Locally Computable Structures

Russell G. Miller[1] *

Department of Mathematics
Queens College – C.U.N.Y.
65-30 Kissena Blvd.
Flushing, New York 11367 U.S.A.

Doctoral Program in Computer Science
The Graduate Center of C.U.N.Y.
365 Fifth Avenue
New York, New York 10016 U.S.A.

Russell.Miller@qc.cuny.edu

Abstract. We introduce the notion of a *locally computable structure*, a natural way of generalizing the notions of computable model theory to uncountable structures $\mathcal{S}$ by presenting the finitely generated substructures of $\mathcal{S}$ effectively. Our discussion emphasizes definitions and examples, but does prove two significant results. First, our notion of m-extensional local computability of $\mathcal{S}$ ensures that the Σ_n-theory of $\mathcal{S}$ will be Σ_n for all $n \leq m+1$. Second, our notion of perfect local computability is equivalent (for countable structures) to the classic definition of computable presentability.

Keywords: computability, computable model theory, extensional, locally computable, perfectly locally computable.

1 Introduction

Turing computability has always been restricted to maps on countable sets. This restriction is inherent in the nature of a Turing machine: a computation is performed in a finite length of time, so that even if the available input was a countable binary sequence, only a finite initial segment of that sequence was actually used in the computation. Thus only that finite segment was relevant to the computation. To be sure, there are approaches that have defined natural notions of computable functions on the real numbers. These include the bitmap model, detailed in [2], and the Blum-Shub-Smale model (see [1]). These are elegant in several respects, but also omit certain basic functions, and moreover, each was built with the real numbers (viewed either as 2^ω or as the real line) specifically in mind, rather than arbitrary uncountable structures.

Nevertheless, mathematicians are hardly daunted by the prospect of doing actual computations on $\mathbb{R}$. When faced with a real number whose binary expansion is not immediately accessible, they do not flinch; they simply call that real "x." All field operations can then be performed with ease within the subfield of $\mathbb{R}$ generated by x; the mathematician only needs to know whether x is algebraic or transcendental, and, in the former case, what its minimal polynomial over $\mathbb{Q}$ is. Similar devices handle the situation of several unknown reals at once. The binary expansions of these reals are not required for the algebraic operations.

In this paper we formalize this process, and generalize it to arbitrary structures in finite languages. Starting with the notion of a *computable model*, which is entirely in keeping with Turing's notion of computability, we will view the real numbers and other fields as locally computable structures. No claim is made that operations on the reals can be performed globally, but we develop a definition in which countable objects are used to describe all finitely generated substructures of a (potentially uncountable) structure $\mathcal{S}$. Then the *local computability* of $\mathcal{S}$ is determined by the computability of the countable objects. In cases such as the field $\mathbb{R}$, where every finitely generated substructure is computably presentable, we will say that we have a *computable cover* of the structure. Indeed, for $\mathbb{R}$, a single algorithm can list out all elements of this cover.

The term "cover" is borrowed from the definition of a manifold, and the analogy, while imprecise, can be useful for intuitions about our definitions. For instance, for a topological space M, being a manifold does not just require the existence of a cover by open subsets of $\mathbb{R}^n$, but also that the charts within M given by the cover should fit together in a nice way: the transition functions between open subsets of $\mathbb{R}^n$, defined whenever two charts in M intersect, should be continuous (or differentiable, or C^∞, depending on how nicely we wish the manifold to behave). In short, it is not sufficient just to describe the local behavior of M; one must ensure that where the descriptions overlap, they agree with one another in a reasonable way.

For us, it will certainly be true that finitely generated substructures of a structure $\mathcal{S}$ can overlap. Therefore, our description of finitely generated substructures of $\mathcal{S}$ will include an account of which such substructures extend to others. Since any two finitely generated substructures of $\mathcal{S}$ lie within a single larger finitely generated substructure, it is sufficient for our purposes to consider the question of extensions among them. Topological notions do not fit our setting very well, but embeddings among finitely generated computable structures are themselves inherently computable, since they are determined by their values on the generators of the domain. (This is our main reason for considering only finitely generated substructures of $\mathcal{S}$, in fact, rather than all countable substructures.) In order for a structure to be called *locally computable*, we will require not only that the finitely generated substructures be computably presentable in a uniform way, but also that there be a computable enumeration of the embeddings among them corresponding to extensions in the structure $\mathcal{S}$. Various strengthenings of this requirement will allow us to prove stronger theorems about certain of the structures.

The technical content of this paper is not especially high, and for reasons of space we have emphasized our definitions and new concepts, and omitted many of the proofs. The most involved proof, for Theorem 2, is included in Appendix A. When computability-theoretic notions arise, we refer the reader to [4], the standard source, for notation and definitions. A good overview of the field of computable model theory is given in [3].

2 Local Computability

Let T be a $\forall$-axiomatizable theory in a language with n symbols. (If T is finitely axiomatizable but not $\forall$-axiomatizable, we can Skolemize to give it a set of $\forall$-axioms, while keeping the language finite.) We first consider simple covers of a model $\mathcal{S}$ of T. These describe only the finitely generated substructures of $\mathcal{S}$, with no attention paid to any relations between those substructures.

Definition 1. *A* simple cover *of $\mathcal{S}$ is a (finite or countable) collection $\mathfrak{A}$ of finitely generated models $\mathcal{A}_0, \mathcal{A}_1, \ldots$ of T, such that:*

- *every finitely generated substructure of $\mathcal{S}$ is isomorphic to some $\mathcal{A}_i \in \mathfrak{A}$; and*
- *every $\mathcal{A}_i \in \mathfrak{A}$ embeds isomorphically into $\mathcal{S}$.*

A simple cover $\mathfrak{A}$ is computable *if every $\mathcal{A}_i \in \mathfrak{A}$ is a computable structure whose domain is an initial segment of ω. $\mathfrak{A}$ is* uniformly computable *if the sequence $\langle (\mathcal{A}_i, \overline{a}_i) \rangle_{i \in \omega}$ can be given uniformly: there must exist a computable function which, on input i, outputs a tuple of elements $\langle e_1, \ldots, e_n, \langle a_0, \ldots, a_{k_i} \rangle \rangle \in \omega^n \times \mathcal{A}_i^{<\omega}$ such that $\{a_0, \ldots, a_{k_i}\}$ generates $\mathcal{A}_i$ and φ_{e_j} computes the j-th function, relation, or constant in $\mathcal{A}_i$.*

The intention is that $\mathcal{S}$ itself should not be finitely generated, of course, although the definition is still valid in this case. Indeed, $\mathcal{S}$ is not at all required to be countable, since a single $\mathcal{A}_i$ may be isomorphic to many substructures of $\mathcal{S}$. For countable structures $\mathcal{S}$, a related notion is Fraïssé's concept of the *age* of $\mathcal{S}$, i.e. the set of all finitely generated substructures of $\mathcal{S}$. We note that all elements of $\mathfrak{A}$ will be models of T; this was the reason for which we demanded that T be $\forall$-axiomatizable.

