. It is a fundamental step in natural language processing (NLP) and other tasks that involve working with text data.

pattern recognition: an automated way to detect regularities or predictable sequences and orders of values within a sample of data.

predictive maintenance: a function of AI using machine learning algorithms trained on industrial sensor data to predict future events, such as the likelihood that machinery or equipment will break down, and determine maintenance scheduling by combining data, statistics, machine learning, and modeling.

proximity sensors: sensors that detect the presence of nearby objects without any physical contact. They are often used in devices such as smartphones, tablets, and laptops to detect when a user's face is close to the screen and turn off the display to conserve power. Proximity sensors can also be used in industrial automation and robotics to detect the presence of objects and obstacles.

Python: a high-level, interpreted programming language that is widely used for web development, scientific computing, data analysis, artificial intelligence, and more. It was created in 1989 by Guido van Rossum and has a simple, easy-to-learn syntax that makes it a popular choice for beginners. Python is also known for its large and supportive community, as well as its extensive library of prebuilt modules and tools.

quantum computing: a type of computing that uses quantum-mechanical phenomena, such as superposition and entanglement, to perform operations on data. Unlike classical computers, which use binary digits (bits) to represent data, quantum computers use quantum bits or qubits, which can exist in multiple states simultaneously. This allows quantum computers to perform certain types of computations much faster than classical computers.

R: a programming language and software environment for statistical computing and graphics. It is widely used among statisticians, data scientists, and researchers for data analysis and visualization. R is an open-source programming language, which means that it is free to use and its source code is available to the public.

radar: a remote sensing technology that uses radio waves to detect and locate objects. Radar, which stands for "radio detection and ranging," works by emitting a radio frequency signal and then measuring the time it takes for the signal to bounce back after it hits an object. The radar system uses this information to determine the distance, speed, and direction of the object.

recurrent neural network (RNN): a type of artificial neural network designed to handle data that comes in sequences, such as words in text or daily temperature readings. It is a way for computers to make predictions or decisions based on data that has an order or pattern over time.

reinforcement learning: a method of trial-and-error learning in which the agent learns from its experiences and improves its performance over time. The agent learns to make decisions by maximizing the expected cumulative reward over the long run.

rule-based AI system: an approach for problem-solving or decision-making in AI that employs predetermined rules or logic. It imitates human expertise by applying these rules to input data and creating output or actions based on specific conditions. See also *learning-based AI system*.

sensory perception: the process of interpreting and organizing information received through the senses, such as sight, sound, touch, taste, and smell. It involves the brain's ability to process and make sense of sensory information to understand and respond to the environment. Sensory perception plays a crucial role in our ability to navigate and interact with the world around us.

stemming: the process of reducing a word to its base or root form. The goal of stemming is to reduce words to their base form so that they can be more easily compared and processed. For example, the words "running," "runner," and "ran" are all forms of the word "run," and stemming would reduce all of these forms to the base form "run."

supervised learning: a type of machine learning in which a model is trained on a labeled dataset to predict an output based on one or more input features. The model is provided with a set of inputs and corresponding desired outputs, and it "learns" to map the inputs to the correct outputs. The goal of supervised learning is to generalize from the training data to new unseen data. See also *unsupervised learning*.

thermal imaging: a technology that detects infrared radiation and converts it into an image that can be viewed on a screen. It is used in a variety of applications, such as night vision, surveillance, and thermography.

tokenization: the process of breaking down a piece of text into smaller units called **tokens**, which are often words or phrases. Tokenization is a fundamental step in natural language processing (NLP) and other tasks that involve working with text data. Tokens are also known as “terms” or “lexical units.”

topic modeling: a machine learning technique used to identify and extract themes or topics from large collections of text documents. It involves analyzing the patterns of word occurrences and co-occurrences within the documents to identify clusters of related words that represent a particular topic. Topic modeling helps in organizing, summarizing, and understanding large amounts of unstructured textual data.

transformer: a type of neural network architecture that is a key component in many state-of-the-art natural language processing (NLP) and AI systems, including OpenAI’s GPT series. They allow models to weigh the importance of different words or tokens in a given input sequence, taking into account their relative positions and relationships.

