Risk Management in the Digital Age

By Rob Ceske, Kelly Combs, Nadim Hraibi and Jamie Hooten[1]

This article first appeared on www.theclearinghouse.org. *It is reprinted here with permission.*

Numerous technological advancements are now available to financial institutions that allow them to increase efficiency and keep up with changing consumer demands. Intelligent automation (IA) can range from simple algorithms to cognitive technologies, which have the ability to "learn" and adapt. (See page 27 for an overview of IA) Each type of automation can drive efficiency and effectiveness but also can introduce unique new risks. Traditional risk management techniques, which attempt to detect bad decisions or "rogue" employees and ensure appropriate lines of defense, must be adapted to address these new risks. With fewer human touchpoints throughout IA processes, the importance of design and appropriate usage, anticipating potential unusual circumstances, testing, and monitoring becomes paramount. Risk management teams will need to adapt their thinking and approaches to these new technologies and be proactive in reducing design risks and detecting unintended consequences of the new digital landscape.

FRAMEWORK

Leveraging IA can help financial services firms to automate processes, increase efficiency and consistency, and allow existing human labor to focus on more strategic activities while also improving customer experiences. However, once IA's strategic mandate is defined, its adoption should include strategic, design, and operating considerations as part of an IA risk management framework (see Figure 1). The risk management of IA is a key part of these considerations. Although existing risk management approaches and disciplines (such as those for model risk management) can be leveraged for IA solutions, machine learning and cognitive tools, in particular, require adaptation to traditional risk management approaches that were developed for human decision makers. Many risk techniques were designed to mitigate and/or detect isolated bad decisions or "rogue" behaviors. However, as Class 2 and 3 IA technologies speed operations and decision making, risk approaches need to adapt to focus on design and monitoring activities.

Because Class 2 and 3 IA solutions have less involvement from humans in their operation, risk strategy will require a heightened focus on design components—for example, technology "fit for purpose" reviews, appropriateness and comprehensiveness of calibration approaches, and planned monitoring mechanisms. Solutions that have not incorporated risk monitoring into design are likely to be more challenging to monitor because decision processes in machine learning and cognitive solutions are much harder to discern than processes from operations run by people (e.g., poor or biased underwriting decisions).

Risk management in design should have much more focus on the technology component versus process and people components than would be the case with more traditional activities. Depending on the class of IA that is being used, there will be relatively more emphasis on programming and scripting (with Class 1 IA), tool configuration, and data used for calibration/machine learning (as with Class 2 and 3 IA tools). Class 1 IA tools require much more careful design planning, particularly regarding technology interaction, than traditional solutions. For example, changes to the color of an "OK" box, email formatting, or latency (for computer response time) can all affect Class 1 IA, but none of these changes would affect processes managed by people.

> Leveraging IA can help financial services firms to automate processes, increase efficiency and consistency, and allow existing human labor to focus on more strategic activities.

Risk management of operations for processes that leverage IA solutions requires a heightened focus on business continuity and contingency planning. This may be uniquely challenging if IA has been used to displace humans or process-driven activities because sufficient staff and contingency processes will be harder to implement on a short-term basis. Depending on how critical the IA process is, if unanticipated outcomes are experienced (e.g., underwriting anomalies), the overall provision of affected services could be impacted because cognitive tool decision making will need to be investigated and retrained.

The training process for machine learning algorithms often is not easy to understand or back solve, giving rise to the added risk that these decision processes need to be more actively inferred from outputs (and with challenger models). In addition, risk management oversight professionals will benefit from more active data/analytic techniques as well as traditional monitoring (e.g., volume, throughput, heat maps, etc.). In addition, programming and process enhancements should include a clear audit trail generated by the machine, exception-handling logic, processes to address exceptions that are "kicked out" of the IA process in a timely manner, and a feedback loop to incorporate these exceptions into the algorithms as adaptations over time.

