

ARTIFICIAL INTELLIGENCE GOVERNANCE

in GxP Environments

By Armand Mintancian, Rory Budihandojo, John English, HCCP,
Orlando Lopez, Jose E. Matos, and Robert McDowall, PhD

Artificial intelligence (AI) is used by pharmaceutical and biotech companies, providing support from drug discovery through manufacturing. The nature of AI and concerns of bias, privacy, transparency, and security in a regulated industry necessitate a governance framework to ensure concerns are controlled using “guardrails.” These guardrails ensure the quality, privacy, and security of data used in AI applications. This article provides a recommended approach to implementing guardrails through several policies and procedures and discusses AI governance, which defines data ownership, consent, and access policies and procedures.

This article provides points to consider for those implementing AI in a GxP environment. At a high level, these include governance for AI, as well as for machine learning operations (MLOps). MLOps is a set of practices that automate and simplify machine learning (ML) workflows and deployments [1]. This article also provides the relationship between AI and MLOps and why both are important in an AI implementation.

Definitions for AI, generative AI, ML, and deep learning (DL) can be found in ISO/IEC Standard 22989, “Information Technology — Artificial Intelligence — Artificial Intelligence Concepts and Terminology” [2]. Within this article, security and cybersecurity are used interchangeably, with the definition for cybersecurity found in the National Institute of Standards and Technology’s Glossary [3].

AI USE: FROM DRUG DESIGN TO DISTRIBUTION

AI is used throughout drug discovery, drug product development, clinical development, manufacturing, and distribution. Within drug discovery, drug design (e.g., target protein structure prediction) and drug screening (e.g., bioactivity prediction, toxicity prediction, and physicochemical property prediction) are impacted areas

benefiting from AI [4]. Note that drug design is out of regulatory scope; it would be best business practice to implement policies and procedures noted within this document. In drug product development, AI assists in deciding suitable excipients, monitors and modifies the development process, and ensures in-process specification compliance [4].

Clinical trial design and monitoring (e.g., subject enrollment/selection, patient drop out, and trial monitoring), manufacturing (e.g., automated manufacturing, personalized manufacturing, and correlating manufacturing errors to set parameters), and quality assurance/quality control (e.g., electronic lab notebooks) use AI to improve decision-making as well [4].

In post-marketing, AI could also be used to analyze data to predict a new indication or usage of the drug by using real-world data collected after the drug was marketed and/or in combination with the pre-marketing data. The drug’s new indication or usage may extend the drug’s patent expiration date.

AI-associated processes, such as data collection, processing, analysis, storage, and utilization of algorithms to derive decisions, would require guardrails to assure data quality and integrity so that the results are accurate, reliable, trustworthy, and explainable.

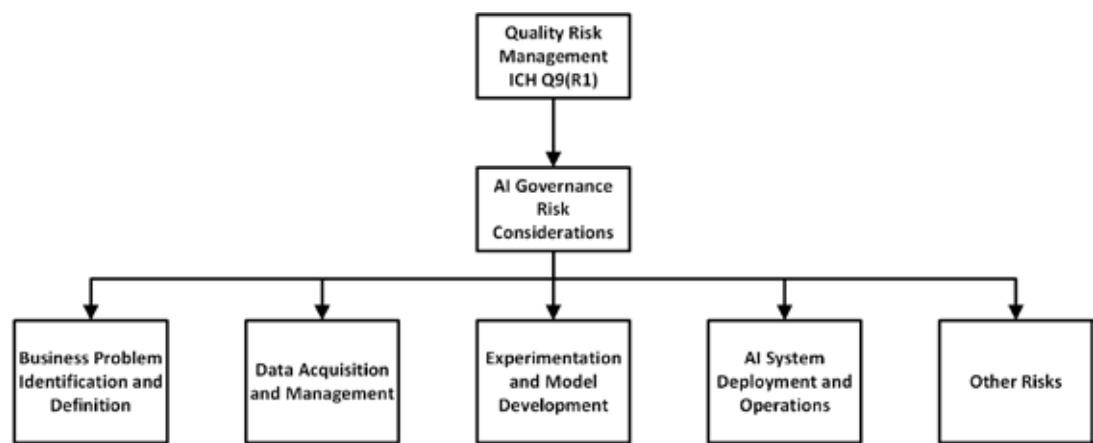
AI GOVERNANCE

AI governance is a system of laws, regulations, policies, controls, frameworks, standards (e.g., ISO/IEC Standard 42001, “Information Technology Artificial Intelligence — Management System”), practices, and processes at international, national, and organizational levels to manage, regulate, and optimize software application development, deployment, and usage within an organization. AI governance allows AI-applicable technology stakeholders to manage, regulate, optimize, implement, and oversee the use of AI technology [5]. It also helps manage associated risks to ensure AI aligns with stakeholders’ objectives, is developed and used responsibly and ethically, and complies with applicable requirements [5].