Trojan horse, or Trojan: a type of malware that disguises itself as a legitimate program or file to trick users into downloading or installing it. Once the Trojan is installed on a computer, it can execute its code and cause harm to the computer system, such as stealing personal information, corrupting or deleting files, or taking control of the computer.

unlabeled data: data that does not have any predefined or assigned output or target variable. See also *labeled data*.

unsupervised learning: a type of machine learning in which a model is trained on an unlabeled dataset without any specific desired output. The goal of unsupervised learning is to find patterns or structure in the data, rather than making predictions. See also *supervised learning*.

virus: a type of malware that attaches itself to a program or file and then replicates itself by copying itself to other programs, files, or other computers. Once a virus infects a computer, it can execute its code and cause harm to the computer system, such as stealing personal information, corrupting or deleting files, or taking control of the computer.

worm: a type of malware that replicates and spreads itself to other computers without the user’s knowledge or consent. They are similar to viruses in that they can cause harm to a computer system, but they differ in that they do not need to attach themselves to a specific program or file to spread. Instead, they typically spread through network vulnerabilities or by exploiting security weaknesses in software.

REFERENCES

- Abduljabbar, Rusul, Hussein Dia, Sohani Liyanage, and Saeed Asadi Bagloee. 2019. [“Applications of Artificial Intelligence in Transport: An Overview.”](#) *Sustainability* 11(1): 189.
- Afzalan, Nader, and Thomas Sanchez. 2017. [“Testing the Use of Crowdsourced Information: Case Study of Bike-Share Infrastructure Planning in Cincinnati, Ohio.”](#) *Urban Planning* 2(3): 33–44.
- Afzalan, Nader, Thomas W. Sanchez, and Jennifer Evans-Cowley. 2017. [“Creating Smarter Cities: Considerations for Selecting Online Participatory Tools.”](#) *Cities* 67(July): 21–30.
- Anderson, Janna, and Lee Rainie. 2017. [“The Future of Truth and Misinformation Online.”](#) Pew Research Center, October 19.
- Andrews, Clint, Keith Cooke, Alexsandra Gomez, Petra Hurtado, Thomas W. Sanchez, Sagar Shah, and Norman Wright. 2022. [“AI in Planning: Opportunities and Challenges and How to Prepare.”](#) Chicago: American Planning Association.
- Antunes, Miguel Eiras, Jean Gil Barroca, and Daniela Guerreiro de Oliveira. 2021. [“Urban Future With a Purpose: 12 Trends Shaping the Future of Cities by 2030.”](#) Deloitte.
- Aqib, Muhammed, Rashid Mehmood, Ahmed Alzahrani, and Iyad Katib. 2020. [“A Smart Disaster Management System for Future Cities Using Deep Learning, GPUs, and In-Memory Computing.”](#) Pages 159–84 in *Smart Infrastructure and Applications*. Springer, Cham.