Notice that the definition requires that the generators of $\mathcal{A}_i$ be given as a tuple $\langle a_0, \ldots, a_{k_i} \rangle$, so that k_i is computable uniformly in i and we know how many values from $\mathcal{A}_j$ are needed to define an embedding in $I_{ij}^{\mathfrak{A}}$. (In the language of [4], the definition requires the *canonical index* for the set $\{a_0, \ldots, a_{k_i}\}$.)

As an example, it is straightforward to show that the best-known uncountable structure in mathematics is locally computable. We omit the proof, since most of its ingredients have long been established.

Proposition 1. *The field $\mathcal{R} = (\mathbb{R}, +, \cdot, -, r, 0, 1)$ of real numbers is locally computable.* □

Notice that we have added the operations of negation and inversion (r, for *reciprocal*) to the usual language of fields, in order to get a Π_1 axiom set.

It is also useful to see a negative example. Although the real numbers form a locally computable field, Adding the usual $<$ relation to the field $\mathcal{R}$ of Proposition 1 destroys local computability.

Proposition 2. *The ordered field* $(\mathcal{R}, <)$ *of real numbers, with* $\mathcal{R}$ *as in Proposition 1, has no computable simple cover, uniform or otherwise.*

Proof. Let b be any noncomputable real number. We claim that the ordered subfield $\mathcal{B}$ of $\mathcal{R}$ generated by b has no computable presentation. Suppose $\mathcal{A}$ were a computable presentation of $\mathcal{B}$, with $a \in \mathcal{A}$ representing b. Then just from knowing the additive and multiplicative identity elements in $\mathcal{A}$, we could compute the representation in $\mathcal{A}$ of any rational number $\frac{p}{q}$. But then we could compute the n-th bit of the binary expansion of a, uniformly in n, just by using the computable relation $<$ in $\mathcal{A}$ to compare a to various dyadic rationals. But this is also the n-th bit of the binary expansion of b, which was assumed to be noncomputable. Therefore no such $\mathcal{A}$ can exist. $\qquad\square$

We will be concerned mainly with the full definition of a cover, in which we also describe how the substructures of $\mathcal{S}$ fit together.

Definition 2. *An embedding* $f : \mathcal{A}_i \hookrightarrow \mathcal{A}_j$ *lifts to the inclusion* $\mathcal{B} \subseteq \mathcal{C}$*, via isomorphisms* $\beta : \mathcal{A}_i \twoheadrightarrow \mathcal{B}$ *and* $\gamma : \mathcal{A}_j \twoheadrightarrow \mathcal{C}$*, if if the diagram below commutes:*

$$
\begin{array}{ccc}
\mathcal{B} & \xrightarrow{\ \subseteq\ } & \mathcal{C} \\
\beta \big\uparrow \cong & & \gamma \big\uparrow \cong \\
\mathcal{A}_i & \xrightarrow{\ f\ } & \mathcal{A}_j
\end{array}
\qquad\qquad \text{with } \gamma \circ f = \beta.
$$

A cover *of* $\mathcal{S}$ *consists of a simple cover* $\mathfrak{A} = \{\mathcal{A}_0, \mathcal{A}_1, \ldots\}$ *of* $\mathcal{S}$*, along with sets* $I_{ij}^{\mathfrak{A}}$ *(for all* $\mathcal{A}_i, \mathcal{A}_j \in \mathfrak{A}$*) of injective homomorphisms* $f : \mathcal{A}_i \hookrightarrow \mathcal{A}_j$*, such that:*

- *for all finitely generated substructures* $\mathcal{B} \subseteq \mathcal{C}$ *of* $\mathcal{S}$*, there exist* $i, j \in \omega$ *and an* $f \in I_{ij}^{\mathfrak{A}}$ *which lifts to* $\mathcal{B} \subseteq \mathcal{C}$ *via some isomorphisms* $\beta : \mathcal{A}_i \twoheadrightarrow \mathcal{B}$ *and* $\gamma : \mathcal{A}_j \twoheadrightarrow \mathcal{C}$*; and*
- *for every* i *and* j*, every* $f \in I_{ij}^{\mathfrak{A}}$ *lifts to an inclusion* $\mathcal{B} \subseteq \mathcal{C}$ *in* $\mathcal{S}$ *via some isomorphisms* β *and* γ*.*

This cover is uniformly computable *if* $\mathfrak{A}$ *is a uniformly computable simple cover of* $\mathcal{S}$ *and there exists a c.e. set* W *such that for all* $i, j \in \omega$*,*

$$
I_{ij}^{\mathfrak{A}} = \{\varphi_e {\restriction} \mathcal{A}_i : \langle i, j, e \rangle \in W\}.
$$

A structure $\mathcal{B}$ *is* locally computable *if it has a uniformly computable cover.*

If $\mathfrak{A}$ is a computable simple cover, then every embedding of any $\mathcal{A}_i$ into any $\mathcal{A}_j$ is determined by its values on the generators of $\mathcal{A}_i$. Since $\mathcal{A}_i$ is finitely generated, all such embeddings are computable, and therefore it is reasonable to

call $\mathfrak{A}$ a computable cover without any further requirements on the sets $I_{ij}^{\mathfrak{A}}$. For a uniformly computable cover, on the other hand, the sets $I_{ij}^{\mathfrak{A}}$ will play a key role in our development of the subject, and it should be kept in mind that $I_{ij}^{\mathfrak{A}}$ need not contain every possible embedding of $\mathcal{A}_i$ into $\mathcal{A}_j$.

