Preparing for the AI Revolution

C-level executives, security professionals, IT technicians and audit managers face constant pressure to assess, assimilate and govern emerging technology, including artificial intelligence (AI).

For enterprises to adopt AI and make the most of it, they should understand what AI is today, what it can become and how it may be useful to the enterprise, now and in the future. Core AI technologies continue to evolve in a very dynamic and fluid marketplace. Several low-cost (or even free) educational options are available, all from reputable sources, to help enterprises learn more about AI. Building on that understanding, enterprises can dispel common misconceptions about AI, learn how to govern it effectively and fully exploit the technology to achieve stakeholder goals.

AI is a collection of cognitive services including natural language processing (NLP), machine learning (ML), and computer vision and indexing, among others, that may help to achieve specific business goals. In the popular imagination, AI tends to be regarded as a virtual intelligence that speaks to end users—think of Amazon Alexa or Microsoft Cortana. While these two examples reflect specific use cases for AI—with very broad penetration into the consumer mobile market—they represent only one small portion of the whole AI market.

AI Market Size and Diversity

Sizing up the AI market can be quite difficult, given the sheer number of market segments, the variety of participating industries and the diversity of related business applications. To complicate matters further, organizations that provide market analysis in the AI space use different methodologies to calculate current and projected market size. Consequently, their projections can vary substantially.

By some estimates, current AI market valuation reaches approximately US$10 billion. By 2024, the total market capitalization of the AI market is projected to grow at an impressive compounded annual growth rate of 37 percent (using current period forecasting models), ultimately to reach US$191 billion.[1] The difference in current and projected market valuation suggests tremendous growth in AI technology and related applications over a very short period of time.[2] As a consequence, job opportunities in AI could expand exponentially in the next decade.

AI Industries and Adoption

Technology organizations such as Google, Microsoft and IBM have entered the AI space. There is extensive adoption of AI, not only in the telecom and automotive industries, but also in professional services, retail, energy and natural gas.[3] AI assists these organizations with marketing and sales, product and service development, and risk and supply chain management, among other disciplines. Machine learning and probabilistic reasoning represent a majority of published research (56 percent as of 2017), followed by neural networks and computer vision.

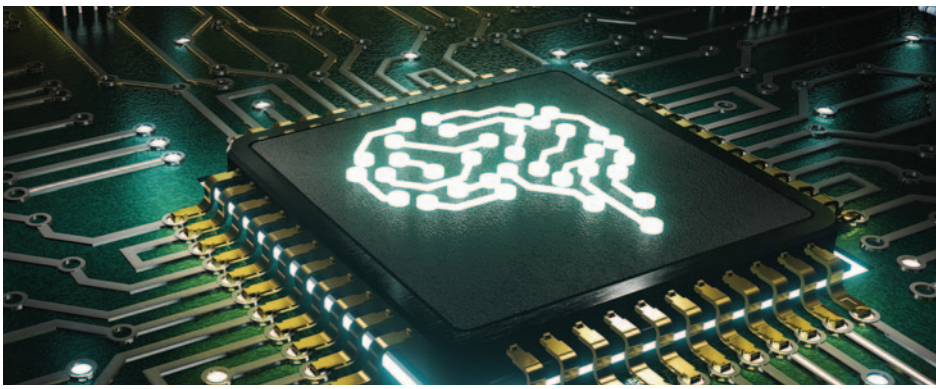

Adam Kohnke, CISA, CISSP
Is currently serving as the senior IT auditor for Total Administrative Services Corp. Kohnke has more than three years of IT audit experience and more than six years in IT operations with various Fortune 500 companies as an incident, change and project manager. As an IT auditor, Kohnke has performed continuous control testing engagements against the US National Institute of Standards and Technology Special Publication (SP) 800-53 rev. 4 standards, performed annual SOC 1 Type 2 and SOC 2 Type 2 engagements, audited Amazon Web Services (AWS) deployments, and been involved with various other technological or operational-based audit engagements.

