iappai
— governance center —

THE AIGP BODY OF KNOWLEDGE (BoK)

VERSION 1.0.0

EFFECTIVE DATE: 6/20/2023

THE AIGP BODY OF KNOWLEDGE (BoK)

The rapid rise of generative artificial intelligence has focused our collective attention on the promise and peril of an AI-fueled society. With equal parts excitement and trepidation, we find ourselves asking how to build a future augmented by the potential benefits of AI, while avoiding its pitfalls.

Every day we hear more about the potential of AI-powered systems to transform how we work, create, solve problems, communicate, and even diagnose and treat illness. The possibilities of advanced AI seem to be unlimited.

But without proper testing, evaluation, validation, and verification at each stage of AI development, foundational AI models could perpetuate biases and amplify other societal challenges that will cascade through later systems and remain for decades.

We must continue to build and refine the governance processes through which trustworthy AI will emerge and we must invest in the people who will build ethical and responsible AI. Those who work in compliance, risk management, legal and governance together with data scientists, AI project managers, model ops teams and others must be prepared to tackle the expanded equities at issue in AI governance.

To meet this demand, the IAPP has developed the Artificial Intelligence Governance Professional (AIGP) certification and training for the emerging AI governance profession. An AIGP trained and certified professional will know how to implement and effectively communicate across teams the emerging best practices and rules for responsible management of the AI ecosystem. We are privileged to grow a community of credentialed AI governance professionals, through which resources and expanding knowledge can be brought together in one place.

THE AIGP BODY OF KNOWLEDGE

UNDERSTANDING THE AIGP BODY OF KNOWLEDGE

The main purpose of the AIGP body of knowledge is to document the knowledge and skills that will be assessed on the AIGP certification exam. The domains of the BoK capture the activities that an AI governance professional should undertake to guide AI's implementation in a manner that mitigates risk and ensures safety and trust. There are six main domains and a seventh that entertains emerging governance and legal issues:

- **Domain 1: "Understanding the Foundations of Artificial Intelligence,"** defines AI and machine learning, provides an overview of the different types of AI systems and their use cases, and positions AI models in the broader socio-cultural context.
- **Domain 2: "Understanding AI Impacts and Responsible AI Principles,"** identifies the risks that ungoverned AI systems can have on humans and society and describes the characteristics and principles that are essential to trustworthy and ethical AI.
- **Domain 3: "Understanding How Current Laws Apply to AI Systems,"** surveys the current laws that govern the use of artificial intelligence.
- **Domain 4: "Understanding the Existing and Emerging AI Laws and Standards,"** outlines the global AI-specific laws (like the EU AI Act and Canada's Bill C-27) and the major frameworks that show how AI systems can be responsibly governed.
- **Domain 5: "Understanding the AI Development Life Cycle,"** broadly outlines the context in which AI risks are managed.
- **Domain 6: "Implementing Responsible AI Governance and Risk Management,"** explains how the major AI stakeholders collaborate in a layered approach, to manage AI risks while fulfilling the potential benefits AI systems have for society.
- **Domain 7: "Contemplating Ongoing Issues and Concerns,"** presents some of the debated issues around AI governance.

The body of knowledge also includes the Exam Blueprint numbers, which show the number of questions from each part of the BoK that will be found on the exam.

The AIGP body of knowledge was developed by a substantial group of experts from the fields of ethics, law, privacy, computer science, sociology and psychology that represents the breadth of responsible AI stakeholders. The BoK will be reviewed (and, if necessary, updated) every six months; changes will be reflected in exam updates and communicated to candidates at least 90 days before the new content appears in the exam.

COMPETENCIES AND PERFORMANCE INDICATORS

The content in the body of knowledge is represented as a series of competencies and connected performance Indicators.

Competencies are clusters of connected tasks and abilities that constitute a broad knowledge domain.

Performance indicators are the discrete tasks and abilities that constitute the broader competence group. Exam questions assess an AI governance professional's proficiency on the performance indicators.

WHAT TYPES OF QUESTIONS WILL BE ON THE EXAM?

For the certification candidate, the performance indicators are guides to the depth of knowledge required to demonstrate competency. The verbs that begin the skill and task statements (identify, evaluate, implement, define) signal the level of complexity of the exam questions and find their corollaries on the Bloom's Taxonomy (see next page).

