Philosophy and Innovation

David Corfield

11 June 2024

Philosophers as innovators rather than system builders?

If computer scientists (and other researchers) can find inspiration in the philosophical writings of the past, may they reasonably expect to profit further by looking to understand the context of these ideas more completely?

Philosophers as innovators rather than system builders?

If computer scientists (and other researchers) can find inspiration in the philosophical writings of the past, may they reasonably expect to profit further by looking to understand the context of these ideas more completely?

Or should they rest content having gleaned some piecemeal ideas from what was presented as a system?

Sections

- Category theory and its uses
- Category theory and AI
- Philosophical precursors

Category theory and its uses

Applications (with approximate dates)

- Mathematics (from the 1940s)
- Logic/Foundations (from the 1960s)
- Computer science (from the 1970s)
- Physics (from the 1980s)
- *'Applied Category Theory'* (from the 2010s)

In mathematics

Today category theory is used everywhere in:

- Algebraic topology
- Algebraic geometry
- Langlands program
- ...

In mathematics

Today category theory is used everywhere in:

- Algebraic topology
- Algebraic geometry
- Langlands program
- ...

Its range is extending...

> *Some time ago I expressed skepticism regarding the possibility of recovering hard analysis categorically. I am still not convinced but it seems soft analysis is already there, in this summer's course by Clausen and Scholze. (Michael Harris, personal communication on Condensed Mathematics and Complex Geometry)*

> *We would like to say that our proofs are proofs by "formal nonsense" and in particular analysis-free. (Clausen and Scholze)*

In logic and computer science

Computational Trinitarianism/Trilogy

1. Idea

A profound cross-disciplinary insight has emerged – starting in the late 1970s, with core refinements in recent years – observing that three superficially different-looking fields of mathematics,

- computation/programming languages
- formal logic/type theory
- ∞-category theory/∞-topos theory (algebraic topology)

are but three different perspectives on a single underlying phenomenon at the foundations of mathematics:

(nLab: computational trilogy)

In physics

- Topological quantum field theory in 1980s
- String diagrams as morphisms in monoidal categories for quantum mechanics, 2000s
- Higher gauge theory, 2000s
- Modal and linear homotopy type theory in Physics, The Quantum Monadology, now

Applied Category Theory today

We see people today applying category theory for:

> *Causality, probabilistic reasoning, statistics, learning theory, deep neural networks, dynamical systems, information theory, database theory, natural language processing, cognition, consciousness, active inference, systems biology, genomics, epidemiology, chemical reaction networks, neuroscience, complex networks, game theory, robotics, quantum computing,...*

Applied Category Theory today

We see people today applying category theory for:

> *Causality, probabilistic reasoning, statistics, learning theory, deep neural networks, dynamical systems, information theory, database theory, natural language processing, cognition, consciousness, active inference, systems biology, genomics, epidemiology, chemical reaction networks, neuroscience, complex networks, game theory, robotics, quantum computing,...*

They're stepping into our territory!

Systems diagrams

A key idea here is Compositionality.

- Plug together systems in parallel.
- Plug together systems in series.
- Plug one system inside another.

A key idea here is Compositionality.

- Plug together systems in parallel.
- Plug together systems in series.
- Plug one system inside another.

This allows the representation of: open graphs, open Petri nets, open chemical reaction with rates, open electrical circuits, open Markov processes, open dynamical systems,... (c.f. Baez)

A cutting edge development along this line is the applied category-theoretic treatment of Stock and Flow diagrams used to model disease spread and control.

- John C. Baez, Xiaoyan Li, Sophie Libkind, Nathaniel D. Osgood and Eric Redekopp, **A categorical framework for modeling with stock and flow diagrams**, to appear in *Mathematics for Public Health*, Springer.
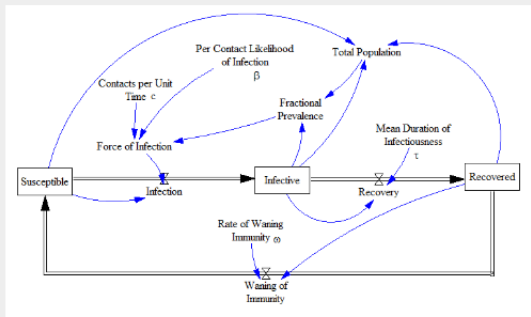
Category theory and AI

