1

Why AI and Data Literacy?

Artificial Intelligence (AI) is quickly becoming infused into the very fabric of our everyday society. AI already influences decisions in employment, credit, financing, housing, healthcare, education, taxes, law enforcement, legal proceedings, travel, entertainment, digital marketing, social media, news dissemination, content distribution, pricing, and more. AI powers our GPS maps, recognizes our faces on our smartphones, enables robotic vacuums that clean our homes, powers autonomous vehicles and tractors, helps us find relevant information on the web, and makes recommendations on everything from movies, books, and songs to even who we should date!

And if that's not enough, welcome to the massive disruption caused by AI-powered chatbots like OpenAI's ChatGPT and Google's Bard. The power to apply AI capabilities to massive data sets, glean valuable insights buried in those massive data sets, and respond to user information requests with highly relevant, mostly accurate, human-like responses has caused fear, uncertainty, and doubt about people's futures like nothing we have experienced before. And remember, these AI-based tools only learn and get smarter the more that they are used.

Yes, ChatGPT has changed everything!

In response to this rapid proliferation of AI, **STEM (Science**, **Technology**, **Engineering**, and **Mathematics)** is being promoted across nearly every primary and secondary educational institution worldwide to prepare our students for the coming AI tsunami. Colleges and universities can't crank out data science and machine learning curriculums, classes, and graduates fast enough.

But AI and data literacy are more than just essential for the young. 72-year-old congressman Rep. Don Beyer (Democrat Congressman from Virginia) is pursuing a master's degree in machine learning while balancing his typical congressman workloads to be better prepared to consider the role and ramifications of AI as he writes and supports the legislation.

Thomas H. Davenport and DJ Patil declared in the October 2016 edition of The Harvard Business Review that data science is the sexiest job in the 21[st] century[1]. And then, in May 2017, The Economist anointed data as the world's most valuable resource[2].

"Data is the new oil" is the modern organization's battle cry because in the same way that oil drove economic growth in the 20[th] century, data will be the catalyst for economic growth in the 21[st] century.

But the consequences of the dense aggregation of personal data and the use of AI (neural networks and data mining, deep learning and machine learning, reinforcement learning and federated learning, and so on) could make our worst nightmares come true. Warnings are everywhere about the dangers of poorly constructed, inadequately defined AI models that could run amok over humankind.

> "*AI, by mastering certain competencies more rapidly and definitively than humans, could over time diminish human competence and the human condition itself as it turns it into data. Philosophically, intellectually — in every way — human society is unprepared for the rise of artificial intelligence.*"
>
> —*Henry Kissinger, MIT Speech, February 28, 2019*

> "*The development of full artificial intelligence (AI) could spell the end of the human race. It would take off on its own and re-design itself at an ever-increasing rate. Humans, limited by slow biological evolution, couldn't compete and would be superseded.*"
>
> —*Stephen Hawking, BBC Interview on December 2, 2014*

> "*Powerful AI systems should be developed only once we are confident that their effects will be positive and their risks will be manageable.*"
>
> —*Future of Life Institute (whose membership includes Elon Musk and Steve Wozniak) on March 29, 2023*

There is much danger in the rampant and untethered growth of AI. However, there is also much good that can be achieved through the proper and ethical use of AI. We have a once-in-a-generation opportunity to leverage AI and the growing sources of big data to power unbiased and ethical actions that can improve everybody's quality of life through improved healthcare, education, environment, employment, housing, entertainment, transportation, manufacturing, retail, energy production, law enforcement, and judicial systems. This means that we need to educate everyone on AI and data literacy. That is, we must turn everyone into **Citizens of Data Science**.

AI and data literacy can't just be for the high priesthood of data scientists, data engineers, and ML engineers. We must prepare *everyone* to become Citizens of Data Science and to understand where and how AI can transform our personal and professional lives by reinventing industries, companies, and societal practices to fuel a higher quality of living for everyone.

