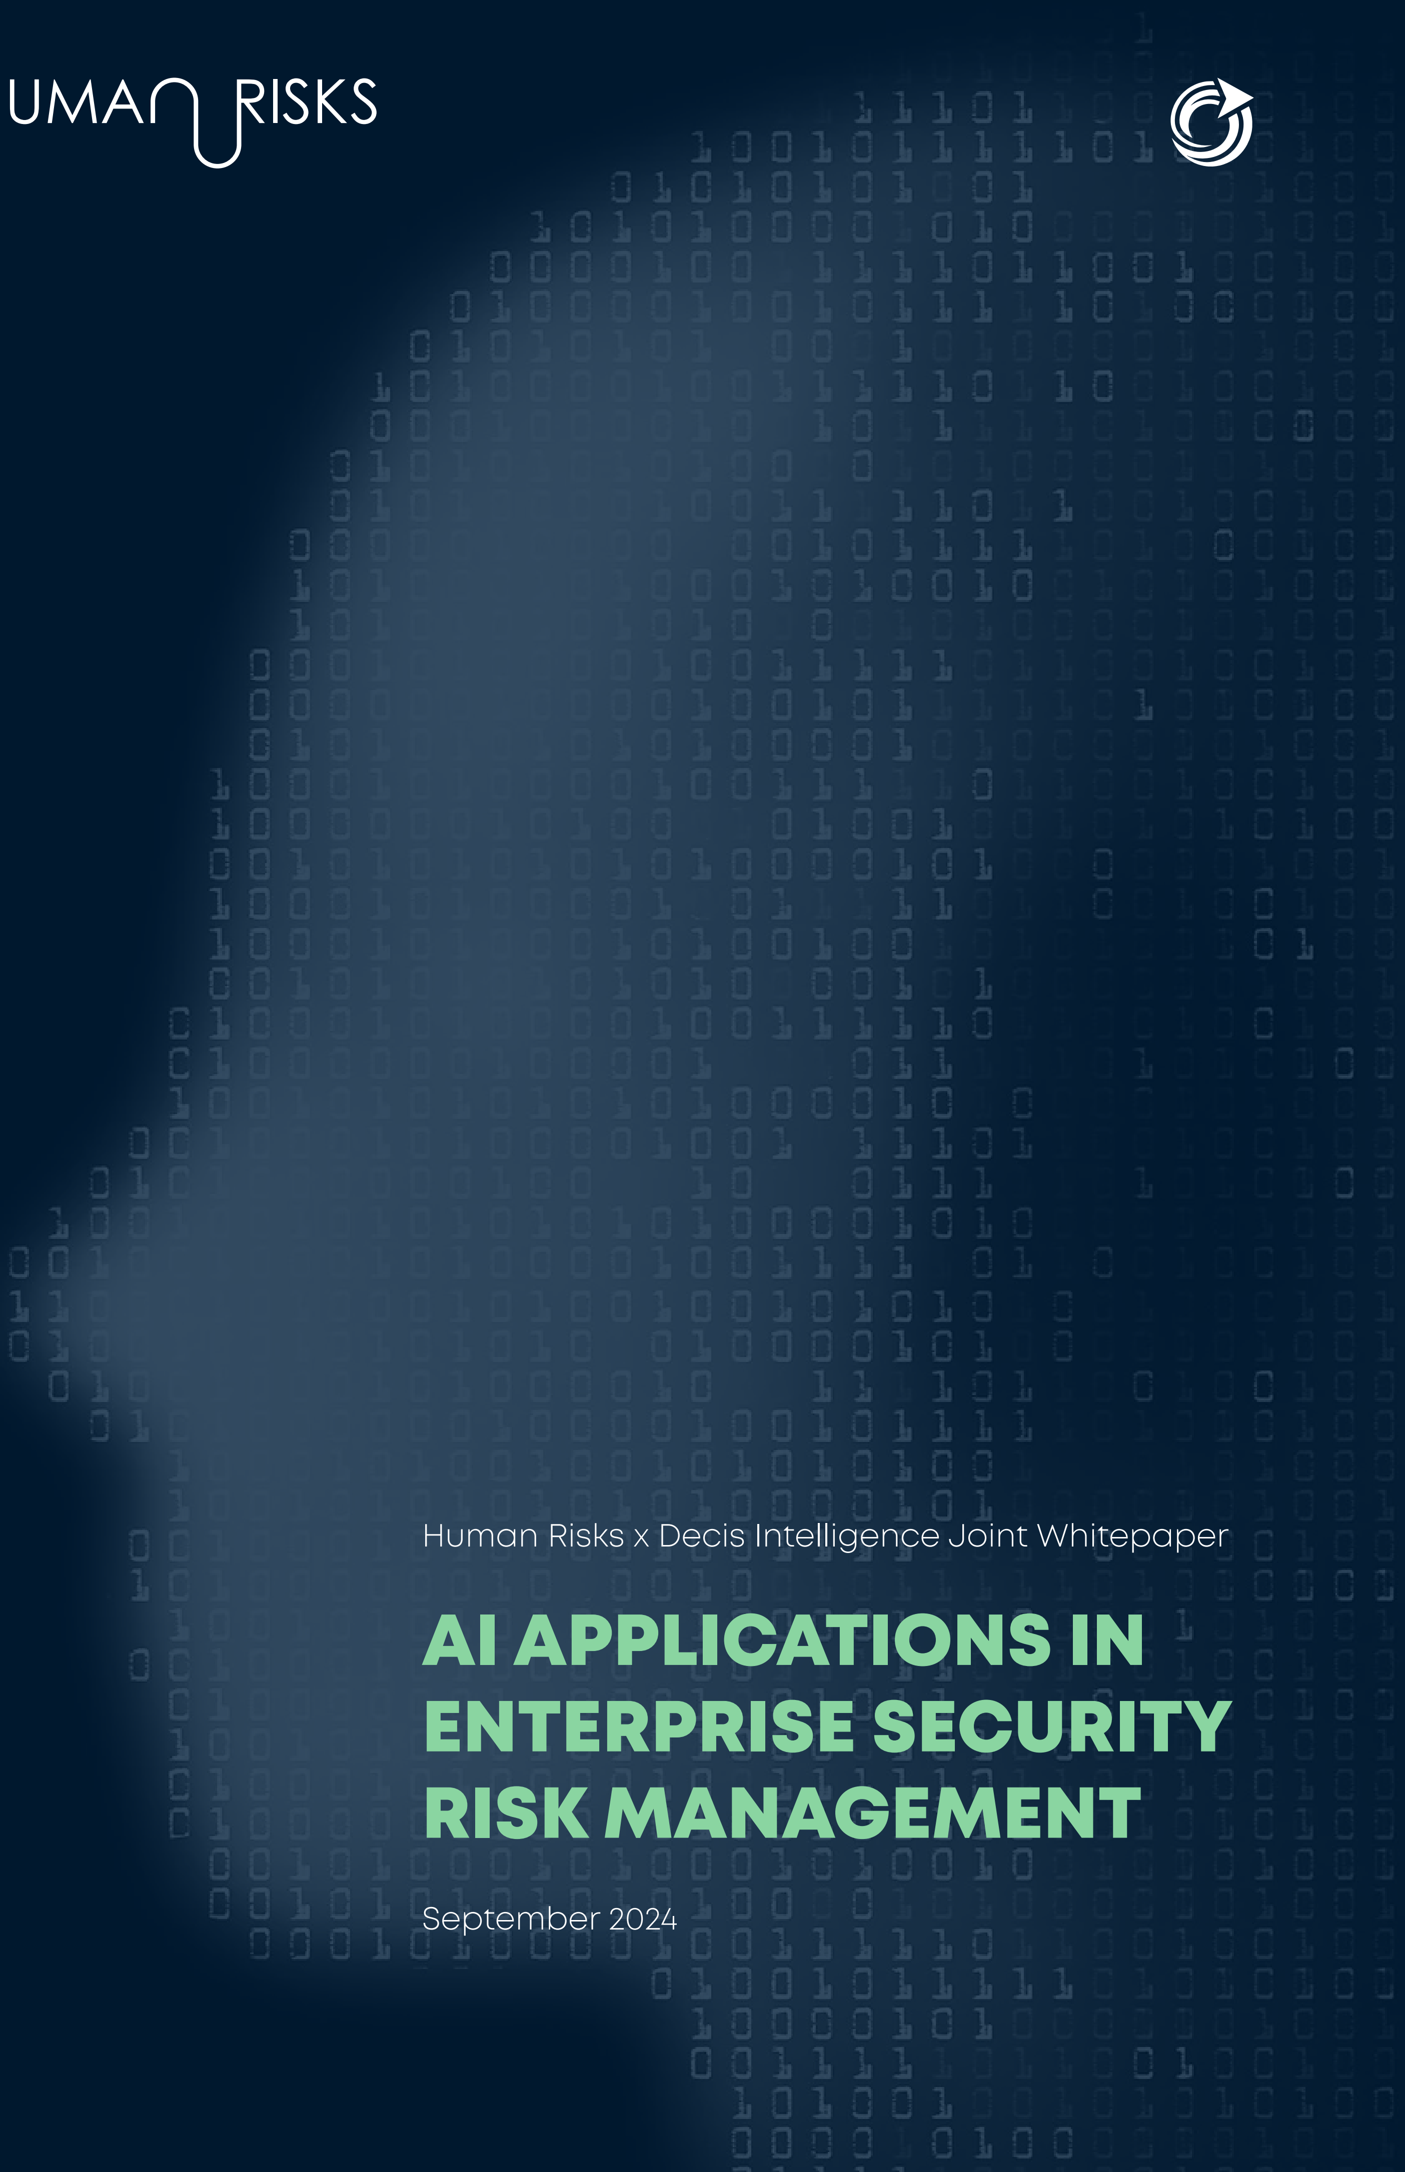

Human Risks x Decis Intelligence Joint Whitepaper

AI APPLICATIONS IN ENTERPRISE SECURITY RISK MANAGEMENT

September 2024

Executive Summary

Following the release of highly accessible AI Large Language Models in late 2022, we have witnessed the widespread and rapid global adoption of AI tools for both personal and workplace applications.

As sufficient time has now passed to allow the impact of these advancements to be looked at and practical benefits to be identified in more detail, this paper seeks to examine where these identified benefits are likely to play a significant role in a sector crucial to organizations globally: Security Risk Management.

The analysis begins with an examination of the benefits of AI in the broader workplace via a review of several well-regarded studies. Next, we compare these benefits to widespread challenges identified by leaders in the security risk management industry, allowing us to contextualize applications of AI tools specifically for security risk management teams. In particular, we focus on teams seeking to improve the efficiency and efficacy of core management processes and procedures with innovative solutions.

Included is a detailed breakdown of AI applications across the Enterprise Security Risk Management process as defined by ASIS International, to clearly identify areas where Security Risk Leaders can consider implementations across their workflows, overcome strategic challenges and drive a more effective security posture across their organizations. Based on this assessment, we conclude that the efficiency and time-saving benefits of AI-augmented processes, in addition to the additional strategic insights AI can support, are well placed to be levered by Security Risk Leaders in order to elevate the strategic value of security activities.

While comprehensive, we stress that the findings included in this paper are intended to be a practical resource for security risk teams and are therefore non-technical, based on both the citations below and observations we have made in our own work.

AI has significant potential to benefit the security risk management sector. It is the opinion of the authors that the industry should move towards adopting these tools with a combination of speed, caution and empathy. Noting that there remains to be security, confidentiality and staff welfare issues that require ongoing consideration.

Key Takeaways From This Paper

General Benefits of AI

1. AI Tools can support significant efficiency improvements across routine tasks
2. Deployment of AI tools in the workplace aids non-expert employees more
3. AI models do however continue to perform poorly when pushed beyond their limits

Applicable Security Management Challenges

1. Security leaders are facing an increasingly complex operating environment
2. Security risk teams deploy a wide range of tools with complex outputs as part of their day-to-day workflows
3. Many security teams continue to be seen as tactical functions within their organizations

Applied AI Benefits for Security Leaders

1. AI tools are well-placed to augment, enhance, and improve security teams' workflows at scale
2. AI pattern recognition, filtering and data synthesis capabilities will support a significant uplift in proactive insights across complex threat environments
3. Efficiency improvements from AI Tools can support leaders in elevating the profile of security within their organizations from tactical to strategic

Contents

Executive Summary Page 1

About the Authors Page 3

Literature Review Page 4

Key Takeaways | AI In The Workplace

Key Takeaways | State of Security Risk Management

AI in Enterprise Security Risk Management Page 7

General Applications

Specific Capabilities

Applied Use Cases | Matrix

Higher Level Benefits Page 13

Development of an Effective AI Roadmap Page 14

Key Considerations for Security Leaders

The Human Factor | AI and Expertise Page 15

Conclusions Page 16

About the Authors

