"All my relations"
Contributing to a fabric of belonging

David I. Spivak

TOPOS INSTITUTE

AGI 2024
2024 August 15

Outline

All my relations

—Lakota Sioux

How this talk will go

I'll start with philosophy and move into math.

- My motivation is personal: How to have a good life?
- I am deeply insecure: Which way is up?
- Math is solid: it grounds the philosophy as accountable claims.

How this talk will go

I'll start with philosophy and move into math.

- My motivation is personal: How to have a good life?
- I am deeply insecure: Which way is up?
- Math is solid: it grounds the philosophy as accountable claims.

I want to engage you.

- I'm here at AGI-24 because I want help tracking something important.
- I think we are under-developed in our ability to track *care*.
- I asked you to consider all your relations to invoke what's at stake.

How this talk will go

I'll start with philosophy and move into math.

- My motivation is personal: How to have a good life?
- I am deeply insecure: Which way is up?
- Math is solid: it grounds the philosophy as accountable claims.

I want to engage you.

- I'm here at AGI-24 because I want help tracking something important.
- I think we are under-developed in our ability to track *care*.
- I asked you to consider all your relations to invoke what's at stake.

No need to commit to excitement or doom about AGI.

- Both of these are rooted in care.
- So we need an open-ended investigation into care itself.

Care tends to actualize potential.

Care:

- Care words: *matters*, *useful*, *important*, *worry*, *problem*, *want*, *value*.
- Pirsig: what shows up externally as quality occurs internally as care.
- Heidegger: the being of Dasein [roughly, humans] is disclosed as care.

Care tends to actualize potential.

Care:

- Care words: *matters, useful, important, worry, problem, want, value*.
- Pirsig: what shows up externally as quality occurs internally as care.
- Heidegger: the being of Dasein [roughly, humans] is disclosed as care.

Potential:

- Potential is possibility, what could become, what's potent.
- This baby, this genius kid from India, this algorithm, this experience.
- What could be, if only...? This is what we care about.

Care tends to actualize potential.

Care:

- Care words: *matters, useful, important, worry, problem, want, value*.
- Pirsig: what shows up externally as quality occurs internally as care.
- Heidegger: the being of Dasein [roughly, humans] is disclosed as care.

Potential:

- Potential is possibility, what could become, what's potent.
- This baby, this genius kid from India, this algorithm, this experience.
- What could be, if only...? This is what we care about.

Actualizing:

- The actual is what is here. Actions speak louder than words.
- Self actualization, becoming all you can be, reaching your potential.

Care tends to actualize potential.

Care:

- Care words: *matters, useful, important, worry, problem, want, value*.
- Pirsig: what shows up externally as quality occurs internally as care.
- Heidegger: the being of Dasein [roughly, humans] is disclosed as care.

Potential:

- Potential is possibility, what could become, what's potent.
- This baby, this genius kid from India, this algorithm, this experience.
- What could be, if only...? This is what we care about.

Actualizing:

- The actual is what is here. Actions speak louder than words.
- Self actualization, becoming all you can be, reaching your potential.

To tend:

- Empirically, how do we *tend* to accomplish stuff? We care.
- Care works by tending, as to a garden. Stay with it, concern yourself.
- To care is to attend, to pay attention.

Why am I here?

Here's the problem I've cared about, which has led me to be here with you.

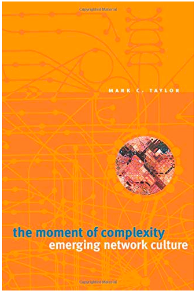

- In 2007 I read *The moment of complexity*.
  - "In ∼1993 the world's brain came online."
  - "Consequence: more and more will be different."
  - It presaged fake news, anomalies becoming the norm.
- I wanted to help, and saw an opportunity.

Why am I here?

Here's the problem I've cared about, which has led me to be here with you.

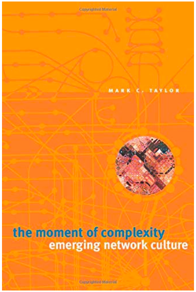

- In 2007 I read *The moment of complexity*.
  - "In ∼1993 the world's brain came online."
  - "Consequence: more and more will be different."
  - It presaged fake news, anomalies becoming the norm.
- I wanted to help, and saw an opportunity.
  - My math PhD relied heavily on Category Theory (CT)
  - It organizes perspectives and translation systems.
  - Maybe it could help us navigate the complexity.

Why am I here?

Here's the problem I've cared about, which has led me to be here with you.

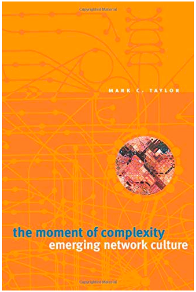

- In 2007 I read *The moment of complexity*.
  - "In ∼1993 the world's brain came online."
  - "Consequence: more and more will be different."
  - It presaged fake news, anomalies becoming the norm.
- I wanted to help, and saw an opportunity.
  - My math PhD relied heavily on Category Theory (CT)
  - It organizes perspectives and translation systems.
  - Maybe it could help us navigate the complexity.
- I've spent the last 17 years working on this.