- Araujo, Theo, Natali Helberger, Sanne Kruijkemeier, and Claes H. de Vreese. 2020. [“In AI We Trust? Perceptions About Automated Decision-Making by Artificial Intelligence.”](#) *AI & Society* 35(3): 611–23.
- Barreda-Luna, Antonio A., Juvenal Rodríguez-Reséndiz, Omar Rodríguez-Abreo, and José Manuel Álvarez-Alvarado. 2022. [“Spatial Models and Neural Network for Identifying Sustainable Transportation Projects with Study Case in Querétaro, an Intermediate Mexican City.”](#) *Sustainability* 14(13): 7796.
- Batty, Michael. 2018. [“Digital Twins.”](#) *Environment and Planning B: Urban Analytics and City Science* 45(5): 817–20.
- . 2021. [“The Digital Transformation of Planning.”](#) *Environment and Planning B: Urban Analytics and City Science* 48(4): 593–97.
- . 2022. [“The Emergence And Evolution of Urban AI.”](#) *AI & Society*, July.
- Belle, Vaishak, and Ioannis Papantonis. 2021. [“Principles and Practice of Explainable Machine Learning.”](#) *Frontiers in Big Data* 39.
- Boeing, Geoff, Michael Batty, Shan Jiang, and Lisa Schweitzer. 2021. [“Urban Analytics: History, Trajectory, and Critique.”](#) arXiv Preprint, arXiv:2105.07020.
- Bostrom, Nick. 1998. [“How Long Before Superintelligence?”](#) *International Journal of Futures Studies* 2.
- . 2014. *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press.
- Brail, Richard K. 1987. *Microcomputers in Urban Planning and Management*. New Brunswick: Rutgers University Center for Urban Studies.
- Brandtzaeg, Petter Bae, and Asbjørn Følstad. 2017. [“Why People Use Chatbots.”](#) Pages 377–92 in *International Conference on Internet Science*. Springer, Cham.
- Buchanan, Bruce G. 2005. [“A \(Very\) Brief History of Artificial Intelligence.”](#) *AI Magazine* 26(4): 53.
- Cai, Qing, Mohamed Abdel-Aty, Yangyang Sun, Jaeyoung Lee, and Jinghui Yuan. 2019. [“Applying a Deep Learning Approach for Transportation Safety Planning by Using High-Resolution Transportation and Land Use Data.”](#) *Transportation Research Part A: Policy and Practice* 127: 71–85.
- Cellan-Jones, Rory. 2014. [“Stephen Hawking Warns Artificial Intelligence Could End Mankind.”](#) *BBC News*, December 2.
- Chang, Victor. 2021. [“An Ethical Framework for Big Data and Smart Cities.”](#) *Technological Forecasting and Social Change* 165(April): 120559.
- Chen, Chen, Alireza Mostafizi, Haizhong Wang, Dan Cox, and Lori Cramer. 2022. [“Evacuation Behaviors in Tsunami Drills.”](#) *Natural Hazards* 112(1): 845–71.
- Corbett-Davies, Sam, and Sharad Goel. 2018. [“The Measure and Mismeasure of Fairness: A Critical Review of Fair Machine Learning.”](#) arXiv Preprint, arXiv:1808.00023.
- Crawford, Kate. 2021. *The Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. Yale University Press.
- Daniel, Claire, and Christopher Pettit. 2021. [“Digital Disruption and Planning—Use of Data and Digital Technology by Professional Planners, and Perceptions of Change to Planning Work.”](#) *Australian Planner* 57(1): 50–64.