Andrew Sheves is a seasoned expert in global risk, crisis, and security management. He is a life-long advocate for leveraging technology to enhance organizational resilience and has developed numerous tools to support this goal.

He currently runs Decis Intelligence, a geopolitical intelligence firm utilizing advanced AI and Machine Learning to speed up and simplify the decision-making cycle.

Based in Washington DC, Andrew began his career in the British Army, holds an MSc in Risk, Crisis, and Disaster Management from Leicester University, the IRM's SIRM designation, and is a member of ASIS.

Douglas Gray is a risk and resilience advisor with experience leading best-practice enterprise risk, resilience, crisis and security management projects across multiple sectors.

A strong advocate of a people-first approach to technology integration, he is a strategy lead on the Human Risks team, an Enterprise Security Risk Management platform focusing on making security risk management smarter.

Based in Paris, Douglas began his career in the New Zealand FMCG and Manufacturing Sector, and holds a MA in War Studies from King's College London alongside multiple industry affiliations.

Special Thanks

In addition to inputs from our colleagues across the security and technology sectors, we would like to extend our thanks to all those who reviewed and offered suggestions on drafts of this paper.

In particular, we would like to give special thanks to Steven Eames of Sigurine, Greg Hutchins of CERM Academy and Quality + Engineering, Andrew Tollinton of SIRV, Ryan Schonfeld of Hivewatch, and Jason Wright of Studitu.io and Google.

Any remaining errors, omissions or inaccuracies are the responsibility of the authors, and a reflection of the fast-paced environment around AI applications in the security risk sector.

Literature Review

As a foundation for our assessment of AI applications across the Security Risk Management Sector, we reviewed several recent studies investigating the impacts of AI integrations in the workplace. We also reviewed the findings of ASIS's recent 2024 State of Security Risk Management Report.

The intent of the review was threefold: first, to identify the most relevant tasks and workflows where AI can enhance workplace efficiencies; second, to identify tasks and workflows where AI applications have been observed to be less suitable for the workplace; and third, to identify common concerns and challenges faced by security risk managers.

Each paper was published or updated in early to mid-2024 and therefore reflects the efficacy of commonly available AI solutions at that time (circa ChatGPT 4 and Claude 2).

The principal papers reviewed were

Mollick | 'Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality', *Dell'Acqua, McFowland III, Mollick, et al.*, Wharton and Harvard Business School 2024

McKinsey | 'The Economic Potential of Generative AI: The Next Productivity Frontier', *Michael Chui, Roger Roberts, Lareina Yee, Eric Hazan, Alex Singla, Kate Smaje, Alex Sukharevsky, Rodney Zemmel,* McKinsey 2024

ASIS | 'The Current State of Security Risk Management: Benchmarks and Effectiveness Measures', *Gigi Agassini CPP, Mark Ashford, Diana Concannon, PCI, Rhys Robinson, Scott Wolford, CPP, Matt Jones,* ASIS International 2024

GitHub | 'Quantifying GitHub Copilot's Impact on Developer Productivity and Happiness', *Eirini Kalliamvakou and the GitHub Next team,* GitHub 2024

(N.B. The preceding name is the lead author or publisher and will be used throughout this paper)

Key Takeaways | AI in the Workplace

1

AI Tools Improve the Efficiency of Routine Tasks

All studies that measured the effectiveness of AI tools found that workers were able to complete routine tasks between 25% (Mollick) and 55% (GitHub) more efficiently.

In the future, AIs are expected to be able to complete an established 60% to 70% of routine process tasks via automations (McKinsey).

2

AI Tools Benefit Lower Skilled Workers More

All studies found that lower skilled workers benefitted from AI integrations more than their higher skilled counterparts, 43% to 17%.

