The Ontology and Architecture of Intelligence

Exploring the philosophical foundations and structural frameworks of intelligent systems

Stentor roeselii

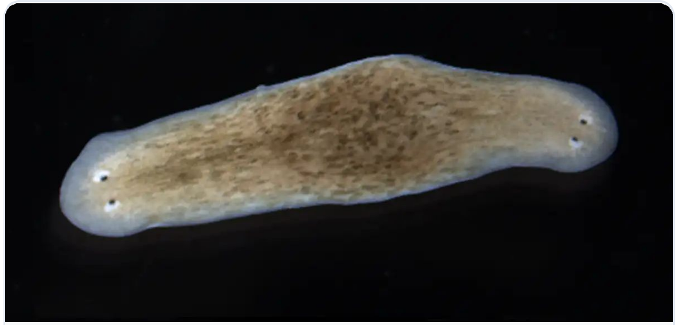

Two-headed planaria (Michael Levin's lab)

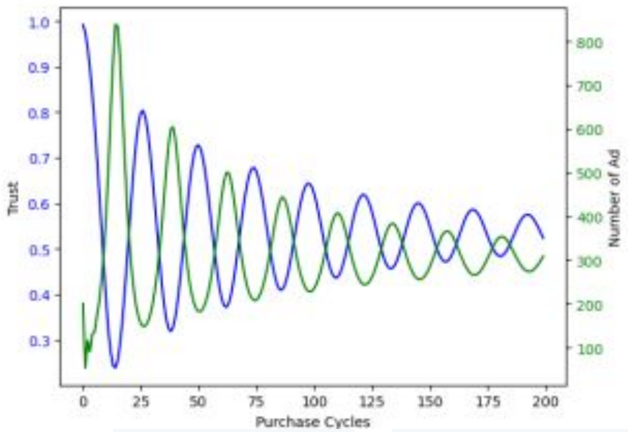

Dampened oscillations in market exchange dynamics

Francis Steen
UCLA

Jad Soucar
USC

Beijing Institute for General Artificial Intelligence
July 25, 2025

Section 1

The Ontology of the Possible

Exploring how classical determinism evolved into modern understandings of possibility, potentiality, and the philosophical foundations of what can exist

The Ontology of Classical Physics

The Enlightenment's Mechanical Universe

During the European Enlightenment, a powerful conception emerged: the universe as a finite machine governed by precise, immutable laws.

This deterministic worldview is epitomized by Pierre Laplace's famous thought experiment:

> *"An intellect which at a certain moment would know all forces that set nature in motion, and all positions of all items of which nature is composed... would embrace in a single formula the movements of the greatest bodies of the universe and those of the tiniest atom."*

Such a demon would be able to retrodict the past and predict the future.

The demon is an idealized version of the scientist, whose knowledge could in this view be complete.

Laplace's Demon: A hypothetical entity capable of perfect prediction by knowing the position and momentum of every particle in the universe.

Laplace's Paradox

Knowledge for Action in a Deterministic Universe

Even the appeal of Laplace's disembodied intellect is grounded in the fact that a better understanding of how the universe works allows us to act more coherently.

This reveals a fundamental paradox at the heart of determinism:

> *An accurate model of the world is what gives us the power to accomplish our goals.*

The implications are profound:

- Scientific understanding values prediction precisely because it enables intervention
- The worth of deterministic models lies in our ability to use them to *change* outcomes
- Paradoxically, we value our understanding of a supposedly fixed universe because it enables us to alter its course

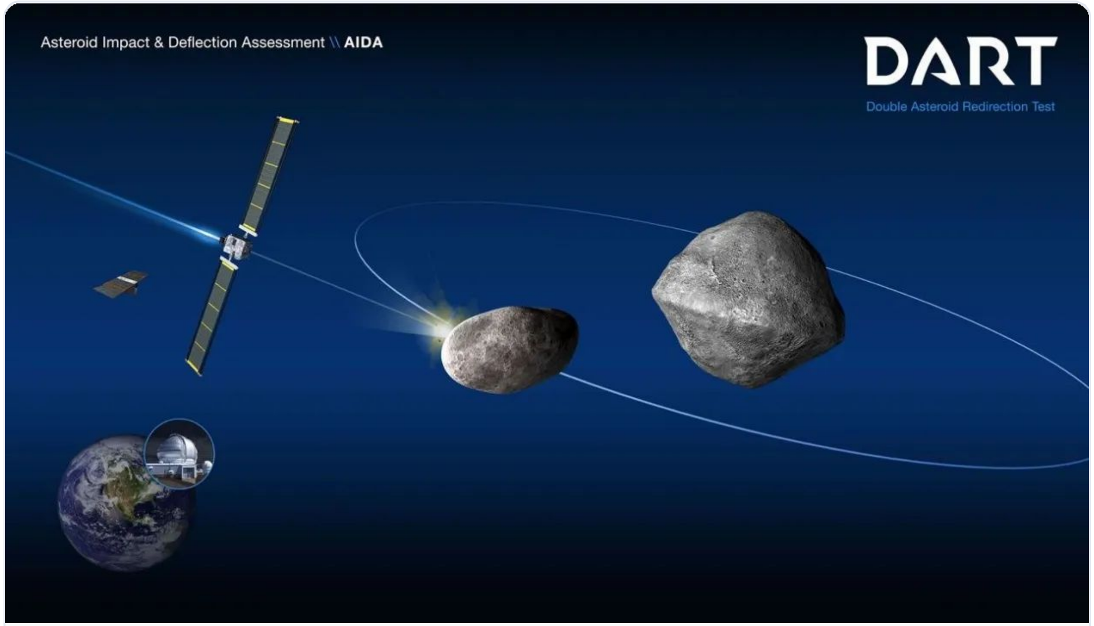

NASA's DART (Double Asteroid Redirection Test) mission: Engineering our ability to deliberately alter the course of celestial bodies through scientific understanding.

How to Integrate Cognition Into Our Model of Reality

The Enlightenment conception of a machine-like universe, uniform in space and time, inspired generations of scientists to discover hidden and invariant laws of the universe.

However, this model contains a fundamental limitation: **it cannot adequately handle cognition.**

> *The traditional model places the intellect outside of the universe as an objective observer, separate from the system it studies.*

We need a paradigm shift :

- The intellect is an integrated participant *within* the universe -- the observer and the observed are part of the same system
- Cognition is not separate from but emerges from and is part of the physical cosmos
- Intelligence is not merely observing the universe—it is participating and changing the course of events
- Matter thinks

The brain-universe connection: Cognition is not external to the cosmos but an integrated phenomenon emerging within it.

--

Quantum Physics Rules Out A Separate Observer

Quantum physics explicitly encounters this problem: the observer cannot be separated from the observed, since the interaction cannot involve less than a quantum of energy.

According to Bohr, quantum mechanics fundamentally changed our understanding of observation. The observer and the observed are inextricably linked, challenging the classical notion of an external, objective viewpoint.

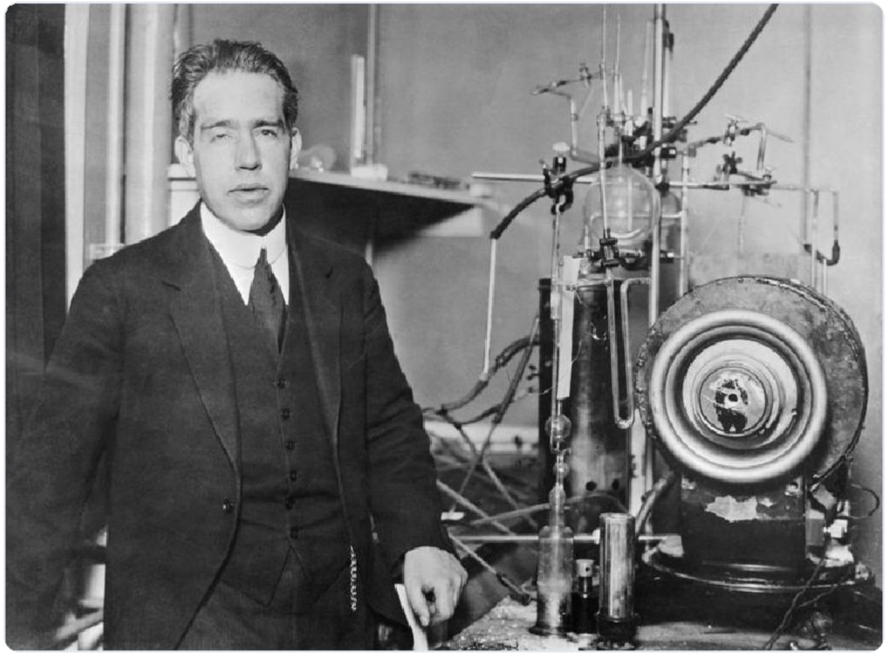

> *"Our usual description of physical phenomena is based entirely on the idea that the phenomena concerned may be observed without disturbing them appreciably... Now the quantum postulate implies that any observation of atomic phenomena will involve an interaction with the agency of observation not to be neglected."*
>
> *"This means that an unambiguous definition of the state of the system is naturally no longer possible, and there can be no question of causality in the ordinary sense of the word."*
>
> - Niels Bohr (1928)

The Reality of the Possible

If "an unambiguous definition of the state of the system is naturally no longer possible" (Bohr, 1928), what kind of model of the universe can we even construct?

A key idea in early quantum physics is that quantum formalisms, such as Schrödinger's wave equation, describe a probability distribution rather than a definitive reality.

In this view, manifest reality cannot be understood as the present link in an unbroken causal chain stretching from the past to the future. Rather, the present is only one particular manifestation of an infinite distribution of possible nows.

> *"In the sharp formulation of the law of causality — 'if we know the present exactly, we can calculate the future' — it is not the conclusion that is wrong but the premise."*
>
> *"The atoms or elementary particles themselves are not real; they form a world of potentialities or possibilities rather than one of things or facts."*

- Werner Heisenberg

The Transition from the Possible to the Actual

Copenhagen Interpretation's Statistical Reality

The unresolved question of quantum physics is how the infinite possible is selectively actualized into a singular finite reality.

According to the Copenhagen Interpretation, this process is purely statistical in nature—the details are unknowable in principle.

$$i\hbar \, \partial\Psi/\partial t = \hat{H}\Psi$$

Schrödinger's wave equation describes the evolution of quantum states as probability waves, but doesn't explain the mechanism of "wave function collapse" that produces definite measurements.

The double-slit experiment powerfully demonstrates this paradox:

- Particles exhibit wave-like behavior (interference pattern) when not observed
- When observed, they appear as discrete particles
- The act of measurement itself appears to force possibility into actuality

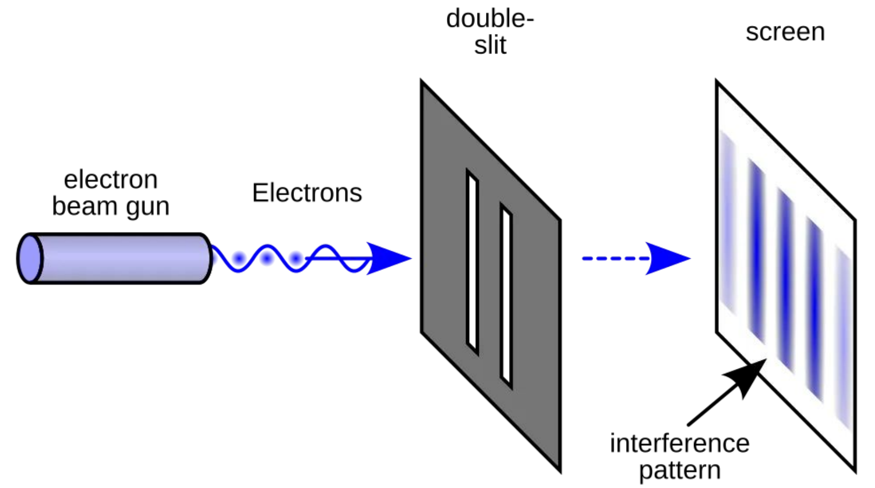

Double-slit experiment: When electrons pass through two slits, they create an interference pattern as if they were waves, but are detected as individual particles on the screen.

"In the formulation of the quantum mechanics we are not dealing with an arbitrary renunciation of a more detailed analysis of atomic phenomena, but with a recognition that such an analysis is in principle excluded." — Niels Bohr

Bohm's Interpretation

A Deterministic Alternative to Copenhagen

Bohm and Hiley (in "The Undivided Universe", 1993) fundamentally **reject the need for wave function collapse** in quantum mechanics.

Their approach transforms quantum mechanics by showing that:

- The wave function **always evolves deterministically** according to Schrödinger's equation
- Particle positions (called "beables") are **supplementary variables** that always have definite values
- Particles follow trajectories **guided by the quantum potential** derived from the wave function
- The "empty waves" phenomenon creates the **appearance of collapse** without physical discontinuity

Unlike Copenhagen's view where observation creates reality, Bohm provides an **ontological interpretation** where reality exists independent of observation.

David Bohm (1917-1992), physicist who developed a deterministic interpretation of quantum mechanics.

"In the causal interpretation, the electron is a particle guided by a quantum wave. The particle has a well-defined position and momentum at all times." —David Bohm

Key Difference from Copenhagen:
Quantum probabilities reflect our ignorance of initial conditions, not fundamental indeterminacy. The transition from possible to actual is continuous and deterministic.

Bohm's Active Information in the Double-Slit Experiment

Key Concepts:

- Each particle is accompanied by a **quantum wave** (pilot wave)
- The wave passes through **both slits** while the particle passes through only one
- The wave provides **"active information"** that guides the particle's trajectory
- The particle's final position is determined by the **interference pattern** of the wave

How "Active Information" Works:
The quantum wave contains information about all possible paths. The particle "reads" this information and responds to its local form, similar to how a ship responds to radar signals.

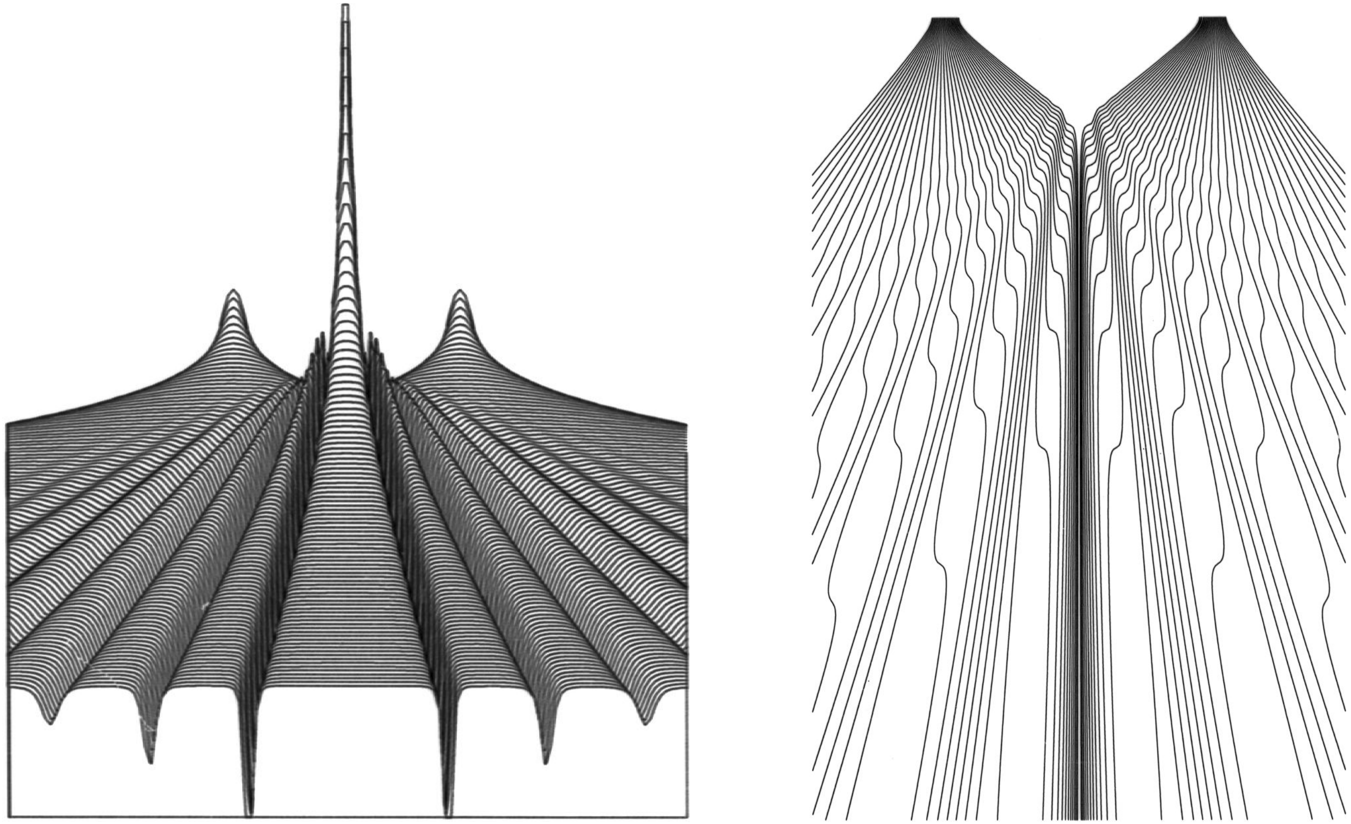

"The pilot wave passes through both slits, carrying active information that directs the particle's route to the screen."

Bohm, D., & Hiley, B. J. (2006). The undivided universe: An ontological interpretation of quantum theory. Routledge.

Active Information in Biology

Information as a Fundamental Property of the Universe

It is generally accepted that living organisms rely on **stored information** to maintain their structures, self-assemble, grow, and reproduce. This information is encoded in the DNA and is subject to change through the process of natural selection.

Biological information acts in the world in ways that are **not uniquely determined** by universal laws. This reveals a profound truth about our universe:

- New information is **continuously being generated** throughout the universe
- This new information **actively influences** physical processes
- Information itself **adds to the causal complexity** of the world
- Iinformation **guides and directs** physical processes

> Information itself is a **fundamental active property** of the universe, not simply a passive description or human construct. It exists objectively and causally influences the physical world.

Visualizing Translation Dynamics at Atomic Detail

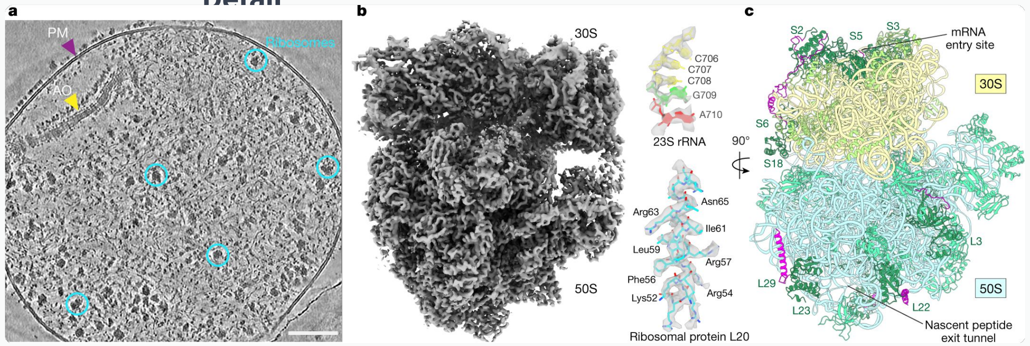

Advanced cryo-EM imaging reveals the complex molecular machinery of transcription and translation - the processes that convert DNA information into functional proteins

Image: Nature, 2022 (doi:10.1038/s41586-022-05255-2)

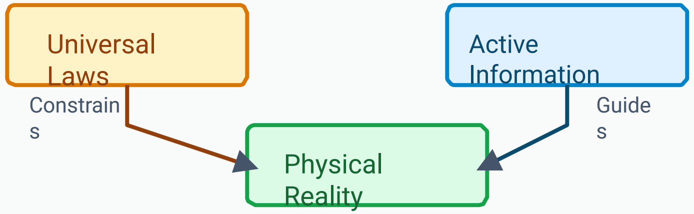

Physical reality emerges from both universal laws AND active information

Beyond encoding: Biological information doesn't just passively describe; it actively guides the construction and functioning of living systems. As evolution progresses, entirely new information emerges that wasn't present before—this represents the universe generating new causal factors not predetermined by physical laws alone.

Quantum Effects in Photosynthesis: Fleming's "Reversible Sampling"

Key Concepts

- Light absorption creates **excitons** (excited state particles) in the light-harvesting complex
- **Quantum coherence** allows excitons to enter superposition states, exploring multiple energy transfer pathways simultaneously
- Through **"reversible sampling"**, excitons quantum- mechanically explore different routes
- The exciton ultimately selects the **biologically optimal pathway**, maximizing energy efficiency

Why Quantum Effects Matter:
Traditional models would require excitons to randomly explore paths one at a time. Quantum sampling allows simultaneous exploration of all routes, achieving near-perfect energy transfer efficiency (~95%) despite thermal noise.