In this first chapter, we'll discuss the following topics:

- History of literacy
- Understanding AI
- Data + AI: Weapons of math destruction
- Importance of AI and data literacy
- What is ethics?
- Addressing AI and data literacy challenges

History of literacy

Literacy is the ability, confidence, and willingness to engage with language to acquire, construct, and communicate meaning in all aspects of daily living. The role of literacy encompasses the ability to communicate effectively and comprehend written or printed materials. Literacy also includes critical thinking skills and the ability to comprehend complex information in various contexts.

Literacy programs have proven to be instrumental throughout history in elevating the living conditions for all humans, including:

- The development of the printing press in the 15th century by Johannes Gutenberg was a pivotal literacy inflection point. The printing press enabled the mass production of books, thereby making books accessible to all people instead of just the elite.
- The establishment of public libraries in the 19th century played a significant role in increasing literacy by providing low-cost access to books and other reading materials regardless of class.

- The introduction of compulsory education in the 19th and early 20th centuries ensured that more people had access to the education necessary to become literate.
- The advent of television and radio in the mid-20th century provided new methods to improve education and promote literacy to the masses.
- The introduction of computer technology and the internet in the late 20th century democratized access to the information necessary to learn new skills through online resources and interactive learning tools.

Literacy programs have provided a wide range of individual and societal benefits too, including:

- **Improved economic opportunities**: Literacy is a critical factor in economic development. It helps people acquire the skills and knowledge needed to participate in the workforce and pursue better-paying jobs.
- **Better health outcomes**: Literacy is linked to better health outcomes, as literate people can better understand health information and make informed decisions about their health and well-being.
- **Increased civic participation**: Literacy enables people to understand and participate in civic life, including voting, community engagement, and advocacy for their rights and interests.
- **Reduced poverty**: Literacy is critical in reducing poverty, as it helps people access better-paying jobs and improve their economic prospects.
- **Increased social mobility**: Literacy can be a critical factor in upward social mobility, as it provides people with the skills and knowledge needed to pursue education, job training, and other opportunities for personal and professional advancement.

We need to take the next step in literacy, explicitly focusing on AI and data literacy to support the ability and confidence to seek out, read, understand, and intelligently discuss and debate how AI and big data impact our lives and society.

Understanding AI

Okay, before we go any further, let's establish a standard definition of AI.

An AI model is a set of algorithms that seek to optimize decisions and actions by mapping probabilistic outcomes to a utility value within a constantly changing environment...with minimal human intervention.

Let's simplify the complexity and distill it into these key takeaways:

- Seeks to optimize its decisions and actions to achieve the desired outcomes as framed by user intent
- Interacts and engages within a constantly changing operational environment
- Evaluates input variables and metrics and their relative weights to assign a utility value to specific decisions and actions (AI utility function)
- Measures and quantifies the effectiveness of those decisions and actions (measuring the predicted result versus the actual result)
- Continuously learns and updates the weights and utility values in the AI utility function based on decision effectiveness feedback

...all of that with minimal human intervention.

We will dive deeper into AI and other advanced analytic algorithms in *Chapters 3* and *4*.

Dangers and risks of AI

The book *1984* by George Orwell, published in 1949, predicts a world of surveillance and citizen manipulation in society. In the book, the world's leaders use propaganda and revisionist history to influence how people think, believe, and act. The book explores the consequences of totalitarianism, mass surveillance, repressive regimentation, propagandistic news, and the constant redefinition of acceptable societal behaviors and norms.

The world described in the book was modeled after Nazi Germany and Joseph Goebbels' efforts to weaponize propaganda – the use of biased, misleading, and false information to promote a particular cause or point of view – to persuade everyday Germans to support a maniacal leader and his criminal and immoral view for the future of Germany.

But this isn't just yesteryear's challenge. In August 2021, Apple proposed to apply AI across the vast wealth of photos captured by their iPhones to detect child pornography[3].