BLOOM'S TAXONOMY

Bloom's Taxonomy (often represented as a pyramid) is a hierarchy of cognitive skills used to establish educational learning objectives. IAPP exam questions mostly focus on the remember/understand and apply/analyze levels.

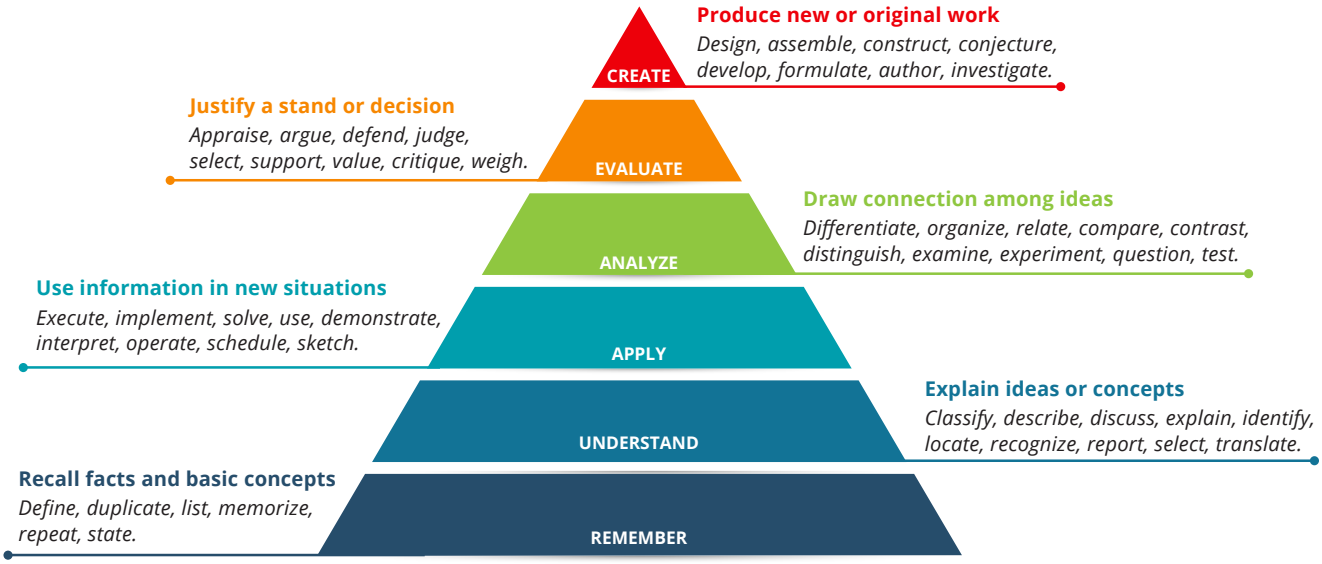

Domain I: Understanding the Foundations of Artificial Intelligence

Domain I – "Understanding the Foundations of Artificial Intelligence," defines AI and ML, provides an overview of the different types of AI systems and their use cases, and positions AI models in the broader socio-cultural context

Competencies	Performance Indicators
Understand the basic elements of AI and ML	Understand widely accepted definitions of AI and ML, and the basic logical-mathematical principles over which AI/ML models operate.
	Understand common elements of AI/ML definitions under new and emerging law: 1. Technology (engineered or machine-based system; or logic, knowledge, or learning algorithm). 2. Automation (elements of varying levels). 3. Role of humans (define objectives or provide data). 4. Output (content, predictions, recommendations, or decisions).
	Understand what it means that an AI system is a socio-technical system.
	Understand the need for cross-disciplinary collaboration (ensure UX, anthropology, sociology, linguistics experts are involved and valued).
	Knowledge of the OECD framework for the classification of AI systems.
	Understand the use cases and benefits of AI (recognition, event detection, forecasting, personalization, interaction support, goal-driven optimization, recommendation).

