

Structure and Power: an Emerging Landscape

Samson Abramsky*

Department of Computer Science

University College London

66-72 Gower St., London WC1E 6EA, U.K.

s.abramsky@ucl.ac.uk

Abstract. In this paper, we give an overview of some recent work on applying tools from category theory in finite model theory, descriptive complexity, constraint satisfaction, and combinatorics. The motivations for this work come from Computer Science, but there may also be something of interest for model theorists and other logicians.

The basic setting involves studying the category of relational structures via a resource-indexed family of adjunctions with some process category - which unfolds relational structures into tree-like forms, allowing natural resource parameters to be assigned to these unfoldings. One basic instance of this scheme allows us to recover, in a purely structural, syntax-free way: the Ehrenfeucht-Fraïssé game; the quantifier rank fragments of first-order logic; the equivalences on structures induced by (i) the quantifier rank fragments, (ii) the restriction of this fragment to the existential positive part, and (iii) the extension with counting quantifiers; and the combinatorial parameter of tree-depth (Nesetril and Ossona de Mendez). Another instance recovers the k-pebble game, the finite-variable fragments, the corresponding equivalences, and the combinatorial parameter of treewidth. Other instances cover modal, guarded and hybrid fragments, generalized quantifiers, and a wide range of combinatorial parameters. This whole scheme has been axiomatized in a very general setting, of arboreal categories and arboreal covers.

Beyond this basic level, a landscape is beginning to emerge, in which structural features of the resource categories, adjunctions and comonads are reflected in degrees of logical and computational tractability of the corresponding languages. Examples include semantic characterisation and preservation theorems, and Lovász-type results on counting homomorphisms.

*Address for correspondence: Department of Computer Science, University College London, 66-72 Gower St., London WC1E 6EA, U.K.

1. Introduction

In recent work [7, 13], a program has been initiated of

Relating Structure and Power

Here we have in mind two major organizing principles in the foundations of computation:

Structure Compositionality and semantics, addressing the question of mastering the size and complexity of computer systems and software.

Power Expressiveness and computational complexity, addressing the question of how we can harness the power of computation and recognize its limits.

A striking feature of the current state of the art is that there are *almost disjoint communities* of researchers studying **Structure** and **Power** respectively, with no common technical language or tools. In our opinion, this is a major obstacle to fundamental progress in Computer Science.

The research program initiated in [7, 13], and developed further in [17, 18, 11], aims at relating categorical semantics, which exemplifies **Structure**, to finite model theory, which exemplifies **Power**. This is the focus of a current joint project with Anuj Dawar.¹ Contributors to this program include Dan Marsden, Luca Reggio, Tomáš Jakl, Tom Paine, Nihil Shah, Adam Ó Conghaile and Yoàv Montacute.

The key initial idea of this program [7, 13] is to encapsulate various forms of model comparison game, in which Spoiler tries to distinguish two structures, and Duplicator tries to show they are the same, as *resource-indexed comonads* on the category of relational structures. Intuitively, these comonads correspond to *co-resources*, which limit the ability of Spoiler to access the structure by bounding resources of some kind. Thus, if C_k is such a comonad, with resource index k , then to have a homomorphism

$$C_k A \rightarrow B$$

means that we only have to check the homomorphism conditions against limited – k -bounded – parts of the structure of A .

Some key features have emerged in the elaboration of this idea in subsequent work:

- This idea has proved to be extremely robust. It can be used to capture a wide range of model comparison games, and the resource-indexed equivalences induced by the corresponding logics. These include the quantifier rank fragments, the finite variable fragments, the modal, guarded, hybrid and bounded fragments, generalized quantifiers, and more.
- The basic idea of coresources $C_k A \rightarrow B$ seems at first blush restricted to forth-only equivalences, and hence to existential positive fragments. However, the ideas in fact extend smoothly to cover the full back-and-forth equivalences, using a refinement of the open maps formulation of bisimulation to *open pathwise embeddings*. Moreover, isomorphism in the coresource setting characterizes extensions of the logics with counting quantifiers.

¹EPSRC-funded project EP/T00696X/1: Resources and Co-resources: a junction between categorical semantics and descriptive complexity.

- Coalgebras $A \rightarrow C_k A$ correspond to resource-bounded *structural decompositions* of A . The existence of such decompositions yields significant combinatorial invariants of A . This allows important combinatorial parameters such as *tree-width* and *tree-depth* to be recovered.
- The whole pattern of comonads, their coalgebras, and the corresponding resolutions into comonadic adjunctions, forms a robust template which recurs throughout finite model theory and descriptive complexity. This template has been axiomatized at a very general categorical level in [11]. This provides a new kind of axiomatic basis for model theory, finite and infinite.

We can regard the axiomatization in [11] as the culmination of a “first wave” of this research program. Results building on this initial phase of development are beginning to emerge. These include:

- General versions of model-theoretic results such as preservation theorems: Rossman’s homomorphism preservation theorems, the van Benthem-Rosen theorem on bisimulation invariance, etc.
- Uniform proofs of preservation theorems in the finite and infinite cases: “model theory without compactness”.
- Structural features of comonads (idempotence, bisimilar companions property), and their significance for computational tractability.
- Lovász-type theorems on counting homomorphisms.
- Combinatorial parameters: concrete cases, and an axiomatic approach via density comonads.
- Feferman-Vaught-Mostowski type theorems, with applications to Courcelle’s theorem.
- New cohomological approximation algorithms for constraint satisfaction and structure isomorphism.

Our aim in the present paper is to give an overview of this line of research, and the new directions currently being developed. We aim to give an accessible account, emphasising conceptual elements, and keeping technicalities to a minimum. In particular, we aim to make this presentation accessible both to people working in finite model theory and descriptive complexity, and to “mainstream” model theorists.

Boaz Trakhtenbrot

This paper is submitted to a special issue in celebration of the centenary of Boaz Trakhtenbrot, one of the founding fathers of theoretical computer science. I had the privilege of meeting Boaz on a number of occasions, and hosted a visit by him to Imperial College in 1990. I warmly remember the fascinating conversations I had with him. He was a unique link to the history and prehistory of our field. Even more importantly, he lived and breathed his science in the present, with passionate commitment.

Boaz was a pioneer of both *Structure* and *Power*. In particular, many of the renowned results from the earlier part of his career relate to *Power*:

- Trakhtenbrot’s theorem

- Büchi-Elgot-Trakhtenbrot theorem
- Borodin-Trakhtenbrot gap theorem

In his later work, *Structure* played an equally prominent role:

- semantics of dataflow
- behaviour structures and nets
- hybrid systems

I believe that Boaz had a unified vision of the field, and that our attempts to relate *Structure* to *Power* are fully in the spirit of his work.

2. Background

2.1. The setting: relational structures

A relational vocabulary σ is a family of relation symbols R , each of some arity $n > 0$. A relational structure for σ is $(A, \{R^A \mid R \in \sigma\})$, where $R^A \subseteq A^n$. We shall not distinguish notationally between a structure and its underlying set. A homomorphism of σ -structures $f : A \rightarrow B$ is a function $f : A \rightarrow B$ such that, for each relation $R \in \sigma$ of arity n and $(a_1, \dots, a_n) \in A^n$:

$$(a_1, \dots, a_n) \in R^A \Rightarrow (f(a_1), \dots, f(a_n)) \in R^B.$$

These notions are pervasive in

- logic (model theory),
- computer science (databases, constraint satisfaction, finite model theory)
- combinatorics (graphs and graph homomorphisms).

Our setting will be $\text{Struct}(\sigma)$, the category of σ -structures and homomorphisms.

2.2. Remarks to model theorists


A first remark is that $\text{Struct}(\sigma)$ is *not* the usual category (implicitly) studied in model theory. Instead, stronger properties are required for morphisms, typically that they are embeddings or elementary embeddings, so that truth of a larger class of first-order sentences is preserved. These categories are “thin”, with rather few morphisms, and consequently lack interesting categorical structure. By contrast, $\text{Struct}(\sigma)$ has a rich categorical structure – it is in fact a quasitopos.² Moreover, the fact that we focus on this category with all homomorphisms does not at all mean that we cannot capture elementary equivalence, as we shall see.

This difference in setting is surely part of the explanation for the strange dissociation which exists between model theory and category theory. Both model theory and category study the structure of

²This observation is due to Dan Marsden.

mathematical theories in general, but they do so in rather different ways, and with strikingly little communication between them.

What makes this dissociation all the stranger is indicated by the following diagram:



Algebraic and arithmetic geometry are prime foci of contemporary applied model theory. Indeed, in an aphorism of Wilfrid Hodges [24]:

$$\text{model theory} = \text{algebraic geometry} - \text{fields.}$$

Yet contemporary arithmetic and algebraic geometry are imbued with the Grothendieck heritage: categories, sheaves, schemes, cohomology – and beyond (stacks, higher categories etc.).

