

Limit Learning Equivalence Structures

Ekaterina Fokina

*Institute of Discrete Mathematics and Geometry
Vienna University of Technology, Austria*

EKATERINA.FOKINA@TUWIEN.AC.AT

Timo Kötzing

*Hasso Plattner Institute
University of Potsdam, Germany*

TIMO.KOETZING@HPI.DE

Luca San Mauro

*Institute of Discrete Mathematics and Geometry
Vienna University of Technology, Austria*

LUCA.SAN.MAURO@TUWIEN.AC.AT

Editors: Aurélien Garivier and Satyen Kale

Abstract

While most research in Gold-style learning focuses on learning formal languages, we consider the identification of computable structures, specifically equivalence structures. In our core model the learner gets more and more information about which pairs of elements of a structure are related and which are not. The aim of the learner is to find (an effective description of) the isomorphism type of the structure presented in the limit. In accordance with language learning we call this learning criterion **InfEx**-learning (explanatory learning from informant).

Our main contribution is a complete characterization of which families of equivalence structures are **InfEx**-learnable. This characterization allows us to derive a bound of \mathbf{O}'' on the computational complexity required to learn uniformly enumerable families of equivalence structures. We also investigate variants of **InfEx**-learning, including learning from text (where the only information provided is which elements are related, and not which elements are not related) and finite learning (where the first actual conjecture of the learner has to be correct). Finally, we show how learning families of structures relates to learning classes of languages by mapping learning tasks for structures to equivalent learning tasks for languages.

Keywords: Gold style learning, learning from informant, equivalence structures, finite separability

1. Introduction

Consider a learner observing (a countably infinite number of) different items to be equivalent or not equivalent. The learner would like to arrive at a conjecture about the *structure* of this equivalence relation, that is, the learner would like to determine the *isomorphism type* of the equivalence structure embodied by the items. For example, the learner could see more and more groups of 5 equivalent objects, and no groups of other sizes, and announces as a conjecture “infinitely many equivalence classes of size 5 and no equivalence classes of other sizes”. If the first guess has to be correct, we call the setting *finite learning* (denoted **Fin**), if the conjecture may be changed an arbitrary (but finite) number of times before stabilizing on a correct conjecture, we call the setting *explanatory learning* (denoted **Ex**). In each case, the data available to the learner is a *complete accurate* list of which elements of the structure are equivalent and which are not. Following standard convention

in inductive inference, we call this *learning from informant* (**Inf**), where both positive and negative information is available.

In general, this style of learning is called *learning in the limit* or *inductive inference*, and dates back to (Gold, 1967). Most work in inductive inference concerns either learning of formal languages or learning of total functions (see the text books (Osherson et al., 1986; Jain et al., 1999)), the case of learning other structures has first been considered by (Glymour, 1985) and is surveyed by (Martin and Osherson, 1998). More recently, researchers investigated the case in which the languages to be learned correspond to substructures of a given structure. For instance, (Stephan and Ventsov, 2001), (Harizanov and Stephan, 2007), and (Merkle and Stephan, 2004) considered learnable ideals of rings, subgroups and submonoids of groups, subspaces of vector spaces and isolated branches on uniformly computable sequences of trees. They showed that different types of learnability of various families of computable or computably enumerable structures can be characterized algebraically.

With the present paper, we want to strengthen the connection between algorithmic learning theory and computable structure theory by developing a learning framework in which one can formalize the intuition of learning an arbitrary structure in the limit (the interested reader can consult (Ash and Knight, 2000) for a classical introduction to computable structure theory). To this end, equivalence structures represent an ideal case-study. Although being fairly basic from an algebraic point of view, equivalence structures exhibit many deep effective properties, and thus they attracted much attention from computable theorists. We offer just few examples: (Calvert et al., 2006) classified computable equivalence structures that possess a unique computable presentation up to computable isomorphism; (Downey et al., 2017) studied the complexity of listing computable equivalence structures with no repetitions; and recently there has been an increasing interest in analyzing the effective content of computably enumerable equivalent relations, as in (Andrews et al., 2014).

In what follows, we denote by ω the set of natural numbers and we use a convenient short hand to denote isomorphism types. For any function $f : \omega \cup \{\omega\} \rightarrow \omega \cup \{\omega\}$ we denote by $[f]$ the isomorphism type of exactly $f(a)$ many equivalence classes of size a , and if only finitely many values of f are non-zero, then we can list all these values as $[n_0 : f(n_0), \dots, n_k : f(n_k)]$. For example, $[5 : \omega]$ denotes the isomorphism type of infinitely many equivalence classes of size 5 and no others. Note that, to learn a structure \mathcal{A} , we must learn any presentation of \mathcal{A} with members of the natural numbers.

Trivially, a single structure, or also a single isomorphism type, is always learnable by a learner which constantly outputs a correct conjecture. Thus, we are interested in which *families* of structures can be learned by a single learner. Thus, we consider families \mathfrak{A} of equivalence relations on ω and ask whether there is a single learner M such that M can learn any $\mathcal{A} \in \mathfrak{A}$ when given more and more (accurate and, in the limit, complete) information about \mathcal{A} . We will mostly consider arbitrary functions as learners, but we will also discuss the computational complexity of learning. The following example illustrates the concept of learning families of structures.

Example 1 $[5 : \omega, 6 : 2]$ and $[5 : \omega, 7 : 1]$ are simultaneously **InfFin**-learnable (finitely learnable from informant): If there is ever an equivalence class of size 7 in the input, then the second structure is the only possibility, while once there are two classes of size 6 are in the input, the first structure is the only possibility.

It is most interesting to see what parts of the learning setting influence the learning power, and in what way. For example, we might wonder whether the ability to change the hypothesis arbitrarily

often (as in explanatory learning) gives an advantage over finite learning. The next example shows that the learning power of finite learning is indeed smaller than that of explanatory learning.

Example 2 $[5 : \omega], [6 : \omega]$ are not simultaneously **InfFin**-learnable, but **InfEx**-learnable: Regarding the negative part, any finite information about an instance of $[5 : \omega]$ can be extended to an instance of $[6 : \omega]$, so at no time can the learner commit to a hypothesis. An explanatory learner on the other hand can conjecture $[5 : \omega]$ until any equivalence class of size 6 appears in the input and then change to $[6 : \omega]$.

This shows that **InfFin**-learning is unreasonably weak, only very restricted families of structures are learnable in this sense (we characterize finite learning in two ways in Theorem 11). Note that a common way of restricting the learner even less than in **InfEx**-learning is by not requiring *syntactic* convergence to a final hypothesis, but only *semantic* convergence; that is, from some point on, all conjectures are correct, just not necessarily the same. The corresponding learning criteria replace **Ex** by **Bc** (behaviorally correct). We are mostly interested in learnability by arbitrary learners which can check for equivalence of conjectures (which is typically undecidable), so for our setting this relaxation does not make a difference. Thus, in this paper, we are interested in understanding the **InfEx** learning criterion.