> ALTHOUGH AI PROMISES TO REVOLUTIONIZE TECHNOLOGY AND BUSINESS, THE FUNDAMENTAL PRINCIPLES OF SECURITY REMAIN INTACT AND RELEVANT—IF NOT EVEN MORE CRITICAL THAN BEFORE.

There is heavy investment and development by governments; for example, since 2007, the Chinese government has increased AI investment by 400 percent (while over the same period, organizations based in China increased investment only by 73 percent). US congressional transcripts mentioned AI and ML fewer than 25 times between 1995 and 2016. Then, in 2017—for the first time—the terms occurred more than 25 times. Finally, in 2018, the terms exceeded 75 documented instances.[4] Based upon previous trends, governments worldwide will likely continue to expand AI investment, investigation and discussion.

Working in AI

AI jobs generally fall within computer and information research scientist, software programmer, and software developer skill sets, and enterprises should focus on building out, cultivating and hiring for these roles internally to meet future AI market demand. In the United States alone, computer and information research scientist roles are expected to grow 19 percent between 2016 to 2026—which far surpasses average overall job growth—while software developer roles are projected to grow at an even more aggressive 24 percent over the same time frame.[5] Based on these projections, opportunity for advancement will expand tremendously for individuals who hold these jobs presently. For those not yet in the roles—but interested in joining the field—education will be critical. New job seekers and incumbents alike require appropriate knowledge and skills not only to help enterprises assess and implement AI today, but also to build and update skills over time so that enterprises can continuously monitor, assess and (re)align business strategy, governance over AI and emerging technology.

Learning AI

Today it is easier than ever before for enterprises to provide their workforce with skills necessary to become AI experts. For example, the following online sources offer training and credentials for AI skills in demand today:

- **edX**—Founded by Harvard University (Cambridge, Massachusetts, USA) and the Massachusetts Institute of Technology (USA) in 2012, edX is a massive online learning platform that provides a variety of substantive classes in AI—mostly free or at very low-cost.[6] EdX also offers courses in a MicroMasters program, whose credits transfer to the master's program in AI at Columbia University (New York, USA). The introduction to AI course from edX takes about 25 hours to complete. It clearly presents ML, cognitive services and relevant business use cases; it enables the student to understand core concepts of AI, using a free version of Microsoft Azure. The course requires students to complete knowledge checks and hands-on lab exercises. Students must achieve a 70 percent overall score to receive a verifiable certificate that may be presented to employers or posted on social media for online verification.
- **Southern New Hampshire University (SNHU)**—SNHU offers a 120+ credit bachelor of science degree in robotics and AI that includes courses in ML, scripting, robotics, and calculus, among others.[7] The program is priced at US$960 per class. Prior college credit and certifications from accredited institutions are accepted to reduce overall time and resources required to complete the program.
- **Amazon Web Services (AWS)**—AWS currently offers six different instructional paths for ML.[8] All classes, videos and curricula can be viewed for free after signing up for an AWS certification and training account.[9] The number of classes ranges from four classes for the business decision-maker learning path to 12 for a specialty

certification in ML. The exam costs US$300 and can be scheduled at any time with no requirement to take any of the classes. AWS recommends at least one year of experience with ML services in AWS and offers two months of free usage with AWS SageMaker, one of the core ML services in AWS.[10]

Securing AI

Although AI promises to revolutionize technology and business, the fundamental principles of security remain intact and relevant—if not even more critical than before. Enterprises should understand the following principles and methodologies as they evaluate and adopt AI.[11]

Role-Based Awareness Training

IT managers, employees and stakeholders alike should understand their respective roles in securing data and safeguarding its privacy across all enterprise AI solutions. Role-specific training should focus on relevant regulations, data protection, classification, retention and authentication, and it should advance a common terminology for security and privacy across the enterprise. An example of role-based training for IT managers may also focus more on the specific tools deployed in the environment, the available features, their specific and potential use cases within the enterprise, how to frame the business problems using the tools, properly classifying the business problem, managing the data needed to solve business problems, and selecting performance data to measure outputs produced by AI solutions. Role-based training for standard employees may revolve around understanding common AI terminology, loading data sets, training ML models associated to AI tool sets and interpreting performance data produced by the tools.

