ARTIFICIAL INTELLIGENCE FOR CLIMATE CHANGE MITIGATION ROADMAP (SECOND EDITION)

CHAPTER 1:

INTRODUCTION TO AI

Alp Kucukelbir

November 2024

Innovation for Cool Earth Forum

CHAPTER 1:

INTRODUCTION TO AI

Alp Kucukelbir

Artificial intelligence (AI) is part of our everyday lives. Email providers use AI to filter spam. Postal services apply AI to route hand-written envelopes. Technology companies leverage AI to identify faces in photographs, while radiologists reach for AI to interpret medical scans. Economists use AI to forecast elections, and digital retailers turn to AI to optimize prices.[1,2]

The release of ChatGPT in November 2022 generated extraordinary public attention to AI. ChatGPT quickly became the most rapidly adopted product in human history, with more than 100 million users by January 2023. Its operator claims that 200 million people are now using ChatGPT on a weekly basis.[3] This increased attention has led to questions about how AI could help address major global challenges, including climate change—the topic of this report.

A. What Is AI?

AI is the science of making computers perform complex tasks typically associated with human intelligence. Modern AI relies on a branch of computer science called machine learning (ML). ML refers to a set of algorithms that detects patterns from large and sometimes messy data without explicit programming (i.e., without a human-crafted description of each pattern). This is a task often associated with human learning—for example, learning to walk, speak or identify objects.

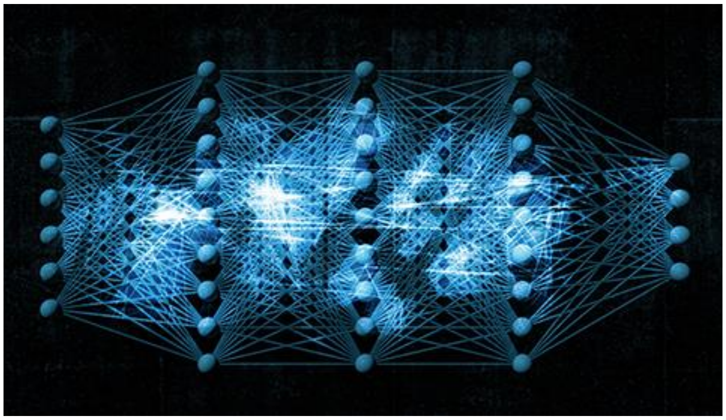

Figure 1-1. A visualization of a deep neural network, a type of AI model that powers popular AI systems such as ChatGPT.

How does AI differ from traditional computation? Consider a computer program that plays chess. The traditional approach to building such an algorithm involves explicitly programming the rules of chess, encoding basic principles of good game play, and specifying a method to search over all possible moves to pick the best one. Even in a game as seemingly simple as chess, this is an enormous task for a computer—the number of chess positions is about the same as the number of atoms on Earth.[4] (For the curious, that is about 1,000,000,000,000,000,000,000,000,000,000,000, 000,000,000,000,000,000 positions.) No existing computer, even the most powerful supercomputing clusters, can efficiently play chess this way.

Now consider an AI approach to playing chess. The core idea is to replace human input on what constitutes good strategy with a system that only uses the rules of the game to play against itself to find good strategies. Leveraging clever mathematics that significantly reduce the need to search over all possible moves, an AI system can efficiently simulate games against itself millions of times. This repeated simulation enables the AI system to "learn" the principles of good play, in a way that exceeds the ability of human programmers to explicitly encode them in software. This approach to AI uses branches of ML known as deep neural networks (see Figure 1-1) and reinforcement learning, which are ideally suited to problems where simulation plays a prominent role.[5] Table 1-1 summarizes the key difference between AI and traditional computation.

Supervised and unsupervised ML are two other ways to build AI systems—both rely on historical data to "learn" patterns.

- Supervised learning requires historical data with labels or explicit targets. One common example includes handwritten digit recognition—used by many postal services around the world—which pairs many thousands of scanned pictures of written digits with their corresponding number to "train" the AI system.
- Unsupervised learning only requires historical data, without any corresponding labels. The AI system is trained to search for patterns and associations hidden in the data. This form of AI is commonly used in recommendation engines, which can suggest movies you might like based on movies you have previously watched and historical patterns of the likes and dislikes of other people watching similar movies.