So far we have seen examples of families that are **InfEx**-learnable. In order to establish that some families are not learnable, we turn to the concept of *locking sequences*, which is used extensively to show nonlearnability in the setting of learning formal languages. A locking sequence is a sequence of inputs σ for the target concept such that the learner does not change its mind regardless of how σ is extended with information for the target concept. Since we only want to learn up to isomorphism and the original concept of locking sequences is adjusted to exact learning, we get a slightly different notion of locking sequence (see the appendix for details). With this we can get the following result: there are two structures which are bi-embeddable (i.e., there is an embedding from any of the two structures into the other), but not simultaneously learnable. Clearly, if we only required learning up to bi-embeddability, then the two structures would be simultaneously learnable (by a constant learner). We use **InfEx** $_{\cong}$ to denote learning up to isomorphism and **InfEx** $_{\approx}$ for learning up to bi-embeddability¹. This result is summarized in the following example.

Example 3 $[5 : \omega, 2 : 1], [5 : \omega]$ are not simultaneously **InfEx** $_{\cong}$ -learnable, but **InfEx** $_{\approx}$ -learnable: Since the structures are bi-embeddable, a constant learner can **InfEx** $_{\approx}$ -learn both. Proving that these two structures are not **InfEx** $_{\cong}$ -learnable is a bit harder: intuitively, the problem is that any finite fragment of the two structures can be extended into a finite fragment of the other, in such way that a potential learner would be forced to change its mind infinitely often (see Theorem 5 and Theorem 12).

Interestingly, Example 3 shows that there is a class containing only finitely many different learning targets (in this case two), but it is still not **InfEx** $_{\cong}$ -learnable; this is in contrast to language learning, where learning finitely many different learning targets can always be distinguished.

Learning up to bi-embeddability is also of independent interest. In recent years, the relation of bi-embeddability received much attention in descriptive set theory and computable structure theory. In particular, bi-embeddability is of fundamental interest for classifying the complexity of equivalence relations in terms of Borel reducibility, as in (Friedman and Ros, 2011), and the effective counterpart of Borel reducibility discussed, e.g., in (Fokina et al., 2012). See also (Bazhenov

1. We use \cong to denote isomorphism of structures and \approx for bi-embeddability of structures.

et al., 2018) for a full classification of computable presentations of equivalence structures up to bi-embeddability.

In our examples so far we have never exploited the information that two elements are not related (the negative information). Learning without negative information is called *learning from text* (as opposed to learning from informant) and is denoted by **Text** instead of **Inf**. In fact we can show that, for learning structures neither of which has an equivalence class of size ω , informant and text learning are equivalent (see Theorem 13). The following example shows that this does not extend to structures which contain equivalence classes of size ω .

Example 4 $[\omega : 1], [\omega : 2]$ is **InfEx_≅**-learnable, but not **TextEx_≅**-learnable: Regarding the positive part, conjecture $[\omega : 1]$ until two elements are known to not be in the same equivalence class, then conjecture $[\omega : 2]$. The negative part is based on the concept of locking sequences, see Theorem 14.

These examples already give a good impression of what is learnable and what is not learnable with which kind of strategies. To further extend our intuition on what is **InfEx_≅**-learnable and what is not, we consider the following example.

Example 5 The infinite class of structures $\{[5 : n, 1 : \omega] \mid n \in \omega\} \cup \{[5 : \omega]\}$ is not **InfEx_≅**-learnable: Intuitively, when a learner tries to learn $[5 : \omega]$ and stops making mind changes after having seen some finite number n of equivalence classes of size k , the learner cannot successfully learn $[5 : n, 1 : \omega]$, since any extension to a finite part of $[5 : n, 1 : \omega]$ can be extended to $[5 : \omega]$.

We use these intuitions in Section 3 to give a characterization of **InfEx_≅**-learning which we call *finite separability*. Intuitively, a class of structures \mathfrak{A} is finitely separable if nonisomorphic structures of \mathfrak{A} that are finitely bi-embeddable can be distinguished by looking at some finite fragment of them. This characterization now completely informs about which families of structures are learnable. This simplifies proofs and furthermore allows us to give a bound on the complexity of learning c.e. enumerations of structures (see Section 3.1). Note that the notion of finite separability is similar to the existence of *tell-tales* as used for characterizing learnable classes of languages (Angluin, 1980).

With this characterization we were able to approach the complexity of **InfEx_≅**-learning and show that, for uniformly enumerable sets of **InfEx_≅**-learnable structures, $0''$ -computable learners are sufficient for learning, while computable learners are not.

For the reader familiar with language learning we provide two embeddings of learning of equivalence structures into the setting of language learning in Section 5. For both **InfEx_≅** and **InfEx_≈** we can map families of structures in an intuitive way to classes of languages such that a class of structures can be learned iff its image under the corresponding map is **TextEx**-language-learnable, provided that the structure of learning equivalence classes is in some sense a substructure of learning languages from text.

Note that, compared with language learning, learning of structures provides interesting new settings in which targets do not have to be learned exactly, but only up to some equivalence relation on structures (such as isomorphism). Furthermore, learning up to isomorphism has the advantage that “coding tricks” from language learning (making classes of languages learnable by having each language “encode” a correct hypothesis artificially in the data) are somewhat avoided. See (Jain et al., 1999, §13) for a discussion on coding tricks.

2. Learning of Structures

Our object of study is the class \mathbb{E} of the equivalence structures \mathcal{A} on natural numbers. For the benefit of exposition, we assume that all our equivalence structures are of the form (ω, E) , where E is an equivalence relation on ω . We say that \mathcal{A} is an ω -*presentation* of \mathcal{M} if $\mathcal{A} \cong \mathcal{M}$ and \mathcal{A} has domain ω . Note that the choice of limiting our focus to ω -presentations of equivalence structures is not a strong restriction: given any infinite equivalence structure $\mathcal{M} = (M, E)$, with $M = \{m_0, m_1, \dots\}$, one can use the bijection $i \mapsto m_i$ to get an equivalence structure \mathcal{A} that is an ω -presentation of \mathcal{M} .

The *atomic diagram* of \mathcal{A} is the set of atomic formulas and negations of atomic formulas true of \mathcal{A} .

