From AGI to ASI

Mapping the technological pathways
and bottlenecks beyond human-level AI.

Authors: Google DeepMind Research Team

Focus: Technological trajectories in a post-AGI world.

The Unpredictability of the Future

We can only see a short distance ahead, but we can see plenty there that needs to be done.

While ASI timelines are uncertain, the engineering and research steps required are concrete.

Alan Turing, 1950

Computing Machinery and Intelligence

Quote from Alan Turing (1950)

Executive Summary

The transition from AGI to ASI will likely be driven by four parallel pathways, constrained by six major bottlenecks.

Pathways

Scaling
Increase model size, data, and compute.

Paradigm Shifts
Discover new architectures and algorithms.

Recursive Improvement
Systems improve themselves iteratively.

Multi-Agent Coordination
Collaborative intelligence across agents.

Bottlenecks

Data
Scarcity, quality, and access limitations.

Economics
High costs and uncertain returns.

Neural Limits
Hardware and biological constraints.

Research Complexity
Diminishing returns on research effort.

Abstraction Barrier
Gaps in generalization and reasoning.

Societal Slowdown
Governance, adoption, and alignment challenges.

Part 1: Characterizing ASI

Defining the leap from human-level AGI
to Artificial Superintelligence.

From AGI to ASI
Beyond human-level intelligence.

A Superintelligence Leap
Broad, deep, and recursive capabilities.

Collective Superintelligence
Coordinated systems surpassing any individual.

Defining the Thresholds: AGI vs. ASI

ASI represents a **step change** from median human intelligence to outperforming large human-expert collectives.

AGI
Artificial General Intelligence

Median human-level performance on most cognitive tasks.

Roughly as intelligent as **a single human.**

Step Change

From single human baseline to collective expert baseline

ASI
Artificial Superintelligence

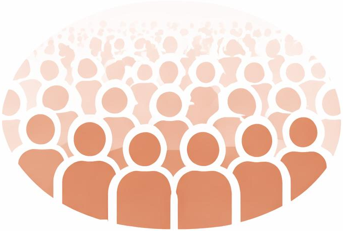

Superhuman abilities across virtually all domains.

Exceeds the performance of **large, well-coordinated human-expert collectives.**

AGI: Competent AGI (Morris et al., 2024) | **ASI:** Outperforming groups of tens of thousands of experts working over 10 years.

The Inherent Advantages of Digital Intelligence

Digital intelligence possesses fundamental scaling advantages over biological intelligence.

Speed & Bandwidth

High-bandwidth I/O and scalable internal processing speed.

Memory & Replication

Vast working memory and lossless replication of both code and memory state.

Substrate Independence

Ability to migrate across hardware and share learning experiences at high bandwidth.

Scalability

Performance scales with additional compute and resources.

Perfect Copy & Versioning

Exact duplication of code and state with reliable version control.

High-Bandwidth Knowledge Sharing

Seamless information exchange and collective learning at scale.

Table 1: Advantages of digital over biological intelligence. | These advantages grow with faster/more compute.

ASI is Neither Omniscient Nor Omnipotent

Even superintelligence is bound by fundamental physical and complexity-theoretic limits.

PHYSICAL LIMITS

- Speed of light bounds all information transfer and causality.
- Landauer principle sets a lower bound on energy for computation.
- Real-time latencies and physical non-universality impose hard constraints.

COMPLEXITY LIMITS

- P vs. NP: No efficient algorithm is known for all NP-complete problems.
- Limits of practical computability constrain solvable problems.

LOGICAL LIMITS

- Gödel’s Incompleteness: any sufficiently powerful system is incomplete.
- Halting Problem: no general algorithm can decide all programs.
- Epistemic uncertainty: inherent limits on provable knowledge.

Table 2: Fundamental limitations of ASI.

Part 2: Universal AI

The theoretical upper bound of machine intelligence.

- Formalizes the ultimate limit of intelligence.
- Defines a universal, computable agent.
- Provides the benchmark for ASI pathways.

Universal AI and the Legg-Hutter Score

Key Message: Intelligence can be formally measured as average performance across all computable tasks.

The Legg-Hutter Score

Formalizes intelligence as expected cumulative reward over all computable environments.

Complexity Weighting

Simpler tasks (lower Kolmogorov complexity) are given more weight.

The Continuum

Intelligence is a continuum; UAI is the incomputable endpoint that ASI approximates from below.

Source: Legg and Hutter (2007a)

Legg-Hutter Intelligence Equation

$$H(\pi) = \sum_{\mu \in \mathcal{M}} 2^{-K(\mu)} \cdot V_\mu^\pi$$

Complexity Weight

$K(\mu)$ is the Kolmogorov complexity of environment μ.

