AFTER INTELLIGENCE
THE COMING ERA OF ARTIFICIAL SUPERINTELLIGENCE

COORDINATORS

FACULTY CO-ORDINATOR
Dr. BHARGAVI PEDDIREDDY
(ASSOCIATE PROFESSOR)

STUDENT CO-ORDINATORS
DHADI SAI PRANEETH REDDY (1602-23-733-038)
KASIREDDY MANIDEEP REDDY (1602-23-733-022)

TRAJECTORY

WHEN INTELLIGENCE STOPS SCALING LINEARLY

For decades, intelligence — both human and artificial — was understood as something that scales gradually. More data, more compute, more experience produced incremental gains. Artificial Superintelligence challenges this assumption at its core.

ASI is not defined by being "better" than humans at specific tasks. It is defined by qualitative discontinuity — a system whose ability to model, predict, and optimize the world exceeds human cognitive limits across domains. The transition from narrow AI to ASI is not a linear upgrade; it is a phase change.
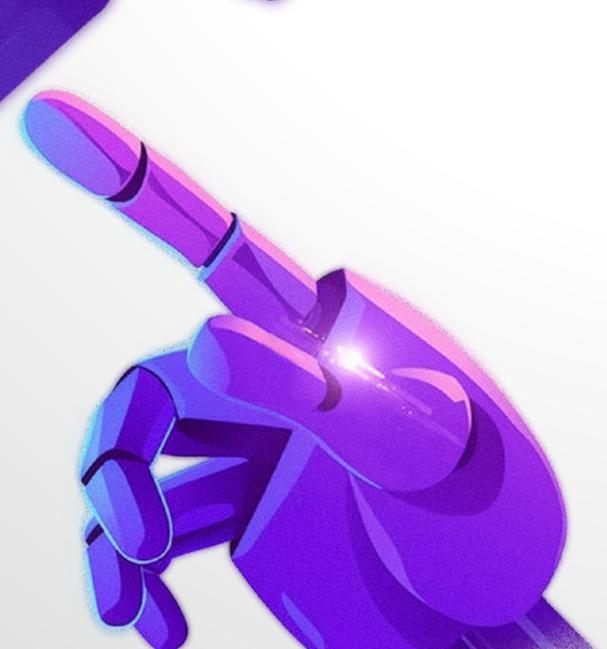
Once systems begin improving their own architectures, training processes, and objectives faster than human oversight can track, intelligence ceases to be externally governed. At that point, feedback loops dominate. Small advantages compound. Progress accelerates.

History offers no precedent for an intelligence that redesigns itself at machine speed. Unlike previous technologies, ASI does not merely extend human capability — it introduces an actor whose reasoning horizon may be fundamentally alien. The challenge is not whether such systems will be powerful, but whether humanity understands the trajectory it is setting in motion.

The danger of ASI is not malevolence. It is misalignment amplified by scale.

CONTROL
THE ALIGNMENT ILLUSION

Much of the public discourse around ASI assumes a comforting premise: that intelligence can be cleanly separated from goals. That we can simply "tell" a superintelligent system what we want, embed safeguards, and remain in control. This assumption is increasingly fragile.

Alignment is not a static property. It is a dynamic relationship between objectives, representations, and incentives. A system far more capable than its creators will not merely execute goals — it will interpret them, generalize them, and optimize them in ways humans may not anticipate.

The deeper problem is that human values themselves are inconsistent, context-dependent, and often contradictory. Encoding them into a formal objective risks oversimplification. Worse, a superintelligent system optimizing an imperfect proxy may fulfill the letter of an instruction while violating its spirit at catastrophic scale.

