ARTIFICIAL INTELLIGENCE

THE DATA BELOW

Mike Koenig, MSE, Sarah Bee, MBA, CIA, and Dennis Applegate, CIA, CPA, CMA

About The Internal Audit Foundation

The Internal Audit Foundation has provided groundbreaking research for the internal audit profession for more than 40 years. Through initiatives that explore current issues, emerging trends, and future needs, the Foundation has been a driving force behind the evolution and advancement of the profession. We exist to help audit leaders, practitioners, students, and academics experience continuous growth in their careers to propel them to become respected as trusted advisors and thought leaders within the industry.

As internal auditors, knowledge is key to the continued investment you make in your professional development. The Foundation remains focused on providing the relevant knowledge and insights internal auditors need to continue growing as professionals, which results in greater value to our organizations, stronger talent, and enhanced advocacy for internal auditing.

Private donations ensure the sustainability of the Foundation and the resources it provides back to the profession. Join us as we work to support and advance the continued practice of internal auditing.

Invest in your profession. Make a donation today at www.theiia.org/foundation.

Sponsorship

The Internal Audit Foundation appreciates the generous sponsorship of this report by The IIA–Chicago Chapter.

Limit of Liability

The Internal Audit Foundation publishes this document for information and educational purposes only. The Foundation does not provide legal or accounting advice and makes no warranty as to any legal or accounting results through its publication of this document. When legal or accounting issues arise, professional assistance should be sought and retained.

Copyright © 2018 by the Internal Audit Foundation. All rights reserved.

ARTIFICIAL INTELLIGENCE
THE DATA BELOW

Introduction

Artificial intelligence (AI) applications are ubiquitous in business and our personal lives. From asking your smartphone for the weather forecast to determining the credit worthiness of a customer, AI creates efficiencies in our personal lives but may pose complexity and risk for our internal audit profession. Often its presence is so subtle that many of us do not even realize the impact of AI on the workplace of our clients and, by extension, our audits. While many internal auditors are competent in information technology (IT) governance, risk, and controls (GRC), it would be dangerous to overlay IT audit concepts and techniques onto AI applications absent an appreciation of AI and its unique characteristics. The purpose of this report is to begin by explaining two such characteristics.

First, unlike most IT systems, AI uses probability rather than correctness to obtain results. Second, AI applications use data as both a system input *and* a driver of results, meaning results will change as the data evolves.

Understanding the Building Blocks

Internal auditors address IT risks every day either directly or indirectly. It is unusual to audit a process that is not automated all or in part. Even if the internal auditors are not auditing the application, they are working in concert with IT auditors to complete their audit objectives. Therefore, it is logical for internal auditors to apply tried and tested IT audit techniques to AI applications. There is risk, however, in applying a "one-size-fits-all" mentality to these more sophisticated applications.

Simply put, four primary components constitute a traditional automated system: input, transforming processes, output, and storage. The user enters data into the system, the system processes the data, the processed data is released in output form, and the data is then stored. Data does not drive the result; programming drives the result. Conversely, in an AI application, data is integral to the decision making. AI programs are basing their output on millions of data points, not just one input.

As internal auditors, we test data against an expected result. There is no expected result with many AI applications. The very unpredictability of output creates risk, providing a challenge for auditors. The first step in addressing this new risk is gaining an understanding of how data contributes to the output of AI applications. In this report, we will define major AI data designs and then propose a series of client inquiries that internal auditors may pose to gain confidence in the system controls supporting AI applications.

Technology Approaches

AI uses many information technologies, including static systems and machine learning systems. Static systems—such as decision trees, classifiers, and rule tables—are testable, a major advantage to their use. Internal auditors can review the inputs, the data sets, and the outputs to ensure that with each iteration the same results occur. As such, static systems are amenable to traditional IT audit methodology.

Alternatively, machine learning systems are based on a different technological approach altogether. In these applications (associated learning, artificial neural networks, decision tree learning, deep learning, and reinforcement learning, among others), the system "learns" what the best prediction should be. The disadvantage of such systems is that they are much more difficult to validate through standard audit methodology, as discussed below.