Driving question: *how can CT help humanity make better collective sense?*

Collective sensemaking

"All intelligence is collective intelligence" –Mike Levin.

- Humans communicate, make sense together, learn from each other.
- But each human is itself a collection of cells acting collectively.
- Intelligence and our sense of things arise out of coordinated activity.

Collective sensemaking

"All intelligence is collective intelligence" –Mike Levin.

- Humans communicate, make sense together, learn from each other.
- But each human is itself a collection of cells acting collectively.
- Intelligence and our sense of things arise out of coordinated activity.

If tweets, texts, and tiktoks are like neurons firing, is the world thinking?

- Videos, blogs, etc. explaining the world's situation in real time.
- Let's call this activity *cultural processing*.
- Cultural processing is improving and getting faster.
- But this is causing intense instabilities that counter the improvements.
- We're going through a phase-transition, a metamorphosis.

Collective sensemaking

"All intelligence is collective intelligence" –Mike Levin.

- Humans communicate, make sense together, learn from each other.
- But each human is itself a collection of cells acting collectively.
- Intelligence and our sense of things arise out of coordinated activity.

If tweets, texts, and tiktoks are like neurons firing, is the world thinking?

- Videos, blogs, etc. explaining the world's situation in real time.
- Let's call this activity *cultural processing*.
- Cultural processing is improving and getting faster.
- But this is causing intense instabilities that counter the improvements.
- We're going through a phase-transition, a metamorphosis.

We need something to help us take care of ourselves.

- We don't know which way is up or how to get right with ourselves.
- To survive and flourish, we must tend to the transition carefully.
- Where can we turn to help us organize and balance the coll'tive sense?

Collective sensemaking

"All intelligence is collective intelligence" –Mike Levin.

- Humans communicate, make sense together, learn from each other.
- But each human is itself a collection of cells acting collectively.
- Intelligence and our sense of things arise out of coordinated activity.

If tweets, texts, and tiktoks are like neurons firing, is the world thinking?

- Videos, blogs, etc. explaining the world's situation in real time.
- Let's call this activity *cultural processing*.
- Cultural processing is improving and getting faster.
- But this is causing intense instabilities that counter the improvements.
- We're going through a phase-transition, a metamorphosis.

We need something to help us take care of ourselves.

- We don't know which way is up or how to get right with ourselves.
- To survive and flourish, we must tend to the transition carefully.
- Where can we turn to help us organize and balance the coll'tive sense?

Math helps organize thinking; can it be of use?

Maybe it's all stored in our heads

How can we imagine the problem of collective sensemaking? A first stab:

- The most import't thing is communication, understanding each other.
- Imagine everyone has a database in their head, organizing everything:
- ...a schematic layout of how things fit together, and tons of examples.

$\mathscr{D} =$ 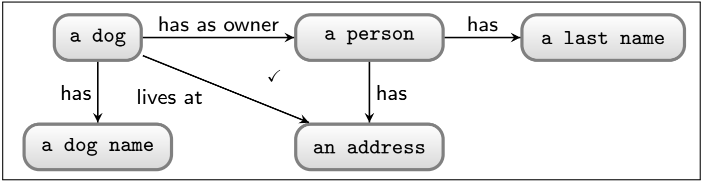

dog			
ID	**name**	**owner**	**address**
D101	Wally	P34	15 Ash St.
D102	Fido	P46	201 5th Ave.
D104	Buster	P17	27 Spring Ln.

person		
ID	**lastName**	**address**
P17	Jones	27 Spring Ln.
P34	Smith	15 Ash St.
P46	D'Angelo	201 5th Ave.

Maybe it's all stored in our heads

How can we imagine the problem of collective sensemaking? A first stab:

- The most import't thing is communication, understanding each other.
- Imagine everyone has a database in their head, organizing everything:
- ...a schematic layout of how things fit together, and tons of examples.

$\mathscr{D} =$

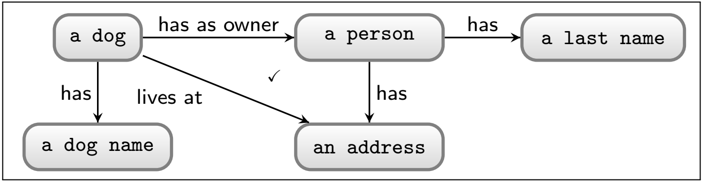

dog			
ID	**name**	**owner**	**address**
D101	Wally	P34	15 Ash St.
D102	Fido	P46	201 5th Ave.
D104	Buster	P17	27 Spring Ln.

person		
ID	**lastName**	**address**
P17	Jones	27 Spring Ln.
P34	Smith	15 Ash St.
P46	D'Angelo	201 5th Ave.