These findings are consistent with other studies such as a 2023 US National Bureau of Economic Research (NBER) study, which found that *"access to the tool increases productivity, as measured by issues resolved per hour, by 14% on average, including a 34% improvement for novice and low-skilled workers but with minimal impact on experienced and highly skilled workers."*

3

AI Models Perform Poorly When Pushed Beyond Their Core Capabilities

The 'jagged frontier' as referred to by Mollick et al. highlights the point where an AI exceeds its capabilities and fails to work effectively, often performing far worse than human counterparts. There is no clear division of this point across tasks – hence the jagged nature of the frontier – but in general terms, AI models studied performed well at routine, codified, well documented tasks (e.g. writing computer code) but performed poorly at highly creative problem-solving tasks. Note this assessment is of model performance alone, not of human-model collaboration as is described in the other studies.

4

AI Tools Improve Creative Problem Solving

The study by Mollick et al. of consultants at Boston Consulting Group found that those equipped with AI tools were 12% more productive in completing creative tasks.

5

AI Tools Make Routine Tasks More Tolerable and Reduce Fatigue

The GitHub study specifically surveyed user satisfaction and happiness finding that:

- 60% to 75% repored an increase in job satisfaction using an AI copilot.
- 73% were able to maintain a 'flow' state often associated with developer happiness.
- 87% were less fatigued by the routine work when using an AI copilot.

Key Takeaways | State of Security Risk Management

1

Security Risk Managers Face an Increasingly Complex Environment

No single event type dominates the threat environment faced by those surveyed. Eight of the 12 common security threats surveyed with industry professionals fell into the 'posing most risk' category, reflecting a broad range of threats, compared to a narrow-tail distribution where one of two events dominate the threat landscape.

Those surveyed noted that previously firms faced a narrower range of threats. Moreover, many of these threats lead to complex, blended events of significant complexity.

2

Security Risk Managers Utilize a Wide Range of Tools

Given a range of six threat modelling techniques or tools, the majority of respondents deployed each to at least some degree. This blended approach has benefits as each tool brings strengths, but the need to combine the outputs of different tools and systems likewise adds significant complexity to team workflows.

Respondents also noted that management of threat modelling outputs, such as risk control monitoring, has became increasingly complex.

3

Security Response Plans Remain Effective for Security-Focused Events

Response plans have a high success ratio for security focused events, including complex situations such as wars and natural disasters. Plans were assessed as less successful for 'blended' events such as compliance failures where security forms part of the response, rather than leading.

4

Many Security Risk Management Teams are Seen as Tactical Functions

43% of respondents to the ASIS study indicated their organization as viewing security as either entirely or mostly tactical versus the 29% who were seen as mostly or entirely strategic.

This perspective is not however associated with the influence security has in an organization. 73% of respondents rated security as important or very important to the organization's objectives.

Notably, all respondents reported that they spent much less time on strategic planning than they considered ideal, suggesting that the tactical elements of security risk management workloads absorb a disproportionately large share of team capacity.

Only 34% of respondents saw enterprise security risk management (ESRM) systems, where security is a much more strategic function within an organization – as essential. Although a majority of respondents indicated ESRM frameworks are beneficial.

(N.B. The ASIS Report contained fewer comparative metrics than the AI studies included above. The findings however remain a useful input on the relevancy of AI applications across the security sector)

AI in Enterprise Security Risk Management

General Applications

Bringing together the strategic and operational challenges faced by Security Risk Leaders across industries alongside the benefits from application of AI in the workplace, we are able to see a number of immediate practical applications for utilizing AI tools in the security risk management industry.

The immediate benefits of these applications can be broadly summarized as:

1. **Reducing repetitive workloads to increase team inefficiency and free up capacity for further value-adding activities.**
2. **Improving the value of proactive insights generated by security risk teams to support greater engagement in security objectives.**
3. **Decreasing learning curves for stakeholders to provide high-value inputs into the security risk management processes.**

These benefits already hint at the value of Security Risk Leaders seeking to integrate AI tools in the workflows of their teams. However, to further assess the return on investment from onboarding AI-enabled tools and systems, it is important to further map potential benefits against security-specific processes and functional challenges faced by security teams.

Specific Capabilities

To support more effective analysis on the value from investment in AI tools, we have further identified ten specific AI capabilities with potential applications to the security risk management industry in line with key challenges faced by security teams:

Specific Capabilities | AI in ESRM

1. Process Efficiency

Automation of resource-intensive and repetitive processes such as complex asset mapping and external threat risk intelligence analysis to speed up the flow of information across organizations and free up team capacity.

2. Standardization and Synthesis

Utilization of AI and machine learning to standardize data flows against consistent terminologies and rating libraries. The creation of consistent 'apples to apples' data feeds to support cross-functional reporting, analysis and decision making.

3. Filtering

Machine learning pattern matching is highly effective at filtering for terms of importance beyond strict Boolean search, including advanced semantic search functionalities to identify and filter critical trends that a strict or manual pattern matching process would miss.

4. Pattern Recognition

Similar to filtering, AI semantic search functionalities can aid in recognizing patterns across often-disparate data sets in large site risk portfolios, allowing for review of high volumes of data at a significantly greater speed than manual analysis.

5. Predictive Analysis

Deployment of AI for automation of threat analysis and intelligence gathering activities to speed up the generation of preventative insights across connected sites, and resulting response activities.

6. Content Generation

Utilization of LLM content generation capabilities to efficiently produce content inputs for draft reports, contextual guidance, real-time translation in support of stakeholder engagements such as training and exercising.

7. Summarization

Similar to content generation, use of LLM capabilities to summarize large reports and qualitative datasets in order to speed up processes for cross-functional reporting.

8. Proactive / Tailored Messaging

Combined use of standardized data sets, pattern recognition capabilities and LLM content generation to provide staff with focused insights and tailored proactive alerts that drive deeper engagement in security risk management objectives.

9. Brainstorming / Red Teaming

Deployment of secure AI chat bots as an alternative for otherwise scarce or unavailable strategic thought partners to security risk teams and stakeholders providing inputs to strategic planning.

10. Data Processing

Deployment of AI models to manage and merge large streams of information (both internal and external) across site footprints into usable datasets.

Applied Use Cases Across the Enterprise Security Risk Management Process

The Enterprise Security Risk Management (ESRM) process, as described by ASIS International, is a strategic approach to security management integrating security practices with an organization's overall strategy using globally established risk management principles.

The primary goal of ESRM is to shift security from a reactive to a proactive stance by focusing on risk-based decisions, fostering constructive partnerships with asset owners and moving security from a necessary expense to a valuable business enabler.

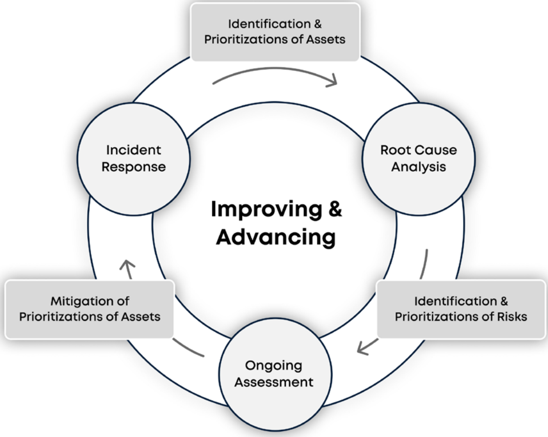
Identification & Prioritizations of Assets
Root Cause Analysis
Identification & Prioritizations of Risks
Ongoing Assessment
Mitigation of Prioritizations of Assets
Incident Response
Improving & Advancing

Fig 1: ESRM Continuous Process Flow (ASIS).

The lifecycle is structured around four key interconnected management processes:

1. Identifying and prioritizing assets.
2. Identifying and prioritizing risks.
3. Mitigating prioritized risks.
4. Underpinning security risk management as a QC function of continuous improvement.

By expanding on the core management activities under each of these four elements of the ESRM lifecycle, we are able to map specific management challenges and corresponding AI benefits across a typical end-to-end process flow. In doing so, we can likewise clearly identify areas where Security Risk Leaders should consider AI applications in their workflows to overcome strategic challenges and support a more effective security posture.

The below matrix provides an overview of the practical use cases for AI mapped against the key components of the ESRM Process and common problems identified by security risk teams. Supporting the matrix, is a high-level summary of the assessment against an amended ESRM Process Flow to enable quick identification of potential benefits.

