# A Message from Rich Apostolik, President and CEO

To our Risk and AI Candidates:

Advances in technology have increased access to and the scalability of artificial intelligence (AI) and machine learning (ML) applications across every industry sector. AI and ML are likely to be truly transformational.

As organizations evaluate how to integrate AI into their day-to-day operations, it's critical that organizational leaders and risk managers understand the potential risks, regulations, and ethical considerations associated with its use.

Developed in consultation with world-leading AI experts and senior risk practitioners, GARP's Risk and AI (RAI®) Certificate provides insight on the evolution of AI/ML methodologies and knowledge of current tools and techniques for leveraging AI in support of business decision making, identifies potential risks that arise through its use, and discusses governance frameworks that can mitigate exposure and ensure AI is deployed responsibly within organizations. Exclusive case studies further support the program's practice-driven approach, supplementing the RAI curriculum to enhance a candidate's experience, and providing insight on how AI techniques are being successfully applied by companies to solve real-world problems.

GARP's cutting-edge and globally unique RAI program responds to a growing need for organizations to understand the risks associated with AI as they attempt to realize its considerable benefits. A candidate successfully earning the Certificate will have objectively demonstrated the knowledge and understanding needed to assess and manage AI-related challenges.  Regardless of industry, RAI Certificate Holders, and the organizations that employ them, will be seen as responsible leaders in the use of AI. And, by earning the RAI Certificate, the successful candidate will join a diverse, global community of professionals who recognize the accelerating impact of AI and the importance of developing knowledge to effectively manage its risks.

We wish you the very best in your pursuit of the RAI Certificate and in achieving your individual or company-related objectives.

Thank you and best regards.

Sincerely,

Rich Apostolik

President & CEO

# Navigating This Course

GARP's Risk and AI (RAI®) Certificate program provides a holistic overview of AI tools and techniques, the issues managers need to consider so that Artificial Intelligence (AI) is deployed responsibly, and the potential risks posed by AI and how to manage them. Born of the growing pressure felt by firms across industries to incorporate AI into operations and business practices, our cutting-edge program symbolizes the urgency for organizations to address relevant issues driven by the deployment of AI.

All technologies expose the companies that use them to risk of various sorts. But AI introduces a novel set of risks that were quite unfamiliar until recently. The **Risk & AI – Introduction & Overview** module offers a historical perspective on AI, an overview of both machine learning (ML) methodologies and generative AI (GenAI), and an introduction to the risks associated with AI/ML.

In our **Tools and Techniques** module, we take a deeper dive into the core building blocks of AI. We categorize machine learning algorithms for supervised, unsupervised, and reinforcement learning tasks. You'll gain an understanding of data preparation techniques, transforming raw information into insights fit for modeling. We cover feature selection, visualization, and the nuanced art of balancing model accuracy with explainability – a vital skill for ensuring AI aligns with critical business decision-making.

Although understanding individual tools is important, our **Risks and Risk Factors** module takes a broader perspective. We dissect the potential unintended consequences of AI, from the amplification of existing biases to new threats to autonomy and safety. You'll gain strategies for mitigating risks related to fairness, transparency, and the complex ways errors can translate into real-world harm. As AI applications evolve, we confront the reputational risks stemming from public mistrust and the importance of proactive AI governance for building a trustworthy organization in the face of potential missteps. Finally, we examine governance considerations and challenges in the context of GenAI.

In the domain of **Ethical and Responsible AI**, we move beyond risk mitigation toward principled implementation. Here, we analyze frameworks from consequentialism, deontology, and virtue ethics, applying them to real-world scenarios with AI at their core. You'll gain an understanding of the evolving legal landscape, including the landmark EU AI Act and emerging US regulatory initiatives. This will equip you to be proactive as regulations evolve, not merely reactive – a key leadership quality in a field where the rules are still being shaped.

Governance will be a recurring theme – not as a bureaucratic burden, but a critical lever for successful implementation. In the **Data and AI Model Governance** module, you'll gain a holistic picture of both data and model governance. You'll learn strategies like designing a robust data strategy, ensuring data quality, and establishing clear processes for model development, testing, and ongoing monitoring. These principles ensure your AI efforts yield reliable insights, not unexpected liabilities or operational weaknesses.