- Our brains are very different and organize the world differently.
- Let's make math about that: how diff'nt databases can communicate.

Maybe it's all dynamics

But databases are static, whereas the world is dynamic. Second stab:

- Each thing is a dynamical system, taking inputs & producing outputs.
- Lots of little dynamic systems interact to form a larger one.

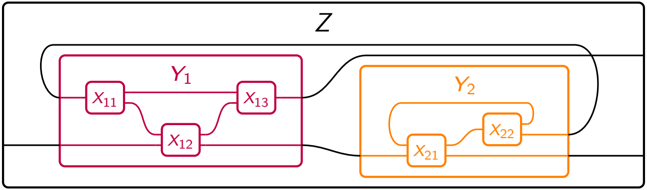

Maybe it's all dynamics

But databases are static, whereas the world is dynamic. Second stab:

- Each thing is a dynamical system, taking inputs & producing outputs.
- Lots of little dynamic systems interact to form a larger one.

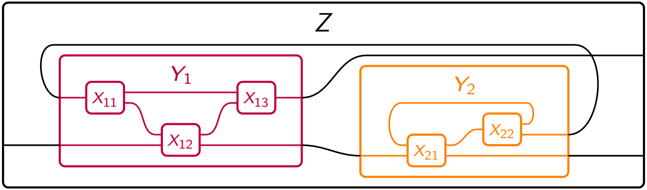

We could consider the math of interacting dynamical systems.

- Is sensemaking just some algorithm that these systems are running?
- Perhaps the flowing information causes the wiring diagram to change?

Maybe it's all dynamics

But databases are static, whereas the world is dynamic. Second stab:

- Each thing is a dynamical system, taking inputs & producing outputs.
- Lots of little dynamic systems interact to form a larger one.

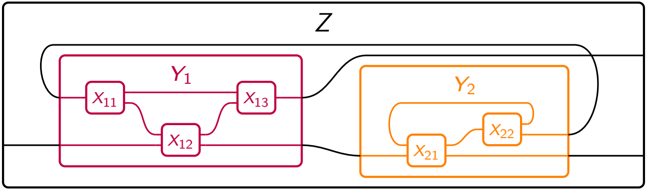

We could consider the math of interacting dynamical systems.

- Is sensemaking just some algorithm that these systems are running?
- Perhaps the flowing information causes the wiring diagram to change?

There's beautiful math for all this. But it's still missing some'ng *important.*

Plan for the talk

In the remainder of the talk, I'll discuss:

- Motivation for sensemaking
- Actual math with potential to help
- Where I think we need to go from here.

Outline

The computer

In the 1930s a *computer* was a person who calculated professionally.

- Often women, they carefully performed meticulous calculations.
- Plotting trajectories, e.g. of missiles, required making very few errors.

The computer

In the 1930s a *computer* was a person who calculated professionally.

- Often women, they carefully performed meticulous calculations.
- Plotting trajectories, e.g. of missiles, required making very few errors.

Turing extracted the mechanical essence to define his *automatic machine.*

- Instructions for reading/writing symbols on a limitless supply of paper.
- This is what "computer" has become, but it misses a key element.

The computer

In the 1930s a *computer* was a person who calculated professionally.

- Often women, they carefully performed meticulous calculations.
- Plotting trajectories, e.g. of missiles, required making very few errors.

Turing extracted the mechanical essence to define his *automatic machine.*

- Instructions for reading/writing symbols on a limitless supply of paper.
- This is what "computer" has become, but it misses a key element.

What made the women do such a good job? Care.

- They needed care to do the high-quality work.
- Whether family, mission, interest, etc., they brought care to their work.
- Turing skipped this, like Shannon skipped "meaning" in defining info.

The computer

In the 1930s a *computer* was a person who calculated professionally.

- Often women, they carefully performed meticulous calculations.
- Plotting trajectories, e.g. of missiles, required making very few errors.

Turing extracted the mechanical essence to define his *automatic machine.*

- Instructions for reading/writing symbols on a limitless supply of paper.
- This is what "computer" has become, but it misses a key element.

What made the women do such a good job? Care.

- They needed care to do the high-quality work.
- Whether family, mission, interest, etc., they brought care to their work.
- Turing skipped this, like Shannon skipped "meaning" in defining info.

Since then we've scaled up the uncaring mechanism by ~ 20 OOMs.

- But it's still the living who care. CPUs and LLMs don't *care*.
- This essential piece—the interest, the drive—is going un-tracked.

We've lost track of what matters most. Let's consider how it works.

What is sense?

Here's the power formula in physics: $P = IV$ (current times voltage)

- Let's think of *intelligence* as current: how much flows per unit time.
- Let's think of *care* as voltage: it's the potential difference.
- Without the other, intelligence is dormant and care is impotent.

What is sense?

Here's the power formula in physics: $P = IV$ (current times voltage)

- Let's think of *intelligence* as current: how much flows per unit time.
- Let's think of *care* as voltage: it's the potential difference.
- Without the other, intelligence is dormant and care is impotent.