Figure 1
Intelligent Automation Risk Management Framework

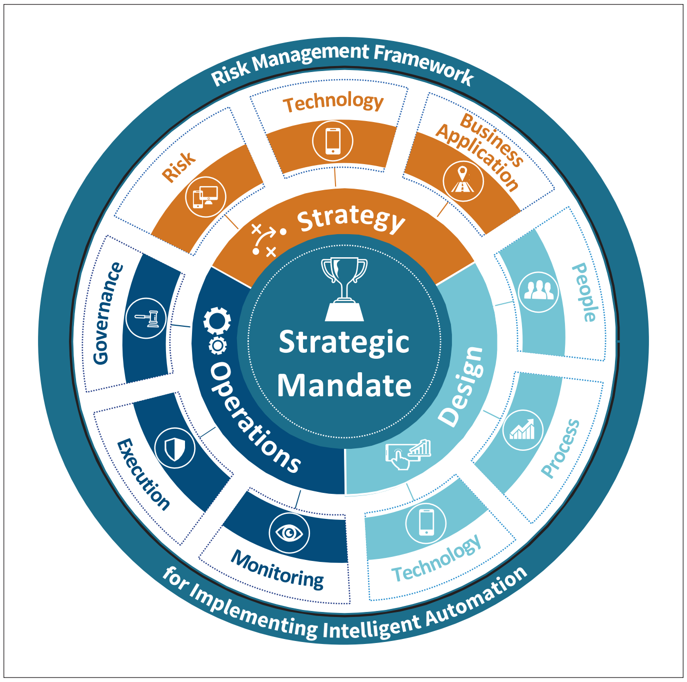

Monitoring approaches may also be candidates for intelligent automation (e.g., leveraging IA tools to perform independent validation, file reviews, and oversight of call center activities). Financial services firms that have not implemented Class 2 and 3 IA may find that existing review activities for second and third lines of defense are good candidates for "training" machine learning and cognitive tools because they are likely to have relatively greater amounts of data and be less time sensitive because the reviews and decisions already will have been made by humans (allowing the "right" responses to be known).

CAPABILITY CONSIDERATIONS

In adopting IA capabilities, executives should consider the current operational environment, governance, change management, resourcing, and integration with existing technologies. Adopting an evolutionary approach will lessen the risk inherent in technological disruption. Companies may wish to start IA implementation in lower-risk areas where results are easier to observe and verify. Firms may also wish to run IA and traditional processes in parallel and slowly transition from human- to IA-based processing. After a company gains experience with IA, implementation can progress to higher-risk, higher-reward areas.

For an organization to advance to Class 2 or enhanced process automation, technology teams should have the ability to analyze structured and unstructured data. Intelligent automation technology required for Class 2 should support a built-in knowledge repository, from which it can perform some elements of machine learning.

In adopting IA capabilities, executives should consider the current operational environment, governance, change management, resourcing, and integration with existing technologies.

Regardless of complexity, all IA technologies consume data to complete tasks in a more efficient manner. As organizations progress through the classes of automation and data becomes increasingly more important, so does the need for effective data management and governance. A model is only as good as the underlying data. It is important that roles, responsibilities, and ownership are clearly established related to data.

IA implementation should follow traditional model validation processes. The model must be clearly documented and independently reviewed and tested. Model documentation becomes more important as human touchpoints are removed from the process. A monitoring function is required to review IA results and ensure that the model is operating as intended (in addition to the exception-handling process referenced previously). For example, the monitoring function could analyze input data to evaluate whether new patterns or conditions are prevalent in this data that was not anticipated in model development or training. In addition, an automation Center of Excellence can serve as a central point of contact for organizations to share knowledge and best practices.

> In adopting IA capabilities, executives should consider the current operational environment, governance, change management, resourcing, and integration with existing technologies.

Companies must be equipped with data scientists who will have the ability to train and evaluate the model, as well as transform the data as the model evolves. In instances where anomalies are removed or data is modified to enhance the outcome and accuracy of the model, documented approval and justification for why this has occurred should be in place.

Risk managers of the future will need to use more sophisticated data analytics to monitor artificial intelligence and have direct involvement with process owners to do root-cause analyses of issues. Risk managers will need to understand the implications of their models and be agile enough to respond to model corrections, understand the output, and evaluate risk of the model as it evolves over time.

ETHICAL RISK IMPLICATIONS

Modelers have been building statistical models used for predicting outcomes for decades. So why is artificial intelligence different? What are the ethical implications that need to be considered?

Artificial intelligence models need to consider the availability of historical and current data, be able to identify and correlate patterns in data, and be able to predict complex outcomes based on the same indicators as the human brain. Humans have inherent biases, however, so how is it possible to build a model that thinks like a human without the societal bias? And how does the model determine what bias is considered good within the appropriate context?

> The projections out of algorithms are only as good as the data entered into the system. If the data is skewed or biased, then a destructive feedback loop can ensue, only worsening with time.

Because machine learning in itself is theoretically unbiased, the designers of the model need to be explicit and thoughtful about the design to help ensure that unintended bias is not created from unanticipated sources (e.g., data or flawed logic in the algorithm design). Think of machine learning in the context of a parent: Did you raise the child (build your model) well enough to ensure he or she has good morals (i.e., a low propensity for bias)? Poorly designed or managed machine learning models can have detrimental effects on individual stakeholders (e.g., through credit scoring or mortgage/loan decisions) as well as enterprises.

As machines continue to learn, they alter and develop their own algorithms so complex that the engineers who designed the system may not be able to identify the reasoning behind a single output. Therefore, the disconnection between humans and artificial intelligence opens up risks for predicting when failures might happen.[2] A model that is transparent—when the design of the model can be understood and the factors that attribute the outcome are known—allows the user of the model to understand what influences the outcome of the model.

