The Twenty-Second International Conference on
Autonomic and Autonomous Systems
ICAS 2026
March 08, 2026 to March 12, 2026 - Valencia, Spain

Agentic AI and the Autonomic Computing Vision: Continuity, Evolution, and Beyond

From MAPE-K to LLM Cognitive Control

Roy Sterritt
Ulster University
r.sterritt@ulster.ac.uk
@RoysterUlster

Google Scholar
ScholarGPS®
POWERING SCHOLARLY ANALYTICS

"Short" Bio

Roy Sterritt,

Ulster University, Northern Ireland.

r.sterritt@ulster.ac.uk @RoysterUlster

- Biography—Roy Sterritt is a member of Faculty in Computing and Engineering at Ulster University. He spent several years in industry with IBM, first at their UK headquarters in Portsmouth, and then at the IBM Hursley Labs in Winchester. Initially he was a Software Developer in their KBS department but then became a Product Development Manager with responsibility for tools to support risk assessment and project management in personal and mobile environments which were used widely in the UK and US. Roy's academic research career began in 1996 when he was appointed to the first of a series of joint University of Ulster and Nortel research projects investigating parallel, automated and intelligent approaches to the development and testing of fault management telecommunications systems.

- Roy's main focus of research is Systems and Software Engineering of Autonomic (Self-Managing Computer-Based) Systems, essentially a research area developed from a call from industry to deal with the complexity and total cost of ownership of our systems of systems (IBM 2001). To date he has 250+ publications in the field including research collaborations with NASA, IBM TJ Watson Center, BT, SAP, HP and Core Systems as well as many academic partners. The research with NASA also lead to 16 US patents. He was the founding chair of the IEEE Task Force and subsequently Technical Committee on Autonomous & Autonomic Systems and elected chair of IEEE Technical Committee on Engineering of Computer-Based Systems. He has held many other IEEE roles such as; IEEE CS Publications board member, chair of the Conference Publications Operations Committee (CPOC); served on the IEEE CS Technical & Conferences Activities Board (T&C Excom and Opcom) and chaired the Conference Advisory Committee (CAC). He has been appointed to the many editorial boards including the NASA Journal on Innovations in Systems and Software Engineering, ACM Transactions on Autonomous and Adaptive Systems (TAAS), AIAA Journal of Aerospace Computing, Information, and Communication, Journal of Autonomic and Trusted Computing, and Multiagent and Grid Systems - An International Journal; and served on steering and/or program committees of the majority of the conferences in his field at some stage during the last 20+ years.

- This extensive research community activity and NASA collaboration during the *noughties* lead to the common query from colleagues and management "are you ever at home?". As such Roy scaled back his international activity during the *tens* and took on institutional roles such as; Placement Coordinator-looking after 400+ students while seeking and on year-long industrial placement; Manager of CPPD (Continuous Personal and Professional Development) including developing outreach courses and summer schools in both Computer Science and Space Science. Yet with 16 patents with NASA, Roy also took the opportunity to explore spinning out the Autonomic Research as well as continuing that research into the 2020s, in particular, through his PhD researchers.

Ulster University

IARIA

IARIA Fellow

Systems Complexity Has Outpaced Human Management

- Cloud-native, distributed, AI-infused systems
- Dynamic configuration and orchestration
- Governance & safety concerns

- We have seen this problem before

The Original Answer: Autonomic Computing

- The vision was self-managing systems.

Self-CHOP

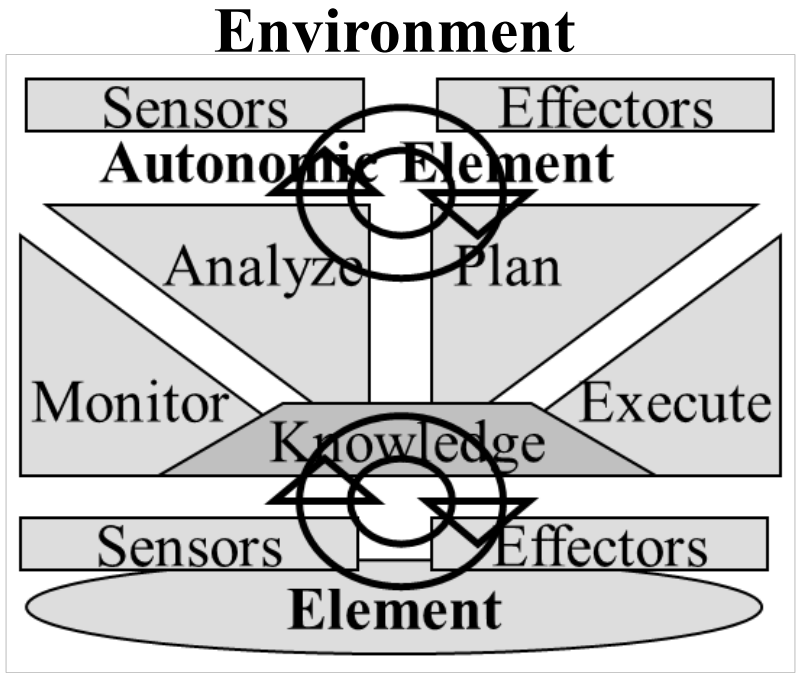

MAPE-K

- Why didn't it fully generalize?