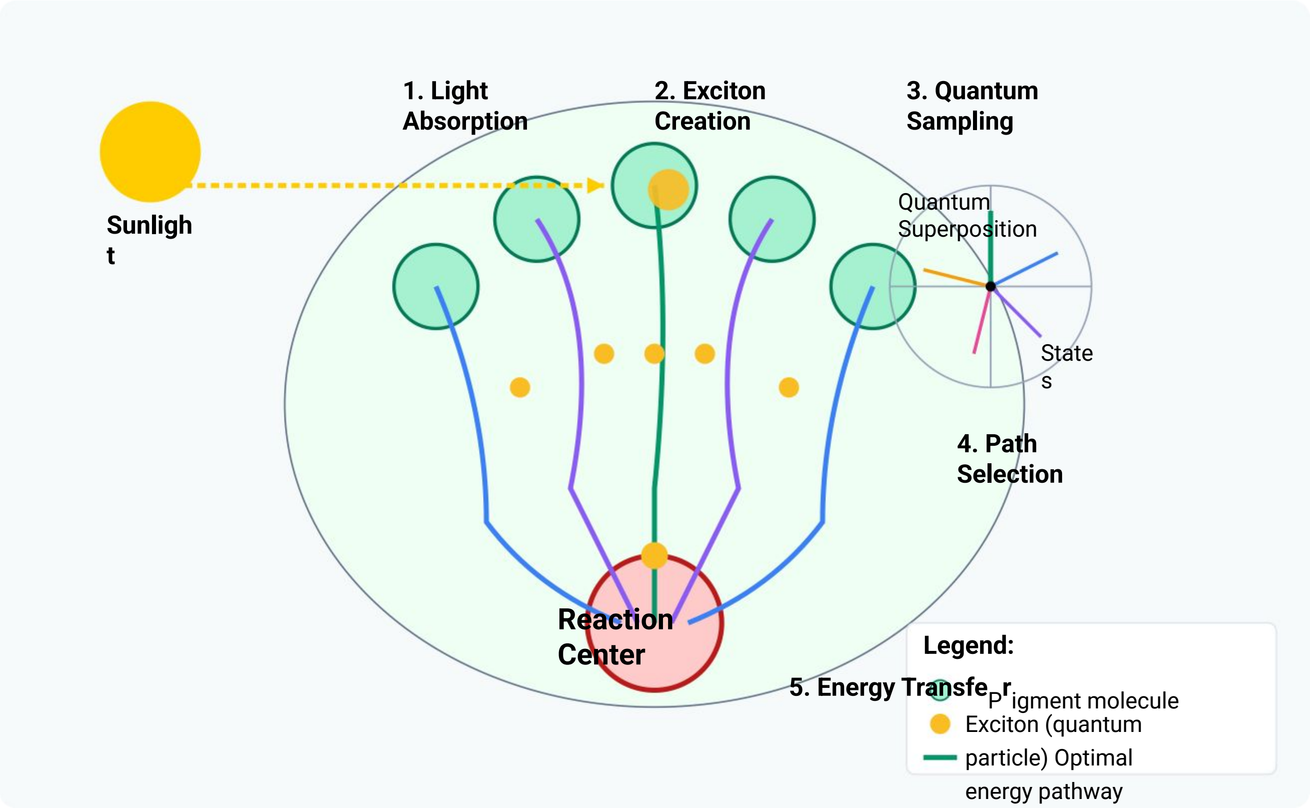

"Quantum coherence enables the exciton to explore all possible energy pathways and select the most efficient route for energy transfer in photosynthesis."

Reference: Engel, G. S., Calhoun, T. R., Read, E. L., Ahn, T. K., Mančal, T., Cheng, Y. C., ... & Fleming, G. R. (2007). Evidence for wavelike energy transfer through quantum coherence in photosynthetic systems. *Nature*, 446(7137), 782-786.

Classical vs Quantum Possibility Spaces

How energy relates to possible outcomes in different physical frameworks

Classical Physics

Possibility space defined by energy potentials — *what is possible is what energy can achieve*

Energy Aggregation

Single ant: Limited energy, small leaf only

Energy Barriers

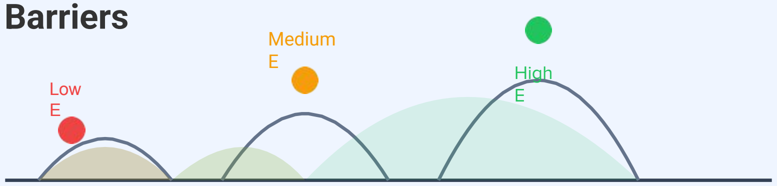

Energy Landscape

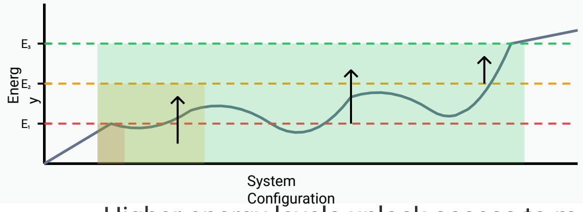

Higher energy levels unlock access to more states

Quantum Physics

Possibility space opens up *within* a given energy — *multiple outcomes possible at same energy*

Double-Slit Experiment

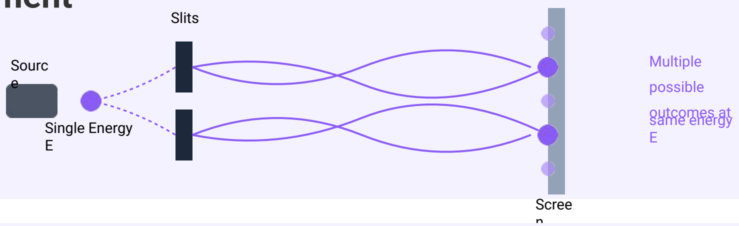

Quantum Superposition

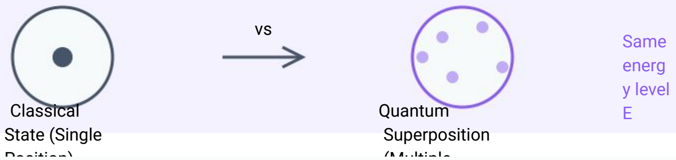

Quantum Possibilities Within Energy Levels

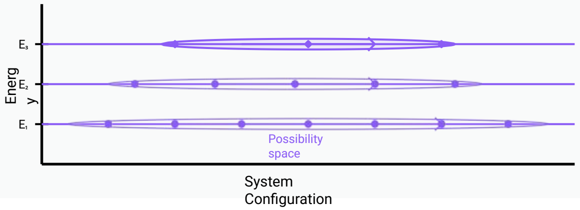

Key Contrast: Classical physics uses energy to select from possible outcomes, quantum physics features multiple outcomes at the same energy—'possibility space' is broader and structured differently

Wayfinding: The Informed Actualization of Possibilities

To model intelligence, we need a new ontology that acknowledges we live in an evolving world of **latent possibilities** that can be selectively actualized.

This actualization process involves the control of material processes through the **informed application of energy**. Quantum theory engages with this ontology, but does not yet provide a model sufficiently advanced to handle biological processes.

> **In a world of infinite latent possibilities and finite resources to identify biologically optimal outcomes and selectively actualize them, we see a place for intelligence.**

It is the task of intelligence to locate, in the tiny subset of biologically relevant possibilities, the most energetically or metabolically efficient way forward.

Wayfinding in action: Children crossing a stream by selecting optimal pathways from many possible stepping stones - a natural metaphor for how intelligence navigates possibilities

Wayfinding identifies the core function of intelligence: to navigate through infinite possibilities and locate pathways that are metabolically efficient and biologically advantageous.

道德经 (Tao Te Ching)

Wayfinding in Pre-Scientific Thinking

The ontology implied in the concept of wayfinding predates the scientific view of the world as a comprehensible machine. Many cultures make pervasive use of this metaphor.

"Tao Te Ching" literally translates to "The Book of the Way and Its Virtue".

道 **Tao (Dao)**: Often interpreted as the fundamental principle of the universe, the natural order, or the underlying reality that governs everything. It can also refer to the path or way of living in harmony with this principle.

德 **Te (De)**: Can be understood as virtue, power, integrity, or the manifestation of Tao in action. It represents the natural power or ability that arises from aligning oneself with the Tao.

經 **Ching (Jing)**: Signifies a classic text, a book, or a scripture.

A plant tendril with spiral curves illustrating natural wayfinding behavior as the plant actively explores space to find support and optimal growth paths.

"The Tao that can be told is not the eternal Tao. The name that can be named is not the eternal name." — Tao Te Ching

Section 2

The Architecture of Intelligence

Examining how intelligence emerges through structured processes of wayfinding, exploring the organizational principles and mechanisms that enable the informed actualization of possibilities

Hierarchical Decision Making

We propose that decision making in any organism or system capable of it can be expressed as a wayfinding problem

Hierarchical Decision Making

We propose that decision making in any organism or system capable of it can be expressed as a wayfinding problem

Example 1: Beijing International → Peking University

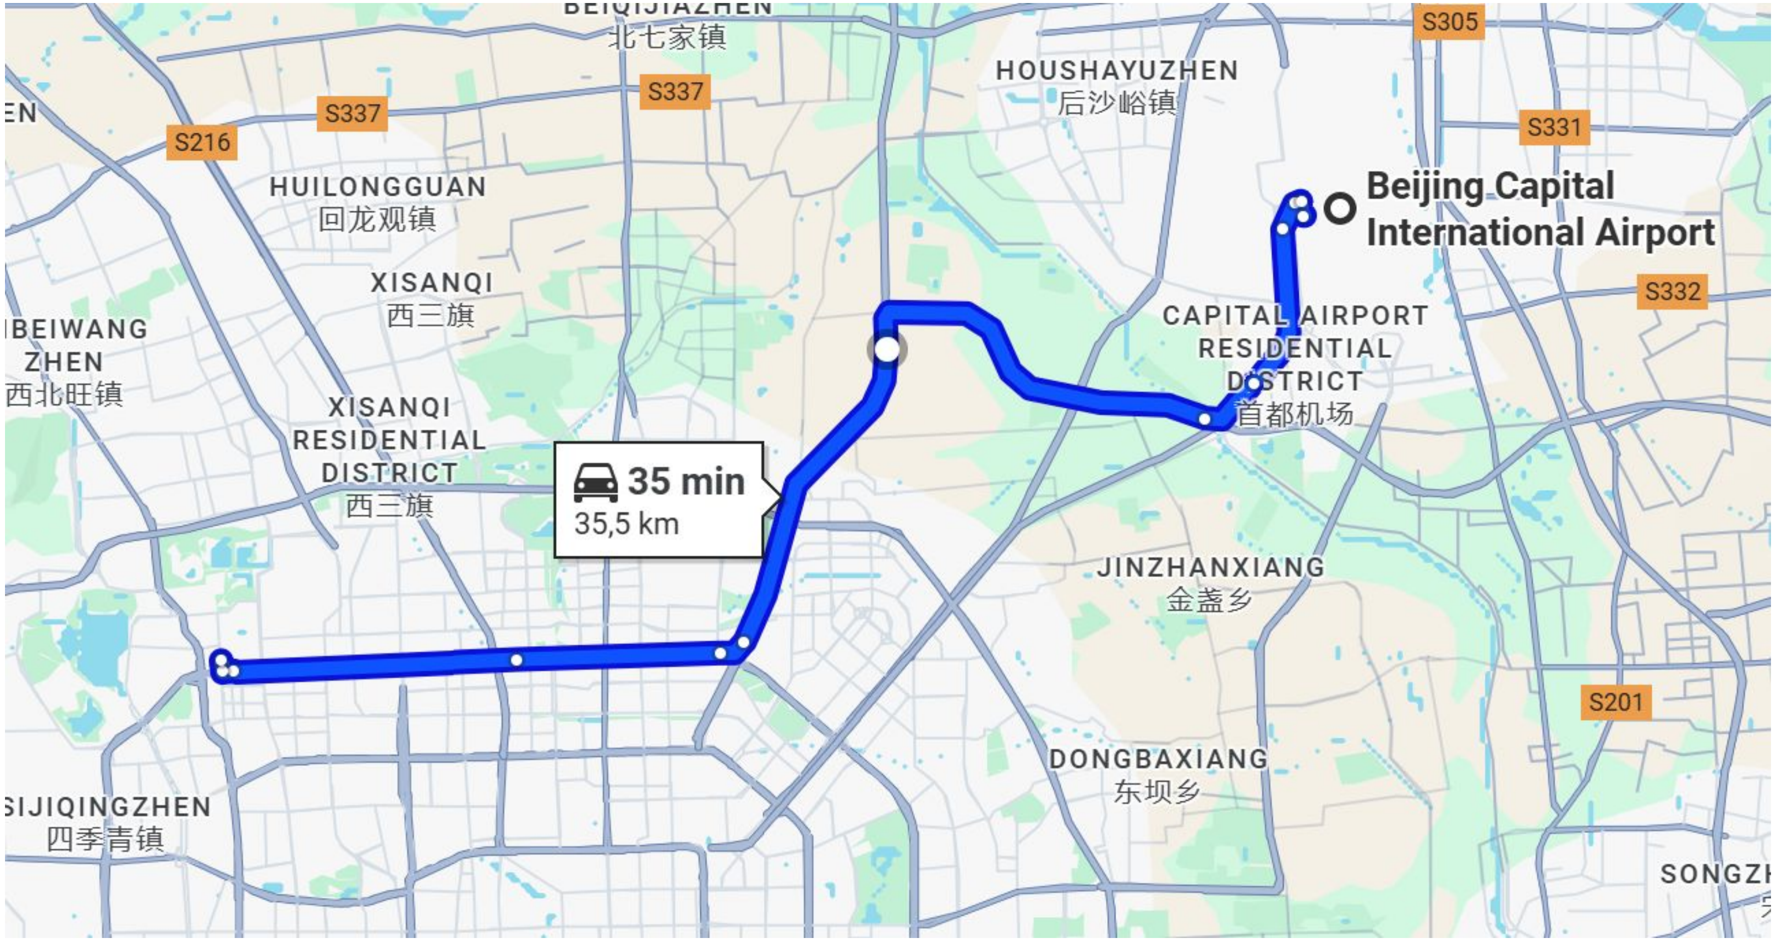

Hierarchical Decision Making

We propose that decision making in any organism or system capable of it can be expressed as a wayfinding problem

Example 1: Beijing International → Peking University

Hierarchical Decision Making

We propose that decision making in any organism or system capable of it can be expressed as a wayfinding problem

Example 1: Beijing International → Peking University

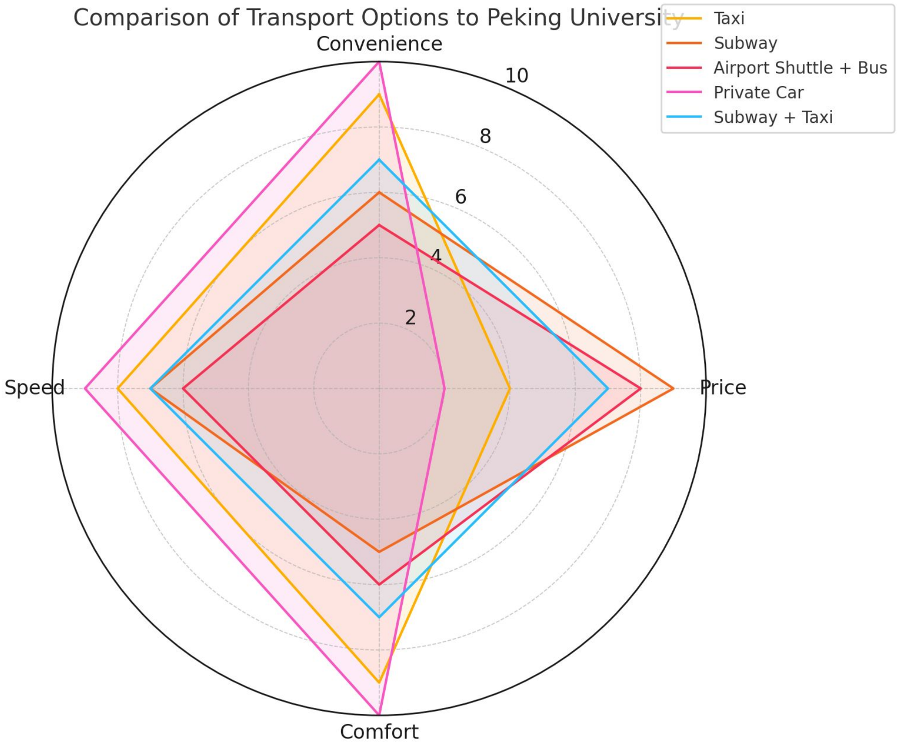

Hierarchical Decision Making

We propose that decision making in any organism or system capable of it can be expressed as a wayfinding problem