- Dimitrakopoulos, George J., Lorna Uden, and Iraklis Varlamis. 2020. *The Future of Intelligent Transport Systems*. Elsevier.
- Dimock, Wai Chee. 2022. "What AI Can Do for Climate Change, and What Climate Change Can Do for AI." *Scientific American*, April 5.
- Dyson, George B. 2012. *Darwin Among the Machines: The Evolution of Global Intelligence*. Basic Books.
- Eubanks, Virginia. 2018. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. St. Martin's Press.
- European Commission. 2018. *Reform of EU Data Protection Rules*.
- Everett, Michael, Yu Fan Chen, and Jonathan P. How. 2018. "Motion Planning Among Dynamic, Decision-Making Agents With Deep Reinforcement Learning." Pages 3052–9 in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE.
- Fairfax County, Virginia. 2022. *FY2023 Information Technology Plan*.
- Fang, Zhou, Ying Jin, and Tianren Yang. 2022. "Incorporating Planning Intelligence Into Deep Learning: A Planning Support Tool for Street Network Design." *Journal of Urban Technology* 29(2): 99–114.
- Frem, Joe, Vineet Rajadhyaksha, and Jonathan Woetzel. 2018. "Thriving Amid Turbulence: Imagining the Cities of the Future." McKinsey & Company, October 11.
- Gaur, Loveleen, and Biswa Mohan Sahoo. 2022. "Introduction to Explainable AI and Intelligent Transportation." Chapter 1 in *Explainable Artificial Intelligence for Intelligent Transportation Systems*. Springer, Cham.
- Gharaibeh, Anne, Abdulrazzaq Shaamala, Rasha Obeidat, and Salman Al-Kofahi. 2020. "Improving Land-Use Change Modeling by Integrating ANN With Cellular Automata-Markov Chain Model." *Heliyon* 6(9): e05092.
- Goertzel, Ben. 2014. "Artificial General Intelligence: Concept, State of the Art, and Future Prospects." *Journal of Artificial General Intelligence* 5(1): 1–48.
- Gomez, Alessandra, and Joseph DeAngelis. 2022. *Digitalization and Implications for Planning*. Chicago: American Planning Association.
- Hindle, Kim. 2021. "The Time is Now! The Adoption of New Technology is the Way of Tomorrow." *Carego Blog*, July 14.
- Hollander, Justin B., Ruth Potts, Maxwell Hartt, and Minyu Situ. 2020. "The Role of Artificial Intelligence in Community Planning." *International Journal of Community Well-Being* 3(4): 507–21.
- Hosseini, Ali, Ahmad Pourahmad, Ali Taeab, Milad Amini, and Sara Behvandi, S. 2017. "Renewal Strategies and Neighborhood Participation on Urban Blight." *International Journal of Sustainable Built Environment* 6(1): 113–21.
- Howe, Jeff. 2006. "The Rise of Crowdsourcing." *Wired*, June 1.
- Hurtado, Petra, with Benjamin G. Hitchings and David C. Rouse. 2021. *Smart Cities: Integrating Technology, Community, and Nature*. PAS Report 599. Chicago: American Planning Association.
- Hurtado, Petra, Sagar Shah, Joseph DeAngelis, and Alessandra Gomez. 2022. *2022 Trend Report for Planners*. Chicago: American Planning Association.
- . 2023. *2023 Trend Report for Planners*. Chicago: American Planning Association.
- Ibrahim, Mohamed R., James Haworth, and Tao Cheng. 2020. "Understanding Cities With Machine Eyes: A Review of Deep Computer Vision in Urban Analytics." *Cities* 96(January): 102481.
- Iyer, Lakshmi Shankar. 2021. "AI Enabled Applications Towards Intelligent Transportation." *Transportation Engineering* 5(September): 100083.
- Jargowsky, Paul A., and Mohamed El Komi. 2011. "Before or After the Bell? School Context and Neighborhood Effects on Student Achievement." Chapter 4 in *Neighborhood and Life Chances: How Place Matters in Modern America*, edited by Harriet B. Newburger, Eugenie L. Birch, and Susan M. Wachter. Philadelphia: University of Pennsylvania Press.
- Kaku, Michio. 2015. *The Future of the Mind: The Scientific Quest to Understand, Enhance, and Empower the Mind*. Anchor Books.
- Kelly, Kevin. 2017. "The Myth of a Superhuman AI." *Wired*, April 25.
- Khan, Sikander, Adnan Adnan, and Naveed Iqbal. 2022. "Applications of Artificial Intelligence in Transportation." Conference paper, 2022 International Conference on Electrical, Computer and Energy Technologies (ICECET). July.
- Klein, William. 2011. "The Five Strategic Points of Planning Intervention." *PAS QuickNotes* 31. Chicago: American Planning Association.
- Kurzweil, Ray. 2013. *How to Create a Mind: The Secret of Human Thought Revealed*. Penguin.
- Kusner, Matt J., Joshua R. Loftus, Chris Russell, and Ricardo Silva. 2017. "Counterfactual Fairness." *Advances in Neural Information Processing Systems* 30.
- Massey, Douglas S., and Nancy A. Denton. 1993. *American Apartheid: Segregation and the Making of the Underclass*. Harvard University Press.
- Mayor, Adrienne. 2018. *Gods and Robots: Myths, Machines, and Ancient Dreams of Technology*. Princeton, N.J.: Princeton University Press.
- McCarthy, J., M.L. Minsky, N. Rochester, and C.E. Shannon. 1955. "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence." August 31.
- Mitchell, Melanie. 2019. *Artificial Intelligence: A Guide for Thinking Humans*. Penguin UK.
- Moallem, Abbas, editor. 2018. *Human-Computer Interaction and Cybersecurity Handbook*. CRC Press.
- Mobley, R. Keith. 2002. *An Introduction to Predictive Maintenance*. Elsevier.