Why Autonomic Computing Stalled

- Handcrafted policies
- Rigid models
- Limited semantic reasoning
- Poor generalization

Agentic AI: A Different Route to Autonomy

- Learned representations
- Multi-step reasoning
- Tool orchestration

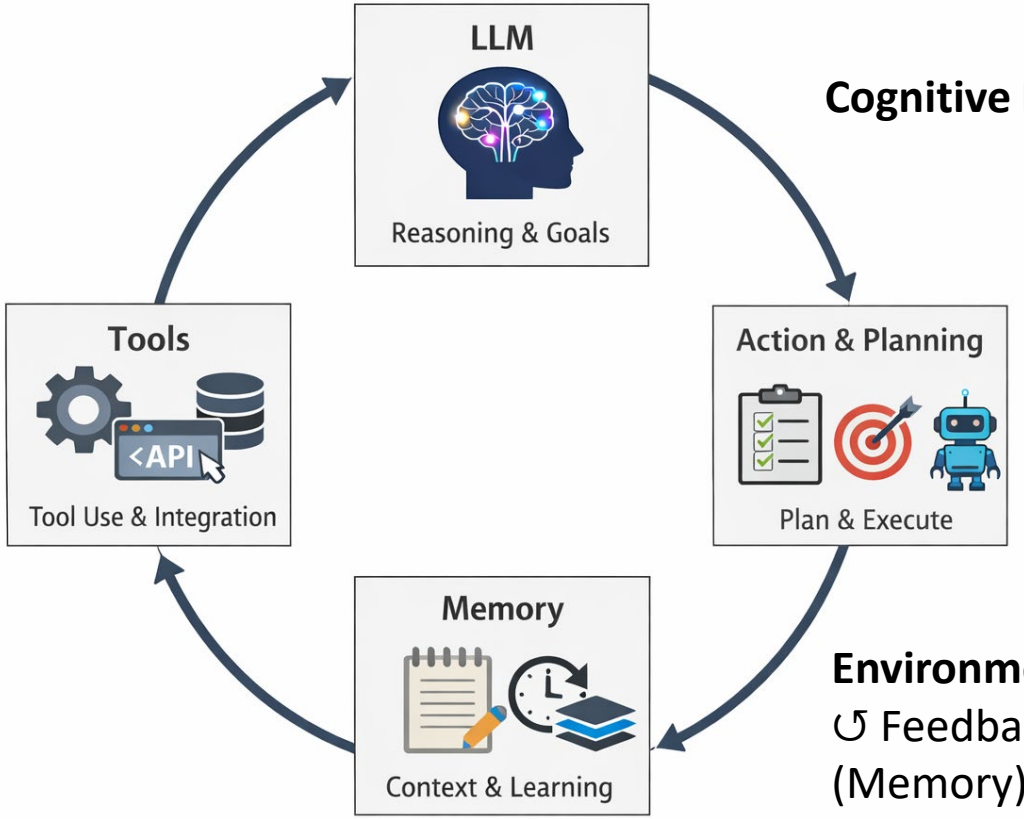

Cognitive Layer (LLM Core)

Execution Layer (Tools/APIs)

Environment
↺ Feedback into Knowledge (Memory)