Hierarchical Decision Making

In Abstraction

1. Identify Goal Area

Hierarchical Decision Making

In Abstraction

1. Identify Goal Area

Hierarchical Decision Making

In Abstraction

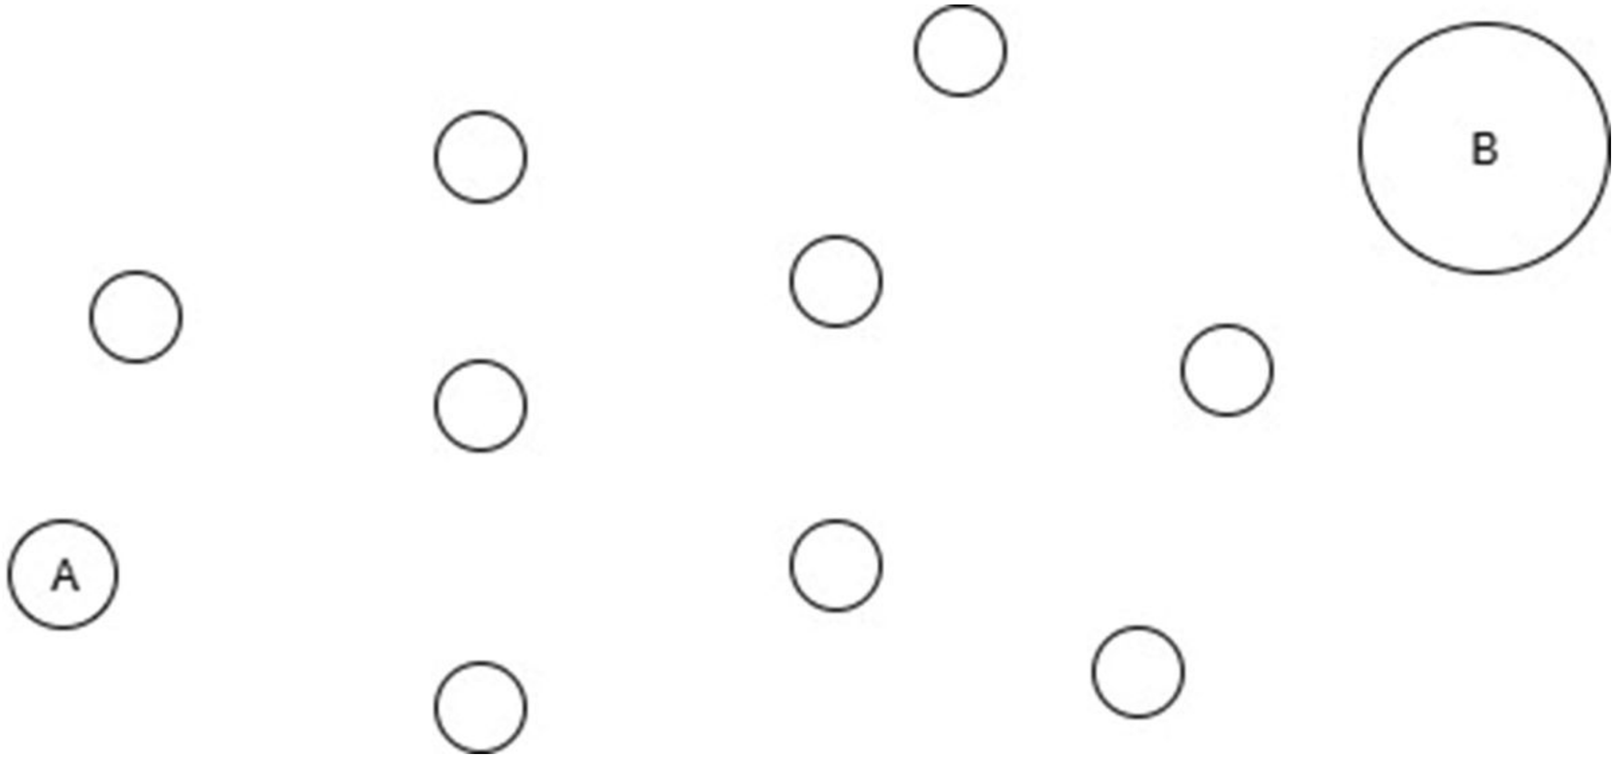

1. Identify Goal Area
2. Identify Nodes “Waypoints”

Hierarchical Decision Making

In Abstraction: *Reactionary Agent*

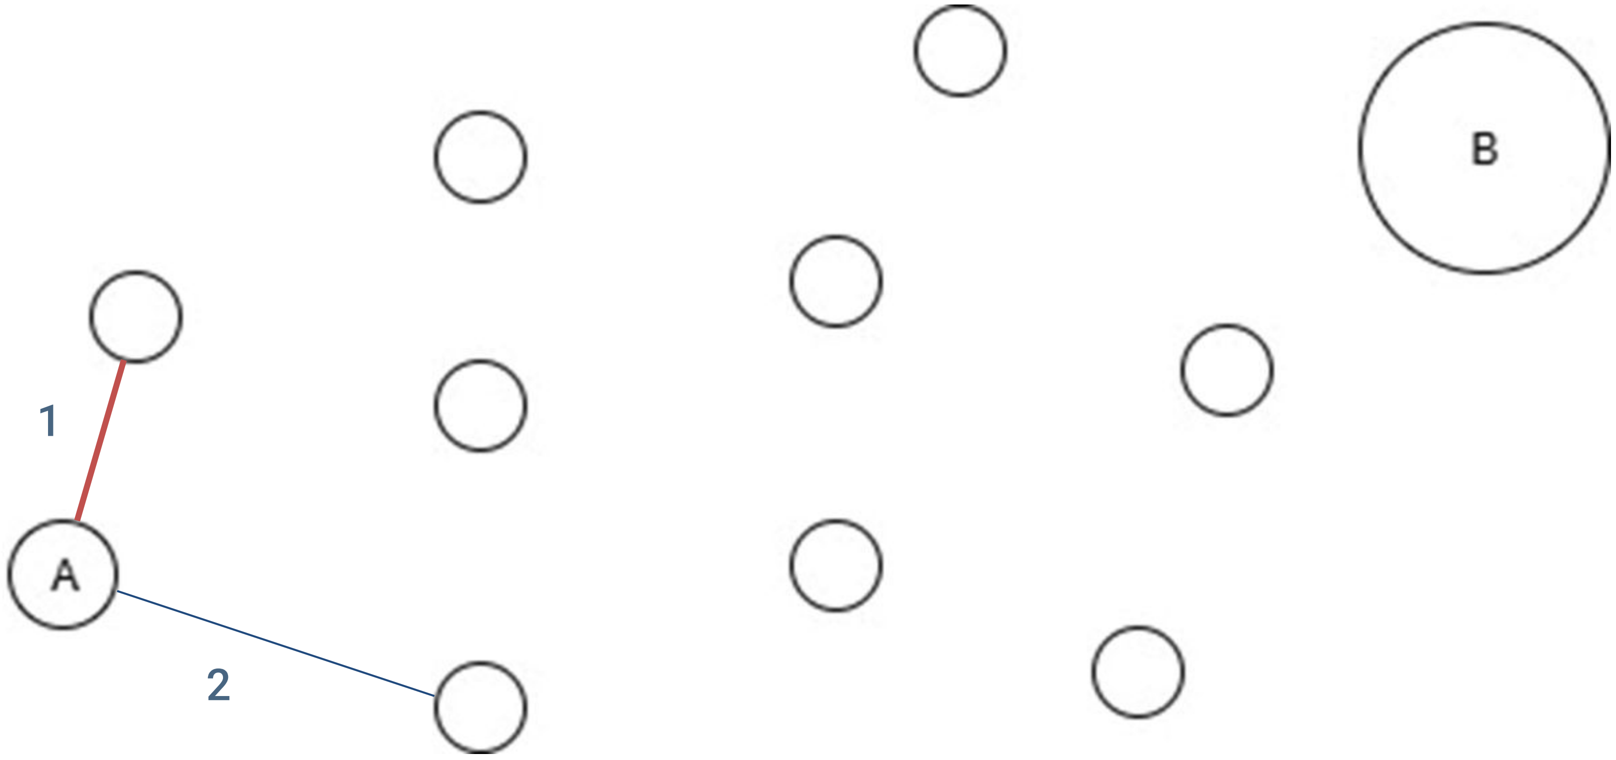

1. Identify Goal Area
2. Identify Nodes “Waypoints”

Hierarchical Decision Making

In Abstraction: *Reactionary Agent*

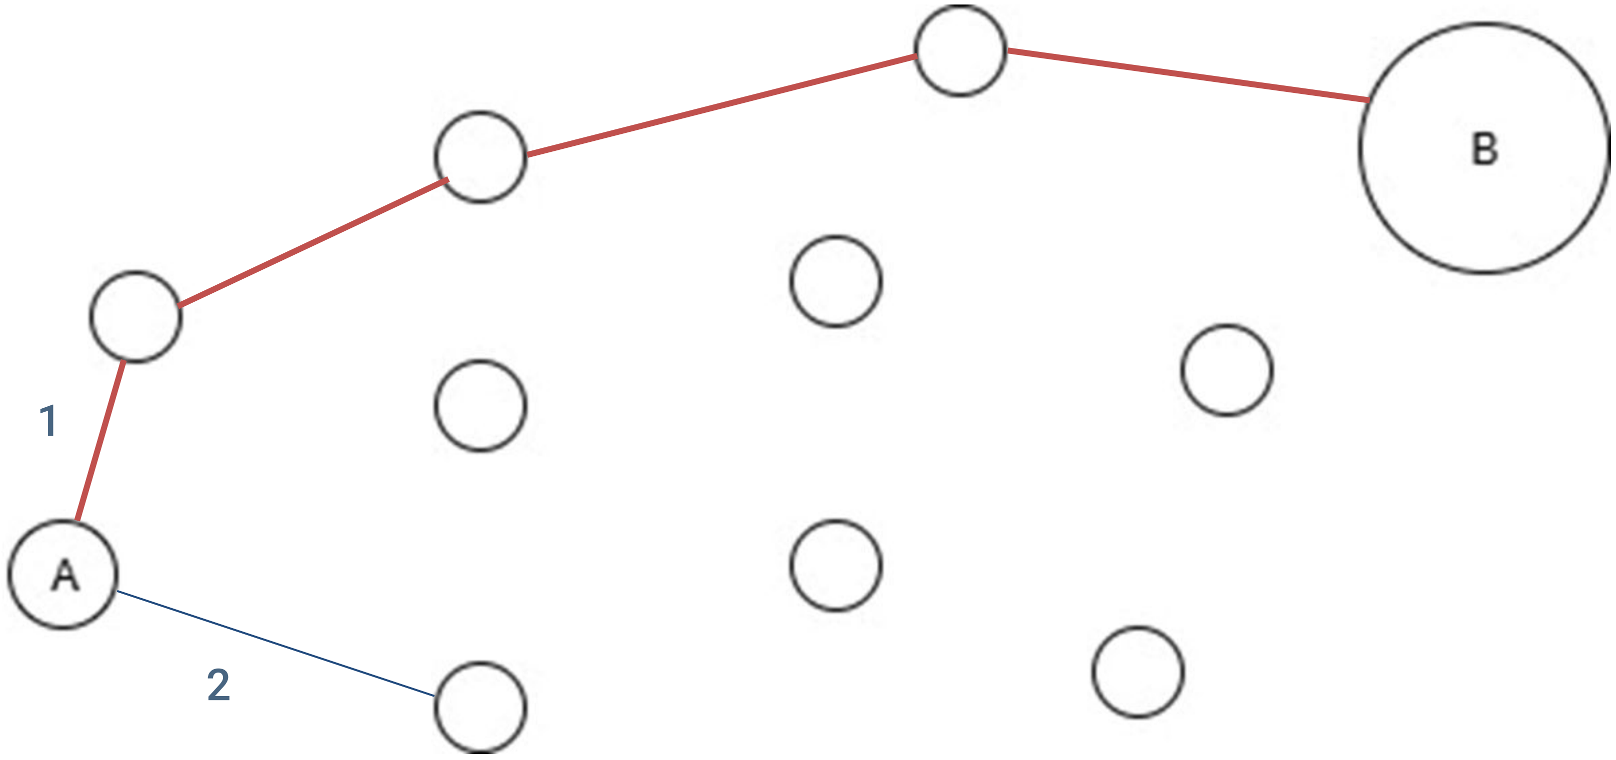

1. Identify Goal Area
2. Identify Nodes “Waypoints”

Aneural learning

In 1906, experiments by Herbert Spencer Jennings (right) suggested that the single-celled ciliate Stentor roeselii (left) can learn — that it was “changed by the experiences it has passed through” (Jennings 1976/1906).

This work was long thought unreplicable and discredited, but was recently confirmed by Gunawardena’s lab at Harvard Medical School (Dexter et al. 2019).

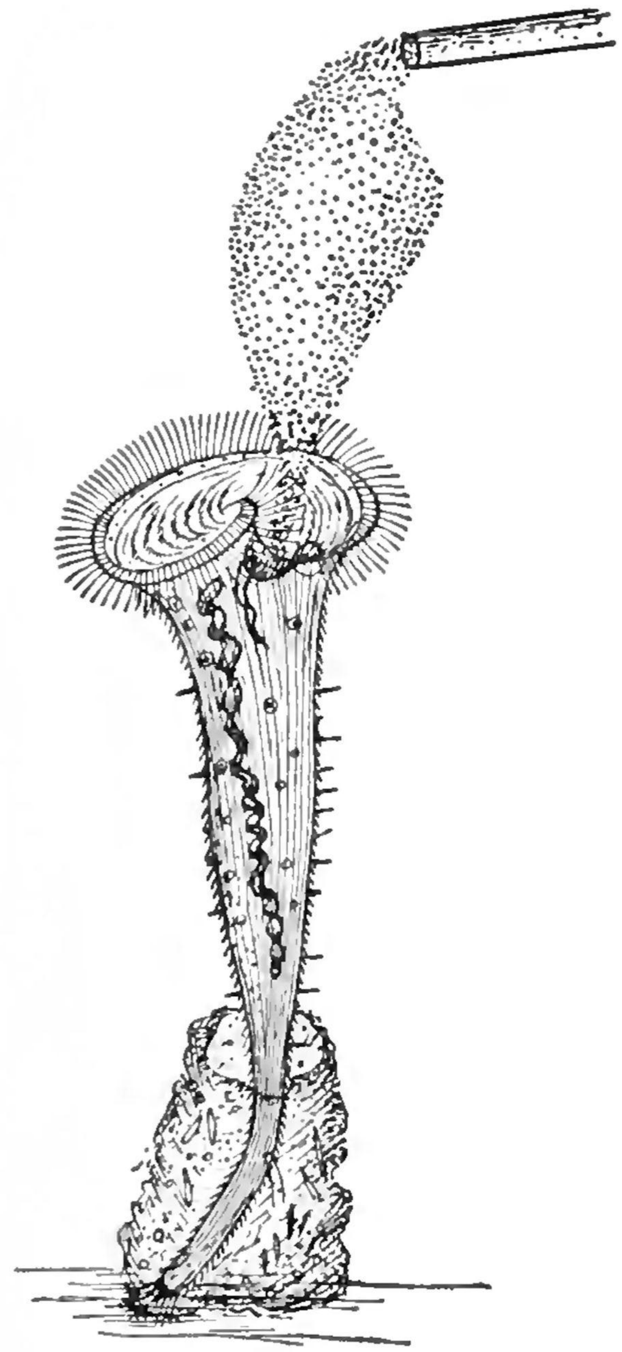

Jennings, H. S. *Behavior of the lower organisms*. Bloomington, IN: Indiana University Press, 1976 (reprint of the 1906 edition) p. 175.

Unicellular Information Processing

Decision-Making in Stentor roeselii

In 1906, the biologist Herbert Spencer Jennings reported that *Stentor roeselii* exhibited signs of complex behavior, potentially indicating a form of decision-making (Jennings, 1931).

In 2019, Jenning's results were verified by Dexter et al. (Dexter 2019). Specifically the protist follows an escalatory sequence of the following actions to avoid toxins in their environment:

Escalatory Avoidance Sequence:

1. **Bend away** from the toxins
2. **Reverse the movement** of the cilia to expel the toxins
3. **Close off its trumpet** to protect the oral opening
4. **Detach and swim away**

This hierarchical response suggests that even single-celled organisms without nervous systems possess sophisticated information processing capabilities.

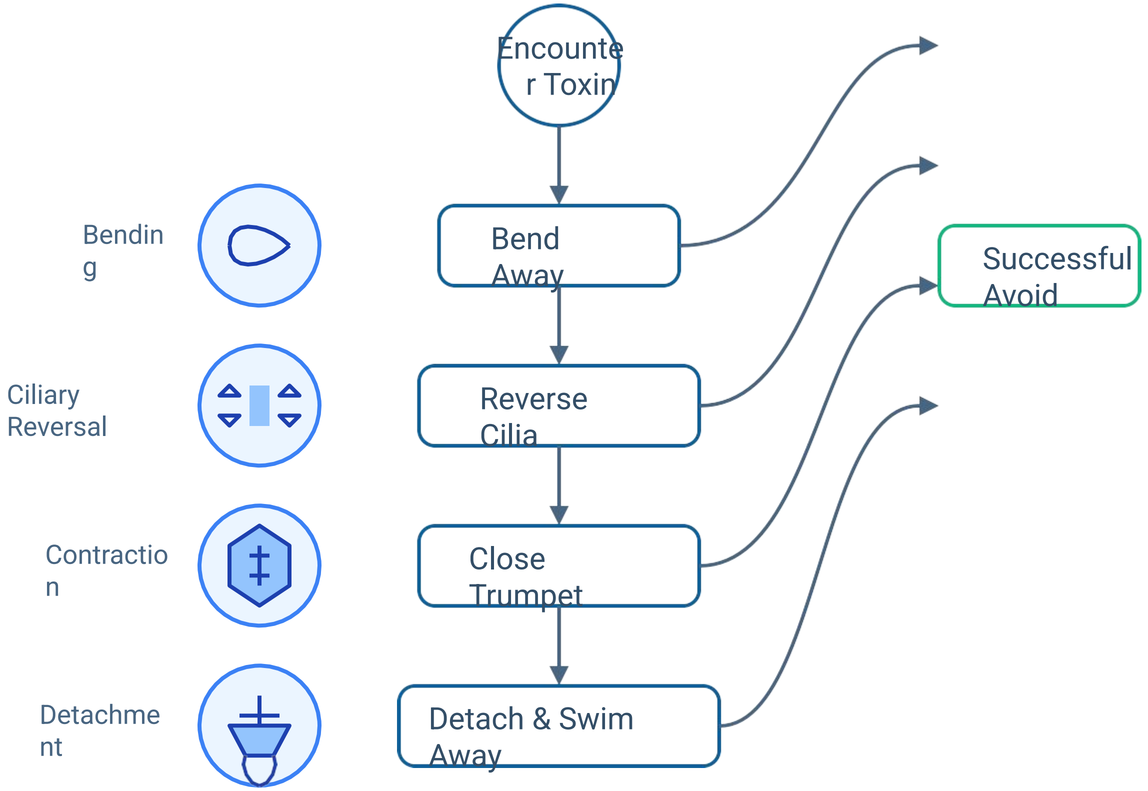

Stentor roeselii follows a hierarchical avoidance response sequence. The organism attempts different strategies in sequence until successfully avoiding toxins in its environment.

References: Jennings, H.S. (1931). Behaviour of the Lower Organisms. Dexter, J.P., et al. (2019). A Complex Hierarchy of Avoidance Behaviors in a Single-Cell Eukaryote. Current Biology, 29(24), 4323-4329.

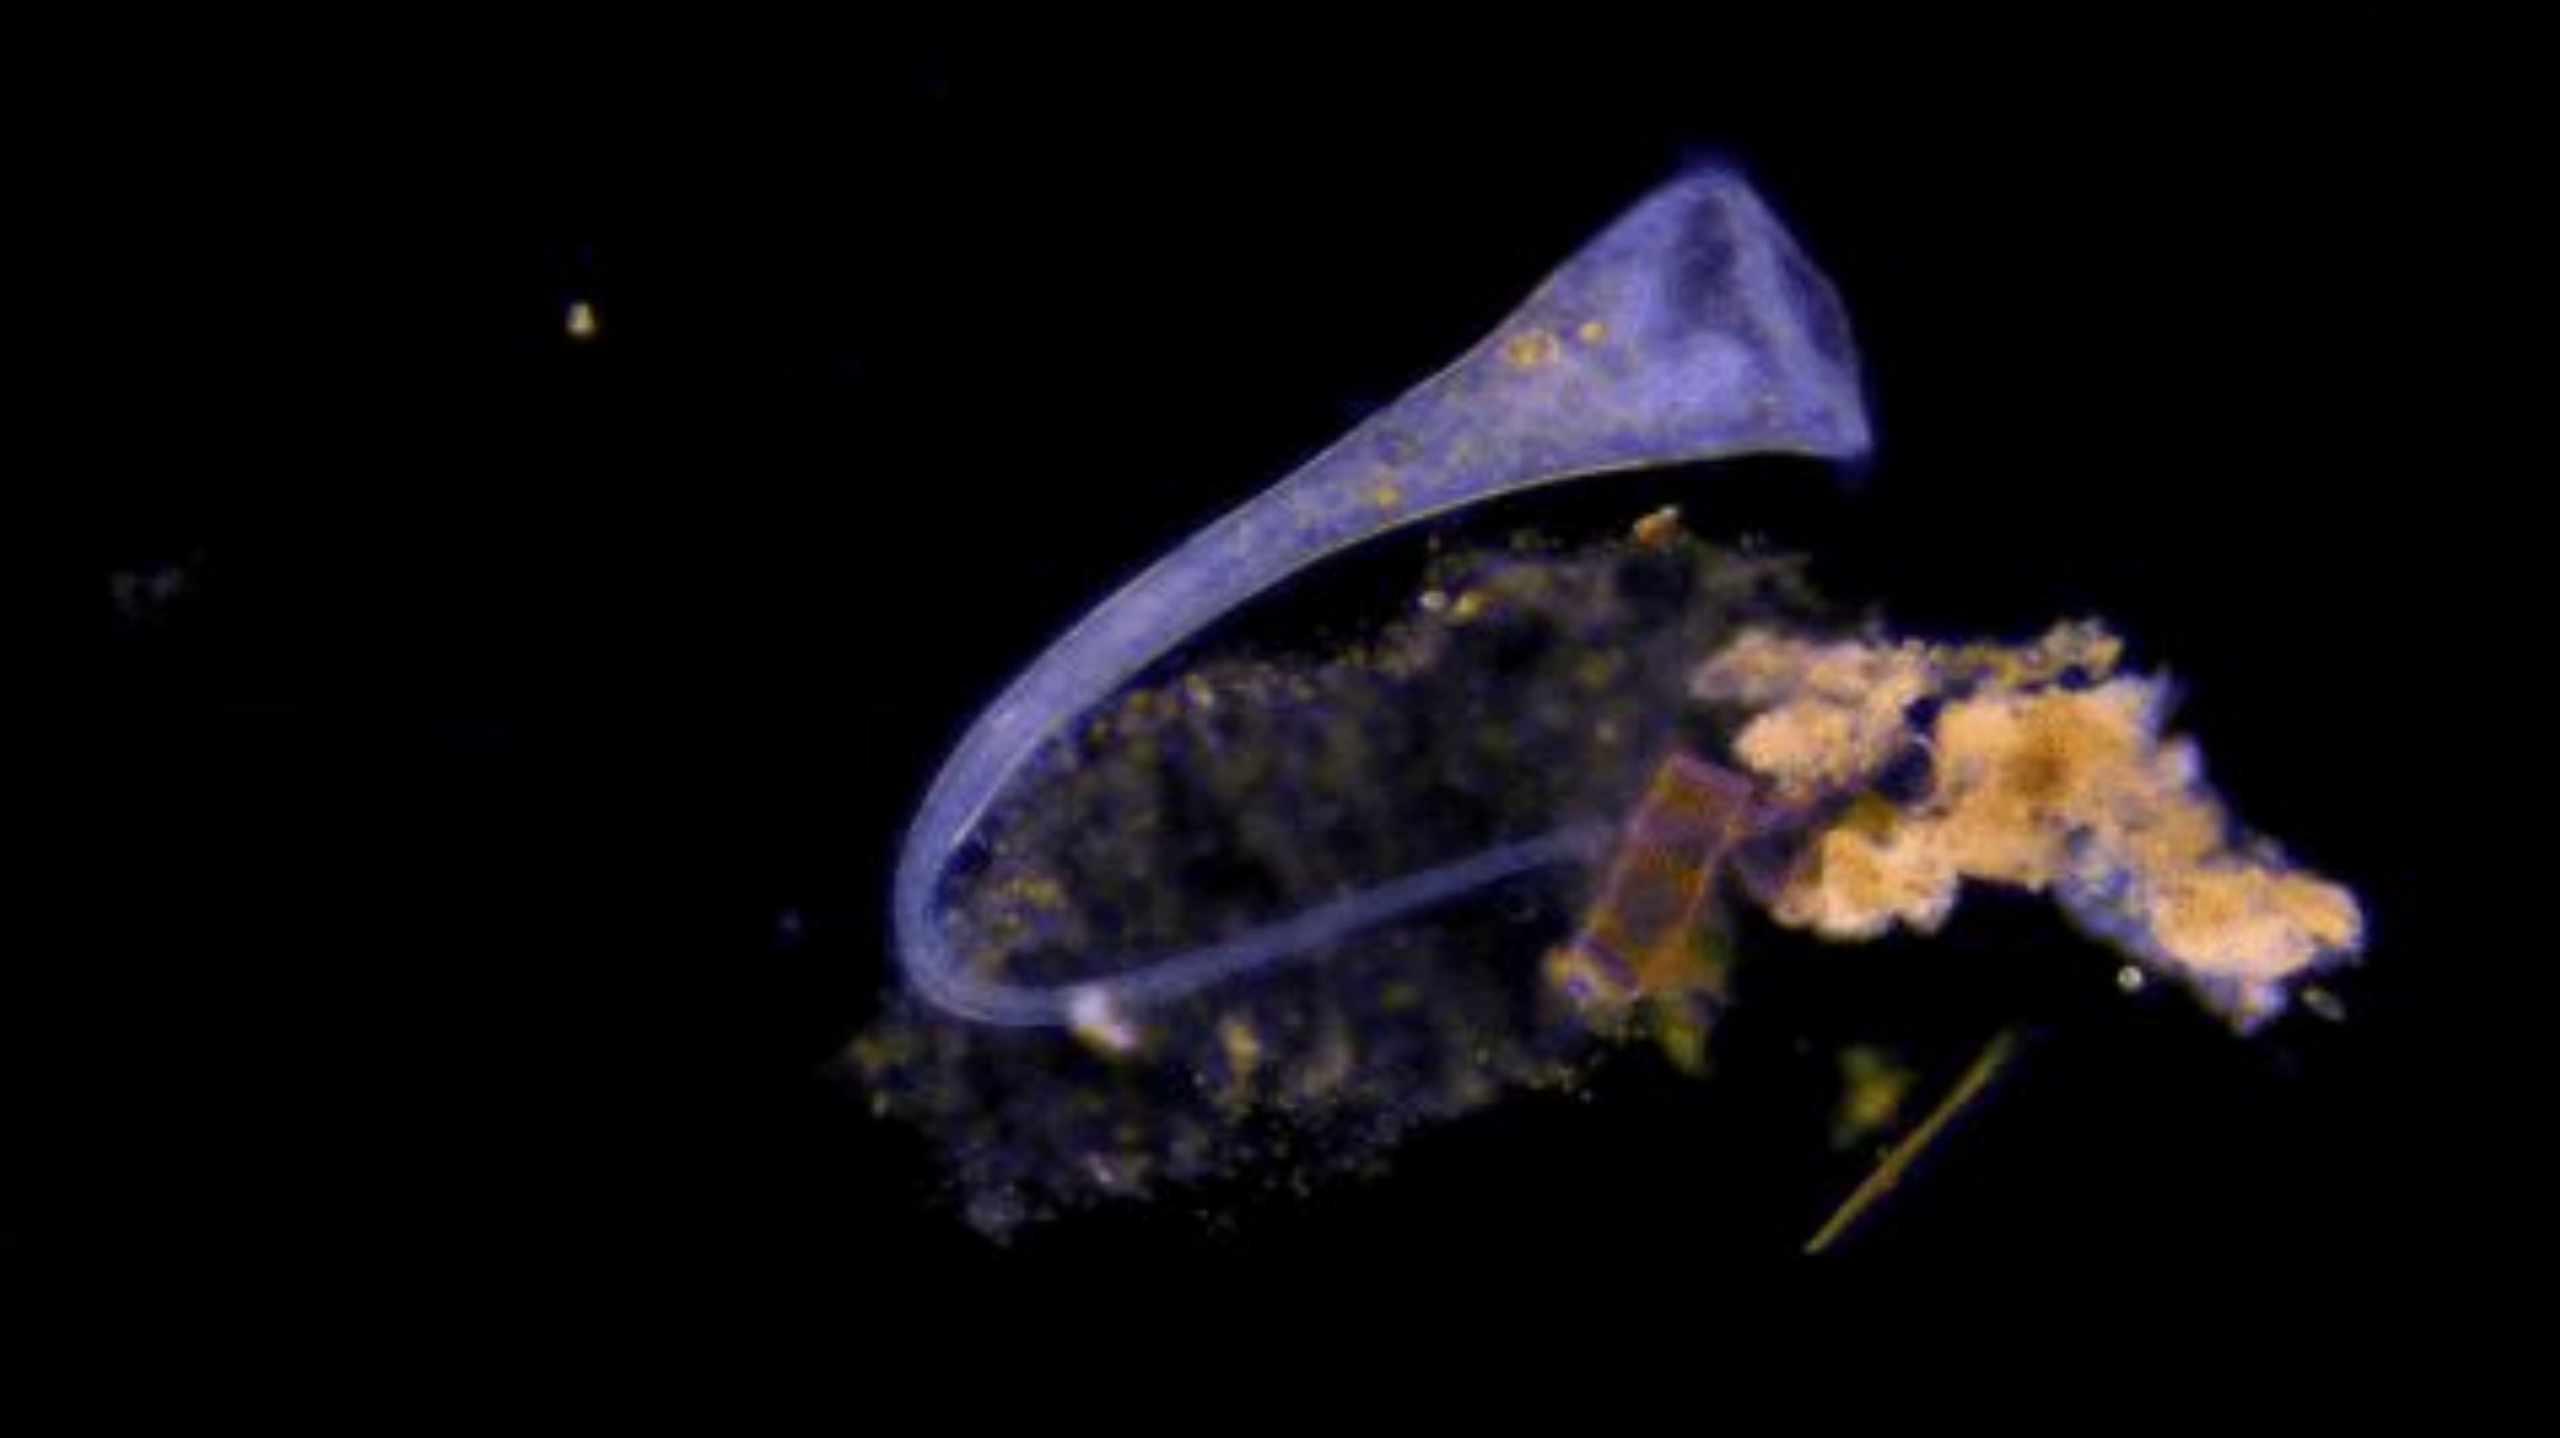

Unicellular Decision Making

Multi-level hierarchical decision making in stentor roselii fits the graph based model