An equivalence structure \mathcal{A} is *computably presentable* if \mathcal{A} is isomorphic to a computable equivalence structure. To formally define our learning framework, we rely on some effective enumeration of the computable structures from \mathbb{E} , up to isomorphism. This can be done in many ways; for instance, (Downey et al., 2017) showed, with a rather involved proof, that such an enumeration can be constructed which is a *Friedberg* enumeration, i.e., with no repetitions of the isomorphism types. For our interests, it is enough to fix some computable sequence of equivalence relations $(E_i)_{i \in \omega}$, where E_i is an equivalence relation on ω , such that any infinite computable equivalence structure is isomorphic to $\mathcal{M}_i = (\omega, E_i)$, for some i (see (Downey et al., 2017) for more details). We say that e is a *conjecture* for \mathcal{M}_e .

Recall that we aim at modeling a learner that receives larger and larger finite pieces of information about some equivalence structure \mathcal{A} .

A *text* is a function $T : \omega \rightarrow \omega^2 \cup \{\#\}$, where $\#$ is a special symbol denoting a *pause*, that is, no new information. We let $\text{content}(T) = \text{range}(T) \setminus \{\#\}$ be the *content* of T . For any text T and equivalence structure $\mathcal{A} = (\omega, E)$, we say T is a *text for* \mathcal{A} iff $\text{content}(T) = E$ (that is, $\text{content}(T)$ is all and only the positive information about \mathcal{A}). Note that pause symbol is required so that the structure $[1 : \omega]$, where no element is related to any other, has a text. By $\mathbf{Text}(\mathcal{A})$ we denote the set of all texts for \mathcal{A} .

An *informant* is a function $I : \omega \rightarrow \omega^2 \times \{0, 1\}$ such that, for any $(x, y) \in \omega^2$, either $((x, y), 0) \in \text{range}(I)$ or $((x, y), 1) \in \text{range}(I)$ (but never both). We denote the *positive content* of I by $\text{content}^+(I) = \{(x, y) \mid ((x, y), 1) \in \text{range}(I)\}$. Intuitively, an informant eventually lists, for each pair of elements, whether they are related or whether they are unrelated. For any informant I and equivalence structure $\mathcal{A} = (\omega, E)$, we say I is an *informant for* \mathcal{A} iff $\text{content}^+(I) = E$. By $\mathbf{Inf}(\mathcal{A})$ we denote the set of all informants for \mathcal{A} .

For any function f defined on natural numbers (such as texts and informants) and $n \in \omega$ we let $f[n]$ denote the finite sequence $f(0), \dots, f(n-1)$.

A learner is a function mapping initial segments of texts or informants to conjectures (elements of $\omega \cup \{?\}$). The *learning sequence* of a learner M on a text or informant f is $p : \omega \rightarrow \omega \cup \{?\}$ such that $p(n) = M(f[n])$.

For any finite sequence σ which is an initial part of an informant, we let \mathcal{A}_σ be the finite structure encoded by σ by using as universe all elements mentioned either positively or negatively in σ , taking the transitive closure of all positively mentioned pairs and assuming all other relations to be negative. We denote by $\mathcal{A}[s]$ the finite substructure of \mathcal{A} with domain $\{0, \dots, s-1\}$.

Any predicate on learning sequences and \mathbb{E} is called a *learning restriction*. We use the following learning restrictions.

Definition 1 Let \sim be any equivalence relation on structures.² We define the learning restriction \mathbf{Ex}_{\sim} on a learning sequence p and $\mathcal{A} \in \mathbb{E}$ such that

$$\mathbf{Ex}_{\sim}(p, \mathcal{A}) \Leftrightarrow \exists e \forall^{\infty} n : p(n) = e \wedge \mathcal{A} \sim \mathcal{M}_e.$$

Further, we define the learning restriction corresponding to finite learning by

$$\mathbf{Fin}_{\sim}(p, \mathcal{A}) \Leftrightarrow \exists e : \{e\} \subseteq \text{range}(p) \subseteq \{e, ?\} \wedge \mathcal{A} \sim \mathcal{M}_e.$$

Definition 2 A learning criterion is a triple of a set \mathcal{C} of (partial) functions $\omega \rightarrow \omega$ (the set of admissible learners), an operator α returning, for a given structure, a set of presentations for that structure (either all texts or all informants) and a learning restriction δ . We also write $\mathcal{C}\text{-}\alpha\delta$ for the learning criterion $(\mathcal{C}, \alpha, \delta)$. The class of $\mathcal{C}\text{-}\alpha\delta$ -learnable structures contains all those sets \mathfrak{A} of structures such that there is a learner $M \in \mathcal{C}$ such that, for all $\mathcal{A} \in \mathfrak{A}$, all ω -presentations \mathcal{A}^* of \mathcal{A} , and all $f \in \alpha(\mathcal{A}^*)$, $\delta(n \mapsto M(f[n]), \mathcal{A}^*)$.

Note, in the above definition, that to learn a structure \mathcal{A} , a learner should learn all the ω -presentations of \mathcal{A} . Sometimes we will denote by $M(\mathcal{A})$ the limit conjecture (if exists) of the learner M on input \mathcal{A} and by $M(\mathcal{A}[s])$ the conjecture M produces when given a string encoding $\mathcal{A}[s]$.

2.1. Notation Regarding Equivalence Structures

The character $\text{char}(\mathcal{A})$ of \mathcal{A} is

$$\text{char}(\mathcal{A}) = \{ \langle k, i \rangle : \mathcal{A} \text{ has at least } i \text{ equivalence classes of size } k, \text{ for } k \in \omega \cup \{\omega\} \}.$$

We call any element of $\text{char}(\mathcal{A})$ a *component* of \mathcal{A} . Sometimes, we will approximate the character of \mathcal{A} as follows

$$\text{char}(\mathcal{A}[s]) = \{ \langle k, i \rangle : \mathcal{A}[s] \text{ has at least } i \text{ equivalence classes of size } k, \text{ for } k \in \omega \}.$$

Let $\mathcal{A} = (\omega, E_{\mathcal{A}})$ and $\mathcal{B} = (\omega, E_{\mathcal{B}})$ be in \mathbb{E} . \mathcal{A} *embeds* into \mathcal{B} (notation: $\mathcal{A} \hookrightarrow \mathcal{B}$) if there is an injection $f : \omega \rightarrow \omega$ such that, for all $i, j \in \omega$, $i E_{\mathcal{A}} j \Leftrightarrow f(i) E_{\mathcal{B}} f(j)$. \mathcal{A} *finitely embeds* into \mathcal{B} (notation: $\mathcal{A} \hookrightarrow_{\text{fin}} \mathcal{B}$) if $\mathcal{A}[s] \hookrightarrow \mathcal{B}$, for all $s \in \omega$. \mathcal{A} and \mathcal{B} are *bi-embeddable* (notation: $\mathcal{A} \approx \mathcal{B}$) if they embeds in each other. \mathcal{A} and \mathcal{B} are *finitely bi-embeddable* (notation: $\mathcal{A} \approx_{\text{fin}} \mathcal{B}$) if they finitely embeds in each other. \mathcal{A} and \mathcal{B} are *isomorphic* (notation: $\mathcal{A} \cong \mathcal{B}$) if $\mathcal{A} \hookrightarrow \mathcal{B}$ via a bijection $f : \omega \rightarrow \omega$.