Agentic AI as a Continuous Cognitive Control Loop

Agentic AI is not a replacement — it is an architectural evolution

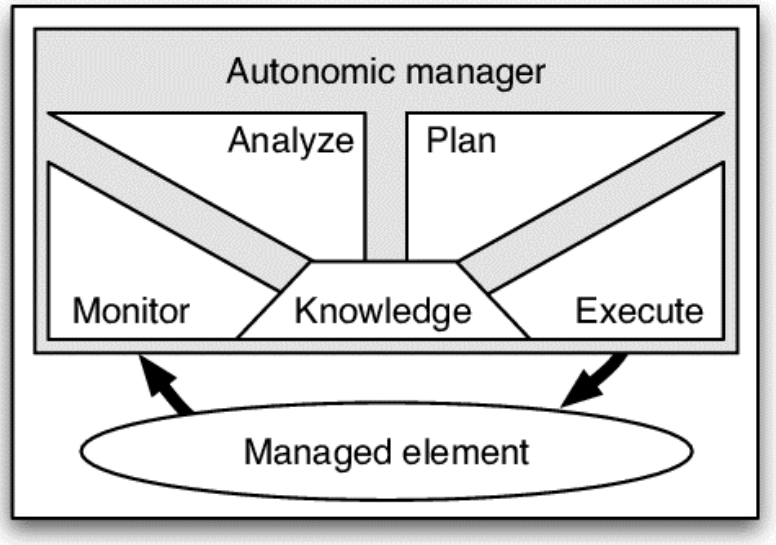

→

Implicit Generalisation

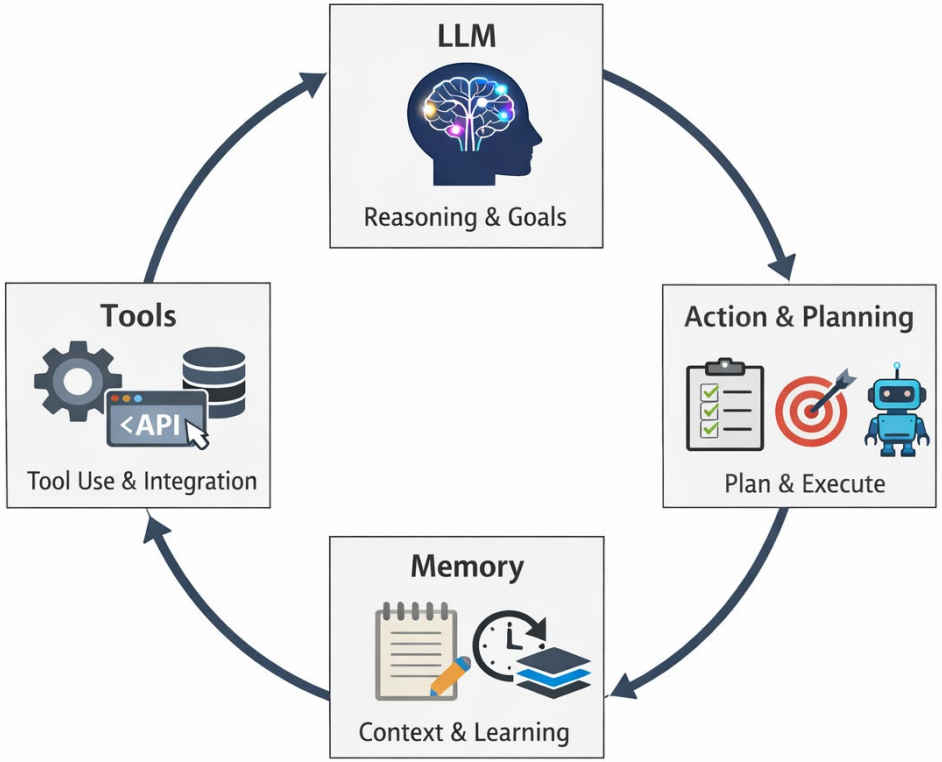

Knowledge Externalized → Knowledge Internalized

Agentic AI is not a replacement — it is an architectural evolution

MAPE-K	Agentic AI
Monitor	Tool outputs / observations
Analyze	LLM reasoning
Plan	LLM planning
Execute	Tool invocation
Knowledge	Memory (context + embeddings)

This is MAPE-K with a cognitive core.

Comparative

Dimension	Autonomic	Agentic
Control	Explicit loops	Implicit reasoning
Knowledge	Structured	Learned
Autonomy	Policy-bounded	Flexible
Governance	Rule verification	Alignment & safety