Apple has a strong history of protecting user privacy. Twice, by the Justice Department in 2016 and the US Attorney General in 2019, Apple was asked to create a software patch that would break the iPhone's privacy protections. And both times, Apple refused based on their beliefs that such an action would undermine the user privacy and trust that they have cultivated.

So, given Apple's strong position on user privacy, Apple's proposal to stop child pornography utilizing their users' private data was surprising. Controlling child pornography is certainly a top social priority, but at what cost? The answer is not black and white.

Several user privacy questions arise, including:

- To what extent are individuals willing to sacrifice their personal privacy to combat this abhorrent behavior?
- To what extent do we have confidence in the organization (Apple, in this case) to limit the usage of this data solely for the purpose of combating child pornography?
- To what extent can we trust that the findings of the analysis will not fall into the hands of unethical players and be exploited for nefarious purposes?
- What are the social and individual consequences linked to the AI model's *false positives* (erroneously accusing an innocent person of child pornography) and *false negatives* (failing to identify individuals involved in child pornography)?

These types of questions, choices, and decisions impact everyone. And to be adequately prepared for these types of conversations, we must educate everyone on the basics of AI and data literacy.

AI Bill of Rights

The White House **Office of Science and Technology Policy (OSTP)** released a document called *Blueprint for an AI Bill of Rights*[4] on October 4, 2022, outlining five principles and related practices to help guide the design, use, and deployment of AI technologies to safeguard the rights of the American public in light of the increasing use and implementation of AI technologies.
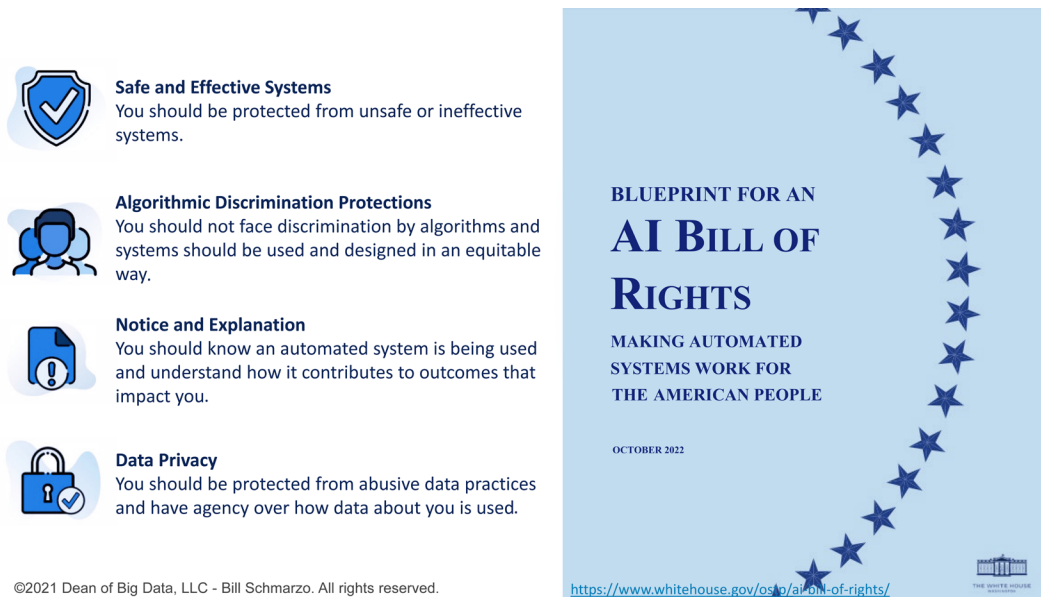
To quote the preface of the extensive paper:

