

Introducing the Generative AI Function in the KOS



Letter from our Founder & CEO

Consumer chatbots run on large language models trained on data scraped from the web have triggered a paradigm shift, instantly creating completely new types of risks and opportunities never before faced by CEOs.

These and other risks can be mitigated or eliminated, and opportunities captured, by employing a one-of-a-kind enterprise operating system designed to optimize the knowledge yield curve in the digital workplace.¹

Twenty-six years of R&D in knowledge engineering, data physics, and artificial intelligence has resulted in our KOS, which is the world's first and to my awareness only distributed enterprise AI OS. The KOS provides a simple-to-use natural language interface with strong system-wide governance.

Our digital assistant provides personalized learning, security, prevention, and productivity, all deliverable in a single, efficient, and cohesive system, which can be extended to partners and customers—enterprise and consumer.

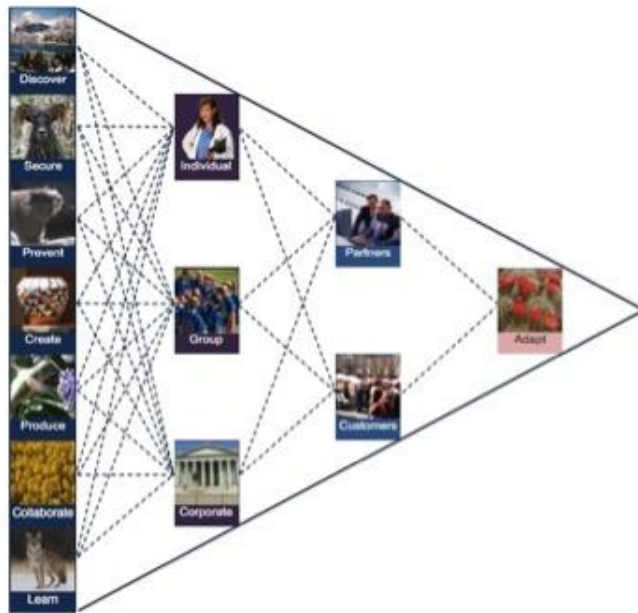
We are now offering generative AI as an additional function in the KOS. Early indications suggest productivity improvements can be achieved in the range of 15-40% across the enterprise.

Please email me to set up a time to discuss in more detail.

Sincerely,

Mark Montgomery
Founder & CEO
KYield, Inc.
markm@kyield.com





KOS: Inspired by Nature. Managed by Humans. Assisted by AI

Table of Contents

Letter from our Founder & CEO.....	1
Introduction.....	3
Generative AI: Necessary but Insufficient	4
Rules-based Governance	5
Metamorphic Transformation: GIGO to QIQO	6
Learning Networks.....	7
KOS Process.....	8
Affordability & ROI.....	9
Conclusion	10
References	11



Introduction

Although the average citizen wouldn't realize it by reading media headlines, or watching the inaugural U.S. Senate Judiciary Committee hearing on AI, AI systems consist of much more than large language models (LLMs).

Indeed, a variety of machine learning (ML) and neural network (NN) models have been applied to a great many use cases in recent years, ranging from voice chatbots to accelerating drug development. However, the concept of enterprise-wide AI systems is still relatively new, despite more than a decade of deep-dive discussions by KYield with leading organizations.

When we first approached the Department of Defense, major oil companies, banks, and pharmaceutical companies, among others, about a decade ago, all were familiar with AI for narrow use cases, but none had yet considered applying AI across the digital workplace. An enterprise AI operating system (EAI OS) seemed futuristic a decade ago.

Fast forward to today and LLM chatbots have been unleashed prematurely to the public, resulting in a historic level of adoption, hype, and massive funding.

However, LLMs run on data scraped from the web and other unreliable sources generate false information. Advances in LLMs notwithstanding, the problem with LLM chatbots is lack of data governance and provenance.

Unfortunately, NN models that scrape vast amounts of data from the web also include bits of data that can be very dangerous when processed and delivered in response to a prompt.² One of several plausible scenarios we've developed internally involves bioterrorism to develop a virus as infectious as COVID-19 and more deadly than smallpox—a virus that could risk half the world's population.

Using consumer LLM chatbots without benefit of a strong EAI OS is comparable to a body builder who takes steroids to pump up his upper body while ignoring his cardiovascular and neurological

systems, which are of course essential for achieving a long, healthy life.