Care tends to actualize potential by making sense of things.

- By *sense*, I don't mean raw perception.
- A spidey sense, a sense of danger or belonging, a sense of direction.
- These are not mere perception. They're tracking the right variables.
- When we have a good sense of something, we navigate it with ease.

What is sense?

Here's the power formula in physics: $P = IV$ (current times voltage)

- Let's think of *intelligence* as current: how much flows per unit time.
- Let's think of *care* as voltage: it's the potential difference.
- Without the other, intelligence is dormant and care is impotent.

Care tends to actualize potential by making sense of things.

- By *sense*, I don't mean raw perception.
- A spidey sense, a sense of danger or belonging, a sense of direction.
- These are not mere perception. They're tracking the right variables.
- When we have a good sense of something, we navigate it with ease.

Consider a snapshot of two math students, both wanting to succeed:

- Student A is faithfully copies down what the teacher says.
- Student B seems to be doing the opposite: ...
- ...clearly frustrated, arguing with the teacher, "but then why XYZ??"
- Suddenly student B says "Oh!! Is it because ABC??"
- B relaxes, having made sense; A may feel distant and stop caring.

What is sense?

Here's the power formula in physics: $P = IV$ (current times voltage)

- Let's think of *intelligence* as current: how much flows per unit time.
- Let's think of *care* as voltage: it's the potential difference.
- Without the other, intelligence is dormant and care is impotent.

Care tends to actualize potential by making sense of things.

- By *sense*, I don't mean raw perception.
- A spidey sense, a sense of danger or belonging, a sense of direction.
- These are not mere perception. They're tracking the right variables.
- When we have a good sense of something, we navigate it with ease.

Consider a snapshot of two math students, both wanting to succeed:

- Student A is faithfully copies down what the teacher says.
- Student B seems to be doing the opposite: ...
- ...clearly frustrated, arguing with the teacher, "but then why XYZ??"
- Suddenly student B says "Oh!! Is it because ABC??"
- B relaxes, having made sense; A may feel distant and stop caring.

Making sense of things takes work, but it produces sense!

Tracking what we care about

What do we mean by *tracking*?

- Tracking criminals, tracking bears: we care about finding them.
- Tracking trains, keeping them on rails. Traction of tires on road.
- Tracking someone's logic, not getting distracted, getting back on track.
- To track something is to follow it, to stay with it as it moves.

Tracking what we care about

What do we mean by *tracking*?

- Tracking criminals, tracking bears: we care about finding them.
- Tracking trains, keeping them on rails. Traction of tires on road.
- Tracking someone's logic, not getting distracted, getting back on track.
- To track something is to follow it, to stay with it as it moves.

We make sense so we can track what we care about.

- In order to track something, we need a sense of how it operates.
- To care effectively for something requires us to have a sense of it.
- So what is this all-important sensemaking?

Tracking what we care about

What do we mean by *tracking*?

- Tracking criminals, tracking bears: we care about finding them.
- Tracking trains, keeping them on rails. Traction of tires on road.
- Tracking someone's logic, not getting distracted, getting back on track.
- To track something is to follow it, to stay with it as it moves.

We make sense so we can track what we care about.

- In order to track something, we need a sense of how it operates.
- To care effectively for something requires us to have a sense of it.
- So what is this all-important sensemaking?

I believe that sensemaking may come down to accounting.

- Student B was trying to figure out how to account for XYZ.
- We seek an abstraction with which to overlay our experience.
- Once we find it, we can "put things in their proper place".

Sensemaking may be finding systematic ways of accounting for experience.

Abstraction: extraction and application

Abstracting is tracking the structural essence of something.

- We care about a particular case of a more general thing.
- As we'll see, abstracting is actually a two-way process.

Abstraction: extraction and application

Abstracting is tracking the structural essence of something.

- We care about a particular case of a more general thing.
- As we'll see, abstracting is actually a two-way process.

What happens when we make use of an abstraction?

- Noticing "the garden is full of weeds," *extracts* concept from exper'nce.
- Now, on the concept level, we choose an action, e.g. "use trowel".
- This abstract action *applies* as experience: we use the trowel.
- The abstract guided our exper'nce by offering a menu of known options

Abstraction: extraction and application

Abstracting is tracking the structural essence of something.

- We care about a particular case of a more general thing.
- As we'll see, abstracting is actually a two-way process.

What happens when we make use of an abstraction?

- Noticing "the garden is full of weeds," *extracts* concept from exper'nce.
- Now, on the concept level, we choose an action, e.g. "use trowel".
- This abstract action *applies* as experience: we use the trowel.
- The abstract guided our exper'nce by offering a menu of known options

Abstracting is not a one-way street; its usefulness comes from application.

- We will use this in two ways: to motive CT and to discuss **Poly**.
  - Category theory is perhaps the most abstract math.
  - **Poly** is a part of CT that holds many of the ideas we're discussing.
- Goal: to see how well the abstraction applies to what we care about.