Table 1-1. *AI differs from traditional software in its requirements and its outputs.*

	TRADITIONAL SOFTWARE	ARTIFICIAL INTELLIGENCE (AI)
Requirements	• No historical data needed • Explicit programming of domain knowledge • No "training" needed (everything is explicitly programmed)	• Historical data or simulator • Implicit programming of expectations of patterns from data • Need to "train" the AI algorithm to extract patterns
Outputs	• Deterministic results • Can efficiently solve simpler problems	• Statistical results: can sometimes make mistakes • Can offer solutions to more complex problems

B. What Can AI Do?

Modern AI systems have far-reaching capabilities in at least four areas.

Detection. AI can detect patterns and anomalies in vast and complex data sets. This capability enables AI to perform tasks such as detecting faces in images and pinpointing greenhouse gas (GHG) leaks from satellite data. Monitoring combines continuous detection with alerting capabilities. In this context, AI facilitates continuous detection of unusual patterns or anomalies within data sets, which is different from traditional monitoring methods that involve periodic checks and human intervention. Classic examples of monitoring include tracking financial transactions for signs of fraud and surveying gas extraction asset data to detect methane leaks, both of which benefit from AI-powered detection and monitoring.

Prediction. AI systems can learn from historical patterns to make predictions and forecasts about how a system might behave in the future. This capability enables AI to perform tasks such as guessing what movies you might like to watch and forecasting complex weather patterns for the upcoming week. Forecasting typically implies a prediction over time (almost always in the future). But the ability to make predictions is a fundamental part of AI systems, one that enables the capabilities below.

Optimization. AI systems can leverage their predictions to optimize systems and recommend actions that achieve specific goals. For instance, AI can identify the minimum amount of fertilizer needed for a particular crop by predicting its effect on production yield. Similarly, AI can optimize steel production by predicting how different recipes will impact its final strength properties. The output of AI-based optimization are action recommendations, which are typically implemented by human experts. (See the example below on AI's potential.)

Simulation. AI systems can create complex simulations and scenario plans, allowing organizations to test hypotheses in situations where running real-world experiments are not practical. AI-powered simulations can sift through millions of new material candidates, helping identify promising candidates for empirical validation. Scenario planning identifies future "what-if" situations, assesses risks and provides actionable insights for strategic decision-making. This can enable energy providers to plan for supply and demand scenarios that their grids may have never experienced before, minimizing operational costs and risks.

Many AI systems offer capabilities that fall into more than one category above.

C. How Does AI Work?

With AI, there is no longer a need to explicitly program every detail of how to solve a problem. Instead, we rely on data, a model and computation.

Data. To replace explicit programming, supervised and unsupervised AI methods require historical data—observations and measurements that pertain to the problem at hand. In postal routing, these are images of handwritten letters and digits mapped to their correct digital representations. In facial recognition, these are many photographs of the same individual, labeled with their name. Access to high-quality data is essential for AI training. More data directly improves the odds of finding useful patterns—up to a point, after which more data provide diminishing benefits. (In reinforcement learning, data sets are typically simulated.)

Model. AI methods require implicit programming of the types of patterns that lie hidden in data. This part of an AI program is called the "model"—a mathematical description of pattern types expected in data. For example, if a sequence of chess moves appears frequently in winning games, the model should pick this up as a successful strategy. If some people write the letter "t" with a straight line and others with a curve at the bottom, the model should identify both as valid forms of a "t." The scientific community has been steadily developing increasingly sophisticated models over the past several decades.

Computation. Models by themselves are useless—they provide nonsense answers—until they are "trained" on data. Collectively, the various statistical approaches to achieving this goal and the hardware that enables such algorithms fall under the term "computation"—a set of mathematical methods to use a model to find and evaluate the quality of patterns ("training"), while simulating multiple scenarios. In chess, this involves making thousands of clever hypothetical moves to evaluate a particular strategy. In postal routing, this involves quantifying the uncertainty in differentiating a "3" from an "8" to recognize such digits reliably. Computation integrates the idea that AI programs do not contain explicitly programmed rules; rather, computation is the mechanism by which AI unravels and leverages implicit patterns from data (Figure 1-2).