Knowledge representation:

> *The ethos of this research program is that category theory can serve as a general purpose modeling language for science and engineering. Having internalized this perspective, it is but a short step to contemplate a general-purpose knowledge representation system based on category theory.*

> *Our philosophy is that category theory is a universal modeling language enabling a more expansive understanding of knowledge representation. (Evan Patterson, Knowledge Representation in Bicategories of Relations, p. 49)*

Category theory and AI

Category theory is well-suited to inferential and learning processes which are

- Compositional
- Structural
- Presentable diagrammatically
- Hybrid mixes

Symbolica – $31 million

Approach

All current state of the art large language models such as ChatGPT, Claude, and Gemini, are based on the same core architecture. As a result, they all suffer from the same limitations.

Extant models are expensive to train, complex to deploy, difficult to validate, and infamously prone to hallucination. Symbolica is redesigning how machines learn from the ground up.

We use the powerfully expressive language of category theory to develop models capable of learning algebraic structure. This enables our models to have a robust and structured model of the world; one that is explainable and verifiable.

It's time for machines, like humans, to think symbolically.

ARIA – safeguarded AI - £59 million

TA1.1 **Theory** shall research and construct computationally practical mathematical representations and formal semantics for world-models, specifications, proofs, neural systems, and "version control" (incremental updates or patches) thereof.
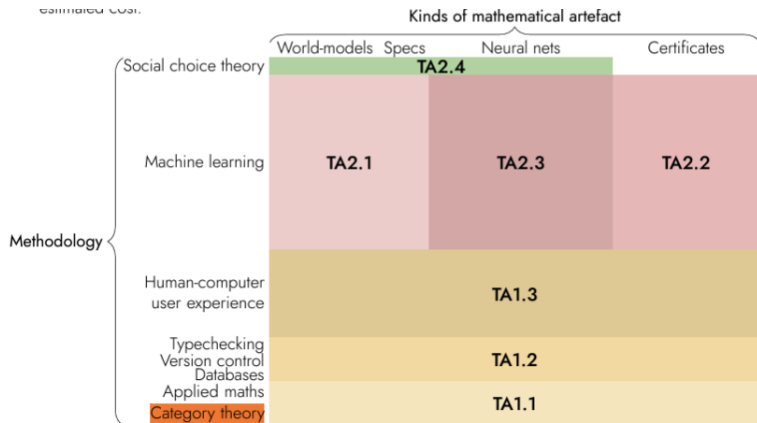
Quantinuum – worth lots

www.quantinuum.com/news/quantinuum-is-developing-new-frameworks-for-artificial-intelligence

Quantinuum is developing new frameworks for artificial intelligence

We do this in part by using a type of math called "**category** theory" that has been used in everything from classical computer programming to neuroscience....
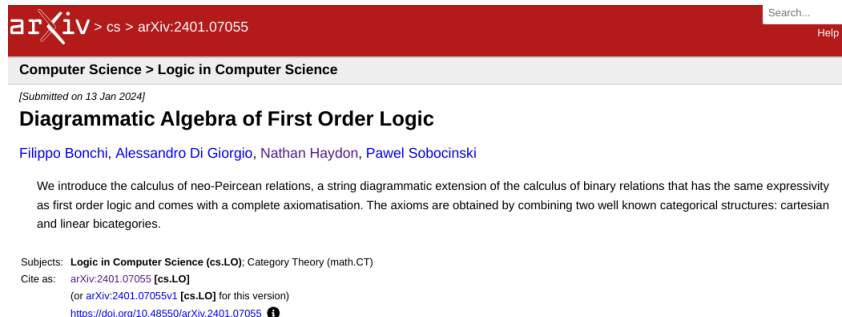
View →

Philosophical precursors: Peirce

Charles Peirce developed his *existential graphs* over many years in three systems: alpha, beta, gamma. The beta system corresponds to first-order logic.

Philosophical precursors: Peirce

Charles Peirce developed his *existential graphs* over many years in three systems: alpha, beta, gamma. The beta system corresponds to first-order logic.

A recent response:

Calculus of relations

With the rise of first order logic, Peirce's calculus was forgotten until Tarski, who in [80] recognised its algebraic flavour. In the introduction to [81], written shortly before his death, Tarski put much emphasis on two key features of CR: (a) its lack of quantifiers and (b) its sole deduction rule of substituting equals by equals. The calculus, however, comes with two great shortcomings: (c) it is strictly less expressive than FOL and (d) it is *not* axiomatisable.