We believe the optimal solution for the greater challenges facing CEOs in AI is the KOS—an EAI OS developed over a long period of rigorous R&D, which provides multiple functions in a single cohesive system, including strong, simple-to-use governance. The KOS was pre-optimized as an AI system from inception over two decades ago, guided by our 15 EAI management principles.

Every customer of our enterprise KOS receives a digital assistant called DANA for every employee that tailors precision data to each person within the parameters provided by the corporate admin in the CKO Engine.

In addition to GenAI functions, the KOS and DANA work together to provide personalized learning, multiple types of security, prevention of several types of crises, work-related networking, messaging, prescient search, and enhanced productivity—all in one highly efficient, easy to use, cohesive system.



GenAI Risks for Organizations

IP: Although chatbots offer an opt out function for their users to store data, few exercise that option. The risk of IP disclosed in data stores is high and may be impossible to remove.

Inaccuracy: AI models are only as accurate as the data they are trained on. Consumer LLMs are trained on vast data scraped from the web, resulting in error rates much higher than is acceptable for most organizations.

Commoditization: LLM chatbots were instantly commoditized, quickly followed by competitive open-source models. A competitive advantage in AI will require much more than consumer chatbots.

Security: Consumer LLM chatbots examine code and deliver vulnerabilities to hackers in response to prompts, successful phishing text, behavioral methods, workarounds for physical security, and even catastrophic bioterrorism weapons.

GenAI Opportunities for Organizations

Productivity: Early studies have indicated improved productivity in GenAI range between 14% and 37%.^{3,4}

Enterprise-wide accelerant: GenAI can be efficiently applied to most tasks in the enterprise by the KOS to accelerate objectives tailored to each entity, and do so without reckless risks, including sovereignty.⁵

Competitive advantage: Within the KOS, the GenAI functions can provide additional competitive advantages whereas consumer GenAI alone was instantly commoditized.

New Opportunities: The option to extend DANA to end customers is an example of a large new opportunity for many enterprises.

Generative AI (GenAI) is a term that has become widely adopted to describe machine learning (ML) models that generate text, images, and code in response to prompts, or user instructions. The most common ML models employed in GenAI are generative adversarial networks (GANs) and generative pre-trained transformers (GPTs).

When models reach billions of parameters, they are considered large language models (LLMs). Due to the unique characteristics of consumer LLMs, such as scraping data from the web, consumer chatbots represent considerable risk for the enterprise, resulting in bans or restrictions by many leading companies.

CEOs and boards are faced with several choices, including assuming risks when employees work with consumer chatbots, use only commoditized GenAI offered by tech giants, build custom GenAI system similar to BloombergGPT, or some combination thereof.

KYield offers a better option with the KOS. We take an efficient end-to-end systems approach that is continuously adaptive to rapidly changing environments, and tailored to the needs of each organization by the organization and its employees. The KOS provides a competitive advantage, strong governance, efficiency, and high performance at small fraction of the cost of a custom system.

Rules-based Governance

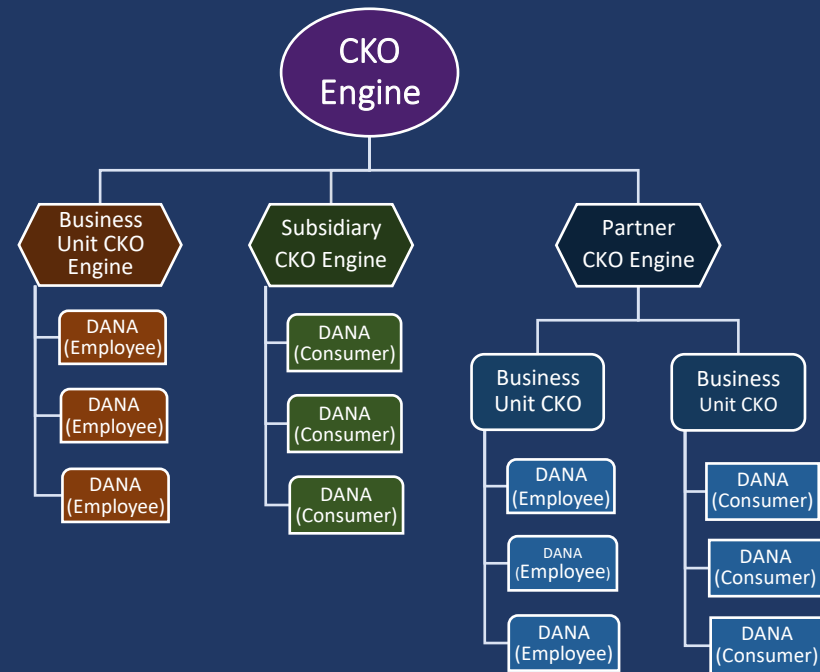


Nations and businesses are governed by the rule of law, by-laws, and organizational policies, whereas neural networks (NNs) cannot yet be governed directly by rules in an accurate manner by algorithms alone. To date, NNs have only achieved rough representation of rules, not precision governance necessary in the enterprise.⁶