- MohanRajan, Sam Navin, Agilandeewari Loganathan, and Prabukumar Manoharan. 2020. [“Survey on Land Use/Land Cover \(LU/LC\) Change Analysis in Remote Sensing and GIS Environment: Techniques and Challenges.”](#) *Environmental Science and Pollution Research* 27(24): 29900–26.
- National Association of State Chief Information Officers (NASCIO). 2020. [“Chat with Us: How States are Using Chatbots to Respond to the Demands of COVID-19.”](#)
- National League of Cities. 2020. [“Using AI and Machine Learning to Proactively Address Urban Blight.”](#)
- . 2022. [Future of Cities.](#)
- Newell, Mark. 2022. [“The Social Justice of Intelligent Transportation Systems: Deep Learning-based Autonomous Driving Technologies, Cooperative Navigation Algorithms, and Vehicle and Pedestrian Detection Tools.”](#) *Contemporary Readings in Law and Social Justice* 14(2): 9–26.
- Newman, Kathe, and Elvin K. Wyly. 2006. [“The Right to Stay Put, Revisited: Gentrification and Resistance to Displacement in New York City.”](#) *Urban Studies* 43(1): 23–57.
- Noble, Safiya Umoja. 2018. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York University Press.
- Obermeier, Klaus K. 1989. *Natural Language Processing Technologies in Artificial Intelligence: The Science and Industry Perspective*. Prentice-Hall, Inc.
- O’Neil, Cathy. 2016. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. Crown.
- Organisation for Economic Co-operation and Development (OECD). 2020. [“Using Artificial Intelligence to Help Combat COVID-19.”](#) Updated April 23.
- Ortiz, Ben, Laura Kahn, Marc Bosch, Philip Bogden, Viveca Pavon-Harr, Onur Savas, and Ian McCulloh. 2020. [“Improving Community Resiliency and Emergency Response With Artificial Intelligence.”](#) arXiv Preprint, arXiv:2005.14212.
- Owens, Ann. 2017. [“The Consequences of Income Segregation Between School Districts for Economic and Racial Achievement Gaps.”](#) Conference paper, Population Association of America 2017 Annual Meeting, April 28.
- Pignataro, Anthony. 2021. [“Long Beach Using AI Software to Monitor How Residents Feel About COVID-19 Policies.”](#) *Long Beach Post News*, May 18.
- Plass, Lian. 2023. “Additional Prompt Ideas.” In “A Planning Chat With ChatGPT,” PowerPoint presentation by Ann Dilleuth, Tom Sanchez, Lian Plass, and Jennifer Fierman. HUB session, APA National Planning Conference, April 1.
- Rahnama, Hossein, and Alex “Sandy” Pentland. 2022. [“The New Rules of Data Privacy.”](#) *Harvard Business Review*, February 5.
- Rainie, Lee, Cary Funk, Monica Anderson, and Alec Tyson. 2022. [AI and Human Enhancement: Americans’ Openness Is Tempered by a Range of Concerns.](#) Pew Research Center, March.
- Rich, Elaine, Kevin Knight, and Shivashankar B. Nair. 2009. *Artificial Intelligence*. Tata McGraw Hill.
- Rogers, Everett M., Arvind Singhal, and Margaret M. Quinlan. 2014. [“Diffusion of Innovations.”](#) Pages 182–86 in *An Integrated Approach to Communication Theory and Research*, edited by Don W. Stacks, Michael B. Salwen, and Kristen C. Eichorn. 3rd Edition. Routledge.
- Roh, Yuji, Geon Heo, and Steven Euijong Whang. 2019. [“A Survey on Data Collection for Machine Learning: A Big Data-AI Integration Perspective.”](#) *IEEE Transactions on Knowledge and Data Engineering* 33(4): 1328–47.
- Roselli, Drew, Jeanna Matthews, and Nisha Talagala. 2019. [“Managing Bias in AI.”](#) In *Companion Proceedings of The 2019 World Wide Web Conference*, May.
- Russell, Ben. 2017. *Robots: The 500-Year Quest to Make Machines Human*. Scala Arts Publishers, Inc.
- Russell, Stuart J., and Peter Norvig. 2016. *Artificial Intelligence: A Modern Approach*. Malaysia: Pearson Education Limited.
- Sanchez, Thomas W., and Marc Brenman. 2013. [“Public Participation, Social Equity, and Technology in Urban Governance.”](#) Chapter 3 in *Citizen E-Participation in Urban Governance: Crowdsourcing and Collaborative Creativity*, edited by Carlos Nunes Silva. IGI Global.
- Sanchez, Thomas W., Hannah Shumway, Trey Gordner, and Theo Lim. 2022. [“The Prospects of Artificial Intelligence in Urban Planning.”](#) *International Journal of Urban Sciences* 1–16.
- Sarker, Iqbal H. 2021. [“Deep Learning: A Comprehensive Overview on Techniques, Taxonomy, Applications and Research Directions.”](#) *SN Computer Science* 2(6): 1–20.
- Sarker, Md Nazirul Islam, Min Wu, and Md Altab Hossin. 2018. [“Smart Governance Through Bigdata: Digital Transformation of Public Agencies.”](#) In *2018 International Conference on Artificial Intelligence And Big Data (ICAIBD) Proceedings*, May. Institute of Electrical and Electronics Engineers.
- Schilling, Joseph, and Jimena Pinzón. 2016. [“The Basics of Blight.”](#) Vacant Property Research Network Research & Policy Brief no. 2.
- Shu, Kai, Amy Sliva, Suhang Wang, Jiliang Tang, and Huan Liu. 2017. [“Fake News Detection on Social Media: A Data Mining Perspective.”](#) *ACM SIGKDD Explorations Newsletter* 19(1): 22–36.