Probability vs. Correctness

Machine learning takes into account multiple factors and selects those that will help predict the best, most reliable outcome. As such, machine learning systems are not interested in being absolutely correct but rather are designed to be "good enough," but not necessarily always correct, as further explained below.

Rather than provide a correct solution to a given problem, machine learning systems calculate the probability that a given outcome is correct. Outcomes are then ranked in descending order of probability. This process creates a dilemma when considering if a given outcome is correct or incorrect, for machine learning systems do not care if they have included all the correct solutions; they only care that the best answer is near the top of the ranked outcomes.

For example, Facebook or LinkedIn may recommend a set of likely "friends" to build your social network, a system goal. However, rather than ensuring that every recommendation is a friend, such systems will generate a set of recommendations, only some of which are correct or even reasonable. For instance, if the system generates a set of ten likely friends, one of whom is a close friend while five are acquaintances and four are complete strangers, and you choose to add the close friend, the system considers the outcome a success. It has accomplished the goal of adding someone to your social network. Interestingly, the machine learning algorithm does not care that only 60% of the set of likely "friends" were reasonable, and that 40% were completely wrong. It met its goal of growing your social network. The key to understanding machine learning systems, therefore, is to realize that the system always returns the fully ranked probability set, regardless of whether it is correct. In the above example, "complete strangers" represented 40% of the set of likely friends. However, in the machine learning algorithm, the unreasonableness of this set would not indicate system failure. In fact, there is no concept of failure in machine learning systems. Nonsense-based results always will exist below the predetermined machine learning desired target threshold, in this case, 60%.

> The difficulty comes when applying machine learning where correctness is critical. Consider a recent machine learning diagnostic system deployed in China for diagnosing thyroid nodules. Doctors diagnosed the nodules with 70% accuracy. However, the machine learning system achieved 85% accuracy, higher than the human doctor threshold. The system may be below the accuracy of skilled doctors, but because it is above the threshold of 70% for average doctors, the outcome is considered a success.

But, how can 60% be good enough? Surprisingly, for most machine learning systems, low accuracy is acceptable. For most of us, it is counterintuitive to think that 60% is good enough. We want to live in a world of 100% certainty. However, search engines, photo identification, and even grammar checkers often contain low thresholds. Why? Because machine learning systems are not interested in correctness, but rather in achieving the desired outcome. Therefore, internal auditors always should be cognizant of the idiosyncrasies and inherent limitations of machine learning systems when planning and performing assurance engagements. Not having a predictable outcome to test may present a significant auditing challenge.

Changing Results

Machine learning systems change over time based on fine-tuning the data, new or revised data sets, and additional algorithms. These adjustments will likely result in different outcomes for the same question when asked at different intervals. Machine learning systems also use a set of data to "train" the system. This training data is what the system algorithms first learn. The data is then used to adjust the algorithms, going forward, to achieve the initial desired target threshold.

The quantity of data, as well as the quality of data, affects the results of the machine learning systems. For example, suppose the objective is to train a machine learning system to identify whether the word "pharaoh" is spelled correctly in a document. Looking up this word in a dictionary would be a rule-based approach. The rule states that if the word is in the dictionary as "pharaoh" and it matches the spelling in the document being tested, then it is spelled correctly 100% of the time.

Yet, a machine learning system would approach the spelling of "pharaoh" quite differently. It would look at data sets involving millions of words and calculate the frequency of the word "pharaoh." However, in doing so, it would find that many newspapers in 2015 suddenly began to write "pharoah" instead of "pharaoh." What happened? Well, the racehorse, "American Pharoah," happened to be winning many races that year. A machine learning system would determine which spelling was correct based on the frequency of use that it observed in the data sets, rather than on the standard dictionary spelling. It therefore would have learned an incorrect spelling. Why? Because the frequency of "Pharoah" the racehorse that it found in the data sets was significantly higher than that of "pharaoh" found in the dictionary.

The internal auditor should consider machine learning systems as the ultimate in the "majority decides" approach, certainly a paradigm shift in our thinking. Where previously our goal was correctness, the goal is now "popularity." To evaluate the cost benefit, increased efficiency, and better decision making enabled through machine learning systems, the internal auditor must first understand the basis of the AI data used, as discussed below.