To further support decision making around opportunities to implement AI tools within organizations, required data flows to enable AI and Machine Learning and ongoing security manager inputs have likewise been included in the mapping. Both of these components highlight critical considerations for any leader looking to develop an AI roadmap in order to innovate their teams' processes in the future. Namely, that the success of any AI implementation project is underpinned by high quality data, and that applied expertise will continue to play a crucial role in effective risk management processes – albeit with changing competencies.

Key Points To Interpret This Assessment

1. The below breakdown of applications is intended to provide a practical, illustrative assessment of AI capabilities to inform decision making and is therefore non-technical.

2. While detailed, the analysis is likewise not intended to be exhaustive. Similar assessments have been completed illustrating opportunities to apply AI across universal risk frameworks such as ISO 31000 or COSO Risk Management Systems, giving us further confidence on the strategic benefits for security leaders.

3. The roadmap to integration of AI tools can take multiple forms, including in-house development where capacity and capability allows, or through working with partners to leverage innovation specific to the security risk management applications.

4. No matter the appetite or pathway to implementation, the most critical component of any AI roadmap is a clearly defined problem statement specific to organizational context, including existing data quality and wider tools available across the organization to leverage AI capabilities.

Applied AI Use Cases | AI Across the Enterprise Security Risk Management Process

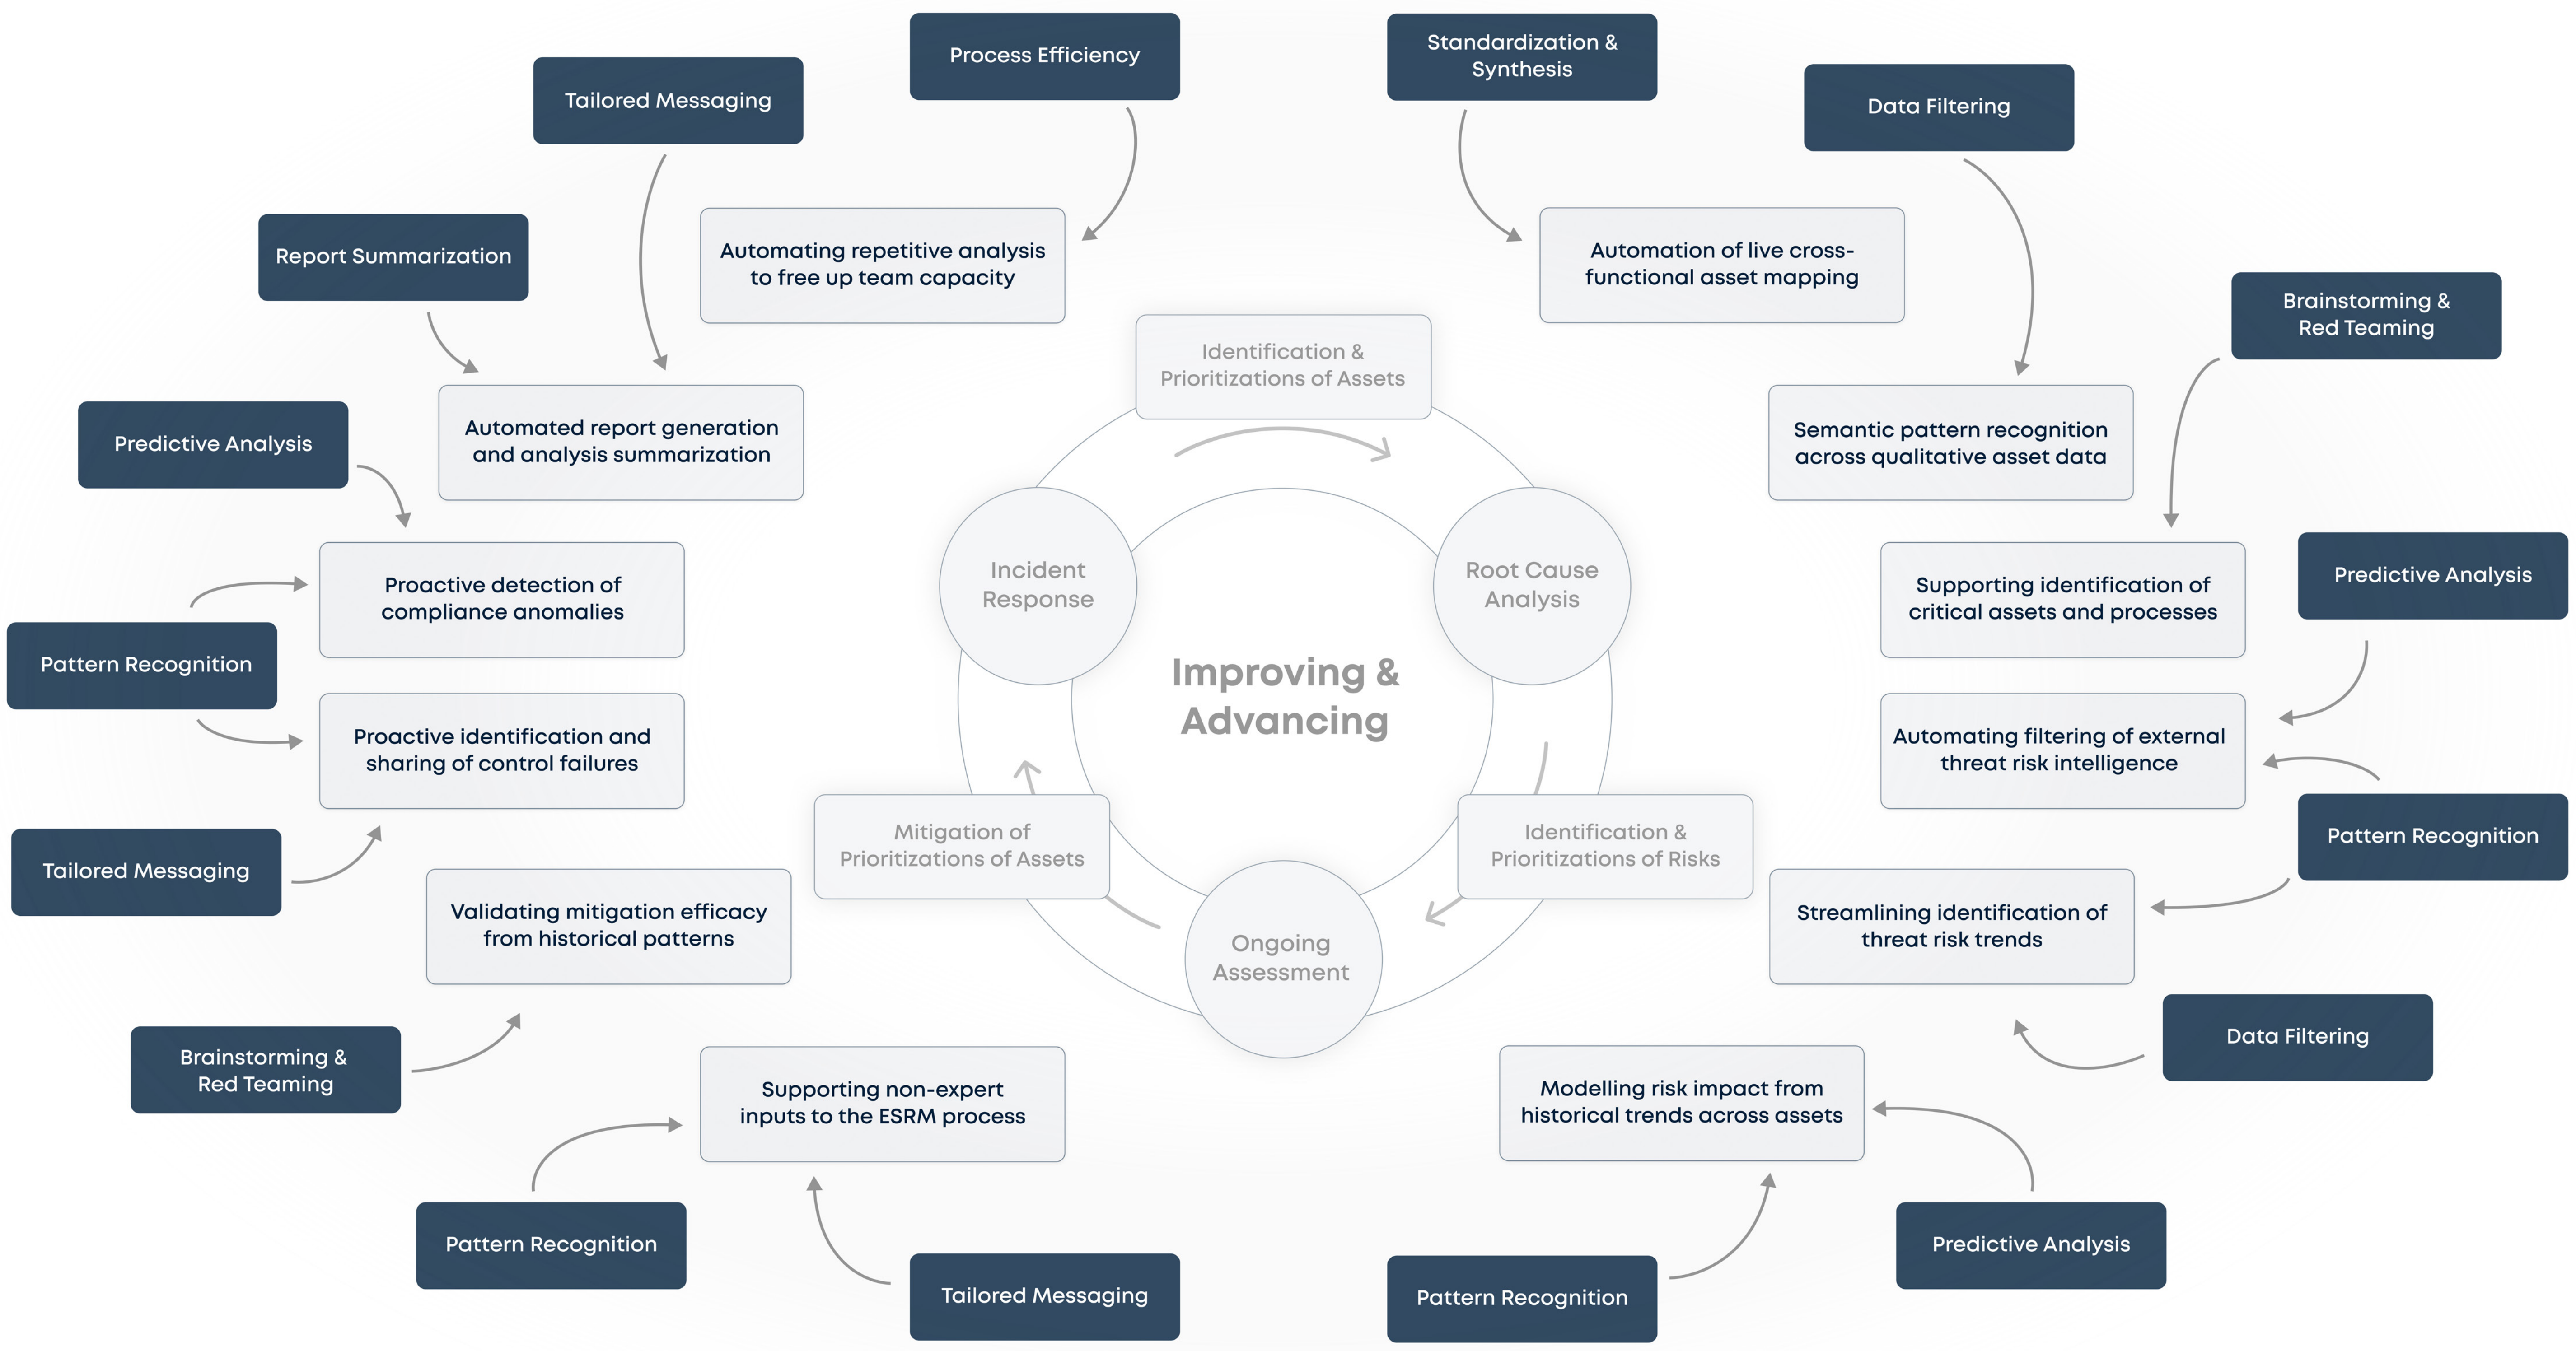

Improving & Advancing

Identification & Prioritizations of Assets

Root Cause Analysis

Identification & Prioritizations of Risks

Ongoing Assessment

Mitigation of Prioritizations of Assets

Incident Response

Process Efficiency

Automating repetitive analysis to free up team capacity

Standardization & Synthesis

Automation of live cross-functional asset mapping

Data Filtering

Semantic pattern recognition across qualitative asset data

Brainstorming & Red Teaming

Supporting identification of critical assets and processes

Predictive Analysis

Automating filtering of external threat risk intelligence

Pattern Recognition

Streamlining identification of threat risk trends

Data Filtering

Modelling risk impact from historical trends across assets

Predictive Analysis

Pattern Recognition

Supporting non-expert inputs to the ESRM process

Tailored Messaging

Pattern Recognition

Validating mitigation efficacy from historical patterns

Brainstorming & Red Teaming

Tailored Messaging

Proactive identification and sharing of control failures

Pattern Recognition

Proactive detection of compliance anomalies

Predictive Analysis

Automated report generation and analysis summarization

Report Summarization

Tailored Messaging

Fig 2: Applied AI Capability Across the ESRM Continuous Process Flow (Adapted: ASIS).

Applied AI Use Cases | AI Across the Enterprise Security Risk Management Process

ESRM Lifecycle	Management Processes	Critical Data Inputs	Key Management Challenges	Applied AI Capability	Strategic Focus	Ongoing Security Manager Inputs
Identification & Prioritization of Assets	**Strategic & Operational Asset Mapping** **Contextual Analysis & Objectives Mapping** **Threat Intelligence Gathering** **Asset Categorization & Prioritization**	Site profile data (e.g. GIS data) Established taxonomy of site and asset categories Baseline business impact analysis Sub asset and site process mapping Critical supplier risk profiles Logged security insights	Reducing the inherent resource-intensive nature of maintaining up-to-date maps of interconnected assets and sub assets across large site profiles Increasing the efficiency of adding additional