<u>S. Roselii's Heuristically selected "waypoints" and promising path</u>

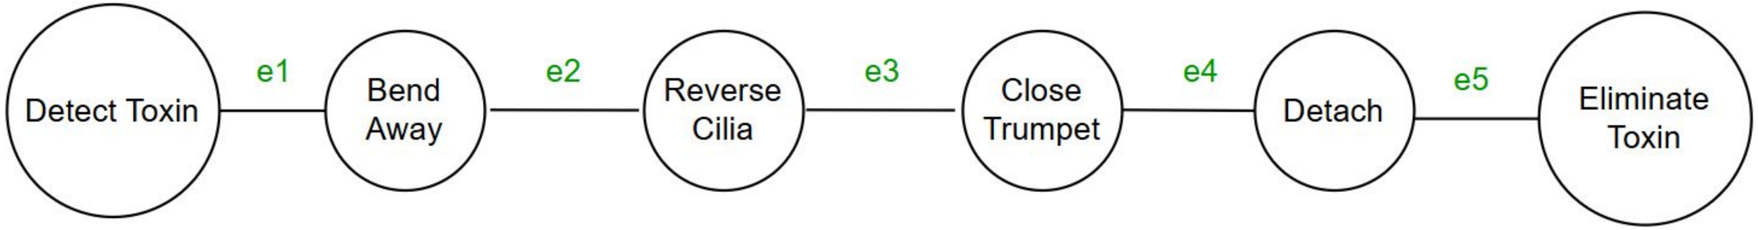

Dexter, J. P., Prabakaran, S., & Gunawardena, J. (2019). A complex hierarchy of avoidance behaviors in a single-cell eukaryote. Current biology, 29(24), 4323–4329

Unicellular Decision Making

Multi-level hierarchical decision making in stentor roselii fits the graph based model

S. Roselii's Heuristically selected "waypoints" and promising path

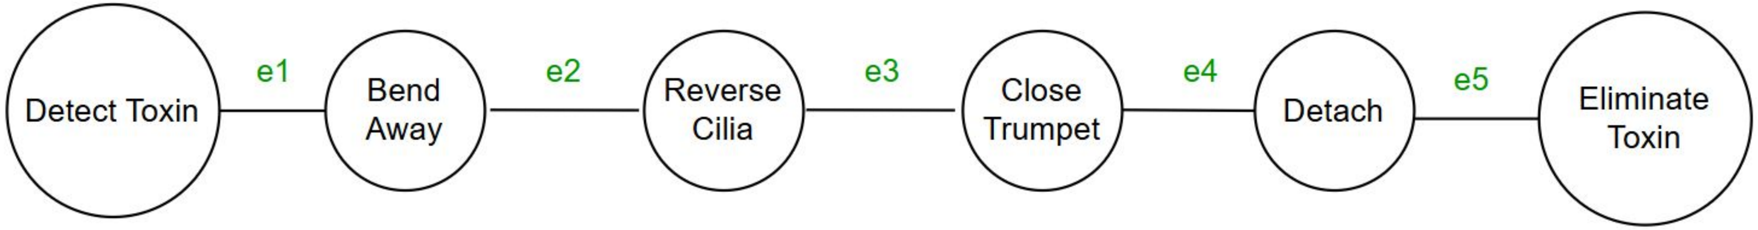

$$e_i = \begin{bmatrix} \Delta\text{Energy} \\ \Delta\text{Temperature} \\ \Delta\text{Algae} \\ \Delta\text{Bacteria} \\ \Delta\text{Oxygen} \\ \Delta\text{Salinity} \\ \Delta\text{pH} \end{bmatrix}$$

$$\underset{S \subseteq \{1,\ldots,d\}}{\text{minimize}} \sum_{i \in S} e_i^1, \quad \text{s.t} \quad \left| \sum_{i \in S} e_i^n \right| \leq M_n \quad \forall n$$

Dexter, J. P., Prabakaran, S., & Gunawardena, J. (2019). A complex hierarchy of avoidance behaviors in a single-cell eukaryote. Current biology, 29(24), 4323–4329

Modular Competencies

Ciliary architecture in Stentor roeselii

Each of *S. roeselii*'s strategies involves hundreds of cilia working in coordination. The organism's structural organization includes:

- **Kineties** — rows of cilia arranged longitudinally from the anterior to posterior end
- **Membranelles** — the oral region is encircled by a spiraling row of specialized ciliary structures

> **Cilium Structure:**
>
> Each cilium consists of a 9+2 arrangement of microtubules controlled by individual rotors at the base and dynein proteins along the body.

The coordination of all these components into a coherent avoidance response requires highly precise orchestration. Strategy selection likely involves simple signals that call on existing lower-level competencies.

The ciliary system demonstrates how complex behaviors can emerge from the coordination of simpler modular components, each with specific functions:

- Individual rotors provide the base motion
- Dynein proteins adjust the flexibility and direction
- The 9+2 microtubule structure provides stability and form

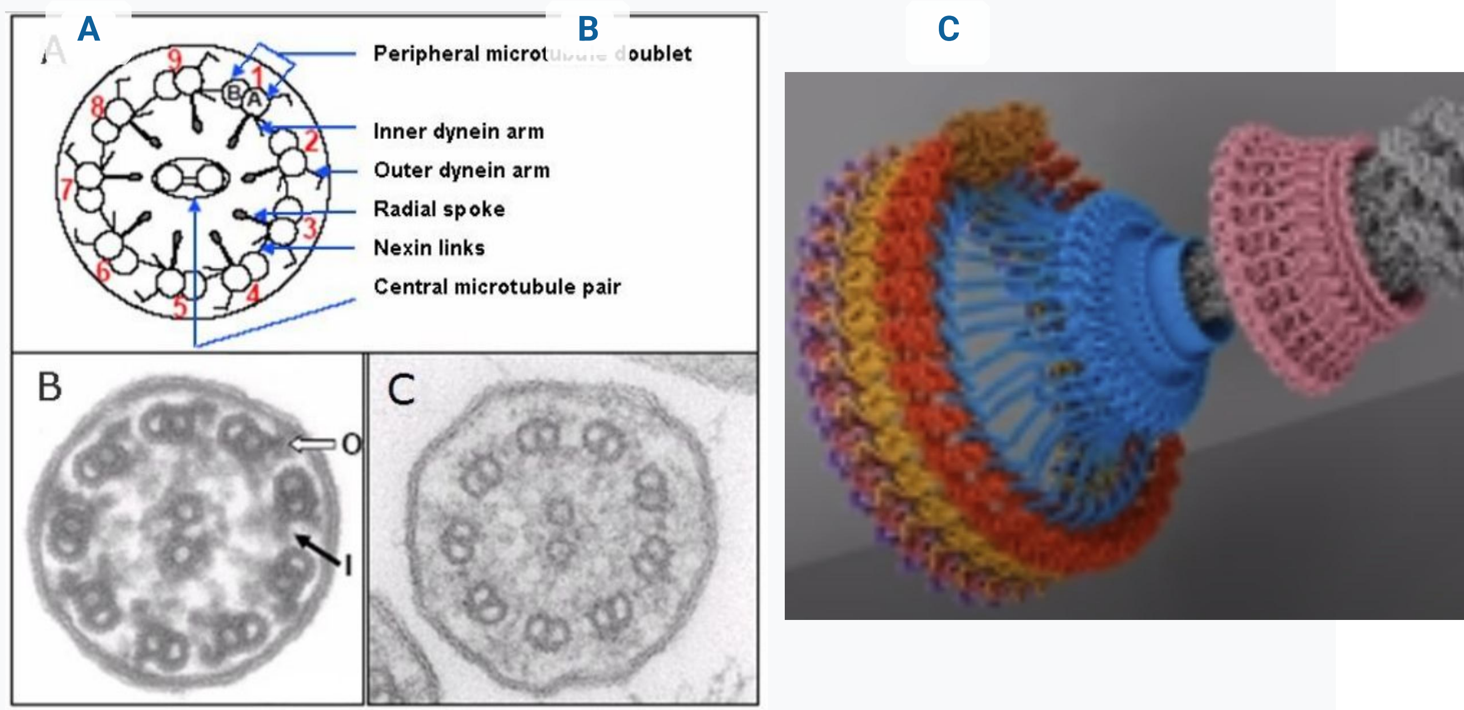

*Cross-sectional view of cilium structure. **A:** Diagram showing the 9+2 arrangement with peripheral microtubule doublets, dynein arms, and central pair. **B & C:** Electron microscopy images of cross-section. Inset: Molecular model of the microtubule rotor base structure responsible for ciliary movement in Stentor roeselii.*

The Goldilocks Zone

Environmental parameters for Stentor roeselii survival

Life depends on a suite of variables remaining within a given range. Together, these variables define the zone of survival or the **Goldilocks Zone**.

For *Stentor roeselii*, the organism can survive in an environment that is within a certain range of temperature, salinity, acidity, oxygen concentration, and nutrient availability.

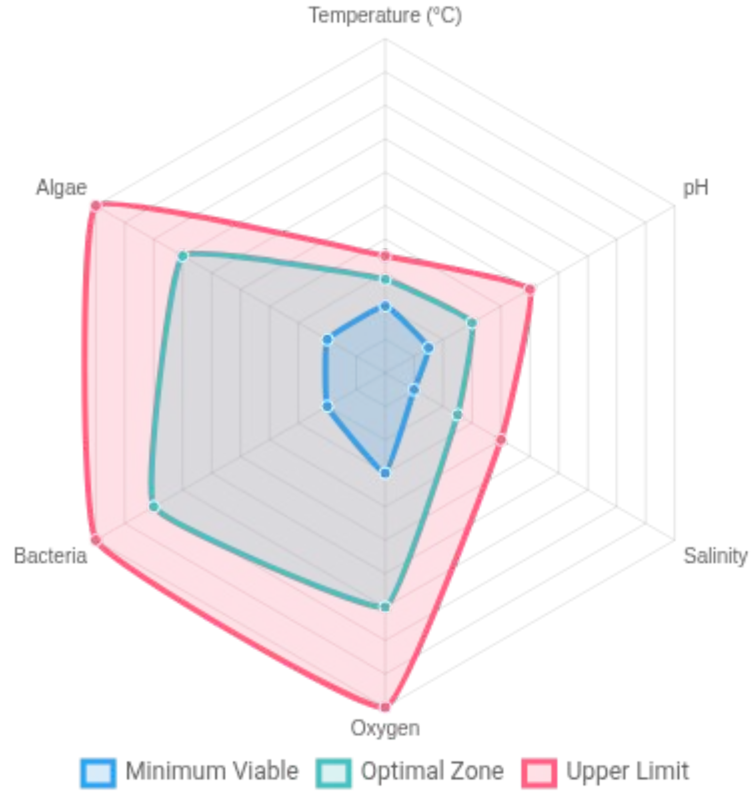

Key Environmental Parameters:

- **Temperature:** Optimal 25-30°C
- **pH:** Optimal 7-9 (slightly alkaline)
- **Salinity:** Freshwater (low salinity)
- **Oxygen:** Moderate to high levels
- **Food sources:** Bacteria, small algae
- **Organic matter:** Moderate levels

Outside of these ranges, Stentor roeselii cannot maintain homeostasis and struggles to survive. Within the optimal ranges (the Goldilocks Zone), it thrives and reproduces.

Temperature

10° C — 35° C

Optimal: 25-30° C

pH Level

4 — 11

Optimal: 7-9

Salinity Tolerance

0 ppt — 10 ppt

Optimal: 0-5 ppt

Multi-level Decision-making

Responding to multiple environmental dimensions

Jennings (1906) and Dexter (2019) show that *S. roeselii* is capable of responding to the presence of an irritant by an escalatory sequence of avoidance strategies, demonstrating advanced information-processing within a single cell.

Consider that similar strategies may be required to respond to other environmental dimensions to remain within the Goldilocks zone.

> **Multiple Dimension Reaction Model:**
>
> A simple model of how the organism accomplishes this is to assume it has the information-processing capacity to respond to only one survival dimension at a time, triggered by that particular value moving outside of the Goldilocks Zone.

This approach suggests a prioritization system, where the most critical environmental threats receive attention first, followed by secondary concerns when the primary threat is neutralized.

Current Research: Mayank Palan is working on modeling such a system to understand how single-celled organisms make complex, multi-dimensional decisions despite limited information processing capacity.

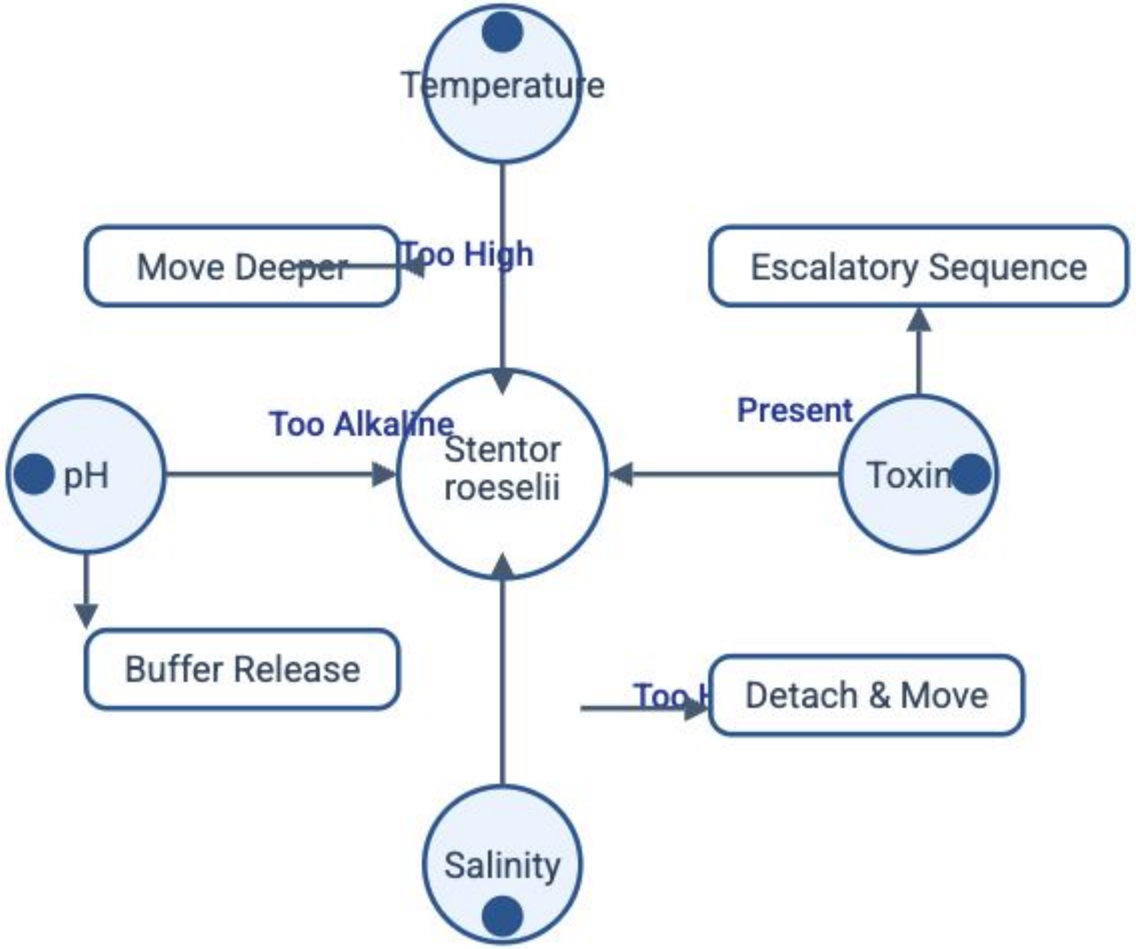

Conceptual model of multi-dimensional decision-making in S. roeselii. The organism prioritizes responses to one environmental dimension at a time

Bioelectrical Switches

Morphogenesis through Electrical Signaling

Michael Levin's lab has shown that headless sections of a planarium flatworm can be prompted to regenerate two heads through simple electrical signals.

These signals contain no information about how to make heads.

Key Findings:

- Bioelectrical signals act as switches that activate complex morphogenetic programs
- A small number of cells prompted to perform a large task will recruit additional cells to complete it
- They not only have the competency to build the target body part, but also know when to stop.

What we see in these experiments is modular competencies flexibly navigating in vast morphospaces and reliably reaching functional phenotypes along pathways that have no evolutionary precedence.

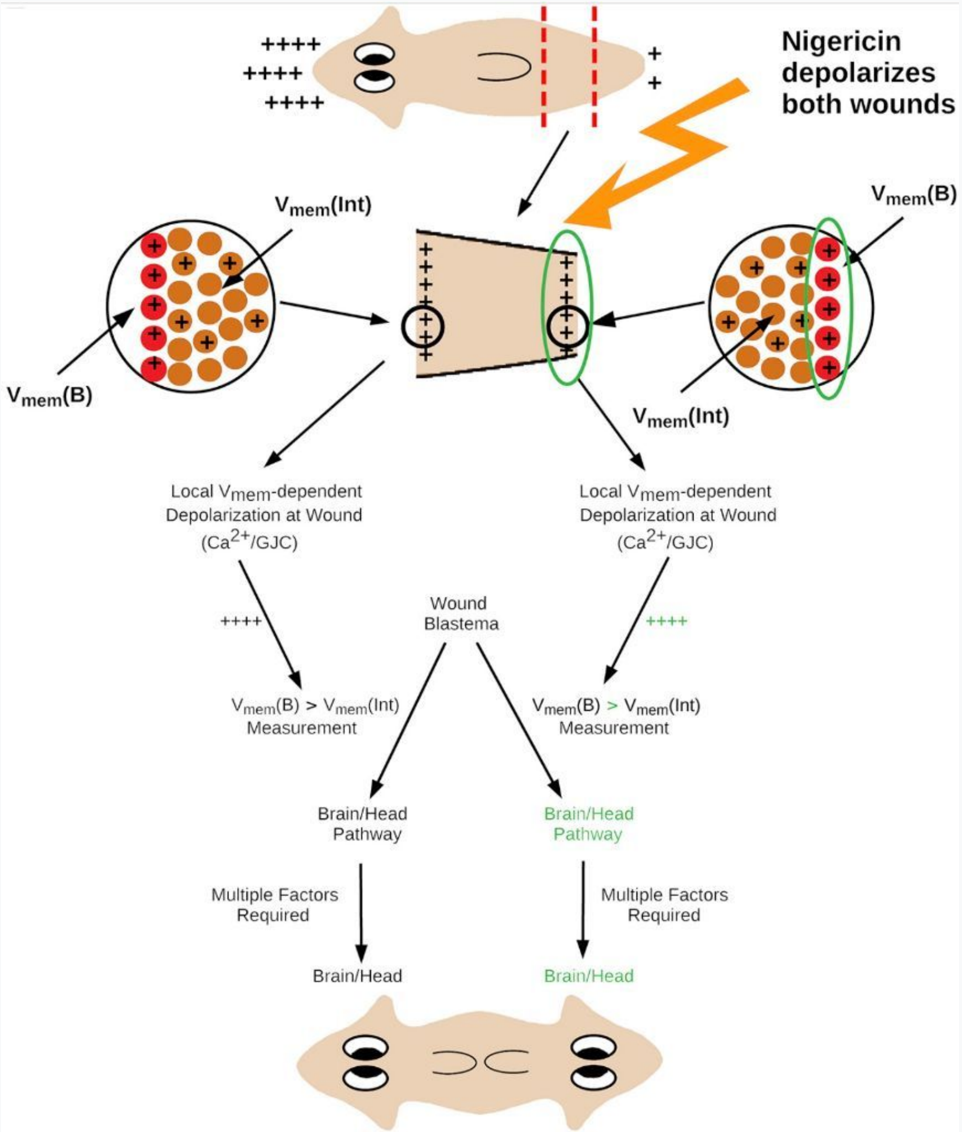

Diagram showing bioelectrical signaling pathways during planarian regeneration. The voltage gradients (Vmem) determine developmental outcomes without containing detailed instructions for body parts.