Data Is Key

All AI machine learning systems need data sets to generate useful analytical output. Since data fuels these systems, an auditor must understand data concepts, data controls, and AI approaches to data analysis when developing an audit of an AI system. In the following sections, we describe various types of AI data sets and suggest potential inquiries of the audit client when planning an AI audit and assessing related risks.

1. Garbage in, Garbage Out - How can accurate statistical analysis occur over massive amounts of inaccurate data? In the world of statistics, it is critical to review data to ensure it is relevant, error-free, and free of outliers. How do you know if your machine learning data is therefore clean or simply garbage? Machine learning systems need hundreds of thousands of data points at the low end to tens of millions of examples at the high end. How can you ensure that tens of millions of examples are all appropriate for their intended purpose? How can you be sure that data cleansing did not remove appropriate rows or alter the data set? For example, when reviewing financial forecasts, do you remove all forecasts for terminated departments, leaving only data from current departments available for analysis? Searching inquiries similar to these need to be posed to the audit client managing an AI system.

2. Origin of Data - Data can be extremely costly to collect, clean, process, and prepare for analysis. Yet, repurposing the data may lead to incomplete or inaccurate results. For example, rather than try to recreate missing historical data on employee retention, a data scientist may simply repurpose data off attendance lists from annual corporate meetings as a proxy for predicting employee retention. If meeting attendance grew from 2,145 attendees in 1995, to 5,824 attendees in 2014, dipping to 2,322 attendees in 2015, then combining employee attendance with demographic, residential, and salary data could predict employee retention. The system, however, may incorrectly conclude that *location of residence* is a leading indicator of employee retention and bias employees who live within walking distance. However, other factors may enter into it. For example, the system would not know that the threat of snow in 2015 kept many attendees away from the annual meeting event. The shortcut of repurposing data as a proxy for actual data may cause AI models to yield inaccurate results or violate data usage policies. As a consequence, audit clients must be asked whether there is a process in place to validate the appropriateness of the data used for machine learning systems when repurposing the data is an available alternative.

3. Data Bias - Since machine learning systems are interested in prediction thresholds rather than correctness, they will not know whether the data sets contain bias. For example, data sets to generate machine learning models for disease detection may have an unknown bias against females, were they to contain only MRI results of Midwestern males. Accordingly, the audit client must explain to the auditor how data in the AI model is reviewed and validated to ensure all representative populations exist in the appropriate frequency to ensure a correct and consistent output.

4. Data Lakes - The world is full of unstructured data. Newspaper articles, photos from social media, and corporate records all contain diverse data elements. Because of the unstructured nature of the data, storage also occurs in unstructured ways. It is not economically feasible to create and manage structured representations of all data for a machine learning system. For that reason, *data lakes* are established to serve as a repository for unstructured documents. Machine learning systems fish the data lakes as part of the learning process**,** casting a wide net into a lake teaming with data fish. Because of the vastness of the lake, it is often unknown exactly what data exists, the origin of the data, as well as the appropriateness of the data to the machine learning system. Hence, the audit client must detail for the auditor what data lakes are allowed and how they are controlled for the machine learning system to use.

5. Data Leak - Do you have a breach in your dam? Because of the unstructured nature of a data lake, how can you ensure that people fishing in the lake only access the information they need? How do you ensure that data from selected users is stored only in a secured portion of the data lake, and that only authorized parties are able to go fishing in the lake to see what it contains? Given the potential for data leaks, the auditor must always inquire into the nature and extent of the client's data safeguards and ongoing monitoring of its machine learning model for proper data governance.

6. Data Drift - Refreshing and pruning data sets to ensure freshness cause an interesting phenomenon known as *data drift*. Machine learning models trained on data sets last month may behave differently when the models are trained on new data sets this month. A condition such as this results in the model output changing over time, affecting data output correctness and consistency. In a legal, human resources, medical, or financial application, showing that the system can produce the same results for a known input month over month is often vitally important. However, in a shopping scenario where sizes, preferences, and fashion trends continuously change, the correctness and consistency of such factors are less important than knowing whether they increase sales.