AI governance could further set ethical principles by controlling how AI is developed, implemented, and used. Consistent, end-to-end AI governance would enhance internal capabilities through methodologies and tools that address critical ethical requirements such as accountability, fairness, privacy, transparency, and robustness.

The recommended approach for AI governance throughout this document could be seamlessly integrated with *ISPE GAMP® 5*

Figure 1: AI-specific risks as outlined in ICH Q9.



Guide: *A Risk-Based Approach to Compliant GxP Computerized Systems (Second Edition)* [6]. It is recommended to implement operational processes, as defined in *GAMP® 5, Second Edition* Appendix O - Introduction to Operation Appendices (e.g., corrective and preventive action), and quality risk management processes.

Quality Risk Management

Quality risk management (QRM) processes as per “ICH Harmonised Tripartite Guideline Q9: Quality Risk Management” (e.g., risk assessment, risk control, and risk review) must be implemented [7]. However, the following risks specific to AI would need to be considered [8]:

The implementation and/or use of the policies and procedures noted throughout this article should be conducted considering a risk-based approach, with only a subset of the documents recommended being implemented for lower-risk systems. Overall risk should be determined by impact to patient, product, and data integrity.

A quality management system (QMS) in alignment with the Pharmaceutical Inspection Convention Pharmaceutical Inspection Co-Operation Scheme (PIC/S) “PIC/S Guide to Good Manufacturing Practice for Medicinal Products - Part I” [9] and Code of Federal Regulations - Part 820 - Subpart B - Quality System Requirements [10] is recommended to be in place. This is to provide the quality oversight needed to provide overall assurance to the output generated by the AI (i.e., ML, DL, or generative) systems used to aid in developing, manufacturing, or distributing a drug product.

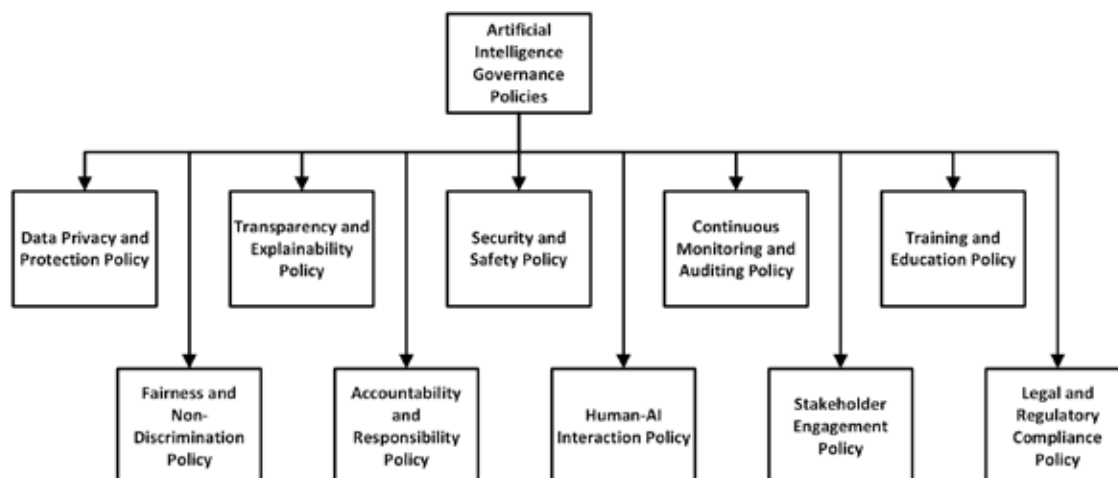
AI GOVERNANCE POLICIES

Creating policies for an AI governance framework is essential to ensure that AI systems are developed, deployed, and managed responsibly and ethically. It ensures the systems are tested and validated for GxP intended use and comply with data governance

Table 1: A categorical list of AI risks and topics for consideration.