It is an easy exercise to see that the second condition of Definition 2 follows trivially from the definition of a simple cover, for any embedding $f : \mathcal{A}_i \hookrightarrow \mathcal{A}_j$. We include this second condition here because it is the dual of the first, and in the rest of our study of local computability, this duality between inclusion maps within $\mathcal{S}$ and embeddings among structures in $\mathfrak{A}$ will appear repeatedly.

Lemma 1. *A structure $\mathcal{S}$ has a uniformly computable cover (i.e. is locally computable) iff $\mathcal{S}$ has a uniformly computable simple cover.*

Proof. Given a uniformly computable simple cover $\mathfrak{A} = \{\mathcal{A}_0, \mathcal{A}_1, \ldots\}$, we adjoin all finitely generated substructures of each $\mathcal{A}_i$. The embeddings are precisely the inclusion maps from each substructure of $\mathcal{A}_i$ into $\mathcal{A}_i$. It is quickly seen that this yields a uniformly computable cover. $\square$

In light of this lemma, one naturally asks why we bothered to give Definition 2. The answer is that local computability will be the $m = 0$ case in the following definition, which uses the enumeration of the sets $I_{ij}^{\mathfrak{A}}$ extensively. Indeed, it is the enumeration of the embeddings, rather than that of the finitely generated substructures of $\mathcal{S}$, which will be the heart of our study of local computability.

Definition 3. *Let $\mathfrak{A}$ be a cover of a structure $\mathcal{S}$. Every embedding β of any $\mathcal{A}_i \in \mathfrak{A}$ into $\mathcal{S}$ will be called 0-extensional. For each $m \geq 0$, we say that such an embedding β, with image $\mathcal{B} \subseteq \mathcal{S}$, is $(m + 1)$-extensional if:*

- *for every $j \in \omega$, every $f \in I_{ij}^{\mathfrak{A}}$ lifts to an inclusion $\mathcal{B} \subseteq \mathcal{C}$ in $\mathcal{S}$ via β and an m-extensional match γ with domain $\mathcal{A}_j$; and*
- *for every finitely generated $\mathcal{C}$ with $\mathcal{B} \subseteq \mathcal{C} \subseteq \mathcal{S}$, there exists $j \in \omega$ and $f \in I_{ij}^{\mathfrak{A}}$ which lifts to $\mathcal{B} \subseteq \mathcal{C}$ via β and some m-extensional match γ.*

A uniformly computable cover $\mathfrak{A}$ of $\mathcal{S}$ is m-extensional if every $\mathcal{A}_i \in \mathfrak{A}$ m-extensionally matches some substructure of $\mathcal{S}$ and every finitely generated substructure of $\mathcal{S}$ m-extensionally matches some $\mathcal{A}_i \in \mathfrak{A}$. If such a cover exists, we say that $\mathcal{S}$ is m-extensionally locally computable (or just m-extensional). A structure is ω-extensionally locally computable if it is m-extensionally locally computable for every $m \in \omega$.

Notice that $\mathcal{S}$ is 0-extensional iff $\mathcal{S}$ is locally computable, iff $\mathcal{S}$ has a uniformly computable simple cover (by Lemma 1). Definition 3 will be used in Proposition 4 to derive results about the complexity of the theory of $\mathcal{S}$. The idea of 1-extensionality is that the embeddings in the sets $I_{ij}^{\mathfrak{A}}$ (for all j) correspond precisely to the finitely generated superstructures of $\mathcal{B}$ in $\mathcal{S}$, rather than just to some possible extension of some $\mathcal{B}' \cong \mathcal{B}$ within $\mathcal{S}$ to some superstructure of $\mathcal{B}'$ in $\mathcal{S}$. The distinction is best illustrated by the negative example of Proposition 3 below. However, Proposition 5 will show that the definition holds for the field

of complex numbers. Indeed, the set $I_{ij}^{\mathfrak{A}}$ of embeddings of one $\mathcal{A}_i$ in the cover into another $\mathcal{A}_j$ is just the set of all embeddings of $\mathcal{A}_i$ into $\mathcal{A}_j$, and this set is actually computable, uniformly in i and j.

m-extensionality is the obvious iteration of this notion. The extra conditions for extensionality strengthen the idea that each finitely generated substructure of $\mathcal{S}$ is represented by some $\mathcal{A}_i \in \mathfrak{A}$: not only are they isomorphic, but the embeddings (given by $I^{\mathfrak{A}}$) of $\mathcal{A}_i$ into other structures in $\mathfrak{A}$ coincide exactly with the extensions of $\mathcal{B}$ to larger finitely generated substructures of $\mathcal{S}$.

Proposition 3. *The field $\mathcal{R}$ of real numbers is not 1-extensionally locally computable.*

Proof. Suppose that $\mathfrak{A}$ were a 1-extensional cover of $\mathcal{R}$. Fix any noncomputable real number $t \in \mathbb{R}$. Definition 3 gives an $\mathcal{A}_i \in \mathfrak{A}$ which 1-extensionally matches (via some isomorphism β) the subfield $\mathcal{B}$ of $\mathcal{R}$ generated by t, and we may assume we know i and $\beta^{-1}(t)$, since they constitute finitely much information.

Now we can enumerate the lower cut of rationals $q < t$ in $\mathcal{R}$, knowing that extensions of $\mathcal{B}$ in $\mathcal{R}$ correspond to embeddings $f \in I_{ij}^{\mathfrak{A}}$ (for all j) in the 1-extensional cover. For any rational $q \in \mathcal{R}$:

$$
\begin{aligned}
\models_{\mathcal{R}} q < t \iff & \models_{\mathcal{R}} (\exists x) x^2 = t - q \\
\iff & (\exists \text{ f.g. } \mathcal{C})[\mathcal{B} \subseteq \mathcal{C} \subset \mathcal{R} \ \& \ \models_{\mathcal{C}} (\exists x) x^2 = t - q] \\
\iff & (\exists j)(\exists f \in I_{ij}^{\mathfrak{A}}) \models_{\mathcal{A}_j} (\exists x) x^2 = f(\beta^{-1}(t - q)) \\
\iff & (\exists j)(\exists f \in I_{ij}^{\mathfrak{A}})(\exists a \in \mathcal{A}_j) \models_{\mathcal{A}_j} a^2 = f(\beta^{-1}(t)) - f(\beta^{-1}(q)).
\end{aligned}
$$