3. Characterizing \mathbf{InfEx}_{\cong}

Our main focus is on the class \mathbf{InfEx}_{\cong} . To help the reader get acquainted with our framework, we stress that some $\mathfrak{A} \in \mathbf{InfEx}_{\cong}$ if there is a learner M (of arbitrary complexity) such that, for any ω -presentation \mathcal{A}^* of a structure $\mathcal{A} \in \mathfrak{A}$, $M(\mathcal{A}^*) \cong \mathcal{A}$. In this section we characterize which families of equivalence structures are \mathbf{InfEx}_{\cong} -learnable. For the ease of presentation, we focus on equivalence structures with no infinite classes. Yet, it is not hard to modify the forthcoming analysis in order to obtain a full characterization of \mathbf{InfEx}_{\cong} ; this will be done in future work.

2. The equivalence relation \sim intuitively defines that $[\mathcal{A}]_{\sim}$ is the target at which a learner M is supposed to aim (typically \sim is \cong or \approx).

Definition 3 Let \mathfrak{A} be a family of equivalence structures with no infinite classes, and \mathcal{S} an equivalence structure.

1. If \mathfrak{A} has finitely many isomorphism types, then \mathcal{S} is a limit of \mathfrak{A} if there is $\mathcal{A} \in \mathfrak{A}$ such that

$$\mathcal{A} \not\cong \mathcal{S} \wedge \mathcal{A} \hookrightarrow_{\text{fin}} \mathcal{S},$$

$$\text{and } \text{char}(\mathcal{S}) \subseteq \text{char}(\mathcal{A}).$$

2. If \mathfrak{A} has infinitely many isomorphism types, \mathcal{S} is a limit of \mathfrak{A} if

$$(\forall \mathcal{A} \in \mathfrak{A})(\mathcal{A} \not\cong \mathcal{S} \wedge \mathcal{A} \hookrightarrow_{\text{fin}} \mathcal{S})$$

$$\text{and } \text{char}(\mathcal{S}) \subseteq \{\langle k, i \rangle : \mathfrak{A} \text{ contains infinitely many } \mathcal{B}'\text{s with } \langle k, i \rangle \in \text{char}(\mathcal{B})\}.$$

\mathfrak{A} is finitely separable if, for all $\mathfrak{B} \subseteq \mathfrak{A}$, \mathfrak{B} has no limits in \mathfrak{A} .

To clarify the above notion of limit, together with the corresponding one of finite separability, let us compare it with known examples of failure of \mathbf{InfEx}_{\cong} -learning.

Condition 1. is designed to deal with cases such as Example 3: It says that \mathcal{S} cannot be finitely separated from \mathcal{A} if \mathcal{A} finitely embeds in \mathcal{S} and any component of \mathcal{S} is a component of \mathcal{A} . Intuitively, this makes impossible to \mathbf{InfEx}_{\cong} -learn $\{\mathcal{A}, \mathcal{S}\}$ because we can build an ω -presentation \mathcal{S}^* of \mathcal{S} such that \mathcal{S}^* has arbitrarily large fragments that resemble \mathcal{A} , and this forces any potential learner M to have infinitely many mind changes if attempting to learn \mathcal{S}^* .

The infinite case is handled by Condition 2. and turns out to be more delicate. But the idea is the same and it aims at formalizing cases such as Example 5: If \mathcal{S} is the limit of an infinite family of pairwise nonisomorphic structures \mathfrak{A} , then we can construct an ω -presentation of \mathcal{S} that, for infinitely many initial segments, looks like some structure of \mathfrak{A} . In doing so, we eventually obtain a structure isomorphic to \mathcal{S} because each component of \mathcal{S} is witnessed by infinitely many structures of \mathfrak{A} .

We make the latter observations more precise by proving that finite separability coincides with \mathbf{InfEx}_{\cong} -learnability.

Remark 4 In the following proof, we introduce the formal notion of separators (of a given structure) that justify the terminology “finite separability” and the underlying intuition that, if \mathfrak{A} is finitely separable, then structures from \mathfrak{A} can be distinguished by looking at only finite fragments of them.

Theorem 5 Let \mathfrak{A} be a family of equivalence structures with no infinite classes. We have that

$$\mathfrak{A} \text{ is finitely separable} \Leftrightarrow \mathfrak{A} \in \mathbf{InfEx}_{\cong}.$$

Proof (\Rightarrow) : Assume that $\mathfrak{A} = \{\mathcal{A}_i\}_{i \in \omega}$ is finitely separable. To show that $\mathfrak{A} \in \mathbf{InfEx}_{\cong}$, we first prove that there is a learner M that learns the \approx_{fin} -type of any given $\mathcal{A} \in \mathfrak{A}$. Assume that M reads \mathcal{A}_z and, for all stages s , let $\mathfrak{B}_s \subseteq \mathfrak{A}$ be the class of equivalence structures in which $\mathcal{A}_z[s]$ is finitely embeddable and that are minimal with respect to $\hookrightarrow_{\text{fin}}$. We define M as follows: at stage s , M outputs the \approx_{fin} -type of the equivalence structure in \mathfrak{B}_s with least index in the enumeration $\{\mathcal{A}_i\}_{i \in \omega}$.

We claim that via this procedure M learns the \approx_{fin} -type of \mathcal{A}_z . In particular, we prove that $[\mathcal{A}_z]_{\approx_{\text{fin}}} \supseteq \bigcap_{i \in \omega} \mathfrak{B}_i$,

It is clear that $\mathcal{A}_z \in \bigcap_{i \in \omega} \mathfrak{B}_i$. Towards a contradiction, suppose that there is $\mathcal{A}_w \not\approx_{\text{fin}} \mathcal{A}_z$ such that $\mathcal{A}_w \in \bigcap_{i \in \omega} \mathfrak{B}_i$. We distinguish three cases.