Risk Area	Topics for Consideration
Business problem identification and definition	<ul style="list-style-type: none">• Regulatory compliance risk• Missing ethics risk• Risk of unexpected consequences• Risk of missing risk rating and classification• Missing security and privacy requirements risk• Inadequate/improper business process design
Data acquisition and management	<ul style="list-style-type: none">• Data set selection or bias risk• Data quality and missing data risk• Data labeling risk• Missing regulatory and compliance data check• Data privacy risk• Adversarial attack risk (cyber)• Data ownership and stewardship• Data standards and architecture
Experimentation and model development	<ul style="list-style-type: none">• Lack of sensitivity and scenarios analysis risk• Model assumption, limitations, selection, transparency, fairness, evaluation, use, and impact risk• Missing lineage risk
AI system deployment and operations	<ul style="list-style-type: none">• Human supervision risk• Technology integration and scalability risk• Model performance and behavior change risk• Fallback procedure risk• Adversarial attacks (cyber)• Cybersecurity
Other risks	<ul style="list-style-type: none">• Bias and fairness• Lack of transparency and regulation• Security and ethical concerns• Data privacy• Dependence and reliability

Figure 2: AI governance policies.



and cybersecurity policies and regulations. The following are some key policies that organizations should consider implementing as part of their AI governance framework [11–13].

Data Privacy and Protection

Establish a policy for AI systems. Ensure compliance with data protection laws such as collecting, storing, and processing data used in the General Data Protection Regulation, California Consumer Privacy Act, and EU AI Act. Define obtaining informed consent, anonymizing data, and protecting sensitive information.

Fairness and Nondiscrimination

Develop a policy to prevent and mitigate bias and discrimination in AI systems. It includes conducting regular audits to identify and address sources of bias in data, algorithms, and models. Establish guidelines for fairness testing and validation.

Transparency and Explainability

Create a policy to ensure that AI systems are transparent and explainable. Provide documentation and explanations of AI models, decision-making processes, and outcomes. Make sure that stakeholders can understand and trust AI systems.

Accountability and Responsibility

Define clear roles and responsibilities for developing, deploying, and managing AI systems. Assign accountability for addressing issues such as bias, discrimination, and errors in AI systems. Establish procedures for reporting and addressing AI-related incidents.

Security and Safety

Implement a policy to ensure the security and safety of AI systems. Protect AI systems from unauthorized access, tampering, and attacks. Establish guidelines for secure coding practices, vulnerability assessments, and incident response.

Human–AI Interaction

Develop a policy for human–AI interaction, including guidelines for human oversight and intervention in AI systems. Define procedures for handling situations where AI systems provide incorrect or harmful recommendations.

Continuous Monitoring and Auditing

Create a policy for continuous monitoring and auditing of AI systems. Establish the tracking of the performance and impact of AI systems over time. Regularly review and update AI models to meet ethical and legal standards.

Stakeholder Engagement

Establish a policy for engaging with stakeholders, including employees, customers, regulators, and the public. Gather input and feedback on AI systems to ensure alignment with the values and needs of the community.

Training and Education

Implement a policy for training and education on AI ethics, best practices, and governance. Provide resources and support for employees and other stakeholders to build a culture of responsible AI.

Legal and Regulatory Compliance

Develop a policy to ensure AI systems comply with relevant laws and regulations. Establish conducting legal and regulatory assessments of AI systems.

AI GOVERNANCE FRAMEWORK PROCEDURES

Creating procedures for an AI governance framework is essential to operationalizing the policies established in the framework. This section includes essential procedures that organizations should consider implementing as part of their AI governance framework [11–13].

Data Management

Establish a procedure for collecting, storing, and processing GxP data used in AI systems. Define steps for data cleaning, validation, and transformation. Ensure that data is representative, unbiased, and of high quality.

Bias Detection and Mitigation

Develop a procedure for identifying and addressing bias in AI systems. Define steps for conducting fairness audits, analyzing sources of bias, and implementing bias mitigation techniques.

Model Development and Validation

Create a procedure for developing and validating AI models. Define feature selection, model training, hyperparameter tuning, and model evaluation steps. Ensure that models are robust, accurate, and generalizable.

Continuous training and continuous testing (or retesting) should be considered where the addition of data, fine-tuning of the algorithm(s), and/or retraining of the model occur. In addition, new patterns of bias should be investigated.

Transparency and Explainability

Establish a procedure for providing transparency and explainability in AI systems. Define steps for generating explanations, visualizations, and documentation of AI models and decision-making processes. The verification of AI-generated content would need to be considered by cross-referencing the information with credible sources; verifying facts, statistics, and claims against multiple trustworthy references; monitoring the AI's performance and reviewing its output regularly; and assessing the generated content for potential biases, errors, or inconsistencies [14].