Despite these limitations, CR had —and continues to have— a great impact in computer science, e.g., in the theory of databases [20] and in the semantics of programming languages [2, 38, 45, 47, 74]. Indeed, the lack of quantifiers avoids the usual burden of bindings, scopes and capture-avoid substitutions (see [25, 30, 33, 40, 68, 70] for some theories developed to address specifically the issue of bindings). This feature, together with purely equational proofs, makes CR particularly suitable for proof assistants [43, 71, 72].

In this paper, we introduce the calculus of *neo-Peircean relations*, a string diagrammatic account of FOL that has several key features:

(1) Its diagrammatic syntax is closely related to Peirce's EGs, but it can also be given through a context free grammar equipped with an elementary type system;

(2) It is quantifier-free and, differently than FOL, its compositional semantics can be given by few simple rules: see (8);

(3) Terms and predicates are not treated as separate syntactic and semantic entities;

(4) Its sole deduction rule is substituting equals by equals, like CR, but differently, it features a complete axiomatisation;

(5) The axioms are those of well-known algebraic structures, also occurring in different fields such as linear algebra [11] or quantum foundations [21];

(6) It allows for compositional encodings of FOL, CR and PFL;

(7) String diagrams disambiguate interesting corner cases where traditional FOL encounters difficulties. One perk is that we allow empty models —forbidden in classical treatments— leading to (slightly) more general Gödel completeness;

(8) The corner case of empty models coincides with *propositional* models and in that case our axiomatisation simplifies to the deep inference Calculus of Structures [15, 34].

By returning to the algebraic roots of logic we preserve CR's benefits (a) and (b) while overcoming its limitations (c) and (d).

Let R be a symbol with arity 2 and coarity 0. The two diagrams on the right correspond to FOL

formulas $\exists x.\forall y.\,R(x,y)$ and $\forall y.\exists x.\,R(x,y)$: see § 9 for a dictionary of translating between FOL and diagrams. It is well-known that $\exists x.\forall y.\,R(x,y) \models \forall y.\exists x.\,R(x,y)$, i.e. in any model, if the first formula evaluates to true then so does the second. Within our calculus, this statement is expressed as the above inequality. This can be proved by mean of the axiomatisation we introduce in this work:

$$(1)$$

Features taken over from Peirce

- Variables being captured by wires or strings
- Conjunction as absorbed into the syntax
- Focus on triadic relations
- Reduced set of inference rules
- Emphasis on relational composition

But Peirce's existential graphs were devised as central to his system of philosophy

> *The System of Existential Graphs may be characterized with great truth as presenting before our eyes a moving picture of thought... A picture is visual representation of the relations between the parts of its object; a vivid and highly informative representation, rewarding somewhat close examination. (MS, 1908)*

> *Now let us see how the diagram entrains its consequence. The Diagram sufficiently partakes of the percussivity of a Percept to determine, as its Dynamic, or Middle, Interpretant, a state of activity in the Interpreter, mingled with curiosity. (The New Elements, Vol. 4, p. 318)*

Should we look to understand his innovations in their full context...

...or does Peirce merely provide a potentially useful series of ideas?

- (One of my mine): Deduction, induction, abduction as composition, extension, lifting (see here)

...or does Peirce merely provide a potentially useful series of ideas?

- (One of my mine): Deduction, induction, abduction as composition, extension, lifting (see here)
- But with Peirce there's always a larger story:

II. Cenoscopy, or Philosophy. (Philosophia Prima, First Philosophy.) About positive phenomena in general, such as are available to every person at every waking moment, and not about special classes of phenomena. Does not resort to special experiences or experiments in order to settle theoretical questions.	Epistémy (1902 classification only).[18]	A. Phenomenology (or Categorics or Phaneroscopy). (Includes the study of the cenopythagorean categories: Firstness, Secondness, and Thirdness.)		
		B. Normative Science.	i. Esthetics. (Study of the good, the admirable. Peirce reserved the spelling "aesthetics" for the study of artistic beauty.)	
			ii. Ethics. (Study of right and wrong).	
			iii. Logic (or Semiotic or Formal Semiotic). (Study of true and false.) (The presuppositions of reason are the locus of Peirce's truth theory and his fallibilism.)	1. Speculative Grammar (or Philosophical or Universal Grammar) (or Stecheology) (Includes the classification of signs).
				2. Critic (or Logical Critic, Critical Logic, or Logic Proper). (Includes study of the modes of inference: abduction, induction, and deduction).
				3. Methodeutic (or Speculative Rhetoric, or Universal or Philosophical Rhetoric). (Is the locus of Peirce's Pragmatism, and includes study of scientific method).
		C. Metaphysics.	i. Ontology or General. (Locus of Peirce's Scholastic Realism.)	
			ii. Psychical or Religious.	1. God.
				2. freedom (& destiny).
				3. immortality.
			iii. Physical. (Locus of Peirce's Objective Idealism and his doctrines	