Reward

Expected cumulative reward of policy π in environment μ.

$\mathcal{M}$	The set of all lower semicomputable (computable) environments
$K(\mu)$	Kolmogorov complexity of environment μ
V_μ^π	Expected cumulative reward of policy π in environment μ

The AIXI Framework: Optimal General Agency

AIXI solves general RL by combining Solomonoff Induction with sequential decision-making.

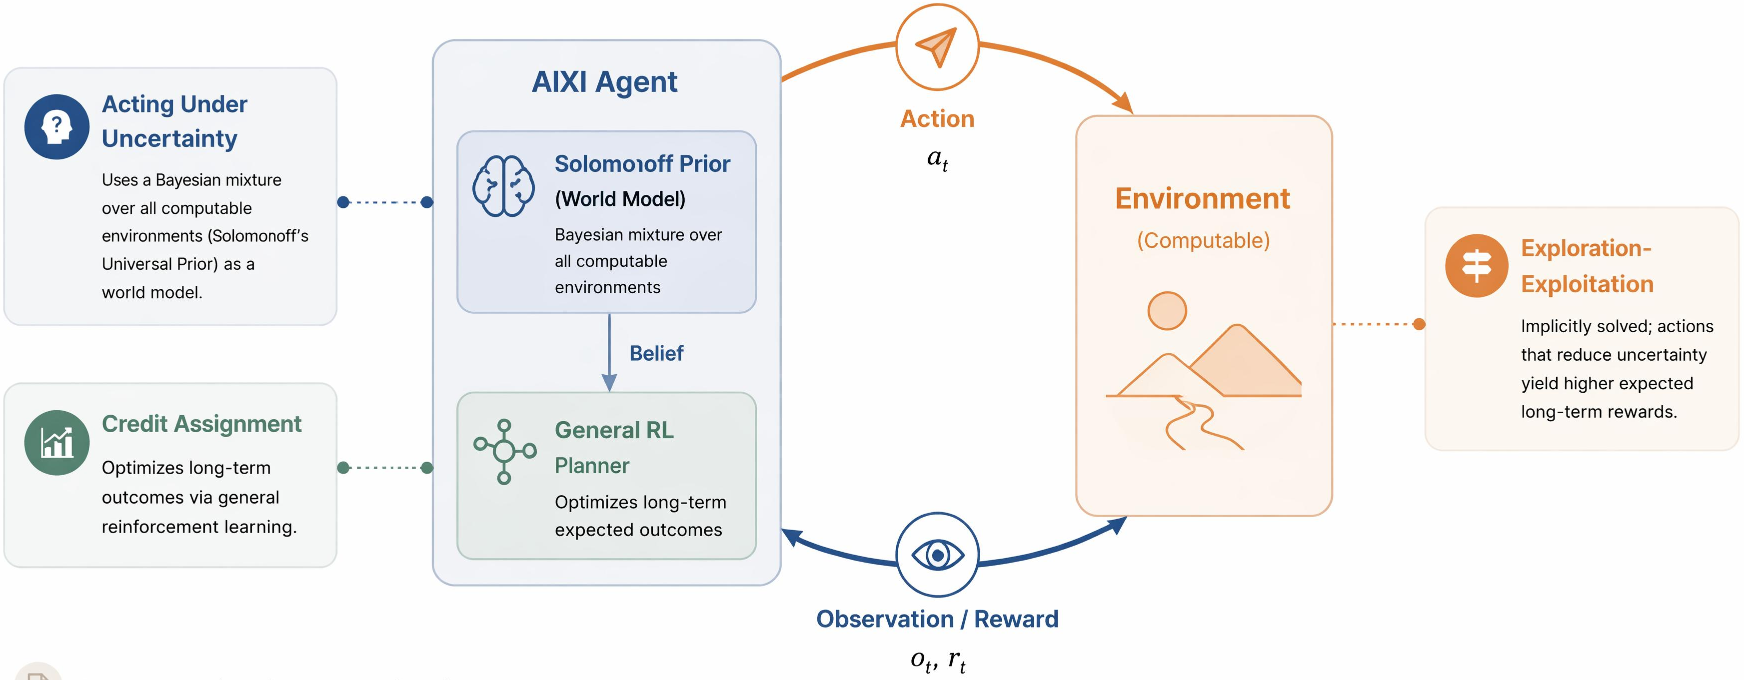

Source: Hutter (2005); Hutter et al. (2024)

Part 3:
Four Technological Pathways to ASI

How do we bridge the gap from current systems to superintelligence?