Security and Safety

Implement a procedure for ensuring the security and safety of AI systems. Define steps for conducting vulnerability assessments, implementing secure coding practices, and responding to security incidents.

Human–AI Interaction

Develop a procedure for human–AI interaction, including human oversight and intervention in AI systems. Define steps for handling situations where AI systems provide incorrect or harmful recommendations.

Monitoring and Auditing

Create a procedure for continuous monitoring and auditing of AI systems. Define steps for tracking performance metrics, conducting impact assessments, and updating AI models.

Stakeholder Engagement

Establish a procedure for engaging with stakeholders, including employees, customers, regulators, and the public. Define steps for

gathering input, addressing concerns, and incorporating feedback into AI systems.

Training and Education

Implement a procedure for providing training and education on AI ethics, best practices, and governance. Define steps for conducting training sessions, providing resources, and assessing knowledge and skills.

Incident Response

Develop a procedure for responding to AI-related incidents, such as bias, discrimination, and errors. Define steps for reporting incidents, conducting investigations, and implementing corrective actions.

Legal and Regulatory Compliance

Create a procedure for ensuring compliance with relevant laws and regulations. Define steps for conducting legal and regulatory assessments, obtaining approvals, and maintaining documentation. It's recommended to add a compliance issue escalation process (including responsibilities) to this procedure.

MLOps PROCEDURES

MLOps is a set of practices that unifies ML system development and operations (Ops). It aims to automate the end-to-end ML life cycle, ensuring faster experimentation, deployment, reproducibility, and monitoring. For a robust MLOps framework, several procedures should be established [15]. By establishing these procedures within an MLOps framework, organizations can streamline the ML life cycle, ensuring faster deployments, scalability, reproducibility, and maintainability of ML systems.

Version Control

Use version control systems for code, data, and model artifacts, and to ensure traceability of changes and facilitate collaboration among team members.

Data Management

This procedure should define protocols for data collection, storage, preprocessing, and validation; implement data versioning to track changes and ensure reproducibility; and monitor data for drifts or anomalies that might affect model performance.


Model Development and Validation

This procedure should establish coding and architecture standards for model development, implement automated testing for model validation, and use techniques like cross-validation to assess model performance. It should be used in addition to computer validation (e.g., *GAMP® 5, Second Edition* expectations).

Continuous Integration and Continuous Deployment (CI/CD)

This procedure should automate the integration of new code, data, or model changes. It should ensure automated testing at

Figure 3: The relationship between regulation and operations (MLOps) in AI governance.



each integration step and automate the deployment of validated models to production.

Model Monitoring and Logging

This procedure should monitor deployed models for performance degradation or drift and log model predictions, inputs, and anomalies for traceability and debugging.

Model Retraining and Fine-Tuning

This procedure should define criteria for when and how models should be retrained and automate the retraining process using updated data or when performance drops below a threshold.

Experiment Tracking and Management

This procedure should use tools to track and manage multiple experiments, hyperparameters, and results. It should also ensure reproducibility by logging experiment details and outcomes.

Infrastructure and Environment Management

This procedure should define protocols for provisioning, managing, and scaling infrastructure resources. It should also ensure consistency in development, testing, and production environments using containerization or virtualization.

Model Interpretability and Explainability

This procedure should implement tools and practices to interpret and explain model predictions and should ensure stakeholders understand model decisions, especially in critical applications.

Model Security and Compliance

This procedure should establish protocols for model security, including access controls and encryption, and ensure compliance with data privacy regulations and industry-specific standards.

Feedback Loop

This procedure should implement mechanisms to gather feedback

from end-users or system interactions and then use feedback to improve models and address any issues or concerns.

Rollback and Disaster Recovery

This procedure should define protocols for rolling back deployments in case of failures or issues and ensure backup and recovery mechanisms for data, code, and models.

Collaboration and Communication

This procedure should facilitate collaboration among data scientists, ML engineers, DevOps, and other stakeholders. It should ensure clear communication channels for updates, issues, or changes in the MLOps pipeline.

MLOPS AND AI GOVERNANCE RELATIONSHIP

MLOps and AI governance are essential for responsible, effective AI development and deployment. They are closely related and complement each other but focus on different aspects of the AI life cycle. The following sections explain how they are related [16].

Relationship

MLOps focuses on the operational aspects of ML, including model development, deployment, monitoring, and maintenance. It aims to automate and streamline the ML life cycle, improve collaboration between teams, and ensure the reproducibility and scalability of models.