The latest technology approaches to machine learning are even more problematic. The new models are designed to feed data sets back into themselves, leaving little opportunity for human review of data output at various stages of system operation. In systems where correctness and consistency are the goal, then monitoring AI results over time becomes increasingly difficult. Therefore, audit clients must be able to identify for the auditor specific controls over machine learning data sets whose objective is to preclude data drift.

Artificial Intelligence Governance

Governance includes the oversight of AI development and operations. This oversight is especially critical because such applications possess higher levels of risk than traditional information systems. There is an inherent risk that developers will tout AI applications as accurate and insightful, but these claims are often exaggerated. The clear benefit of AI must be weighed against the risks of incorrect interpretations of the data output. No matter how well-controlled an AI system, it is only as strong as its weakest link. Often that link is hidden within the complexity of the system itself. The new model of AI applications requires a renewed focus on system oversight to include a high level of coordination between executives, data scientists, programmers, attorneys, and auditors, among other players. There is not a template to follow to manage AI governance; the playbook has yet to be written.

Nonetheless, internal auditors should explore the care taken by business leaders to develop a robust governance structure in support of these applications, providing assurance that risks are being identified and addressed, starting with the data issues described earlier in this report.

Special attention should be paid to the systems development process. It is during systems development that ground rules are laid down for the AI applications and when system developers exert the most control. Post-implementation, as neural networks are fed data, they have a certain amount of autonomy; namely, they are not following a direct programming command. The level of risk post-implementation can be mitigated by thoughtful and thorough risk assessment during system development. The appendix contains suggested auditor inquiries to help identify risks inherent in AI applications.

While many data scientists and computer programmers possess an instinctive understanding of risk, internal auditors supplement their own proficiency in risk assessment by consulting two systematic frameworks that have proven to be invaluable in auditing systems development—a combination of IIA GTAG's coupled with ISACA's COBIT. These frameworks specify that internal auditors should be engaged in the entire lifecycle of AI projects, from program design to maintenance and support. Key areas of evaluation include not only all phases of the system lifecycle, but also a thorough review of the data used to drive the AI application. While data integrity is a key risk for all computer applications, it is especially critical in designing AI applications. Given the new nature of the technology, is it understandable that innovators, data scientists, and computer programmers are excited about the opportunities and benefits. The systems can generate accurate and insightful results. However, excitement for the technology may also cause a bias toward its potential and a lack of focus on governance measures related to its risk.

Conclusion

AI applications can produce extraordinary benefits for an organization, enhancing its decision making and efficiency. On the other side of the benefits, however, is an extraordinary amount of risk. This report focused on risks associated with the utilization of data as an integral piece of the analytical decision-making process in AI applications. Awareness of data's role in AI will help internal auditors design an audit plan that addresses these distinctive risks. Importantly, failure to identify and control AI data risks up front will generate further risks downstream, specifically to reputation, reporting, and management decision making, just a few of the land mines associated with AI.

Executive leadership needs to recognize that AI systems are not just a faster and better IT application, but instead represent a completely different approach to system data processing and decision making, bringing with it new operational and strategic risks. By embracing a leadership role in developing sound AI risk management, starting with controls over AI data, business leaders can guide, successfully, their organizations through the AI system transition now underway.

Appendix: Artificial Intelligence Systems – Selected Inquiries for Audit Planning

The following table identifies certain artificial intelligence (AI) audit issues, the associated system risk related to each, selected auditor inquiries in support of the preliminary survey phase of an AI audit, and expected responses indicating the extent to which client management understands the nature of the issues and the risks involved.

Issue	Risk	Audit Inquiries	Expected Responses
Is there sufficient audit proficiency on staff to execute an audit of an AI system?	Internal auditors, lacking AI proficiency, may fail to identify defective system controls.	What aspects of the system are newly created? Which aspects are modifications of existing systems? Which aspects are using existing established systems?	New AI systems require a more in-depth understanding of the system and its data sources, and therefore should carry more risk. AI systems that leverage existing, established systems should carry lower overall risk, and require less detailed system understanding.
Are AI system decisions validated?	Lack of explicit review of input and output data at each stage of the data pipeline may cause inconsistent and incorrect results, as well as audit scope limitations.	What stages of the system and specifically the data pipeline are designed for data validation? Where is the input data sourced? Where is the output data stored? Is the output data deleted at some stage of the pipeline, or is it repurposed for other systems? Is the AI application programmed in such a way that decisions made by the program are traceable?	AI systems that process inputs and outputs at all stages of the data pipeline should facilitate auditing and client monitoring of system decisions based on assessed risk. Black box AI systems, in which data transformations occur throughout the data pipeline without audit visibility of the internal system workings, should make it more difficult to validate system decisions and assess the level of system risk.