Data Governance

AI services are often highly data-driven and data-intensive; enormous volumes of data may be required to realize the benefits of AI applications. As data are collected, modified, processed, transmitted and consumed by AI services, all individuals who are accountable and responsible for data at each stage of a given workflow should understand their responsibilities in terms of security, privacy and compliance. As AI is primarily data-driven, the primary responsibility of implementing and managing data governance for AI is likely to fall on the shoulders of the chief information officer (CIO), but may be the responsibility of others within the enterprise. Elements of a successful AI data governance framework may include the following high-level areas with these sub-elements and require the enterprise to define these areas further internally. Examples for the performance domain are shown in **figure 1**.

Threat Modeling

Threat modeling from a security perspective is an assessment process where system or data security risk is documented and analyzed, enabling the enterprise to understand a given system's threat profile as seen through the eyes of potential malicious users. Effort should be taken to define security objectives, identify vulnerabilities and design appropriate countermeasures, which is key to protecting sensitive data—not only at a component level, but also holistically across the AI ecosystem, especially at interfaces between networks or AI services and at the enterprise network boundary. Security should be integrated into the overall process and treated as a continuous discipline—not just revisited after a breach. The five elements of threat modeling are:

1. **Identify threats**—Using either data flow diagrams (DFDs)[12] or Unified Modeling Language (UML) deployment diagrams,[13] the enterprise should identify assets of interest (e.g., AI service account credentials, data sets, access tokens) and potential entry points into AI systems using application programming interfaces (APIs), web services or user interfaces. Malicious users will seek to access AI systems via available entry points; they are the starting points for

Figure 1—High-Level Overview of a Successful AI Data Governance Framework

Performance	Security	Privacy	Transparency
• Accuracy • Bias • Completeness	• Adaptability of data models • Adversary robustness	• IP capture • User impact	• Intent • Assurance level

understanding potential threats to AI solutions and their data.

2. **Understand the threat(s)**—To understand the potential threats at an entry point, the enterprise should identify security-critical activities that may occur and determine what malicious users might do to attack or misuse the system. Ask questions such as "How could malicious users leverage specific assets to modify how control over the AI solution or its data works, retrieve sensitive information, manipulate system data, escalate rights, or cause the system to fail or be unavailable?" In this way, the enterprise can determine the chances of malicious users accessing assets without being audited, bypassing access control checks or masquerading as another user.
3. **Categorize the threats**—An industry-accepted model for threat categorization is the spoofing, tampering, repudiation, information disclosure, denial of service and elevation of privilege (STRIDE) model.14 Threat classification is the first step toward effective mitigation. For example, if it is known that there is a risk that someone could make a change to a data set used for input with the AI solution, then that person later denies responsibility for the change, effective monitoring over data input sources with logging details of IP address, username, date/time, machine name, access token used and other details will help support nonrepudiation, spoofing and other items in the STRIDE model.

> "LEVERAGING THREAT TREES IS AN EFFECTIVE TOOL FOR DEVELOPING MITIGATION STRATEGIES."