Finally, as a supplement to the course curriculum, we provide a range of optional resources aimed at contextualizing the course content. Found throughout the curriculum as well as in the optional **Module 6**, these resources include **case studies,** articles, webcasts, **practitioner perspective videos**, and links to recommended papers, books, regulatory standards and guidance, and data resources. The case studies and other resources help bring the tools and techniques discussed in the course to life. Here, we examine specific applications in finance and risk management. We analyze how machine learning transforms traditional risk modeling, while discussing practical strategies to ensure explainability and robustness, particularly in high-stakes, regulated environments.

Within Modules 1 to 5 of the curriculum, optional resources will be identified by these icons:

| ***Icon*** | ***Descriptor*** |
|---|---|
| *X* | Case Study |
| *Y* | Video |
| *Z* | Reading |

Throughout this course, the focus is on application and action. Your goal is not to become a data scientist – but rather a business leader who can speak their language and bridge the gap between technical possibilities and real-world business needs. Mastering the interplay among AI, ethics, and risk management is critical to unlocking the transformative potential of this powerful technology in a sustainable, responsible, and ultimately profitable way.

# Author Attributions

## Module 1: AI and Risk – Introduction and Overview

**Peter Millican**, Ph.D, Professor of Philosophy, Head of Education and Outreach, Institute for Ethics in AI University of Oxford

## Module 2: Tools and Techniques

**Chris Brooks**, Ph.D, Professor of Finance, University of Bristol

**Manuela Pedio**, Ph.D, Senior Lecturer in Finance, University of Bristol

**Michael Dowling**, Ph.D, Full Professor of Finance, DCU Business School

**Jonathan Hakim**, Fintech Entrepreneur

**Amine Aboussalah**, Ph.D, Industry Assistant Professor, Department of Finance and Risk Engineering, NYU Tandon School of Engineering

## Module 3: Risk and Risk Factors

**Carina Prunkl**, Ph.D, Assistant Professor for Ethics of Technology, Utrecht University

## Module 4: Responsible and Ethical AI

**Carissa Veliz**, Ph.D, Associate Professor, Institute for Ethics in AI, University of Oxford

## Module 5: Data and AI Model Governance

**Joseph Breeden**, Ph.D, CEO, Deep Future Analytics, President, Model Risk Management International Association

**Peter Quell**, Ph.D, Board Member, Model Risk Management International Association

## Module 6: Case Studies and Practitioner Perspectives

**Amine Aboussalah**, Ph.D, Industry Assistant Professor, Department of Finance and Risk Engineering, NYU Tandon School of Engineering

**Chris Brooks**, Ph.D, Professor of Finance, University of Bristol

**Manuela Pedio**, Ph.D, Senior Lecturer in Finance, University of Bristol

**Peter Quell**, Ph.D, Board Member, Model Risk Management International Association

**Kishore Yalamanchili**, Ph.D, Director, Risk Content, GARP

**Rama Cont**, Ph.D, Head of the Oxford Mathematical and Computational Finance Group, University of Oxford

**Alexander Denev**, Ph.D, FRM, Co-Founder, Turnleaf Analytics

**Bo Xu**, FRM, Principal, AI Lead, BCG

**Thomas Dahlin**, Chief Model Risk Officer, Centennial Bank Board Member, Model Risk Management International Association

# Reviewer Attributions

## Module 1: AI and Risk – Introduction and Overview

**Michael Imerman**, Ph.D, Assistant Professor of Teaching in Finance, Paul Merage School of Business, UC Irvine

## Module 2: Tools and Techniques

**Paul Feehan,** Ph.D, Professor of Mathematics and Director, Mathematical Finance Master's Program, Rutgers University

**Samuel Po-Shing Wong**, Ph.D, Independent Researcher, Machine Learning, Mathematics and Statistics

**Michael Dowling**, Ph.D, Full Professor of Finance, DCU Business School, Dublin City University

**Qiuyan Xu**, Ph.D, FRM, Managing Director, Gravitate AI

**Amine Aboussalah**, Ph.D, Industry Assistant Professor, Department of Finance and Risk Engineering, NYU Tandon School of Engineering