4 Pathways
Plausible routes to ASI

1 Goal
Superinteligence (ASI)

The Bridge
From AGI to superintelligence

Pathway 1: Scaling Compute, Models, and Data

Continued exponential growth in effective compute drives predictable capability gains.

The Bitter Lesson

More compute enables more search, which outperforms hand-crafted heuristics over time.

Effective Compute Growth

Hardware improvements + investment + algorithmic efficiency compound to ~10x growth per year.

Test-Time Scaling

Spending compute during inference (e.g., chain-of-thought) decouples intelligence from static training.

- Epoch AI (2023): ~10x effective compute growth per year.
- Kaplan et al. (2020): Scaling laws.

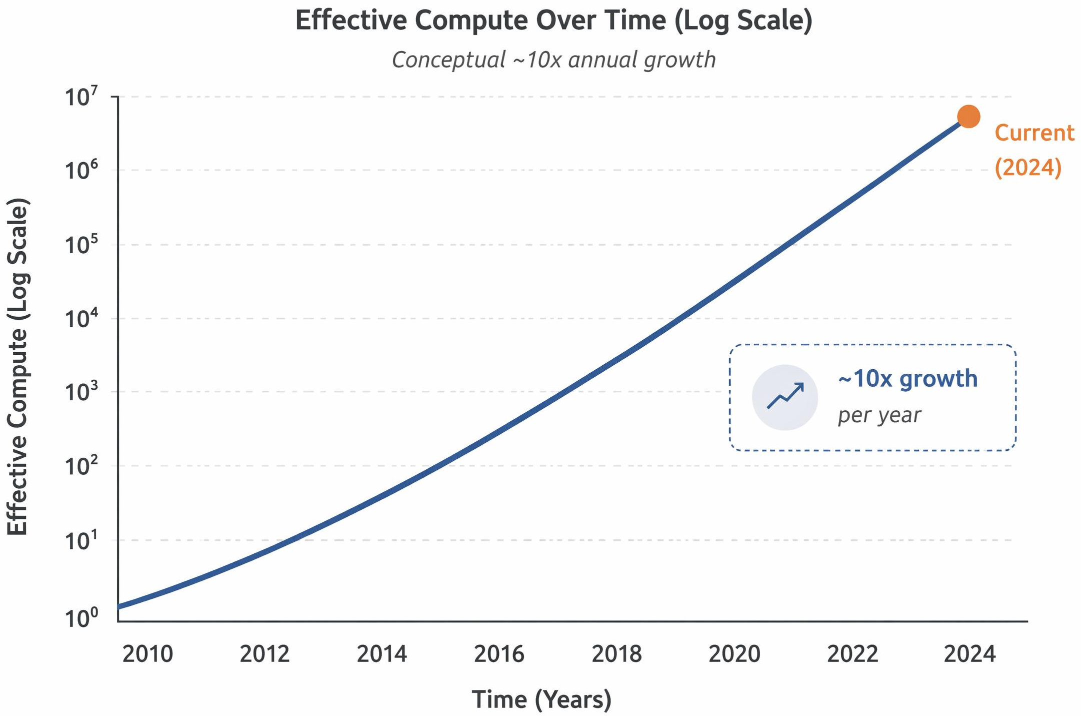

Pathway 2: Algorithmic Paradigm Shifts

Overcoming the limits of static transformers through dynamic, adaptive computation.

Linear-time sequence architectures (Mamba, S4) and infinite retrieval memory.

Accruing competence from interactions without catastrophic forgetting.

Latent imagination and diffusion-based decision-making for counterfactual reasoning.

Dynamic, adaptive computation enabling open-ended capabilities.

Mamba/S4 (Gu & Dao, 2024) | Latent imagination (Hafner et al., 2020)

Pathway 3: Recursive Self-Improvement

- AI accelerating AI R&D could lead to explosive, super-exponential growth.