Abstraction: extraction and application

Abstracting is tracking the structural essence of something.

- We care about a particular case of a more general thing.
- As we'll see, abstracting is actually a two-way process.

What happens when we make use of an abstraction?

- Noticing "the garden is full of weeds," *extracts* concept from exper'nce.
- Now, on the concept level, we choose an action, e.g. "use trowel".
- This abstract action *applies* as experience: we use the trowel.
- The abstract guided our exper'nce by offering a menu of known options

Abstracting is not a one-way street; its usefulness comes from application.

- We will use this in two ways: to motive CT and to discuss **Poly**.
  - Category theory is perhaps the most abstract math.
  - **Poly** is a part of CT that holds many of the ideas we're discussing.
- Goal: to see how well the abstraction applies to what we care about.

To actualize our own potential, we need to be clearer about all our relations

Category theory $\approx$ relational math

Category theory is math about structures and how they relate.

- Mathematical definitions are composed of interlocking structures.
- Category theory tracks the layers of structure and their connections.
- This makes analogies—similarities of structure—into formal objects.

Category theory ≈ relational math

Category theory is math about structures and how they relate.

- Mathematical definitions are composed of interlocking structures.
- Category theory tracks the layers of structure and their connections.
- This makes analogies—similarities of structure—into formal objects.

It was invented in the 1940s to account for *naturality*.

- Mathematicians considered certain transformations to be "natural".
- Setting up the theory to ground it required: categories and functors.
- Grothendieck used CT to prove imp't conjectures in number theory.
- It has largely become the language of pure mathematics.

Category theory ≈ relational math

Category theory is math about structures and how they relate.

- Mathematical definitions are composed of interlocking structures.
- Category theory tracks the layers of structure and their connections.
- This makes analogies—similarities of structure—into formal objects.

It was invented in the 1940s to account for *naturality*.

- Mathematicians considered certain transformations to be “natural”.
- Setting up the theory to ground it required: categories and functors.
- Grothendieck used CT to prove imp’t conjectures in number theory.
- It has largely become the language of pure mathematics.

It is increasingly applied outside of math.

- Some math people laugh at the idea of “applied category theory”:
- Roughly, “how could anything so abstract have applications??”
- But abstraction is tracking the structural essence of something...
- ...and that’s by definition very applicable!
- Despite the learning curve, interest has been consistently growing.

Category theory ≈ relational math

Category theory is math about structures and how they relate.

- Mathematical definitions are composed of interlocking structures.
- Category theory tracks the layers of structure and their connections.
- This makes analogies—similarities of structure—into formal objects.

It was invented in the 1940s to account for *naturality*.

- Mathematicians considered certain transformations to be “natural”.
- Setting up the theory to ground it required: categories and functors.
- Grothendieck used CT to prove imp’t conjectures in number theory.
- It has largely become the language of pure mathematics.

It is increasingly applied outside of math.

- Some math people laugh at the idea of “applied category theory”:
- Roughly, “how could anything so abstract have applications??”
- But abstraction is tracking the structural essence of something...
- ...and that’s by definition very applicable!
- Despite the learning curve, interest has been consistently growing.

Next up, I’ll give you a taste.

Outline

Miracles

In mathematics, there are a few *miracles*.

- The complex numbers are a miracle. You adjoin $\sqrt{-1}$ to $\mathbb{R}$.
- The result is not only that all polynomials factor completely, ...
- ...but that if a function has one derivative, it has infinitely many.
- You seem to get out more than you put in.

Miracles

In mathematics, there are a few *miracles*.

- The complex numbers are a miracle. You adjoin $\sqrt{-1}$ to $\mathbb{R}$.
- The result is not only that all polynomials factor completely, ...
- ...but that if a function has one derivative, it has infinitely many.
- You seem to get out more than you put in.

I'd say that polynomial functors have that flavor.

- They're simple and I'd say miraculous in the above sense.
- It seems to be at the core of a great deal of computer science.
- I won't talk about them in detail but I'll give the main idea.
- Namely: interfaces and delegation are surprisingly useful abstractions.

Miracles

In mathematics, there are a few *miracles*.

- The complex numbers are a miracle. You adjoin $\sqrt{-1}$ to $\mathbb{R}$.
- The result is not only that all polynomials factor completely, ...
- ...but that if a function has one derivative, it has infinitely many.
- You seem to get out more than you put in.

I'd say that polynomial functors have that flavor.

- They're simple and I'd say miraculous in the above sense.
- It seems to be at the core of a great deal of computer science.
- I won't talk about them in detail but I'll give the main idea.
- Namely: interfaces and delegation are surprisingly useful abstractions.

What goal do they serve, especially in terms of this talk?

- Polynomial functors are the simplest framework I've found...
- ...in which we can talk about agents, interactions, dynamics,...
- ...and aim to understand ourselves: what's inside and how we relate.
- **Poly** is a good language with which to ask very basic questions.