4. **Identify mitigation strategies**—Leveraging threat trees is an effective tool for developing mitigation strategies. The potential threat sits at the root of the threat tree, and all potential children (or leaves) are the conditions that must be true for malicious users to realize the specific threat. Conditions may possess subconditions. For example, under the condition that malicious users can generate and provision access tokens needed to view and/or change data sets used in the AI solution, the fact that the malicious user must leverage an available system entry point is a subcondition. For each leaf condition, the enterprise should identify acceptable mitigation strategies and, in this example, effective access controls, multifactor authentication, detailed logging and monitoring for access token creation/assignment. Each path through the threat tree that does not end with a mitigation strategy should be viewed as a vulnerability.
5. **Test**—The threat model becomes a road map for effective penetration testing. Penetration testing evaluates threats by directly attacking a system and the data housed in the system, and it may be executed in an informed (white box) or uninformed (black box) manner.

Software Development

Industry-accepted best practices and principles for application development include code simplicity and review, continuous testing, coherence, realistic project estimation using Agile methodology, etc. At the project level, example best practices for development include documenting the objective function of the AI solution—what business purpose should the tool achieve? Another example includes avoiding unnecessary data cleansing from the model. ML models typically seek to detect patterns in the data, and a "dirty" data set may be more likely to provide business value vs. one that has been cleansed to the point where the data provide limited business foresight. At the technical security level, best practices include a review of authentication cookies or tokens used by AI services to ensure that they are protected from known vulnerabilities such as cross-site request forgery (CSRF) and cross-site scripting (XSS) attacks.

Change Management

Change management controls should be in place for all AI solution components and subsystems, including virtual machines, before deploying AI services. Controls should restrict who may

introduce changes and to what extent. Examples include preventing personnel with development responsibilities from having rights allowing deployment of code to the production AI solution. From an enterprise preparedness perspective, AI is a relatively new solution being rolled out into many enterprises that will change how the organization achieves certain goals. By focusing less on the feature requirements that will be provided by the AI solution, the enterprise adopts a public relations-style communication effort, which helps individuals at all levels of the enterprise understand how the solution will benefit them in their daily roles and how personnel can get specific questions they have answered and the training they need to successfully change over to successfully leverage the product.

Security Program

The enterprise should define and adopt a common security framework for all AI solutions as mainstream security frameworks for AI do not presently exist. Application of a security framework and the controls provided by the framework will provide the enterprise a common approach for securing the AI solution platform and the data housed. Ongoing maintenance should include vulnerability and penetration testing, strict patch management, and monitoring and logging of usage and changes to AI configuration and algorithms. The enterprise should review existing security incident response plans and ensure that those plans fully address responses to incidents affecting AI services. Personnel should be adequately trained through routine incident tabletop exercises, red team/blue team exercises or simulated attacks to respond in the event of AI security incidents or breaches.

Governing AI

Governing AI requires meticulous planning, continuous oversight and proper allocation of resources to ensure that adopted AI services fully meet the needs of the enterprise stakeholders. When considering the adoption of AI cognitive services, enterprises should ask the following questions to help anticipate, design and implement good governance practices:

- **How does the enterprise secure and validate the integrity of information that AI cognitive services access and ingest?**—Just like human beings, AI cognitive services typically require business files and online documents to examine, correlate and train themselves. These sources can include Excel files, Access databases, etc. AI cognitive services should use complete and accurate data to perform processes or they may reach faulty conclusions and trigger a multitude of undesirable consequences. Methods to ensure completeness and accuracy of data may include reconciliations of source data after they are introduced to the AI solution or after the AI solution modifies the data in its models.
- **How can the enterprise comply with regulatory authorities and customer expectations as data pass through AI cognitive services?**—The enterprise may deploy AI cognitive services on premise, in the cloud or somewhere in between. AI cognitive services may communicate with several external endpoints or report data directly to them. AI should only communicate with trusted sources (e.g., customers, business partner networks) that are designated by the enterprise, using industry standard protocols and should adhere to applicable data privacy laws such as the EU General Data Protection Regulation (GDPR), the US Health Insurance Portability and Accountability Act (HIPPA) and the US Gramm-Leach-Bliley Act (GLBA). Understanding input sources such as Excel spreadsheets, flat database files and web-based endpoints involved in AI cognitive service interactions are critical for compliance.
- **How does the enterprise authorize, monitor and maintain access to outputs from AI cognitive services?**—With Microsoft Azure, access tokens can be generated and embedded in various locations to allow consumption of data produced by a given AI cognitive service. This functionality appears to increase risk—most notably risk of unauthorized disclosure, use and/or monetization of enterprise data. Controlling individual access required to generate the tokens—along with revoking access tokens from the central console periodically and refreshing this information at end points—may sufficiently reduce associated risk.