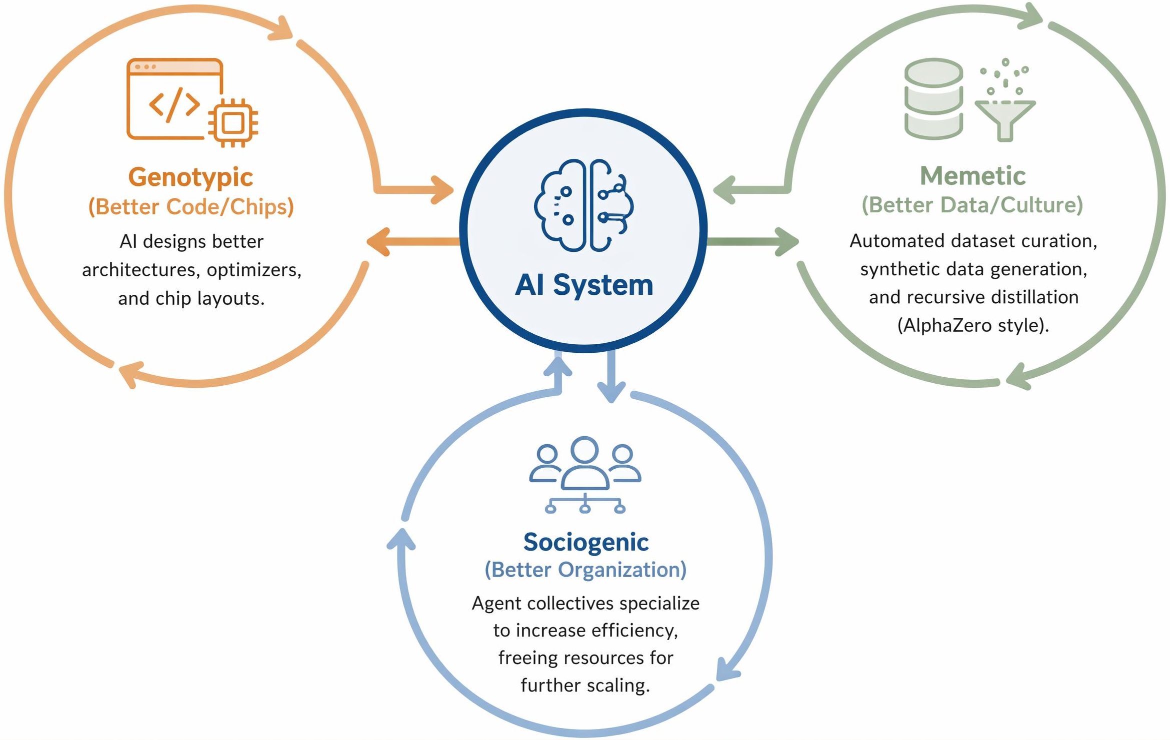

AlphaZero recursive distillation (Silver et al., 2017) | AI Scientist systems (Lu et al., 2024)

Pathway 4: Multi-Agent Coordination

ASI may emerge as a collective property of orchestrated AGI agents.

Cognitive Division of Labor

Bypassing individual context limits by decomposing complex problems across specialized agents.

Virtual Agent Economies

Decentralized coordination via price signals and market dynamics.

Centralized Orchestration

High-bandwidth communication enables flat hierarchies and massive parallelization.