AI has been steadily improving since its inception in the early days of computing. A combination of better access to rich data sources, better models for complex applications and better computing technology (software and hardware) for simulation has led to AI's proliferation.

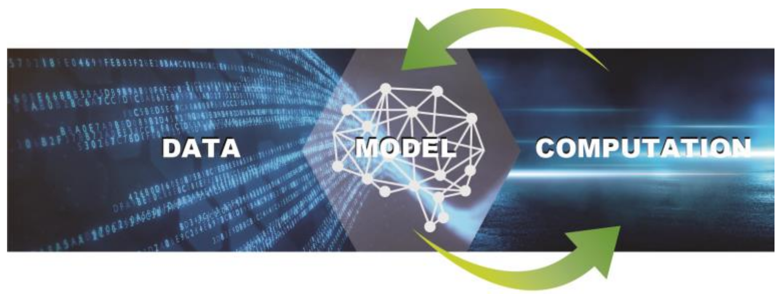

Figure 1-2. *AI systems work by using a model to identify patterns in data. Models by themselves are not useful and must be "trained" on data through computation. Computation integrates the idea that AI systems do not contain explicit information, rather computation is the mechanism by which AI unravels and leverages implicit patterns from data.*

D. What Is AI's Potential?

While chess contrasts AI to traditional software, it does not fully capture AI's potential; a chess program is effectively playing a game. To dive deeper into a practical discipline that is evolving with AI,[6] we turn to radiology—a branch of medicine in which specialist doctors use medical imaging (data) to diagnose and treat diseases.

Radiologists are experts at pattern recognition. After years of training, these doctors spend much of their time detecting anatomical and physiological deviations from blurry and noisy medical scans—which are themselves proxies for tissue and biology, not the real thing itself. AI can provide an important boost to performing this task.

In cancer medicine, for instance, medical imaging data sets with expert-verified labeling of the location and type of tumors are increasingly available. Armed with these data sets, AI systems can be trained to detect patterns in the medical images that expert humans have labeled as a tumor. Once trained, these systems can be directed to examine new medical images, searching for similar patterns in the data that would imply the existence of a tumor.

Once a tumor has been identified, an AI system can begin to simulate various treatment scenarios. How big would the tumor be after one session of radiotherapy? How about after the second? What if the parameters of the radiotherapy are slightly different? Do we end up with a better outcome? These are the types of questions radiologists can explore using AI to assist them in designing a treatment plan, which they execute using tested traditional software that operates medical equipment. The AI outputs a series of outcome probabilities, which themselves recommend treatment actions.

AI technologies not only help radiologists in their practice but also help push the scientific boundaries of their field. AI is enabling radiologists to process and search for patterns across huge databases, paving the way toward personalized treatments. This movement is so significant, it has its own name: radiomics.[7]

The rise of AI in radiology has neither usurped traditional software nor displaced its practitioners. But it highlights a particular type of AI success story. When AI is combined with traditional software and human domain experts, the results are stronger than what AI can produce alone. Keeping “humans in the loop” is key to using AI to solve many real-world problems (Figure 1-3).

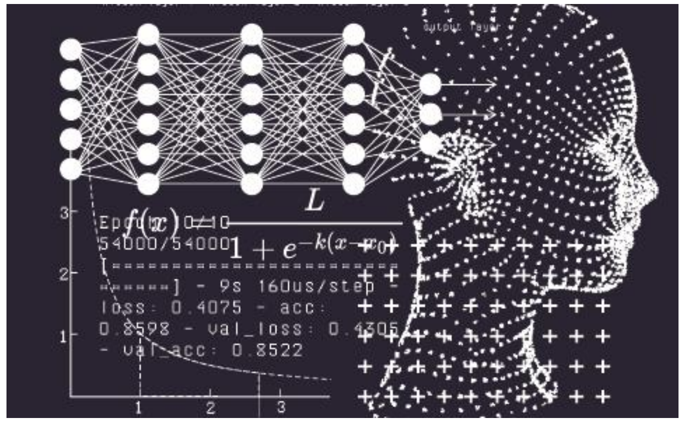

Figure 1-3. *Keeping "humans in the loop" is essential to using AI to solve real-world problems.*

