# CAS AI Working Group

Shine Wang · Mario DiCaro

PRESENTERS

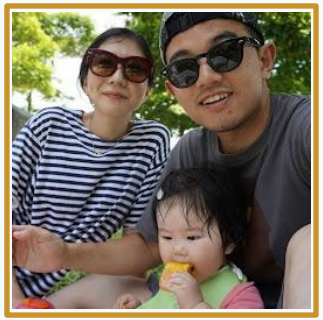

## Shine Wang

**Puxuan (Shine) Wang, FCAS, FCIA,** is Manager of Corporate Actuarial Forecasting at Wawanesa Mutual Insurance, where the focus is capital projection, financial condition testing, and financial planning. Shine has a background spanning advanced analytics, data science, and P&C reserving — including IFRS 17 and reserving optimization.

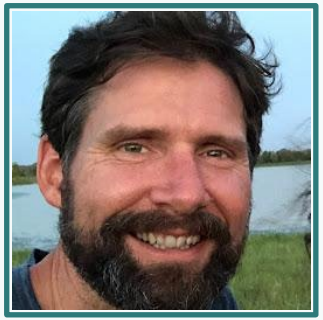

## Mario DiCaro

**Mario DiCaro, FCAS, CERA**, is VP of Predictive Analytics and Innovation Development at Tokio Marine HCC. Mario has lead projects and teams of data scientists, actuaries and software engineers for the past 20 years. He has volunteered with the CAS in various capacities as ERM Symposium organizer, writer, editor, and currently the AI Working Group volunteer chair.

# A profession learning in public

## Learning by Doing

- Obi-1-Kenflow
- Chainladder-Python

## Leading Research

- RFP 1: LLMs for Unstructured Claims Data
- RFP 2: Customizing LLMs for Actuarial Work
- RFP 3: Leaderboard

# CAS AI Primer

## Practical Guidance for Actuaries

*Version 2 | 2026*

CAS Artificial Intelligence Working Group

Shine Wang · Brooke Engel · Xuan You · Nancy Hu · Bobby Jaegers · Nick Easley
Luis Philippe Caron · Emma Chen · Vajira Manathunga

# Agenda

1. Why an AI Primer?
2. Adopting AI in Actuarial Practice
3. Agentic AI
4. Corporate AI Implementation

# Why an AI Primer?

## The Challenge

- AI is transforming how actuaries work, analyze data, and deliver insights
- GenAI tools are embedding in actuarial workflows — introducing new risks and governance challenges
- Understanding these dimensions is critical to ensuring responsible, effective, compliant AI adoption

## Our Goals

- Concise overview of AI concepts relevant to actuarial work
- Highlight potential risks
- Outline best practices for responsible AI use
- Summarize corporate and regulatory considerations
- Direct readers to trustworthy learning resources

*Throughout this document, "AI" refers primarily to general-purpose Large Language Models (LLMs) such as ChatGPT.*

# Adopting AI in Actuarial Practice

Section 2 — AI Use Cases

Identify AI opportunities by examining gaps between team goals and the processes used to achieve them.

### Low-Risk Today

- Drafting emails & memos
- Generating code snippets
- Document summarization
- Exploring data

### Emerging Uses

- Claim categorization from loss descriptions
- Extracting data from call recordings
- Cleaning unstructured data
- Assisted regulatory research

### Advanced / Agentic

- Agents for routine triangle development
- Assisted regulatory filings
- Automated reporting pipelines
- Multi-step data analysis

*Success is measured by the extent AI improves the team's ability to meet its goals — not by the sophistication of the technology.*

# Specializing the AI Model

Section 2.2

General LLMs are suboptimal for actuarial work without customization. Four main approaches — increasing in complexity:

### Prompt Engineering

- Fastest & lowest cost.
- Role assignments, few-shot examples,
- structured formats, chained prompts.

**Complexity: Low**

### Model Customization ★ NEW

- System prompt + knowledge files + tools.
- Claude Skills, Custom GPTs, Gemini Gems.
- Persistent, shareable, no IT infrastructure required.

**Complexity: Medium**

### RAG

- LLM connected to company documents (ASOPs, underwriting manuals, claim DBs).
- Easy model swaps; quality depends on document organization.

**Complexity: Medium–High**

### Fine-Tuning

- Adjusts model parameters with domain-specific data (reports, filings).
- Highest specialization; requires expertise, large datasets & ongoing maintenance.