- a) If \mathcal{A}_w and \mathcal{A}_z are $\hookrightarrow_{\text{fin}}$ -incomparable, then there must be a component of \mathcal{A}_z that is not a component of \mathcal{A}_w . So, there exists a stage s such that $\mathcal{A}_z[s]$ does not embed in \mathcal{A}_w . Therefore, $\mathcal{A}_w \notin \bigcap_{i \geq s} \mathfrak{B}_i$.
- b) If $\mathcal{A}_w \hookrightarrow_{\text{fin}} \mathcal{A}_z$, then there must be a component of \mathcal{A}_z witnessing the fact that $\mathcal{A}_z \not\hookrightarrow_{\text{fin}} \mathcal{A}_w$. So, there exists again a stage s such that $\mathcal{A}_z[s]$ does not embed in \mathcal{A}_w . Therefore, \mathcal{A}_w will be eventually outside from the \mathfrak{B}_s 's.
- c) If $\mathcal{A}_z \hookrightarrow_{\text{fin}} \mathcal{A}_w$, then, for all s , $\mathcal{A}_w \in \mathfrak{B}_s$ only if $\mathcal{A}_z \notin \mathfrak{B}_s$. This is because \mathfrak{B}_s contains only structures that are minimal with respect to $\hookrightarrow_{\text{fin}}$. But we already know that $\mathcal{A}_z \in \bigcap_{i \in \omega} \mathfrak{B}_i$. Therefore, $\mathcal{A}_w \notin \bigcap_{i \in \omega} \mathfrak{B}_i$.

This shows that M correctly learns the \approx_{fin} -type of any given \mathcal{A}_z . Call j the limit conjecture of M .

Now we construct a learner M^* that learns the isomorphism type of \mathcal{A}_z . To choose among the structures in $[\mathcal{M}_j]_{\approx_{\text{fin}}}$, M^* adopts the following procedure: For $\mathcal{C} \in [\mathcal{M}_j]_{\approx_{\text{fin}}}$, let $\text{sep}(\mathcal{C})$ be the *separator of \mathcal{C}* defined as

$$\text{sep}(\mathcal{C}) = \bigcup_{\mathcal{S} \in [\mathcal{M}_j]_{\approx_{\text{fin}}}} \{\min(\text{char}(\mathcal{C}) \setminus \text{char}(\mathcal{S}))\}.$$

First, note that separators of pairwise nonisomorphic structures form an anti-chain with respect to \subseteq . This follows from the fact that $\text{char}(\mathcal{C}) \setminus \text{char}(\mathcal{S}) \neq \emptyset$, for all $\mathcal{C} \not\cong \mathcal{S}$ in $[\mathcal{M}_j]_{\approx_{\text{fin}}}$. Otherwise, \mathcal{C} would be a limit of $\{\mathcal{S}\}$, since $\text{char}(\mathcal{C}) \subseteq \text{char}(\mathcal{S})$ and $\mathcal{S} \hookrightarrow_{\text{fin}} \mathcal{C}$, and this would contradict the finite separability of \mathfrak{A} .

Second, we claim that separators are all finite. Towards a contradiction assume that, for given $\mathcal{C} \in [\mathcal{M}_j]_{\approx_{\text{fin}}}$, $\text{sep}(\mathcal{C})$ is infinite. Denote by $\text{char}(\mathcal{C}) \upharpoonright_i$ the finite set consisting of the first i elements of $\text{char}(\mathcal{C})$. Let

$$\mathfrak{S}_i = \{\mathcal{S} \not\cong \mathcal{C} \in [\mathcal{M}_j]_{\approx_{\text{fin}}} : \text{char}(\mathcal{C}) \upharpoonright_i \subseteq \text{char}(\mathcal{S})\}.$$

We have that \mathfrak{S}_i is nonempty, for all i . Otherwise, if k is the least such $\mathfrak{S}_k = \emptyset$, we would obtain that $\text{sep}(\mathcal{C}) \subseteq \text{char}(\mathcal{C}) \upharpoonright_k$, against the hypothesis that $\text{sep}(\mathcal{C})$ is infinite. Suppose that $\bigcap_{i \in \omega} \mathfrak{S}_i$ is nonempty and contains a structure \mathcal{C}^* . It is immediate to see that $\text{char}(\mathcal{C}) \subseteq \text{char}(\mathcal{C}^*)$. Recall that $\mathcal{C}^* \hookrightarrow_{\text{fin}} \mathcal{C}$, since $\mathcal{C}^* \approx_{\text{fin}} \mathcal{C}$. It follows that \mathcal{C} is a limit of $\{\mathcal{C}^*\}$, which is impossible because of the finite separability of \mathfrak{A} .

But if all the \mathfrak{S}_i 's are nonempty, it follows that

$$\text{char}(\mathcal{C}) \subseteq \bigcup_{i \in \omega} \bigcap_{\mathcal{S} \in \mathfrak{S}_i} \text{char}(\mathcal{S}),$$

and this contradicts the finite separability of \mathfrak{A} . Thus, we conclude that, for all $\mathcal{C} \in [\mathcal{M}_j]_{\approx_{\text{fin}}}$, $\text{sep}(\mathcal{C})$ is finite.

We aim at making use of the finiteness of separators to show that M^* can eventually learn the isomorphism type of \mathcal{A}_z . We say that a finite equivalence structure $\mathcal{S}[t]$ *realizes* a given separator $\text{sep}(\mathcal{S})$ if $\text{sep}(\mathcal{S}) \subseteq \text{char}(\mathcal{S}[t])$. Denote by $\text{SEP}(\mathcal{S}[t])$ the class of the separators realized by $\mathcal{S}[t]$. Note that \mathcal{A}_z eventually realizes its own separator, i.e., there is s such that, for all $t \geq s$, $\text{sep}(\mathcal{A}_z) \in \text{SEP}(\mathcal{A}_z[t])$. The problem is that $\text{SEP}(\mathcal{A}_z[t])$ might consist of more than one separator for infinitely many stages, i.e., for infinitely many t , there might be $\mathcal{A}_w \not\cong \mathcal{A}_z$ and $\mathcal{A}_w \approx_{\text{fin}} \mathcal{A}_z$ such that $\text{sep}(\mathcal{A}_w) \in \text{SEP}(\mathcal{A}_z[t])$. Nonetheless, if any such $\text{sep}(\mathcal{A}_w)$ is realized by some $\mathcal{A}_z[t]$, there must be $t^* > t$ such that $\text{sep}(\mathcal{A}_w) \notin \text{SEP}(\mathcal{A}_z[t^*])$ (otherwise, by definition of separator, $\min(\text{char}(\mathcal{A}_w) \setminus \text{char}(\mathcal{A}_z))$ would belong to $\text{char}(\mathcal{A}_z)$, which is impossible). So, at some given stage t , M^* can choose the *oldest* separator in $\text{SEP}(\mathcal{A}_z[t])$, i.e., the separator that belongs to

$$\bigcap_{i \leq s \leq t} (\text{SEP}(\mathcal{A}_z[s])) \quad (1)$$

for the least i .