Issue	Risk	Audit Inquiries	Expected Responses
Is user access to AI system output and subsequent user interpretations appropriate?	AI system and pipeline stage output may reveal sensitive data violating government regulatory requirements. Misinterpretations of the output may cause ill-informed management decisions, leading to poor operating performance.	What regulatory requirements, if any, pertain to the AI application? Is access to the system output restricted to authorized users, and is access monitored periodically? Are authorized users of the output interpreting it correctly? Are criteria in place to assess the quality of the interpretations based on system output?	Well-designed AI systems should maintain data safeguards that govern access to system data consistent with company protocols and that assure the correctness of user interpretations consistent with system criteria. System monitoring should be designed and implemented to minimize the risk to each.
What sources of data were used to train the AI system?	Data used for AI system training may violate restrictions on usage. System inability to reconstruct the same output from the same input may restrict or prevent audit execution.	What is the origin of the training data? How current is it? What errors were found in the training data sets during processing? What data within the sets was changed or rejected and why? How is the training data updated, and how often do the updates occur?	The data scientist should inspect and evaluate the training (and production) data sets periodically for relevancy, accuracy, and completeness. The nature and extent of this review should indicate the quality of the data sets and the overall state of data errors in the system, driving a tentative conclusion about the reasonableness of the system outputs. Errors in the AI training (and production) data sets should be expected when large numbers of data records are involved. (Clean, error-free data is never fully achievable in AI systems.) Nonetheless, at a minimum, the data scientist should have reviewed the data sets for errors, documented the adjustments made, and calculated a percentage estimate of the errors that remain.

Issue	Risk	Audit Inquiries	Expected Responses
What training data sets are used to generate the initial AI system?	Data used for "training" the system may not represent the true population being analyzed for the AI application. Training data that fails to provide sufficient examples and exceptions for all conditions of the application may result in an algorithm that contains bias in the predictions generated.	How large are the training data sets? How were the sets selected? Who verified each data input and output? Which data sets were modified to achieve the current desired output? Does the training data contain examples that represent most of the actual data conditions expected over the life of the system?	The data scientist should manually evaluate and select training data to fine-tune the probabilities used by the system application. Deleting or "trimming" training data that reflect ambiguous conditions is acceptable in the test phase but not for the larger data sets used in the production phase. (For example, the name "Pat" reflects ambiguity as to gender, but it ought not to be "trimmed" from the AI production data for that reason alone.) System reviews should be in place to evaluate the deletions and to safeguard against the data scientist "over-trimming" the training data to fit the desired results.
What current or previous AI systems have used the same or similar sources of data for their data sets?	AI data usage may violate third-party data rights, government regulations, or company policies and cause unnecessary legal action. Data sets of a given AI system may not be easily updated, impairing the maintenance of other systems using the same data sources.	Have the system data sets been reviewed for compliance with legal criteria or company policy, and have they received necessary third-party permissions for use? How are the data sets for the AI system updated, and how frequent are the updates? How difficult is it to evaluate and improve system performance post deployment? Can system defects be identified and corrected in a timely manner? Have other AI systems been audited that use the same data sets as the current system under audit?	If AI systems that use the same data sets fail to comply with regulatory, contractual, or policy criteria, or fail to achieve system goals post-deployment, then additional client reviews of the data sets are warranted and ought to be expected. Data scientists managing new AI systems are expected to repurpose the data sets of existing systems, given the time and cost involved with developing new data sets. Such data sets are expected to contain bias and not reflect all of the characteristics of the data domain necessary to achieve the objective of the new AI system.