...or does Peirce merely provide a potentially useful series of ideas?

- (One of my mine): Deduction, induction, abduction as composition, extension, lifting (see here)
- But with Peirce there's always a larger story:

II. Cenoscopy, or Philosophy. (Philosophia Prima, First Philosophy.) About positive phenomena in general, such as are available to every person at every waking moment, and not about special classes of phenomena. Does not resort to special experiences or experiments in order to settle theoretical questions.	Epistémy (1902 classification only).[19]	A. Phenomenology (or Categorics or Phaneroscopy). (Includes the study of the cenopythagorean categories: Firstness, Secondness, and Thirdness.)		
		B. Normative Science.	i. Esthetics. (Study of the good, the admirable. Peirce reserved the spelling "aesthetics" for the study of artistic beauty.)	
			ii. Ethics. (Study of right and wrong).	
			iii. Logic (or Semiotic or Formal Semiotic). (Study of true and false.) (The presuppositions of reason are the locus of Peirce's truth theory and his fallibilism.)	1. Speculative Grammar (or Philosophical or Universal Grammar) (or Stechiology) (Includes the classification of signs).
				2. Critic (or Logical Critic, Critical Logic, or Logic Proper). (Includes study of the modes of inference: abduction, induction, and deduction).
				3. Methodeutic (or Speculative Rhetoric, or Universal or Philosophical Rhetoric). (Is the locus of Peirce's Pragmatism, and includes study of scientific method).
		C. Metaphysics.	i. Ontology or General. (Locus of Peirce's Scholastic Realism.)	
			ii. Psychical or Religious.	1. God.
				2. freedom (& destiny).
				3. immortality.
			iii. Physical. (Locus of Peirce's Objective Idealism and his doctrines	

- In any case, Patterson (cited earlier) claims that the path to knowledge representation via the category or bicategory of relations is flawed and that double categories are the answer.
- Where then Peirce?

Philosophical precursors: Geometric deep learning and Cassirer

Cassirer's 1944 paper The concept of group and the theory of perception (translated from an earlier article in French) brings together Klein's *Erlangen Program* with the findings of the *Gestalt school* of psychology.

When we perceive an object or one of its properties we are detecting an invariant feature of our sensations. E.g., an object of constant size although it moves; an object of constant colour, although lighting changes; a tune, although we transpose it to a different key.

Cassirer

If perception is to be compared to an apparatus at all, the latter must be such as to be capable of "grasping intrinsic necessities." Such intrinsic necessities are encountered everywhere. It is only with reference to such "intrinsic necessity" that the "transformation" to which we subject a given form is well defined, inasmuch as the transformation is not arbitrary and executed at random but proceeds in accordance with some rule that can be formulated in general terms. (p. 26) ("grasping intrinsic necessities" is due to Max Wertheimer)

By their reference to such "good points", the particular impressions receive a new kind of determination. They lose, so to speak, their "atomicity", their uniqueness as mere particular items; they unite into groups and totals. (p. 28)

The "images" that we receive from objects, the "impressions" which sensationalism tried to reduce perception to, exhibit no such unity. Each and every one of these images possesses a peculiarity of its own; they are and remain discrete as far as their contents are concerned. But the analysis of perception discloses a formal factor which supersedes this particularity and disparity. Perception unifies and, as it were, concentrates the manifolds of particular images with which we are supplied at every moment...Each invariant of perception is ... a scheme toward which the particular sense-experiences are orientated and with reference to which they are interpreted. (p. 32)

Erlangen Program

Felix Klein (1872) brings order to the panoply of geometries, organising classification around group actions – 'group' in the mathematical sense (nLab).