Domain I: Understanding the Foundations of Artificial Intelligence	
Competencies	**Performance Indicators**
Understand the differences among types of AI systems	Understand the differences between strong/broad and weak/narrow AI.
	Understand the basics of machine learning and its training methods (supervised, unsupervised, semi-supervised, reinforcement).
	Understand deep learning, generative AI, multi-modal models, transformer models, and the major providers.
	Understand natural language processing: text as input and output.
	Understand the difference between robotics and robotic processing automation (RPA).
Understand the AI technology stack	Platforms and applications.
	Model types.
	Compute infrastructure: software and hardware (servers and chips).
Understand the history of AI and the evolution of data science	1956 Dartmouth summer research project on AI.
	Summers, winters and key milestones.
	Understand how the current environment is fueled by exponential growth in computing infrastructure and tech megatrends (cloud, mobile, social, IOT, PETs, blockchain, computer vision, AR/VR, metaverse).

Domain II: Understanding AI Impacts on People and Responsible AI Principles

Domain II – "Understanding AI Impacts on People and Responsible AI Principles," identifies the risks that ungoverned AI systems can have on humans and society and describes the characteristics and principles that are essential to trustworthy and ethical AI

Competencies	Performance Indicators
Understand the core risks and harms posed by AI systems	Understand the potential harms to an individual (civil rights, economic opportunity, safety).
	Understand the potential harms to a group (discrimination towards sub-groups).
	Understand the potential harms to society (democratic process, public trust in governmental institutions, educational access, jobs redistribution).
	Understand the potential harms to a company or institution (reputational, cultural, economic, acceleration risks).
	Understand the potential harms to an ecosystem (natural resources, environment, supply chain).
Understand the characteristics of trustworthy AI systems	Understand what it means for an AI system to be "human-centric."
	Understand the characteristics of an accountable AI system (safe, secure and resilient, valid and reliable, fair).
	Understand what it means for an AI system to be transparent.
	Understand what it means for an AI system to be explainable.
	Understand what it means for an AI system to be privacy-enhanced.
Understand the similarities and differences among existing and emerging ethical guidance on AI	Understand how the ethical guidance is rooted in Fair Information Practices, European Court of Human Rights and Organization for Economic Cooperation and Development principles.
	OECD AI Principles; White House Office of Science and Technology Policy Blueprint for an AI Bill of Rights; High-level Expert Group AI; UNESCO Principles; Asilomar AI Principles; The Institute of Electrical and Electronics Engineers Initiative on Ethics of Autonomous and Intelligent Systems; CNIL AI Action Plan.

Domain III: Understanding How Current Laws Apply to AI Systems

Domain III – "Understanding How Current Laws Apply to AI Systems," surveys the current laws that govern the use of artificial intelligence

Competencies	Performance Indicators
Understand the existing laws that interact with AI use	Know the laws that address unfair and deceptive practices.
	Know relevant non-discrimination laws (credit, employment, insurance, housing, etc.).
	Know relevant product safety laws.
	Know relevant IP law.
	Understand the basic requirements of the EU Digital Services Act (transparency of recommender systems).
	Know relevant privacy laws concerning the use of data.
Understanding key GDPR intersections	Understand automated decision making, data protection impact assessments, anonymization, and how they relate to AI systems.
	Understand the intersection between requirements for AI conformity assessments and DPIAs.
	Understand the requirements for human supervision of algorithmic systems.
	Understand an individual's right to meaningful information about the logic of AI systems.
Understanding liability reform	Awareness of the reform of EU product liability law.
	Understand the basics of the AI Product Liability Directive.
	Awareness of U.S. federal agency involvement (EO14091).

Domain IV: Understanding the Existing and Emerging AI Laws and Standards

Domain IV – "Understanding the Existing and Emerging AI Laws and Standards," identifies and describes global AI-specific laws and the major frameworks that show how AI systems can be responsibly governed