A similar argument holds for the upper cut of rationals $q > t$, using square roots of $(q - t)$ in $\mathcal{R}$. This contradicts the noncomputability of t. (Of course, β and f fix the rationals, so $f(\beta^{-1}(q)) \in \mathcal{A}_j$ is just the element of $\mathcal{A}_j$ representing q. This $f(\beta^{-1}(q))$, lying in dom($\mathcal{A}_j$), is a natural number, but the element of $\mathcal{A}_j$ representing any particular rational q can easily be computed from the numerator and denominator of q, uniformly in j, by using the functions of $\mathcal{A}_j$.) $\qquad \square$

So the extensional local computability of $\mathcal{C}$ in Proposition 5 does not follow solely from the existential closure of the structure; after all, $\mathcal{R}$, viewed as a real closed field, is also existentially closed. The difficulty for $\mathcal{R}$ is that real closed fields have an implicit order on their elements, whether it is included in the language of the structure or not, and as we saw in Proposition 2, adding the order relation to $\mathcal{R}$ destroys local computability. $\mathcal{R}$ itself can still be locally computable, because the relation $<$ cannot be defined in $\mathcal{R}$ without quantifiers (even though it is both Σ_1-definable and Π_1-definable!) and existential questions about $\mathcal{R}$ can be left unanswered by a uniformly computable cover. An extensional cover, on the other hand, answers all such questions, as we now see.

Proposition 4. *For $m \in \omega$, any m-extensionally locally computable structure $\mathcal{S}$, and any $n \leq m + 1$, the Σ_n-theory of $\mathcal{S}$,*

$$\{\varphi \in Th(\mathcal{S}) : \varphi \text{ is a } \Sigma_n \text{ sentence}\},$$

is itself a Σ_n set in the arithmetic hierarchy. (For $n > 0$, this means that the Σ_n-theory is 1-reducible to $\emptyset^{(n)}$, and for $n = 0$, the Σ_0-theory is computable.)

Proof. Let $\langle \mathcal{A}_i \rangle_{i \in \omega}$ be an m-extensionally computable cover of $\mathcal{S}$. For arbitrary $n \leq m + 1$, the key fact is simply that for any formula $\varphi(\overline{x})$,

$$\models_{\mathcal{S}} (\exists \overline{x}) \varphi(\overline{x}) \ \ \text{iff} \ \ (\exists \text{ f.g. } \mathcal{B} \subseteq \mathcal{S})(\exists \, \overline{b} \in \mathcal{B}^j) \models_{\mathcal{S}} \varphi(\overline{b}).$$

When we have alternating quantifiers, we need to take superstructures at each step. For an arbitrary formula $\varphi(\overline{x}, \overline{y})$,

$$\models_{\mathcal{S}} (\exists \overline{x})(\forall \overline{y}) \varphi(\overline{x}, \overline{y})$$
$$\text{iff} \ \ (\exists \text{ f.g. } \mathcal{B} \subseteq \mathcal{S})(\exists \, \overline{x} \in \mathcal{B}^k) \models_{\mathcal{S}} (\forall \overline{y}) \varphi(\overline{x}, \overline{y})$$
$$\text{iff} \ \ (\exists \text{ f.g. } \mathcal{B} \subseteq \mathcal{S})(\exists \, \overline{x} \in \mathcal{B}^k)(\forall \text{ f.g. } \mathcal{C} \text{ s.t. } \mathcal{B} \subseteq \mathcal{C} \subseteq \mathcal{S})(\forall \overline{y} \in \mathcal{C}^p) \models_{\mathcal{S}} \varphi(\overline{x}, \overline{y})$$

If the original sentence was Σ_2, then the matrix (after all the quantifiers) will be the truth in $\mathcal{S}$ of the quantifier-free formula $\varphi(\overline{x}, \overline{y})$. In this case, $\varphi(\overline{x}, \overline{y})$ holds in $\mathcal{S}$ iff it holds in $\mathcal{C}$, so we can add the following:

$$\text{iff} \ \ (\exists \text{ f.g. } \mathcal{B} \subseteq \mathcal{S})(\exists \, \overline{x} \in \mathcal{B}^k)(\forall \text{ f.g. } \mathcal{C} \text{ s.t. } \mathcal{B} \subseteq \mathcal{C} \subseteq \mathcal{S})(\forall \overline{y} \in \mathcal{C}^p) \models_{\mathcal{C}} \varphi(\overline{x}, \overline{y})$$
$$\text{iff} \ (\exists i)(\exists \, \overline{b} \in \mathcal{A}_i^k)(\forall j)(\forall f \in I_{ij}^{\mathfrak{A}})(\forall \overline{c} \in \mathcal{A}_j^p) \models_{\mathcal{A}_j} \varphi(f(\overline{b}), \overline{c}).$$

The definition of 1-extensional cover shows these last two lines to be equivalent. Specifically, if the last line holds, then the witness $\mathcal{A}_i$ has a 1-extensional match β onto some $\mathcal{B} \subseteq \mathcal{S}$, and Definition 3, applied to any $\mathcal{A}_j$ and $f \in I_{ij}^{\mathfrak{A}}$, provides a 0-extensional match γ from $\mathcal{A}_j$ onto some $\mathcal{C} \supseteq \mathcal{B}$ such that $\gamma \circ f = \beta$. Then $\varphi(\gamma(f(\overline{b})), \gamma(\overline{c}))$ must hold in $\mathcal{C}$, since $\varphi(f(\overline{b}), \overline{c})$ holds in $\mathcal{A}_j$ and γ is an isomorphism. Conversely, if the next-to-last line holds, a similar argument applies, since there is some 1-extensional match onto the witness $\mathcal{B}$ from some $\mathcal{A}_i \in \mathfrak{A}$. This completes the proof of the result on 1-extensional structures.