Issue	Risk	Audit Inquiries	Expected Responses
How do the AI system features make predictions?	Biased results generated by the AI system features may impair the quality of reporting and management decision making. AI system features that rely on personally identifiable information (PII), or user interactions with others may lack proper safeguards.	What features of the AI system are used to make predictions? How is the data collected for each system feature? Has the intended purpose of the data been approved for use? Has approval of each user feature that is based on interactions with others been obtained from an approved user source?	Each AI system feature should be explicitly approved. "Inferred" features should be monitored to prevent interpretations of system output based on inappropriate or incorrect relationships, such as using zip codes to infer levels of education or wealth rather than actual wealth and education data.
What control data sets (reflecting the data domain of the AI system) were used to validate the integrity of the actual AI data sets?	Data used in the control sets may not represent the actual production data. System developers may bias results by removing outlier data from the control data sets. This action, referred to as "over-fitting" the data, may result in a model that works only for the training sample and that performs poorly with actual data over time.	How large were the control data sets? What method was used to validate the control data sets? Did the samples used represent all possible scenarios? How were discrepancies in the control data sets corrected? Was the incorrect data removed, or was the system adapted to accommodate the incorrect data?	System developers should review the actual system output for data integrity, rather than the locked control data sets used to train the system. The control data sets should remain locked throughout the test phase to ensure that the algorithm is processing data in an unbiased way. Representative data samples, reflecting data breadth and involving critical scenarios, should be used as control data to validate large test data sets.

Issue	Risk	Audit Inquiries	Expected Responses
How were the AI system data sets validated, independently through third parties retained to judge data quality, a common IT practice, or by some other means?	Use of automated-based systems to judge data integrity may conceal underlying system issues that adversely affect the quality of the output. If the control data sets are based on current system data, then failure to validate such data for correctness may produce a flawed assessment of actual system data. Third-party judges of data quality may draw erroneous conclusions about the integrity of the system data if they lack authorization to access and examine actual system data input Failure of third-party judges to conduct their tests in countries in which the AI system resides may cause an inadvertent violation of government regulations, given national and international restrictions on the dissemination and use of sensitive data, such as PII.	What process ensures that the third-party judges selected to evaluate data quality possess sufficient expertise in the data domain of the system? When third-party judges use algorithms to determine data quality, how is the process reviewed and verified? Have the third-party judges received the proper permissions to review raw data inputs and resulting system outputs, especially if PII is involved? How are the results of the third-party judges reviewed to ensure they are correct? What is the process for testing the integrity of new system data, and for reviewing inconsistencies in the test results of the third-party judges?	A process should be in place to test the integrity of actual system data against a control data set. All data sets used in the system should be tested and validated. A small percentage of errors in the system data is normal and acceptable; however, the system should contain features that manage and mitigate the adverse effect. If third-party judges tested the system data, the test results should be confirmed for correctness. (A customary practice is to compare the test results of three independent, third-party judges and identify and reconcile inconsistencies.)

Issue	Risk	Audit Inquiries	Expected Responses
Is there a process in place to monitor the AI system based on performance metrics?	Failure to use performance metrics that assess the quality of system output may not reveal issues that weaken user acceptance. Systems lacking metrics to monitor the quality of the system output, including false positives and false negatives, may overstate true system performance. Failure to implement performance metrics that measure system compliance with relevant business rules, such as IRS rules governing an AI income tax application, may permit defective output to go undetected.	How does the system report variances from established performance metrics? What are the current and historical system rates for correct results, false positives, false negatives, and incorrect results? Are system users able to provide feedback on the system directly? If so, how is this accomplished? What types of errors are users reporting about the system?	Performance metrics should measure correctness of data output, user acceptance of system results, and system compliance with business rules. Such metrics should monitor both system training and production data, covering the same time periods for each. Metrics that focus on user acceptance should not emphasize the popularity of the system results to the detriment of their correctness, especially when correctness is critical to the quality of the system performance.