data points to asset prioritization (e.g. third-party supplier data) Maintaining real-time visibility of evolving asset profiles across large multi-site organizations Increasing the speed of knowledge sharing across siloed departments and functions	**Standardization & Synthesis |** Automating ongoing collation of relevant asset information across multiple data sources to enable baseline security analysis **Brainstorming |** Supporting on-site identification and mapping of critical (sub)assets and processes with AI guidance **Data Filtering |** Streamlining the filtering of critical asset information to inform prioritization **Pattern Recognition |** Supporting a shift to more advanced quantitative asset trend analysis using sematic search capabilities **Process Efficiency |** Automating ongoing maintenance of asset maps across large site profiles	Decreasing tactical workloads to free up capacity for further strategic analysis Standardizing asset categorization across risk functions and disparate datasets Driving more effective cross-functional insights from asset trend analysis and threat intelligence Reducing training requirements and language barriers for staff with automated process guidance	Designing asset prioritization methodologies Setting guidelines and assessment criteria Designing and continuously reviewing model filters LLM Model training Continuous validation of model outputs Creating lexicons to translate key terms between systems Redefining workflows to incorporate AI and training staff to integrate with new systems
Identification & Prioritization of Risks	**Threat Intelligence Gathering** **Root Cause Analysis & Threat Identification** **Impact & Probability Assessment** **Risk Prioritization & Acceptability Analysis**	Organizational risk taxonomy External threat intelligence feeds Raw sensor data Internal incident and near miss reports Historical incident and loss data indicating the impact of threats materializing	Standardizing risk datasets against a unified taxonomy to support cross-functional information sharing Reducing resource-intensive intelligence filtering workloads for security risk teams Automating highly repetitive threat analysis reporting tasks Increasing the speed of threat analysis reporting to inform proactive decision making Automating threat detection and protection from raw sensors (e.g. video surveillance) Identifying threat intelligence correlations that are not easily observable via manual analysis Improving identification of cross-functional impacts	**Standardization & Synthesis |** Standardizing live threat intelligence information across data feeds to consistent taxonomies **Content Generation & Summarization |** Automating the creation of real-time intelligence reports **Pattern Recognition |** Automating identification of historical trends, root causes and critical dependencies across large asset profiles **Data Filtering |** Streamlining the filtering threat risk intelligence to inform triggers for (re)assessment of key risks **Predictive Analysis |** Estimating the probability and impact of events based on historical patterns **Process Efficiency |** Automating repetitive tasks in the threat intelligence gathering process to support a shift to more proactive activities	Automating external threat intelligence collection and filtering against baseline organizational data sets Centralizing oversight and reducing siloed analysis of root causes and historical patterns Supporting effective stakeholder engagement on security objectives via robust proactive insights Automating suggestive risk analysis to support strategic and operational planning	Threat risk assessment methodology development Designing and continuously reviewing model filters Continuous validation of model outputs Security report and template development to embed AI content generation Intelligence trigger identification and validation Redefining workflows to incorporate AI and training staff to integrate with new systems