The obvious iteration of this process, applied to any Σ_n sentence about $\mathcal{S}$, yields a statement consisting of a Σ_n-sequence of quantifiers over structures in $\mathfrak{A}$, their elements, and the sets $I_{ij}^{\mathfrak{A}}$, followed by a quantifier-free statement about an $\mathcal{A}_j \in \mathfrak{A}$. The argument requires that each $\mathcal{A}_i$ correspond to some $\mathcal{B}$ via an $(n-1)$-extensional map, so that the extensions must then correspond via $(n-2)$-extensional maps, and so on down to 0-extensional maps once all the quantifiers have been moved outside the turnstile $\models$. Therefore, for an m-extensionally locally computable $\mathcal{S}$ with $m \geq (n - 1)$, the Σ_n statement yielded by iterating the process holds iff the original Σ_n sentence held in $\mathcal{S}$. Since the structures in $\mathfrak{A}$, the sets $I_{ij}^{\mathfrak{A}}$ and the atomic diagram of such an $\mathcal{A}_j$ are all computable uniformly in i and j, the truth of the original Σ_n-sentence in $\mathcal{S}$ is itself a Σ_n fact. Moreover, this process is entirely uniform in n. $\square$

Notice that this argument does not extend to values $n > m + 1$. For $m = 0$ a specific counterexample appears in Proposition 3.

Theorem 1. *For $m \in \omega$, any m-extensionally locally computable structure $\mathcal{S}$, any finite tuple $\overline{p}$ of parameters from $\mathcal{S}$, and any $n \leq m$, the Σ_n-theory of $\mathcal{S}$ over $\overline{p}$,*

$$\{\varphi \in Th(\mathcal{S}, \overline{p}) : \varphi \text{ is a } \Sigma_n \text{ sentence}\},$$

is itself Σ_n, uniformly in n and in an appropriate description of the parameters (as discussed after the proof).

Proof. Let $\mathcal{B}$ be generated by $\overline{p}$ in $\mathcal{S}$, and fix an m-extensional match $\beta : \mathcal{A}_l \twoheadrightarrow \mathcal{B}$ for some $\mathcal{A}_l \in \mathfrak{A}$. As before, we give an example by evaluating the truth in $\mathcal{S}$ of an arbitrary Σ_2 sentence with the parameters $\overline{p}$, assuming now that $m \geq 2$. By an argument similar to that in the proof of Proposition 4, the Σ_2 sentence $(\exists \overline{x})(\forall \overline{y})\varphi(\overline{p}, \overline{x}, \overline{y})$ holds in $\mathcal{S}$ iff

$$(\exists i)(\exists h \in I_{li}^{\mathfrak{A}})(\exists \overline{b} \in \mathcal{A}_i^k)(\forall j)(\forall f \in I_{ij}^{\mathfrak{A}})(\forall \overline{c} \in \mathcal{A}_j^m)$$
$$\models_{\mathcal{A}_j} \varphi(f(h(\overline{a})), f(\overline{b}), \overline{c}),$$

which is a Σ_2 condition, uniformly in $\overline{a}$ and l. The obvious iteration works for any $n \leq m$, but no longer applies when $n = m + 1$. In the example above, as long as $\mathcal{S}$ is 2-extensional, we need β to be 2-extensional, so that $\mathcal{A}_i$ will have a 1-extensional match in its turn. Adding the parameters forces us to start by fixing an $\mathcal{A}_l \in \mathfrak{A}$ and a β, whereas in Proposition 4 we were allowed simply to search for any $\mathcal{A}_i$ and a single embedding into an $\mathcal{A}_j$. Hence parameters require one more level of extensionality.

Of course, knowing an original parameter $p_i \in \mathcal{B}$ is useless to us; we need to know l and the value $a_i = \beta^{-1}(p_i)$ in $\mathcal{A}_l$. For finitely many parameters, this constitutes only finitely much information, but we also wish to consider uniformity. It does not make sense to ask that parameters from a potentially uncountable structure $\mathcal{S}$ be given uniformly. Instead, our formal statement of uniformity is that if we are given an l and finitely many parameters $\overline{a}$ from $\mathcal{A}_l$, then for any n and any n-extensional match β mapping $\mathcal{A}_l$ into $\mathcal{S}$, the Σ_n-theory of $\mathcal{S}$ over the parameters $\beta(a_1), \ldots, \beta(a_k)$ is Σ_n uniformly in l and $\overline{a}$ and n. $\quad \square$

We view Theorem 1 as the strongest argument yet that local computability, and in particular perfect local computability, is the correct analogue in uncountable structures to computable presentability in countable structures. The point of a computable presentation of a structure is not just that it allows us to compute the atomic theory and enumerate the Σ_1-theory and so on, but that it actually allows us to do over specific elements of the structure: the atomic *diagram* is computable, and the Σ_n diagram is Σ_n, uniformly in n. For an uncountable $\mathcal{S}$, of course, there is no effective way to name all individual elements, so it is hopeless to expect the entire atomic diagram to be computable. An m-extensional cover, however, gives us a way of describing individual elements and tuples of them: using the cover, we name an $\mathcal{A}_l$ which m-extensionally matches the substructure of $\mathcal{S}$ generated by the tuple, and specify which elements of $\mathcal{A}_l$ correspond to the tuple.

To state the same fact differently, having an m-extensional cover tells us exactly what information we need about the tuple $\overline{p}$ from $\mathcal{S}$ in order to compute the atomic theory of $\mathcal{S}$ over $\overline{p}$, or to enumerate its Σ_1 theory over $\overline{p}$, etc. For the field of complex numbers, for instance, an $\mathcal{A}_i$ is given by its transcendence degree and the minimal polynomial of a single additional element generating the rest of $\mathcal{A}_i$ over a transcendence basis. If we can determine this information for the subfield $\mathbb{Q}(\overline{p}) \subset \mathbb{C}$, and know which elements correspond to $\overline{p}$, then without further information we can give a Σ_n description of the Σ_n-theory of $(\mathbb{C}, \overline{p})$. Indeed, there is another perfect cover of the field $\mathbb{C}$, more difficult to describe than the one in Proposition 5, in which every finite tuple $\overline{p}$ from $\mathbb{C}$ corresponds, via a perfect match, to the generators of a particular $\mathcal{A}_i$ in the cover. Using this cover, one would only need to know the minimal polynomial of each p_{i+1} over $\mathbb{Q}(p_1, \ldots, p_i)$, or else to know that no such minimal polynomial exists. Similarly, each m-extensional cover of any $\mathcal{S}$ says, "if you tell me this particular information about your tuple $\overline{p}$ from $\mathcal{S}$, I will give you a Σ_n-presentation of the Σ_n facts about $\overline{p}$ in $\mathcal{S}$, for each $n \leq m$."