- Silva, Carlos N. 2010. "[The E-Planning Paradigm—Theory, Methods And Tools: An Overview](#)." Chapter 1 in *Handbook of Research on E-Planning: ICTs For Urban Development And Monitoring*, edited by Carlos Nunes Silva. IGI Global.
- Swingler, Kevin. 1996. *Applying Neural Networks: A Practical Guide*. Morgan Kaufmann.
- Tate, W., Shaw, S., and Taylor, D. 1997. "The Historical Roots of GIS in Urban Planning." *Journal of Housing Research* 8(2): 203–25.
- Tu, Sunan, and Ming Zhang. 2022. "[The Big Data Model for Urban Road Land Use Planning Is Based on a Neural Network Algorithm](#)." *Computational Intelligence & Neuroscience* 2022: 2727512.
- Tuscaloosa (Alabama), City of. 2018. "[West Tuscaloosa Community Inventory: Analysis of Blight Conditions](#)."
- UN Habitat. 2022. *World Cities Report 2022: Envisaging the Future of Cities*. Nairobi, Kenya: United Nations Human Settlements Programme.
- U.S. Census Bureau. 2021. *Disclosure Avoidance for the 2020 Census: An Introduction*. U.S. Government Publishing Office, Washington, D.C.
- U.S. Government Accountability Office. 2020. *Facial Recognition Technology: Privacy and Accuracy Issues Related to Commercial Uses*. GAO-20-522.
- Villasenor, John. 2020. "[How to Deal With AI-Enabled Disinformation](#)." The Brookings Institution, November 23.
- Wasserman, David. 2020. "[The Art of Learning by Example](#)." *Planning*, October.
- World Economic Forum. 2022. "[A Blueprint for Equity and Inclusion in Artificial Intelligence](#)." White Paper, June.
- Wright, Norman. 2019. "[Applying Algorithms to Land-Use Decision Making](#)." *Zoning Practice*, March.
- Yigitcanlar, Tan, Duzgun Agdas, and Kenan Degirmenci. 2022. "[Artificial Intelligence in Local Governments: Perceptions of City Managers on Prospects, Constraints and Choices](#)." *AI & Society* (2022): 1–16.
- Yigitcanlar, Tan, Rita Yi Man Li, Tommi Inkinen, and Alexander Paz. 2022. "[Public Perceptions on Application Areas and Adoption Challenges of AI in Urban Services](#)." *Emerging Sciences Journal* 6(6): 1199–1236.
- Zheng, Hao, and Philip F. Yuan. 2021. "[A Generative Architectural and Urban Design Method Through Artificial Neural Networks](#)." *Building and Environment* 205(November): 108178.
- Zheng, Helen W., Geoffrey Q.P. Shen, Yan Song, Bingxia Sun, and Jingke Hong. 2017. "[Neighborhood Sustainability in Urban Renewal: An Assessment Framework](#)." *Environment and Planning B: Urban Analytics and City Science* 44(5): 903–24.
- Zheng, Yunhan, Shenhao Wang, and Jinhua Zhao. 2021. "[Equality of Opportunity in Travel Behavior Prediction With Deep Neural Networks and Discrete Choice Models](#)." *Transportation Research Part C: Emerging Technologies* 132(November): 103410.