Interfaces for tasks and outcomes

Let's talk about "agents" in terms of their *interfaces.*

- Let's hold off on awareness, preferences, etc.
- The interface of an agent is what you can give it and receive from it.
- What's the interface for an agent, like James Bond or a travel agent?

Interfaces for tasks and outcomes

Let's talk about "agents" in terms of their *interfaces.*

- Let's hold off on awareness, preferences, etc.
- The interface of an agent is what you can give it and receive from it.
- What's the interface for an agent, like James Bond or a travel agent?

Let's imagine it like this: you give it a *task*, and it returns an *outcome*.

- Different agents have different sets of tasks they can do.
- And each task has its own set of possible outcomes.
- The task "flip coin" has two possible outcomes: $\{\mathtt{Heads}, \mathtt{Tails}\}$.
- The task "pick number" or "tell joke" has infinitely many outcomes.

Interfaces for tasks and outcomes

Let's talk about "agents" in terms of their *interfaces.*

- Let's hold off on awareness, preferences, etc.
- The interface of an agent is what you can give it and receive from it.
- What's the interface for an agent, like James Bond or a travel agent?

Let's imagine it like this: you give it a *task*, and it returns an *outcome*.

- Different agents have different sets of tasks they can do.
- And each task has its own set of possible outcomes.
- The task "flip coin" has two possible outcomes: $\{\mathtt{Heads}, \mathtt{Tails}\}$.
- The task "pick number" or "tell joke" has infinitely many outcomes.

Let's move into the math. What is the interface for an agent?

- An interface consists of a set T of tasks and,...
- ...for each task $t : T$, a set O_t of possible outcomes for it.

Delegation: mapping between interfaces

Delegation has a lot in common with what we called abstraction.

- Suppose we have two agent interfaces.
- To specify a way that agent 1 *delegates* to agent 2:
  - for each agent-1 task $t_1 : T_1$, specify an agent-2 task $t_2 : T_2$
  - for each agent-2 outcome $o_2{:}O_{t_2}$, specify an agent-1 outcome $o_1{:}O_{t_1}$
- For abstraction, we went from concrete to abstract.
  - Starting with an experience, we *extract* some concept, and...
  - ...given an conceptual action, we *apply* it to the experience.
  - We could say abstraction delegates from experiential to conceptual.

Delegation: mapping between interfaces

Delegation has a lot in common with what we called abstraction.

- Suppose we have two agent interfaces.
- To specify a way that agent 1 *delegates* to agent 2:
  - for each agent-1 task $t_1 : T_1$, specify an agent-2 task $t_2 : T_2$
  - for each agent-2 outcome $o_2 \!:\! O_{t_2}$, specify an agent-1 outcome $o_1 \!:\! O_{t_1}$
- For abstraction, we went from concrete to abstract.
  - Starting with an experience, we *extract* some concept, and...
  - ...given an conceptual action, we *apply* it to the experience.
  - We could say abstraction delegates from experiential to conceptual.

You can encode all this in "polynomial functor" form. Here's an interface:

$$p \coloneqq \sum_{t:T} y^{O_t}.$$

- This actually specifies a *functor*, in the sense of category theory.
- Above "delegation" is precisely a *natural transformation* $p_1 \to p_2$.
- So this back-and-forth mapping is *natural* in the mathematical sense.

The algebra of interfaces

Polynomials let us do algebra with interfaces. Suppose given poly's p, q.

- Each is considered as interface: tasks and their possible outcomes.
- We can combine them in many ways to get new interfaces

$$p + q, \quad p \times q, \quad p \otimes q, \quad p \triangleleft q, \quad [p, q]$$

What are the tasks and outcomes of each of these new interfaces?

- $p + q$: pick a task from p or q; return an outcome of it.
- $p \times q$: pick a task from p and q; return an outcome of either.
- $p \otimes q$: pick a task from p and q; return an outcome of both.
- $p \triangleleft q$: pick a task from p and, for each outcome, pick a task from q;...
- ...return an outcome of each.
- $[p, q]$: pick a delegation from p to q; return a task from p...
- ...and an outcome from its delegate.

The algebra of interfaces

Polynomials let us do algebra with interfaces. Suppose given poly's p, q.

- Each is considered as interface: tasks and their possible outcomes.
- We can combine them in many ways to get new interfaces

$$p + q, \quad p \times q, \quad p \otimes q, \quad p \triangleleft q, \quad [p, q]$$

What are the tasks and outcomes of each of these new interfaces?

- $p + q$: pick a task from p or q; return an outcome of it.
- $p \times q$: pick a task from p and q; return an outcome of either.
- $p \otimes q$: pick a task from p and q; return an outcome of both.
- $p \triangleleft q$: pick a task from p and, for each outcome, pick a task from q;...
- ...return an outcome of each.
- $[p, q]$: pick a delegation from p to q; return a task from p...
- ...and an outcome from its delegate.