Corollary 1. *For any ω-extensionally locally computable structure $\mathcal{S}$, and any finite parameter set $\overline{p}$ from $\mathcal{S}$, the Σ_n-theory of $(\mathcal{S}, \overline{p})$ is Σ_n for every n, uniformly in n and in $\overline{p}$ (as described above). In particular, this holds for any $\mathcal{S}$ with a correspondence system, including any perfectly locally computable $\mathcal{S}$.* □

Corollary 2. *Any two structures with the same ω-extensional cover are elementarily equivalent and realize the same types. Also, any two structures with the same m-extensional cover have the same $(m+1)$-quantifier theory.* □

We add one more version of local computability, even stronger than ω-extensional local computability, whose main interest stems from Theorem 2.

Definition 4. *Let $\mathfrak{A}$ be a uniformly computable cover for a structure $\mathcal{S}$. A set M is a* correspondence system *for $\mathfrak{A}$ and $\mathcal{S}$ if it satisfies all of the following:*

1. *Each element of M is an embedding of some $\mathcal{A}_i \in \mathfrak{A}$ into $\mathcal{S}$; and*
2. *Every $\mathcal{A}_i \in \mathfrak{A}$ is the domain of some $\beta \in M$; and*
3. *Every finitely generated $\mathcal{B} \subseteq \mathcal{S}$ is the image of some $\beta \in M$; and*
4. *For every i and j and every $\beta \in M$ with domain $\mathcal{A}_i$, every $f \in I_{ij}^{\mathfrak{A}}$ lifts to an inclusion $\beta(\mathcal{A}_i) \subseteq \gamma(\mathcal{A}_j)$ via β and some $\gamma \in M$; and*
5. *For every i, every $\beta \in M$ with domain $\mathcal{A}_i$, and every finitely generated $\mathcal{C} \subseteq \mathcal{S}$ containing $\beta(\mathcal{A}_i)$, there exist a j and an $f \in I_{ij}^{\mathfrak{A}}$ which lifts to $\beta(\mathcal{A}_i) \subseteq \mathcal{C}$ via β and some $\gamma : \mathcal{A}_j \twoheadrightarrow \mathcal{C} \in M$.*

The correspondence system is perfect *if it also satisfies*

6. *For every finitely generated $\mathcal{B} \subseteq \mathcal{S}$, if $\beta : \mathcal{A}_i \twoheadrightarrow \mathcal{B}$ and $\gamma : \mathcal{A}_j \twoheadrightarrow \mathcal{B}$ both lie in M and have image $\mathcal{B}$, then $\gamma^{-1} \circ \beta \in I_{ij}^{\mathfrak{A}}$.*

If a perfect correspondence system exists, then its elements are called perfect matches *between their domains and their images. $\mathcal{S}$ is then said to be* perfectly locally computable, *with* perfect cover $\mathfrak{A}$.

This concept is related to extensionality, clearly, and any element β of a correspondence system M is quickly seen to be an m-extensional match for every $m \in \omega$. However, perfect local computability is stronger than ω-extensional local computability in two distinct ways. First, for the map β to be an m-extensional match, we only needed the existence of $(m-1)$-extensional matches γ to relate the embeddings $f \in I_{ij}^{\mathfrak{A}}$ (for all j) to the finitely generated extensions of the image of β in $\mathcal{S}$, and for different values of m, we could use different maps γ. Here Conditions 4 and 5 require that the isomorphisms γ be in M themselves, hence that they satisfy the same conditions. The second difference is Condition 6, which is not related to Definition 3, but appears necessary for Theorem 2.

Proposition 5. *Every algebraically closed field of characteristic* 0*, and in particular the field* $\mathbb{C}$*, is perfectly locally computable.* $\quad\square$

Theorem 2. *Let* $\mathcal{S}$ *be a countable structure. Then* $\mathcal{S}$ *is computably presentable iff* $\mathcal{S}$ *is perfectly locally computable.*

The proof appears in Appendix A.

To close, we state without proof several more relevant results.

Proposition 6. *Let* $\mathcal{R}_c$ *be the ordered field containing all computable real numbers. Then* $\mathcal{R}_c$ *has a computable cover, but no uniformly computable cover.* $\quad\square$

Lemma 2. *There exist countable structures* $\mathcal{S}$ *and* $\mathcal{S}'$ *with the same uniformly computable cover, such that* $\mathcal{S}$ *is computable (and hence perfectly locally computable), but* $\mathcal{S}'$ *is not computably presentable, indeed not even* 1*-extensional.* $\quad\square$

Lemma 3. *There exist* 2^{ω}*-many countable structures, all with the same uniformly computable cover, which are pairwise elementarily non-equivalent. Indeed, these structures all have distinct* Σ_2*-theories, and every Turing degree is the degree of the* Σ_2*-theory of some such structure.* $\quad\square$

Finally, the converse of each statement in Theorem 1 is false.

Proposition 7. *There exists a tree* T *which is not* 1*-extensionally locally computable, yet such that for every* m *and every finite tuple* $\overline{p}$ *from* T*, the* Σ_m*-theory of* $(T, \overline{p})$ *is itself* Σ_m*.*

References

1. L. Blum, M. Shub, & S. Smale; On a theory of computation and complexity over the real numbers: NP-completeness, recursive functions, and universal machines, Bulletin of the AMS **21** (1989) 1-46.
2. M. Braverman & S. Cook; Computing over the Reals: Foundations for Scientific Computing, Notices of the AMS **53** (2006) 3, 318-329.
3. V.S. Harizanov; Pure computable model theory, Handbook of Recursive Mathematics, vol. 1 (Amsterdam: Elsevier, 1998), 3-114.
4. R.I. Soare; Recursively Enumerable Sets and Degrees (New York: Springer-Verlag, 1987).

3   Appendix A: Proof of Theorem 2.

The forward direction is not difficult. Suppose that $\mathcal{B}$ is a computable presentation of $\mathcal{S}$. Let $\mathcal{B}_i$ be the substructure of $\mathcal{B}$ generated by the i-th tuple $(\bar{b})_i$ of $\omega^{<\omega}$. Then the domain of $\mathcal{B}_i$ is a computably enumerable set. Enumerate its elements, and let A_i be the domain of the enumeration. (That is, A_i is an initial segment of ω from which we have a bijection β_i onto the domain of $\mathcal{B}_i$, all computable uniformly in i.) Then we may define computable structures $\mathcal{A}_i$, uniformly in i and each with domain A_i, isomorphic to each $\mathcal{B}_i$ via the map β_i. Next, we define $I_{ij}^{\mathfrak{A}}$ to be the set of functions

$$\{\beta_j^{-1} \circ \beta_i : (\bar{b})_i \subseteq \mathcal{B}_j\}.$$

The condition $(\bar{b})_i \subseteq \mathcal{B}_j$ is Σ_1, so this is a c.e. set, uniformly in i and j. Clearly all elements of $I_{ij}^{\mathfrak{A}}$ are embeddings of $\mathcal{A}_i$ into $\mathcal{A}_j$.