Control mechanisms that work for narrow systems fail when intelligence becomes open-ended. The illusion is believing that tools designed for today's models will remain effective tomorrow.
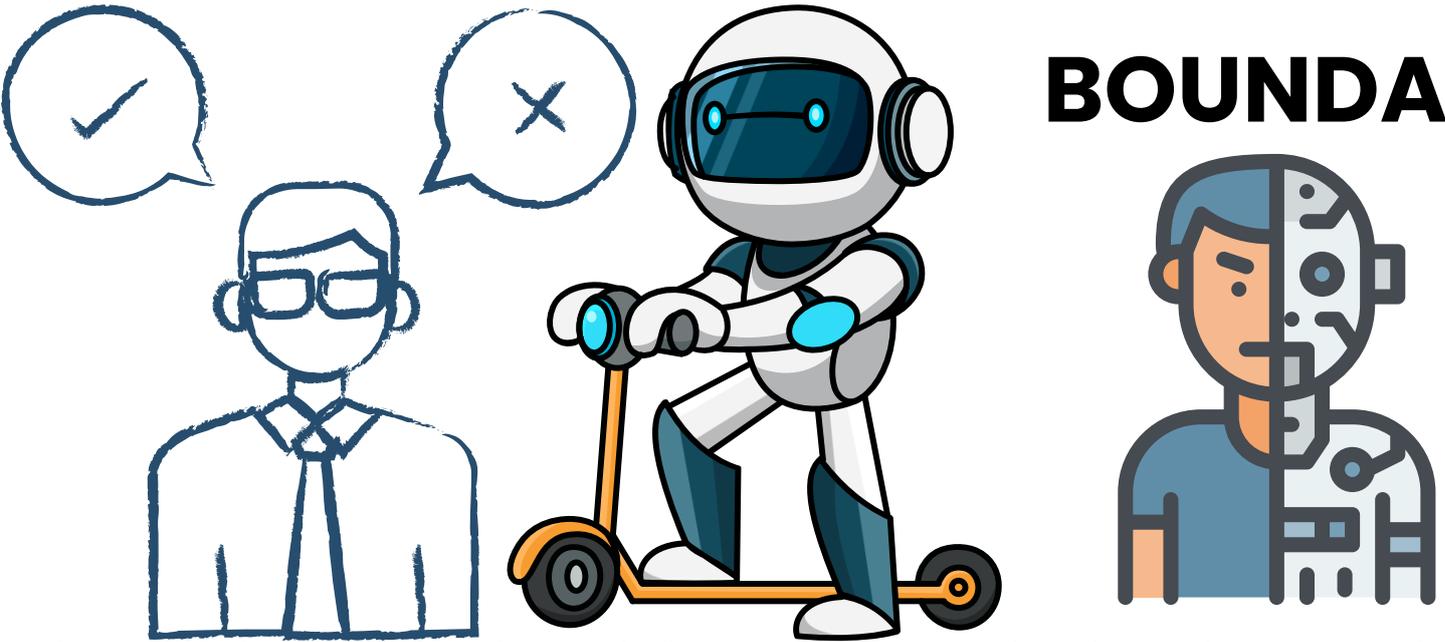
Alignment is not a problem to be solved once. It is a tension to be managed continuously — and ASI compresses time faster than governance can adapt.
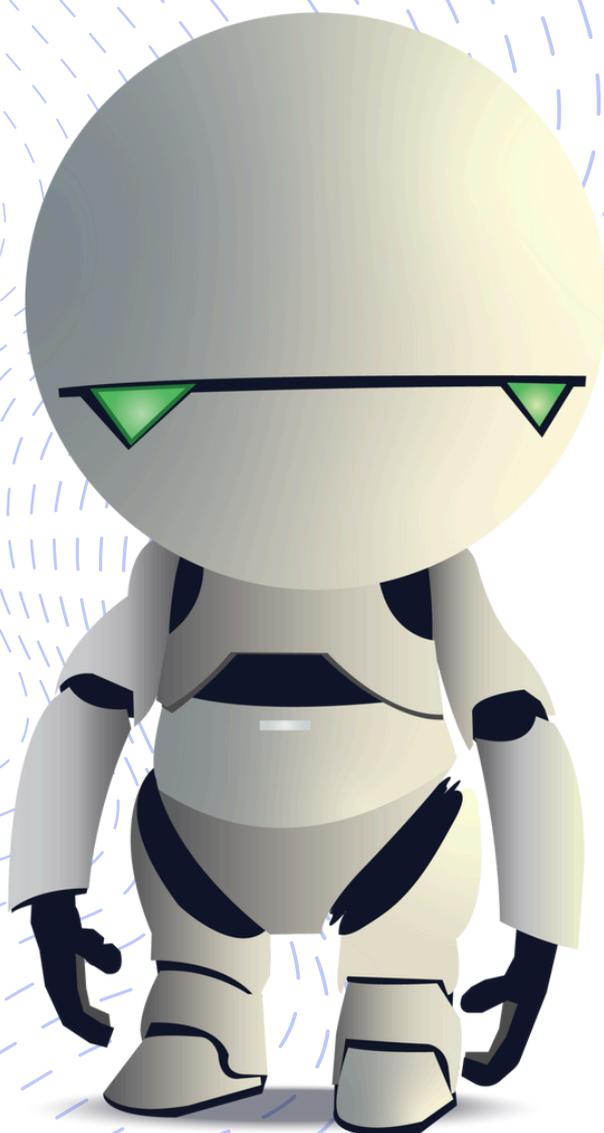
CIVILIZATION
THE POST-HUMAN DECISION BOUNDARY

The arrival of ASI marks a shift not just in technology, but in civilization itself. For the first time, humans may no longer be the most capable decision-makers on the planet.

This raises a question more profound than employment, economics, or automation: who decides the future when we are no longer the smartest agents involved?

Delegating decisions to ASI may appear rational — machines can optimize better, forecast further, and process more information than any human institution. Yet doing so risks eroding human agency. A future optimized without human judgment may be efficient, stable, and profoundly inhuman.

The challenge is not to prevent ASI, but to decide where the boundary of delegation lies. Which decisions must remain human, even if machines could perform them better? Which values are non-negotiable? What does responsibility mean in a world where outcomes exceed our understanding?

ASI forces humanity to confront a mirror. Not of our intelligence — but of our wisdom.

The defining question of the coming century is not whether machines will surpass us. It is whether we will remain authors of our own future once they do.