In our view, just as **Structure** and **Power** can benefit from being brought together and intertwined, so category theory and model theory can both benefit from some rapprochement.

One existing locus of interaction is between accessible categories and abstract elementary classes, see e.g. [14]. The connections made between these theories have a heavily set-theoretic flavour.

By contrast, we aim for a resource-sensitive structure theory “down below”, capturing fine structure of resource-indexed fragments in a form relevant to computational tractability and complexity.

2.3. Brief review of adjunctions and comonads

We recall the notion of adjunction through an example which hopefully will be familiar.

Given a commutative ring R , the category of R -modules is denoted $R\text{-Mod}$. There is an evident forgetful functor $U : R\text{-Mod} \rightarrow \text{Set}$, and an adjunction

$$\text{Set} \begin{array}{c} \xrightarrow{R(\bullet)} \\ \perp \\ \xleftarrow{U} \end{array} R\text{-Mod}$$

$R^{(X)}$ is the free module generated by X (formal finite R -linear combinations over X). The unit of the adjunction is the map $\eta_X : X \rightarrow UR^{(X)}$ which sends x to $1 \cdot x$.

Freeness is captured by a *universal mapping property*:

$$\begin{array}{ccc} \text{Set} & & R\text{-Mod} \\ X & \xrightarrow{\eta_X} & UR^{(X)} \\ & \searrow f & \downarrow Uf \\ & & UM \end{array} \qquad \begin{array}{c} R^{(X)} \\ \downarrow \hat{f} \\ M \end{array}$$

which says that every map $f : X \rightarrow UM$ factors through η_X via a unique R -module morphism $\hat{f} : R^{(X)} \rightarrow M$.

Given any adjunction

$$\mathcal{C} \begin{array}{c} \xrightarrow{L} \\ \perp \\ \xleftarrow{R} \end{array} \mathcal{D}$$

there is an associated *monad* RL on \mathcal{C} , and *comonad* LR on \mathcal{D} .³ These notions are pervasive in mathematics [28]:

- monads occur in topology, and finitary monads on **Set** subsume universal algebra.
- comonads feature in descent theory.
- adjunctions, monads and comonads specialize to familiar notions on posets: adjunctions specialize to Galois correspondences, monads to closure operators, and comonads to coclosures.

In our context, we can think of a resource-indexed comonad C_k as a *modality*:

$$C_k A \rightarrow B$$

means we have a homomorphism which only needs to be checked against a limited part of the structure of A . This is exactly what logical languages do in model theory! They calibrate limited means for accessing structures.

2.4. The general scheme: resource-indexed adjunctions

In our approach, we build tree-structured covers of a given, purely extensional relational structure. Such a tree cover will in general not have the full properties of the original structure, but be a “best approximation” in some resource-restricted setting. More precisely, this means that for each integer $k > 0$ we have an adjunction

$$\mathcal{A}_k \begin{array}{c} \xrightarrow{L_k} \\ \perp \\ \xleftarrow{R_k} \end{array} \text{Struct}(\sigma)$$

between the category of σ -structures, and the resource-bounded category \mathcal{A}_k . This adjunction yields the corresponding comonad $C_k = L_k R_k$.

The objects of the category where the approximations live have an intrinsic tree structure, which can be captured axiomatically, as *arboreal categories*. The tree encodes a process for generating (parts of) the relational structure, to which resource notions can be applied. This allows us to apply these resource notions to the objects of the extensional category via the adjunction.

We shall now see what all this means in terms of a simple but fundamental example.

³Monads and comonads have additional structure, of a unit and multiplication, and counit and comultiplication, respectively. We will give additional details later when we come to coalgebras, but prefer to keep the exposition simple for now.

3. First example

We need a few notions on posets. A *forest* is a poset (F, \leq) such that, for all $x \in F$, the set of predecessors of x is a finite chain.⁴ The roots of a forest are the minimal elements. The *height* of a forest is the supremum of the cardinalities of its chains (note that the branches of the forest are its maximal chains). A *tree* is a forest with a unique root. A forest morphism is a map preserving roots and the *covering relation*: $x \prec y$ if $x < y$, and for all z , $x \leq z \leq y$ implies $z = x$ or $z = y$.

The *Gaifman graph* of a σ -structure A has set of vertices A , and a adjacent to b if $a \neq b$ and a, b both occur in some tuple of one of the relations R^A , $R \in \sigma$.

Now a *forest-ordered σ -structure* (A, \leq) is a σ -structure A with a forest order \leq on A . This must satisfy the following condition:

(E) elements of A which are adjacent in the Gaifman graph of A must be comparable in the order.

The *minimum height* of such a forest order on A is the *tree-depth* of A (Nešetřil and Ossona de Mendez [30]). This is an important combinatorial parameter, used extensively by Rossman in his Homomorphism Preservation Theorems.

Morphisms of forest-ordered σ -structures are σ -homomorphisms which are also forest morphisms. This gives rise to a category $\mathcal{R}^E(\sigma)$, and an evident forgetful functor $U : \mathcal{R}^E(\sigma) \rightarrow \text{Struct}(\sigma)$. $\mathcal{R}^E(\sigma)$ is our first example of an *arboreal category*.

There is a natural resource-indexing by height. For each $k > 0$, if we restrict to forest orders of height $\leq k$, we get a sub-category $\mathcal{R}_k^E(\sigma)$, and a functor $U_k : \mathcal{R}_k^E(\sigma) \rightarrow \text{Struct}(\sigma)$. This functor has a right adjoint G_k , giving rise to a comonad $\mathbb{E}_k = U_k G_k$ on $\text{Struct}(\sigma)$.

We describe the construction for $G_k A$, which builds a forest-ordered σ -structure from any σ -structure A in a “cofree” way.⁵

- Given a structure A , the universe of $G_k A$ is $A^{\leq k}$, the non-empty sequences of length $\leq k$.
- This is forest-ordered by the prefix order.
- The counit map $\varepsilon_A : A^{\leq k} \rightarrow A$ sends a sequence $[a_1, \dots, a_n]$ to a_n .
- The key question is: how do we lift the σ -relations on A to $G_k A$? Given e.g. a binary relation R , we define $R^{G_k A}$ to be the set of pairs of sequences (s, t) such that

- $s \sqsubseteq t$ or $t \sqsubseteq s$ (in the prefix order)
- $R^A(\varepsilon_A(s), \varepsilon_A(t))$.

This generalizes straightforwardly to n -ary relations. Given an n -ary relation R , $R^{G_k A}$ is the set of tuples of sequences (s_1, \dots, s_n) such that for all i, j , s_i is comparable with s_j in the prefix order, and $R^A(\varepsilon_A(s_1), \dots, \varepsilon_A(s_n))$.

⁴A chain in a poset P is a linearly ordered subset $C \subseteq P$.

⁵This is cofree since, as we will see, it has a couniversal mapping property.

We verify the couniversal property:

$$\begin{array}{ccc}
 \text{Arboreal category} & & \text{Extensional category} \\
 G_k A & & U_k G_k A \xrightarrow{\varepsilon_A} A \\
 \uparrow \hat{f} & & \uparrow U_k \hat{f} \quad \nearrow f \\
 \mathcal{T} & & U_k \mathcal{T}
 \end{array}$$

This says that any σ -homomorphism $f : U_k \mathcal{T} \rightarrow A$ factors uniquely through the counit ε_A via a morphism $\hat{f} : \mathcal{T} \rightarrow G_k A$ of forest-ordered σ -structures. To see this, note that for each $x \in \mathcal{T}$, its predecessors in the forest order form a covering chain $x_1 \prec \cdots \prec x_n = x$, with $n \leq k$, and x_1 a root. We define $\hat{f}(x) = [f(x_1), \dots, f(x_n)]$. It is easily verified that this gives a forest morphism. Moreover, by property (E), any instance $\vec{x} \in R^{\mathcal{T}}$ must occur along a chain in \mathcal{T} , and hence the image of this tuple under \hat{f} will be a tuple of sequences \vec{s} pairwise related in the prefix order. Since f is a homomorphism, by the definition of $R^{G_k A}$ we will have $\vec{s} \in R^{G_k A}$. Thus \hat{f} is a morphism in $\mathcal{R}^E(\sigma)$.

Consequences

Recall that the right adjoint is uniquely determined by the forgetful functor, and the comonad by the adjunction. So everything follows from the delineation of the arboreal category $\mathcal{R}_k^E(\sigma)$, and the evident forgetful functor $U_k : \mathcal{R}_k^E(\sigma) \rightarrow \text{Struct}(\sigma)$.