To help improve the accuracy and integrity of IA-driven decisions or predictions, organizations may want to consider implementing feedback loops. This process allows for better monitoring of conclusions reached by the algorithm against factual data sets (expected outputs) to identify degradation of the model, which, in some cases, may require model retraining.

DATA IS THE NEW OIL

When companies use cognitive solutions, they will also need to recognize that "data is the new oil"—that is, data will be the most vital component of a cognitive model—and companies will need to evaluate whether the company has appropriate historical data to feed the cognitive algorithms. Organizations will be challenged with evaluating whether competitors have better data or more accurate data sources than they do.

The projections out of algorithms are only as good as the data entered into the system. If the data is skewed or biased, then a destructive feedback loop can ensue, only worsening with time. Because cognitive systems learn from patterns,[3] it is detrimental if they do not identify errors early. Therefore, when exposing a system to data, there must be a balance between the overfitting and underfitting of data. Data that can be directly attributed to the model outcome should be used where possible. Where proxy data is utilized, or data that indirectly is correlated to the outcome, this data should be understood and evaluated for its influence on the outcome over time.

In order to train algorithms, enough training data must be available. The more data variables that can be evaluated, the better the overall model. However, with every new dimension added to the model, the more computational power and storage is needed. As this computational volume increases, the available data to support the validity of the model decreases.

REGULATORY OVERSIGHT

Regulatory oversight of financial services firms (particularly oversight of risk management processes) will need to evolve with the increasing use of intelligent automation, particularly with Class 2 and 3 tools. A particular challenge will be in regulation and supervision that is designed to combat human bias in sales practices, extensions of credit, and similar financial decisioning for retail customers. Although it will not be acceptable for financial services executives to just say "the computer made the decision," supervisors will need to adapt oversight techniques and approaches to combat intentional (or directly embedded) bias in IA decision making—and not assume outcomes result from programmed bias. While correlation does not mean causation (or design, in this case) and supervisors should focus attention on intent/design, there will still be a need for risk managers to prevent unintended bias and to detect issues based on outcomes.

Intelligent Automation technologies create opportunities for improved efficiency and effectiveness in financial services firms, but they can also create risks that need to be managed. By expanding existing data and model risk management techniques as part of a comprehensive IA risk framework, companies may benefit greatly from these new technologies, while managing their risk. IA is here to stay. Let's get the greatest net benefit from it!

OVERVIEW OF IA

Intelligent automation solutions can be broken down into three classes:

1. Basic process automation
2. Enhanced process automation
3. Machine learning/cognitive automation

Basic process automation (Class 1) addresses transactional work activities that are rules based and primarily repetitive in nature and typically completed in existing IT applications. This includes screen scraping, macros, incorporating workflows, and basic design capabilities. This is the simplest form of IA, where macro-based applets synthesize structured data to complete a noncomplex, limited judgment task or job function. Class 1 automation is used where there is no ambiguity in the processes and uses structured and standardized input data. Common types of basic process automation include robotic process automation (RPA) and screen scraping.

> **Example usage**: *Systematic form population or bank account reconciliations*

Enhanced process automation (Class 2) enables the recognition of unstructured data and aids in adapting to the business environment. It builds upon basic process automation by incorporating a knowledge base and repository (RPA with the addition of a simple script/API add-on). The knowledge base is an important part of Class 2 automation, which allows the script and other capabilities to handle minor variations in input (e.g., date, address, business acronym). Such scripts can structure subprocesses or manual work that is not fully incorporated into the IT applications. Additionally, a key role of historical data

