

Relative toposes and meta-learning

Olivia Caramello

(University of Insubria - Como, and Grothendieck Institute)

Computing and Artificial Intelligence Summit
London, 3-4 June 2024

The key features of true intelligence

Human intelligence enjoys a number of fundamental features that are responsible for its effectiveness:

- the ability to identify **invariants**, giving a qualitative, rather than numerical or quantitative, understanding of complex phenomena;
- the ability to combine and synthesize **different points of view** on a given theme, corresponding to different invariants;
- the ability to function at **different levels of abstraction**, and to build on previously accumulated knowledge;
- the use of **languages** of different kinds to describe and represent pieces of reality, so that findings about the world can be expressed in them, stored and communicated to other agents.

Despite the spectacular advances in deep learning systems of the last years, we are still far from a form of artificial intelligence enjoying these features. In fact, **new mathematical foundations** are needed to achieve these goals.

The next-generation AI

Relative toposes
and
meta-learning

Olivia Caramello

General
principles for a
future AI

The relativity
method
Relative toposes
Relevance for AI

Classifying
toposes

Syntactic
learning via
relative toposes

Modelling the
stratified nature of
knowledge

The logic of games
A financial
application

The new generation of artificial learning systems should be based on computations that are not blind, but conceptually inspired and **meaningful** to the human mind.

In particular, we should shift to a new conception of **information** which is based on **semantics** and **invariants** rather than on the classical set-theoretic foundations for mathematics.

Topos theory is bound to play a key role in this kind of developments.

In this talk, I will focus on one particular aspect of the theory, that is, the possibility of studying toposes in relation to each other, and discuss its relevance for AI. In particular, I will argue about the role that relative toposes can play in connection with the development of systems for **meta-learning** (i.e. learning taking place at different levels of abstraction constructed on top of each other).

The relativity method

- In Mathematics the **relativity method** consists in trying to state notions and results in terms of **morphisms**, rather than objects, of a given category, so that they can be ‘relativized’ to an arbitrary base object.
- One works in the new base universe as if it were the ‘classical’ one, and then interprets the obtained results from the point of view of the original universe. This process is usually called *externalization*.
- Relativity techniques can be thought as general ‘**change of base techniques**’, allowing one to choose the universe relatively to which one works according to one’s needs.
- The relativity method has been pioneered by Grothendieck, in particular for **schemes**, in his categorical refoundation of Algebraic Geometry, and has played a key role in his work.
- Since the theory of toposes is more general and technically better-behaved than the theory of schemes, it is particularly important to develop relativity techniques for **toposes**.

Applications to mathematics

- By definition, a **relative topos** is a morphism between toposes. The codomain topos is regarded as the base over which the domain topos is defined.
- The relativity method provides an extremely high degree of **technical flexibility**, resulting from the possibility of ‘encapsulating’ part of the complexity of a situation in the base topos, so that the given notions acquire a simpler (e.g. a lower-degree) expression with respect to it.
- For example, a first-order theory becomes a propositional one when regarded relative to a suitable base topos.
- Several mathematicians have already started using **relativity techniques for toposes**:
 - Scholze and Clausen’s condensed mathematics;
 - Tao and Janneshan’s topos-theoretic measure theory;
 - Tomasic’s topos-theoretic difference algebra.

The theory of relative toposes is currently being systematically developed, from the geometric, algebraic and logical points of view, by our research team.

Relative toposes for AI

- Since toposes are spaces which embody **information**, in a way which is **invariant** with respect to the different syntactic presentations, the theory of relative toposes is also crucial in connection with the development of a theory of **semantic information**.
- A **relative topos** corresponds to a way of organizing a given piece of information by incorporating part of it in the base topos and the ‘remaining part’ in the topos defined over it.
- Modelling learning through chains of relative toposes thus allows us to formalize the crucial ability of human intelligence to reason at different levels of abstraction and learn on top of existing knowledge, that is, as we shall see, to realize a form of **meta-learning**.

Classifying toposes

Relative toposes
and
meta-learning

Olivia Caramello

General
principles for a
future AI

The relativity
method

Relative toposes
Relevance for AI

Classifying
toposes

Syntactic
learning via
relative toposes

Modelling the
stratified nature of
knowledge

The logic of games
A financial
application

A fundamental connection between topos theory and logic is realized by the theory of classifying toposes.