As we shall now see, this structure gives us directly:

- The Ehrenfeucht-Fraïssé game, which characterizes the elementary equivalences \equiv_k induced by the quantifier-rank indexed fragments of FOL.
- This leads directly to syntax-free, purely structural descriptions of the equivalences of structures induced by:
 - the full fragment of quantifier rank $\leq k$
 - the existential positive part of the fragment
 - the extension of the fragment with counting quantifiers.
- Note that model classes of formulas of quantifier rank k are exactly the unions of equivalence classes of \equiv_k . Thus first-order definability, graded by quantifier rank, is also captured in a syntax-free fashion.
- We also recover the important tree-depth combinatorial parameter from the coalgebras of the comonad.

Moreover, this template can be used to give similar analyses of a wealth of other logical and combinatorial notions.

The Ehrenfeucht-Fraïssé game ([20])

Model comparison games in general are especially important in finite model theory, where the compactness theorem is not available.

The EF-game between A and B . In the i 'th round, Spoiler moves by choosing an element in A or B ; Duplicator responds by choosing an element in the other structure. Duplicator wins after k rounds if the relation $\{(a_i, b_i) \mid 1 \leq i \leq k\}$ is a partial isomorphism.

In the *existential EF-game*, Spoiler only plays in A , and Duplicator responds in B . The winning condition is that the relation $\{(a_i, b_i) \mid 1 \leq i \leq k\}$ is a partial homomorphism.

The *Ehrenfeucht-Fraïssé Theorem* [22, 21] says that a winning strategy for Duplicator in the k -round EF game characterizes the equivalence $\equiv^{\mathcal{L}_k}$, where \mathcal{L}_k is the fragment of first-order logic of formulas with quantifier rank $\leq k$.

CoKleisli maps are strategies

Intuitively, an element of $A^{\leq k}$ represents a play in A of length $\leq k$. A coKleisli morphism $\mathbb{E}_k A \rightarrow B$ represents a Duplicator strategy for the existential Ehrenfeucht-Fraïssé game with k rounds:

- Spoiler plays only in A , and $b_i = f[a_1, \dots, a_i]$ represents Duplicator's response in B to the i 'th move by Spoiler.
- The winning condition for Duplicator in this game is that, after k rounds have been played, the induced relation $\{(a_i, b_i) \mid 1 \leq i \leq k\}$ is a partial homomorphism from A to B .

Theorem 3.1. The following are equivalent:

1. There is a homomorphism $\mathbb{E}_k A \rightarrow B$.
2. Duplicator has a winning strategy for the existential Ehrenfeucht-Fraïssé game with k rounds, played from A to B .
3. For every existential positive sentence φ with quantifier rank $\leq k$, $A \models \varphi \Rightarrow B \models \varphi$.

Open pathwise embeddings and back-and-forth equivalences

How do we capture back-and-forth equivalences, and hence the whole logic rather than just the existential positive part? The idea is to work in the arboreal category $\mathcal{R}_k^E(\sigma)$, where we have enough process structure to make game and bisimulation notions meaningful. The key notions are

- *paths*, i.e. objects of $\mathcal{R}_k^E(\sigma)$ in which the order is linear (so the forest is a single branch), and
- *path embeddings*, i.e. forest morphisms with paths as domains which are embeddings of relational structures.

These are special cases of notions which are axiomatised in the arboreal categories setting in great generality in [11].

A morphism $f: X \rightarrow Y$ in $\mathcal{R}_k^E(\sigma)$ is a *pathwise embedding* if, for all path embeddings $m: P \rightarrow X$, the composite $f \circ m$ is a path embedding.

To capture the “back” or p-morphism condition, we introduce a notion of *open map* [26] that, combined with the concept of pathwise embedding, will allow us to define an appropriate notion of bisimulation. A morphism $f: X \rightarrow Y$ in $\mathcal{R}_k^E(\sigma)$ is said to be *open* if it satisfies the following path-lifting property: Given any commutative square

$$\begin{array}{ccc} P & \xrightarrow{\quad} & Q \\ \downarrow & \dashrightarrow & \downarrow \\ X & \xrightarrow{f} & Y \end{array}$$

with P, Q paths, there exists a diagonal filler $Q \rightarrow X$ (i.e. an arrow $Q \rightarrow X$ making the two triangles commute). (Note that, if it exists, such a diagonal filler must be an embedding.)

If we read the embeddings from P as the current positions reached in X and Y , and the extension from P to Q as Spoiler playing a new move in Y , then the diagonal filler witnesses the ability of Duplicator to find a matching move in X .

A *bisimulation* between objects X, Y of $\mathcal{R}_k^E(\sigma)$ is a span of open pathwise embeddings

$$\begin{array}{ccc} & R & \\ & \swarrow & \searrow \\ X & & Y \end{array}$$

If such a bisimulation exists, we say that X and Y are *bisimilar*.

Theorem 3.2. $G_k A$ and $G_k B$ are bisimilar in $\mathcal{R}_k^E(\sigma)$ iff Duplicator has a winning strategy in the k -round Ehrenfeucht-Fraïssé game between A and B .

Note that we use the *resource category* $\mathcal{R}_k^E(\sigma)$ to study logical properties of objects of the *extensional category* $\text{Struct}(\sigma)$.

Connection to logic

We consider the following fragments of first-order logic:

- \mathcal{L}_k is the fragment of quantifier-rank $\leq k$.
- $\exists \mathcal{L}_k$ is the existential positive fragment of \mathcal{L}_k
- $\mathcal{L}_k(\#)$ is the extension of \mathcal{L}_k with counting quantifiers $\exists_{\geq n}$ (“there exists at least n ”).

We define three resource-indexed equivalences on structures A, B of $\text{Struct}(\sigma)$ using the resource categories $\mathcal{R}_k^E(\sigma)$:

- $A \rightleftarrows_k^G B$ iff there are morphisms $G_k A \rightarrow G_k B$ and $G_k B \rightarrow G_k A$.
Note that there need be no relationship between these morphisms.

- $A \leftrightarrow_k^G B$ iff $G_k A$ and $G_k B$ are bisimilar in $\mathcal{R}_k^E(\sigma)$.
- $A \cong_k^G B$ iff $G_k A$ and $G_k B$ are isomorphic in $\mathcal{R}_k^E(\sigma)$.

Theorem 3.3. For structures A and B :

- (1) $A \equiv \exists \mathcal{L}_k B \iff A \rightleftarrows_k^G B.$
- (2) $A \equiv \mathcal{L}_k B \iff A \leftrightarrow_k^G B.$
- (3) $A \equiv \mathcal{L}_k(\#) B \iff A \cong_k^G B.$

Coalgebra number and tree-depth

Another fundamental aspect of comonads is that they have an associated notion of *coalgebra*. To explain this, we need to recall the detailed definition of a comonad (G, ε, δ) on a category \mathcal{C} . As well as an endofunctor $G : \mathcal{C} \rightarrow \mathcal{C}$, we have a natural transformation $\varepsilon_A : GA \rightarrow A$ (the counit), and a natural transformation $\delta_A : GA \rightarrow GGA$ (the comultiplication). These are required to make the following diagrams commute:

$$\begin{array}{ccc}
 GA & \xrightarrow{\delta_A} & GGA \\
 \delta_A \downarrow & & \downarrow G\delta_A \\
 GGA & \xrightarrow{\delta_{GA}} & GGGA
 \end{array}
 \qquad
 \begin{array}{ccc}
 GA & \xrightarrow{\delta_A} & GGA \\
 \delta_A \downarrow & \searrow \text{id}_{GA} & \downarrow G\varepsilon_A \\
 GGA & \xrightarrow{\varepsilon_{GA}} & GA
 \end{array}$$

A coalgebra for a comonad (G, ε, δ) is a morphism $\alpha : A \rightarrow GA$ such that the following diagrams commute:

$$\begin{array}{ccc}
 A & \xrightarrow{\alpha} & GA \\
 \text{id}_A \searrow & & \downarrow \varepsilon_A \\
 & & A
 \end{array}
 \qquad
 \begin{array}{ccc}
 A & \xrightarrow{\alpha} & GA \\
 \alpha \downarrow & & \downarrow \delta_A \\
 GA & \xrightarrow{G\alpha} & G^2A
 \end{array}$$

Note that whereas a homomorphism $\mathbb{E}_k A \rightarrow B$ is generally *easier* to construct than a homomorphism $A \rightarrow B$ (fewer conditions to check for the domain), a homomorphism $A \rightarrow \mathbb{E}_k B$ will be *harder* (fewer available properties of the codomain to show that relations are preserved). We should only expect a coalgebra structure $\alpha : A \rightarrow \mathbb{E}_k A$ to exist when the k -local information on A is sufficient to determine the structure of A .