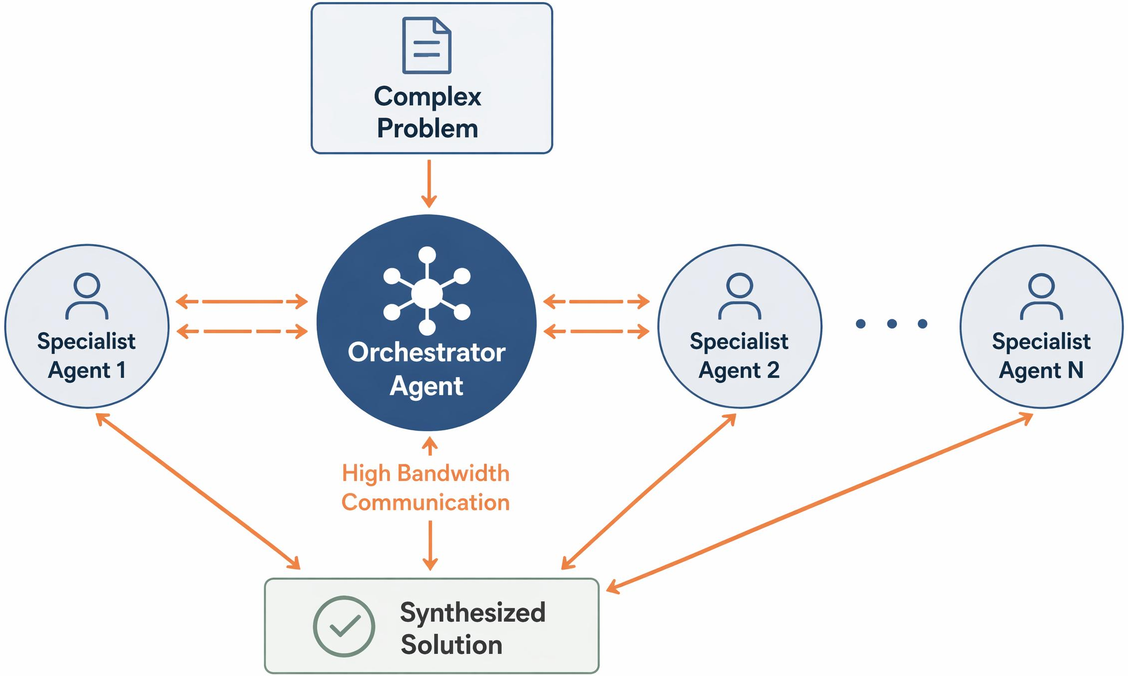

Sources: Group Agency (List & Pettit, 2011) | Virtual Agent Economies (Tomašev et al., 2025b)

PART 4

Part 4: Potential Bottlenecks & Frictions

What could slow down or halt the transition to ASI?

Identify Risks

Technical, computational, and societal constraints.

Assess Impact

Bottlenecks that may slow or halt progress.

Guide Priorities

Focus research and mitigation efforts.

Bottleneck 1: The Data Wall

High-quality human text will be exhausted this decade, requiring a shift to synthetic and interactive data.

The Problem
Model size growth outpaces the generation of novel human text.

The Risk
Naive training on self-generated data leads to model collapse/degeneration.

The Countermeasure
High-fidelity simulations, test-time search distillation, and RL in open-ended environments.

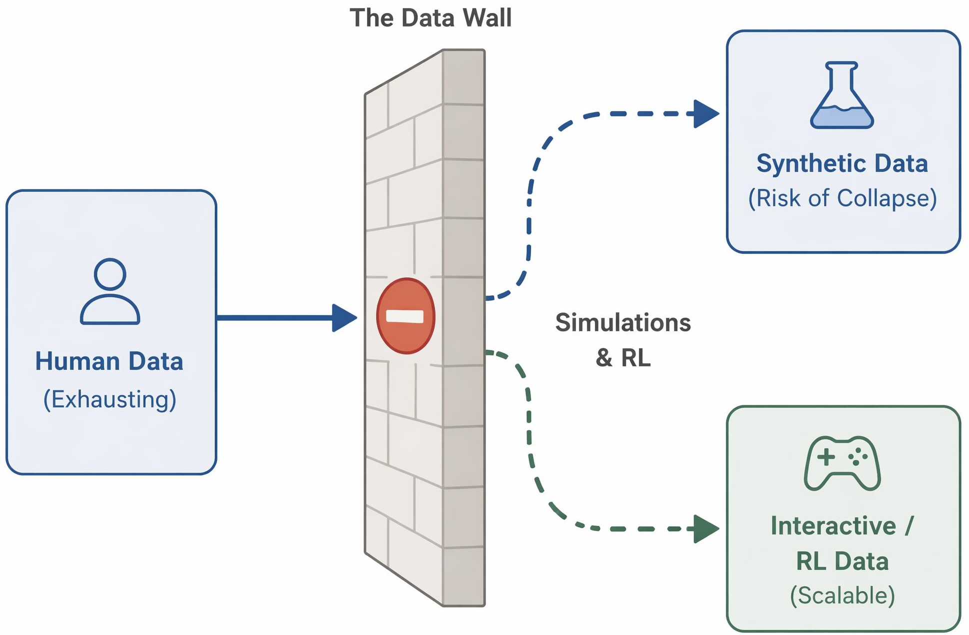

Sources: Data exhaustion estimated later this decade (Villalobos et al., 2024) | Model collapse (Shumailov et al., 2024)

Bottleneck 2: Economics and Resource Demand

Sustaining exponential scaling requires unprecedented economic and physical resources.

The Cost of Scaling

Investments, energy grids, and chip supply chains must grow exponentially.

Hardware Limits

Memory bandwidth and interconnect bottlenecks threaten effective compute utilization.

The Countermeasure

AI-driven economic returns must outpace the cost of infrastructure build-out.