> *"Among the great challenges posed to democracy today is the use of technology, data, and automated systems in ways that threaten the rights of the American public. Too often, these tools are used to limit our opportunities and prevent our access to critical resources or services. These problems are well documented. In America and around the world, systems supposed to help with patient care have proven unsafe, ineffective, or biased. Algorithms used in hiring and credit decisions have been found to reflect and reproduce existing unwanted inequities or embed new harmful bias and discrimination. Unchecked social media data collection has been used to threaten people's opportunities, undermine their privacy, or pervasively track their activity—often without their knowledge or consent."*

The paper articulates your rights when interacting with systems or applications that use AI (*Figure 1.1*):

- You should be protected from unsafe or ineffective systems.
- You should not face discrimination by algorithms; systems should be used and designed equitably.
- You should be protected from abusive data practices via built-in protections and have agency over how data about you is used.
- You should know that an automated system is being used and understand how and why it contributes to outcomes that impact you.

Safe and Effective Systems
You should be protected from unsafe or ineffective systems.

Algorithmic Discrimination Protections
You should not face discrimination by algorithms and systems should be used and designed in an equitable way.

Notice and Explanation
You should know an automated system is being used and understand how it contributes to outcomes that impact you.

Data Privacy
You should be protected from abusive data practices and have agency over how data about you is used.

BLUEPRINT FOR AN
AI BILL OF RIGHTS
MAKING AUTOMATED SYSTEMS WORK FOR THE AMERICAN PEOPLE

OCTOBER 2022

https://www.whitehouse.gov/ostp/ai-bill-of-rights/

THE WHITE HOUSE

Figure 1.1: AI Bill of Rights blueprint

The AI Bill of Rights is a good start in increasing awareness of the challenges associated with AI. But in and of itself, the AI Bill of Rights is insufficient because it does not provide a pragmatic set of actions to guide an organization's design, development, deployment, and management of their AI models. The next step should be to facilitate collaboration and adoption across representative business, education, social, and government organizations and agencies to identify the measures against which the AI Bill of Rights will be measured. This collaboration should also define the AI-related laws and regulations – with the appropriate incentives and penalties – to enforce adherence to this AI Bill of Rights.

Data + AI: Weapons of math destruction

In a world more and more driven by AI models, data scientists cannot effectively ascertain on their own the costs associated with the unintended consequences of false positives and false negatives. Mitigating unintended consequences requires collaboration across diverse stakeholders to identify the metrics against which the AI utility function will seek to optimize.

As was well covered in Cathy O'Neil's book *Weapons of Math Destruction*, the biases built into many AI models used to approve loans and mortgages, hire job applicants, and accept university admissions yield unintended consequences severely impacting individuals and society.

For example, AI has become a decisive decision-making component in the job applicant hiring process[5]. In 2018, about 67% of hiring managers and recruiters[6] used AI to pre-screen job applicants. By 2020, that percentage had increased to 88%[7]. Everyone must be concerned that AI models introduce bias, lack accountability and transparency, and aren't even guaranteed to be accurate in the hiring process. These AI-based hiring models may reject highly qualified candidates whose resumes and job experience don't match the background qualifications, behavioral characteristics, and operational assumptions of the employee performance data used to train the AI hiring models.

The good news is that this problem is solvable. The data science team can construct a feedback loop to measure the effectiveness of the AI model's predictions. That would include not only the false positives – hiring people who you thought would be successful but were not – but also false negatives – not hiring people who you thought would NOT be successful but, ultimately, they are.

We will deep dive into how data science teams can create a feedback loop to learn and adjust the AI model's effectiveness based on the AI model's false positives and false negatives in *Chapter 6*.

So far, we've presented a simplified explanation of what AI is and reviewed many of the challenges and risks associated with the design, development, deployment, and management of AI. We've discussed how the US government is trying to mandate the responsible and ethical deployment of AI through the introduction of the AI Bill of Rights. But AI usage is growing exponentially, and as citizens, we cannot rely on the government to stay abreast of these massive AI advancements. It's more critical than ever that, as citizens, we understand the role we must play in ensuring the responsible and ethical usage of AI. And that starts with AI and data literacy.