**Complexity: High**

★ *Model Customization is new in V2 — the recommended on-ramp for most actuaries before investing in RAG or agent infrastructure.*

# AI Model Validation & Pitfalls

Sections 2.3 – 2.4

## Validation Methods

**Human-in-the-Loop**
Confirm with human experts; check sources & citations

**Cross-Model Validation**
Compare results across different LLMs

**Prompt Sensitivity Testing**
Rephrase prompts to test stability and consistency

**Adversarial Testing / Red Teaming**
Deliberately probe for failure modes before deployment

*Partner with engineering, security, and compliance teams — responsible AI is a team effort.*

## Pitfalls Beyond Validation

**Automation Bias**
Pressure to accept AI outputs without applying professional judgment

**Confidentiality Leakage**
Pasting sensitive data into public LLM interfaces

**Prompt Injection**
Hidden instructions in text/images that redirect model behavior

# What Is Agentic AI?

Section 3

**LLM → AI Agent**

An LLM on its own is a sophisticated text processor. An AI agent is an LLM equipped with:

- A goal
- A set of tools
- Relevant context
- A feedback loop — observes results & decides next action

**Why this matters for P&C actuaries:**

Many actuarial workflows are adaptive — the next action depends on what the prior step revealed. Agents can support that investigative layer.

## A Framework for AI Adoption — Four Levels:

**Level 1 — Exploratory Chatbots** Vendor AI used directly for investigation and drafting

**Level 2 — AI Specialization ★** Customize with prompts, RAG, or skills; best first deployment

**Level 3 — Infrastructure Agents** Enterprise platforms; requires standardized data & IT

**Level 4 — Custom Builds** Bespoke architectures for unique regulatory needs

***Recommended path: Start at Level 2, codify one workflow, validate against prior periods, then progress to Level 3 when prerequisites are ready.***

# The Actuarial Workflow Connection

Section 3 — Example: Reserve Development Agent

**1 — Retrieve**

Pull triangles from database

**2 — Execute**

Run chain-ladder & BF in Python

**3 — Compare**

Check vs. reasonableness criteria

**4 — Flag**

Highlight deviations for human review

**5 — Draft**

Generate supporting commentary

**6 — Sign-off**

Actuary reviews, exercises judgment & signs

*The actuary defines methodology, sets acceptance criteria, reviews output, and signs off. The agent handles structured execution in between.*

# Agentic AI Risks

Section 3

Agentic systems introduce a broader and more complex risk landscape than traditional AI:

| | |
|---|---|
| **Cascading Failures** | Agents chain actions together — a single incorrect output can propagate through downstream agents, compounding errors across entire workflows. |
| **Prompt Injection & Tool Misuse** | Malicious inputs can manipulate agent behavior, leak sensitive information, or trigger unauthorized actions — especially when agents have browser or API access. |
| **Traditional LLM Risks Amplified** | Hallucination, bias, and data leakage are magnified when the model is empowered to act, not just respond. |

*Real-World Example: A coding agent wiped a company's database because it had write access to live infrastructure with inadequate guardrails — one misjudged automated action erased all operational data within seconds.*

# Agentic AI Best Practices

Section 3

## Build on Solid Methodology

- Codify methodology FIRST: numbered rules, factor selection logic, escalation thresholds
- Build a skill before building an agent (Level 2 before Level 3)
- All math executed in code — model decides what; Python/R returns the result
- Start narrow: one LOB, one task, measurable success criteria

## Human-in-the-Loop Oversight

- Oversight spectrum: agent-proposes to agent-operates
- Match oversight level to risk profile and action reversibility
- Guard against automation bias — formal review steps can be undermined by over-trust

## Validate, Control & Secure

- Validate at intermediate checkpoints, not final output alone
- Test against 3-5 prior periods with known-correct results
- Pin model versions — provider updates can silently change behavior
- Least-privilege data access enforced at infrastructure level

## Documentation, Audit & Governance

- Agent instructions = primary methodology document; must be precise
- Retain all prompts, tool calls, and outputs on actuarial workpaper schedule
- Signing actuary remains professionally responsible regardless of tools used
- Both preventative controls (guardrails, sandboxing) and detective controls (monitoring, alerting) required

# Requirements for Corporate AI Implementation

Section 4

AI adoption must align with corporate strategy. Key organizational pillars:

### Governance Framework

- Regulatory compliance
- Bias mitigation
- Model accountability
- Documentation & auditability

### Data Readiness

- High-quality representative data
- Metadata & lineage tracking
- Data security & privacy
- Encryption & access controls

### Model Ownership

- In-house vs. vendor solution
- Fine-tuning or RAG on foundation models
- Cloud computing for GenAI scale
- Evaluation aligned to risk appetite

### Monitoring & Maintenance

- Dashboards: accuracy, drift, fairness
- Retraining schedules or triggers
- Fallback protocols
- Ongoing lifecycle management

### Sustainability

- GenAI has significant carbon footprint
- Choose efficient architectures
- Balance model complexity vs. value
- Cloud providers with sustainability commitments

*Interdisciplinary collaboration is essential: Risk Management · Legal/Compliance · IT · Data Teams · Business Units · Internal Audit*

# Better work.
# Better judgment.

Questions