ESRM Lifecycle	Management Processes	Critical Data Inputs	Key Management Challenges	Applied AI Capability	Strategic Focus	Ongoing Security Manager Inputs
Mitigation of Prioritized Risks	**Risk Control Identification & Implementation** **Control Design & Efficacy Analysis** **Residual Risk Analysis** **Task & Corrective Action Management**	Organizational control taxonomy Design and operating efficacy data of mitigation measures Raw sensor data Regulatory measures and internal minimum control requirements	Reducing manual requirements to maintain living risk assessments against group asset profiles Achieving faster organizational reaction times to emerging threats and control failures before they materialize Driving more advanced threat detection and protection from raw sensors Supporting a shift from reactive to preventative mitigation measures Streamlining the capture of critical insights from incident response processes	**Pattern Recognition |** Validating the efficacy of mitigation procedures based on historical patterns **Predictive Analysis |** Proactively suggesting corrective actions and control review requirements based on incidents and mitigation failures across multi-site risk profiles **Brainstorming |** Supporting non-expert inputs into the ESRM process with automated guidance on mitigation measures based on historical data **Pattern Recognition |** Automated detection of anomalies in control efficacy and mitigation processes	Centralizing oversight and validation of mitigation measures / risk-based compliance reviews Automating suggestive guidance on mitigation efficacy analysis Supporting effective stakeholder engagement in strategic objectives via robust proactive insights Reducing training requirements and language barriers for staff with automated process guidance	Building and maintaining standardized mitigation libraries Validating model outputs and historic analysis and analyzing lessons learned Defining business as usual operating guidance across site risk profiles to support pattern recognition and trend analysis Redefining workflows to incorporate AI and training staff to integrate with new systems
Improving & Advancing	**Incident Response** **Ongoing Threat Intelligence Analysis** **Assurance & Compliance Monitoring** **Reporting & Consultation**	All data collected across the ESRM process Qualitative and quantitative analysis from security risk teams Continuous training inputs to feed aggregation models	Increasing the efficiency of report generation and contextualization of insights to inform decision making Proactively identifying and mapping emerging strategic changes in the external threat environment Identifying strategic trends that are not easily observable via manual analysis Reducing routine workloads from information sharing Supporting cross-functional data sharing across business units	**Pattern Recognition |** Automating suggestive triggers for re-assessment of security risk profiles based on strategic trends **Pattern Recognition |** Automated detection of anomalies in compliance processes **Content Generation & Summarization |** Automating report drafting to speed up the sharing of strategic insights **Customized Messaging |** Automating generation of tactical alerts and data sharing to support a shift to more value-add strategic analysis **Process Efficiency |** Decreasing training and personnel requirements to achieve effective ESRM implementation	Freeing up capacity for security risk teams to enable a greater shift towards strategic risk analysis Supporting effective stakeholder engagement on security objectives via robust proactive insights Driving improved scalability across the ESRM process to embed security procedures further into operational management	Managing the overall design and delivery of the Enterprise Security Risk Management Framework across the organization LLM Model training Expert validation of security risk profiles and mitigation in line with the shifting threat landscape Maintaining distribution lists and tailored briefing formats for specific stakeholders Redefining workflows to incorporate AI and training staff to integrate with new systems

Higher-Level Benefits | Tactical Improvements to Drive Strategic Activity

A lack of strategic influence across organizations remains a key concern for security risk leaders in the industry. At its core, this lack of strategic influence is driven by two factors. First is time allocation, as the majority of team capacity is spent on overly tactical activities. And second is the perception that security analysis does not add sufficient strategic value in the eyes of organizational leadership.

In integrating AI across security risk management processes, many of the principal benefits are likewise tactical in nature, freeing up time from mundane, repetitive, tactical tasks. Applied well, this supports solving the time-allocation issue, while the additional time and attention should allow Security Risk Managers to focus on strategic initiatives, elevating SRM to a more strategic, enterprise activity.

By returning to the higher-level benefits from the integration of AI tools in the workplace, as outlined in the applications section above, we can see three key areas emerging:

1 Improving the Value of Strategic & Proactive Security Insights

Deploying AI capabilities to help separate signal from noise and identify patterns in otherwise unstructured datasets will enable security risk management teams to more effectively monitor threat risk environments. Supporting analysis of risks in a more holistic, aggregate way. In turn, the resulting analysis can drive improved insights by identifying changes in the threat environment that would otherwise go unnoticed or unheard.

In better automating analysis of historical patterns, organizations can identify emerging threats and begin to take action much more quickly compared to existing manual processes. This gain in team efficiency translates into a material benefits for organizational responsiveness to emerging threats - and likewise supports lifting team thinking from the tactical to the strategic.

2 Reducing Repetitive Workloads

The outsourcing of routine processes to AI models will further free up significant capacity for security risk teams. These capacity savings will allow a focus on less acute but equally important strategic analysis activities – or allow for reduced capacity in non-strategic areas.

A practical example is the outsourcing of data collection, collation, and first-pass analysis of geopolitical news. This allows analysts to focus on strategic analysis as the centerpiece of their roles, more effectively utilizing their expertise. In turn, the resulting insights can provide decision-makers with greater actionable intelligence due to the increased quality of analysis. Moreover, this work can be done in a fraction of the time, accelerating decision-action cycles.

3 Driving More Effective Stakeholder Engagement

Effective partnerships with a wide network of stakeholders to achieve buy-in on the importance of embedding security in strategic decision making is core to the Enterprise Security Risk Management process. To this end, the deployment of AI is set to play a key role in driving better stakeholder engagement in security management objectives, both upwards by enhancing executive reporting and across organizations with more tailored, situation-specific insights.

Furthermore, AI tools can provide tailored security guidance and planning advice for non-specialist teams, and support the standardization of reporting outputs. This further reduces ongoing training requirements and improves the effectiveness of each touchpoint between central teams and on-site leaders (particularly in multi-lingual settings). These improvements will drive an increased return on investment from both AI and wider security program activities - in particular across large operating profiles.

Development of an Effective AI Roadmap

It is important to note that while this paper has highlighted many of the new and emerging benefits AI tools offer the Security Risk Management industry, we have not sought to provide detailed guidance on the development of an effective roadmap to integrate these tools. Extensive resources are readily available on best-practice AI project management, and contexts for the application of new technologies and systems will always be organization-specific.

There are however four universal AI project considerations for security leaders seeking to position themselves as innovative value drivers within their organizations:

Clearly defining use cases

Prior to any new integration project, it is essential to clearly identify the challenges being solved and/or additional value being sought from changes to existing processes.

Specific to AI, this extends to ensuring that those involved in the decision-making process have a good understanding of both the benefits and limitations of AI models, in addition to a strong understanding of specific use cases for security programs.

Identifying Data Requirements

As highlighted across the critical data inputs above, the success of any AI integration project relies on well-maintained datasets that allow models to learn and generate value from advanced analytical capabilities within a secure environment.

Full evaluation of data quality and management procedures must be included in the delivery of an integration project, and creation of a structured knowledge repository needs to be considered as early as possible.

Choosing the right partners | Both Internal & External

Capacity and resources to deliver on AI integration projects differ significantly between organizations. From the outset, effective partnerships between internal departments (e.g. Security and IT) are vital to address internal capability gaps and identifying the best mix of solutions. In many cases, a key consideration will be whether purchasing an off-the-shelf solution from a security solutions provider is a more effective path than building bespoke tools in-house.

