



MiraiSolutions

smarter analytics - better decisions



contact@mirai-solutions.com
www.mirai-solutions.com

Agentic AI Applications in Insurance

Gustavo Martinez
Senior Manager, Mirai Solutions

Insurance

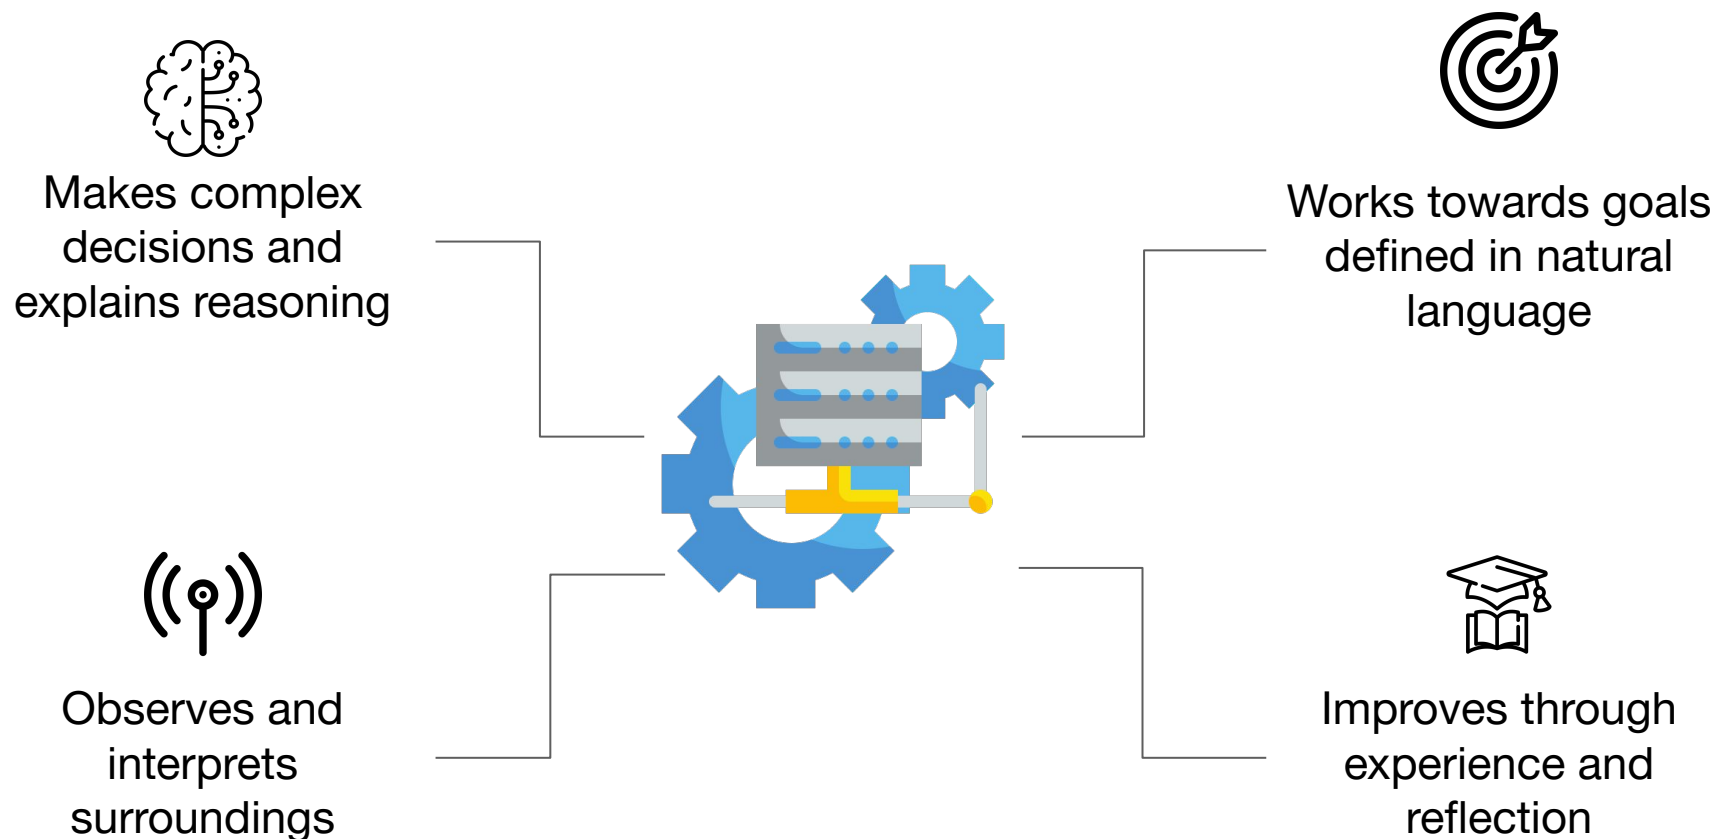
Data

Science

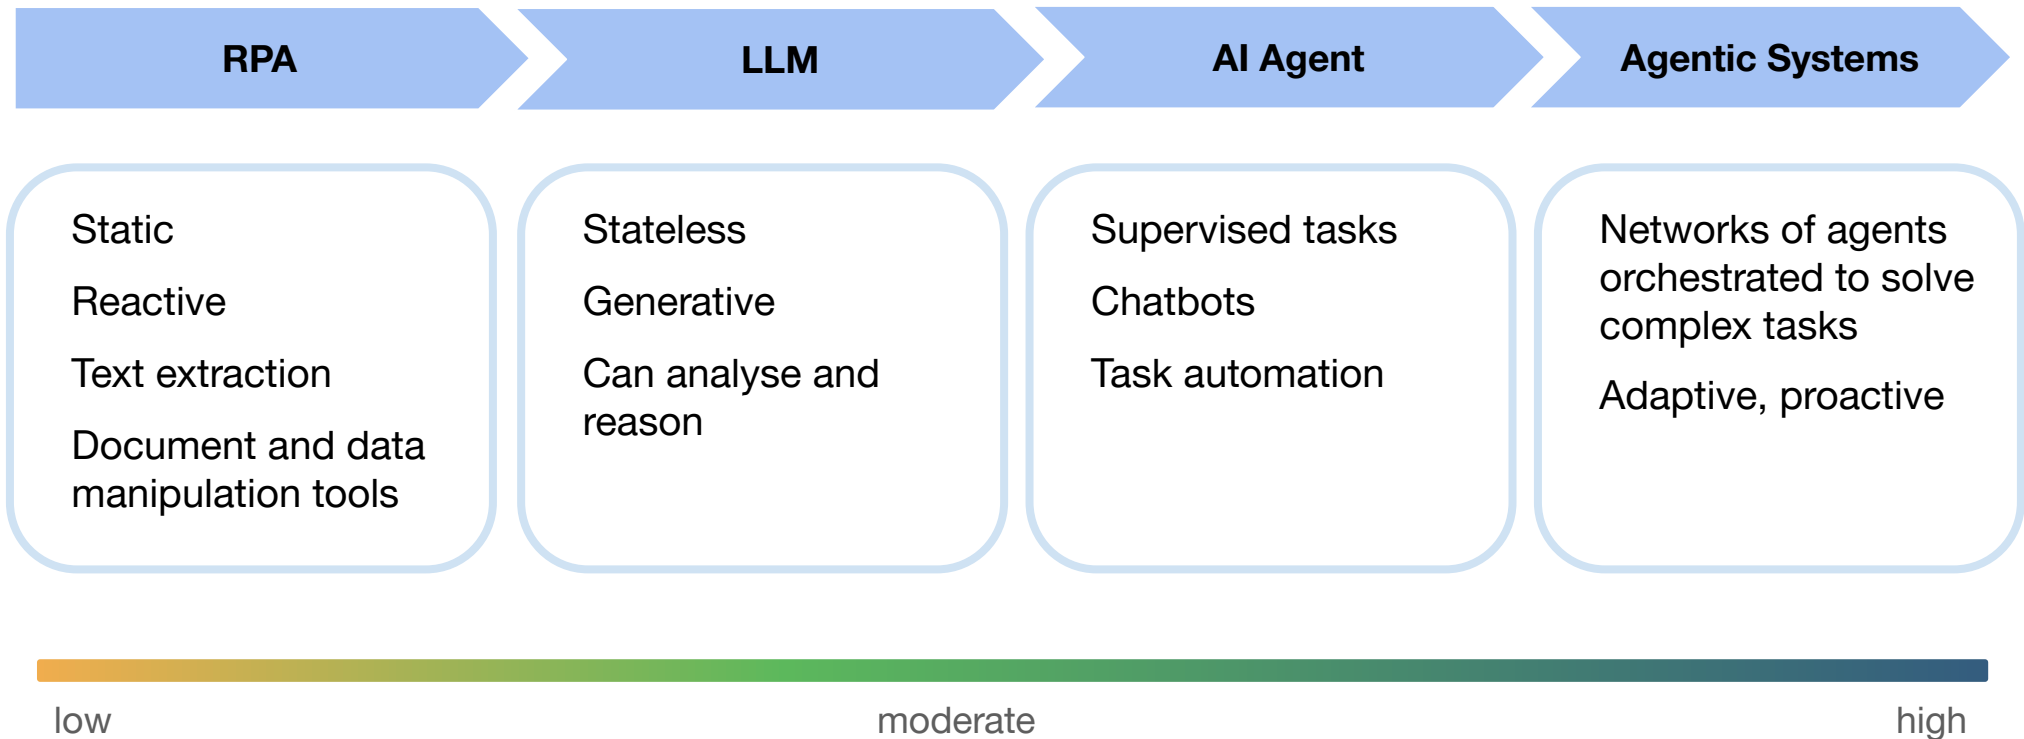
20.06.2025

Defining Agentic AI

Systems that can accomplish complex tasks by reasoning, planning and acting autonomously



The Agency Spectrum



Benefits and Applications in Insurance

Operational Efficiency

Improved Customer Experience

Enhanced Data Insights

Scalability




Claims processing

Assessment, documentation review, payments




Audit & Compliance

Regulatory reporting and documentation



Marketing

Content generation, lead management,



Customer Service

Sentiment analysis, product recommendation



Fraud Detection

Pattern recognition across claims, suspicious activities



Risk Assessment

Assessment, documentation review, payments

Reliability Challenges

Systems powered by AI must deliver **reliable outcomes** based on **unreliable AI models**

Stochastic nature

LLMs are not deterministic.

Insurance decisions require consistency and predictability.



Compound unpredictability

Combining models and tools for retrieval and generation results in systems with compound stochasticity.

Each step adds noise.



Infinite feedback loops

Agents that are unable to reflect on their findings may find themselves repeatedly invoking infinite feedback loops.



Prompting is not programming