What's the point?

- We can combine agent interfaces in all sorts of controlled, lawful ways.

Machines

Moore machines and Mealy machines are used throughout the sciences.

- Each requires a set A of poss. inputs and a set B of poss. outputs.
- Also a set S of "internal states", updated by A's & reading out Bs.

$$A - \boxed{\text{machine}} - B$$

- Moore: update $u \colon A \times S \to S$, readout $r : S \to B$.
- Mealy: combined $f \colon A \times S \to B \times S$.
- Either of these will transform streams of A's into streams of B's.

Machines

Moore machines and Mealy machines are used throughout the sciences.

- Each requires a set A of poss. inputs and a set B of poss. outputs.
- Also a set S of "internal states", updated by A's & reading out Bs.

$$A - \boxed{\text{machine}} - B$$

- Moore: update $u\colon A \times S \to S$, readout $r : S \to B$.
- Mealy: combined $f\colon A \times S \to B \times S$.
- Either of these will transform streams of A's into streams of B's.

We can think of both Moore and Mealy machines in terms of delegations.

- Moore: it's exactly a natural transformation $Sy^S \to By^A$.
- Mealy: it's exactly a natural transformation $Sy^S \to [Ay, By]$.
- In each case, Sy^S is an interface with S as tasks and outcomes.
- A simple agent is in a state $s : S$ and wants a new state $s' : S$.
- It delegates this problem to the Moore interface By^A.
- Moore reveals its task $b : B$ and waits for an outcome $a : A$.

Dynamic arrangements

You can interconnect machines, and the result is another machine.

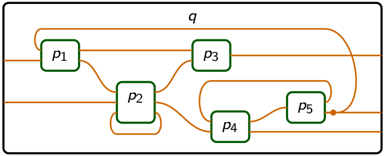

One can even make the interaction pattern *dynamic*:

- Data flowing on the wires can cause the interaction pattern to change.
- This happens in deep learning: the loss changes the weights.
- One can formulate the gradient-backprop algorithm in these terms.

Dynamic arrangements

You can interconnect machines, and the result is another machine.

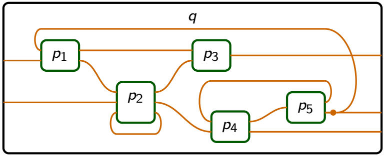

One can even make the interaction pattern *dynamic*:

- Data flowing on the wires can cause the interaction pattern to change.
- This happens in deep learning: the loss changes the weights.
- One can formulate the gradient-backprop algorithm in these terms.

One can also define hierarchical agents that call on others.

- One agent delegates tasks to a host of subagents.
- Define arbitrary flowcharts for how subagents pass data around.
- After finite time, an outcome is returned to the original agent.
- Finally, the flowchart itself is updated in response to the outcome.
- All this is easy to specify (low K-complexity) in the poly. formalism.

Application

Now that we've looked at all this abstraction, how do we apply it?

- We can bring anything from this math back into our experience.
- To do so, someone builds it. Much has already been built (ANNs).
- Other stuff is being built along these lines as we speak.

Application

Now that we've looked at all this abstraction, how do we apply it?

- We can bring anything from this math back into our experience.
- To do so, someone builds it. Much has already been built (ANNs).
- Other stuff is being built along these lines as we speak.

But we might also learn about ourselves from this.

- Are we agents in this sense? Maybe we're more somehow.
- What is this meeting? What makes a body hang together?
- What is eating, reproducing, sensing, healing?
- Where do care and intention come from?

Application

Now that we've looked at all this abstraction, how do we apply it?

- We can bring anything from this math back into our experience.
- To do so, someone builds it. Much has already been built (ANNs).
- Other stuff is being built along these lines as we speak.

But we might also learn about ourselves from this.

- Are we agents in this sense? Maybe we're more somehow.
- What is this meeting? What makes a body hang together?
- What is eating, reproducing, sensing, healing?
- Where do care and intention come from?

I think we can ground philosophical questions in mathematics like this.

- In fact, doing so *creates* new math.
- And that math will apply by helping us make sense of our lives.

Story time

Let's try to do what Turing did, but trying to take care into account.

- He found a memetically-fit abstraction for the *mechanism* of computers
- It's harder but more *important* to understand how the care part works.
- Let's try to tell the story of how we got here and where we're going.
- And let's do so in a way that is plausibly mathematizable.

A story of care

Actualizing potential is already part of physics.

- A hurricane, a lightning bolt: these actualize a potential difference.
- Metabolism-first origins: life as the simplest chemistry lab.
- The tendency to actualize potential was somehow encapsulated.

A story of care

Actualizing potential is already part of physics.

- A hurricane, a lightning bolt: these actualize a potential difference.
- Metabolism-first origins: life as the simplest chemistry lab.
- The tendency to actualize potential was somehow encapsulated.

Actualizing potential locally organizes.