Geometric deep learning

In this text, we make a modest attempt to apply the Erlangen Programme mindset to the domain of deep learning, with the ultimate goal of obtaining a systematisation of this field and 'connecting the dots'. We call this geometrisation attempt 'Geometric Deep Learning', and true to the spirit of Felix Klein, propose to derive different inductive biases and network architectures implementing them from first principles of symmetry and invariance.

(Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges, Michael M. Bronstein, Joan Bruna, Taco Cohen, Petar Veličković, p.2)

Encoding prior knowledge

Such a 'geometric unification' endeavour in the spirit of the Erlangen Program serves a dual purpose: on one hand, it provides a common mathematical framework to study the most successful neural network architectures, such as CNNs, RNNs, GNNs, and Transformers. On the other, it gives a constructive procedure to incorporate prior physical knowledge into neural architectures and provide principled way to build future architectures yet to be invented. (p. 4)

Equivariance

Two ways to encode group invariance or *equivariance*:

- Algebras for monads of group actions.
- Dependence over delooped groups, BG, (see pp. 124-5 of my MHoTT book relating modal type operators to fixed points and orbits of actions). (An equivariant map is a morphism in the slice over BG.)

A generalised form of the first is developed an article, Categorical Deep Learning: An Algebraic Theory of Architectures timed to launch with Symbolica.

We present our position on the elusive quest for **a general-purpose framework for specifying and studying deep learning architectures**. *Our opinion is that the key attempts made so far lack a coherent bridge between specifying constraints which models must satisfy and specifying their implementations. Focusing on building a such a bridge,* **we propose to apply category theory** *– precisely, the universal algebra of monads valued in a 2-category of parametric maps – as a single theory elegantly subsuming both of these flavours of neural network design.*

A generalised form of the first is developed an article, Categorical Deep Learning: An Algebraic Theory of Architectures timed to launch with Symbolica.

> *We present our position on the elusive quest for* **a general-purpose framework for specifying and studying deep learning architectures**. *Our opinion is that the key attempts made so far lack a coherent bridge between specifying constraints which models must satisfy and specifying their implementations. Focusing on building a such a bridge,* **we propose to apply category theory** *– precisely, the universal algebra of monads valued in a 2-category of parametric maps – as a single theory elegantly subsuming both of these flavours of neural network design.*

The dream

> *we hypothesise neural networks that can learn not merely conservation laws... but verifiably correct logical argument, or code.*

Extended Erlangen

Category theory may be "regarded as a continuation of the Klein Erlangen Programme, in the sense that a geometrical space with its group of transformations is generalized to a category with its algebra of mappings" (Eilenberg and Mac Lane 1945)

(Note that the second approach to equivariance may also be enormously expanded, cf. Higher Klein geometry, which leads to uses in physics–higher Cartan geometry.)

Here the writings of this precursor, the neo-Kantian Cassirer, on invariance and the 'productive imagination" have not been tapped by current researchers, nor his immense body of work on 'Symbolic Forms' (language, myth, art, science, religion, history, and later, economy, technology, law, social life, etc.).

> *we never find naked sensation as a raw material to which some form is given: all that is tangible and accessible to us is rather the concrete determinacy, the living multiformity, of a world of perception, which is dominated and permeated through and through by definite modes of formation. (Cassirer, The Philosophy of Symbolic Forms. Vol. III: Phenomenology of Knowledge., 1955, pp. 14–5)*

Opportunities for pieces of inspiration or something more systematic?

(Note an affinity between Peirce's three categories and Cassirer's triples –
Basisphänomene (Ich – Du – Es), functions (expression - presentation - significance),
expression (mimetic - analogical - symbolic) (cf. Krois here), both focus on semiosis.)

Philosophical precursors: Quantum AI

Quantum cognition meets category theory from Quantinuum:

- Sean Tull, Razin A. Shaikh, Sara Sabrina Zemljič and Stephen Clark, From Conceptual Spaces to Quantum Concepts: Formalising and Learning Structured Conceptual Models

> In this article we present a new modelling framework for concepts based on the mathematical formalism used in quantum theory, and demonstrate how the conceptual representations can be learned automatically from data, using both classical and quantum-inspired models. A contribution of the work is a thorough category-theoretic formalisation of our framework, following Bolt et al. (2019) and Tull (2021). Formalisation of conceptual models is not new (Ganter & Obiedkov, 2016), but we claim that the use of category theory (Fong, 2019), and in particular the use of string diagrams to describe quantum processes (Coecke & Kissinger, 2017), helps elucidate some of the most important features of our approach to concept modelling. This aspect of our work also fits with the recent push to introduce category theory into machine learning and AI more broadly. The motivation is to make deep learning less ad-hoc and less driven by heuristics, by viewing deep learning models through the compositional lens of category theory (Shiebler et al., 2021).