Prompts do not work consistently across workflows.

Poor prompts can break results.

Expectations vs reality

Models can generate plausible but false information (hallucinations).

Generating results over possibly irrelevant context.

Compliance risks

Privacy regulations and data sovereignty concerns.

Traceability, explainability and accountability challenges.

Engineering Challenges

Agents give **structure** to LLMs, agentic systems give **scale** to agents.

Testing and debugging

Randomness cannot be eliminated.
Heuristic metrics are flawed.
Human evaluations are inconsistent and expensive.
Prompt tweaks can cause large results variations.
LLMs remain black boxes to a large extent

Complexity

Multi-agent systems are distributed and orchestrated.
LLMs are stateless, context is assembled.
Systems must be auditable, traceable and be highly scalable and resilient.

Security horror show

LLMs are leaky and create many security holes.



Evaluating and monitoring

Reasoning must be made visible and measurable.
Comprehensive and consistent synthetic evaluation sets are required to test reasoning steps.
Monitoring against heuristics.

Scaling makes it even harder

As natural language interfaces eliminate user barriers, the number of user interactions and system transactions will increase dramatically.
LLMs are computationally intensive and have very poor latencies.

Design Patterns for Success



Domain-Driven Development

Produce context maps and boundaries.
Develop localized workflows for each boundary.