Box 1-1

LARGE LANGUAGE MODELS (LLMS) AND THE FUTURE OF AI

Large language models (LLMs), such as ChatGPT, are one type of AI system. LLMs analyze vast amounts of text data and can string together responses to queries by predicting the most likely next word in a sentence. The user interface is similar to conversing with a human, expanding the potential user base for such technology to anyone who can type a question into a mobile phone or computer.

The success of these systems has revived questions around the future capabilities of AI. ML and AI experts are divided on the transformational potential of LLMs and the best balance between rapid innovation and caution. Chapter 11 of this Roadmap discusses LLMs in greater detail.

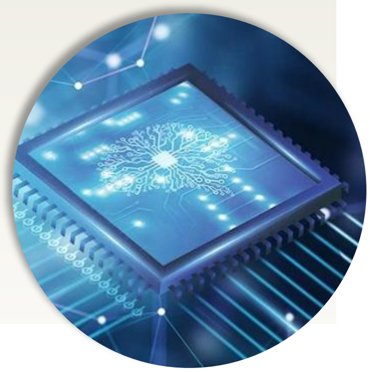

E. How Much Energy Does AI Need?

In the past year the energy needs of generative AI models, such as ChatGPT, have received considerable media attention.[8-10] But not all AI systems require as much computing power as generative AI. Some types of AI, such as simple statistical models, neural networks and reinforcement learning, require much less energy. The amount of energy an AI system needs is dictated by its model type and how frequently it is trained and used.

Most AI models require relatively modest energy inputs, even with large data sets. However recently popular generative AI models, such as LLMs and image/video diffusion models, are far more energy-intensive than other AI systems. This is because they require substantial energy both to train and to use.

Training AI Models. In general, training is the most energy-intensive part of building an AI system. Yet for most AI systems, energy demands are not enormous. Some AI systems that analyze medical data, forecast manufacturing sensor outputs and process agricultural drone imagery can be trained on a laptop, often in a matter of minutes. In generative AI systems, however, the type of model (e.g., LLMs) and scale of data (e.g., billions of web pages) can require enormous amounts of computation. Training can become a weeks-long energy-intensive task, executed on supercomputers housed in data centers.

Using AI Models. Once an AI system has been trained, it becomes ready for use—detecting patterns, predicting the future, optimizing systems and simulating “what-if” scenarios. In most AI systems, this is reasonably cheap to do. However, generative AI systems have introduced a new dynamic: they are energy-intensive both to train *and* to use. Enormous (and fiercely secretive) training costs have effectively priced out all but the largest technology firms from developing popular systems, like LLMs. But using such models is also very expensive, with ChatGPT rumored to cost OpenAI $700,000 per day.[11]

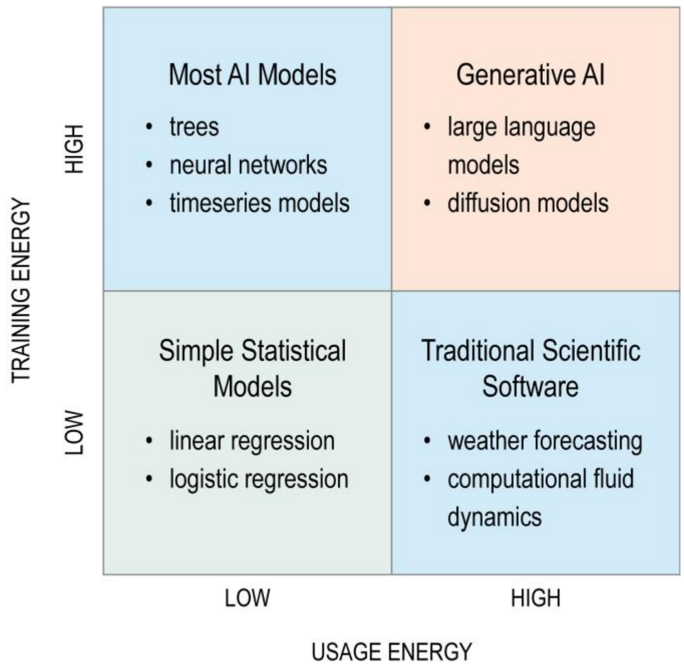

Figure 1-4. *AI systems differ in how much energy they require to train and to use.*