## Module 3: Risk and Risk Factors

**Kay Firth-Butterfield**, CEO, Good Tech Advisory LLC; TIME Impact Honouree 2024; Former Head of AI and Member of the Executive Committee at World Economic Forum

**Michael Imerman**, Ph.D, Assistant Professor of Teaching in Finance, Paul Merage School of Business, UC Irvine

**Raghurami Etukuru**, Ph.D, FRM, Founder and Principal AI Scientist, AISCIENCES.AI

## Module 4: Responsible and Ethical AI

**Kay Firth-Butterfield**, CEO, Good Tech Advisory LLC; TIME Impact Honouree 2024; Former Head of AI and Member of the Executive Committee at World Economic Forum

**Raghurami Etukuru**, Ph.D, FRM, Founder and Principal AI Scientist, AISCIENCES.AI

## Module 5: Data and AI Model Governance

**Michael Imerman**, Ph.D, Assistant Professor of Teaching in Finance, Paul Merage School of Business, UC Irvine

**Evan Sekeris**, Ph.D, Chief Model Risk Officer, Capital One

**Manish Chakrabarti**, Ph.D, Head of Model Governance, Vanguard

**Jennifer Voitle**, FRM, Model Risk Management Expert

## Module 6: Case Studies and Practitioner Perspectives

**Satyajit Karnick**, Ph.D, Lecturer in Finance, Scheller College of Business, Georgia Tech

**Paul Feehan**. Ph.D, Professor of Mathematics and Director, Mathematical Finance Master's Program, Rutgers University

**Samuel Po-Shing Wong**, Ph.D, Independent Researcher, Machine Learning, Mathematics and Statistics

**Jennifer Voitle**, FRM, Model Risk Management Expert

**Glenn Perillo**, Chief Risk Officer, Libremax Capital

# Risk and AI Advisory Committee

**Joseph Breeden**, Ph.D, CEO, Deep Future Analytics; President Model Risk Management International Association

**Agostino Capponi**, Ph.D, Professor, IEOR, Director of the Center for Digital Finance and Technologies, Columbia University

**Manish Chakrabarti**, Ph.D, Head of Model Governance, Vanguard

**Rama Cont**, Ph.D, Head of the Oxford Mathematical and Computational Finance Group, University of Oxford

**Lucy da Piedade**, Head of the Group Controls Office, Deutsche Bank

**Charles Currat**, Ph.D, FRM, Head of Quantitative Modeling, Wells Fargo Investment Institute

**Alexander Denev**, Ph.D, FRM, Co-Founder, Turnleaf Analytics

**Chris Donohue**, Ph.D, Managing Director, GARP Benchmarking Initiative (GBI), GARP

**Kenneth Doucet**, Senior Vice President, Content, Information and Continuing Education, GARP

**Raghurami Etukuru**, Ph.D, FRM, Founder and Principal AI Scientist, AISCIENCES.AI

**Kay Firth-Butterfield**, CEO, Good Tech Advisory LLC

**Lukas Kölbl**, Ph.D, FRM, Managing Director Data and AI, Accenture DACH

**William May**, Managing Director, Global Head, Certification and Educational Programs, GARP

**Peter Millican**, Ph.D, Professor of Philosophy, Head of Education and Outreach, Institute for Ethics in AI, University of Oxford

**Yogesh Mudgal**, Head of Risk, Runtime Compute, JPMorgan Chase

**Carina Prunkl**, Ph.D, Researcher, INRIA

**Anand Rao**, Ph.D, Distinguished Service Professor of Applied Data Science and Artificial Intelligence, Carnegie Melon University; Global AI Lead, PwC (retired)

**Alberto Rossi**, Ph.D, Professor of Finance and Director of the AI, Analytics and Future of Work Initiative, Georgetown University

**Rajesh Shekhar**, FRM, Founder, Numanic

**Stephen Slade**, Ph.D, FRM, Senior Lecturer in Computer Science, Yale University

**Agus Sudjianto**, Ph.D, Senior Vice President, Risk and Technology, H2O.ai

**Bo Xu**, FRM, Principal, AI Lead, BCG

**Qiuyan Xu**, Ph.D, FRM, CEO, Gravitate AI