To sum up, to learn any structure $\mathcal{S} \in \mathfrak{A}$ the learner M^* does the following: At any given stage t , M^* takes the output of $M(\mathcal{S}[t])$ as the current guess of the \approx_{fin} -type of \mathcal{S} . Within the latter type M^* considers only the structures whose separators are realized by $\mathcal{S}[t]$ and outputs the isomorphism type of the one with the oldest realized separator. By this procedure, M^* **InfEx \cong** -learns \mathfrak{A} .

(\Leftarrow): This implication can be proved via locking sequences. But instead of crudely applying Theorem 18, we take here the opportunity of illustrating with some details how to dynamically build structures that serve as a counterexample to a given learning condition.

Assume that \mathfrak{A} is not finitely separable. We show that $\mathfrak{A} \notin \text{InfEx}_{\cong}$. Towards a contradiction, suppose that M learns \mathfrak{A} and let $\mathcal{A} \in \mathfrak{A}$ be a limit of \mathfrak{A} . We construct an ω -presentation \mathcal{B} of some structure in \mathfrak{A} that M cannot learn. We start by constructing \mathcal{B} as an ω -presentation of \mathcal{A} and we continue until M , on input $\mathcal{B}[s]$ for some s , outputs an index of \mathcal{A} . If this never happens, we obviously win: M fails to learn an ω -presentation of \mathcal{A} , and therefore M does not **InfEx \cong** -learn \mathfrak{A} . Otherwise, let s be a stage such that $\mathcal{A} \cong M(\mathcal{B}[s])$. Since \mathcal{A} is a limit of \mathfrak{A} , there must be some $\mathcal{S}_0 \in \mathfrak{A}$ such that $\mathcal{S}_0 \hookrightarrow_{\text{fin}} \mathcal{A}$, $\mathcal{S}_0 \not\cong \mathcal{A}$, and $\text{char}(\mathcal{B}[s]) \subseteq \text{char}(\mathcal{S}_0)$. We now extend $\mathcal{B}[s]$ as an ω -presentation of \mathcal{S}_0 , with the caution of not expanding the equivalence classes already defined in $\mathcal{B}[s]$ (this can always be done since $\text{char}(\mathcal{B}[s]) \subseteq \text{char}(\mathcal{S}_0)$). We continue building \mathcal{B} as an ω -presentation of \mathcal{S}_0 until we find some stage t such that M correctly guesses our plan, i.e., $M(\mathcal{B}[t]) \cong \mathcal{S}_0$. If this happens, we now go back to \mathcal{A} and extend $\mathcal{B}[t]$ as an ω -presentation of \mathcal{A} . We can do that because $\mathcal{S}_0 \hookrightarrow_{\text{fin}} \mathcal{A}$.

By iterating this reasoning, and possibly defining many \mathcal{S}_i 's, it is not hard to see that we can force M to either fail at learning some structure in \mathfrak{A} (this structure being either \mathcal{A} or some \mathcal{S}_i) or have infinitely many mind changes. If the latter case happens, the above caution of never expanding already defined \mathcal{B} -classes when we need to make \mathcal{B} an ω -presentation of some \mathcal{S}_i guarantees that we eventually obtain an ω -presentation of \mathcal{A} . \blacksquare

A nice consequence of our characterization theorem is that we can separate families of equivalence structures consisting of finitely many isomorphism types from those consisting of infinitely many isomorphism types by means of the following partial analogous of Compactness.

Corollary 6 *The following hold.*

1. If \mathfrak{A}/\cong is finite, then $\mathfrak{A} \in \mathbf{InfEx}_{\cong}$ if and only if the structures of \mathfrak{A} are pairwise \mathbf{InfEx}_{\cong} -learnable.
2. There is $\mathfrak{A} \notin \mathbf{InfEx}_{\cong}$ with \mathfrak{A}/\cong is infinite such that, for all $\mathfrak{B} \subseteq \mathfrak{A}$, if \mathfrak{B}/\cong is finite, then $\mathfrak{B} \in \mathbf{InfEx}_{\cong}$.

Proof (1) follows immediately from item (1) of Definition 3 and Theorem 5.

For (2), let $\mathcal{A}_i = [j : 1 - \delta_{ij}]$, where $\delta_{ij} = \begin{cases} 1, & \text{if } i = j; \\ 0, & \text{otherwise,} \end{cases}$ and let $\mathfrak{A} = \{\mathcal{A}_i\}_{i \in \omega}$. We have

that, for $i \neq j$, \mathcal{A}_i and \mathcal{A}_j can be finitely separated, by using $[j : 1]$ as a separator for \mathcal{A}_i and $[i : 1]$ as a separator for \mathcal{A}_j . From item (1) of the present corollary, it then follows that any subset of \mathfrak{A} which consists of finitely many isomorphism types is \mathbf{InfEx}_{\cong} -learnable. Yet, by Theorem 5, $\mathfrak{A} \notin \mathbf{InfEx}_{\cong}$, since \mathcal{A}_i is a limit of $\mathfrak{A} \setminus \mathcal{A}_i$ for all i . \blacksquare

3.1. Bounding the Complexity of the Learners

The procedure described above for learning any finitely separable family is obviously noneffective. It is natural to ask how much information is needed to perform it. More generally, we want to investigate what can be learned by learners of fixed complexity.

Definition 7 A family \mathfrak{A} of computably presentable structures is uniformly enumerable by f if f is a total computable function such that $\{\mathcal{M}_{f(i)}\}_{i \in \omega}$ is a one-to-one enumeration of all structures of \mathfrak{A} , up to isomorphism.

The next theorem shows that computable learners fail to learn all finitely separable families, even if we restrict to uniformly enumerable ones.

Theorem 8 There is a uniformly enumerable $\mathfrak{A} \in \mathbf{InfEx}_{\cong}$ such that $\mathfrak{A} \notin \mathbf{0-InfEx}_{\cong}$.

Proof We construct in a uniform way a family $\mathfrak{A} = \{\mathcal{A}_e, \mathcal{B}_e : e \in \omega\}$ of computable structures which is finitely separable, but cannot be \mathbf{InfEx}_{\cong} -learned by any computable learner.

INFORMAL STRATEGY

For all e , we want to diagonalize against the learner φ_e by building in stages a pair of equivalence structures $\mathcal{A}_e, \mathcal{B}_e$ that satisfy the following properties at all stages s ,

- there exists n such that $\mathcal{A}_e[s]$ has isomorphism type $[e : n, 1 : t_0]$ and $\mathcal{B}_e[s]$ has isomorphism type $[e : n + 1, 1 : t_1]$,
- and $\mathcal{A}_e[s] \subseteq \mathcal{B}_e[s]$.