F. What Kind Of Data Does AI Need?

Unlike traditional software, AI requires access to historical data. These data can come in many different forms and be hosted by different types of entities. The availability and accessibility of these data are both important considerations for their potential role in AI systems.

i. Data types

AI systems can work with many different data types.

- Tabular data. Measurements that follow a generic row and column structure. Often associated with spreadsheet applications, tabular data can represent multiple measurements (rows) of a set of things (columns). Common across many applications.
- Time-series data. Measurements that have a time ordering and can be plotted over time. While small time-series data sets can also be considered tabular, they are often stored in database software that can handle large volumes of data. Common in signal processing (audio, remote sensing), finance and econometrics.
- Geospatial and raster data. Measurements that have a spatial ordering and can often be viewed as images. This kind of data no longer looks tabular; they are often stored as files or in special databases. Common in satellite imaging and climate science.
- Network data. Measurements that come with a graph of nodes and edges. This kind of data is often stored in special graph databases. Common in power systems and social networks.
- Text and sequential data. Measurements that comprise sequences of symbols, such as words. This kind of data is typically stored as text files but can also be encoded in databases. Common in language applications.

Box 1-2

HOW MUCH DATA DOES AI NEED?

The answer to this important question depends on the “resolution” of the problem AI is solving. In chess, the number of moves in each game in a data set has no effect—the “resolution” of the task is at the game-level. The more games, the better.

In time-series tasks, if a common event is being studied, a few days of data may be sufficient. But for rare events, years if not decades of historical measurements will be needed. In general, data size is not a useful metric—the amount of data to drive successful AI applications can range from megabytes[12] to terabytes.[13]

ii. Data hosts and owners

Data that can be used for AI applications may be hosted by different organizations and entities. Public sector data hosts include government agencies, state-owned enterprises, public universities, national research laboratories and multilateral institutions. Private sector data hosts include for-profit companies, not-for-profit organizations (e.g., private universities, think tanks, private research laboratories) and individuals. For both public sector and private sector organizations, data can have varying degrees of availability and accessibility.

iii. Data availability

The term "data" loosely refers to some amount of measured information. But for AI applications, the way in which data are measured and digitized matters (Figure 1-5).

- Measured and well-digitized. Properly designed and deployed instrumentation will provide high-quality data that can power AI applications. Such data typically exhibit a high degree of spatial and temporal resolution, covering relevant areas in sufficient precision over an appropriate number of experiments and amount of time. Examples include industrial production data, high-fidelity weather data and fine-resolution satellite data.
- Measured but poorly digitized. Data where instrumentation is either insufficient or improperly configured may not be able to drive successful AI applications. These cases can occur in underfunded application areas (biodiversity studies), rapidly changing application areas (agriculture) and broader geographies (weather data in developing nations). For example, digitizing the monthly total energy usage at the building-level is not sufficient to drive AI-based individual household energy optimization.
- Measured but not digitized. Measurements that could support AI applications may be measured but not digitized. Digital instruments without connectivity, analog instrumentation and manual observations constitute much of this category. Examples include digital thermometers without internet connectivity, analog pressure gauges and visual observations of local weather.
- Not measured. Facts and quantities that would be required to drive an AI application may not be measured at all. In these situations, the ideal outcome is to leapfrog to measured and well-digitized data.

iv. Data accessibility

Data that are measured and (ideally well) digitized may have varying levels of accessibility (Figure 1-5).