Competencies	Performance Indicators
Understanding the requirements of the EU AI Act	Understand the classification framework of AI systems (prohibited, high-risk, limited risk, low risk).
	Understand requirements for high-risk systems and foundation models.
	Understand notification requirements (customers and national authorities).
	Understand the enforcement framework and penalties for noncompliance.
	Understand procedures for testing innovative AI and exemptions for research.
	Understand transparency requirements, i.e., registration database.
Understand other emerging global laws	Understand the key components of Canada's Artificial Intelligence and Data Act (C-27).
	Understand the key components of U.S. state laws that govern the use of AI.
	Understand the Cyberspace Administration of China's draft regulations on generative AI.
Understand the similarities and differences among the major risk management frameworks and standards	ISO 31000:2018 Risk Management – Guidelines.
	United States National Institute of Standards and Technology, AI Risk Management Framework (NIST AI RMF).
	European Union proposal for a regulation laying down harmonized rules on AI (EU AIA).
	Council of Europe Human Rights, Democracy, and the Rule of Law Assurance Framework for AI Systems (HUDERIA).
	IEEE 7000-21 Standard Model Process for Addressing Ethical Concerns during System Design
	ISO/IEC Guide 51 Safety aspects – guidelines for their inclusion in standards.
	Singapore Model AI Governance Framework.

Domain V: Understanding the AI Development Life Cycle

Domain V – "Understanding the AI Development Life Cycle," describes the AI life cycle and the broad context in which AI risks are managed

Competencies	Performance Indicators
Understand the key steps in the AI system planning phase	Determine the business objectives and requirements.
	Determine the scope of the project.
	Determine the governance structure and responsibilities.
Understand the key steps in the AI system design phase	Implement a data strategy that includes: • Data gathering, wrangling, cleansing, labeling. • Applying PETs like anonymization, minimization, differential privacy, federated learning.
	Determine AI system architecture and model selection (choose the algorithm according to the desired level of accuracy and interpretability).
Understand the key steps in the AI system development phase	Build the model.
	Perform feature engineering.
	Perform model training.
	Perform model testing and validation.
Understand the key steps in the AI system implementation phase	Perform readiness assessments.
	Deploy the model into production.
	Monitor and validate the model.
	Maintain the model.

Domain VI: Implementing Responsible AI Governance and Risk Management

Domain VI – "Implementing Responsible AI Governance and Risk Management," explains how the major AI stakeholders collaborate, in a layered approach, to manage AI risks while fulfilling the potential benefits AI systems have for society

Competencies	Performance Indicators
Ensure interoperability of AI risk management with other operational risk strategies	Ex. security risk, privacy risk, business risk.
Integrate AI governance principles into the company	Adopt a pro-innovation mindset.
	Ensure governance is risk-centric.
	Ensure planning and design is consensus-driven .
	Ensure team is outcome-focused.
	Adopt a non-prescriptive approach to allow for intelligent self-management.
	Ensure framework is law-, industry-, and technology-agnostic.

Domain VI – Implementing Responsible AI Governance and Risk Management	
Competencies	**Performance Indicators**
Establish an AI governance infrastructure	Determine if you are a developer, deployer (those that make an AI system available to third parties) or user; understand how responsibilities among companies that develop AI systems and those that use or deploy them differ; establish governance processes for all parties.
	Establish and understand the roles and responsibilities of AI governance people and groups including, but not limited to, the chief privacy officer, the chief ethics officer, the office for responsible AI, the AI governance committee, the ethics board, architecture steering groups, AI project managers, etc.
	Advocate for AI governance support from senior leadership and tech teams by: • Understanding pressures on tech teams to build AI solutions quickly and efficiently. • Understanding how data science and model operations teams work. • Being able to influence behavioral and cultural change.
	Establish organizational risk strategy and tolerance.
	Develop central inventory of AI and ML applications and repository of algorithms.
	Develop responsible AI accountability policies and incentive structures.
	Understand AI regulatory requirements.
	Set common AI terms and taxonomy for the organization.
	Provide knowledge resources and training to the enterprise to foster a culture that continuously promotes ethical behavior.
	Determine AI maturity levels of business functions and address insufficiencies.
	Use and adapt existing privacy and data governance practices for AI management.
	Create policies to manage third party risk, to ensure end-to-end accountability.
	Understand differences in norms/expectations across countries.