Our use of indexed comonads \mathbb{C}_k opens up a new kind of question for coalgebras. Given a structure A , we can ask: what is the least value of k such that a \mathbb{C}_k -coalgebra exists on A ? We call this the *coalgebra number* of A .

Theorem 3.4. For the Ehrenfeucht-Fraïssé comonad \mathbb{E}_k , the coalgebra number of A corresponds precisely to the *tree-depth* of A .

This follows directly from the following result.

Theorem 3.5. The category of coalgebras for \mathbb{E}_k is isomorphic to $\mathcal{R}_k^E(\sigma)$, the category of forest-ordered σ -structures.⁶

This says that \mathbb{E}_k -coalgebras on A correspond bijectively to witnesses for the tree-depth of A being $\leq k$.

4. Establishing the template

The rich connections we found in the example of the Ehrenfeucht-Fraïssé comonad are no accident. They instantiate a pattern which recurs throughout model theory, capturing in particular the resource-sensitive aspects important in finite model theory and descriptive complexity.

The pebbling comonad

To show how the same pattern occurs in a significantly different example, we look at the *pebbling comonad*.

Pebble games are similar but subtly different to EF-games. Spoiler moves by placing one from a fixed set of k pebbles on an element of A or B ; Duplicator responds by placing their matching pebble on an element of the other structure. Duplicator wins if after each round, the relation defined by the current positions of the pebbles is a partial isomorphism. Thus there is a “sliding window” on the structures, of fixed size. It is this size which bounds the resource, not the length of the play.

Whereas the k -round EF game corresponds to bounding the quantifier rank, k -pebble games correspond to bounding the number of *variables* which can be used in a formula.

Just as for EF-games, there is an existential-positive version, in which Spoiler only plays in A , and Duplicator responds in B .

The pebbling adjunction

We define a k -pebble forest-ordered σ -structure (A, \leq, p) to be a forest-ordered σ -structure (A, \leq) together with a pebbling function $p : A \rightarrow \mathbf{k}$, where $\mathbf{k} := \{1, \dots, k\}$. In addition to condition (E), it must also satisfy the following condition:

- (P) if a is adjacent to b in the Gaifman graph of A , and $a < b$ in the forest order, then for all $x \in (a, b]$, $p(a) \neq p(x)$.

Morphisms of these structures are morphisms of forest-ordered structures which additionally preserve the pebbling function.

This defines a category $\mathcal{R}_k^P(\sigma)$ (where k bounds the number of pebbles, rather than the height of the forest order), and there is an evident forgetful functor $V_k : \mathcal{R}_k^P(\sigma) \rightarrow \text{Struct}(\sigma)$.

Theorem 4.1. For each $k > 0$, the functor V_k has a right adjoint H_k .

The corresponding comonad is \mathbb{P}_k , the pebbling comonad.

⁶More specifically, the *comparison functor* from $\mathcal{R}_k^E(\sigma)$ to the category of coalgebras for \mathbb{E}_k is an isomorphism. Thus the adjunction $U_k \dashv G_k$ is *comonadic* [28].

The pebbling comonad

Given a structure A , the universe of $\mathbb{P}_k A$ is $(\mathbf{k} \times A)^+$, the set of finite non-empty sequences of moves (p, a) . Note this will be infinite even if A is finite. The counit map $\varepsilon_A : \mathbb{P}_k A \rightarrow A$ sends a sequence $[(p_1, a_1), \dots, (p_n, a_n)]$ to a_n .

To lift the relations on A to $\mathbb{P}_k A$ we have the following condition in addition to those for \mathbb{E}_k :

- If $s \sqsubseteq t$, then the pebble index of the last move in s does not appear in the suffix of s in t .

We can now run exactly the same script as for the Ehrenfeucht-Fraïssé case:

- We can define paths, pathwise embeddings, open maps, bisimilarity in $\mathcal{R}_k^P(\sigma)$ in exactly the same fashion as we did for $\mathcal{R}_k^E(\sigma)$.
- Hence we can define bisimulations between object of the extensional category $\text{Struct}(\sigma)$ using the resource category $\mathcal{R}_k^P(\sigma)$.
- We can define the equivalence relations $A \overset{H}{\rightleftarrows}_k B$, $A \overset{H}{\leftrightarrow}_k B$, $A \overset{H}{\cong}_k B$ with respect to $\mathcal{R}_k^P(\sigma)$.

We now take \mathcal{L}_k to be the k -variable fragment of first-order logic. $\exists \mathcal{L}_k$ is the existential-positive part of this fragment, $\mathcal{L}_k(\#)$ the extension of \mathcal{L}_k with counting quantifiers. With this notation, we get the same result as Theorem 3.3:

Theorem 4.2. For structures A and B :

- (1) $A \equiv^{\exists \mathcal{L}_k} B \iff A \overset{H}{\rightleftarrows}_k B$.
- (2) $A \equiv^{\mathcal{L}_k} B \iff A \overset{H}{\leftrightarrow}_k B$.
- (3) $A \equiv^{\mathcal{L}_k(\#)} B \iff A \overset{H}{\cong}_k B$.

Coalgebra number and tree-width

We can define the coalgebra number for the pebbling comonad exactly as done before for the Ehrenfeucht-Fraïssé comonad.

A slightly more subtle argument is needed to show:

Theorem 4.3. For the pebbling comonad \mathbb{P}_k , the coalgebra number of A corresponds precisely to the *tree-width* of A .⁷

The modal comonad

We briefly summarize another important example, for basic modal logic. This exemplifies both *variation*, since it has significantly different properties to the Ehrenfeucht-Fraïssé and pebbling cases, and also the underlying structural *uniformity*, since again we can run exactly the same script.

- In this case, the modal comonad \mathbb{M}_k corresponds to k -level *unravelling* of a Kripke structure.

⁷Strictly speaking, to treewidth + 1, since by convention 1 is subtracted from the “natural” measure in defining treewidth.

- Open pathwise embedding bisimulation recovers standard modal bisimulation.
- The logical equivalences are the modal versions of those previously considered:
 - full modal logic of depth $\leq k$,
 - the diamond-only positive fragment, and
 - *graded modal logic* [19] for the counting case.
- The coalgebra number in this case recovers the *property* of being a synchronization tree of height $\leq k$.
- The fact that it is a property rather than a structure in this case follows from the fact that this comonad is *idempotent*, and hence corresponds to a coreflective subcategory (the tree-structured models as a full sub-category of the category of all Kripke structures).

Examples galore

Subsequent work has made detailed studies of a considerable number of examples of game comonads corresponding to various logic fragments and corresponding model comparison games:

- pebbling games and finite variable fragments [7]
- Ehrenfeucht-Fraïssé games and quantifier rank fragments [13]
- basic modal logic and modal depth fragments [13]
- guarded quantifier fragments (atom, loose and clique guards) [9]
- generalized quantifier fragments [17]
- hybrid and bounded fragments [10]
- bounded conjunction finite variable logic (motivated by pathwidth) [29]

\mathbb{C}_k	Logic	$\kappa^{\mathbb{C}}$	$\rightarrow_k^{\mathbb{C}}$	$\leftrightarrow_k^{\mathbb{C}}$	$\approx_k^{\mathbb{C}}$
\mathbb{E}_k	FOL w/ $qr \leq k$	tree-depth	✓	✓	✓
\mathbb{P}_k	k -variable logic	treewidth +1	✓	✓	✓
\mathbb{M}_k	ML w/ $md \leq k$	sync. tree-depth	✓	✓	✓
\mathbb{G}_k^g	g -guarded logic w/ width $\leq k$	guarded treewidth	✓	✓	?
$\mathbb{H}_{n,k}$	k -variable logic w/ \mathbf{Q}_n -quantifiers	n -ary general treewidth	✓	✓	✓
\mathbb{PIR}_k	k -variable logic restricted- \wedge	pathwidth +1	✓	?	✓

Figure 1. Summary table

In each case, we have tight connections with logical fragments, and with combinatorial invariants, following exactly the same pattern we have already seen. We get direct descriptions of the coalgebras in terms of *comonadic forgetful functors*. These are important both for formulating bisimulation, and for the connection with combinatorial invariants.

The situation is summarized in Figure 1.

5. Arboreal categories

Arboreal categories and arboreal covers provide an axiomatic framework which can be instantiated to yield the wide range of examples described in the previous section. More broadly, they provide an axiomatic setting for model theory in general, with particular support for fine-grained analysis of resource-bounded aspects and combinatorial invariants.

In this section, we provide a brief overview of these notions, minimising technical details while emphasising the main concepts and results. For details, see [11].

The key notion in arboreal categories is that of *path*. This can be formulated in any category \mathcal{C} equipped with a reasonable factorization system.⁸ We refer to the monomorphisms in this factorization system as *embeddings*. If X is an object of \mathcal{C} , we write $\mathbb{S}X$ for the poset of subobjects determined by embeddings.