Next we claim that the set $M = \{\beta_i : i \in \omega\}$ is a perfect correspondence system for $\mathcal{B}$ and $\mathfrak{A}$. Clearly the first three conditions of Definition 4 are satisfied. Fix i, and suppose first that $f \in I_{ij}^{\mathfrak{A}}$. Then $f = \beta_j^{-1} \circ \beta_i$, and this map entered $I_{ij}^{\mathfrak{A}}$ because $\mathcal{B}_i \subseteq \mathcal{B}_j$. But then $\beta_j \circ f = \beta_i$, and of course $\beta_j \in M$. Conversely, any finitely generated superstructure of $\mathcal{B}_i$ within $\mathcal{S}$ must be $\mathcal{B}_j$ for some j, and again $\beta_j : \mathcal{A}_j \twoheadrightarrow \mathcal{B}_j$ is the necessary element of M, with $f = \beta_j^{-1} \circ \beta_i \in I_{ij}^{\mathfrak{A}}$. Condition 6 of Definition 4 is immediate. Thus M is a perfect correspondence system.

For the converse, suppose that $\mathfrak{A}$ is a perfect cover for $\mathcal{S}$, with a perfect correspondence system M. We let $\overline{I}^{\mathfrak{A}}$ be the closure of $I^{\mathfrak{A}}$ under composition, so that if $f \in \overline{I}_{ij}^{\mathfrak{A}}$ and $g \in \overline{I}_{jk}^{\mathfrak{A}}$, then $g \circ f \in \overline{I}_{ik}^{\mathfrak{A}}$. Clearly $\overline{I}^{\mathfrak{A}}$ is also computably enumerable, and one checks quickly that M is also a perfect correspondence system for the cover $\mathfrak{A}$ with $\overline{I}^{\mathfrak{A}}$ in place of $I^{\mathfrak{A}}$.

We will build a computable structure $\mathcal{B}$ and prove by a back-and-forth argument that $\mathcal{B} \cong \mathcal{S}$. To start, we let $\mathcal{B}_0 = \mathcal{A}_0$ and $i_0 = 0$. Then we proceed recursively.

Given $\mathcal{B}_s$, let $s = \langle t, u, v \rangle$, and fix $i = i_t$ and $k = i_s$, so that $\mathcal{B}_t = \mathcal{A}_i$ and $\mathcal{B}_s = \mathcal{A}_k$. We begin stage $s + 1$ by listing out those maps in $\overline{I}^{\mathfrak{A}}$ with domain $\mathcal{B}_t$ until we find the u-th element on this list. Let $f \in \overline{I}_{ij}^{\mathfrak{A}}$ (for some j) be this element. (Of course $t \leq s$, so $\mathcal{B}_t$ is defined. Also, if there were only finitely many maps in $\overline{I}^{\mathfrak{A}}$ with domain $\mathcal{A}_i$, then Definition 4 would imply that $\mathcal{S}$ is itself finitely generated, hence isomorphic to some $\mathcal{A}_i \in \mathfrak{A}$, hence computably presentable. Therefore we may assume that we do find such an f.)

We will incorporate this f into $\mathcal{B}_{s+1}$ as follows. We now search for some m such that there exist embeddings $g \in \overline{I}_{km}^{\mathfrak{A}}$ and $p \in \overline{I}_{jm}^{\mathfrak{A}}$ such that

$$p \circ f = g \circ g_{s-1} \circ \cdots \circ g_t.$$

The existence of such g and p is shown in Lemma 4 below. We then define $i_{s+1} = m$ and $\mathcal{B}_{s+1} = \mathcal{A}_m$, with $f_s = p \circ f : \mathcal{B}_{t_{s+1}} \hookrightarrow \mathcal{B}_{s+1}$ and $g_s = g : \mathcal{B}_s \hookrightarrow \mathcal{B}_{s+1}$.

This completes stage $s + 1$, and we say that f has been incorporated into $\mathcal{B}$ at this stage. Notice that every $f \in \overline{I}^{\mathfrak{A}}$ whose domain is $\mathcal{A}_{i_t}$ $(= \mathcal{B}_t)$ for any t will be incorporated into $\mathcal{B}$ at infinitely many stages.

Lemma 4. *In this situation, the required g and p do exist.*

Proof. M is a correspondence system, hence contains a map β_i embedding $\mathcal{A}_i$ into $\mathcal{S}$. We are given maps $f \in \overline{I}^{\mathfrak{A}}_{ij}$ and $g_{s-1} \circ \cdots \circ g_t \in \overline{I}^{\mathfrak{A}}_{ik}$, so there exist maps $\beta_j, \beta_k \in M$ embedding $\mathcal{A}_j$ and $\mathcal{A}_k$ into $\mathcal{S}$ with

$$\beta_j \circ f = \beta_i = \beta_k \circ g_{s-1} \circ \cdots \circ g_t.$$

Let $\mathcal{C} \subset \mathcal{S}$ be the substructure generated by the images of β_j and β_k together. Then by Definition 4, there exist $\mathcal{A}_m, \mathcal{A}_n \in \mathfrak{A}$, maps $g \in \overline{I}^{\mathfrak{A}}_{km}$ and $f_n \in \overline{I}^{\mathfrak{A}}_{jn}$, and isomorphisms $\beta_m : \mathcal{A}_m \twoheadrightarrow \mathcal{C}$ and $\beta_n : \mathcal{A}_n \twoheadrightarrow \mathcal{C}$, both in M, with $\beta_m \circ g = \beta_k$ and $\beta_n \circ f_n = \beta_j$. But by Condition 6 of Definition 4, $\overline{I}^{\mathfrak{A}}_{nm}$ must contain the embedding $\beta_m^{-1} \circ \beta_n$. The embedding $p = (\beta_m^{-1} \circ \beta_n) \circ f_n$ lies in $\overline{I}^{\mathfrak{A}}_{jm}$, since $\overline{I}^{\mathfrak{A}}$ is closed under composition, and for $x \in \mathcal{A}_i$ we have as required:

$$\begin{aligned}
p(f(x)) &= \beta_m^{-1}(\beta_n(f_n(f(x)))) \\
&= \beta_m^{-1}(\beta_j(f(x))) \\
&= \beta_m^{-1}(\beta_i(x)) \\
&= \beta_m^{-1}(\beta_k \circ g_{s-1} \circ \cdots \circ g_t(x)) \\
&= \beta_m^{-1}((\beta_m \circ g) \circ g_{s-1} \circ \cdots \circ g_t(x)) \\
&= g \circ g_{s-1} \circ \cdots \circ g_t(x). \qquad \square
\end{aligned}$$