Figure 2
Classes of Intelligent Automation

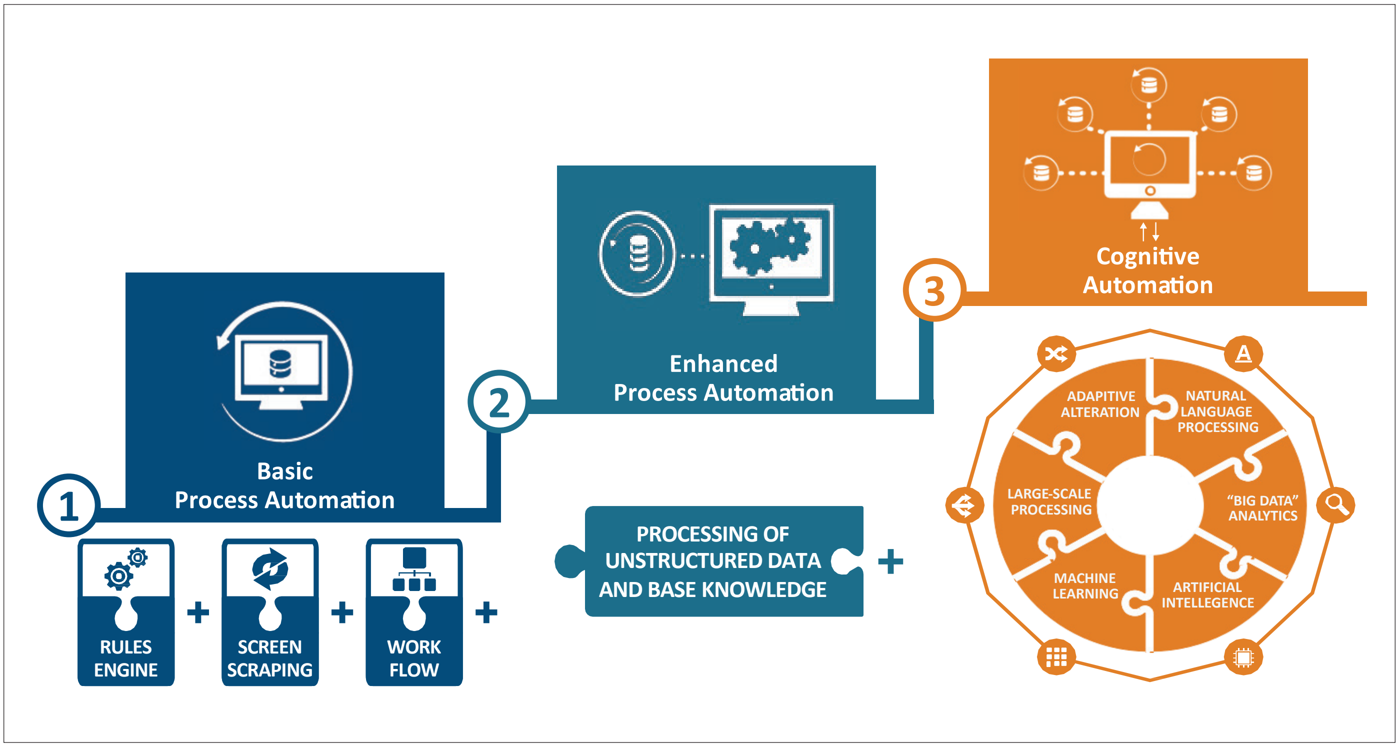

includes use in performance evaluations. Class 2 automation requires moderate to heavy involvement from business users to structure requirements along with structuring rules to build computational algorithms and knowledge base.

> **Example usage**: *Level 1 sanctions screening or cash flow forecasting*

Cognitive automation (Class 3) enables decision support with the help of advanced algorithms. The evolution of these tools is generally linked with advances in artificial intelligence, natural language processing, big data analytics, and evidence-based learning (machine learning). Machine learning is best defined as the ability of computer systems to learn and improve performance by exposure to data without explicit programming. Computer systems observe and recognize patterns, save the patterns in a knowledge repository, and later build on patterns to make predictions and offer solutions. Cognitive automation is the most advanced type of automation and can be used to automate tasks that require a relatively high level of human judgment. Cognitive technologies have the ability to mimic human reasoning and adapt as they self-learn. Cognitive solutions combine natural language processing, big data and predictive analytics, machine learning, and artificial intelligence. Class 3 automation is probabilistic and does not require business users to structure algorithms or logic; instead, models are typically "trained" by leveraging historical data. Additionally, key business users play a big role with evaluating model performance and enhancements. Further, historical data is used in model building and performance evaluations. (Note: It is important to split training and testing data in order to avoid overfitting.)

> **Example usage**: *Level 2 sanctions screening, email classification automation, or cash positioning and investments* ■

Disclaimer: Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates and related entities. The information contained herein is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavor to provide accurate and timely information, there can be no guarantee that such information is accurate as of the date it is received or that it will continue to be accurate in the future. No one should act upon such information without appropriate professional advice after a thorough examination of the particular situation.

Rob Ceske is a principal at KPMG LLP. He can be contacted at *rceske@kpmg.com*.

Kelly Combs is a director at KPMG LLP. She can be contacted at *kcombs@kpmg.com*.

Nadim Hraibi is a director at KPMG LLP. He can be contacted at *nhraibi@kpmg.com*.

Jamie Hooten is an associate at KPMG LLP. He can be contacted at *jhooten@kpmg.com*.

ENDNOTES

[1] The authors would like to gratefully acknowledge Junghoon Woo, Steve Gaeta, and Mayuresh Kulkarni for their valuable assistance.

[2] Will Knight, "The Dark Secret at the Heart of AI," *MIT Technology Review*, April 11, 2017.

[3] Cathy O'Neil, "When Not to Trust the Algorithm," *Harvard Business Review*, Oct. 6, 2016.