- Open-source data. These are the most easily accessed data. Open-source data sets are often hosted on public websites or other public data services. While open-source data sets are widely accessible, they may be subject to licensing[9] agreements that limit their use. Such data may also lack the specificity required in AI applications, as they may have been anonymized to protect individual privacy or trade secrets. Examples include government databases, academic data repositories and data sets shared for data-science competitions.

- Data at cost. These are data that are governed by some sort of usage agreement at a cost dictated by their host. Such data are often high-quality and specific to application areas and may also be governed by additional licensing agreements. Examples include imaging data sold by satellite-operating corporations, curated data for self-driving vehicle development and transportation data from shipping corporations.

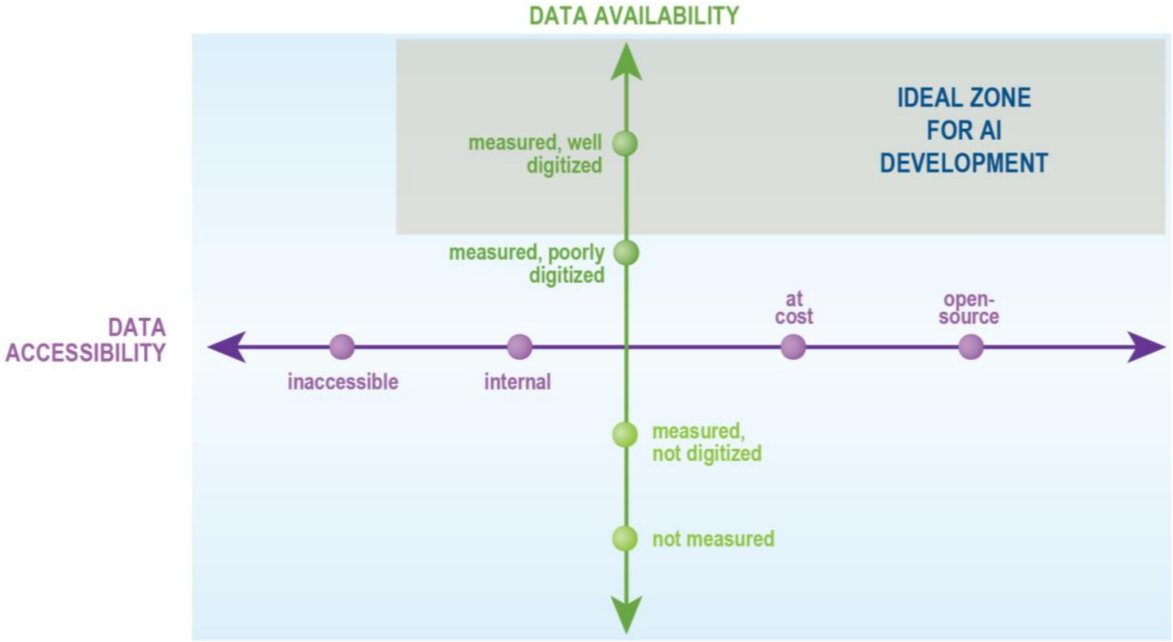

Figure 1-5. Data availability and accessibility are key aspects of enabling AI applications. The ideal zone for AI development relies on accessible, measured and well digitized data.

- Internal data. These data are kept by their hosts to be used internally. Such data are typically proprietary, containing confidential or private information. Examples include industrial production data, material-science research and development records, and GPS location data at the individual level.
- Inaccessible data. These data are generated but not stored. Such data are often temporarily created by computer programs and used in some way. Derived results may be stored, but the raw data are frequently discarded. Examples include physical system simulators and intermediate data used in the processing of other data. Inaccessible data prevents AI development.

G. Why Is AI Developing So Rapidly?

The speed and scale of recent AI development and deployment are remarkable. Improvements in computational technology and exponential reductions in cost are fueling larger and more complex AI systems.[10] The sharing of pre-trained models has also lowered costs by enabling transfer learning instead of building AI systems from scratch. These decreasing costs are enabling more widespread use of advanced AI like large language models for chatbots.