Emerging Changes in Competencies for Security Risk Teams

As highlighted in the matrix above, core security risk management skills will continue to remain essential to delivering effective security frameworks, albeit with new competency requirements for AI-enabled teams.

Security risk leaders seeking to embrace AI need to proactively consider the emerging minimum competencies for modern teams seeking to make best use of AI tools – such as a strong understanding of model training and validation procedures – and ensure these are reflected in the target skills matrices for their future teams.

The Human Factor | AI & Expertise

Building on these themes of human-AI interaction, it is likewise important to address concerns that AI will replace humans in the workplace. These concerns are both valid and pressing, and need to be addressed by leaders at every level prior to implementing AI systems.

Low-skilled, repetitive tasks, such as basic text editing, simple data processing or producing rough document drafts are becoming increasingly easy to replace by autonomous processes including AI. This change has been underway for many years, but AI is undoubtedly accelerating the transformation.

However, while there may be cause for real concern in some cases, particularly for non-expert knowledge workers, there remain very few instances where the latest advancements in AI are replacing jobs altogether. As at the time of writing, the Swedish fintech company Klarna is the only firm which has publicly replaced a significant number of staff with AI agents, primarily in call centers.

As the 'jagged frontier' expands, the tasks and roles that AI tools can undertake will increase. Nevertheless, the dichotomy of human or AI in Security Risk Management remains unwarranted for three key reasons:

1 AI Models Still Have Narrow Training Libraries to Rely On

Many professions, including security risk management, lack a standard set of reference materials, relying instead on individual expertise. It will therefore remain much harder to train AI models to outperform the experience that an effective Security Risk Manager brings for years to come.

2 AI Models Will Not be Trusted as Decision Makers for Some Time

Running models without human validation or intervention for significant decisions is not foreseeable in the near future. This will continue to be particularly relevant in fields managing material impacts with respect to safety, security and personnel management.

3 Many Tasks Still Require Human Intuition and System Two Thinking

Intuitive system two thinking is an area where AI models will continue to struggle for some time. It therefore remains hard to foresee how or when an AI tool will have the capability to replicate the intuition of an experienced thought leader.

These ongoing capability gaps suggest that it will be some time before expertise-based security risk management tasks will be impacted by AI to a significant degree (noting that the pace of change may still accelerate).

However, where we do foresee significant change is where the efficiency and consistency of AI tools can be paired with the capabilities of experienced professionals. As noted above, integration of AI tools will require ongoing inputs from security risk teams to design, develop, train and interact with new models. This reinforces both the importance of existing expertise in the security domain and the need to upskill teams for Security Risk Management in the 21st century.

Therefore, instead of a future where the choice is AI or humans, we strongly believe that the best results will be where we pair AI and humans, blending AI capabilities with human experience and intuition.

Leveraging these combined capabilities ensures organizations get the best of both worlds: human expertise and mechanical efficiency. Meanwhile, it also enables security risk professionals to augment their own best practices and adapt to the emerging new normal in our workplaces, further securing their expertise for the future.

Conclusions

At the time of writing, we are almost two years into the new 'AI era' heralded by the release of widely available AI tools, and we have already seen both the widespread adoption and impacts of AI tools in the workplace.

Consistent with this trend, the advantages and use cases identified where AI has already benefited knowledge workers are analogous to the challenges faced by security risk leaders and, in many instances, specifically match the specific strategic issues associated with effective ESRM implementation. We are therefore comfortable concluding that there are substantial benefits in the rapid, albeit cautious, adoption of AI tools across the Security Risk Management industry.

Detailed benefits from adoption will always be specific to individual organizations. However, as our mapping of AI solutions across the ESRM cycle identifies, AI tools are well-placed to augment, enhance, and improve existing processes at scale. These improvements will allow Security Risk Leaders to run their teams more efficiently and, in many cases, free up capacity to allow them to elevate the profile of security within their organizations from tactical to strategic. In turn, these benefits are likely to also drive strategic advantage in the market for organizations that take better advantage of the strategic insights and initiatives the security risk teams can provide.

However, while the potential benefits from implementation are significant, caution and care will always be required for any major organizational change. The dangers of exceeding the capabilities of any AI (i.e. breaching Mollick's 'jagged frontier') remain significant. Moreover, the potential security and confidentiality challenges of exposing sensitive data to non-secure workflows are notable, and trusted partners will always be required in this regard. These reliability, security and confidentiality challenges warrant further in-depth analysis, both for individual teams and by industry leaders.

Finally, the rapid velocity of AI advancement means this paper is a snapshot in time. It is clear that we are very much at the beginning of the widespread adoption of technological advancements from AI, meaning that future systems may offer even greater advantages than those identified. Meanwhile, more advanced AI models will also introduce new threat vectors and challenges that will have to be addressed, warranting review and revision of this analysis in the future.

Despite these limitations however, we remain of the opinion that the benefits of AI adoption across the security risk management sector will far outweigh the disadvantages. AI presents a unique opportunity for security leaders to take the lead on innovation and drive further value across their organizations.