Knowledge Plane → LLM Cognitive Layer

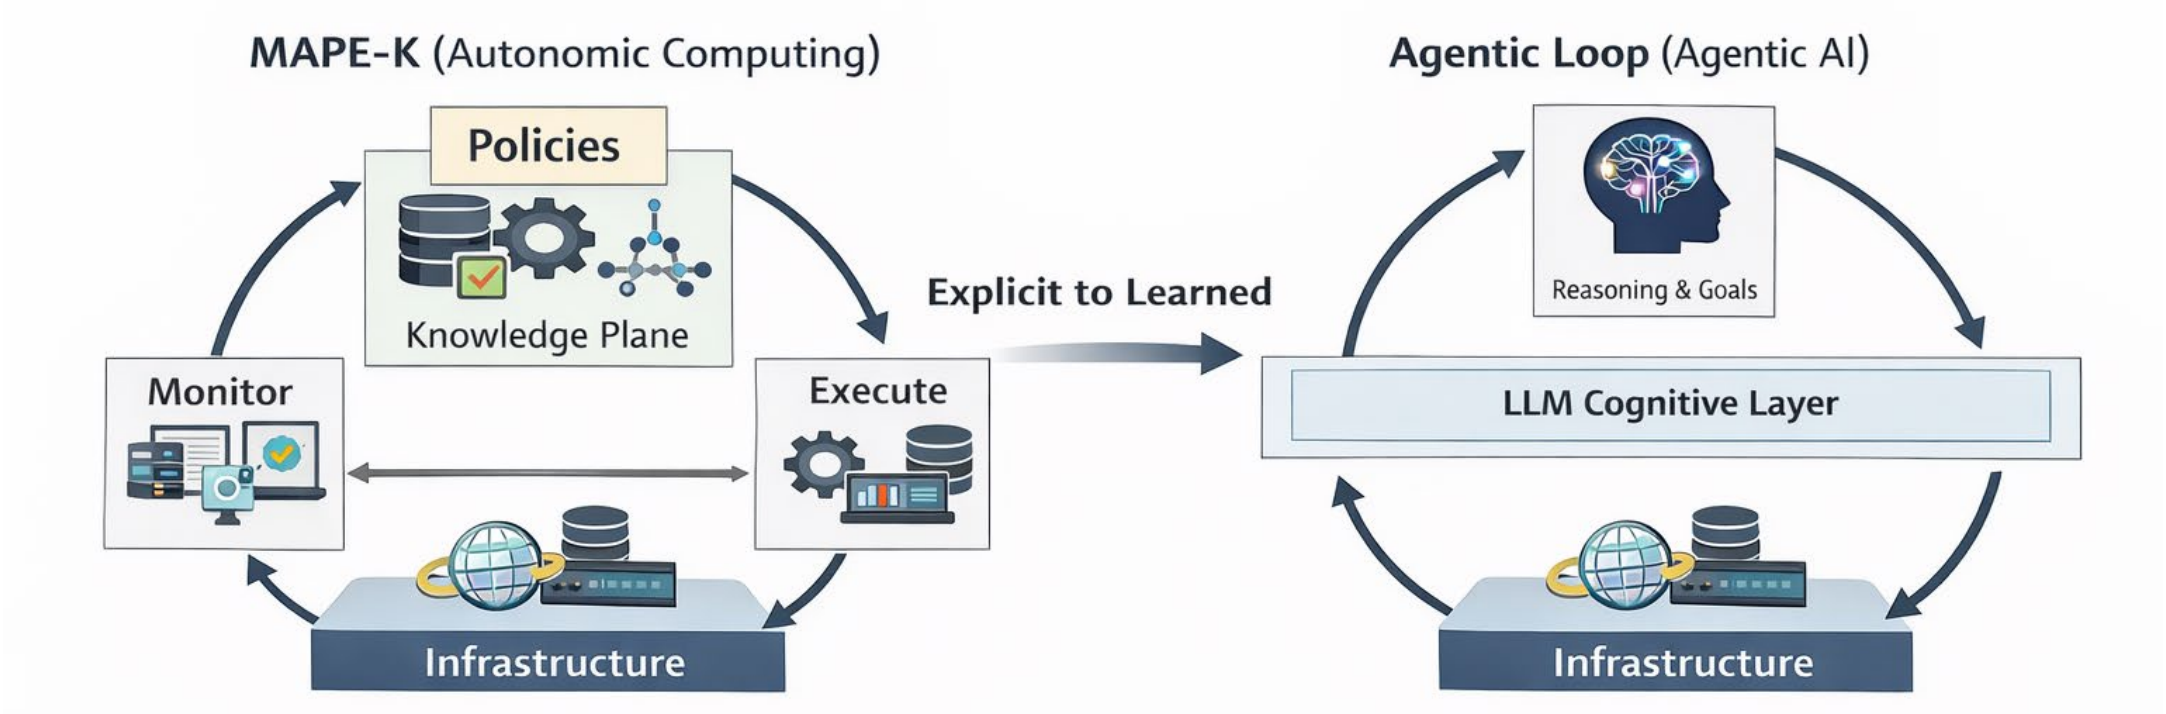

- What was explicit and engineered is now internalized and learned.

