

Artificial Intelligence

A Field Guide for Public Sector Enterprises

EDITION 2 JUNE 2025

(EDITION 1 RELEASED JUNE 2021)

Executive Summary



Digital transformation has led to an “always-on” society that continuously generates new data, connects users instantaneously and delivers seamless, frictionless experiences. Humans alone cannot deal with the massive amount of data that surrounds us on this new competitive playing field for critical decision-making, whether that decision is necessary to protect our nation’s infrastructure or thwart the next cyber attack. The federal government and industry alike are investing heavily in artificial intelligence (AI), harnessing the computational ability to learn and mimic the cognitive functions of humans capable of deriving insights at scale—key to our nation’s digital dominance.

In our experience, many organizations are unsure how to best use their data to power their mission, enable effective and rapid decision-making and ensure the digital fluency of their workforce. It is no longer enough to “find the needle in the haystack;” data needs to be both trusted and actionable and used in a way that provides predictive insights at the speed of need. This Field Guide provides practical guidance and best practices for public sector organizations seeking to adopt AI within their enterprise.

While generative AI solutions on the market today may provide instant gratification through rapid “human-like” responses, most of these solutions currently lack the ground truth necessary for accurate, unbiased and critical decisionmaking to solve complex federal mission challenges. Conceptually, the goal of AI is straightforward, but the range of potential risks and benefits to any organization attempting to apply can be daunting and highly complex. For organizations seeking to perform tasks at a scale unachievable before, in this new era of AI and ubiquitous computing, they must consider factors beyond data, technology and infrastructure to encompass ethical considerations, talent, training and more. As AI becomes more powerful, responsible and ethical AI frameworks will be foundational to ensure the trustworthiness, integrity, privacy and security of data and algorithms. In this guide, our authors share current challenges, opportunities and risks based on real-world experience offering practical advice, use cases and strategies for harnessing AI in the public sector.

- Mile Corrigan, *Noblis President and Chief Executive Officer*

Introduction

The Noblis Artificial Intelligence (AI) Field Guide provides a wide-ranging overview of AI technologies, adoption considerations and responsible AI practices. The Guide consists of four chapters:

1. **AI Capabilities:** A primer on AI-related capabilities, including machine learning (ML), Natural Language Processing (NLP) and machine decision making.
2. **The AI Playbook:** Best practice guidelines for implementing AI programs within an organization.
3. **Responsible AI:** Guidelines for organizations developing and adopting AI technologies and services in a secure, unbiased and ethical AI manner, with an example framework to enable users in doing so.
4. **Special Topics in Emerging AI Technologies:** Various short articles covering new research and developments in the field.

The Field Guide informs readers about the rapidly evolving AI technology landscape, provides a roadmap for organizations to harness the benefits of AI technology and promotes AI adoption—responsibly—across the public sector.

Chapter One - AI Capabilities

Chapter One is a primer on AI-related capabilities, describing a selection of technologies, trends and adoption considerations. **Figure 1** outlines the capabilities discussed in Chapter One.

Machine Learning

The AI Field Guide explores ML, a core component of data science that enables computers to learn from data, discover patterns and make decisions with minimal human intervention. ML algorithms learn from novel




Figure 1: AI Capabilities Addressed in Chapter One

data types and can dynamically train AI models, which can be applied to various problems, such as classifying, grouping, ranking, or predicting data points. ML's flexibility has led to its widespread adoption in both the public and private sectors. Despite potential resource constraints, modern techniques can optimize model performance and adapt ML to specific use cases. ML algorithms are categorized into supervised learning, unsupervised learning and reinforcement learning. Deep learning, a key application of ML, is seeing widespread adoption across the public sector for a variety of computer vision and natural language processing use cases. As ML continues to evolve, it is transforming industries and becoming a critical enabler for law enforcement, national security and citizen services. Adoption considerations include ensuring explainable AI (i.e., enabling human understanding in how the system reached its conclusions), addressing potential bias in the ML models and training data and addressing security concerns inherent in any complex software system.

Perceiving the Environment

Perception of the environment is a critical aspect of data science that enables machines to understand and adapt to their surroundings. Deep learning, with its powerful neural networks, is vital for processing imagery and language, while transfer learning allows for problem-specific adaptations. Effective sensor networks are achieved through the deployment and fusion of data from various sensors with perception delivered at the edge using ML algorithms. Ad-hoc sensor networks with AI-enhanced information collection and dissemination, while fusing sensor data overcomes limitations and perception at the edge addresses hardware limitations. The design and deployment of these systems are influenced by factors such as trustworthiness, timeliness, accessibility, completeness and cost efficiency.