With any first-order mathematical theory \mathbb{T} (of a very general form) one can canonically associate a topos $\mathcal{E}_{\mathbb{T}}$, called its **classifying topos**, which represents its **'semantical core'**.

The classifying topos embodies the semantic 'essence' of the theory, an essence which is **invariant** with respect to its different syntactic (axiomatic) presentations.

The classifying topos of a theory is constructed by means of **completion process** of the theory, with respect, in a sense, to all the concepts that it is potentially capable to express.

It is at the level of these completed objects that **invariants** actually live.

From a sketch to reality

Relative toposes
and
meta-learning

Olivia Caramello

General
principles for a
future AI

The relativity
method

Relative toposes
Relevance for AI

Classifying
toposes

Syntactic
learning via
relative toposes

Modelling the
stratified nature of
knowledge

The logic of games

A financial
application

Every language, in its attempt to express a reality that is much richer, can be compared to a sketch; the transition from a linguistic expression to its **meaning** is thus a kind of **completion**, similar to the automatic one performed by our brain as we watch the artist drawing it:

The classifying topos construction provides a mathematical formalization of this completion process. Indeed, the classifying topos makes *explicit* whatever is *implicit* in the theory.

Empowering learning systems with logic

Artificial learning systems, in particular deep neural networks, should be empowered by building on the notions of mathematical **theory** and **proof**:

- Any intelligent agent must, in order to get an effective understanding of (aspects of) the world, derive knowledge starting from certain 'sensory' inputs, which play a similar role to that of **axioms** for a mathematical theory, by following certain dynamical rules, which correspond to the **inference rules** of the logical system inside which the mathematical theory is formulated.
- As every mathematical theory can be enriched by the addition of new axioms, so the functioning of an agent can be **updated** by the integration of new information which becomes available to it.
- The functioning of a learning system would thus correspond to a **sequence of a mathematical theories**, each of which richer or more refined than the previous ones.

Learning processes via proofs

Relative toposes and meta-learning

Olivia Caramello

General principles for a future AI

The relativity method

Relative toposes
Relevance for AI

Classifying toposes

Syntactic learning via relative toposes

Modelling the stratified nature of knowledge

The logic of games

A financial application

- In principle, any learning sequence, as any process ‘approximating truth’, is infinite, but for practical purposes one is normally satisfied by the result of a learning process when the last theory in the sequence leaves a degree of **ambiguity** which is sufficiently low for the desired applications.
- All the theories in the sequence should extend (that is, be defined over) the basic theory of the agent formalizing its essential features. More generally, every theory in the sequence can be seen as theory defined *over* each of the preceding ones. This corresponds, through the classifying topos construction, to a **sequence of relative toposes**.
- Note that any **constraints** embedded in the logical formalism (or integrated at some step of the sequence of theories) will allow to **significantly reduce the space of parameters** that the agent has to explore and hence correspondingly decrease the computational complexity of the learning process.

The stratified nature of knowledge

Relative toposes
and
meta-learning

Olivia Caramello

General
principles for a
future AI

The relativity
method

Relative toposes
Relevance for AI

Classifying
toposes

Syntactic
learning via
relative toposes

Modelling the
stratified nature of
knowledge

The logic of games
A financial
application

Learning is a process that goes through a sequence of steps: any new, learned information is stored in the brain according to pre-existing structures.

In fact, human learning is essentially **relational**, i.e. it consists in relating the given new piece of information with the previously accumulated knowledge. This is actually what makes it robust; in fact, only structured forms of learning can enjoy strong forms of robustness and resilience.

The realisation of the 'relative nature' of learning also allows us to understand that one can arrive at arbitrarily high degrees of complexity and abstraction in one's learning/understanding by means of a **structured composition of successive fundamental steps**, whose nature can be more easily investigated.

The stratified nature of knowledge

Relative toposes
and
meta-learning

Olivia Caramello

General
principles for a
future AI

The relativity
method

Relative toposes
Relevance for AI

Classifying
toposes

Syntactic
learning via
relative toposes

Modelling the
stratified nature of
knowledge


The logic of games

A financial
application

Think, for instance, about **learning a language**: before being able to write good **prose** or **poetry**, one needs to have a perfect command of the grammar. Now, in order to be able to properly learn a **grammar**, one needs to have a familiarity with the **vocabulary** of the language, which in turn requires a knowledge of the **alphabet** in which it is written.