Safe and effective governance that meets the needs of organizations and regulatory compliance must be achieved with other types of systems, including precision data management systems, symbolic AI, and neurosymbolic AI.⁷

The KOS employs a precision data management system for strong governance and high performance. NNs such as GenAI are compartmentalized in different data stores than the rest of the KOS for safety and accuracy. A program called the CKO Engine provides system-wide governance for the KOS in a semi-automated manner, administered by a simple natural language interface for approved personnel.

Administrators for the CKO Engine manage access for subservient CKO Engines, and DANA (Digital Assistant with Neuroanatomical Analytics), which is provided for every employee in the organization. The KOS and DANA can be extended to business partners and consumers.



Settings in the CKO Engine

- Roles and permissions
- Access to all modules
- Access to data file types
- Multiple types of security
- Menu of algorithms

Metamorphic Transformation: GIGO to QIQO

(Garbage-in / Garbage-out to Quality-in / Quality-out)



Precision Data Management with the Modular KOS

1. Natural language governance for entire system
2. Every file contains rich descriptive metadata
 - Data provenance from the source
 - End-to-end data lineage
3. Accurate rating system for every file
4. Data tuner allows simple control of quality and quantity
5. DANA can be extended to end-consumers with GenAI trained on much safer and more accurate data

Initially conceived in 1997 by Mark Montgomery in his knowledge systems lab where he built and operated leading learning networks, the KYield theorem faced multiple obstacles before it could be realized.

Solving information overload would require an end-to-end precision data management system (DMS). The DMS provides structure to data as individuals use DANA throughout the course of their work. Every file in the KOS is rated, thus providing the mathematical basis for achieving yield management of knowledge.

An integral component of the patented DMS is the data tuner, which enables users of DANA to increase quality and decrease quantity simply by clicking their mouse.

High quality data requires provenance, verified profiles, an accurate rating system, and strong security, all of which is lacking in consumer chatbots that scrape data from the web, resulting in false information and “hallucinations”.

As the KOS becomes rapidly populated with data during the normal course of work, employees can prompt the GenAI function to produce high quality, multi-media knowledge work based on a combination of their own relevant data, licensed data, synthetic data, and data sourced from their knowledge network. The result is metamorphic transformation of the digital workplace.⁸

Learning Networks




KYield was conceived while our founder was operating the premier knowledge network in the 1990s called GWIN (Global Web Interactive Network).

GWIN was thought to have hosted the greatest concentration of senior thought leaders in one network at any time until that point in history, including professors from most leading universities, multiple Nobel laureates, editors in most leading business and economic media, entire boards of Fortune 500 companies, government analysts, senior managers from most of the global 2,000, and top-tier investors, among others.

Human knowledge networks empowered by computer networks have proven to be a powerful force for rapid improvement, risk management and prevention, and increased competitive advantage in organizations and ecosystems. They can be much more powerful when augmented with advanced AI systems.

However, real-world tests of state-of-the-art knowledge networks and AI systems have also demonstrated that a very specific design in end-to-end data management systems is essential, complete with highly specific governance functions, multiple types of security, trustworthy policies based on ethical principles, and an **Quot**enticing design that encourages serious engagement.