The Trade-Off: Flexibility vs Verifiability

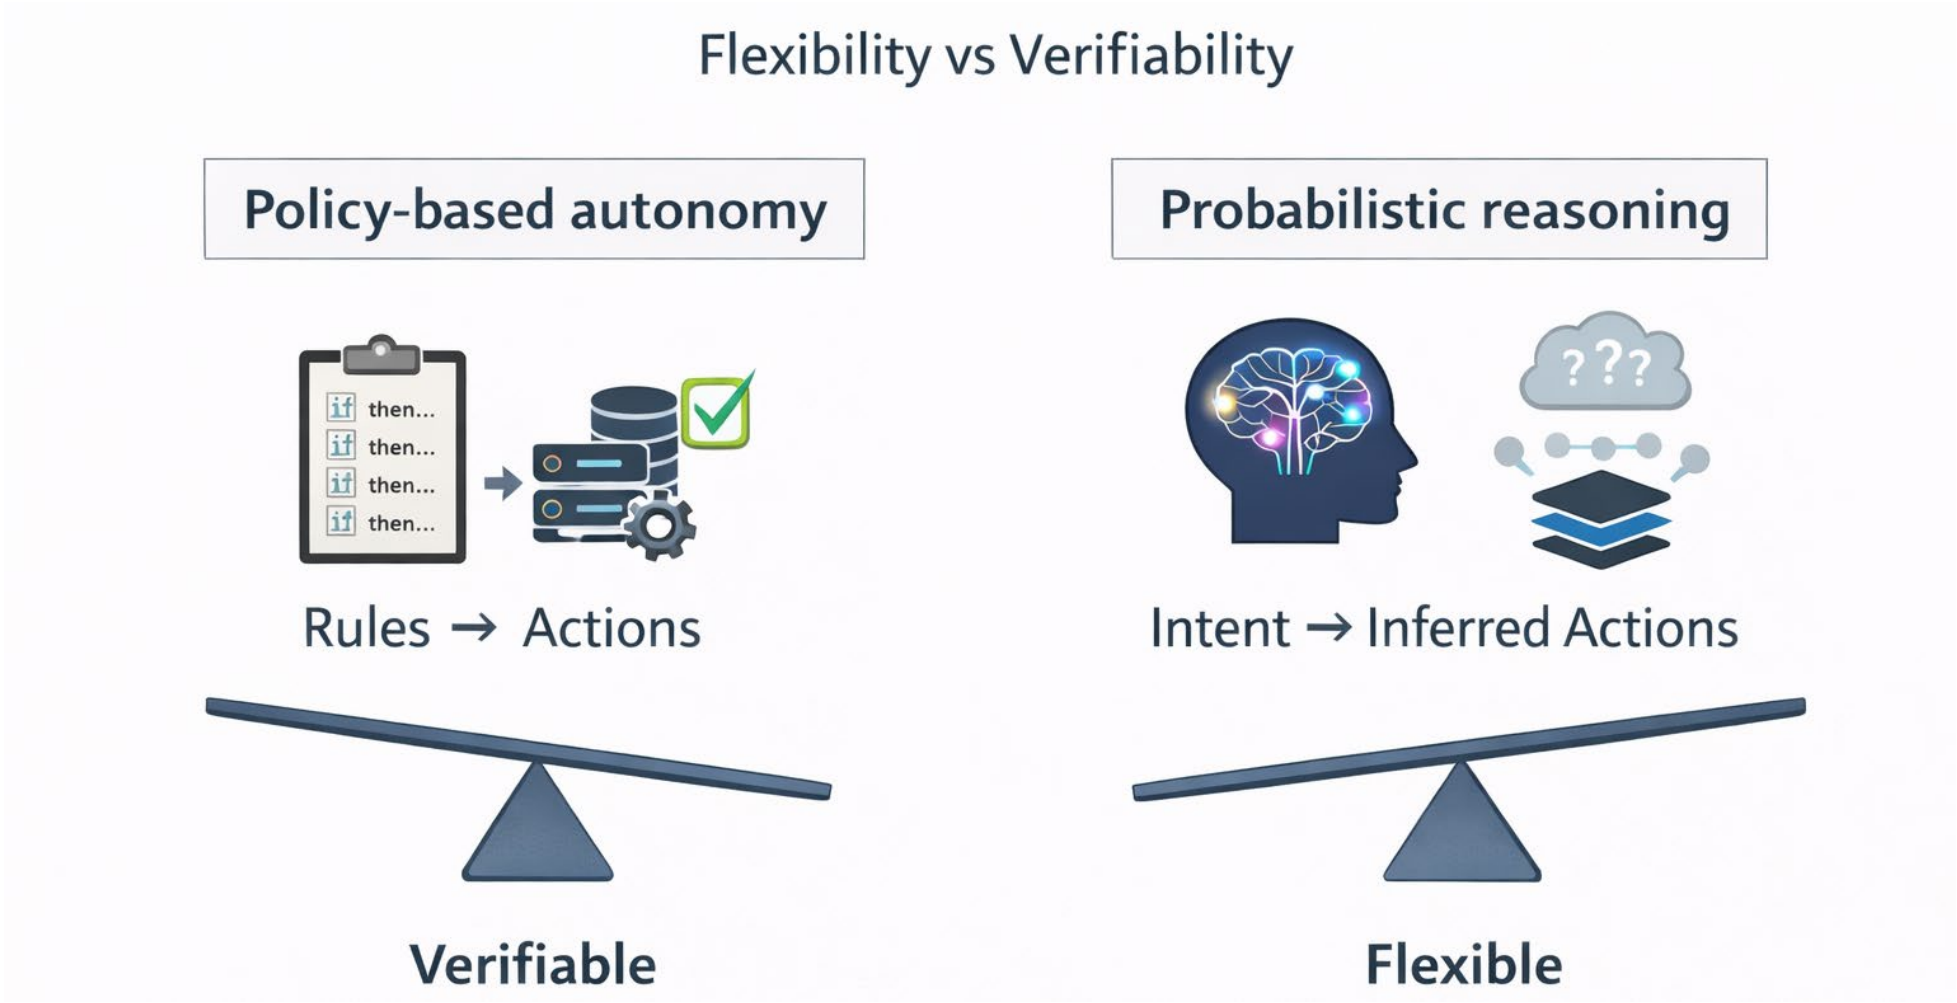

New Challenge: Goal Interpretation & Conflict

- Natural language goals
- Ambiguity & drift
- Multi-agent conflicts
- No formal conflict resolution

- Governable Agentic AI requires conflict detection mechanisms.