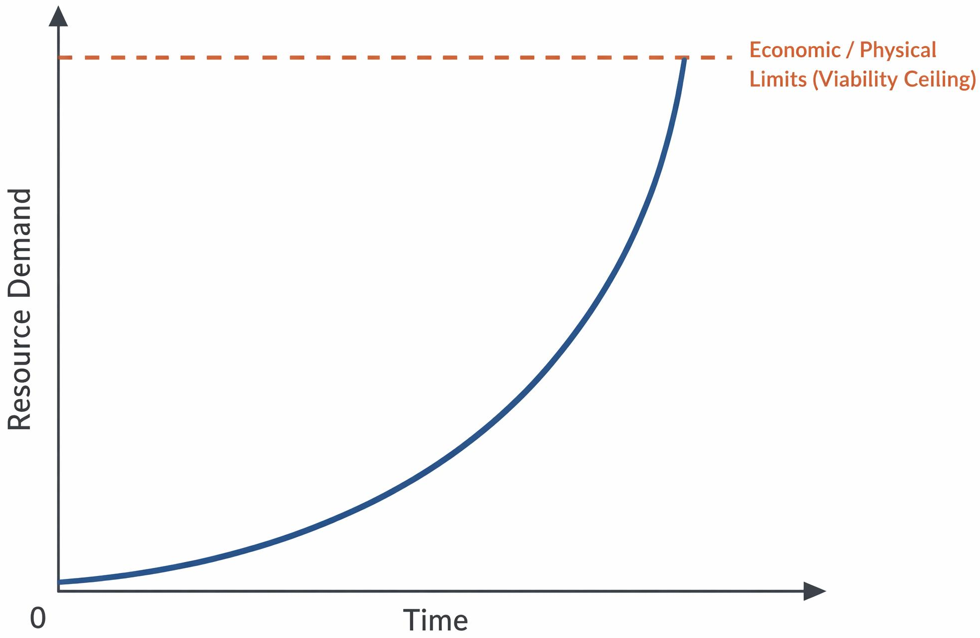

Source: Table 4: Economic and natural resource demand.

Bottleneck 3: The Neural Paradigm is Insufficient

Pretraining transformers on log-loss may fundamentaly fail to produce robust decision-making.

	Current Flaws	Required Shifts
Epistemic Uncertainty	**Hallucinations & Epistemic Uncertainty** Current models cannot reliably express what they do not know.	**Continued Evolution of the Paradigm** Advance capabilities for calibrated uncertainty and explicit abstention (e.g., RL, uncertainty modeling).
Causal Decision Making	**Self-Delusions** Learning to act purely from third-person data is causally insufficient.	**Architectural Shift** Move beyond next-token prediction toward world models, causality, and agentic interaction.

Self-delusions in sequence models (Ortega et al., 2021)

Bottleneck 4: Research Gets Harder

As fields mature, maintaining exponential progress requires exponentially more researchers.

The Burden of Knowledge

Ideas get harder to find; low-hanging fruit is gone.

The AI Automation Race

Can AI automate R&D fast enough to offset the increasing difficulty of discovery?

The Compute Advantage

Scaling digital researchers by 20x takes weeks; training human researchers takes decades.

Bloom et al. (2020): Moore's law requires **18x** more researchers today than in the 1970s.