H. Readings

There is a vast literature on AI, including many books and articles introducing computation, ML and AI to non-experts. The following sources may be helpful:

1. Stuart Russell. *Human Compatible: Artificial Intelligence and the Problem of Control*. (Penguin Publishing Group, London, UK, 2020)
2. Nate Silver. *The Signal and the Noise: Why So Many Predictions Fail-but Some Don't*. (Penguin Publishing Group, London, UK, 2012)
3. Judea Pearl & Dana Mackenzie. *The Book of Why: The New Science of Cause and Effect*. (Basic Books, New York, NY, 2018)
4. Brian Christian & Tom Griffiths. *Algorithms to Live By: The Computer Science of Human Decisions*. (Henry Holt and Company, New York, NY, 2016)

The following textbooks may be helpful to those seeking additional technical depth in AI and ML:

1. Kevin P. Murphy. *Probabilistic Machine Learning: An Introduction*. (MIT Press, Cambridge, Massachusetts, 2022)
2. Moritz Hardt & Benjamin Recht. *Patterns, Predictions, and Actions: Foundations of Machine Learning*. (Princeton University Press, Princeton, New Jersey, 2022)
3. Stuart Jonathan Russell & Peter Norvig. *Artificial Intelligence: A Modern Approach*. (Pearson, London, UK, 2020)
4. Richard S. Sutton & Andrew G. Barto. *Reinforcement Learning, second edition: An Introduction*. (MIT Press, Cambridge, Massachusetts, 2018)

I. References

1 Postal Facts. *Optical character recognition;* US Postal Service (USPS), Washington, D.C., https://facts.usps.com/98-percent-of-hand-addressed-letters/ (Accessed August 2024).

2 Bibb Allen *et al.* 2020 ACR Data Science Institute Artificial Intelligence Survey. *Journal of the American College of Radiology* 18, 1153-1159 (2021). https://doi.org/10.1016/j.jacr.2021.04.002.

3 Ina Fried. *OpenAI says ChatGPT usage has doubled since last year;* Axios, Arlington, Virginia, https://www.axios.com/2024/08/29/openai-chatgpt-200-million-weekly-active-users?utm_source=newsletter&utm_medium=email&utm_campaign=newsletter_axiosam&stream=top (2024).

4 Stefan Steinerberger. On the number of positions in chess without promotion. *International Journal of Game Theory* 44, 761-767 (2015). https://doi.org/10.1007/s00182-014-0453-7.

5 David Silver *et al.* A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science* 362, 1140-1144 (2018). https://doi.org/10.1126/science.aar6404.

6 Garry Choy *et al.* Current Applications and Future Impact of Machine Learning in Radiology. *Radiology* 288, 318-328 (2018). https://doi.org/10.1148/radiol.2018171820.

7 Kaustav Bera *et al.* Predicting cancer outcomes with radiomics and artificial intelligence in radiology. *Nature Reviews Clinical Oncology* 19, 132-146 (2022). https://doi.org/10.1038/s41571-021-00560-7.

8 Brian Calvert. *AI already uses as much energy as a small country. It's only the beginning;* Vox, Washington, D.C., https://www.vox.com/climate/2024/3/28/24111721/climate-ai-tech-energy-demand-rising (2024).

9 Seth Fiegerman & Matt Day. *Why AI Is So Expensive;* Bloomberg, New York, New York, https://www.bloomberg.com/news/articles/2024-04-30/why-artificial-intelligence-is-so-expensive (2024).

10 Ariel Cohen. *AI Is Pushing The World Toward An Energy Crisis;* Forbes, Jersey City, New Jersey, https://www.forbes.com/sites/arielcohen/2024/05/23/ai-is-pushing-the-world-towards-an-energy-crisis/ (2024).

11 Aaron Mok. *ChatGPT could cost over $700,000 per day to operate. Microsoft is reportedly trying to make it cheaper.;* Business Insider, New York, New York, https://www.businessinsider.com/how-much-chatgpt-costs-openai-to-run-estimate-report-2023-4 (2023).

12 Creative Commons. *Creative Commons;* Mountain View, CA, https://creativecommons.org/ (Accessed August 2024).

13 OpenAI. *AI and compute;* San Francisco, California, https://openai.com/research/ai-and-compute (2018).