Natural Language Understanding

This section delves into the rapidly evolving field of NLP, with a focus on Natural Language Understanding (NLU) and Natural Language Generation (NLG). It considers the context and data representativeness in NLP applications to ensure effective and socially responsible adoption.

Gathering & Using Knowledge

The Field Guide emphasizes the importance of a robust data strategy for AI/ML adoption, discussing the Five V's of data and crucial data management aspects. This section provides a comprehensive view of considerations for implementing AI/ML systems in modern organizations.

Automating Action

This section examines the integration of sensors, computing systems and communication devices, enabling machines to operate autonomously and address complex challenges. It covers Environmental Perception, Autonomous Actions, Human-Machine Interaction, and Adoption Considerations in autonomous machine development.

Making Decisions

This section explores the role of AI systems in decisionmaking processes. The focus is on highlighting their benefits in enhancing decision speed, accuracy, scalability, and consistency. It emphasizes factors to consider when designing decision-making systems and the importance of AI in improving government agency performance.

Harnessing Creativity and Generative AI

The Field Guide provides a concise overview of creative AI and its five main areas, which, when combined, generate innovative and functional outputs. Large language models (LMMs), such as ChatGPT, have demonstrated their usefulness across a range of applications by skillfully integrating imitation and optimization. However, concerns regarding intellectual property, information sensitivity and bias must be addressed. The quality of training data is crucial to mitigate potential risks and ensure the responsible development and deployment of these AI systems.

Unleashing the Power of Large Language Models

This section describes the concepts behind, and explores the use of, generative AI, specifically LLMs and their role in the modern enterprise. This portion of the guide includes key advancements in the field, and applications of LLMs for mission and business impact.

Chapter Two: The AI Playbook—Roadmap for Adoption and Management

Chapter Two is a guide for leaders seeking to integrate AI into their organizations, featuring insights and best practices and experiences from across various industries. The content is divided into four key areas (**Figure 2**) that provide guidance on creating an adoption strategy, optimizing workforce investments, developing a data architecture that aligns with organizational goals and ensuring long-term success in building an AI-enabled organization.



Figure 2: AI Playbook Key Areas

Domain Strategy and Organization

This section outlines the importance of a clear AI adoption strategy that covers processes, people and technology. It covers key elements to consider when designing an AI adoption roadmap, addressing success factors for gaining organizational support and adapting to changing circumstances. It also describes common challenges faced by organizations entering the AI field, such as strategic misalignment, unclear AI value and lack of vision, strategy and objectives. It organizes essential questions and considerations for formulating an AI strategy into three main areas: Assessing Organizational Readiness, the Tenets of a Successful Strategy and Communicating Change.

Building The AI-Ready Workforce

This section highlights the critical role of a prepared workforce in AI transformation. It delivers strategies for creating an AI-ready workforce, focusing on investing in people, fostering collaboration and establishing success metrics. It covers challenges in acquiring skilled talent and the need for upskilling existing employees. Additionally, it emphasizes the importance of diverse teams, an agile mindset, and clear objectives and key results to track progress.

Building Your Data Architecture

This section explores the importance of aligning technology investments with organizational objectives and readiness. A successful AI adoption relies on an appropriate data architecture tailored to the organization's needs, maturity and overall capabilities. It delves into the evaluation process for migrating traditional applications to AI models to address aspects such as governance and scaling of data architecture. It also encourages continuous improvements in data architecture as organizational experience and technical evolution progress. Key takeaways include evaluating application suitability, conducting limited pilots and prioritizing security and trust, while considering various pertinent questions throughout the process.

Driving Further AI Adoption

This section highlights the importance of aligning AI investments and initiatives with core business or mission objectives for long-term success as an AI-enabled organization. Factors contributing to achieving the desired scale include focusing on human capital, responsive leadership, recalibrated vendor and partner management, and adopting MLOps approaches. By equipping staff with necessary training, fostering a culture of experimentation and assembling diverse teams, organizations can maintain a competitive edge.

Chapter Three—Ethical and Responsible AI

Responsible AI is a dynamic field focused on managing AI risk while promoting transparent AI practices to prevent undesirable outcomes. Various publications and frameworks have been developed to ensure that AI systems act fairly, operate as intended, and are held accountable for their results. This chapter introduces the concept of Responsible AI and delves into its foundations and associated research, outlines evolving standards and guidance from across the Public Sector and industry, and provides a sample governance framework consisting of principles, a System Development Life Cycle (SDLC), best practices in risk management and an AI System Review Process. These elements work together to promote compliance, provide clear instructions and build trust in AI systems in this developing field. The chapter presents key risk categories organizations may want to address in developing AI solutions (illustrated in **Figure 3** and described in the chapter) and an example Responsible AI Review Process, highlighting the roles of key stakeholders such as the AI Review Committee (ARC) and the AI Review Board (ARB). The ARC and ARB collaborate to ensure that organizations comply with Responsible AI principles. As part of this process, project teams submit AI system proposals for review by the ARB, which assesses adherence to guidelines. Approved projects are then monitored by an ARC that offers technical implementation guidance.