Create reasoning layers to break complex plans into sub-tasks that can be validated



Data Strategy

Form a solid data strategy with focus on unstructured data.
Aim at producing high-quality, datasets for AI analytics.



Evaluation-Driven Development

Implement guardrails and safety boundaries as a core part of the development process.



Evaluation Strategy

Continuous testing and experimentation of different inputs (real-world, synthetic, adversarial)
Track and validate accuracy in different insurance scenarios.



Human in the Loop

Delegate decisions to humans in the workflows.



Specialization

Use dedicated agents that specialize on one area of expertise.
Restrict decisioning.



Check and Balance

Ask for second opinions from other agents and take joint decisions.
Create supervisor agents for quality control.



Data Sovereignty Controls

Implement strict boundaries for sensitive data.
Maintain clear audit trails for all AI decisions.

Adoption Roadmap

Assess Readiness


Target use cases where rule-based approaches fall short and decisions require judgment.

Evaluation Framework

Implement continuous testing across real-world, synthetic, and adversarial inputs.

Agentic Architecture

Build systems with memory, knowledge integration, and uncertainty handling.



Data Strategy

Develop frameworks for unstructured data and establish sovereignty controls.

Select AI Models

Work toward specialized agents with focused expertise areas.

Continuous Improvement

Monitor real-world performance and refine based on feedback loops.