Importance of AI and data literacy

AI and data literacy refers to the holistic understanding of the data, analytic, and behavioral concepts that influence how we consume, process, and act based on how data and analytical assessments are presented to us.

Nothing seems to fuel the threats to humanity more than AI. There is a significant concern about what we already know about the challenges and risks associated with AI models. But there is an even bigger fear of the *unknown unknowns* and the potentially devasting unintended consequences of improperly defined and managed AI.

Hollywood loves to stoke our fears with stories of AI running amok over humanity (remember "Hello, Dave. You're looking well today." from the movie *2001: A Space Odyssey*?), fears that AI will evolve to become more powerful and more intelligent than humans, and humankind's dominance on this planet will cease.

Here's a fun smorgasbord of my favorite *AI-run-amok* movies that all portray a chilling view of our future with AI (and consistent with the concerns of Henry Kissinger and Stephen Hawking):

- *Eagle Eye*: An AI super brain (ARIIA) uses big data and IoT to nefariously influence humans' decisions and actions.
- *I, Robot*: Cool-looking autonomous robots continuously learn and evolve, empowered by a cloud-based AI overlord (VIKI).
- *The Terminator*: An autonomous human-killing machine stays true to its AI utility function in seeking out and killing a specific human target, no matter the unintended consequences.
- *Colossus: The Forbin Project*: An American AI supercomputer learns to collaborate with a Russian AI supercomputer to protect humans from killing themselves, much to the chagrin of humans who seem to be intent on killing themselves.
- *War Games*: The WOPR (War Operation Plan Response) AI system learns through game playing that the only smart nuclear war strategy is "not to play" (and that playing Tic-Tac-Toe is a damn boring game).
- *2001: A Space Odyssey*: The AI-powered HAL supercomputer optimizes its AI utility function to accomplish its prime directive, again, no matter the unintended consequences.

Yes, AI is a powerful tool, just like a hammer, saw, or backhoe (I guess Hollywood hasn't found a market for movies about evil backhoes running amok over the world). It is a tool that can be used for either good or evil. However, it is totally under our control whether we let AI run amok and fulfill Stephen Hawking's concern and wipe out humanity (think of *The Terminator*) or we learn to master AI and turn it into a valuable companion that can guide us in making informed decisions in an imperfect world (think of Yoda).

Maybe the biggest AI challenge is the *unknown unknowns*, those consequences or actions that we don't even think to consider when contemplating the potential unintended consequences of a poorly constructed, or intentionally nefarious, AI model. How do we avoid the potentially disastrous unintended consequences of the careless application of AI and poorly constructed laws and regulations associated with it? How do we ensure that AI isn't just for the *big shots* but is a tool that is accessible and beneficial to all humankind? How do we make sure that AI and the massive growth of big data are used to proactively *do good* (which is different from the passive *do no harm*)?

Well, that's on us. And that's the purpose of this book.

This book is about choosing... errr... umm... "good". And achieving *good* with AI starts with mastering fundamental AI and data literacy skills. However, the foundation for those fundamental AI and data literacy skills is ethics. How we design, develop, deploy, and manage AI models must be founded on the basis of delivering meaningful, responsible, and ethical outcomes. So, what does ethics entail? I've tried to answer that in the next section.

What is ethics?

Ethics is a set of moral principles governing a person's behavior or actions, the principles of *right and wrong* generally accepted by an individual or a social group. Or as my mom used to say, "Ethics is what you do when no one is watching."

Ethics refers to principles and values guiding our behavior in different contexts, such as personal relationships, work environments, and society. Ethics encompasses what society considers morally right and wrong and dictates that we act accordingly.

Concerning AI, ethics refers to the principles and values that guide the development, deployment, and management of AI-based products and services. As AI becomes more integrated into our daily lives, AI ethical considerations are becoming increasingly important to ensure that AI technologies are fair, transparent, and responsible.