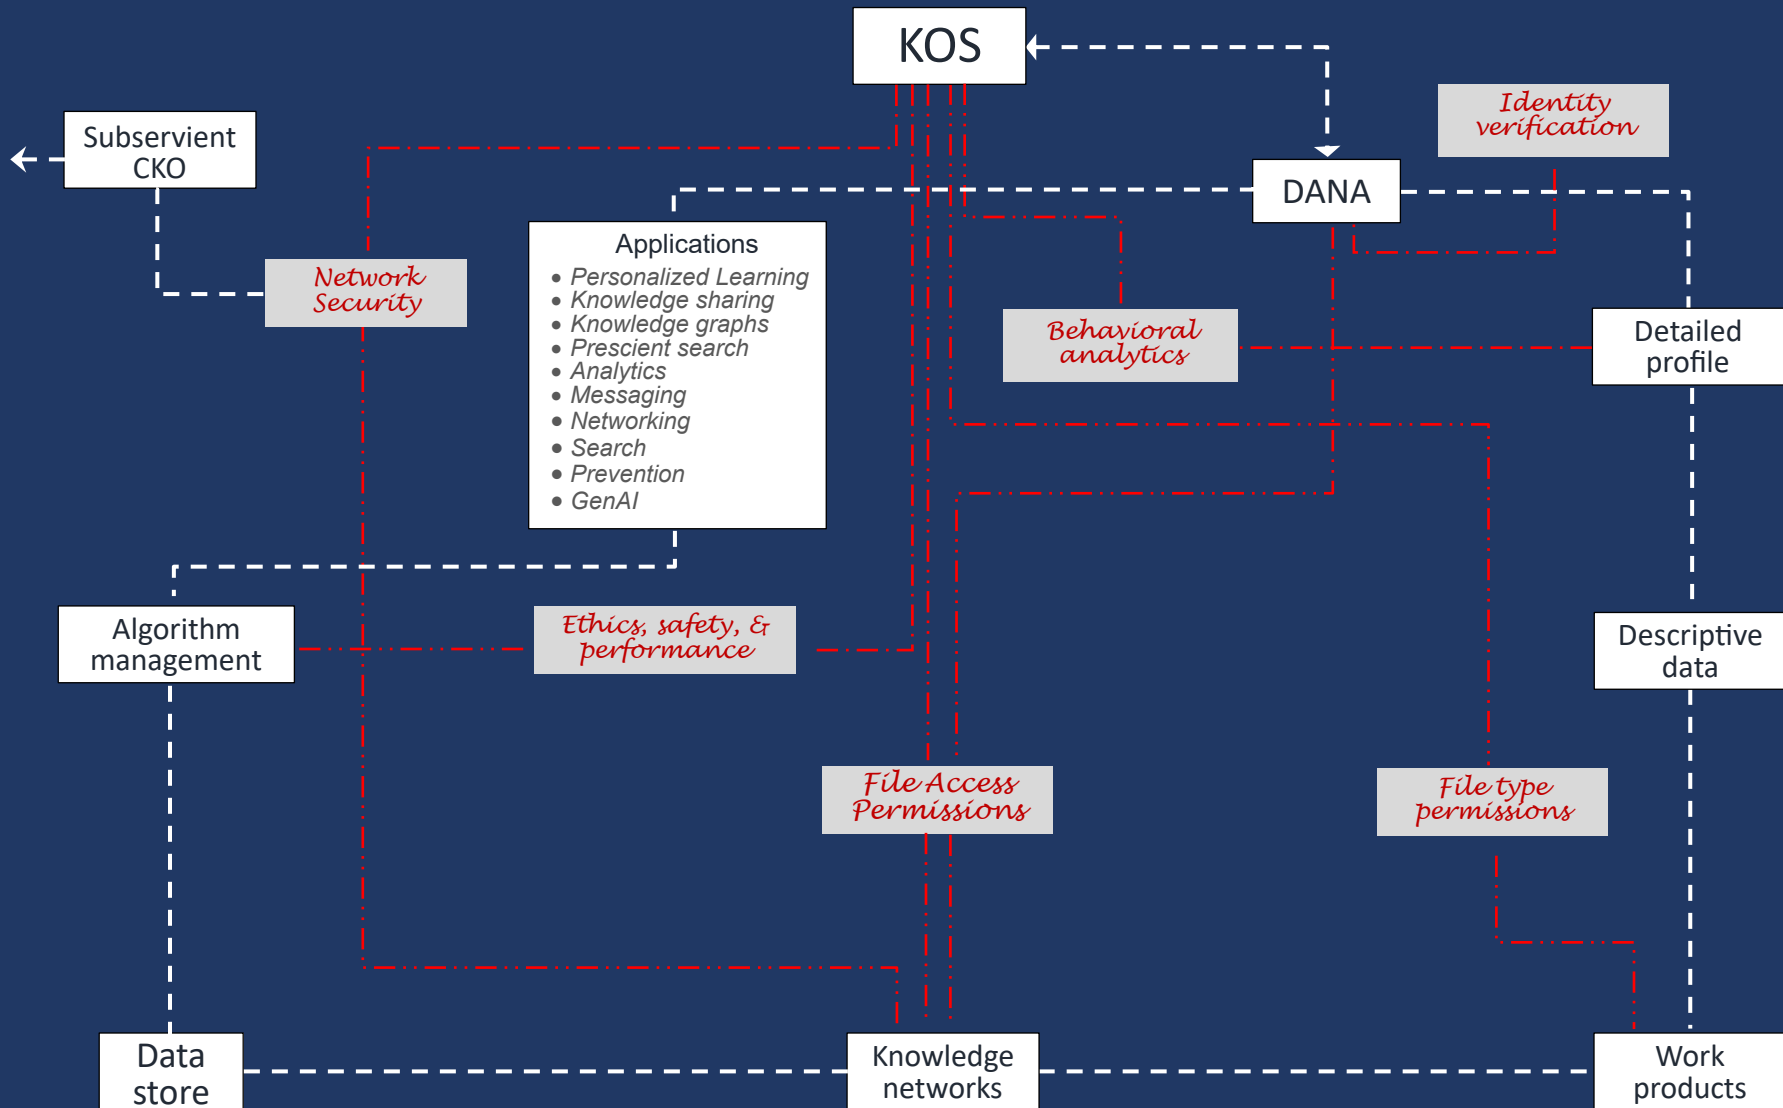


Visual created by a graph database in DANA

- Clickable links to view profiles
- Email invitations to approved individuals
- Corporate approval for networked relationships
- Multiple pre-programmed choices of graphs
- Searchable for rapid discovery of experts
- Automated risks and opportunities in multiple forms



KOS Process





Affordability & ROI

Our first cost estimate to deliver a working prototype of the KOS was in 2012 for the U.S. Army. The rough estimate was \$100 million due to multiple custom components needed, including hardware for the supercomputers necessary to operate the system.

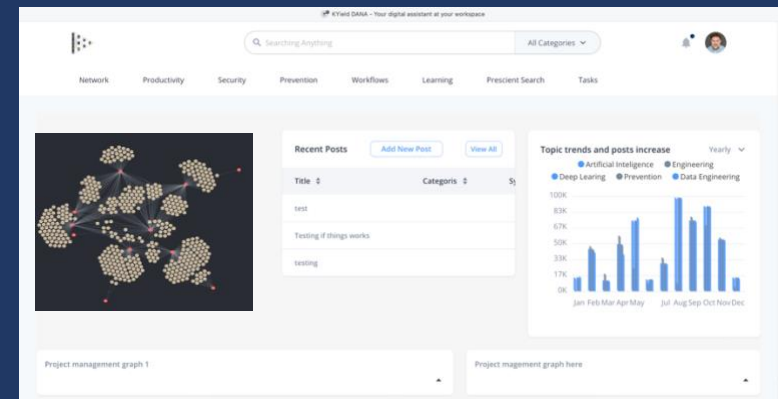
Since that time costs and performance have improved exponentially due to a combination of hundreds of billions of dollars of investment by the industry, KYield R&D, and the breadth and depth of R&D focused on AI, resulting in low-cost components.

The pricing today for the KOS is comparable to most enterprise software vendors. Our model includes an enterprise license for the KOS, a per-person-subscription for individuals, and a data usage cost beyond a reasonable base.

Our research clearly demonstrates the value of including all employees in the system, as well as the risks in not doing so. Our pricing structure therefore favors the inclusion of all employees.

In a recent extensive engagement with a very large company, we estimated the KOS could save over \$200 million per year in operational costs alone, not including captured preventions or other ROI.


Although the KOS can be run on most clouds, on premises, or colocation datacenters, we recommend a hybrid cloud format.



DANA Prototype 2/2023

Top Five Functions for Expected ROI

1. Prevention provides the highest ROI possible
 - A mid-size captured prevention can pay for the entire lifetime cost of the system
2. Personalized learning manifests in improved innovation, lower costs, and new opportunities
3. Enhanced productivity across the enterprise achieves much more with the same people
4. Smart networking reduces risk, improves decision making, and accelerates goal achievement
5. Multiple types of security reduce fraud, insider risk, industrial espionage, and cyber breaches



Conclusion

After nearly 70 years of R&D resulting in a large number of narrow applications, AI is having a societal moment, courtesy of a start-up that began life as a non-profit. Powerful LLM chatbots were created by leading scientists and unleashed prematurely without rigorous testing, presumably due to the need for vast amounts of capital to sustain the model.