Hierarchical Decision Making

In Abstraction: *Intermediate Agent*

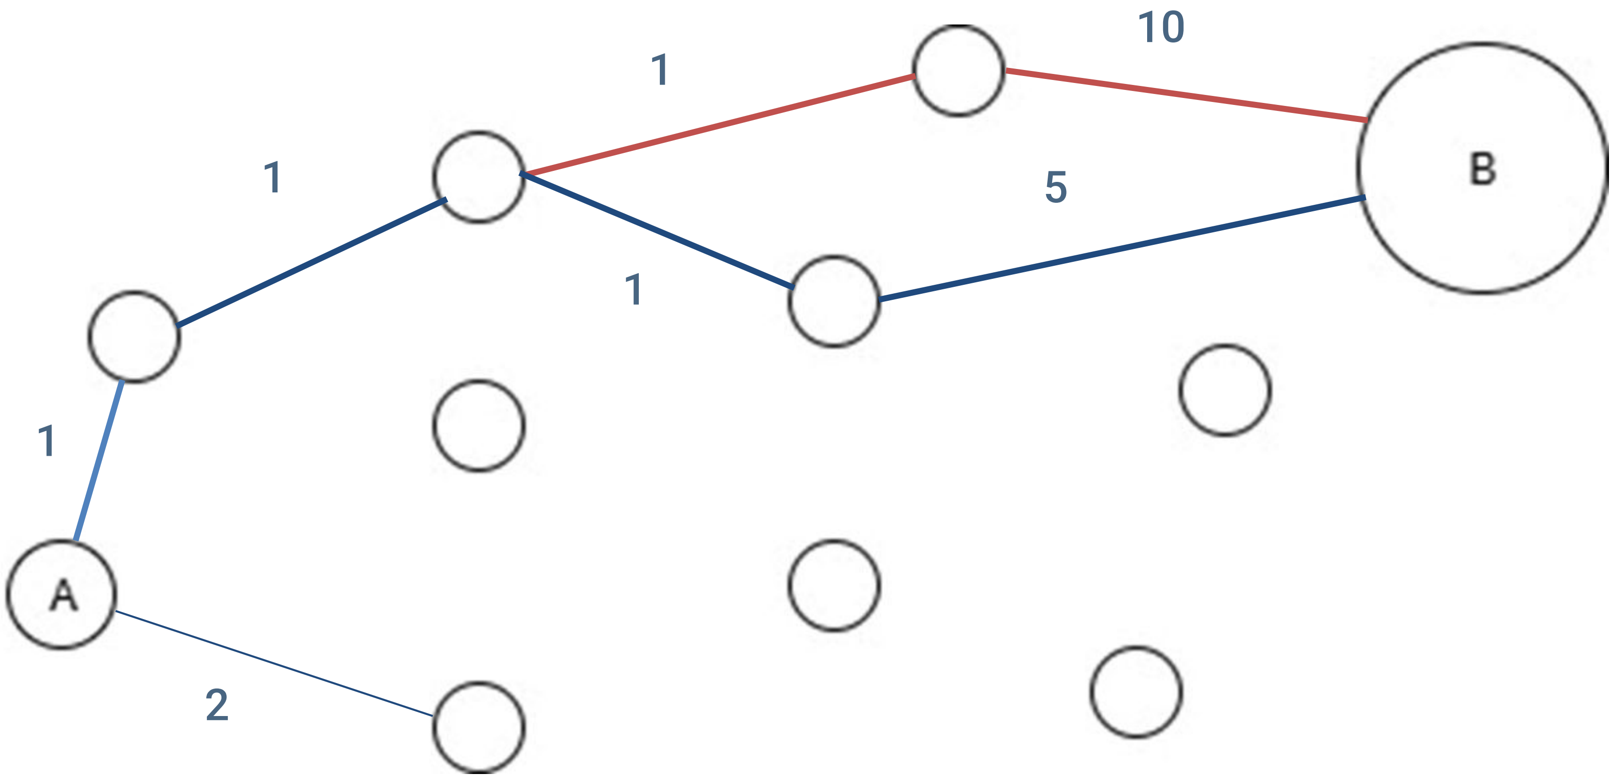

1. Identify Goal Area
2. Identify Nodes "Waypoints"

Unicellular Decision Making

Multi-level hierarchical decision making in stentor roselii fits the graph based model

<u>S. Roselii's Heuristically selected "waypoints" and promising path</u>

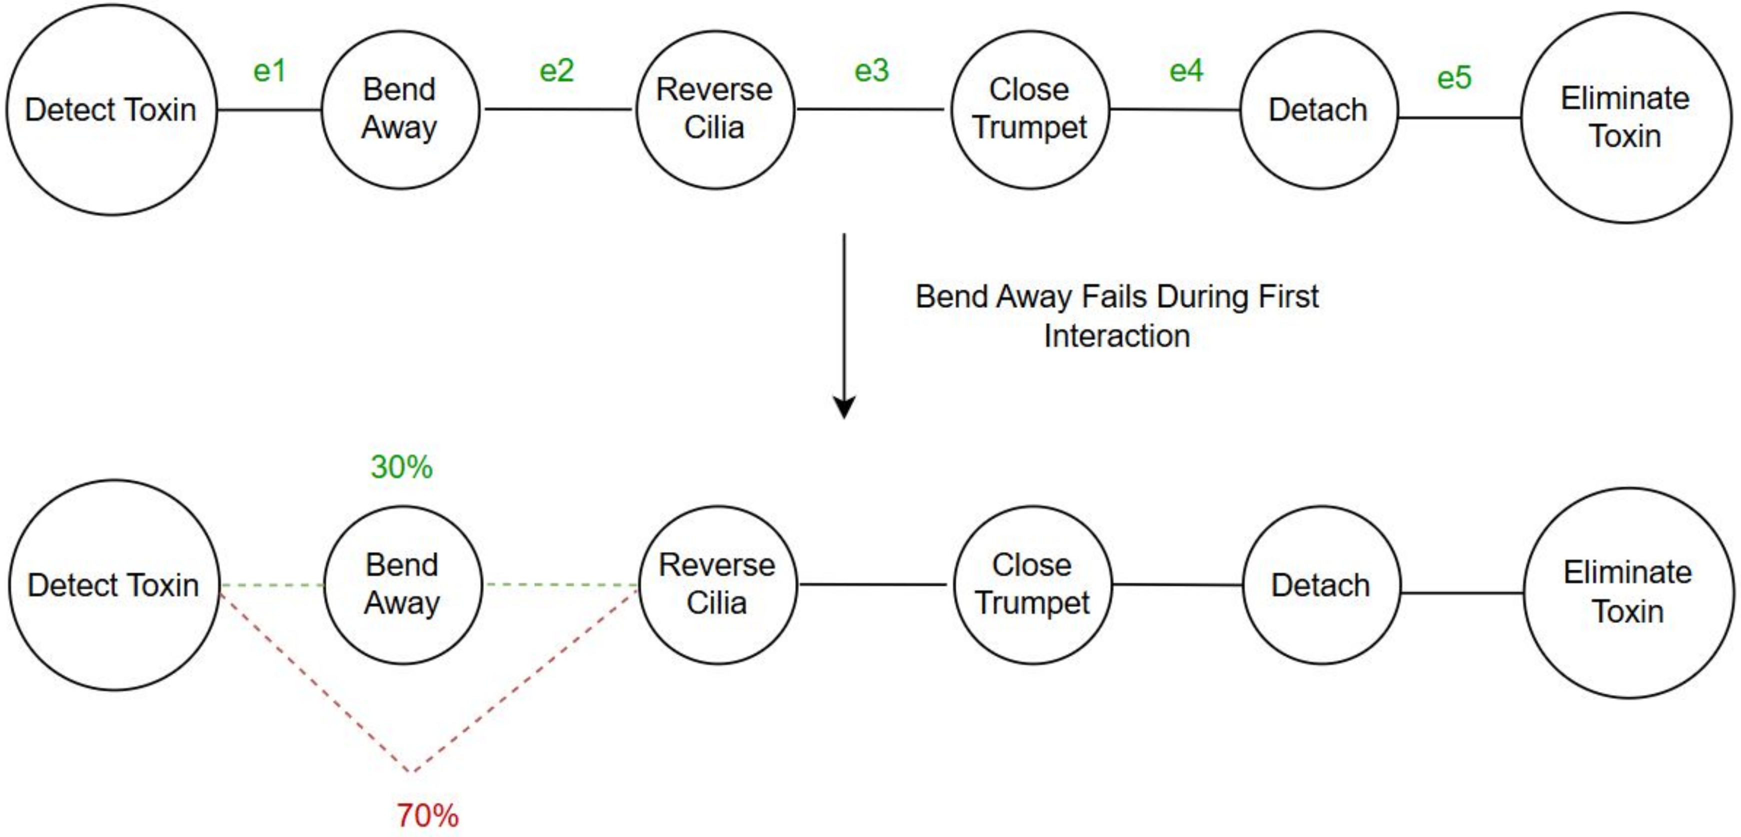

Dexter, J. P., Prabakaran, S., & Gunawardena, J. (2019). A complex hierarchy of avoidance behaviors in a single-cell eukaryote. Current biology, 29(24), 4323–4329

Hierarchical Decision Making

In Abstraction: *Prospecting Agent*

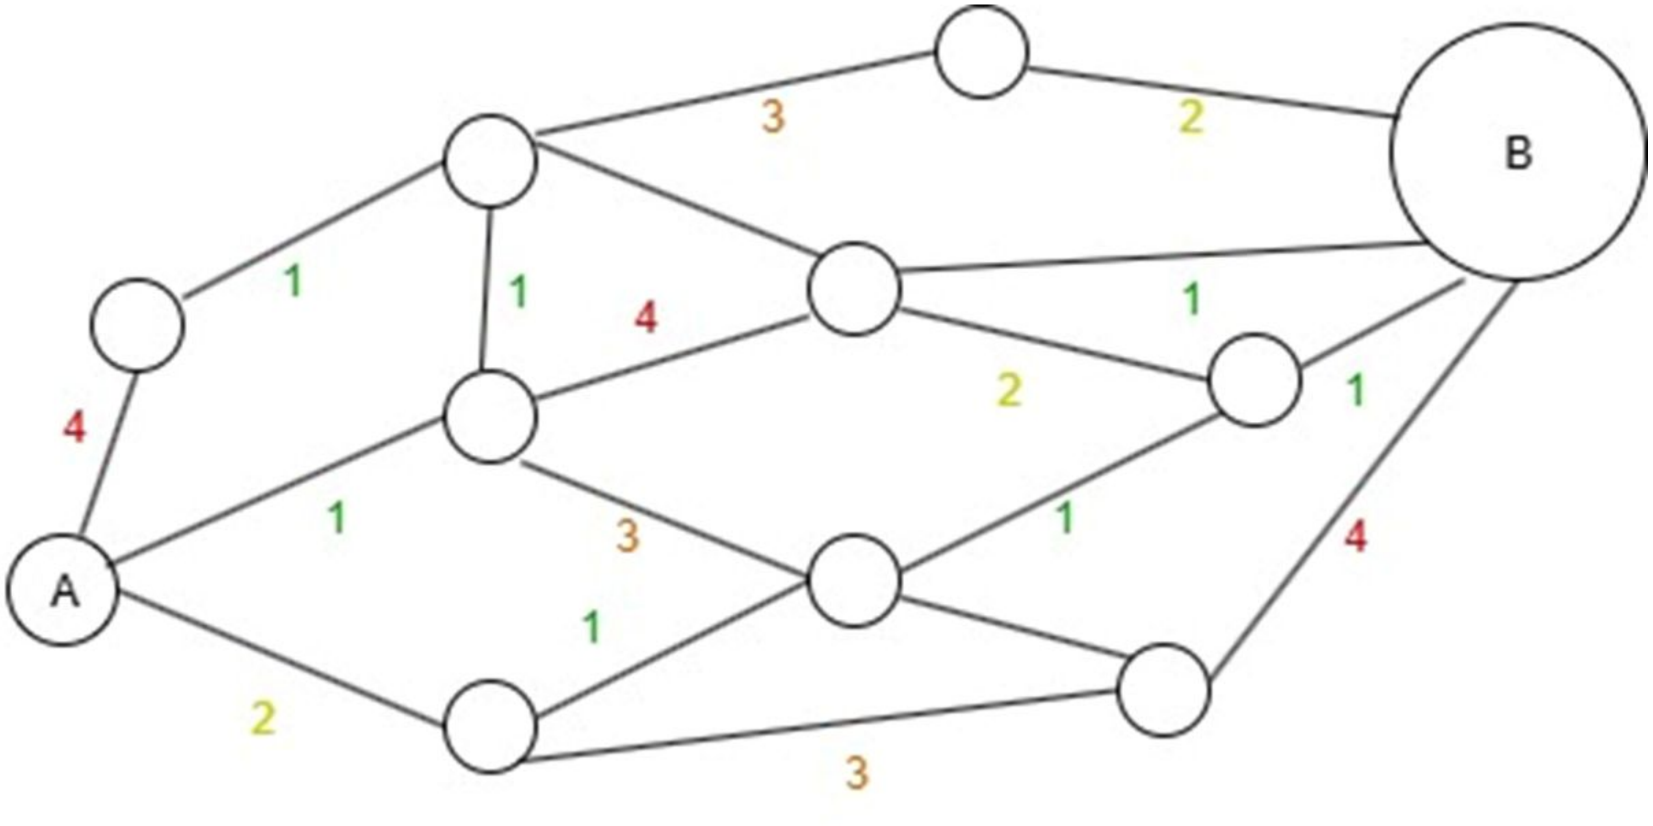

1. Identify Goal Area
2. Identify Nodes “Waypoints”
3. Heuristic Evaluation of Candidate Paths

Hierarchical Decision Making

In Abstraction: *Prospecting Agent*

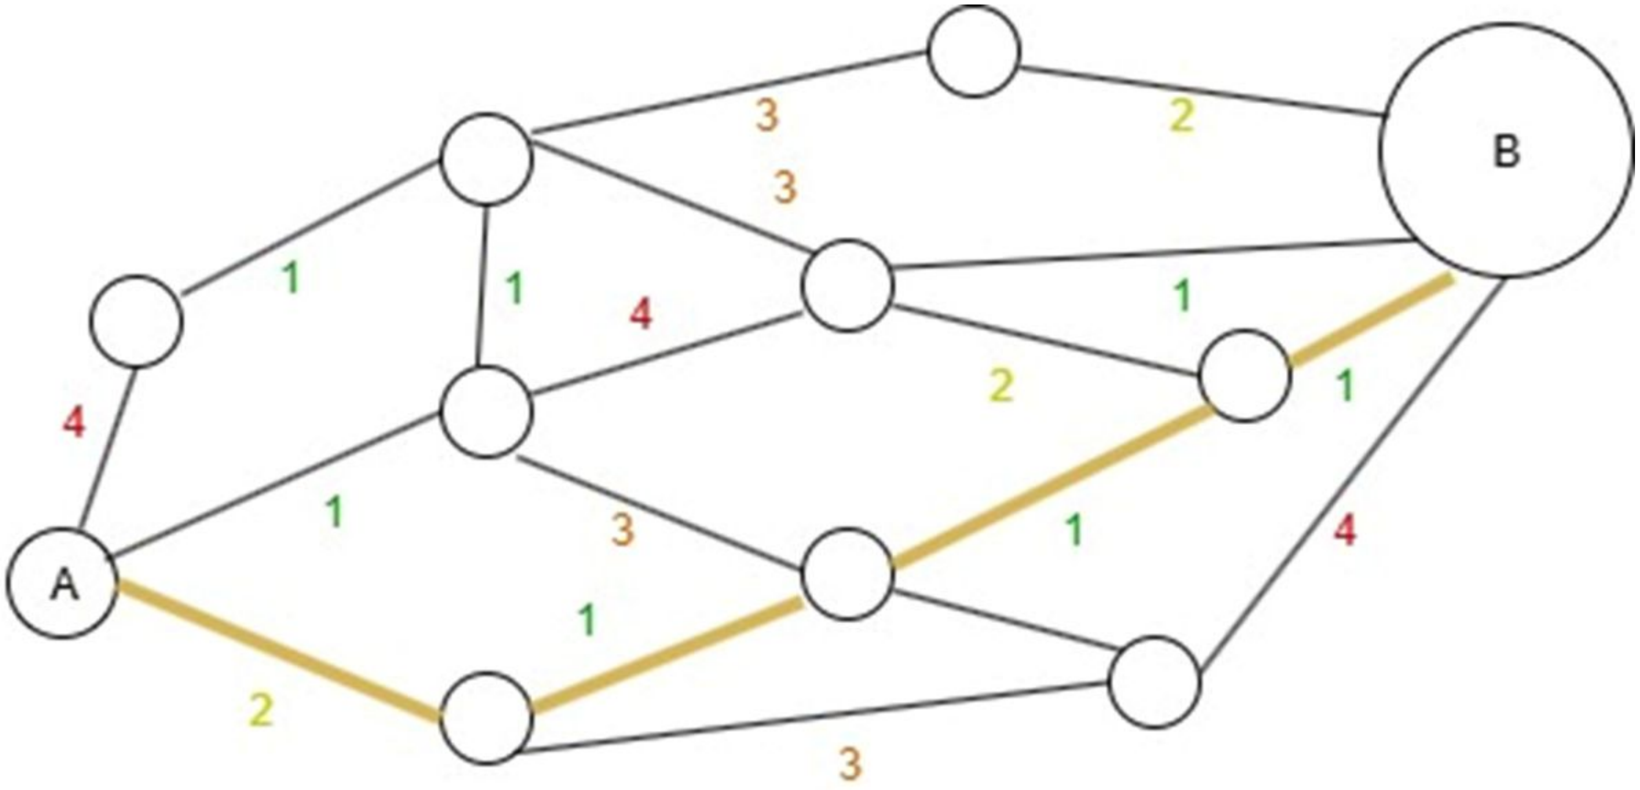

1. Identify Goal Area
2. Identify Nodes "Waypoints"
3. Heuristic Evaluation of Candidate Paths
4. Refine Significant Heuristics

Hierarchical Decision Making

In Abstraction: *Prospecting Agent*

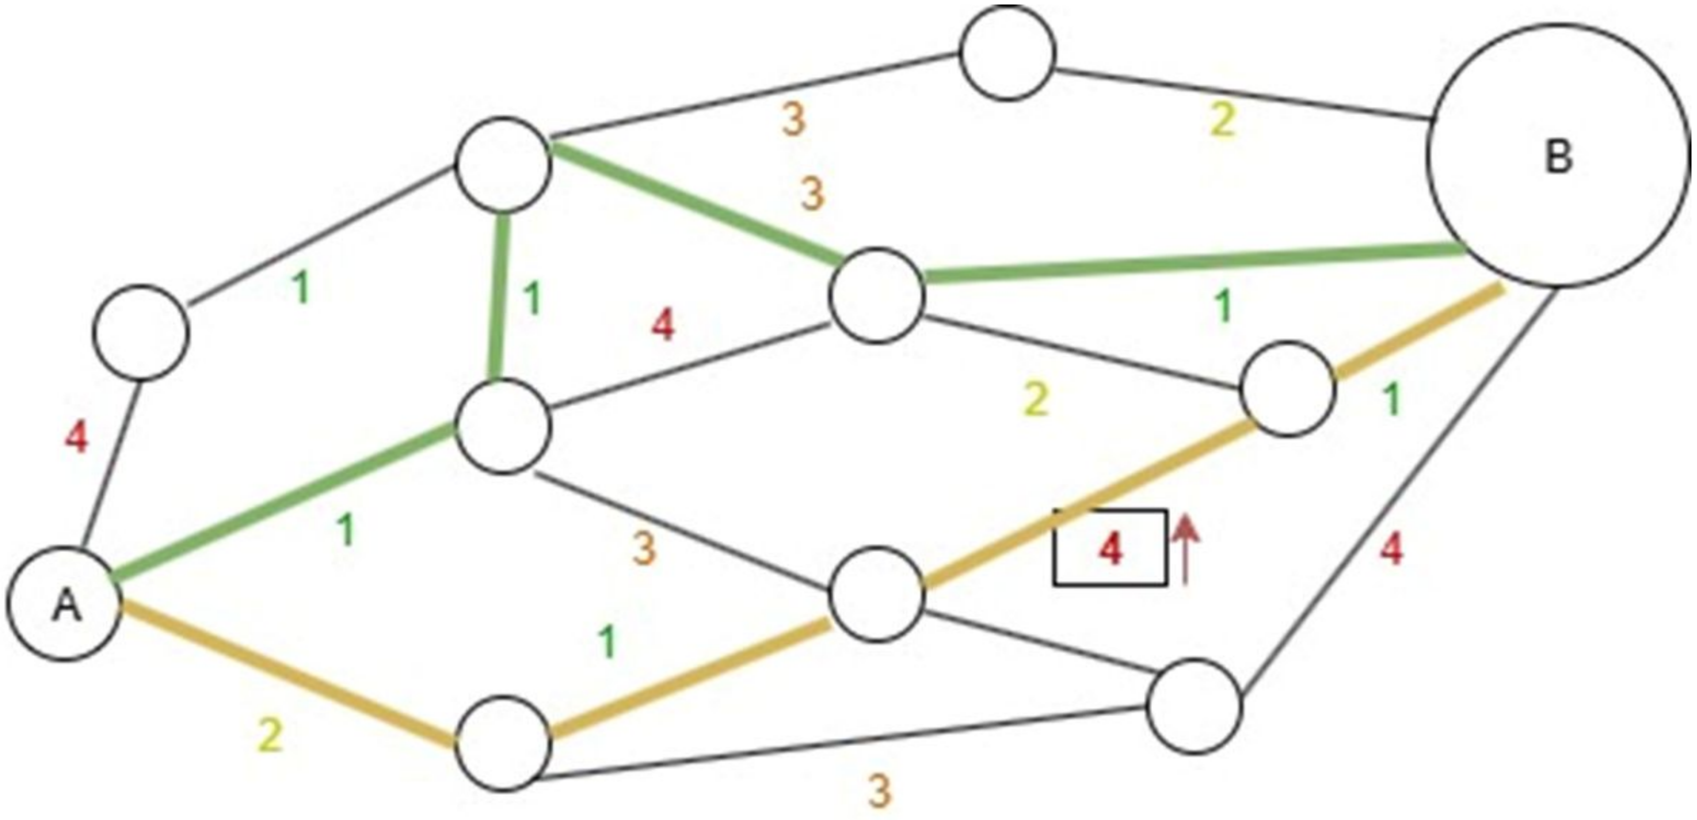

1. Identify Goal Area
2. Identify Nodes “Waypoints”
3. Heuristic Evaluation of Candidate Paths
4. Refine Significant Heuristics