We will dive deep into the topic of AI ethics in *Chapter 8*.

Addressing AI and data literacy challenges

What can one do to prepare themselves to survive and thrive in a world dominated by data and math models? Yes, the world will need more data scientists, data engineers, and ML engineers to design, construct, and manage AI models that will permeate society. But what about everyone else? We must train everyone to become Citizens of Data Science.

But what is data science? Data science seeks to identify and validate the variables and metrics that might be better predictors of performance to deliver more relevant, meaningful, and ethical outcomes.

To ensure that data science can adequately identify and validate those different variables and metrics that might be better predictors of performance, we need to ensure that we democratize the power of data science to leverage AI and big data in driving more relevant, meaningful, and ethical outcomes. We must ensure that AI and data don't create a societal divide where only the high priesthood of data scientists, data engineers, and ML engineers prosper. We must ensure that they don't benefit just those three-letter government agencies and the largest corporations.

We must prepare and empower everyone who is prepared to participate in and benefit from the business, operational, and societal benefits of AI. We must extend the AI Bill of Rights vision by championing a Citizens of Data Science mandate.

However, do you really understand your obligation to be a *Citizen* of Data Science? To quote my good friend John Morley on citizens and citizenship:

> *"Citizenship isn't something that is bestowed upon us by an external, benevolent force. Citizenship requires action. Citizenship requires stepping up. Citizenship requires individual and collective accountability – accountability to continuous awareness, learning, and adaptation. Citizenship is about having a proactive and meaningful stake in building a better world."*

Building on an understanding of the active participation requirements of citizenship, let's define the Citizens of Data Science mandate that will build our AI and data literacy journey:

Citizens of Data Science Mandate

Ensuring that everyone, of every age and every background, has access to the education necessary to flourish in an age where economic growth and personal development opportunities are driven by AI and data.

This Citizens of Data Science mandate would provide the training and pragmatic frameworks to ensure that AI and data's power and potential are accessible and available to everyone, and that the associated risks are understood so that we not only survive but thrive in a world dominated by AI and data.

The AI and Data Literacy Framework

To become a Citizen of Data Science, we need a set of guidelines – a framework – against which to guide our personal and organizational development with respect to understanding how to thrive in a world dominated by AI and data.

A *framework* is a structured approach that provides guidance, rules, principles, and concepts for solving a specific problem or accomplishing a particular task. Think of a framework as providing guard rails, versus railroad tracks, that proactively guide and benchmark our personal and professional development.

The **AI and Data Literacy Framework** provides those guard rails – and the associated guidance, rules, principles, and concepts – to ensure that everyone is aware of and educated on their role in ensuring the responsible and ethical definition, design, development, and management of AI models. It is a framework designed to be equally accessible and understandable by everyone.

We will use the AI and Data Literacy Framework shown in *Figure 1.2* throughout the book to help readers understand the different components of becoming Citizens of Data Science and ensure everyone has the skills and training to succeed in a world dominated by AI and data.

Figure 1.2: AI and Data Literacy Framework

This AI and Data Literacy Framework is comprised of six components:

1. **Data and privacy awareness** discusses how your data is captured and used to influence and manipulate your thoughts, beliefs, and subsequent decisions. This section also covers personal privacy and what governments and organizations worldwide do to protect your data from misuse and abuse.

2. **AI and analytic techniques** focuses on understanding the wide range of analytic algorithms available today and the problems they address. This chapter will explore the traditional, optimization-centric analytic algorithms and techniques so that our Citizens of Data Science understand what each does and when best to use which algorithms. And given the importance of AI, we will also have a separate chapter dedicated to learning-centric AI (*Chapter 4*). We will discuss how AI works, the importance of determining user intent, and the critical role of the AI utility function in enabling the AI model to continuously learn and adapt. This separate chapter dedicated to AI will also explore AI risks and challenges, including confirmation bias, unintended consequences, and AI model false positives and false negatives.