So, we can regard the knowledge of the alphabet as lying at level 0, that of the vocabulary as lying at level 1 and that of grammar lying at level 2, while prose and poetry lie at much higher levels defined over the preceding ones.

The learning of music follows a similar scheme, as well as that of any art or science.



Stratified knowledge and relativity principles

The theory of relative toposes provides a **mathematical formalisation** of this learning architecture.

Given a piece of information embodied by a certain topos, finding a 'base topos' over which the given topos is defined corresponds to an abstraction operation, which consists in organizing the given piece of information by incorporating part of it in the base topos and the remaining part in the topos defined over it.

For example, one can reduce complex notions to simpler ones by changing the base category relative to which one works:


- Topological vector space = vector space relative to the category of topological spaces;
- Lie group = group relative to the category of differentiable manifolds.

The power of relativity techniques is that one can do topos theory relative to an arbitrary base topos *as if* that base topos were the ordinary topos of sets, but the results obtained, once externalized (that is, reinterpreted from the point of view of the 'ground level') are typically much more complex.

Chains of relative theories

Learning thus corresponds to a sequence of relative toposes, each defined over the previous one, with the ground one corresponding to 'raw data'. In fact, each topos in the sequence spatially embodies a certain amount of information, which can also be presented in logical form: while ordinary, ground-base, classifying toposes can only classify first-order theories, by means of relative toposes one can classify theories of arbitrarily high order.

For example, a **theory of order 1** with respect to the ground topos can become a **theory of order 0** with respect to a topos which is richer than the ground topos. More generally, a **theory of order n** can be regarded as a theory of order $n - k$ with respect to a k -chain of relative toposes.



Syntactic learning

In order to practically implement the above ideas, one should, first of all, empower any learning system with **large formal vocabularies** that will serve for expressing the concepts that the system will learn from data.

The idea is to **enforce the learning to take place at the abstract level of syntax**, rather than at that of a particular semantics, as it currently happens.

While the vocabulary should be given at the outset, we could **let the system discover any (non-already embedded) logical rules expressible in that vocabulary** by itself, by using the usual techniques. Also, we should let it suggest enrichments of the vocabulary on the basis of invariances empirically discovered in the data, thereby achieving a form of **'emergence of concepts'**.

In this way, we could obtain systems capable of inferring all sorts of **'syntactic rules'** from data (e.g. the grammar rules of a language from a great amount of samples of texts, or the rules of a game from a big collection of matches, etc.): these rules could, of course, be of different nature and complexity (e.g., propositional, first-order, higher-order, etc.).

Making AI systems 'speak'

Relative toposes
and
meta-learning

Olivia Caramello

General
principles for a
future AI

The relativity
method

Relative toposes
Relevance for AI

Classifying
toposes

Syntactic
learning via
relative toposes

Modelling the
stratified nature of
knowledge

The logic of games

A financial
application

All of this goes in the direction of constructing **logical languages for AI agents**, allowing them, in a sense, to 'speak'. After all, how can we expect a system to learn in a robust way if we do not give it the possibility of reasoning linguistically? Think about how our learning, as human beings, would be impaired without the possibility of expressing, testing and communicating our ideas by using languages.

Think also about the process of learning of a natural language. Little babies have no other possibility to learn a language than to merely rely on data (note, however, that their brain has already a lot of *a priori structures* which are used to organize and categorize the knowledge that they gather from the environment). Still, their knowledge of the language remains fragile until they are brought into contact, notably at school, with **grammar**, which represents syntax in this context.

The role of grammar is crucial in making *explicit* a lot of the *implicit* which had been accumulated in the previous 'bottom-up' learning process, and in bringing the understanding to a higher level, notably in terms of **explainability**.

Uncovering hidden structures and rules

Syntax represents the ‘**skeleton**’ of structures, whose concrete manifestations we access through data. Trying to ‘**lift**’ from a particular concrete setting to the level of syntax represents an abstract form of understanding which enjoys much more resilience and adaptability than the usual forms of statistically-based learning, which do not go to the roots of the reasons for the existence of the given relations.

The deep reasons for the regularities that we may observe in concrete contexts actually live at the syntactic level (think, for instance, of motives in Algebraic Geometry, or to definability and preservation theorems in Logic).

Philosophically speaking, we should teach learning systems to lift from the **phenomenological** level (of concrete manifestations of topos invariants) to the **ontological** one (of toposes themselves, regarded as classifying toposes of logical theories).