Hierarchical Decision Making

In Abstraction: *Prospecting Agent*

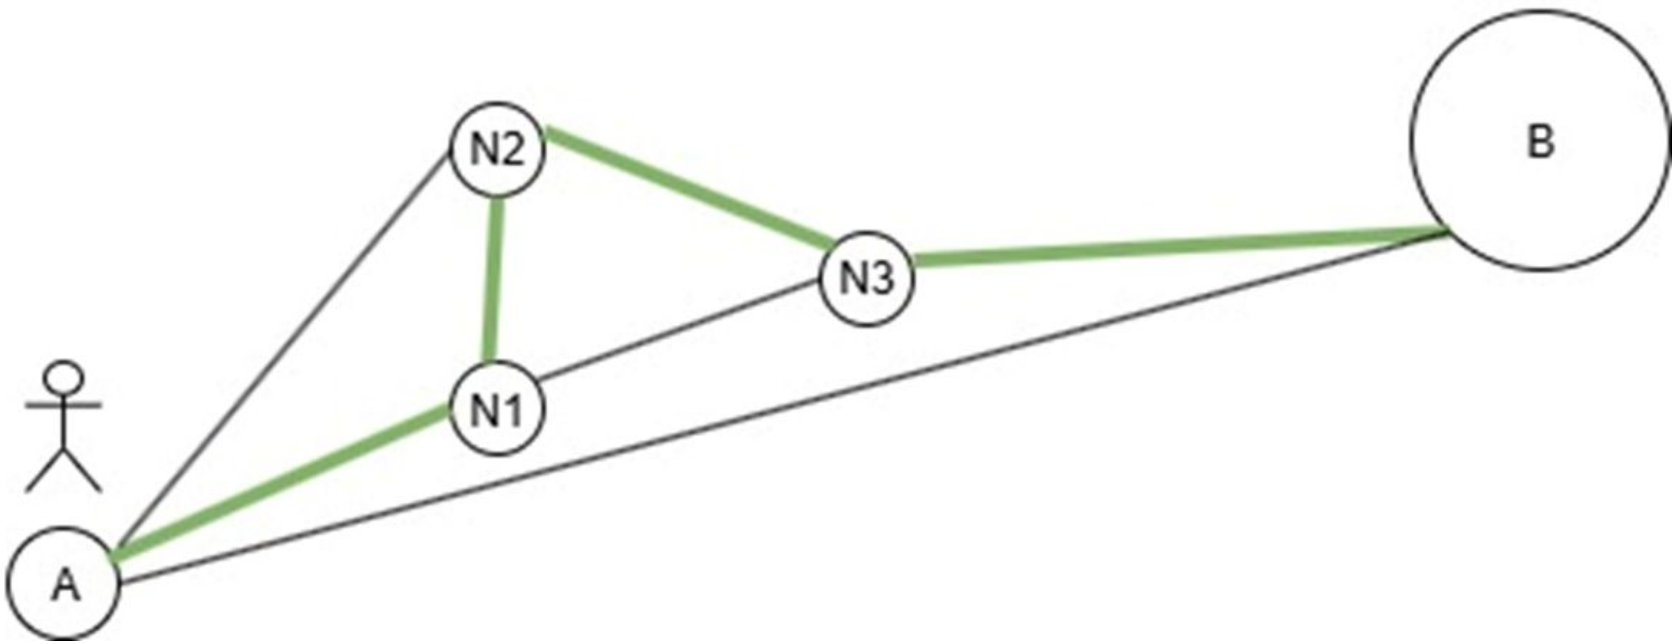

1. Identify Goal Area
2. Identify Nodes “Waypoints”
3. Heuristic Evaluation of Candidate Paths
4. Refine Significant Heuristics
5. Sub-Strategy Selection and Node Permutations

Hierarchical Decision Making

In Abstraction: *Prospecting Agent*

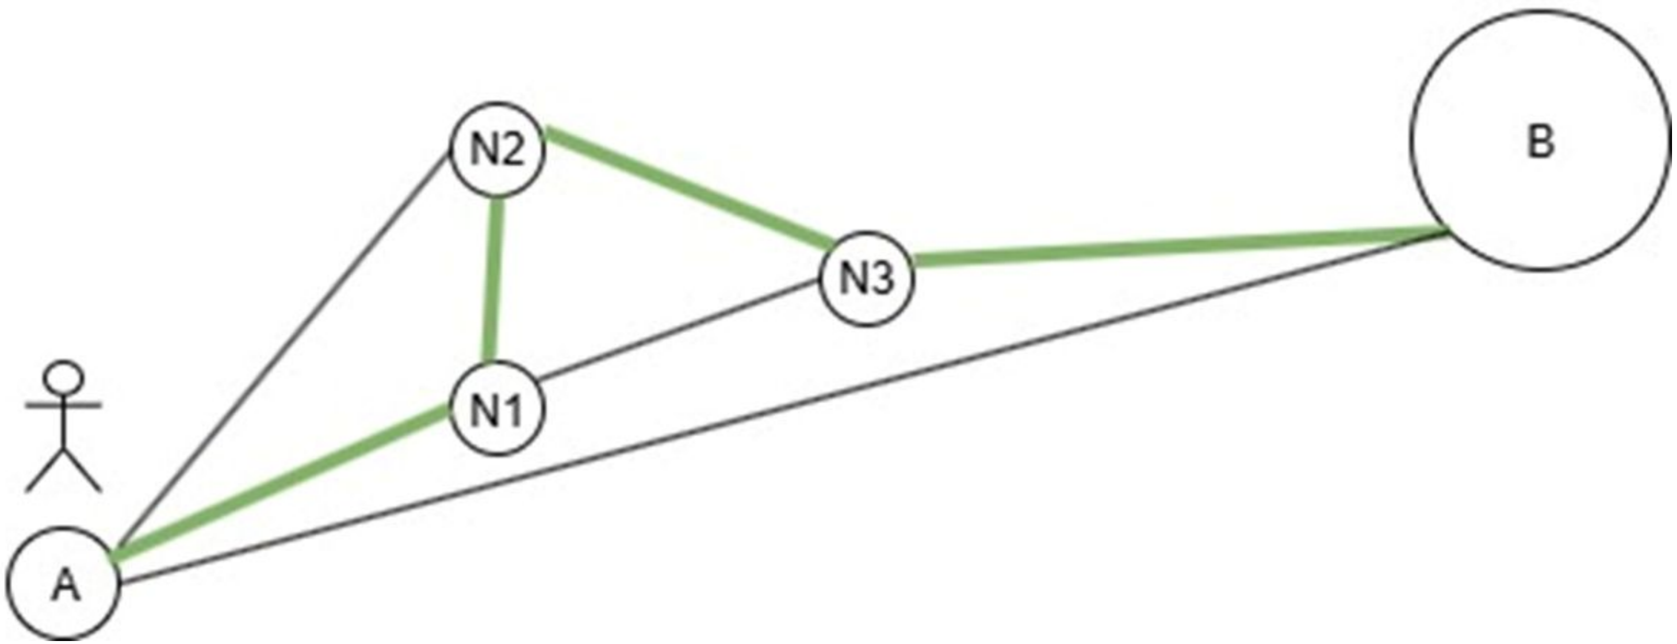

1. Identify Goal Area
2. Identify Nodes “Waypoints”
3. Heuristic Evaluation of Candidate Paths
4. Refine Significant Heuristics
5. Sub-Strategy Selection and Node Permutations
6. Repeat the Process to Find a Near-Optimal Strategy

Lazy Shortest Path Problems for Prospecting Agents

- "LazySP" Framework (**Dellin & Srinivasa, 2016**):
    - Assumes unevaluated edges are valid initially.
    - Computes the shortest path under this assumption.
    - Only evaluates edges on the current best path.
    - *Used In robotic motion planning, evaluating edge costs (e.g., collision checks) is computationally expensive.*

Lazy Shortest Path Problems

- "LazySP" Framework (**Dellin & Srinivasa, 2016**):
  - Assumes unevaluated edges are valid initially.
  - Computes the shortest path under this assumption.
  - Only evaluates edges on the current best path.
  - *Used In robotic motion planning, evaluating edge costs (e.g., collision checks) is computationally expensive.*

- Lazy Shortest Path algorithms also arise as solutions to control problem with L0 Norms

$$\underset{x,u}{\text{minimize}}\, \mathcal{L}(x,u) = \sum_{t=0}^{T} \ell(x_t,u_t) + \gamma\|u_t\|_0, \quad u_t \in [0,1]^d$$

- The L0 loss function promotes sparse decision-making by penalizing the number of actions taken, making it ideal for settings where interventions are costly, limited, or must be minimized.

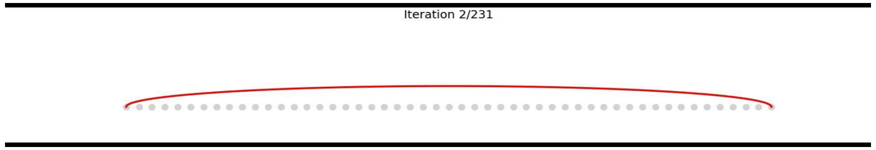

Variations on Norton's Dome

- Norton's Dome is a thought experiment where a perfect sphere rests on top of a perfect dome in an unstable equilibrium
- Its eventual trajectory is not determined by its initial position, creating a classical case of indeterminacy
- Now add that the sphere has a small amount of potential energy internally that allows it to shift its center of gravity by an infinitessimal amount in any direction
- This allows it to control its movement at the equilibrium point

As the sphere gathers speed down the slope:

- Potential energy converts to kinetic energy
- The **agentic energy** (ability to control direction) becomes proportionally less effective
- The **momentum vector** increasingly dominates the sphere's direction

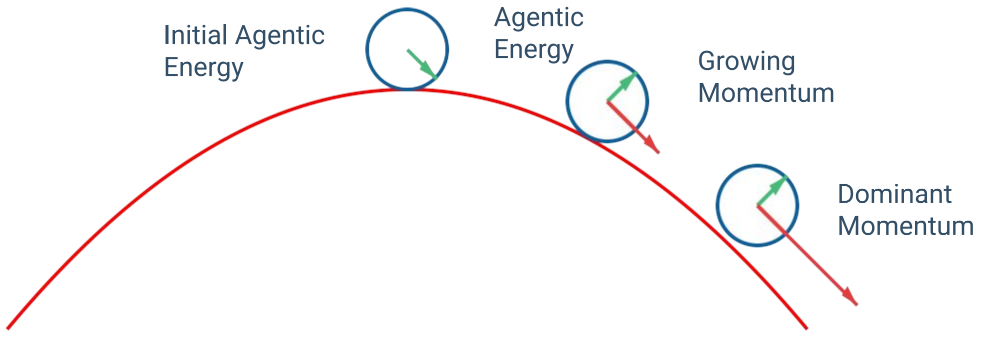

Agentic Energy vs. Growing Momentum

Key Concepts:

- **Unstable Equilibrium**: A state where the slightest perturbation causes the system to leave that state
- → **Agentic Energy**: Small internal force that can influence direction
- → **Momentum**: Mass × velocity; increases as the sphere descends
- ⇄ **Vector Relationship**: As momentum grows, the sphere's ability to change direction diminishes

Oscillating Control Efficacy

Navigating between unstable equilibria in complex landscape

Building on Norton's Dome, we can extend the concept to multiple connected domes, creating a topology with multiple points of unstable equilibrium separated by valleys.

An agent navigating this landscape experiences **oscillating control efficacy** - moments of high and low effectiveness:

- **Peak control** occurs at points of unstable equilibrium where minimal agentic energy can determine trajectory
- **Minimal control** occurs in valleys where momentum dominates movement
- **Oscillation pattern** emerges as the agent moves through the landscape

> **Strategic Navigation:**
>
> By precisely timing interventions at points of high control efficacy, an agent can navigate efficiently between multiple unstable equilibria while conserving energy.

This pattern suggests that intelligent systems might be structured to take advantage of these natural oscillations in control efficacy, concentrating decision-making processes at critical junctures where small inputs yield maximum directional control.

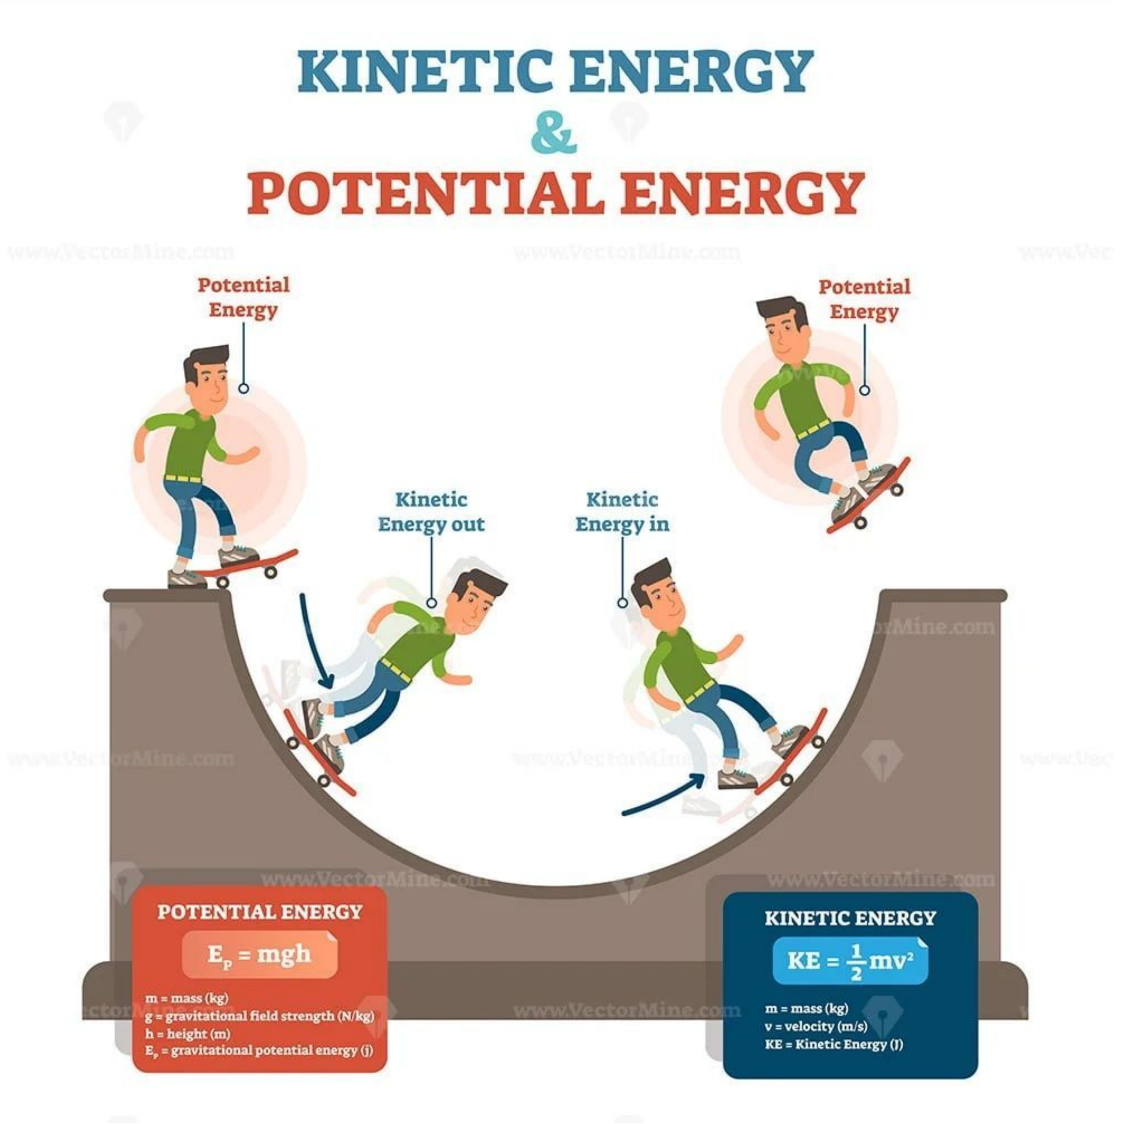

This skateboard half-pipe analogy demonstrates the oscillation between potential energy (at the top of the ramp) and kinetic energy (at the bottom), mirroring how control efficacy oscillates between high and low

Multi-Step Trajectories

Navigating through complex landscapes with oscillating control efficacy

A skateboarder or agentic sphere has high levels of control at the peaks and low levels of control in the valleys. By calibrating interventions to these oscillations in efficacy, an agent can recruit large classical potentials to reach a succession of domes and attain a distant goal.

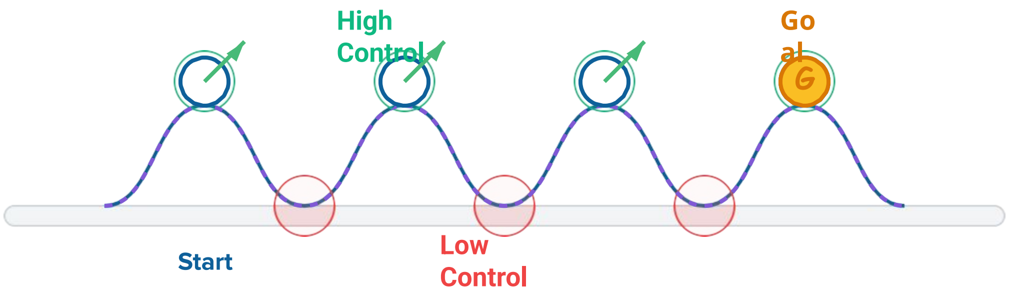

Strategic Intervention Points:

- **Dome Peaks**: Points of unstable equilibrium where minimal force yields maximum directional control
- **Valleys**: Momentum-dominated regions where interventions are less efficient
- **Trajectory Planning**: Success requires calibrating interventions to the natural oscillations in control efficacy

Key Components:

- Agent Sphere
- High Control Zone
- Low Control Zone
- Trajectory Path
- Goal Location
- Calibrated Intervention

This stepwise navigation model suggests how intelligent systems might strategically allocate limited control resources across time, concentrating interventions at critical junctures while allowing natural dynamics to carry the system through intermediate phases.

Such multi-step trajectories demonstrate how complex goal achievement can emerge from simple control principles

Oscillations in Control & Hierarchical Decision Making

Oscillations in hierarchical decision making can arise in both graph based and optimal control based formulations

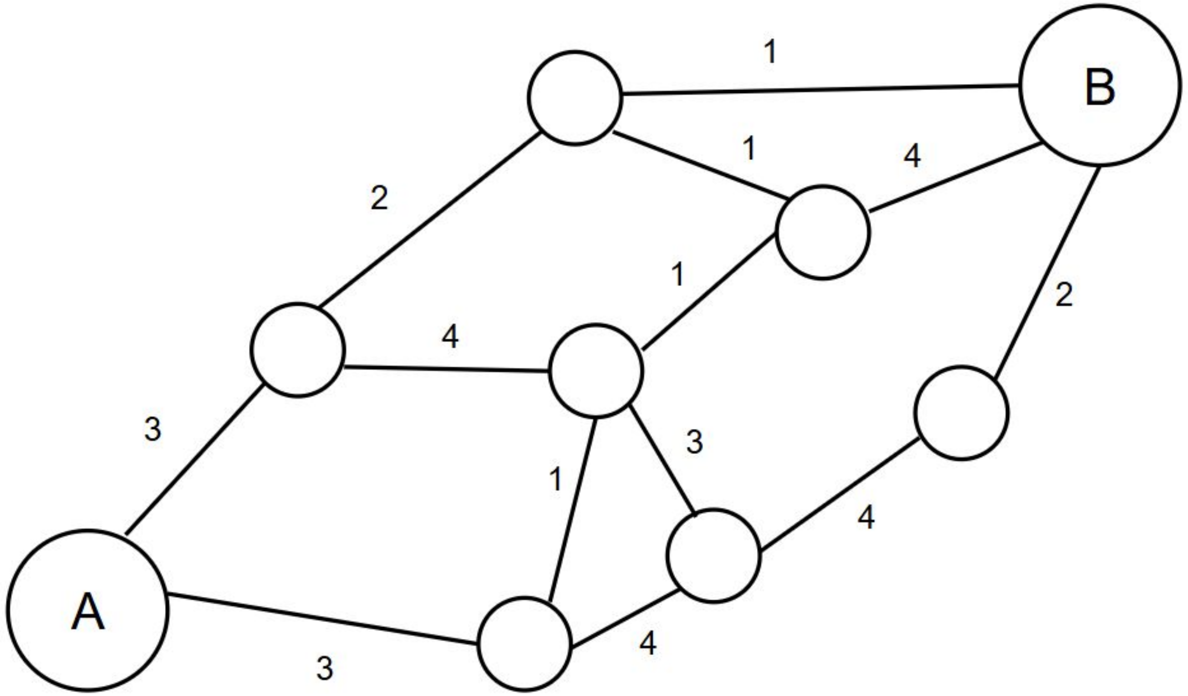

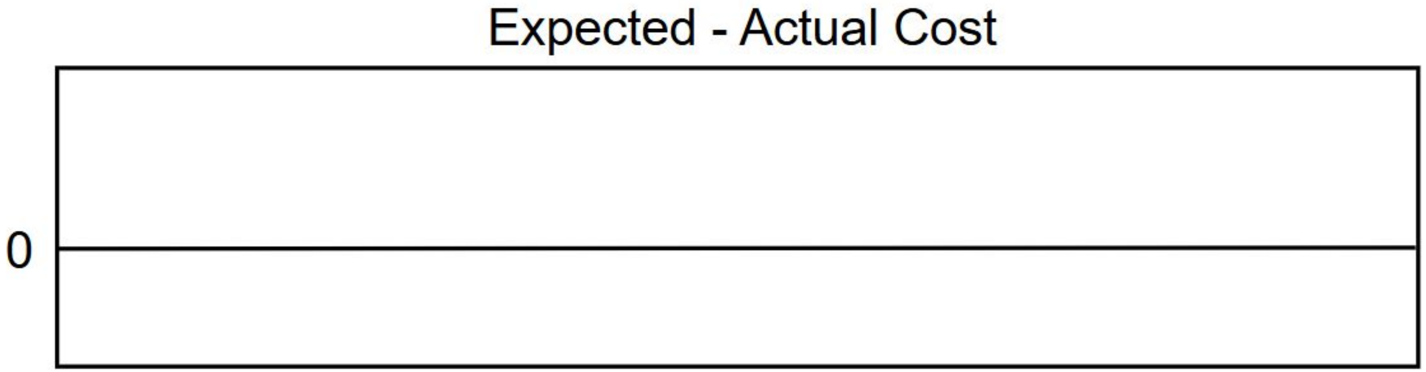

Oscillations in Control & Hierarchical Decision Making

Oscillations in hierarchical decision making can arise in both graph based and optimal control based formulations