Definition 5.1. An object X of \mathcal{C} is called a *path* if the poset $\mathbb{S}X$ is a finite chain.

Paths will be denoted by P, Q, R, \dots . A *path embedding* is an embedding $P \hookrightarrow X$ whose domain is a path. Given any object X of \mathcal{C} , we let $\mathbb{P}X$ be the sub-poset of $\mathbb{S}X$ determined by the path embeddings.

We say that a category \mathcal{C} equipped with a factorization system is a *path category* if it has coproducts of families of paths, as well as satisfying an additional technical condition (“2 out of 3 condition”). It follows from the axioms that, for any object X in a path category, $\mathbb{P}X$ is a tree. (We allow the empty path, given by the initial object (the empty coproduct of paths), which forms the root.)

We can define open pathwise embeddings and bisimulation in any path category. To show that these notions have the expected properties, we need additional axioms, motivated by the fact that trees are the colimits of their branches and the embeddings between them. We say that an object X in a path category is *path-generated* if it is the colimit of its path embeddings.

An *arboreal category* is a path category in which every object is path-generated, and moreover paths are *connected* in the sense that any arrow $P \rightarrow \coprod_i P_i$ factors through one of the coproduct injections $P_i \rightarrow \coprod_i P_i$.

These axioms are sufficient to allow operational or dynamic notions such as games to be defined, and shown to be equivalent to the open pathwise embedding notion of bisimulation.

5.1. Back-and-forth games

Let \mathcal{C} be an arboreal category and let X, Y be any two objects of \mathcal{C} . We define a back-and-forth game $\mathcal{G}(X, Y)$ played by Spoiler and Duplicator on X and Y as follows. Positions in the game

⁸More precisely, a stable proper factorization system [11].

are pairs of (equivalence classes of) path embeddings $(m, n) \in \mathbb{P}X \times \mathbb{P}Y$. The winning relation $\mathcal{W}(X, Y) \subseteq \mathbb{P}X \times \mathbb{P}Y$ consists of the pairs (m, n) such that $\text{dom}(m) \cong \text{dom}(n)$.

Let $\perp_X: P \rightarrow X$ and $\perp_Y: Q \rightarrow Y$ be the roots of $\mathbb{P}X$ and $\mathbb{P}Y$, respectively. If $P \not\cong Q$, then Duplicator loses the game. Otherwise, the initial position is (\perp_X, \perp_Y) . At the start of each round, the position is specified by a pair $(m, n) \in \mathbb{P}X \times \mathbb{P}Y$, and the round proceeds as follows: Either Spoiler chooses some $m' \succ m$ and Duplicator must respond with some $n' \succ n$, or Spoiler chooses some $n'' \succ n$ and Duplicator must respond with $m'' \succ m$. Duplicator wins the round if they are able to respond and the new position is in $\mathcal{W}(X, Y)$. Duplicator wins the game if they have a strategy which is winning after t rounds, for all $t \geq 0$.

Example 5.2. It is shown in [13] that the abstract game $\mathcal{G}(X, Y)$ specialises, in the case of the arboreal categories $\mathcal{R}_k^E(\sigma)$, $\mathcal{R}_k^P(\sigma)$ and $\mathcal{R}_k^M(\sigma)$, to the usual k -round Ehrenfeucht-Fraïssé, k -pebble and k -round bisimulation games, respectively.

Theorem 5.3. Let X and Y be any two objects of an arboreal category such that the product $X \times Y$ exists. Then X and Y are bisimilar if and only if Duplicator has a winning strategy in the game $\mathcal{G}(X, Y)$.

5.2. Arboreal covers

We now return to the underlying motivation for the axiomatic development. Arboreal categories have a rich intrinsic process structure, which allows “dynamic” notions such as bisimulation and back-and-forth games, and resource notions such as the height of a tree, to be defined. A key idea is to relate these process notions to extensional, or “static” structures. In particular, much of finite model theory and descriptive complexity can be seen in this way.

In the general setting, we have an arboreal category \mathcal{C} , and another category \mathcal{E} , which we think of as the extensional category.

Definition 5.4. An *arboreal cover* of \mathcal{E} by \mathcal{C} is given by a comonadic adjunction

$$\mathcal{C} \begin{array}{c} \xrightarrow{L} \\ \perp \\ \xleftarrow{R} \end{array} \mathcal{E}.$$

As for any adjunction, this induces a comonad on \mathcal{E} . The comonad is (G, ε, δ) , where $G := LR$, ε is the counit of the adjunction, and $\delta_a: LRa \rightarrow LRLRa$ is given by $\delta_a := L(\eta_{Ra})$, with η the unit of the adjunction. The comonadicity condition states that the Eilenberg-Moore category of coalgebras for this comonad is isomorphic to \mathcal{C} . The idea is then that we can use the arboreal category \mathcal{C} , with its rich process structure and all the associated notions, to study the extensional category \mathcal{E} via the adjunction.

We now bring resources into the picture.

Definition 5.5. Let \mathcal{C} be an arboreal category, with full subcategory of paths \mathcal{C}_p . We say that \mathcal{C} is *resource-indexed* by a resource parameter k if for all $k \geq 0$, there is a full subcategory \mathcal{C}_p^k of \mathcal{C}_p closed

under embeddings⁹ with

$$\mathcal{C}_p^0 \hookrightarrow \mathcal{C}_p^1 \hookrightarrow \mathcal{C}_p^2 \hookrightarrow \dots$$

This induces a corresponding tower of full subcategories \mathcal{C}_k of \mathcal{C} , with the objects of \mathcal{C}_k those which are the colimit of their path embeddings with domain in \mathcal{C}_p^k .

Example 5.6. One resource parameter which is always available is to take \mathcal{C}_p^k to be given by those paths in \mathcal{C} whose chain of subobjects is of length $\leq k$. We can think of this as a temporal parameter, restricting the number of sequential steps, or the number of rounds in a game. For the Ehrenfeucht-Fraïssé and modal comonads, we recover \mathcal{R}_k^E and \mathcal{R}_k^M as described in sections 3 and 4, corresponding to k -round versions of the Ehrenfeucht-Fraïssé and modal bisimulation games respectively [13]. However, note that for the pebbling comonad, the relevant resource index is the number of pebbles, which is a memory restriction along a computation or play of a game. This leads to \mathcal{R}_k^P as described in section 4.

Proposition 5.7. Let $\{\mathcal{C}_k\}$ be a resource-indexed arboreal category. Then \mathcal{C}_k is an arboreal category for each k .

Definition 5.8. Let $\{\mathcal{C}_k\}$ be a resource-indexed arboreal category. We define a *resource-indexed arboreal cover* of \mathcal{E} by \mathcal{C} to be an indexed family of comonadic adjunctions

$$\begin{array}{ccc} & L_k & \\ \mathcal{C}_k & \xrightarrow{\quad} & \mathcal{E} \\ & R_k & \\ & \perp & \end{array}$$

with corresponding comonads G_k on \mathcal{E} .

Example 5.9. Our key examples arise by taking the extensional category \mathcal{E} to be $\text{Struct}(\sigma)$. For each $k \geq 0$, there are evident forgetful functors

$$L_k^E : \mathcal{R}_k^E(\sigma) \rightarrow \text{Struct}(\sigma), \quad L_k^P : \mathcal{R}_k^P(\sigma) \rightarrow \text{Struct}(\sigma)$$

which forget the forest order, and in the case of \mathcal{R}_k^P , also the pebbling function. These functors are both comonadic over $\text{Struct}(\sigma)$. The right adjoints build a forest over a structure A by forming sequences of elements over the universe A , suitably labelled and with the σ -relations interpreted so as to satisfy the conditions (E) and (P) respectively. In the modal logic case, the extensional category \mathcal{E} is the category $\text{Struct}_*(\sigma)$ of *pointed* σ -structures (A, a) with $a \in A$, and morphisms the σ -homomorphisms preserving the distinguished point. There is a forgetful functor

$$L_k^M : \mathcal{R}_k^M \rightarrow \text{Struct}_*(\sigma)$$

sending $(A, \leq) \in \mathcal{R}_k^M$ to (A, a) , where a is the unique root of (A, \leq) . This functor is comonadic and its right adjoint sends a pointed σ -structure (A, a) to the tree-ordered structure obtained by unravelling

⁹That is, for any embedding $P \hookrightarrow Q$ in \mathcal{C} with P, Q paths, if $Q \in \mathcal{C}_p^k$ then also $P \in \mathcal{C}_p^k$.

the structure A , starting from a , to depth k , and with the σ -relations interpreted so as to satisfy the condition (M), which is a suitable “local” version of condition (E) [13].