3. **Making informed decisions** or decision literacy explores how humans can leverage basic problem-solving skills to create simple models to avoid ingrained human decision-making traps and biases. We will provide examples of simple but effective decision models and tools to improve the odds of making more informed, less risky decisions in an imperfect world.

4. **Prediction and statistics** explains basic statistical concepts (probabilities, averages, variances, and confidence levels) that everyone should understand (if you watch sports, you should already be aware of many of these statistical concepts). We'll then examine how simple stats can create probabilities that lead to more informed, less risky decisions.

5. **Value engineering competency** provides a pragmatic framework for organizations leveraging their data with AI and advanced analytic techniques to create *value*. We will also provide tools to help identify and codify how organizations create value and the measures against which value creation effectiveness will be measured across a diverse group of stakeholders and constituents.

6. **AI ethics** is the foundation of the AI and Data Literacy Framework. This book will explore how we integrate ethics into our AI models to ensure the delivery of unbiased, responsible, and ethical outcomes. We will explore a design template for leveraging economics to codify ethics that can then be integrated into the AI utility function, which guides the performance of AI models.

This book will explore each subject area as we progress, with a bonus chapter at the end! But before we begin the journey, we have a little homework assignment.

Assessing your AI and data literacy

No book is worth its weight if it doesn't require its readers to participate.

Improving our AI and data literacy starts by understanding where we sit concerning the six components of our AI and data literacy educational framework. To facilitate that analysis, I've created the *AI and data literacy radar chart*. This chart can assess or benchmark your AI and data literacy and identify areas where you might need additional training.

A radar chart is a graphical method of capturing, displaying, and assessing multiple dimensions of information in a two-dimensional graph. A radar chart is simple in concept but very powerful in determining one's strengths and weaknesses in the context of a more extensive evaluation. We will use the following AI and data literacy radar chart to guide the discussions, materials, and lessons throughout the book:
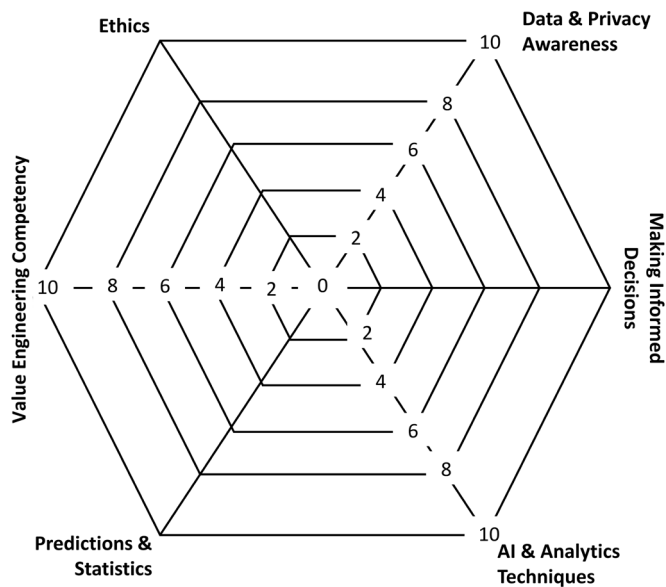

Figure 1.3: AI and data literacy radar chart

This assignment can also be done as a group exercise. A manager or business executive could use this exercise to ascertain their group's AI and data literacy and use the results to identify future training and educational needs.

You will be asked to again complete the AI and data literacy radar chart at the end of the book to provide insights into areas of further personal exploration and learning.