Domain VI – Implementing Responsible AI Governance and Risk Management	
Competencies	**Performance Indicators**
Map, plan and scope the AI project	Define the business case and perform cost/benefit analysis where trade-offs are considered in the design of AI systems. Why AI/ML?
	Identify and classify internal/external risks and contributing factors (prohibitive, major, moderate).
	Construct a probability/severity harms matrix and a risk mitigation hierarchy.
	Perform an algorithmic impact assessment leveraging PIAs as a starting point and tailoring to AI process. Know when to perform and who to involve.
	Establish level of human involvement/oversight in AI decision making.
	Conduct a stakeholder engagement process that includes the following steps: • Evaluate stakeholder salience. • Include diversity of demographics, disciplines, experience, expertise and backgrounds. • Perform positionality exercise. • Determine level of engagement. • Establish engagement methods. • Identify AI actors during design, development, and deployment phases. • Create communication plans for regulators and consumers that reflect compliance/disclosure obligations for transparency and explainability (UI copy, FAQs, online documentation, model or system cards).
	Determine feasibility of optionality and redress.
	Chart data lineage and provenance, ensuring data is representative, accurate and unbiased. Use statistical sampling to identify data gaps.
	Solicit early and continuous feedback from those who may be most impacted by AI systems.
	Use test, evaluation, verification, validation (TEVV) process.
	Create preliminary analysis report on risk factor and proportionate management.

Domain VI – Implementing Responsible AI Governance and Risk Management	
Competencies	**Performance Indicators**
Test and validate the AI system during development	Evaluate the trustworthiness, validity, safety, security, privacy and fairness of the AI system using the following methods: • Use edge cases, unseen data, or potential malicious input to test the AI models. • Conduct repeatability assessments. • Complete model cards/fact sheets. • Create counterfactual explanations (CFEs). • Conduct adversarial testing and threat modeling to identify security threats. • Refer to OECD catalogue of tools and metrics for trustworthy AI. • Establish multiple layers of mitigation to stop system errors or failures at different levels or modules of the AI system. • Understand trade-offs among mitigation strategies.
	Apply key concepts of privacy-preserving machine learning and use privacy-enhancing technologies and privacy-preserving machine learning techniques to help with privacy protection in AI/ML systems.
	Understand why AI systems fail. Examples include: brittleness; embedded bias; catastrophic forgetting; uncertainty.
	Determine degree of remediability of adverse impacts.
	Conduct risk tracking to document how risks may change over time.
	Consider, and select among different deployment strategies.

Domain VI – Implementing Responsible AI Governance and Risk Management	
Competencies	**Performance Indicators**
Manage and monitor AI systems after deployment	Perform post-hoc testing to determine if AI system goals were achieved, while being aware of "automation bias."
	Prioritize, triage and respond to internal and external risks.
	Ensure processes are in place to deactivate or localize AI systems as necessary (e.g., due to regulatory requirements or performance issues).
	Continuously improve and maintain deployed systems by tuning and retraining with new data, human feedback, etc.
	Determine the need for challenger models to supplant the champion model.
	Version each model and connect them to the data sets they were trained with.
	Continuously monitor risks from third parties, including bad actors.
	Maintain and monitor communication plans and inform user when AI system updates its capabilities. Assess potential harms of publishing research derived from AI models.
	Conduct bug bashing and red teaming exercises.
	Forecast and reduce risks of secondary/unintended uses and downstream harm of AI models.

Domain VII: Contemplating Ongoing Issues and Concerns

Domain VII – "Contemplating Ongoing Issues and Concerns," presents some of the current discussions and ideas about AI governance

Competencies	Performance Indicators
Awareness of legal issues	How will a coherent tort liability framework be created to adapt to the unique circumstances of AI and allocate responsibility among developers, deployers and users?
	What are the challenges surrounding AI model and data licensing?
	Can we develop systems that respect IP rights?
Awareness of user concerns	How do we properly educate users about the functions and limitations of AI systems?
	How do we upskill and reskill the workforce to take full advantage of AI benefits?
	Can there be an opt-out for a non-AI alternative?
Awareness of AI auditing and accountability issues	How can we build a profession of certified third-party auditors globally – and consistent frameworks and standards for them?
	What are the markers/indicators that determine when an AI system should be subject to enhanced accountability, such as third-party audits (e.g., automated decision-making, sensitive data, others)?
	How do we enable companies to remain productive using automated checks for AI governance and associated ethical issues, while adapting this automation quickly to the evolving standards and technology?