These constructions yield the comonads described concretely in [7, 13]. The sequences correspond to plays in the Ehrenfeucht-Fraïssé and pebbling games respectively. The work needed to show the correspondence between each of these games and the generic game $\mathcal{G}(X, Y)$ consists only in matching the winning conditions, which is quite straightforward; see [13].

We now show how resource-indexed arboreal covers can be used to define important notions on the extensional category. For a resource-indexed arboreal cover of \mathcal{E} by \mathcal{C} , with adjunctions $L_k \dashv R_k$ and comonads G_k , we define three resource-indexed relations on objects of \mathcal{E} , in terms of their images in \mathcal{C}_k under R_k .

Definition 5.10. Consider a resource-indexed arboreal cover of \mathcal{E} by \mathcal{C} , and any two objects a, b of \mathcal{E} . We define:

- $a \overset{R}{\rightleftarrows}_k b$ iff there are morphisms $R_k a \rightarrow R_k b$ and $R_k b \rightarrow R_k a$ in \mathcal{C}_k .
- $a \overset{R}{\leftrightarrow}_k b$ iff there is a bisimulation between $R_k a$ and $R_k b$ in \mathcal{C}_k .
- $a \overset{R}{\cong}_k b$ iff $R_k a$ and $R_k b$ are isomorphic in \mathcal{C}_k .

Proposition 5.11. Assume \mathcal{E} has binary products. For objects a and b of \mathcal{E} , $a \overset{R}{\leftrightarrow}_k b$ iff Duplicator has a winning strategy in the game $\mathcal{G}(R_k a, R_k b)$.

Once the correspondence between the abstract games $\mathcal{G}(R_k a, R_k b)$ (and variants for the existential positive and counting quantifier cases) and the standard model comparison games has been established, we obtain the connections to logical equivalences as in Theorems 3.3 and 4.2 as corollaries.

6. Homomorphism preservation theorems: towards axiomatic resource-indexed model theory

As we have seen, comonadic semantics has now been given for a number of important fragments of first-order logic, including the quantifier rank fragments, the finite variable fragments, the modal fragment, and guarded fragments. In the landscape emerging from these constructions, some salient properties have come to the fore. These are properties which a comonad, arising from an *arboreal cover* in the sense of [11], may or may not have:

- The comonad may be *idempotent*, meaning that the comultiplication is a natural isomorphism. Idempotent comonads correspond to *coreflective subcategories*, which form the Eilenberg-Moore categories of these comonads. The modal comonads \mathbb{M}_k are idempotent. The corresponding coreflective subcategories are of those modal structures which are tree-models to depth k [13].
- The comonad \mathcal{C} may satisfy the following property: for each structure A , $\mathcal{C}A \overset{\mathcal{C}}{\leftrightarrow} A$, where $\overset{\mathcal{C}}{\leftrightarrow}$ is the bisimulation equivalence defined for \mathcal{C} in terms of open pathwise embeddings. We shall call this the *bisimilar companion* property. Note that an idempotent comonad, such as \mathbb{M}_k , will

automatically have this property. The guarded comonads \mathbb{G}_k from [9] are not idempotent, but have the bisimilar companion property, which is thus strictly weaker.

- Finally, the comonads \mathbb{E}_k and \mathbb{P}_k have neither of the above properties. Unlike the modal and guarded fragments, the quantifier rank and finite variable fragments cover the whole of first-order logic, so we call these comonads *expressive*.

Thus we have a strict hierarchy of comonads in the arboreal categories framework:

$$\text{idempotent} \Rightarrow \text{bisimilar companions} \Rightarrow \text{arboreal.}$$

This hierarchy is correlated with tractability: the modal and guarded fragments are decidable, and have the tree-model property [35, 23], while the expressive fragments do not. We can regard these observations as a first step towards using structural properties of comonadic semantics to classify logic fragments and their expressive power.

6.1. Model classes

Suppose we are given a resourced-indexed arboreal cover between an extensional category \mathcal{E} and an arboreal category \mathcal{C} . As explained in section 5, this induces resource-indexed bisimulation relations $\leftrightarrow_k^{\mathcal{C}}$ on \mathcal{E} . In the examples, \mathcal{E} is typically the category $\text{Struct}(\sigma)$ of relational structures, and the relations $\leftrightarrow_k^{\mathcal{C}}$ coincide with equivalence in logic fragments \mathcal{L}_k . There is also the homomorphism relation $\rightarrow^{\mathcal{C}}$, where $A \rightarrow^{\mathcal{C}} B$ iff $\mathcal{C}A \rightarrow B$, and the resource-indexed version $\rightarrow_k^{\mathcal{C}}$.

Given a formula $\phi \in \mathcal{L}_k$, we consider its “model-class” $\text{Mod}(\phi)$, *i.e.* the full subcategory of $\text{Struct}(\sigma)$ defined by the σ -structures A such that $A \models \phi$. In the next lemma we show that, if \mathcal{L}_k is finite¹⁰, these model-classes can be characterised in terms of the relations $\leftrightarrow_k^{\mathcal{C}}$. For simplicity, we restrict to the case where $\mathcal{E} = \text{Struct}(\sigma)$, but extensions are available, *e.g.* to pointed versions as for the modal fragment.

Lemma 6.1. Consider an arboreal adjunction between $\text{Struct}(\sigma)$ and an arboreal category \mathcal{C} , and let \mathcal{L} be a finite logic fragment such that $\leftrightarrow^{\mathcal{C}}$ coincides with $\equiv^{\mathcal{L}}$. The following statements are equivalent for any full subcategory \mathcal{D} of $\text{Struct}(\sigma)$:

1. \mathcal{D} is saturated under $\leftrightarrow^{\mathcal{C}}$, *i.e.* for all σ -structures A, B , if $A \in \mathcal{D}$ and $A \leftrightarrow^{\mathcal{C}} B$, then $B \in \mathcal{D}$.
2. $\mathcal{D} = \text{Mod}(\phi)$ for some $\phi \in \mathcal{L}$.

For applications to homomorphism preservation theorems, the following variant of the previous lemma will be useful, where $\exists \mathcal{L}$ denotes the existential positive fragment of \mathcal{L} .

Lemma 6.2. Consider an arboreal adjunction between $\text{Struct}(\sigma)$ and an arboreal category \mathcal{C} , and let \mathcal{L} be a finite logic fragment such that $\rightarrow^{\mathcal{C}}$ coincides with $\equiv^{\exists \mathcal{L}}$. The following statements are equivalent for any full subcategory \mathcal{D} of $\text{Struct}(\sigma)$:

1. \mathcal{D} is upwards closed with respect to $\rightarrow^{\mathcal{C}}$, *i.e.* for all σ -structures A, B , if $A \in \mathcal{D}$ and $A \rightarrow^{\mathcal{C}} B$, then $B \in \mathcal{D}$.

¹⁰Up to logical equivalence.

2. $\mathcal{D} = \text{Mod}(\psi)$ for some $\psi \in \exists\mathcal{L}$.

Remark 6.3. The previous lemmas can be extended to the case of *pointed* (or, more generally, n -pointed) structures.

6.2. Homomorphism preservation theorems

Homomorphism preservation theorems relate the syntactic shape of a sentence with the semantic property of being preserved under homomorphisms between structures. We are especially interested in refinements of these results whereby resources (e.g. the quantifier-rank or number of variables in a formula) are preserved. Using the characterisations of model-classes given in Section 6.1, we can capture the content of homomorphism preservation theorems at an abstract level, as we now explain.

Fix an arbitrary resource-indexed arboreal adjunction between an extensional category \mathcal{E} and a resource-indexed arboreal category \mathcal{C} , with adjunctions

$$\mathcal{C}_k \begin{array}{c} \xrightarrow{L_k} \\ \perp \\ \xleftarrow{R_k} \end{array} \mathcal{E} \quad (1)$$

and associated comonads $G_k := L_k R_k$ on \mathcal{E} . Furthermore, let us say that a full subcategory \mathcal{D} of \mathcal{E} is *closed under morphisms* if, whenever there is an arrow $a \rightarrow b$ in \mathcal{E} with $a \in \mathcal{D}$, also $b \in \mathcal{D}$. Note that, when $\mathcal{E} = \text{Struct}(\sigma)$ and $\mathcal{D} = \text{Mod}(\phi)$, the category \mathcal{D} is closed under morphisms precisely when ϕ is preserved under homomorphisms between σ -structures.

Consider the following statement:

(HP) For any full subcategory \mathcal{D} of \mathcal{E} saturated under $\leftrightarrow_k^{\mathcal{C}}$, \mathcal{D} is closed under morphisms iff it is upwards closed with respect to $\rightarrow_k^{\mathcal{C}}$.