Apoptotic Computing: Controlled Failure as Governance

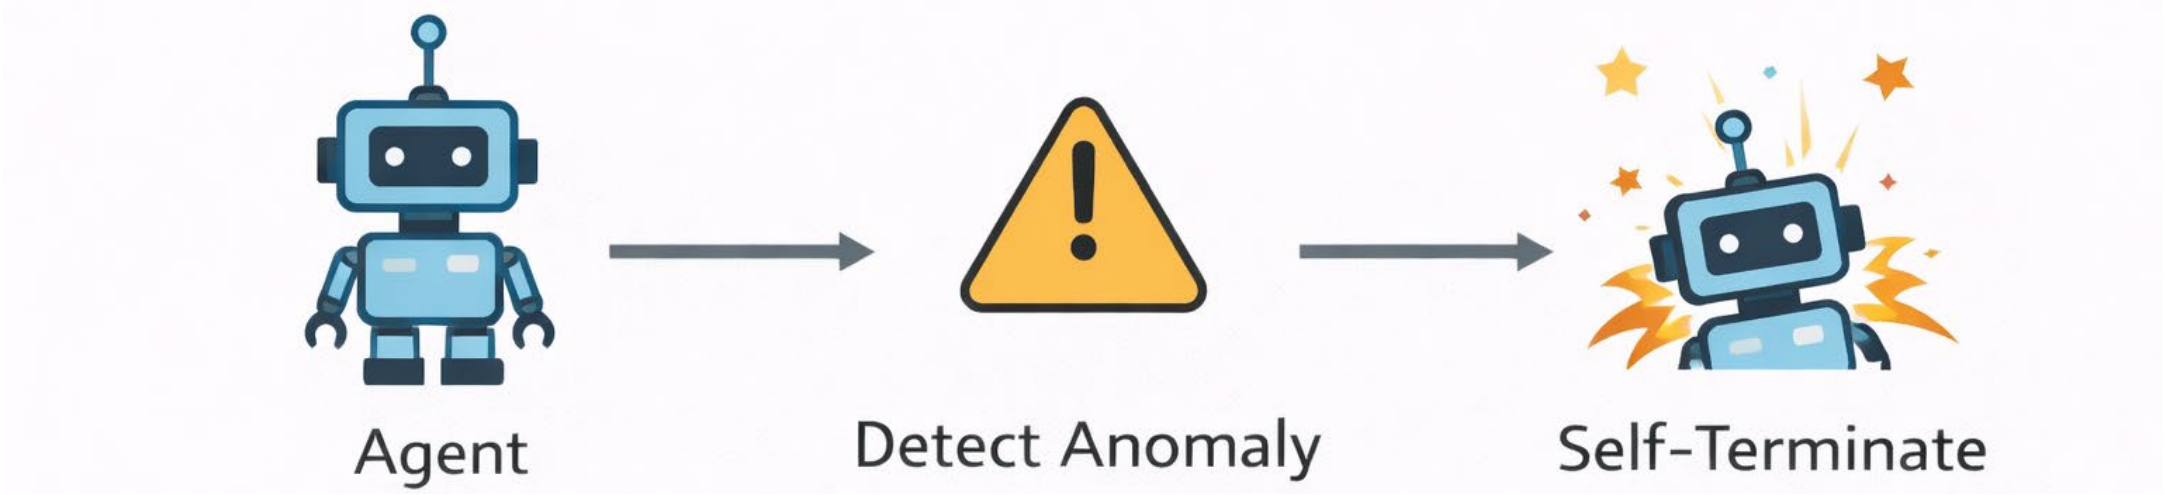

- Inspired by biological apoptosis (controlled cell death)
- Enables safe containment of misbehaving agents
- Provides termination / rollback / isolation mechanisms
- In autonomic systems, termination was a safety primitive. In Agentic AI, we may need it again.

Sterritt & Hinchey (2004, 2009)

Toward Governable Cognitive Infrastructure

- Formal assurance for probabilistic agents
- Goal conflict detection
- Multi-agent oversight
- Energy & scalability

What If We Had Today's LLMs in 2001?