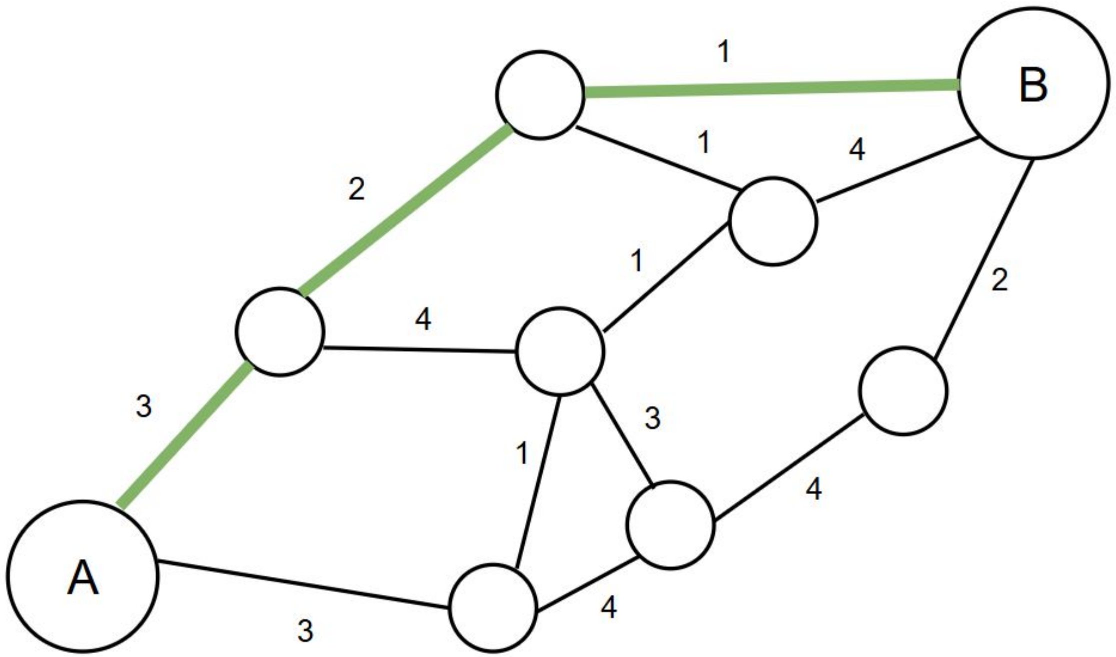

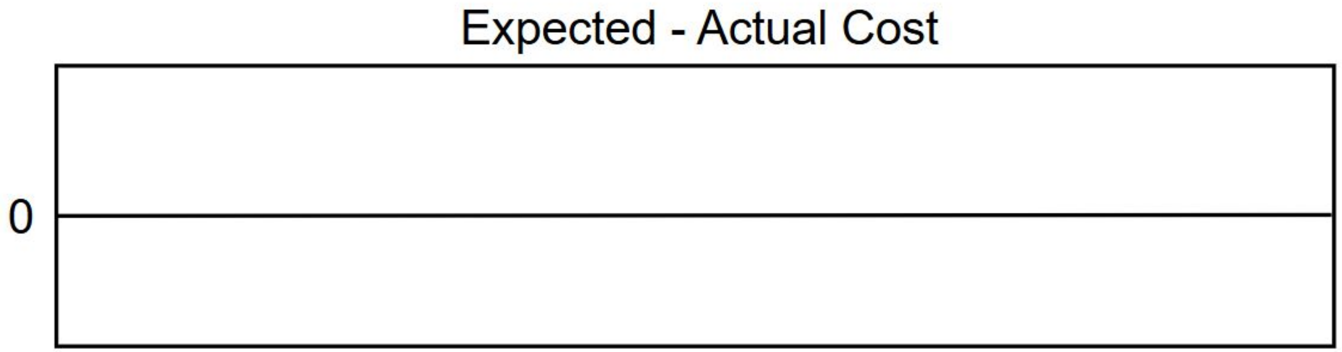

Oscillations in Control & Hierarchical Decision Making

Oscillations in hierarchical decision making can arise in both graph based and optimal control based formulations

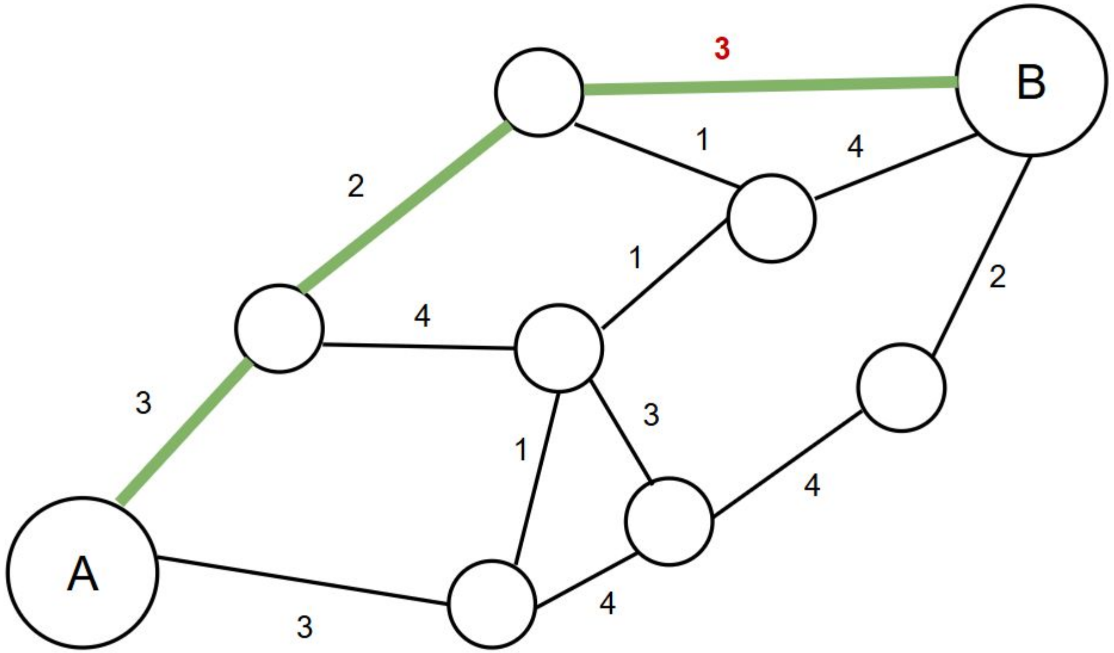

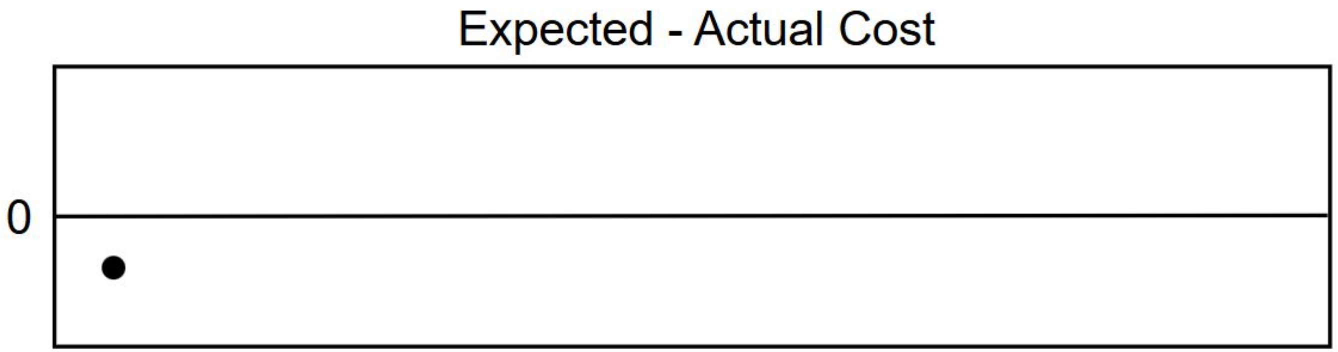

Oscillations in Control & Hierarchical Decision Making

Oscillations in hierarchical decision making can arise in both graph based and optimal control based formulations

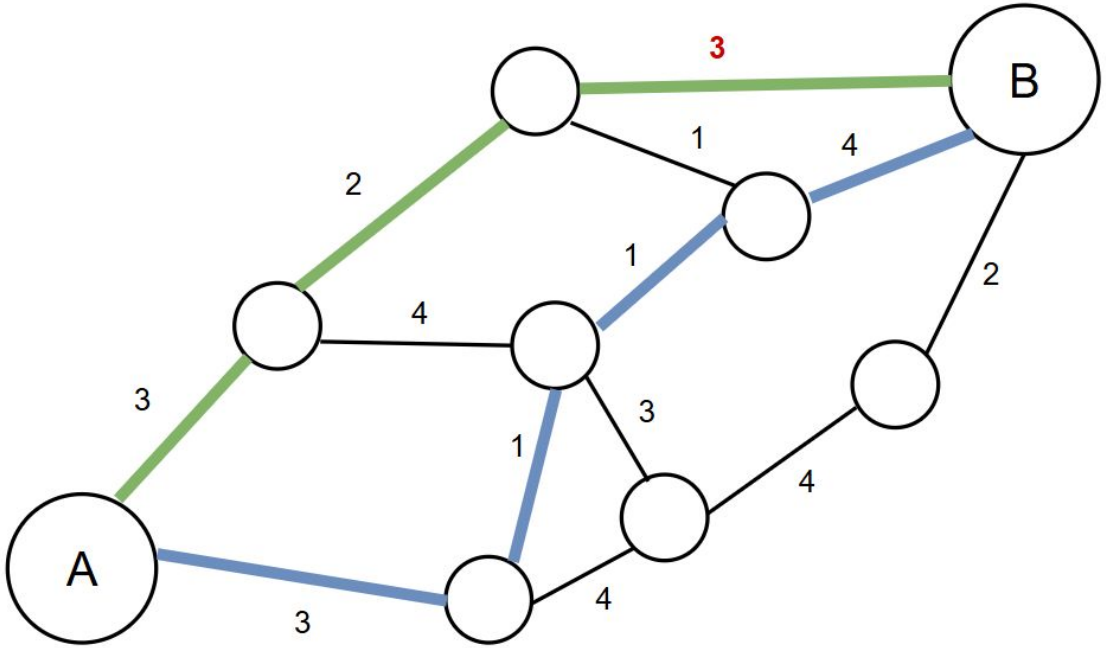

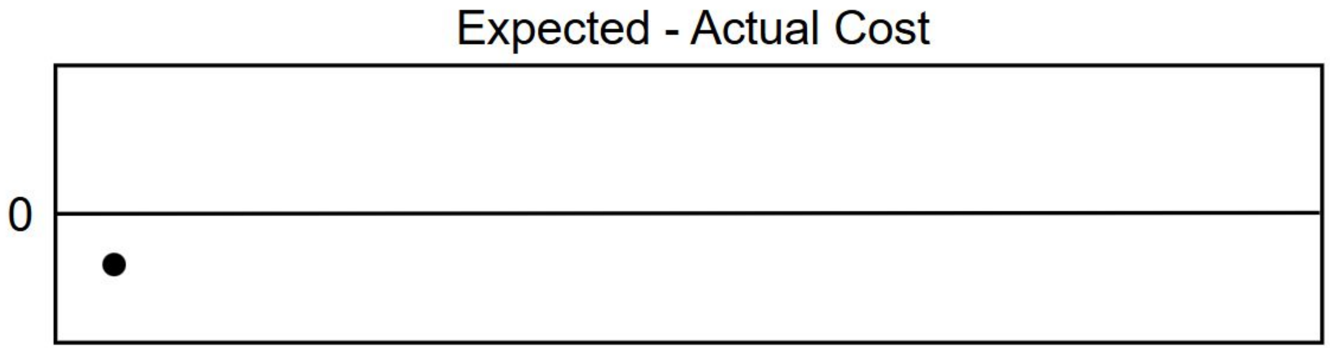

Oscillations in Control & Hierarchical Decision Making

Oscillations in hierarchical decision making can arise in both graph based and optimal control based formulations

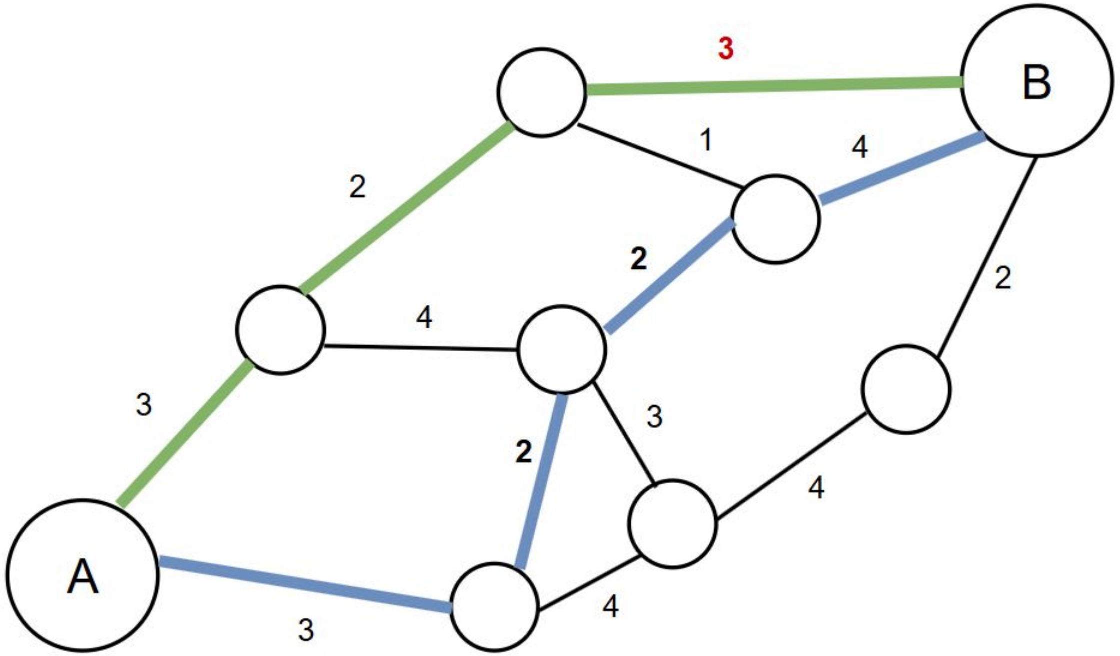

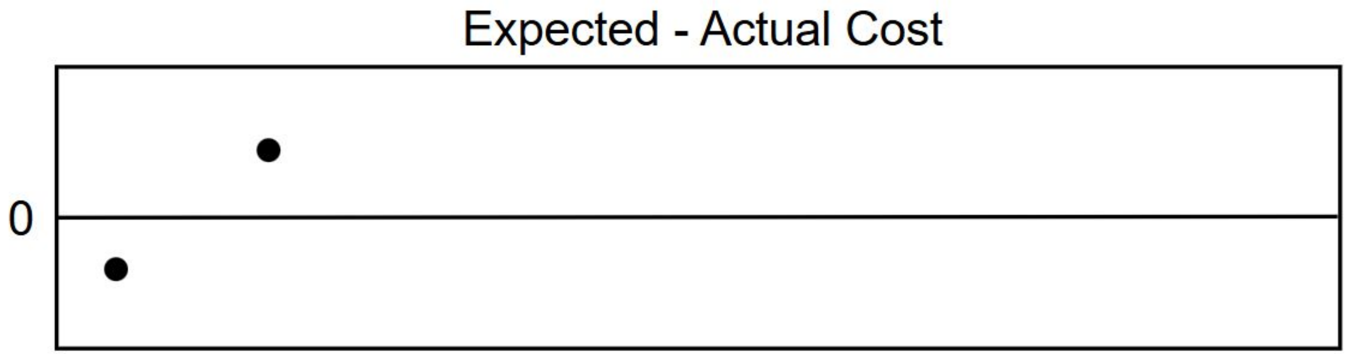

Oscillations in Control & Hierarchical Decision Making

Oscillations in hierarchical decision making can arise in both graph based and optimal control based formulations

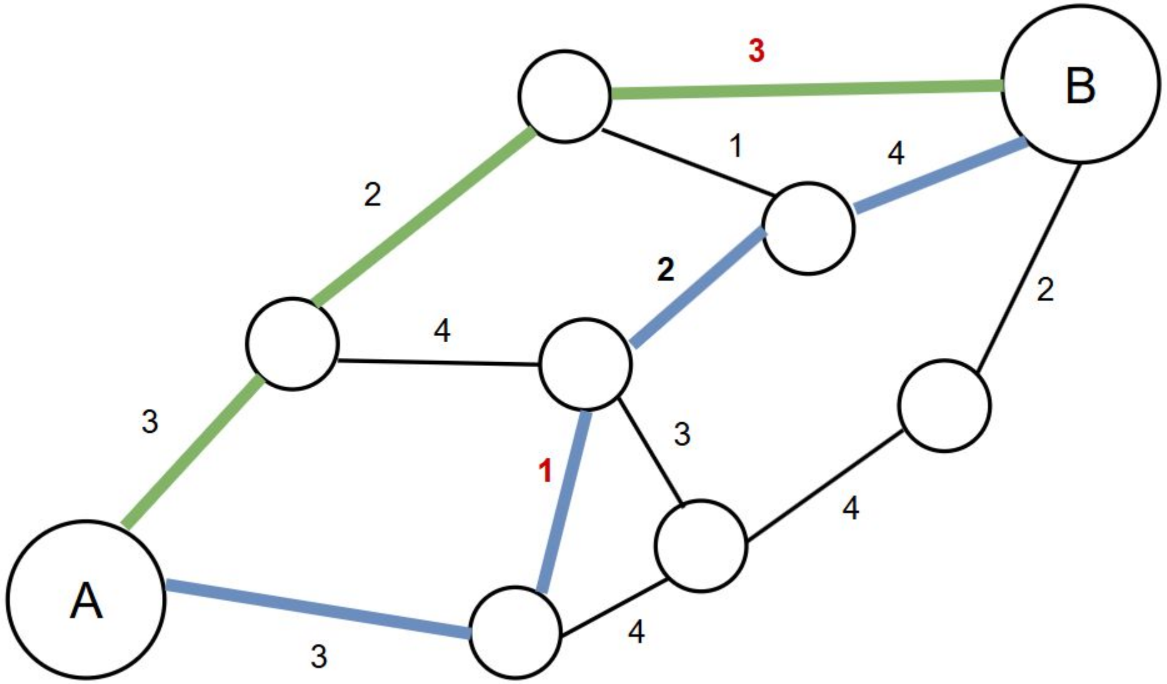

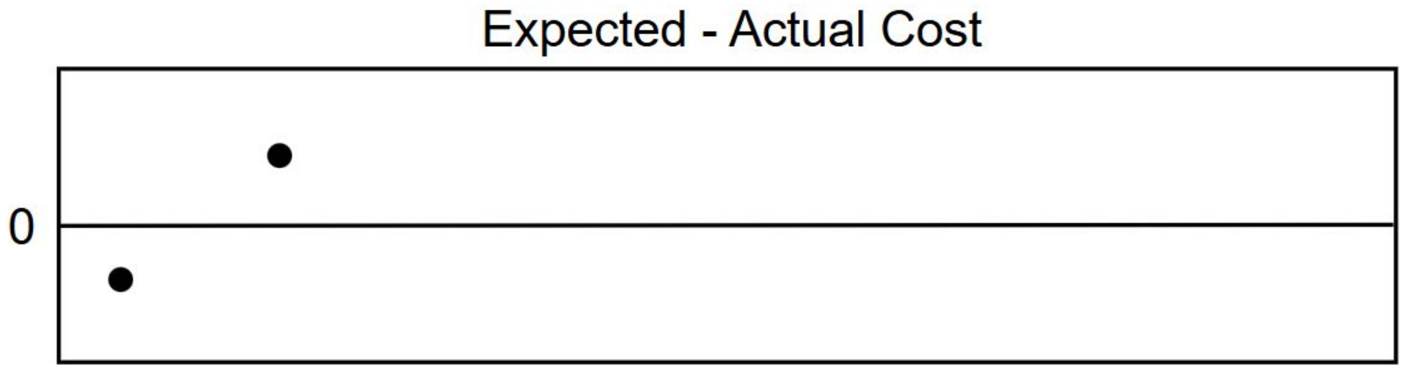

Oscillations in Control & Hierarchical Decision Making

Oscillations in hierarchical decision making can arise in both graph based and optimal control based formulations

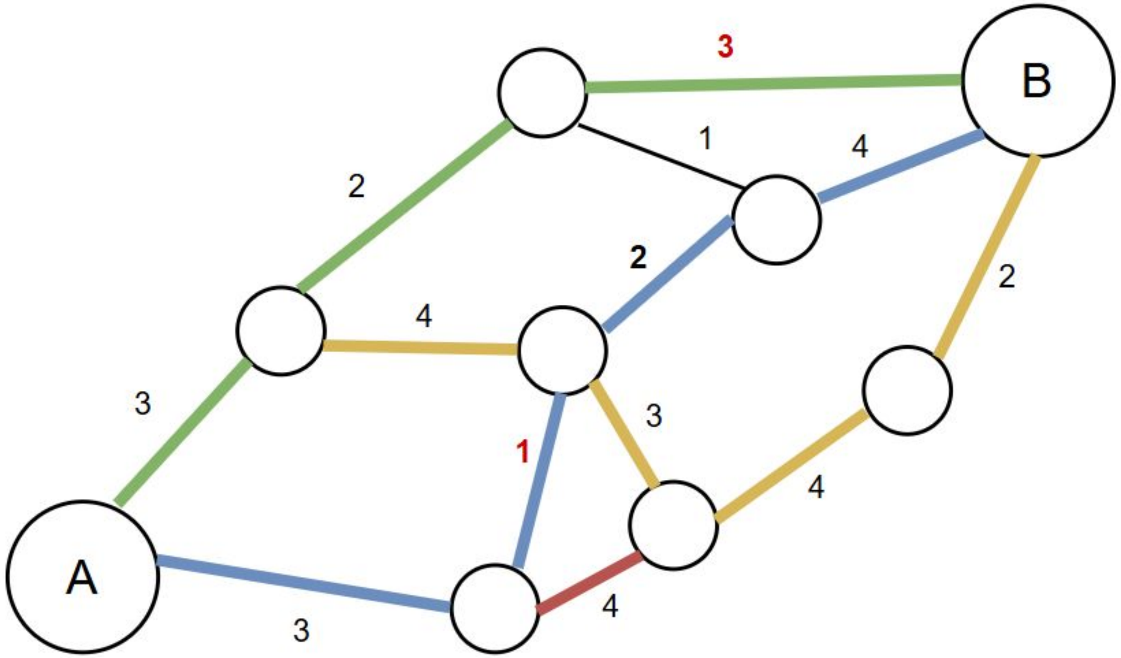

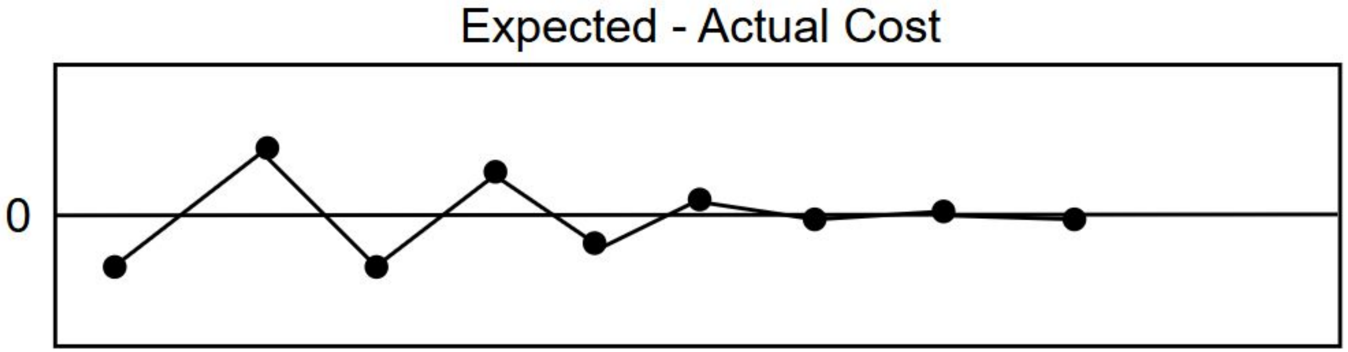

Oscillations in Control & Hierarchical Decision Making

Example 1: Dynamical System of Trust Based Marketplaces

We observe the same oscillatory behavior in decision making in control problems, where perturbing strategies over time may cause stable or unstable oscillatory results