In AWS Sagemaker, managing security access to data and related outputs is handled quite differently than Microsoft Azure. For example, Sagemaker Notebook Instances, which are suited for individual contributors to perform development activities, are Internet-enabled by default and should be reconfigured accordingly to allow only authorized connections. Sagemaker further relies on the appropriate setup on identity and access management (IAM) roles, permission policies and service resource configuration such as managing Simple Storage Service (S3) bucket access, where Sagemaker objects and data may be stored.

- **For voice and facial recognition AI cognitive services, is information being securely stored and deleted after it no longer serves business purposes?**—Storing voice or facial images for periods longer than required by law (and/or the business) introduces unnecessary cost and legal risk, especially related to privacy compliance. Developing data retention and purge directives—along with manual or automatic mechanisms to enforce them—will allow enterprises to retain only information required to meet stakeholder needs. Enterprises can start by identifying and documenting AI data inventories, categorizing those data by importance to the enterprise (i.e., sensitive, non-sensitive), and creating policies and procedures to implement data purge mechanisms that will either inform enterprise personnel to review data before purging it or automatically purge the data after its usefulness has expired.
- **Will the enterprise regularly take action to identify and remove human bias from its AI data processing?**—Human populations are not equally represented or empowered technologically; consequently, information products and services can exclude certain populations from accessing data and services. A mobile application, for example, may show users driving a car as the fastest path to their destination—but what about bus or bike riders? Do they get the same or similar information? If not, there may be an implicit bias toward people who drive cars—and the enterprise's potential revenue may be curtailed due to this bias.

AI learns to make decisions based on data provided by human administrators. Machine bias can be introduced inadvertently when human administrators fail to anticipate certain prejudicial inferences made via data proxies. For example, a user may wish to preclude bias from a data population or sample; therefore, the user decides to exclude the race of individuals in the population, but, without further thought, retains their zip codes. Because zip codes can be used to infer race—and thus constitute a potential data proxy—they may (re)introduce unintentional bias. Routine audits of data input sources, related components and processing algorithms can help identify and remove bias from cognitive services, where possible.

- **How will the enterprise identify, socialize and ensure compliance with emerging regulations on local, state, national and international levels?**—Poor implementation of AI can have legal ramifications. Governance stakeholders across the enterprise should be asking two critical questions:
  - If AI makes the wrong decision or reaches an incorrect conclusion—and the enterprise acts on it—what are the consequences?
  - What laws govern the use of AI?

> "IF THE PACE OF AI REGULATION IN 2017-18 IS ANY INDICATION, THE NEXT FIVE TO 10 YEARS MAY PRESENT A REGULATORY AVALANCHE REQUIRING ENTERPRISES TO EXPAND RESOURCES—BOTH TECHNOLOGICAL AND HUMAN—TO KEEP UP."

Several AI-related bills have been introduced recently by the US Congress.[15] The Self Drive Act (H.R. 3388) addresses safety of automated vehicles; the AV Start Act (S. 1885) aims at proper use of driverless cars; the Future of Artificial Intelligence Act (H.R. 4625) seeks to create an advisory committee on AI-specific issues; the AI Jobs Act (H.R. 4829) requests a US Department of Labor report assessing the impact of AI on the workforce; and the National Security Commission

Artificial Intelligence Act of 2018 (H.R. 5356) seeks to establish a commission that reviews advances in AI, with a focus on promoting US national security. If the pace of AI regulation in 2017-18 is any indication, the next five to 10 years may present a regulatory avalanche requiring enterprises to expand resources—both technological and human—to keep up. There are not many AI-related regulations outside of the United States.