In the following table are some guides to help you complete your AI and data literacy radar chart :

Category	Low	Medium	High
Data and privacy awareness	Just click and accept the website and mobile app's terms and conditions without reading	Attempt to determine the credibility of the site or app before accepting the terms and downloading	Read website and mobile app privacy terms and conditions and validate app and site credibility before engaging
Informed decision-making	Depend on their favorite TV channel, celebrities, or website to tell them what to think; prone to conspiracy theories	Research issues before making a decision, though still overweigh the opinions of people who "think like me"	Create a model that considers false positives and false negatives before making a decision; practice critical thinking
AI and analytic techniques	Believe that AI is something only applicable to large organizations and three-letter government agencies	Understand how AI is part of a spectrum of analytics, but not sure what each analytic technique can do	Understand how to collaborate to identify KPIs and metrics across a wide variety of value dimensions that comprise the AI utility function
Predictions and statistics	Don't seek to understand the probabilities of events happening; blind to unintended consequences of decisions	Do consider probabilities when making decisions but underweigh a thorough assessment of the potential unintended consequences	Actively seek out information from credible sources to improve the odds of making an informed decision

Value engineering competency	Don't understand the dimensions of "value"	Understand the value dimensions but haven't identified the KPIs and metrics against which value creation effectiveness is measured	Understand the value dimensions and have identified the KPIs and metrics against which value creation effectiveness is measured
Ethics	Think ethics is something that only applies to "others"	Acknowledge the importance of ethics but are not sure how best to address it	Proactively contemplate different perspectives to ensure ethical decisions and actions

Table 1.1: AI and data literacy benchmarks

We are already using data to help us make more informed decisions!

Summary

In this chapter, we have seen that with the rise of AI and data, there is a need for AI and data literacy to understand and intelligently discuss the impact of AI and big data on society. The challenges and risks of AI include concerns about surveillance, manipulation, and unintended consequences, which highlight the importance of AI ethics and the need for responsible and ethical AI deployment.

This emphasizes the need to train individuals to become Citizens of Data Science in order to thrive in a data-driven world dominated by AI models. We proposed an AI and Data Literacy Framework consisting of six components, including data privacy, AI techniques, informed decision-making, statistics, value engineering, and AI ethics, to provide guidance and education for responsible and ethical engagement with AI and data.

We ended the chapter with a homework assignment to measure our current level of AI and data literacy. We will do the homework assignment again at the end of the book to measure just how much we learned along the journey, and where one might have to circle back to review selected subject areas.

Each of the next few chapters will cover a specific component of the AI and Data Literacy Framework. The exception is that we will divide the *AI and analytic techniques* component into two parts. *Chapter 3* will focus on a general overview of the different families of analytics and for what problems those analytics are generally used. Then, in *Chapter 4*, we will do a deep dive into AI – how AI works, the importance of the AI utility function, and how to construct responsible and ethical AI models.

Continuously learning and adapting...we're beginning to act like a well-constructed AI model!

References

1. Harvard Business Review. *Data Scientist: The Sexiest Job of the 21st Century* by Thomas Davenport and DJ Patel, October 2016: `https://hbr.org/2012/10/data-scientist-the-sexiest-job-of-the-21st-century`

2. The Economist. *The world's most valuable resource is no longer oil, but data*, May 2017: `https://www.economist.com/leaders/2017/05/06/the-worlds-most-valuable-resource-is-no-longer-oil-but-data?`

3. *Apple Plans to Have iPhones Detect Child Pornography, Fueling Privacy Debate*, Wall Street Journal, Aug. 5, 2021

4. Blueprint for an AI Bill of Rights: `https://www.whitehouse.gov/ostp/ai-bill-of-rights/`

5. *Artificial intelligence will help determine if you get your next job*: `https://www.vox.com/recode/2019/12/12/20993665/artificial-intelligence-ai-job-screen`

6. *LinkedIn 2018 Report Highlights Top Global Trends in Recruiting*: `https://news.linkedin.com/2018/1/global-recruiting-trends-2018`

7. *Employers Embrace Artificial Intelligence for HR*: `https://www.shrm.org/ResourcesAndTools/hr-topics/global-hr/Pages/Employers-Embrace-Artificial-Intelligence-for-HR.aspx`