The idea of the construction is that we want to force φ_e to either fail at learning $\{\mathcal{A}_e, \mathcal{B}_e\}$ or have infinitely mind changes if attempting to learn a particular ω -presentation of $[e : \omega, 1 : \omega]$. To do so, we wait that φ_e produces different outputs on \mathcal{A}_e and \mathcal{B}_e , and while waiting we extend \mathcal{A}_e and \mathcal{B}_e as to make their isomorphism types, in the limit, to be respectively $[e : n, 1 : \omega]$ and $[e : n + 1, 1 : \omega]$. If $\varphi_e(\mathcal{A}_e) \downarrow \neq \varphi_e(\mathcal{B}_e) \downarrow$ never happens, then φ_e does not \mathbf{InfEx}_{\cong} -learn $\{\mathcal{A}_e, \mathcal{B}_e\}$ since $\mathcal{A}_e \not\cong \mathcal{B}_e$.

Otherwise, if at some stage we obtain $\varphi(\mathcal{A}_e) \downarrow \neq \varphi(\mathcal{B}_e) \downarrow$, we add to \mathcal{A}_e two new equivalence classes of size e and to \mathcal{B}_e the same two equivalence classes and an additional one of size e . By iterating this reasoning we obtain that, if \mathcal{A}_e and \mathcal{B}_e have eventually isomorphism type $[e : \omega, 1 : \omega]$, then we can produce an ω -presentation of $[e : \omega, 1 : \omega]$ on which φ_e have infinitely many mind changes.

THE CONSTRUCTION

We build in stages strings σ and τ such that $\mathcal{A}_e = \cup_s \sigma_s$ and $\mathcal{B}_e = \cup_s \tau_s$, and a string ν . During the construction we distinguish between expansionary and nonexpansionary stages and we make use of a counter l that keeps track of the last expansionary stage.

Stage 0: Let $\sigma_0 = \tau_0 = \nu_0 =$ the empty string λ , and set $l := 0$.

Stage $s + 1$: Assume that we have built σ_s and τ_s with $\text{dom}(\sigma_s) = \text{dom}(\tau_s)$ and let z_s be $\max(\text{dom}(\sigma_s)) + 1$. We distinguish two cases.

1. If there exists $l \leq v \leq s$ such that $\varphi_{e,s+1}(\sigma_v) \downarrow \neq \varphi_{e,s+1}(\tau_v) \downarrow$, call $s + 1$ an *expansionary stage* and set $l := s + 1$. For $k \in \{0, 1, 2\}$, let

$$I_k = \{z_s + ke, z_s + ke + 1, \dots, z_s + ke + e - 1\}.$$

Let σ_{s+1} and τ_{s+1} be the following strings with domain $\{x : 0 \leq x \leq z_s + 3e\}$,

$$\sigma_{s+1}(\langle x, y \rangle) = \begin{cases} \sigma_s(\langle x, y \rangle) & x, y \in \text{dom}(\sigma_s), \\ 1 & x = y \vee (\exists k \in \{0, 1\})(x, y \in I_k), \\ 0 & \text{otherwise.} \end{cases}$$

and

$$\tau_{s+1}(\langle x, y \rangle) = \begin{cases} \tau_s(\langle x, y \rangle) & x, y \in \text{dom}(\tau_s), \\ 1 & x = y \vee (\exists k \in \{0, 1, 2\})(x, y \in I_k), \\ 0 & \text{otherwise.} \end{cases}$$

Finally, recall that $\varphi_{e,s+1}(\sigma_v) \downarrow \neq \varphi_{e,s+1}(\tau_v) \downarrow$. Without loss of generality, assume that $\varphi_{e,s+1}(\nu_l) \neq \varphi_{e,s+1}(\sigma_v)$. Define $\nu_{s+1} = \sigma_{s+1}$ (the construction of σ and τ guarantees that $\nu_{s+1} \supseteq \nu_s$).

2. Otherwise, let $\sigma_{s+1} \supseteq \sigma_s$ be the only string such that $\text{dom}(\sigma_{s+1}) = \text{dom}(\sigma_s) \cup \{z_s\}$, $\sigma_{s+1}(\langle z_s, z_s \rangle) = 1$, and $\sigma_{s+1}(\langle x, y \rangle) = 0$ if either $x = z_s$ or $y = z_s$. We define τ_{s+1} in the same way. Finally, let $\nu_{s+1} = \nu_s$.

THE VERIFICATION

It is immediate to see that $\mathcal{A}_e = \cup_s \sigma_s$ and $\mathcal{B}_e = \cup_s \tau_s$ are computable. We now distinguish two cases.

First, suppose that there exist only finitely many expansionary stages. If so, it follows from the construction that there is a number n such that \mathcal{A}_e has isomorphism type $[e : n, 1 : \omega]$ and \mathcal{B}_e has

isomorphism type $[e : n + 1, 1 : \omega]$. Towards a contradiction, suppose that φ_e **InfEx** $_{\cong}$ -learns $\{\mathcal{A}_e, \mathcal{B}_e\}$. If so, we have there are $\mathcal{M}_a \cong \mathcal{A}_e$ and $\mathcal{M}_b \cong \mathcal{B}_e$, with $\mathcal{M}_a \not\cong \mathcal{M}_b$, such that φ_e on input \mathcal{A}_e eventually converges to a , and on input \mathcal{B}_e eventually converges to b . This means that there is a stage s such that, for all $t \geq s$,

$$\varphi_e(\sigma_t) = a \neq b = \varphi_e(\tau_t).$$

But this immediately implies that there are infinitely many expansionary stages, contradicting the initial hypothesis. Hence, φ_e does not **InfEx** $_{\cong}$ -learn $\{\mathcal{A}_e, \mathcal{B}_e\}$.