A few months later the start-up received \$10 billion from a big tech company to support a product full of flaws, based on a model that can't be made safe without high quality data. Soon thereafter congressional hearings arrived on the regulation of AI, with calls for a new federal agency.

Consumer LLM chatbots can also provide many benefits, including productivity enhancement for anyone, and exposure to unrestricted, two-way communications, in a very broad AI system that also has depth, albeit with little safety and no consistent accuracy.

Since most of what we do as humans involves a layer of natural language to execute, natural language processing

(NLP) has been an area of intense interest by AI researchers for decades, including at KYield. Over the past two decades, NLP has rapidly advanced from single letter translation to words, phrases, and voice interpretation. With a combination of transformers and very large-scale models, we now have full content reproduction in multiple media for most languages, including code.

LLM chatbots have also demonstrated universality in AI systems, which is similar to our KOS. AI systems lack the ability to care about specific jobs, interests, or industries. What matters most is the data they are trained on, hence our focus on precision data. AI is really about data management, or more accurately, data management systems for competitive, accurate, and safe AI.⁹

In the enterprise market, the unleashing of consumer LLMs created a new movement called generative AI, triggering a flurry of activity in venture capital funding for specific use cases with little consideration for executable management or governance across the enterprise.

A recent survey by *Fortune* reveals that Fortune 500 CEOs may be less impressed by the LLM hyperbole than others. The survey asked CEOs to rank technologies “in terms of your view of their potential as opportunities for your business over the next ten years”. 58% of CEOs ranked predictive AI highest, whereas GenAI was ranked highest by only 12%.¹⁰

We considered this survey as the latest of many confirmations for our own research and communications on EAI. While prevention provides the highest ROI possible, many areas of potential ROI exist across the enterprise, particularly when managed properly with an efficient, cohesive system. One of those functions embedded in our KOS offering is now GenAI.

Although precision data management preoptimized for AI with proper governance at enterprise scale requires a few months more time than consumer LLM chatbots, and modest investment, we believe the experience of the last few months reconfirms the need and value of the KOS.

References

1. M. Montgomery. “Modular system for optimizing knowledge yield in the digital workplace”. [USPTO #8,005,778, 23 August 2011](#).
2. Unknown security experts discussing vulnerabilities in ChatGPT, reposted on LinkedIn by Mark Montgomery. https://www.linkedin.com/posts/markamontgomery_shocking-ai-chatgpt-activity-7065408586956369920-pldr
3. Brynjolfsson, Erik, Raymond, Lindsey R., and Li, Danielle, “Generative AI at Work”, NBER working paper, April 2023. <http://www.nber.org/papers/w31161>
4. Noy, Shakked and Zhang, Whitney, “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence”, working paper, MIT, March 2, 2023. https://economics.mit.edu/sites/default/files/inline-files/Noy_Zhang_1.pdf
5. Heikkilä, Melissa, “Google is throwing generative AI at everything”, *MIT Tech Review*, March 10, 2023. <https://www-technologyreview-com.cdn.ampproject.org/c/s/www.technologyreview.com/2023/05/10/1072880/google-is-throwing-generative-ai-at-everything/amp/>
6. Seo, Sungyong and Arik, Serkan, “Controlling Neural Networks with Rule Representations”, Google Research, January 28, 2022. <https://ai.googleblog.com/2022/01/controlling-neural-networks-with-rule.html>
7. Montgomery, Mark, “The Power of Neurosymbolic AI”, *EAI Newsletter*, March 29, 2023. <https://www.linkedin.com/pulse/power-neurosymbolic-ai-mark-montgomery>
8. Montgomery, Mark, “Metamorphic transformation with enterprise-wide artificial intelligence”, a recorded talk at ExperienceIT New Mexico conference, September, 2019. <https://youtu.be/yUVNeb6HZCI>
9. Bice, Bill and Montgomery, Mark, “AI is Really About Data Management”, *The Queue Podcast*, Legal Tech & News Analysis, April 14, 2023. <https://nqzw.com/ai-is-really-about-data-management-with-mark-montgomery-founder-ceo-of-kyield/>
10. Murray, Alan, “[Generative AI isn’t for everyone](#)”, “The CEO Daily”, *Fortune*, May 18, 2023.