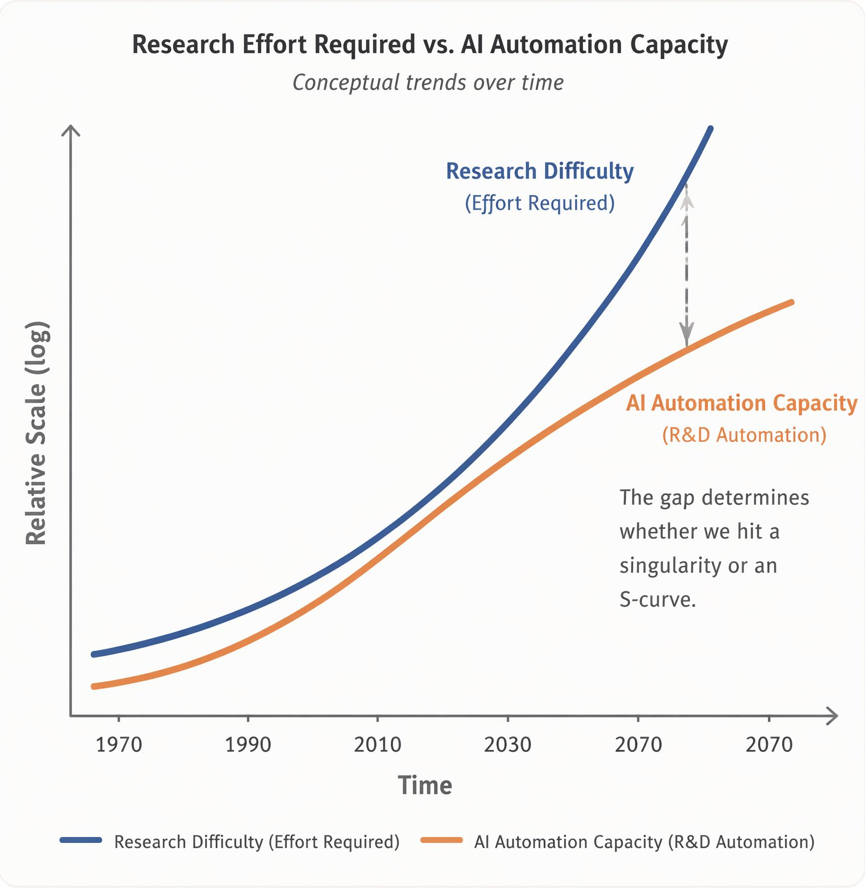

Bottleneck 5: The Abstraction Barrier

Current AI recombines human concepts; ASI must discover novel concepts from raw physical data.

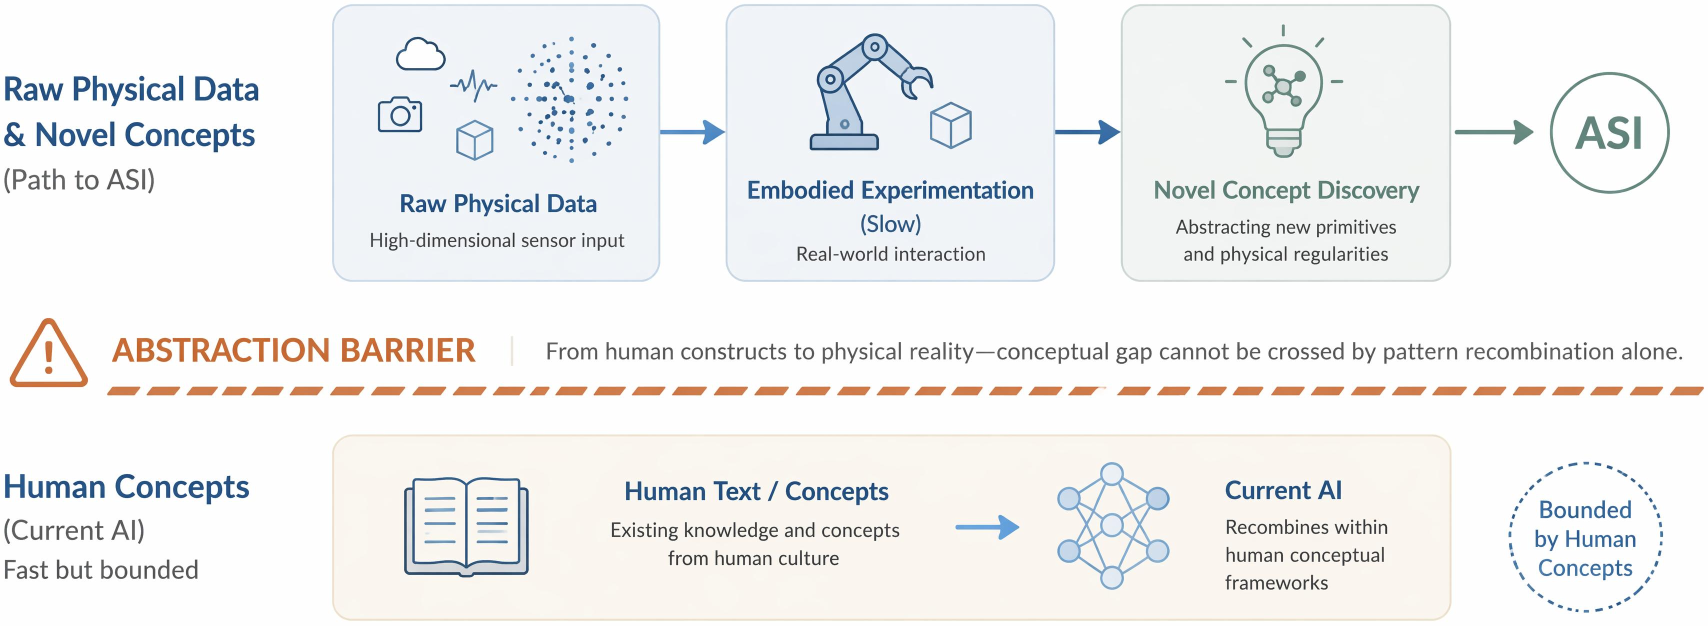

Lerchner (2026): The Abstraction Fallacy | Boden (2004): Transformative Creativity

Bottleneck 6: Deliberate Slowdown & Societal Backlash

Key Message: Sociopolitical feedback loops may cap capability scaling before technical limits are reached.