Naturally, I'd be keen for a variety of modal homotopy type to feature:

Exciting work from the physics/computing end: Linear HoTT and QS

QS – **Quantum Systems language @ CQTS**

⤳ full-blown Quantum Systems language emerges embedded in dLHoTT

Dependent Linear Homotopy Type Theory (dLHoTT)
for universal algorithmic quantum computation

Homotopy Type Theory (HoTT)
for topological logic gates

Quantum Systems Language (QS)
for quantum logic circuits

Topological Quantum Gate Circuits
for realistic quantum computation

ambient dLHoTT	verifies	classically dependent quantum linear types
ambient HoTT	provides	specification of topological quantum gates
ambient dTT	provides	full verified classical control

QS claimed there to improve on existing quantum programming languages:
Quipper, QWIRE, etc.

Philosophers of physics should take note:

Besides these technical properties, the logical language QS is curiously satisfying on quantum-philosophical grounds: For example, the internal language-construct in QS for quantum measurement via the modal logic of necessity is *verbatim* the same as for classical measurement, only now applied to (dependent) linear types where it happens to *imply* the collapse of the wavefunction in the categorical semantics. But in the internal logic this effect is just standard conditioning of expectations. In this sense the notorious "measurement problem" of quantum physics disappears when we speak proper QS. (This is analogous to what happens internal to proper quantum probability theory, see there.) Moreover, the *deferred measurement principle* verified in QS implies that even this collapse of the wavefunction in subsystems may be arbitrarily postponed, by observers who have access to the system at large (the "bath").

Philosophers of physics should take note:

Besides these technical properties, the logical language QS is curiously satisfying on quantum-philosophical grounds: For example, the internal language-construct in QS for quantum measurement via the modal logic of necessity is *verbatim* the same as for classical measurement, only now applied to (dependent) linear types where it happens to *imply* the collapse of the wavefunction in the categorical semantics. But in the internal logic this effect is just standard conditioning of expectations. In this sense the notorious "measurement problem" of quantum physics disappears when we speak proper QS. (This is analogous to what happens internal to proper quantum probability theory, see there.) Moreover, the *deferred measurement principle* verified in QS implies that even this collapse of the wavefunction in subsystems may be arbitrarily postponed, by observers who have access to the system at large (the "bath").

Very much computational trinitarian territory.

Linear HoTT provides a controlled combination of tight type discipline with the vector space-like feel of much current AI.

If Linear HoTT could also be used in 'Quantum AI', we would have a great many philosophical precursors to revisit, including Hegel:

nLab: Science of Logic

To conclude

Three examples out of many possibilities.

- Unless these are merely isolated philosophical ideas which fortuitously happen to play a heuristic role in current data science and AI, there should be interest in understanding their settings.
- Philosophers can then contribute to AI by explaining the contexts of such ideas.

To conclude

Three examples out of many possibilities.

- Unless these are merely isolated philosophical ideas which fortuitously happen to play a heuristic role in current data science and AI, there should be interest in understanding their settings.
- Philosophers can then contribute to AI by explaining the contexts of such ideas.
- But category theory relates and unifies different formalisms, so that perhaps we could use category-theoretic AI to explore connections within philosophy itself.
- Of course, these connections are already there – e.g., Peirce and Cassirer are both well aware of Kant and Hegel.

To conclude

Three examples out of many possibilities.

- Unless these are merely isolated philosophical ideas which fortuitously happen to play a heuristic role in current data science and AI, there should be interest in understanding their settings.
- Philosophers can then contribute to AI by explaining the contexts of such ideas.
- But category theory relates and unifies different formalisms, so that perhaps we could use category-theoretic AI to explore connections within philosophy itself.
- Of course, these connections are already there – e.g., Peirce and Cassirer are both well aware of Kant and Hegel.
- But possibly some new light can be cast on these connections.
- Maybe we should think about which philosophers are relevant to ACT.