ACKNOWLEDGMENTS


The author would like to thank the following individuals who provided input or feedback on the development of this report: Theo Lim, PHD, Virginia Tech; Julio Carillo, AICP, LEED AP, Parkhill; Brendan Moore, AICP, and Jayla Blanke, City of Tuscaloosa, Alabama; Ryan Kurtzman and Lea Eriksen, City of Long Beach, California; Matthew Ensor, FranklyAI; John Duffy, Beca; Xinyue Ye, PHD, Texas A&M University; Soheil Sabri, PHD, University of Melbourne; Max Sokol, AICP, Replica; David Wasserman, AICP, Alta Planning + Design; Sean Koh and Zhongwen Huang, Urban Redevelopment Authority, Singapore; and Eyal Feder-Levy, Zencity.

Thanks also to the following individuals who reviewed the manuscript and provided valuable feedback: Petra Hurtado, PHD, American Planning Association; Claire Daniel, University of New South Wales; Norman Wright, AICP; and Tan Yigitcanlar, PHD, Queensland University of Technology.

Special thanks to editor Ann Dillemath, AICP, for her excellent management and guidance during the preparation of this report.

DIVE DEEPER

with these related PAS resources.



Artificial Intelligence and Planning Practice

PAS Memo 111

AI represents an emerging toolbox for planners that enables a range of new capabilities. Read this article for an overview of AI concepts, their potential uses for practice, and their roles in larger trends connected to digital governance and civic data in planning.



Big Data and Planning

PAS Report 585

Data, analytics, and smart technologies empower planners and decision makers to better understand the present and predict the future. Read this report to learn how developments in data science and technology-based opportunities can advance your community.



ChatGPT: Implications for Planning

PAS QuickNotes 101

The release of ChatGPT by OpenAI in late 2022 brought AI writing capabilities into the mainstream overnight. Read this briefing paper to learn about the implications of these emerging generative AI applications for planning practice.



Smart Cities: Integrating Technology, Community, and Nature

PAS Report 599

If deployed in the right ways, state-of-the-art technologies can help resolve the myriad challenges communities big and small face today. Read this report to learn about the new skills, processes, and tools planners need to plan both with and for smart cities.

APA members and PAS subscribers get full access to the PAS digital toolbox. Learn more at planning.org/pas.