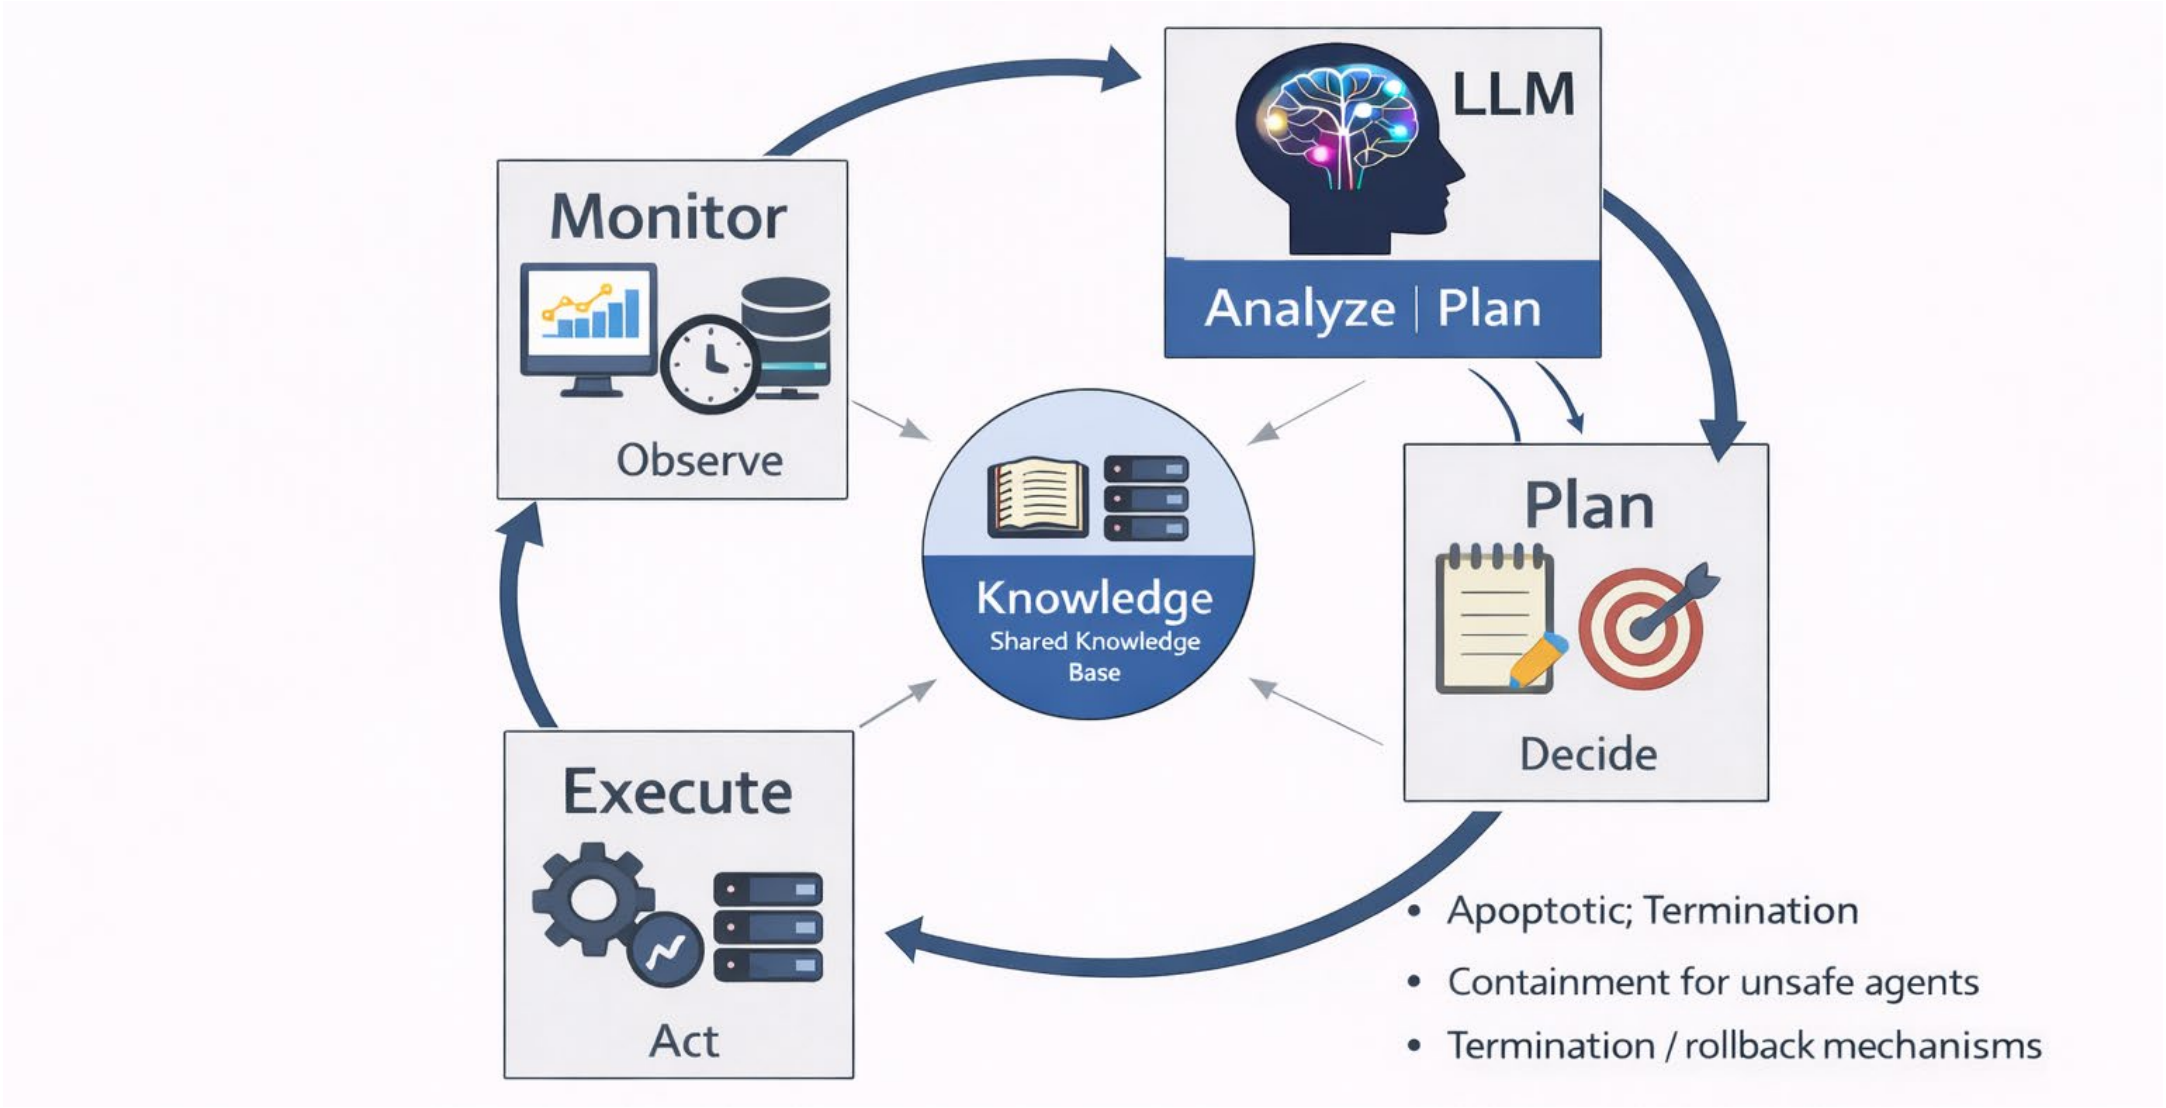

Continuity, Evolution, and Beyond

- Agentic AI realizes autonomic aspirations
- But replaces explicit control with implicit cognition
- Governance must evolve accordingly

2001 …

Thank you.

I've been ending autonomic computing talks with this cartoon (LHS) for over twenty years. The idea was always that systems would eventually manage themselves. Perhaps what we're seeing now with Agentic AI is the cognitive capability that the original vision was waiting for…

Happy 25th / Silver Anniversary Autonomic Computing!

… 2026

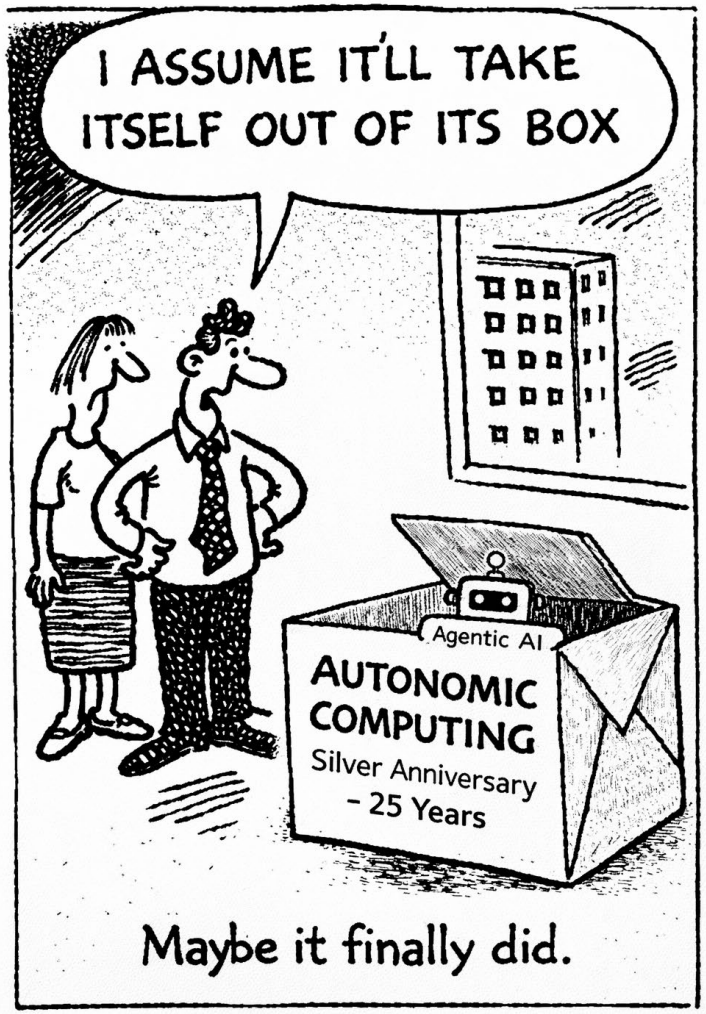

Ulster University

Roy Sterritt
@RoysterUlster
r.sterritt@ulster.ac.uk

Acknowledgements

Fitting with the Agentic AI spirit, the author acknowledges the use of large language model–based tools (specifically ChatGPT by OpenAI) to assist with early-stage drafting, structural organization, and language refinement of this paper & presentation, including generating images. Some retro images from A.C. retained intentionally to contrast. All technical content, interpretations, claims, and final revisions were performed and verified by the author, who takes full responsibility for the work (and mistakes).