In closing, the chapter encourages organizations pursuing Responsible AI to prioritize compliance and actively work to mitigate risks. By establishing a strong ethical foundation and effective governance, organizations can maximize the benefits of AI while minimizing unintended consequences.

Risk Categories in AI/ML

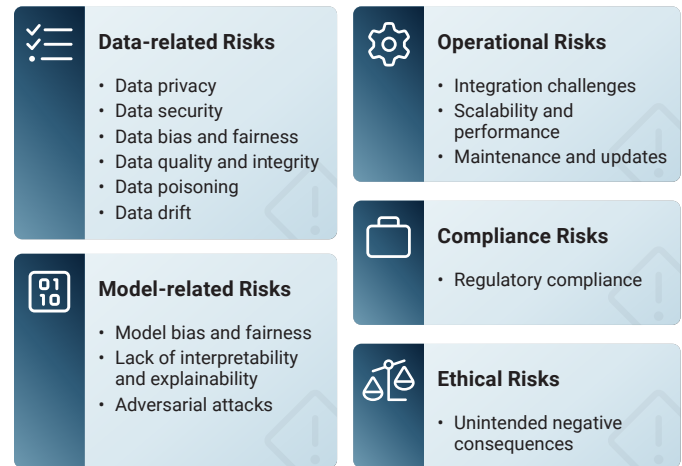



Figure 3: Risk Categories in AI/ML

Chapter Four—Special Topics in Emerging AI Technologies

Developments in the field of AI are accelerating, with the pace of change presenting enormous opportunities to enterprises of all sizes. Chapter Four provides a series of short, frequently updated articles on select advances and research in the field of particular interest to public sector enterprises, and how enterprises can benefit from these advances in the context of their mission.



Chris Barnett
Noblis CTO

Informed application of AI and its subdisciplines holds enormous promise for National Security and Citizen Services alike. While the underlying technologies have been available in various forms for some time, the ubiquity of diverse datasets, affordable processing, and accessible tools

and frameworks make the benefits of AI real to the U.S. Government. Investments in diverse, vendor-agnostic AI systems, tools and training now will enable departments and agencies to take advantage of new sensors for safety and intelligence processing, to deliver accurate and efficient information to citizens, to enable the missions of law enforcement and homeland security personnel, and to augment the effectiveness of our warfighters and National Security partners. NLP techniques will enable faster development and sharing of knowledge, while countering the spread of misinformation; at the same time, machine learning will enable complex machines to take on more repetitive and time-consuming tasks, enabling people to focus on value-added activities instead.

Adoption of AI capabilities requires a multi-pronged approach. First, there is a cultural aspect, where people must be ready to adopt new ways of processing information and interpreting results. Trustworthiness and explainability, two related and very active areas of scientific research by Noblis and other entities, must be present and built into the system to spur widespread adoption. Similarly, data set readiness, preparation, usage and associated training must incorporate a selection of counter bias techniques, ensuring results are not tainted by a preordained outcome. AI system development environments and model frameworks, while still complex, will continue to grow more accessible to more people. This will lead to teams of citizen developers of AI systems, all of whom will benefit from learning techniques to counter bias, enforce trust and collaborate on system security requirements. An AI adoption strategy should focus on one or more achievable use cases, followed by a self-evaluation to understand the risks and benefits.

In this Field Guide, we lay out the foundational organizational requirements and intentional investments for level of AI adoption. The promise of AI is significant across the federal landscape. We hope this document illustrates just some of the potential of this exciting field, demystifying the jargon and encouraging the leap for your own department or agency's adoption of AI strategies.

About Noblis

Noblis is an independent, nonprofit organization with a proud tradition of serving federal clients objectively and with the highest caliber of scientific and technical excellence. Our work impacts the civil, defense, homeland security and intelligence & law enforcement missions that ensure our nation's vitality and security.

We employ objectivity, innovation and a multidisciplinary perspective to solve complex problems in the public interest and are pleased to share our mission impacts and technical insights from our thought leaders.

To learn more about Noblis' research and capabilities, visit noblis.org.



noblis.org

2002 Edmund Halley Drive, Reston, VA 20191 | 703.610.2000 | answers@noblis.org

© 2025 Noblis, Inc. All rights reserved.