Conclusion

AI has introduced a myriad of new challenges. It entails a significant learning curve—not only in assessing, adopting and securing AI solutions, but also in complying with existing regulations and preparing for new or emerging ones. Governments worldwide are expanding requirements, and enterprises must keep pace with a shifting (and often internationally fragmented or discontinuous) regulatory landscape. Despite the challenges, AI promises explosive growth potential across many industries and facets of human life. AI introduces new potential growth regarding labor production and how individuals effectively use their time. Opportunities to optimize AI for business and product innovation, augmentation of the labor force and AI's potential to grow domestic economies are the shining potential resulting from the AI revolution.

Endnotes

1 MarketWatch, "Artificial Intelligence Market Size Is Projected to Be Around US$191 Billion by 2024," 8 August 2018, *https://www.marketwatch.com/press-release/artificial-intelligence-market-size-is-projected-to-be-around-us-191-billion-by-2024-2018-08-08*

2 Market Research Engine, *Artificial Intelligence Market by Technology (Machine Learning, Natural Language Processing, Image Processing, Speech Recognition); by End-User (Media & Advertising, BFSI, IT & Telecom, Retail, Healthcare, Automotive & Transportation) and by Regional Analysis—Global Forecast by 2018-2024*, December 2018, *https://www.marketresearchengine.com/artificial-intelligence-market*

3 Shoham, Y.; R. Perrault; E. Brynjolfsson; J. Clark; J. Manyika; J.C. Niebles; T. Lyons; J. Etchemendy; B. Grosz; Z. Bauer; "The AI Index 2018 Annual Report," AI Index Steering Committee, Human-Centered AI Initiative, Stanford University, California, USA, December 2018, *http://cdn.aiindex.org/2018/AI%20Index%202018%20Annual%20Report.pdf*

4 *Ibid.*

5 Bureau of Labor Statistics, "Computer and Information Technology Occupations," USA, *https://www.bls.gov/ooh/computer-and-information-technology/*

6 edX, "Microsoft Professional Program in Artificial Intelligence," *https://www.edx.org/microsoft-professional-program-artificial-intelligence*

7 Southern New Hampshire University (SNHU), "Online Bachelor's Degree: BS in Information Technologies Robotics & Artificial Intelligence," *https://www.snhu.edu/online-degrees/bachelors/bs-in-information-technologies/robotics-and-artificial-intelligence*

8 AWS, "Machine Learning," *https://aws.amazon.com/training/learning-paths/machine-learning/*

9 AWS, "AWS Training and Certification," *https://www.aws.training/*

10 AWS, "Amazon SageMaker Pricing," *https://aws.amazon.com/sagemaker/pricing/*

11 Prabhu, M.; "Security and Privacy in Artificial Intelligence and Machine Learning—Part 1: Lay of the Land," *Towards Data Science*, 28 July 2018, *https://towardsdatascience.com/security-and-privacy-in-artificial-intelligence-and-machine-learning-part-1-c6f607feb94b*

12 Agile Modeling, "Data Flow Diagram (DFD)s: An Agile Introduction," *www.agilemodeling.com/artifacts/dataFlowDiagram.htm*

13 Agile Modeling, "UML 2 Deployment Diagrams: An Agile Introduction," *www.agilemodeling.com/artifacts/deploymentDiagram.htm*

14 Howard, M.; D. LeBlanc; *Writing Secure Code*, Microsoft Press, USA, 2003

15 Fonzone, C.; K. Heinzelman; "What Congress's First Steps Into AI Legislation Portend," Bloomberg Law Big Law Business, 15 May 2018, *https://biglawbusiness.com/what-congresss-first-steps-into-ai-legislation-portend*