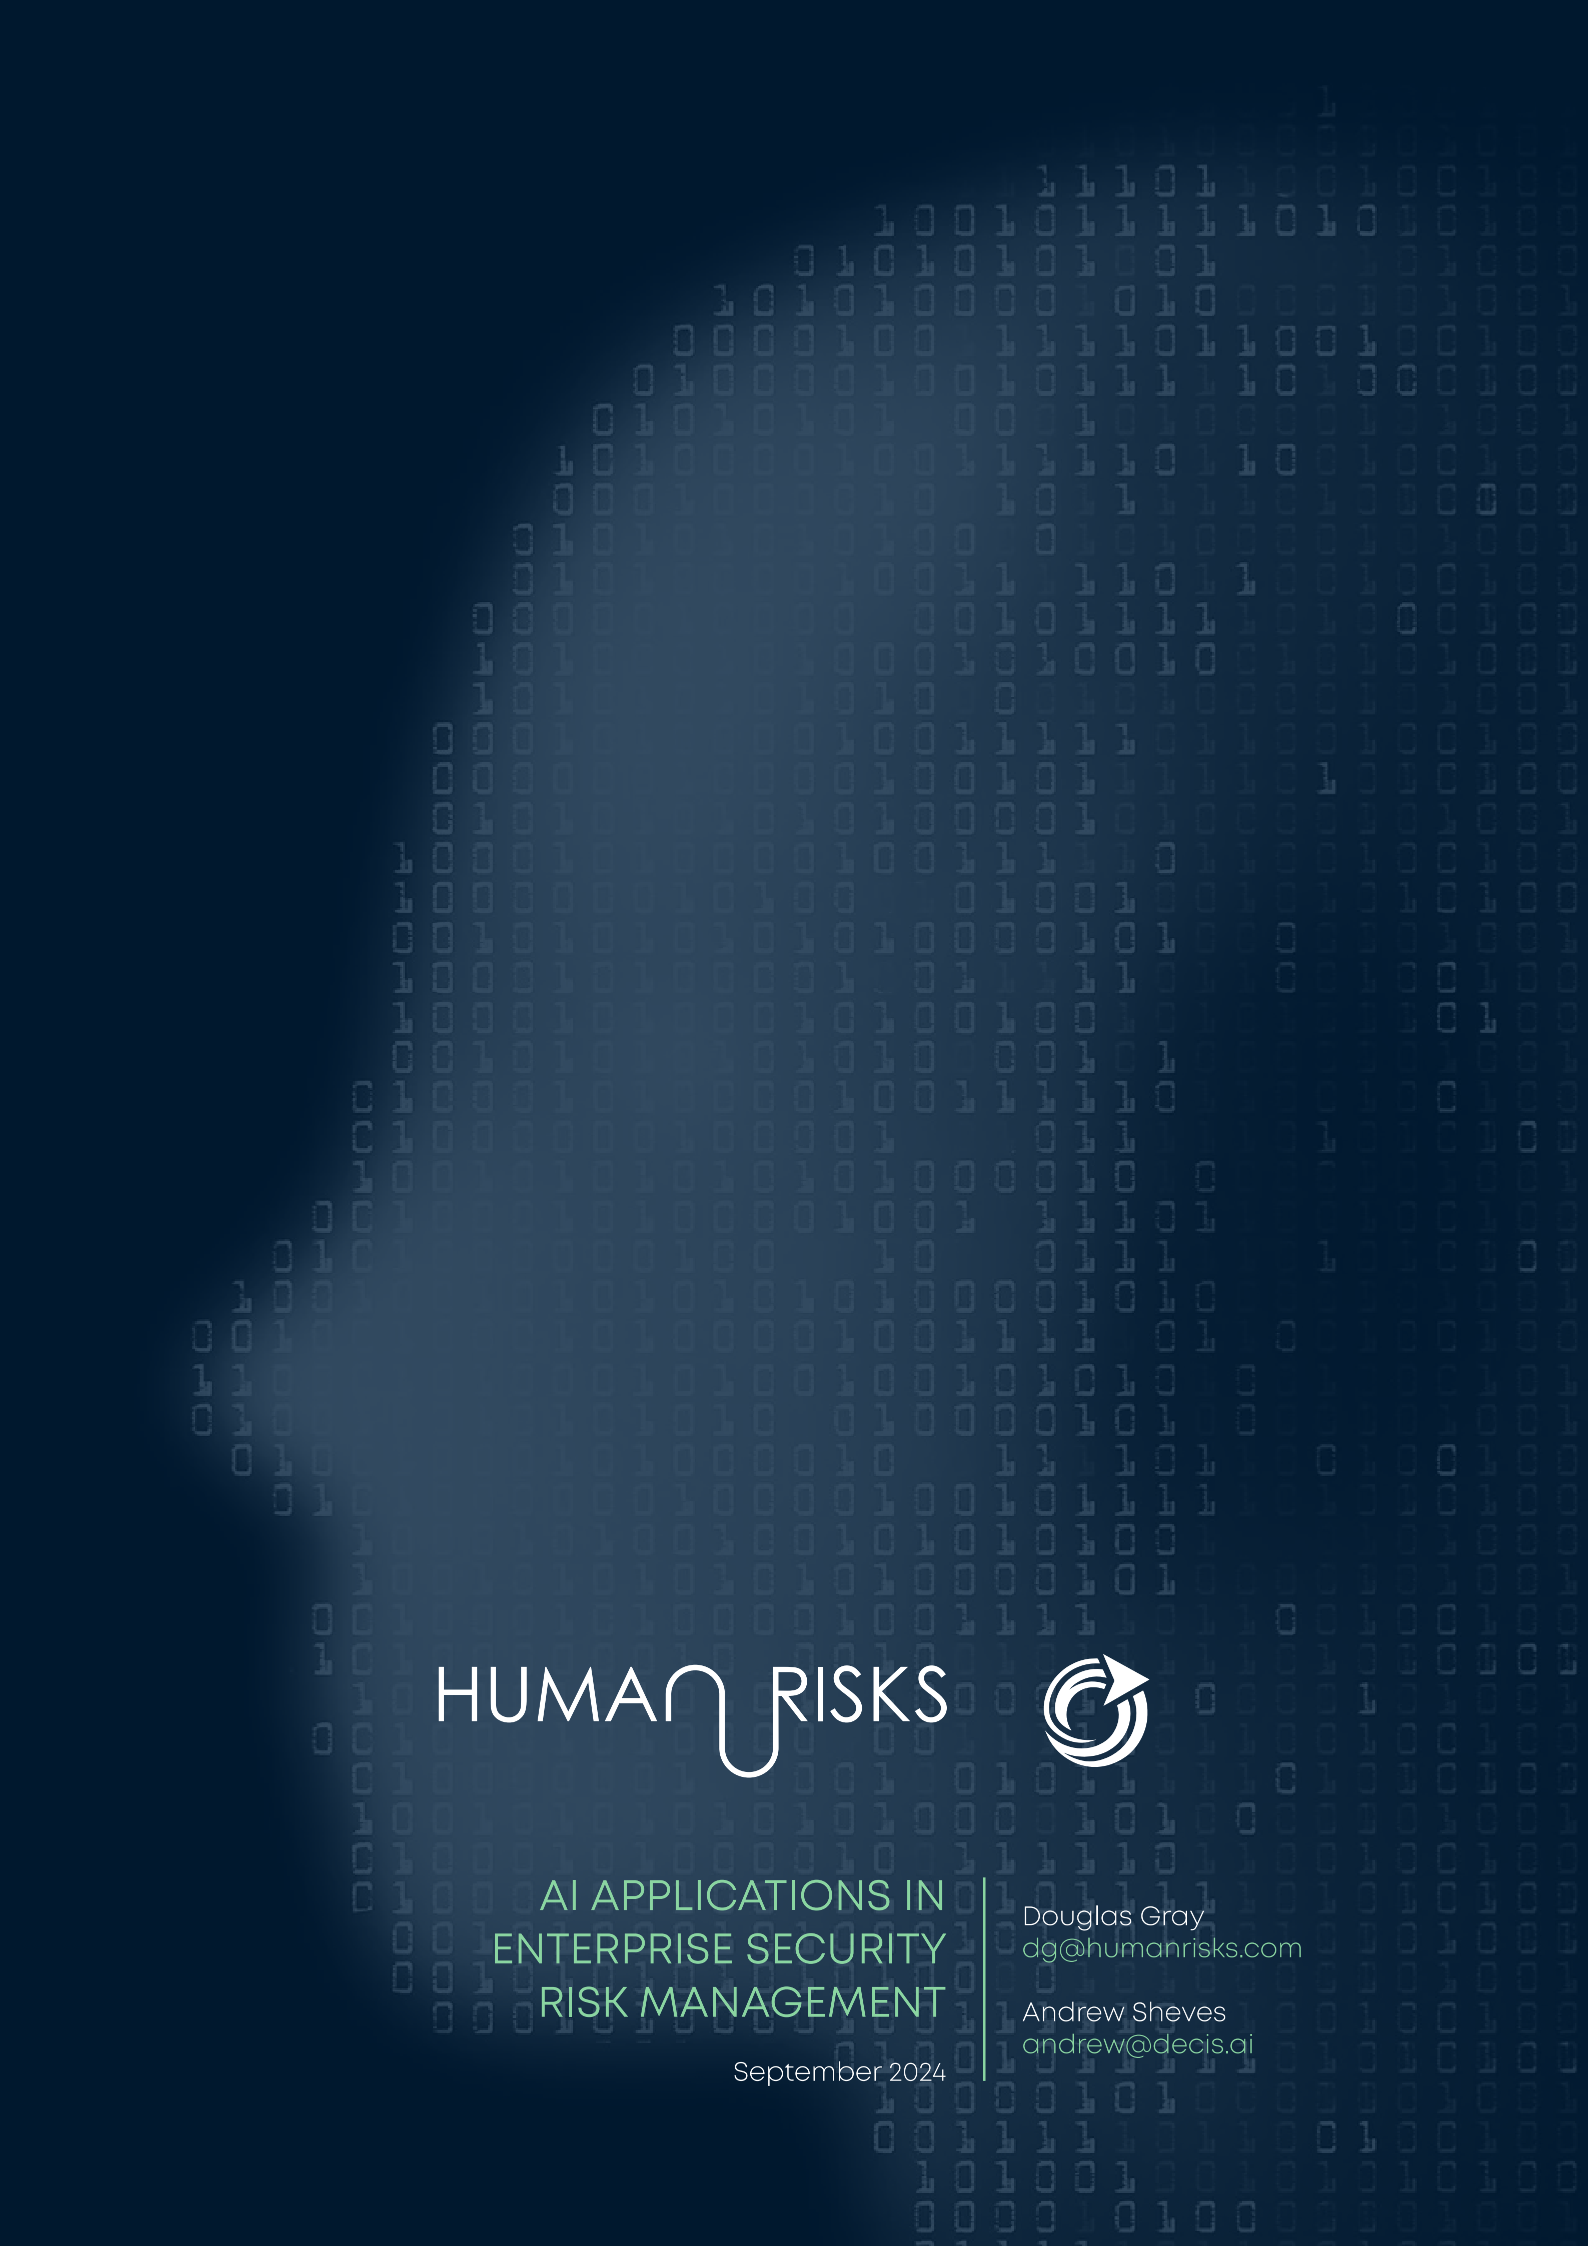

HUMAN RISKS

AI APPLICATIONS IN ENTERPRISE SECURITY RISK MANAGEMENT

September 2024

Douglas Gray
dg@humanrisks.com

Andrew Sheves
andrew@decis.ai