E.g., for the Ehrenfeucht-Fraïssé resource-indexed arboreal adjunction, the statement (HP) is precisely Rossman's equirank homomorphism preservation theorem [34].

Replacing the relation $\leftrightarrow_k^{\mathcal{C}}$ with the equivalence relation $\cong_k^{\mathcal{C}}$, we obtain a strengthening of (HP), namely:

(HP[#]) For any full subcategory \mathcal{D} of \mathcal{E} saturated under $\cong_k^{\mathcal{C}}$, \mathcal{D} is closed under morphisms iff it is upwards closed with respect to $\rightarrow_k^{\mathcal{C}}$.

Remark 6.4. Any full subcategory of \mathcal{E} that is upwards closed with respect to $\rightarrow_k^{\mathcal{C}}$ is closed under morphisms. Hence, the right-to-left implications in (HP) and (HP[#]) are always satisfied.

In Section 6.3 we define the *bisimilar companion property* for resource-indexed arboreal adjunctions, and show that this property entails (HP). This leads to homomorphism preservation theorems for guarded logics.

In Section 6.4 we show that a strengthening of the bisimilar companion property, namely idempotency, implies (HP[#]). We thus obtain equi-depth homomorphism preservation theorems for *graded* modal formulas, which hold uniformly at the general and finite levels.

Let us point out that the bisimilar companion property does not hold for the Ehrenfeucht-Fraïssé resource-indexed arboreal adjunction, so we cannot deduce Rossman’s equirank homomorphism preservation theorem in this way. However, in Section 6.5 we explain how to “force” this property to obtain a refinement of Rossman’s result.

6.3. The bisimilar companion property

Definition 6.5. A resource-indexed arboreal adjunction between \mathcal{E} and \mathcal{C} , with induced comonads G_k , has the *bisimilar companion property* if $a \leftrightarrow_k^{\mathcal{C}} G_k a$ for all $a \in \mathcal{E}$ and $k \geq 0$.

Proposition 6.6. Consider any resource-indexed arboreal adjunction between \mathcal{E} and \mathcal{C} with the bisimilar companion property. Then (HP) holds.

Proof:

Let \mathcal{D} be a full subcategory of \mathcal{E} closed under morphisms and saturated under $\leftrightarrow_k^{\mathcal{C}}$. If $a \rightarrow_k^{\mathcal{C}} b$, there is a morphism $G_k a \rightarrow b$ and so, using the bisimilar companion property,

$$a \leftrightarrow_k^{\mathcal{C}} G_k a \rightarrow b.$$

Therefore, if $a \in \mathcal{D}$ then also $b \in \mathcal{D}$. That is, \mathcal{D} is upwards closed with respect to $\rightarrow_k^{\mathcal{C}}$.

The converse direction follows from Remark 6.4. □

We can now use Proposition 6.6 to obtain a homomorphism preservation theorem for guarded logics.

6.4. Idempotency

Definition 6.7. A resource-indexed arboreal adjunction between \mathcal{E} and \mathcal{C} is *idempotent* if so are the induced comonads G_k , i.e. $G_k G_k a \cong G_k a$ for all $a \in \mathcal{E}$ and $k \geq 0$.

Proposition 6.8. Consider any idempotent resource-indexed arboreal adjunction between \mathcal{E} and \mathcal{C} . Then (HP[#]) holds.

Proof:

Recall that G_k is idempotent if, and only if, ηR_k is a natural isomorphism, where η is the unit of the adjunction $L_k \dashv R_k$. In particular, for any $a \in \mathcal{E}$, the component of ηR_k at a yields an isomorphism $R_k a \cong R_k G_k a$. Hence, $a \cong_k^{\mathcal{C}} G_k a$ for all $a \in \mathcal{E}$. Reasoning as in the proof of Proposition 6.6, it is easy to see that (HP[#]) holds. □

Graded Modal Logic Let σ be an arbitrary finite vocabulary consisting of relation symbols of arity at most 2. We consider the resource-indexed relations \cong_k and \rightarrow_k on the category $\text{Struct}_*(\sigma)$ of pointed Kripke structures induced by the modal resource-indexed arboreal cover, cf. section 4.

Proposition 6.8 entails the following equidepth homomorphism preservation theorem for graded modal formulas:

Theorem 6.9. The following statements are equivalent for any graded modal formula $\phi \in \text{ML}_k(\#)$:

1. ϕ is preserved under homomorphisms between pointed Kripke structures.
2. ϕ is logically equivalent to an existential positive modal formula $\psi \in \exists\text{ML}_k$.

Proof:

Fix an arbitrary formula $\phi \in \text{ML}_k(\#)$. By [12, Proposition 15] and [19, Proposition 3.6], there is an inclusion

$$\cong_k \subseteq \equiv^{\text{ML}_k(\#)}$$

and so $\text{Mod}(\phi)$ is saturated under \cong_k . As the comonad \mathbb{M}_k is idempotent, Proposition 6.8 entails that $\text{Mod}(\phi)$ is closed under morphisms if, and only if, it is upwards closed with respect to \rightarrow_k . Thus, the statement follows from Lemma 6.2 and Remark 6.3, setting $\mathcal{L} = \text{ML}_k$. \square

Remark 6.10. Forgetting about both graded modalities and modal depth, Theorem 6.9 implies that a modal formula is preserved under homomorphisms if, and only if, it is equivalent to an existential positive modal formula. This improves the well known result that a modal formula is preserved under simulations precisely when it is equivalent to an existential positive modal formula (see e.g. [15, Theorem 2.78]).

Since the modal comonad \mathbb{M}_k restricts to finite pointed Kripke structures, the resource-indexed arboreal cover given by \mathcal{R}_k^M restricts to the full subcategory of $\text{Struct}_*(\sigma)$ defined by the finite Kripke structures. Reasoning as above, we obtain a variant of Theorem 6.9 for finite structures:

Theorem 6.11. The following statements are equivalent for any graded modal formula $\phi \in \text{ML}_k(\#)$:

1. ϕ is preserved under homomorphisms between finite pointed Kripke structures.
2. ϕ is logically equivalent over finite pointed Kripke structures to an existential positive modal formula $\psi \in \exists\text{ML}_k$.

6.5. Forcing the bisimilar companion property

When a resource-indexed arboreal adjunction does not have the bisimilar companion property, Proposition 6.6 cannot be applied to obtain an equi-resource homomorphism preservation theorem. This is the case, e.g., for Rossman's equirank homomorphism preservation theorem:

Theorem 6.12. ([34])

The following statements are equivalent for any first-order sentence ϕ with quantifier-rank at most k :

1. ϕ is preserved under homomorphisms between σ -structures.
2. ϕ is equivalent to an existential positive sentence with quantifier rank at most k .

Rossmann's strategy is to construct a bisimilar companion for a structure A . A categorical view of Rossmann's construction as a colimit in a coslice category is outlined in [3]. A "semi-axiomatic" proof of Rossmann's theorem, in the arboreal categories framework but where the extensional category \mathcal{E} is fixed to be $\text{Struct}(\sigma)$, is given in [1]. In [2], a purely axiomatic proof of Rossmann's theorem is given in the arboreal categories setting.

In [32], the following refinement of Rossmann's equirank homomorphism preservation theorem is obtained:

Theorem 6.13. (Equirank-variable HPT)

The following statements are equivalent for any first-order sentence ϕ in n variables with quantifier-rank at most k , where $n \geq k - 2$:

1. ϕ is preserved under homomorphisms between σ -structures.
2. ϕ is equivalent to an existential positive sentence in n variables with quantifier rank at most k .

7. Current developments

The research program we have described in this paper is in a very lively and active state. We shall briefly describe a few of the current lines of research.

- Homomorphism preservation theorems are only one example of semantic characterizations of various types of formulas. Many of the classical results of this kind rely crucially on compactness, and fail for finite models. However, there are results which not only hold in the finite as well as the infinite, but admit *uniform proofs* which apply to both cases. One well-known example is the van Benthem-Rosen theorem, which characterises the modal fragment in terms of bisimulation invariance. A uniform proof of this result was given by Otto [31]. A detailed version, extracting an important general result on Ehrenfeucht-Fraïssé games (the Workspace Lemma), and clarifying the use of the bisimilar companions property, is given in [2]. In [10], game comonads are developed for hybrid logic and the bounded fragment.¹¹ This logic lacks the bisimilar companion property, but it still has a local character. In [10], a uniform proof, covering both finite and infinite cases, is given for a semantic characterization of this fragment. This uses the Workspace Lemma, but replaces the appeal to bisimilar companions made in the van Benthem Rosen theorem by a different argument, again using properties of model comparison games.

These uniform arguments point towards the possibility of a *resource-sensitive model theory without compactness*.