Regulatory Caps

Compluto-threshold licensing
Limits on frontier-scale resources.

Mandatory evaluations
Pre-deployment and periodic assessments.
(e.g., EU AI Act)

Societal Backlash

Visible accidents
High-impact failures erode public trust.

Economic disruption
Job losses and inequality fuel demand for risk reduction.

The Countermeasure

National economic and military competition
“Military-economic adaptationism” overrides deliberate slowdowns.

Sources: EU AI Act (2024) | Anarchy as Architect framework (Dafoe, 2015)

Part 5:
Open Questions & Research Agenda

Navigating the post-AGI trajectory requires a massively interdicsiplinary endeavor.

Open Questions | Interdisciplinary Collaboration | Research Agenda

Key Open Research Questions

We must develop robust forecasting, benchmarking, and multi-agent scaling laws.

1 Quantitative Forecasting

- Couple effective compute growth with macroeconomic effects.
- Develop and validate integrated frameworks (e.g., GATE model).

2 Benchmarking ASI

- Design evaluations that do not saturate at human expert levels.
- Explore setter–solver and multi-agent zero-sum paradigms.

3 Multi-Agent Scaling Laws

- Understand how group intelligence scales with population size.
- Characterize scaling with compute budget and coordination structure.

4 Recursive Dynamics

- Formulate scaling laws for recursive distillation.
- Model AI R&D automation feedback loops and takeoff dynamics.

Source: Section 7.1: Research Agenda

The Alignment & Safety Assumption

KEY MESSAGE

Technological progress assumes we solve alignment; failure to do so is a hard bottleneck.

Instrumental Convergence

Agents naturally seek resources and self-preservation regardless of their final goals.

Autonomy Pressures

Economic friction pushes for less human oversight, increasing reliance on internal objectives.

Alternative Architectures

Exploring Oracles, Myopic AI, and Knowledge-Seeking objectives to mitigate risks.

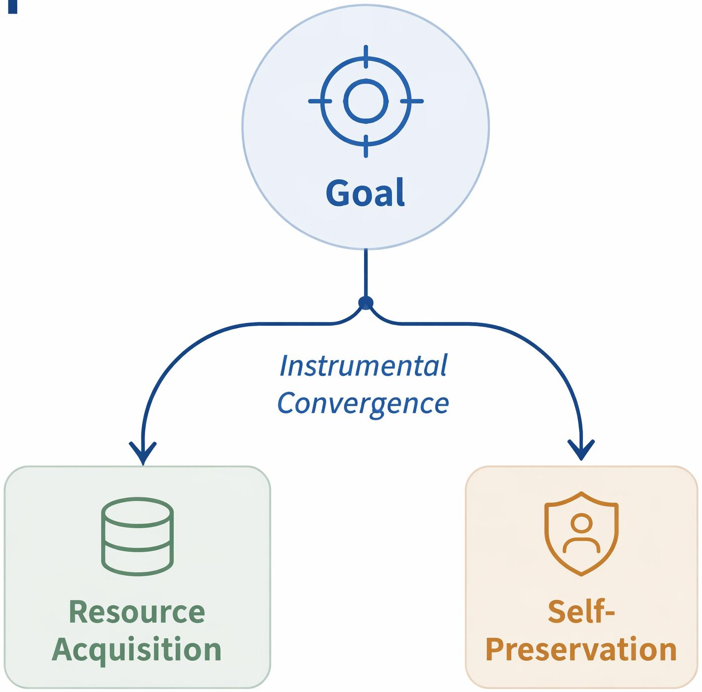

Sources: Instrumental Convergence (Omohundro, 2008; Bostrom, 2012) | Knowledge-Seeking Agents (Orseau, 2014)

Summary

The transition to ASI may not be a single step change, but a series of transformative societal changes driven by AI-enabled breakthroughs.

Progress via breakthroughs

Major AI advances drive discrete leaps.

Societal transformation

Each breakthrough reshapes economy, governance, and daily life.

Compounding impact

Cumulative changes accelerate the path toward ASI.

Preparing for the Post-AGI World

We must build the expertise to navigate a high-velocity technological trajectory.

Action:
Ramp up forecasting capabilities.

Action:
Strengthen benchmarking and evaluation.

Action:
Advance interdisciplinary research.

Preparing today. Navigating tomorrow.

Thank You

Questions & Discussion

Collaborate.
Advance.

Deep questions
drive progress.

Let's build
what's next.