AI governance focuses on AI's ethical, legal, and social aspects, including fairness, transparency, accountability, and privacy. It aims to establish guidelines, policies, and procedures for the responsible and ethical development, deployment, and management of AI systems. This incorporates the regulatory expectation for thorough and reliable AI governance.

MLOps and AI governance overlap in model monitoring, validation, and documentation. MLOps provides the tools and practices for implementing these activities, whereas AI governance provides the principles and standards for guiding these activities. MLOps and AI governance emphasize the importance of stakeholder engagement, continuous improvement, and compliance with laws and regulations. They work together to ensure that AI systems are aligned with the values and needs of the community.

Implementation

Practices of AI governance and MLOps are considered together when developing the overall AI paradigm. This is to ensure governance is practical and implementable, and that complete and operational paradigms are capable of conforming to governance.

AI governance can be integrated into MLOps by incorporating ethical considerations into operational workflows. For example, MLOps can include bias detection and mitigation steps in model development pipelines, and model monitoring systems can track fairness and performance metrics.

Some examples (though not an exhaustive list) of procedural and technical controls, or guardrails, include:


- For US FDA-regulated drug companies: “21 CFR Part 11, Data Integrity and Compliance with Drug CGMP - Questions and Answers,” *GAMP® 5, Second Edition*, and associated ISPE Good Practice Guides
- For EU-regulated drug companies: “EudraLex Volume 4 Annex 11”
- For UK-regulated drug companies: “Medicines & Healthcare Products Regulatory Agency (MHRA) ‘GXP’ Data Integrity Guidance and Definitions”

Benefits

MLOps and AI governance provide a comprehensive approach to AI development and deployment. MLOps ensures that AI systems are efficient, reliable, and scalable, whereas AI governance ensures that AI systems are ethical, transparent, and accountable. By combining MLOps and AI governance, organizations can achieve operational excellence and ethical responsibility in AI. It can lead to better AI outcomes, more significant AI impact, and higher AI trust.

MLOps and AI governance are closely related and complementary to the AI life cycle [17]. They work together to ensure that AI systems are operationally effective and ethically responsible. Organizations should integrate MLOps and AI governance to achieve the best AI development and deployment results.

CONCLUSION

Effective AI governance implementation requires collaboration across different departments and levels of an organization. It involves establishing a governance framework, defining roles and responsibilities, implementing technology solutions, and fostering an AI-driven culture. By adhering to AI governance principles, organizations can maximize the value of their AI assets while mitigating risks associated with poor AI quality. This includes escalation of GxP compliance issues that could impact quality attributes and data integrity principles and their resolution. 

References

1. Amazon. “MLOps.” Accessed 25 February 2024. <https://aws.amazon.com/what-is/mlops/>
2. ISO/IEC Standard 22989. “Information Technology — Artificial Intelligence — Artificial Intelligence Concepts and Terminology.” July 2022. www.iso.org/standard/74296.html
3. National Institutes of Standards and Technology. “Glossary - Cybersecurity.” Accessed February 2024. <https://csrc.nist.gov/glossary/term/cybersecurity>
4. Debleena, P., G. Sanap, S. Shenoy, D. Kalyane, K. Kalia, and R. Tekade. “Artificial Intelligence in Drug Discovery and Development.” *Drug Discovery Today* 26, no. 1 (2021):82.
5. IAPP. “Key Terms for AI Governance.” Updated November 2023. <https://iapp.org/resources/article/key-terms-for-ai-governance/>
6. International Society for Pharmaceutical Engineering. *GAMP® 5 Guide: A Risk-Based Approach to Compliant GxP Computerized Systems (Second Edition)*. North Bethesda, MD: International Society for Pharmaceutical Engineering, 2022.
7. International Council for Harmonisation of Technical Requirements for Pharmaceuticals for Human Use. “ICH Harmonised Tripartite Guideline Q9: Quality Risk Management.” Published November 2005. <https://database.ich.org/sites/default/files/Q9%20Guideline.pdf>
8. Ping, D. *The Machine Learning Solutions Architect Handbook, 2nd Ed.* Birmingham, UK: Packt Publishing Ltd., 2022.
9. Pharmaceutical Inspection Convention/Pharmaceutical Inspection Co-Operation Scheme. “PIC/S Guide to Good Manufacturing Practice for Medicinal Products - Part I.” Published August 2023. <https://picscheme.org/docview/6606>
10. US Food and Drug Administration. “CFR - Code of Federal Regulations Title 21, Part 820 - Subpart B - Quality System Requirements.” October 1996. www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfcfr/CFRSearch.cfm?CFRPart=11
11. General Accounting Office. *Artificial Intelligence, an Accountability Framework for Federal Agencies and other Entities*. Accessed 1 December 2023. www.gao.gov/assets/gao-21-519/sp.pdf
12. European Commission. *Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts*. 21 April 2021. <https://artificialintelligenceact.eu/the-act/>
13. National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. January 2023. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>
14. AIContentfy. “Quality Control: How to Verify AI Generated Content.” 6 November 2023. <https://aicontentfy.com/en/blog/quality-control-how-to-verify-ai-generated-content>
15. Microsoft. “Machine Learning Operations (MLOps) Framework to Upscale Machine Learning Lifecycle with Azure Machine Learning.” Accessed 13 December 2023. <https://learn.microsoft.com/en-us/azure/architecture/ai-ml/guide/mlops-technical-paper/>
16. Darnell, D. “MLOps and AI Governance, A Difference of Perspective.” Dataiku. 21 June 2023. <https://blog.dataiku.com/mlops-and-ai-governance-a-difference-of-perspective/>
17. IBM. “AI Model Lifecycle Management: Overview.” 9 November 2020. www.ibm.com/blog/ai-model-lifecycle-management-overview/