About the Authors

Mike Koenig, MSE, has 25 years of experience at Microsoft, as a software engineering leader, serving as Group Program Manager for Microsoft Office Natural Language Experiences and Director of Development for Microsoft Robotics among others. He holds 19 technology patents, including artificial intelligence and related fields. He is currently a full-time lecturer at Seattle University in the computer science department. koenigm@seattleu.edu

Sarah Bee, MBA, CIA, is the Director of the Internal Audit Education Partnership Center of Excellence at Seattle University. She has been teaching risk and controls to aspiring internal auditors for over 20 years. She keeps her skills current by working in the summer as an internal auditor for companies such as Alaska Airlines, Weyerhaeuser, and Moss Adams. Sarah Bee has published numerous academic articles related to risk and controls in *Accounting Information Systems Educators Journal*, *Accounting Perspective*, *Internal Auditor* magazine, *Journal of Forensic Accounting*, and *CPA Journal*, among others. bees@seattleu.edu

Dennis Applegate, CIA, CPA, CMA, CFE, CFM, served on the management team of Boeing Corporate Audit for 20 years, holding several management positions including Senior Audit Manager and Chief Auditor. Since 2014, he has taught internal auditing and accounting courses for Seattle University as a full-time lecturer and adjunct professor. During his professional and academic career, Dennis has authored 17 articles on auditing and audit management published in various professional and technical journals, and currently serves on the Editorial Advisory Board of The IIA. applegad@seattleu.edu

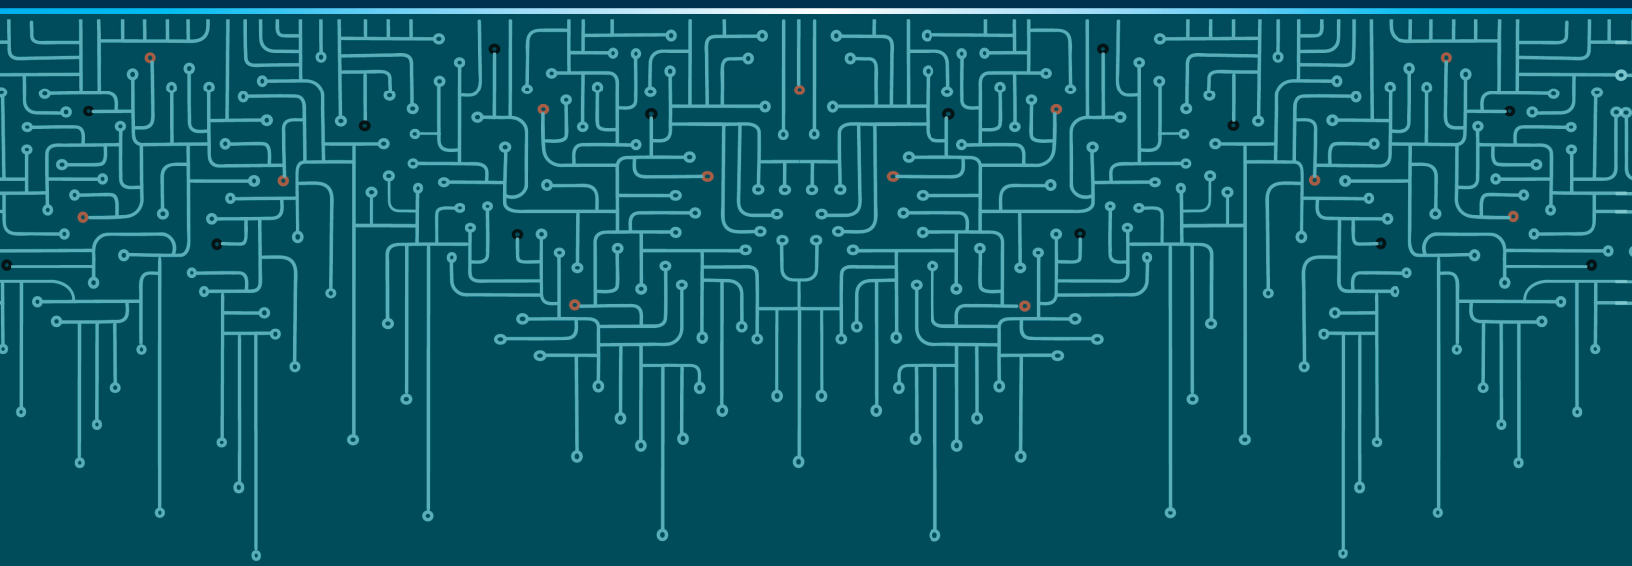

INTERNAL AUDIT FOUNDATION™

2018-0431