In particular, we should aim for an integration of a ‘bottom-up’ approach to artificial learning, such as the one which is dominant today, with a ‘top-down’ one based on **logic** and **topos invariants**.

The logic of games

Relative toposes
and
meta-learning

Olivia Caramello

General
principles for a
future AI

The relativity
method

Relative toposes
Relevance for AI

Classifying
toposes


Syntactic
learning via
relative toposes

Modelling the
stratified nature of
knowledge

The logic of games
A financial
application

Games can be very conveniently modeled in terms of chains of relative toposes, with the ground topos formalizing the “hardware of the game”, the level 1 topos the basic, “first-order rules” of the game, and the higher-level relative toposes the “higher-level rules” (i.e. tactics or strategies).

Each morphism in the chain encodes the way in which the rules at that level are concretely realized in terms of lower-level ones; in particular, the morphism from the level 1 topos to the base topos expresses the way in which the basic rules (e.g. those governing the behaviour of the “pieces” of the game) are “arithmetically encoded” in the game’s hardware.



A topos-inspired trading system

Since 2013, the Italian company EOS S.r.l. has developed an award-winning stock trading system inspired by the above-mentioned topos-theoretic ideas. The most recent developments of the system involve the construction of **meta-models** of order up to 10, which perform significantly better than ordinary (or lower-order) models in terms of robustness and resilience, as shown by the following table (courtesy of EOS):

LEVEL	Performance		Differential with respect to the benchmark 2003-2023		
	Back-test period 2003-2014	Offline period 2015-2023	% positive years	Best annual differential	Worst annual differential
Level 0	18.72	21.70	90.79%	54.28%	-5.45%
Level 1	14.08	14.11	90.48%	36.43%	-3.88%
Level 2	14.04	18.07	94.18%	39.05%	-2.83%
Level 3	14.36	16.94	93.65%	34.63%	-3.61%
Level 4	14.53	17.74	98.41%	36.24%	0.50%
Level 5	14.80	18.46	96.30%	38.14%	-0.32%
Level 6	13.46	18.65	96.30%	33.12%	-0.44%
Level 7	14.18	19.97	97.88%	34.81%	-0.50%
Level 8	14.06	18.56	96.30%	32.60%	-0.03%
Level 9	13.05	19.04	98.94%	31.87%	1.22%
Level 10	13.54	19.68	98.94%	34.85%	1.01%

The table reports the results of theoretical operativity, for the period 2003-2023, of models and meta-models applied to a basket of large-cap stocks traded on the Italian market. The benchmark is the FTSE ITALIA All-Share index. The total period is divided into two sub-periods: back-test (2003-2014) and offline (2015-2023).

International school and conference on topos theory

Relative toposes
and
meta-learning

Olivia Caramello

General
principles for a
future AI

The relativity
method

Relative toposes

Relevance for AI

Classifying
toposes

Syntactic
learning via
relative toposes

Modelling the
stratified nature of
knowledge

The logic of games

A financial
application

**TOPOSES
IN MONDOVI**

SCHOOL
3-6 SEPTEMBER 2024
Aula Magna - Mondovì site of
the Polytechnic University of Turin

School Lecturers
Marc Aiguier
Central Institute, Université de Bourgogne
Olivia Caramello
University of Innsbruck and Institute of Mathematics
Alain Connes
IHES and Collège de France
Denis-Charles Cisinski
University of Regensburg
Dennis Gaitsgory
Max Planck Institute Mathematics
Laurent Lafforgue
Humboldt
Michael Robinson
American University
Isar Stubbe
Université du Littoral-Côte d'Opale

CONFERENCE
9-11 SEPTEMBER 2024
Circolo di Lettura
Palazzo del Governatore - Mondovì

Invited Speakers
Joseph Bernstein
Tel Aviv University
Paolo Giordano
University of Vienna
Peter Haine
University of California, Berkeley
Matthias Hutzler
University of Gothenburg
Asgar Janneshan
IAS, University
Maxim Kontsevich
IHES
Hugo Luiz Mariano
University of San Paulo
Axel Osmond
Institut Grothendieck
Hans Riess
Cornell University
Michael Shulman
University of California, San Diego

Scientific Committee
Olivia Caramello, Alain Connes, Laurent Lafforgue

Information & registration:
<https://www.grothendieck.org/toposesmondovi>
Cooperating:
Istituto Grothendieck

ISTITUTO GROTHENDIECK

Everyone is welcome!