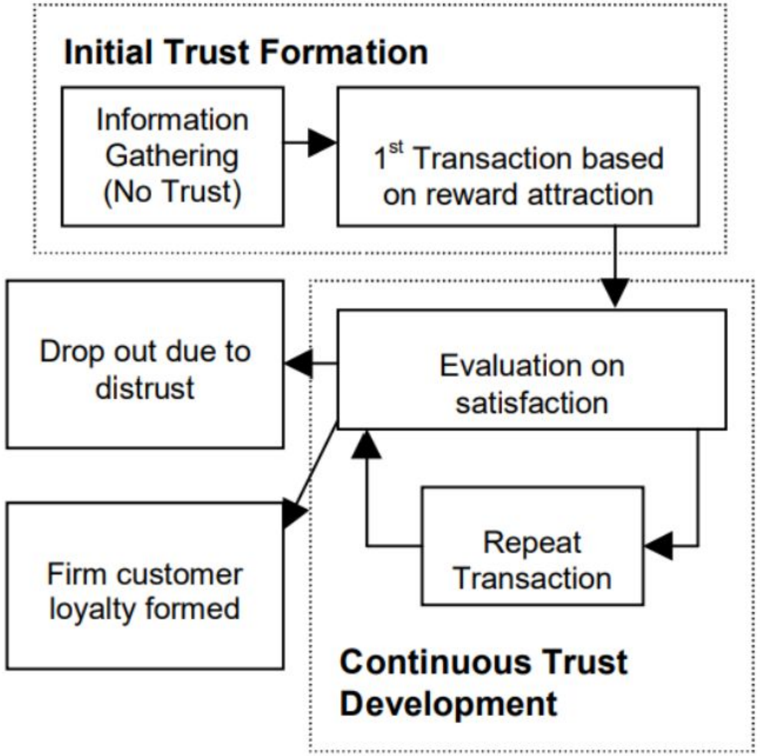

Figure 1: Fung & Lee Trust Development Cycle

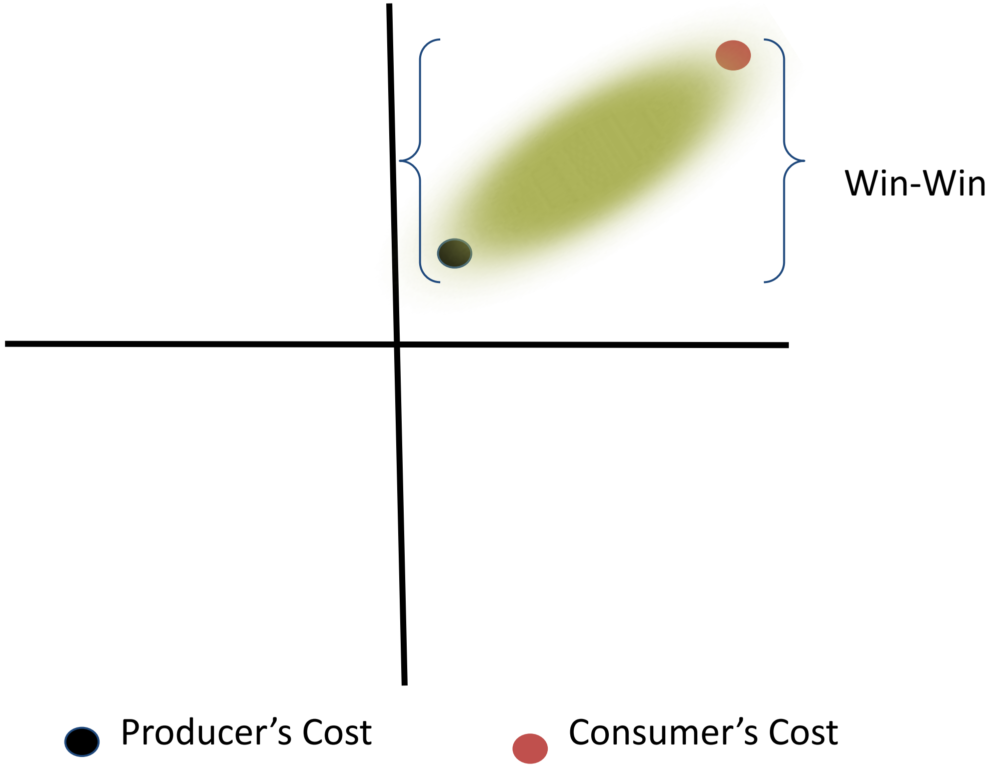

Fung, R. and Lee, M. (1999). EC-Trust (Trust in Electronic Commerce): Exploring the Antecedent Factors. In AMCIS 1999 Proceedings, page 179

Oscillations in Control & Hierarchical Decision Making

Example 1: Dynamical System of Trust Based Marketplaces

We observe the same oscillatory behavior in decision making in control problems, where perturbing strategies over time may cause stable or unstable oscillatory results

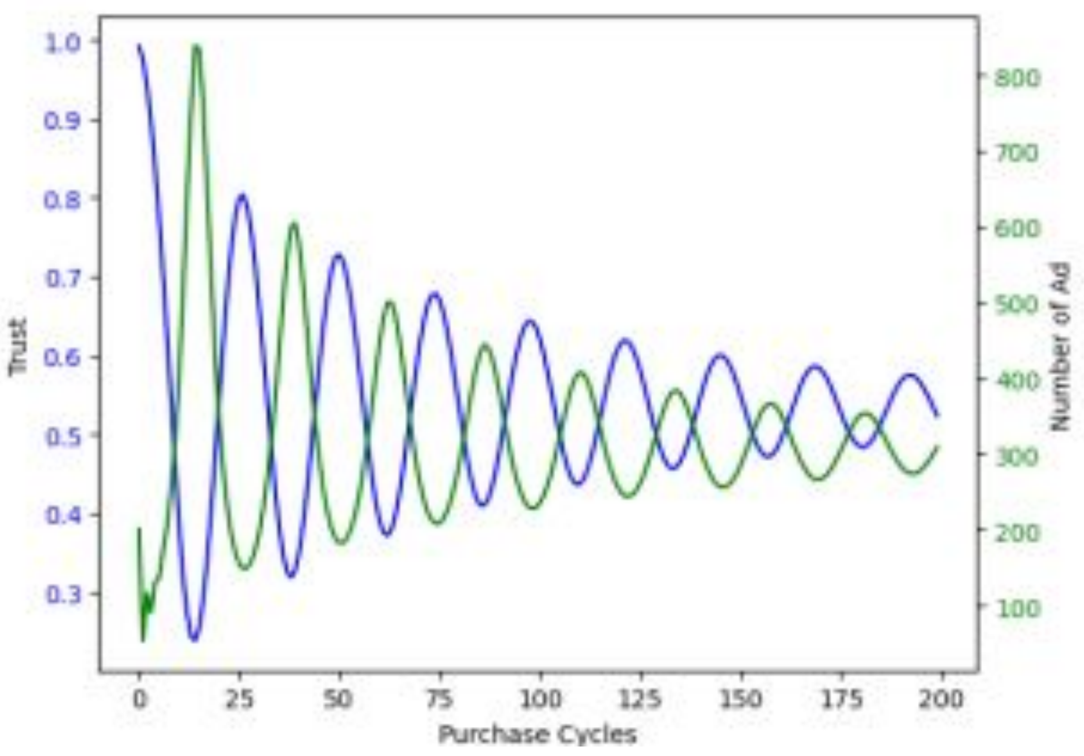

Fung, R. and Lee, M. (1999). EC-Trust (Trust in Electronic Commerce): Exploring the Antecedent Factors. In AMCIS 1999 Proceedings, page 179

Oscillations in Control & Hierarchical Decision Making

Example 2: Commodity Production

We observe the same oscillatory behavior in decision making in control problems, where perturbing strategies over time may cause stable or unstable oscillatory results

$$dP_t = P_t[\eta \log\left(\frac{D_t + \epsilon}{q_t + \epsilon}\right)dt + \sigma_P dW_t^P]$$

$$dD_t = D_{t^-}[(\alpha_D - \beta P_t - \lambda k)dt + \sigma_D dW_t^D + (e^Y - 1)dN_t]$$

$$V(D, P, t) = \sup_q \mathbb{E}^{D,P}\left[\int_0^T e^{\rho(t-T)}(\min\{q_t, D_t\}P_t - C(q_t))\right]$$

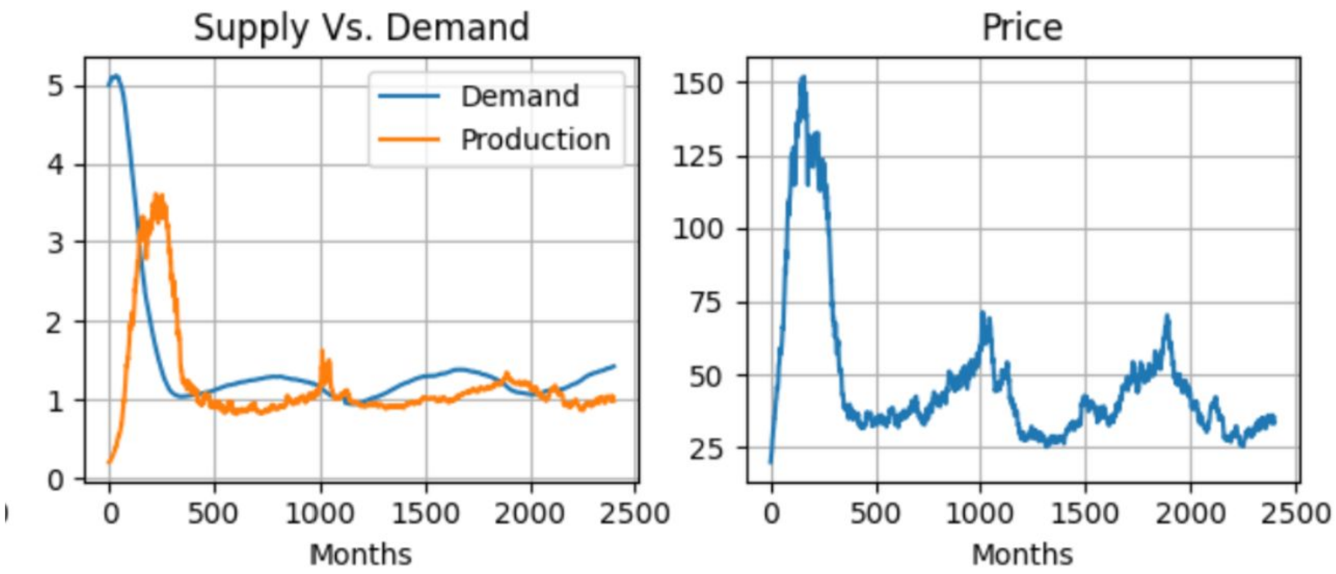

Fung, R. and Lee, M. (1999). EC-Trust (Trust in Electronic Commerce): Exploring the Antecedent Factors. In AMCIS 1999 Proceedings, page 179

Oscillations in Control & Hierarchical Decision Making

Example 2: Commodity Production

We observe the same oscillatory behavior in decision making in control problems, where perturbing strategies over time may cause stable or unstable oscillatory results

$$dP_t = P_t[\eta \log\left(\frac{D_t+\epsilon}{q_t+\epsilon}\right)dt + \sigma_P dW_t^P]$$

$$dD_t = D_{t^-}[(\alpha_D - \beta P_t - \lambda k)dt + \sigma_D dW_t^D + (e^Y - 1)dN_t]$$

$$V(D,P,t) = \sup_q \mathbb{E}^{D,P}\left[\int_0^T e^{\rho(t-T)}(\min\{q_t, D_t\}P_t - C(q_t))\right]$$

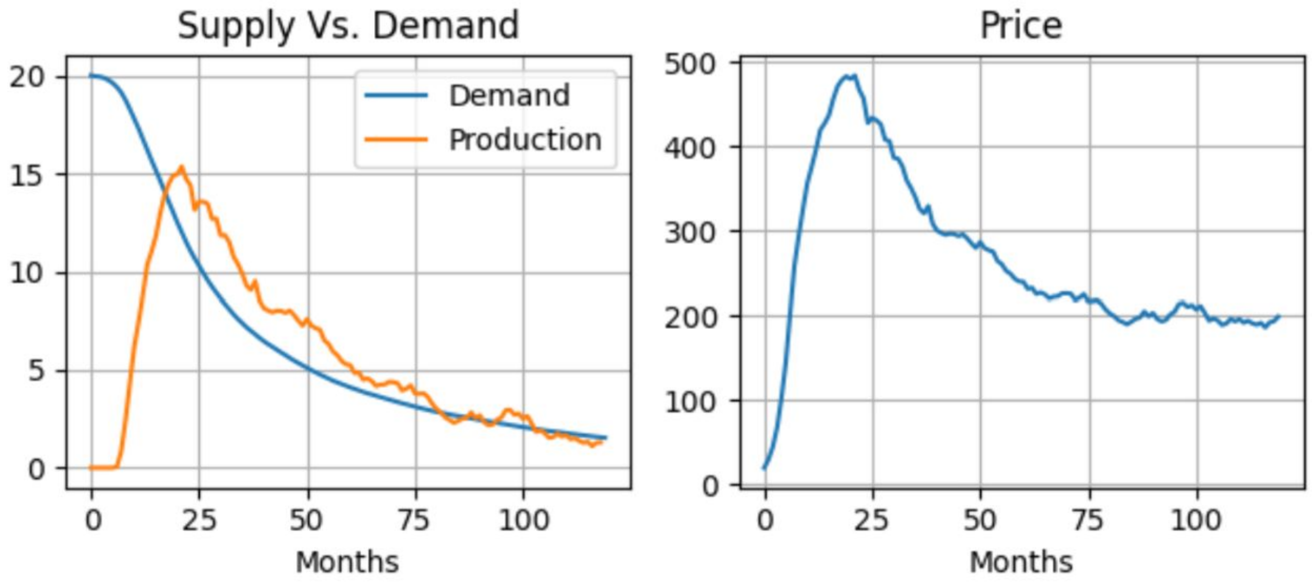

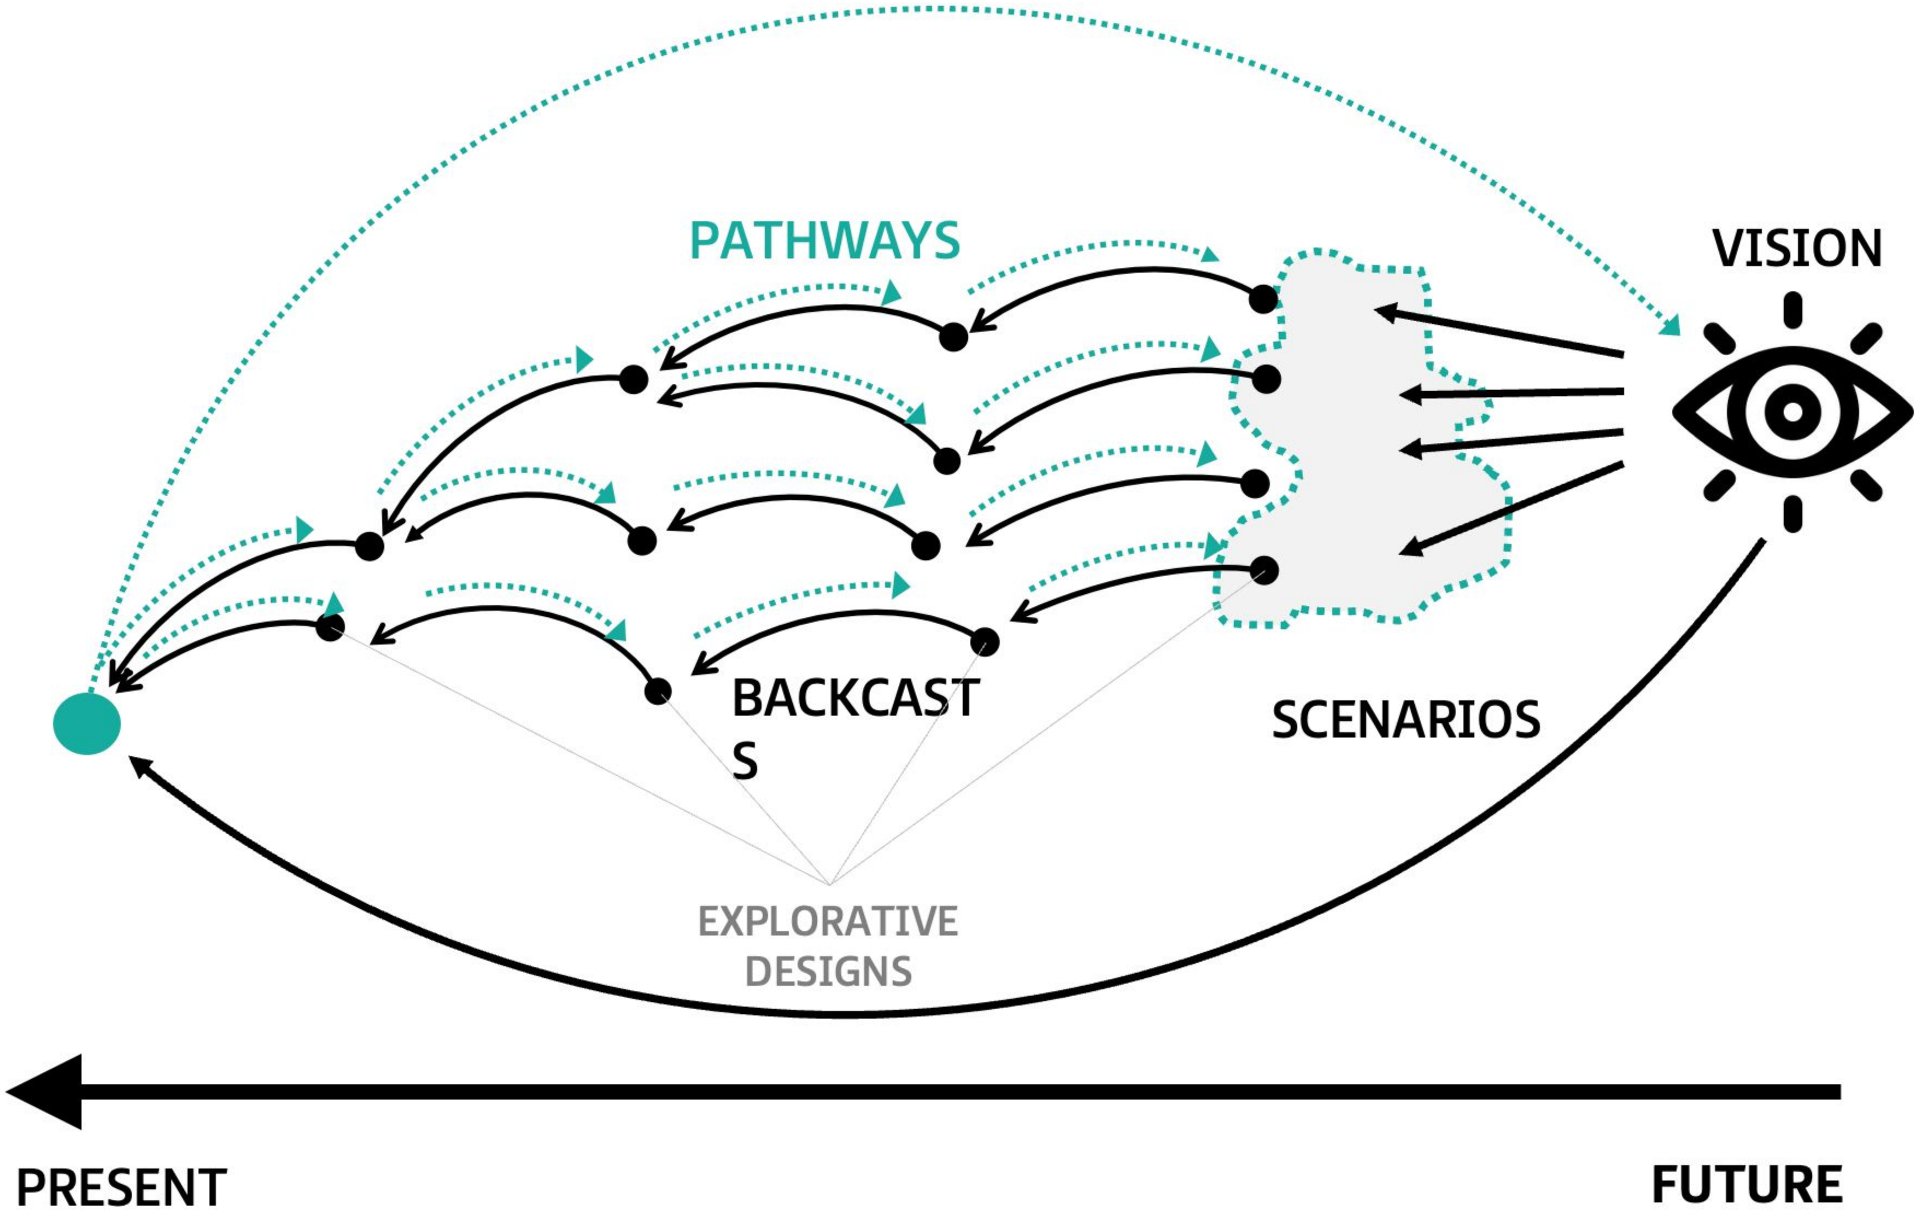

THEORY | PLURALISTIC BACKCASTING

PRACTICE | PLURALISTIC BACKCASTING