- A classic result of Lovász [27] says that two finite structures are isomorphic if and only if they admit the same number of homomorphisms from all finite structures. This result has been extended in many different ways. In one type of generalisation, isomorphisms are replaced by

¹¹In the bounded fragment, quantifiers are relativised to atomic predicates, but by contrast with the guarded fragment, variables in the guarded formulas do not have to appear in the guards. This makes the logic undecidable.

a specified equivalence relation \simeq on finite structures, and the class of all finite structures by a specified class of finite structures Δ . Then a typical Lovász-type theorem expresses that, for finite structures A, B ,

$$A \simeq B \iff \text{hom}(C, A) \cong \text{hom}(C, B) \quad \text{for every } C \in \Delta.$$

In [18], the game comonads framework is used to give an axiomatic account, leading to several new Lovász-type theorems. This is further extended in [33] using Joyal's theory of polyadic spaces, leading to results which apply also in infinite cases.

- The comonadic analysis of both combinatorial parameters such as treewidth, and Lovász-type theorems, leads to the question: when are such characterisations possible? In [8], very general conditions are identified under which there is a comonad which classifies a given quantitative parameter. The same analysis leads to general Lovász-type results. The conditions cover a huge range of concrete examples. The construction which establishes the result is a Kan extension leading to a discrete density comonad construction, which is weakly initial among solutions to the problem.

A question for further work is to relate this construction to the arboreal categories framework, with the aim of establishing general conditions for transfer results, which relate different parameters or counting theorems by varying the comonad.

- In [25], an axiomatic account is given of Feferman-Vaught-Mostowski composition theorems within the game comonad framework, parameterized by the model comparison game. Compositionality results for the corresponding logic, and its positive existential and counting quantifier variants, are obtained in a uniform manner. The game comonads are extended to a monadic second order version, and used to obtain an abstract version of Courcelle's theorem, which can be specialised to yield the standard concrete versions of this result.
- In [16, 4], new cohomological approximation algorithms for constraint satisfaction and structure isomorphism are introduced, and are shown to properly extend standard local consistency and Weisfeiler-Leman algorithms. These results make use of the sheaf-theoretic and cohomological methods originally developed for quantum contextuality in [6, 5].

References

- [1] Abramsky S, and Reggio L. Arboreal categories: An axiomatic theory of resources. Preprint available at <https://arxiv.org/abs/2102.08109>, 2022.
- [2] Abramsky S, and Reggio L. Arboreal Categories and Preservation Theorems. Preprint, 2022.
- [3] Abramsky S. Whither semantics? *Theoretical Computer Science*, 2020. 807:3–14. doi:10.48550/arXiv.2010.13328.
- [4] Abramsky S. Notes on cohomological width and presheaf representations, 2022. Technical Report.
- [5] Abramsky S, Barbosa RS, Kishida K, Lal R, and Mansfield S. Contextuality, cohomology and paradox. In Stephan Kreutzer, editor, *24th EACSL Annual Conference on Computer Science Logic, CSL 2015*,

September 7-10, 2015, Berlin, Germany, volume 41 of *LIPICs*, pages 211–228. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015. doi:10.4230/LIPICs.CSL.2015.211.

- [6] Abramsky S. and Brandenburger A. The sheaf-theoretic structure of non-locality and contextuality. *New Journal of Physics*, 2011. 13(11):113036. doi:10.1088/1367-2630/13/11/113036.
- [7] Samson Abramsky, Anuj Dawar, and Pengming Wang. The pebbling comonad in finite model theory. In *Logic in Computer Science (LICS), 2017 32nd Annual ACM/IEEE Symposium on*, IEEE, 2017 pages 1–12. doi:10.48550/arXiv.1704.05124.
- [8] Abramsky S, Jakl T, and Paine T. Discrete density comonads and graph parameters. Submitted, 2022. doi:10.48550/arXiv.2205.06589.
- [9] Abramsky S, and Marsden D. Comonadic semantics for guarded fragments. In *2021 36th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, IEEE, 2021 pages 1–13. doi:10.1109/LICS52264.2021.9470594.
- [10] Abramsky S, and Marsden D. Comonadic semantics for hybrid logic and bounded fragments. 2021. *arXiv preprint arXiv:2110.09844*.
- [11] Abramsky S, and Reggio L. Arboreal categories and resources. In *48th International Colloquium on Automata, Languages, and Programming (ICALP 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICs.ICALP.2021.115.
- [12] Abramsky S, and Shah N. Relating structure and power: Comonadic semantics for computational resources. In *27th EACSL Annual Conference on Computer Science Logic, CSL 2018, September 4-7, 2018, Birmingham, UK*, 2018 pages 2:1–2:17. doi:10.1007/978-3-030-00389-0_1.
- [13] Abramsky S, and Shah N. Relating structure and power: Comonadic semantics for computational resources. *Journal of Logic and Computation*, 2021. 31(6):1390–1428. doi:10.1093/logcom/exab048.
- [14] Beke T, and Rosický J. Abstract elementary classes and accessible categories. *Annals of Pure and Applied Logic*, 2012. 163(12):2008–2017. doi:10.1016/j.apal.2012.06.003.
- [15] Blackburn P, de Rijke M, and Venema Y. *Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2002. ISBN-13:978-1881526964, 10:9781881526964.
- [16] Conghaile AÓ. Cohomological k -consistency, 2022. Technical Report.
- [17] Conghaile AÓ, and Dawar A. Game comonads and generalised quantifiers. In Christel Baier and Jean Goubault-Larrecq, editors, *29th EACSL Annual Conference on Computer Science Logic (CSL 2021)*, 2021 pages 16:1–16:17. doi:10.48550/arXiv.2006.16039.
- [18] Dawar A, Jakl T, and Reggio L. Lovász-type theorems and game comonads. In *2021 36th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, 2021 pages 1–13. doi:10.1109/LICS52264.2021.9470609.
- [19] de Rijke M. A note on graded modal logic. *Studia Logica*, 2000. 64(2):271–283. doi:10.1023/A:1005245900406.
- [20] Ebbinghaus HD, and Flum J. *Finite model theory*. Springer Science & Business Media, 2005.
- [21] Ehrenfeucht A. An application of games to the completeness problem for formalized theories. *Fund. Math*, 1961. 49(2):129–141. ISSN-0016-2736.
- [22] Fraïssé R. Sur quelques classifications des systèmes de relations. In *Publications Scientifiques, Série A*, Université d’Alger 1954 pages 35–182.

- [23] Grädel E. Why are modal logics so robustly decidable? *Bull. EATCS*, 1999. 68:90–103.
- [24] Hodges W. *A shorter model theory*. Cambridge university press, 1997. ISBN-9780521587136.
- [25] Jakl T, Marsden D, and Shah N. A game comonadic account of courcelle and feferman-vaught-mostowski theorems. 2022. *arXiv preprint arXiv:2205.05387*.
- [26] Joyal A, Nielson M, and Winskel G. Bisimulation and open maps. In *[1993] Proceedings Eighth Annual IEEE Symposium on Logic in Computer Science*, IEEE 1993 pages 418–427. doi:10.1109/LICS.1993.287566.
- [27] Lovász L. Operations with structures. *Acta Math. Acad. Sci. Hungar.*, 1967. 18:321–328. doi:10.1007/BF02280291.
- [28] Lane SM. *Categories for the working mathematician*, volume 5 of *Graduate Texts in Mathematics*. Springer, 2013. doi:10.1007/978-1-4757-4721-8.
- [29] Montacute Y, and Shah N. The pebble-relation comonad in finite model theory. 2021. *arXiv preprint arXiv:2110.08196*.
- [30] Nešetřil J, and De Mendez PO. Tree-depth, subgraph coloring and homomorphism bounds. *European Journal of Combinatorics*, 2006. 27(6):1022–1041. doi:10.1016/j.ejc.2005.01.010.
- [31] Otto M. Bisimulation invariance and finite models. In *Logic Colloquium '02*, volume 27 of *Lect. Notes Log.*, Assoc. Symbol. Logic, La Jolla, CA, 2006 pages 276–298. doi:10.1017/9781316755723.013.
- [32] Paine T. A pebbling comonad for finite rank and variable logic, and an application to the equirank-variable homomorphism preservation theorem. *Electronic Notes in Theoretical Computer Science*, 2020. 352:191–209. doi:10.1016/j.entcs.2020.09.010.
- [33] Reggio L. Polyadic sets and homomorphism counting. 2021. *arXiv preprint arXiv:2110.11061v3*.
- [34] Rossman B. Homomorphism Preservation Theorems. *JACM*, 2008. 55(3):15.
- [35] Vardi MY. Why is modal logic so robustly decidable? Technical Report TR-97-274, Rice University, 1997. URL <https://hdl.handle.net/1911/96465>.