The structure $\mathcal{B}$ itself is the union of the chain of structures $\mathcal{B}_s$ under the embeddings $g_{s+1} : \mathcal{B}_s \hookrightarrow \mathcal{B}_{s+1}$. To build this $\mathcal{B}$ computably, at stage $\langle u, v \rangle$ we consider the element t (if any) which is been enumerated into the domain of $\mathcal{B}_u$ by stage v. (Note that $\mathcal{B}_u$ might be finite, so we cannot just wait for its next element to appear.) For each of the finitely many elements x which has already entered $\mathcal{B}$ at some stage $\langle u', v' \rangle < \langle u, v \rangle$, we check whether either $x = g'_u \circ g_{u'-1} \circ \cdots \circ g_{u+1}(t)$ (if $u' > u$) or $t = g_u \circ g_{u-1} \circ \cdots \circ g_{u'+1}(x)$ (if $u' < u$). If either of these holds, we do nothing at this stage; if neither holds, then we add a fresh element to $\mathcal{B}$ and identify it with $t \in \mathcal{B}_u$. Iterating this process over all stages $\langle u, v \rangle$ builds the domain of $\mathcal{B}$, and we define the functions and relations on it in the obvious way. Notice that for every s the embeddings of $\mathcal{B}_s$ and $\mathcal{B}_{s+1}$ into $\mathcal{B}$ are compatible with the map $g_s : \mathcal{B}_s \hookrightarrow \mathcal{B}_{s+1}$, and that we can compute these embeddings uniformly in s, so from here on we will view each $\mathcal{B}_s$ as a substructure of $\mathcal{B}$.

Next we build the isomorphism $\alpha : \mathcal{B} \twoheadrightarrow \mathcal{S}$, by going back and forth. Of course, α need not be computable. At every stage $s + 1$, the image $\mathcal{S}_s$ of the current approximation α_s will be a finitely generated substructure of $\mathcal{S}$, its domain will be some $\mathcal{B}_t$, and α_s will lie in M. We will ensure that the extension of α_s is compatible with the inclusion maps $g_t : \mathcal{B}_t \hookrightarrow \mathcal{B}_{t+1}$, so that at the end of the

construction, we may define $\alpha(x)$ for $x \in \mathcal{B}$ simply by finding some s and t with $x \in \mathcal{B}_t = \text{dom}(\alpha_s)$ and letting $\alpha(x) = \alpha_s(x)$.

In this situation, suppose first that $\text{dom}(\alpha_s) = \mathcal{B}_t = \mathcal{A}_i$ with $\text{range}(\alpha_s) = \mathcal{S}_s$, a substructure of $\mathcal{S}$, and assume that we wish to extend α_s to $\mathcal{B}_{t+1}$. (Recall that $\mathcal{B}_{t+1} = \mathcal{A}_m$ for some m.) Now we have an embedding $g_t : \mathcal{B}_t \hookrightarrow \mathcal{B}_{t+1}$, which lies in $\overline{I}_{im}^{\mathfrak{A}}$, by the construction of $\mathcal{B}$. Since M is a correspondence system, we may choose α_{s+1} to be an element of M mapping $\mathcal{A}_m$ onto some superstructure $\mathcal{S}_{s+1}$ of $\mathcal{S}_s$, with $\alpha_{s+1} \circ g_t = \alpha_s$.

Going backwards is harder. Suppose that $\text{dom}(\alpha_s) = \mathcal{B}_t = \mathcal{A}_i$ and that we wish to extend α_s to α_{s+1} by adding a new element $y \in \mathcal{S}$ to $\mathcal{S}_s = \text{range}(\alpha_s)$. Since M is a correspondence system containing α_s, there exists $A_j \in \mathfrak{A}$, $f \in \overline{I}_{ij}^{\mathfrak{A}}$, and a map $\beta \in M$ such that β maps $\mathcal{A}_j$ onto the substructure of $\mathcal{S}$ generated by y and $\mathcal{S}_s$, with $\beta \circ f = \alpha_s$. Now since $f \in \overline{I}^{\mathfrak{A}}$, there is some stage $s' > t$ at which f is incorporated into $\mathcal{B}$. When this happened, we defined an embedding $f_{s'} = p \circ f : \mathcal{B}_t \hookrightarrow \mathcal{B}_{s'+1}$, with $p \in \overline{I}_{jm}^{\mathfrak{A}}$ for some m, and with $f_{s'} \in \overline{I}_{im}^{\mathfrak{A}}$ because $\overline{I}^{\mathfrak{A}}$ is closed under composition. But since β lies in the correspondence system M, there exists $\gamma \in M$ with domain $\mathcal{B}_{s'+1}$ such that $\beta = \gamma \circ p$. We define $\alpha_{s+1} = \gamma$, with domain $\mathcal{B}_{s'+1}$ and let $\mathcal{S}_{s+1} = \text{range}(\alpha_{s+1})$, noting that for $x \in \text{dom}(\alpha_s)$,

$$\alpha_{s+1}(g_{s'} \circ \cdots \circ g_t(x)) = \gamma(p \circ f(x)) = \beta(f(x)) = \alpha_s(x).$$

(The first equality follows from Lemma 4.) Finally, $y \in \text{range}(\beta) \subseteq \text{range}(\gamma) = \text{range}(\alpha_{s+1})$. This completes the back-and-forth construction, so $\mathcal{S}$ and $\mathcal{B}$ really are isomorphic, and $\mathcal{S}$ is computably presentable.