- Being at the place where potential is actualized, you're *swirled* by it.
- Maybe this swirl organizes the local area, "storing information."
- Potential actualization skill begets potential actualization skill.
- That is, care—the tending capacity—becomes held.

A story of care

Actualizing potential is already part of physics.

- A hurricane, a lightning bolt: these actualize a potential difference.
- Metabolism-first origins: life as the simplest chemistry lab.
- The tendency to actualize potential was somehow encapsulated.

Actualizing potential locally organizes.

- Being at the place where potential is actualized, you're *swirled* by it.
- Maybe this swirl organizes the local area, "storing information."
- Potential actualization skill begets potential actualization skill.
- That is, care—the tending capacity—becomes held.

In this story, evolution is guided by care the whole way.

- The "random mutations" are just how outsiders see the search process
- Mate selection, sociality, invention, recenter the living carers.
- The distr'n of care throughout the earth has been in dev. for 4B years.

A story of care

Actualizing potential is already part of physics.

- A hurricane, a lightning bolt: these actualize a potential difference.
- Metabolism-first origins: life as the simplest chemistry lab.
- The tendency to actualize potential was somehow encapsulated.

Actualizing potential locally organizes.

- Being at the place where potential is actualized, you're *swirled* by it.
- Maybe this swirl organizes the local area, "storing information."
- Potential actualization skill begets potential actualization skill.
- That is, care—the tending capacity—becomes held.

In this story, evolution is guided by care the whole way.

- The "random mutations" are just how outsiders see the search process
- Mate selection, sociality, invention, recenter the living carers.
- The distr'n of care throughout the earth has been in dev. for 4B years.

Carers—from neurons to humans and beyond—collab. to actualize potential

- The ability to care is the precious thing that's passed down.

Outline

What holds care?

We are now playing with our own lives. “Gain of function” research.

- The math/CS that we create will take us through the phase transition.
- We’ll process and make sense of the world much more powerfully.
- It may cause bizarre hiccups of miscommunication and “fake news”.
- I believe that the more elegant our abstractions, the fewer the hiccups.

What holds care?

We are now playing with our own lives. "Gain of function" research.

- The math/CS that we create will take us through the phase transition.
- We'll process and make sense of the world much more powerfully.
- It may cause bizarre hiccups of miscommunication and "fake news".
- I believe that the more elegant our abstractions, the fewer the hiccups.

The most important thing—by definition—is what we care about.

- Tending to it well may not be easy, but it's our job.
- We need to ask: what are the dynamics of care?
- Is it individualized or does it pass through and between us?
- Can it get lost? Abandoned? Found? Created? Manufactured?

What holds care?

We are now playing with our own lives. “Gain of function” research.

- The math/CS that we create will take us through the phase transition.
- We’ll process and make sense of the world much more powerfully.
- It may cause bizarre hiccups of miscommunication and “fake news”.
- I believe that the more elegant our abstractions, the fewer the hiccups.

The most important thing—by definition—is what we care about.

- Tending to it well may not be easy, but it’s our job.
- We need to ask: what are the dynamics of care?
- Is it individualized or does it pass through and between us?
- Can it get lost? Abandoned? Found? Created? Manufactured?

Purpose of this talk: I’m worried that we’re not properly *tracking* our care.

- Money is the closest thing, and it’s very poor as a tracker of care.
- If we’re going to have a “good” future, we need to remember care.

Contributing to a fabric of belonging

Imagine a fabric, threads knotted together billions of times.

- We're the knots, connected by the threads of which we're made.
- *Belonging* means going together; this is what shapes the fabric.
- Imagine a lead ball sitting on the fabric. So we pull on each other.
- Making sense together distributes the load of whatever comes our way.

Contributing to a fabric of belonging

Imagine a fabric, threads knotted together billions of times.

- We're the knots, connected by the threads of which we're made.
- *Belonging* means going together; this is what shapes the fabric.
- Imagine a lead ball sitting on the fabric. So we pull on each other.
- Making sense together distributes the load of whatever comes our way.

We're knitting this fabric as we speak. Tying knots. Entangling ourselves.

- Evolution and learning change the consistency of this fabric.
- Now our fabric is going through a phase transition, radically changing.
- We put a lot of thought into intelligence. Care is equally important.
- Our motivations guide the future texture of this fabric.

Contributing to a fabric of belonging

Imagine a fabric, threads knotted together billions of times.

- We're the knots, connected by the threads of which we're made.
- *Belonging* means going together; this is what shapes the fabric.
- Imagine a lead ball sitting on the fabric. So we pull on each other.
- Making sense together distributes the load of whatever comes our way.

We're knitting this fabric as we speak. Tying knots. Entangling ourselves.

- Evolution and learning change the consistency of this fabric.
- Now our fabric is going through a phase transition, radically changing.
- We put a lot of thought into intelligence. Care is equally important.
- Our motivations guide the future texture of this fabric.

The most important thing we can possibly put into it is care.

Thank you for attending. Comments and questions welcome...