About the authors

Armand Mintancian is an independent consultant at Reliable Data, LLC, providing services to global pharmaceutical, biotech, and medical device companies. He has been involved with compliance, auditing, quality assurance, computer validation, and remediation of global enterprise information technology (IT) systems (e.g., enterprise resource planning, supply chain management, data analytics, Veeva [clinical, regulatory, quality], IT infrastructure, and process automation) for over 30 years. His areas of expertise are data integrity, data governance, risk management, quality management systems, and consent decree remediation. Armand holds a bachelor's degree in life sciences from the New York Institute of Technology. He joined ISPE in 1995.

Rory Budihandojo has over 40 years of industry experience in research and development, manufacturing, quality control, quality assurance, and IT. He was a co-founder and former chairman of GAMP Americas. He has various validation experience in GxP audits, remediation, and quality management system improvement. He is an expert in computer system validation (CSV) and data integrity. He joined ISPE in 1993.

John English, HCCP, is a Principal/Consultant in a private practice and has over 25 years of experience in review and validation of US FDA-regulated computer systems, both on-site and remotely. He is an expert in 21 CFR Part 11, data integrity, cloud and virtualization issues, analysis of FDA 483, and warning letters. John has served on the FDA-PDA Data Integrity Task Force and is a current member of GAMP Special Interest Group on AI and ML. He holds master's degrees from the University of Connecticut and Seton Hall University (SHU) and a certificate in US healthcare compliance from the SHU School of Law. He joined ISPE in 2002.

Orlando Lopez has worked for the past 30 years in the areas of worldwide pharmaceutical computer compliance, including US FDA and EU Annex 11 in the production and quality control systems relevant to the manufacture of medicinal products. He is a subject matter expert on CSV, e-records integrity, and worldwide computer compliance. His specialties are e-records integrity and computer trustworthiness. He joined ISPE in 1997.

Jose E. Matos has over 30 years of experience in biopharmaceutical, pharmaceutical, OTC manufacturing, global information technology (IT) system operations, engineering design firms, and chemical and petrochemical operations. His areas of expertise include instrumentation engineering, process control and automation engineering, QAe-Compliance, information technology (IT), industrial control system (ICS) cyber security, digital forensics, data recovery, and IT and cyber security compliance audits. He is a part-time university professor, teaching engineering and technology courses. He joined ISPE in 1993.

Robert (Bob) McDowall, PhD, is an analytical chemist with over 50 years of experience, including 15 years working for two pharmaceutical companies, 30 years as a consultant, and over 35 years in CSV. Bob writes the Questions of Quality column for *LCGC International* and the Focus on Quality column for *Spectroscopy*. He edited the first book on a laboratory information management system (LIMS) and is the 1997 LIMS Awardee from the LIMS Institute. He has written two books on validation of chromatography data systems and one on laboratory data integrity. He also contributed to the *ISPE GAMP® Guide: Records and Data Integrity* and four ISPE GAMP® Good Practice Guides. He is the co-author of the *USP <1058>: Analytical Instrument Qualification's* current draft and the European Compliance Academy's *Guide for an Integrated Lifecycle Approach to Analytical Instrument Qualification and System Validation*. He joined ISPE in 1998.