Second, suppose that there are infinitely many expansionary stages. From the construction, it follows that \mathcal{A}_e and \mathcal{B}_e are isomorphic and they have isomorphism type $[e : \omega, 1 : \omega]$. The existence of infinitely many singletons in \mathcal{A}_e and \mathcal{B}_e comes from the fact that, if s is an expansionary stage, then the number $z_s + 3e$ is a singleton of \mathcal{A}_e and a singleton of \mathcal{B}_e . Let $\mathcal{C}_e = \cup_s \nu_s$. We have that also \mathcal{C}_e has isomorphism type $[e : \omega, 1 : \omega]$, since at any expansionary stage new \mathcal{C}_e -classes of size e are defined and they never expand later. If φ_e **InfEx** $_{\cong}$ -learns $\{\mathcal{A}_e, \mathcal{B}_e\}$, then φ_e has to learn \mathcal{C}_e as well, since \mathcal{C}_e is an ω -presentation of \mathcal{A}_e and \mathcal{B}_e . So, there must be some $\mathcal{M}_c \cong \mathcal{C}_e$ such that φ_e on input \mathcal{C}_e eventually converges to c . This means that there is a stage s such that, for all $t \geq s$, $\varphi_e(\nu_t) = c$. But this is impossible, since ν is constructed in such a way that, if t_1 and t_2 are consecutive expansionary stages, then $\varphi_e(\nu_{t_1}) \neq \varphi_e(\nu_{t_2})$. We conclude that φ_e does not **InfEx** $_{\cong}$ -learn \mathcal{C}_e , and thus does not learn $\{\mathcal{A}_e, \mathcal{B}_e\}$. ■

Next, we want to consider the computational complexity of learning. For this we employ the standard notation about relative computability, see (Soare, 1987). In particular, recall that $\mathbf{0}$ denotes the class of computable functions, and $\mathbf{0}^{(n+1)}$ denotes the class of functions that are computable relatively to the halting problem for $\mathbf{0}^{(n)}$ -computable functions.

Although computable learners do not have enough power to grasp all finitely separable families that are uniformly enumerable, two jumps suffice to learn equivalence structures with no infinite classes.

Theorem 9 *Let \mathfrak{A} be uniformly enumerable and such that no equivalence structure from \mathfrak{A} has infinite equivalence classes. The following are equivalent.*

1. \mathfrak{A} is finitely separable.
2. $\mathfrak{A} \in \mathbf{InfEx}_{\cong}$.
3. $\mathfrak{A} \in \mathbf{0}''\text{-InfEx}_{\cong}$

Proof “(1) \Leftrightarrow (2)” is the content of Theorem 5. The direction “(3) \Rightarrow (2)” is trivial. We prove “(2) \Rightarrow (3)”.

Assume that $\mathfrak{A} \in \mathbf{InfEx}_{\cong}$ is uniformly enumerable by f . Let M and M^* be defined as in the proof of Theorem 5; in particular, recall that M^* **InfEx** $_{\cong}$ -learns \mathfrak{A} . We follow the steps of the proof of Theorem 5 to show that there are $\mathbf{0}''$ -computable N and N^* such that N eventually coincides with M and N^* eventually coincides with M^* . Fix $\mathcal{A} \in \mathfrak{A}$.

Note that, for all s , the following set is $\mathbf{0}''$ -computable

$$X_s = \{i \leq s : \mathcal{A}[s] \hookrightarrow_{\text{fin}} \mathcal{M}_{f(i)} \wedge (\forall j \leq s)(\mathcal{M}_{f(j)} \not\hookrightarrow_{\text{fin}} \mathcal{M}_{f(i)})\}.$$

This follows from the fact that, given any two computable structures $\mathcal{A}, \mathcal{B} \in \mathbb{E}$, $\mathcal{A} \hookrightarrow_{\text{fin}} \mathcal{B}$ holds if and only if $(\forall s \exists t)(\mathcal{A}[s] \hookrightarrow \mathcal{B}[t])$, and $\mathbf{0}''$ can decide this Π_2^0 formula. Define $N(n) = \min(X_n)$. By reasoning as in the proof of Theorem 5, it is not difficult to see that $\bigcap_{n \in \omega} \mathcal{M}_{N(n)}$ is contained in the \approx_{fin} -type of \mathcal{A} .

We now observe that the family of separators of computable structures in $[\mathcal{M}_{N(n)}]_{\approx_{\text{fin}}}$ is uniformly c.e. in $\mathbf{0}''$. Indeed, given $\mathcal{M}_{f(z)} \in [\mathcal{M}_{N(n)}]_{\approx_{\text{fin}}}$, by definition of separator we obtain

$$x \in \text{sep}(\mathcal{M}_{f(z)}) \Leftrightarrow (\exists j)[\mathcal{M}_{f(j)} \approx_{\text{fin}} \mathcal{M}_{f(z)} \wedge x = \min(\text{char}(\mathcal{M}_{f(z)}) \setminus \text{char}(\mathcal{M}_{f(j)}))].$$

Since all the $\mathcal{M}_{f(i)}$'s have no infinite equivalence classes, $\text{char}(\mathcal{M}_{f(z)}) \setminus \text{char}(\mathcal{M}_{f(j)})$ is a Σ_2^0 set. Moreover, as already observed, to ask whether $\mathcal{M}_{f(j)} \approx_{\text{fin}} \mathcal{M}_{f(z)}$ holds is a Π_2^0 question. So, the overall condition is Σ_3^0 , and thus c.e. in $\mathbf{0}''$. We can conveniently approximate such separators as follows. At stage s , N^* chooses the structure in $\mathcal{F}_n = [\mathcal{M}_{N(n)}]_{\approx_{\text{fin}}} \cap \{\mathcal{M}_{f(i)}\}_{i \leq n}$ whose separator, restricted to the elements of \mathcal{F}_n , is the oldest realized by \mathcal{A}_n (notice that to check whether a given separator is realized by $\mathcal{A}[s]$ can be done effectively, since any separator is finite). N^* so defined is $\mathbf{0}''$ -computable and eventually outputs the same value of M^* , hence N^* \mathbf{InfEx}_{\cong} -learns \mathfrak{A} . \blacksquare

The next question is left open.

Question 10 *Is there a uniform enumerable family $\mathfrak{A} \in \mathbf{InfEx}_{\cong} \setminus \mathbf{0}'\text{-InfEx}_{\cong}$?*

4. Related Learning Settings

In this section we consider several learning criteria related to \mathbf{InfEx}_{\cong} and show how they compare. This provides us with many examples of families of structures which are learnable in one setting, but not another.

First, we characterize which finite collections of equivalence structures can be finitely learned.

Theorem 11 *Let \mathfrak{A} be a family of equivalence relations such that $\mathfrak{A}/_{\cong}$ is finite. The following are equivalent.*

1. $\mathfrak{A} \in \mathbf{InfFin}_{\cong}$.
2. $\forall \mathcal{A}, \mathcal{B} \in \mathfrak{A} : \mathcal{A} \hookrightarrow_{\text{fin}} \mathcal{B} \Rightarrow \mathcal{A} \cong \mathcal{B}$.
3. $\mathfrak{A}/_{\cong}$ is an anti-chain with respect to $\hookrightarrow_{\text{fin}}$.

Proof Suppose first $\mathfrak{A} \in \mathbf{InfFin}_{\cong}$, witnessed by some learner M . Let $\mathcal{A}, \mathcal{B} \in \mathfrak{A}$ with $\mathcal{A} \hookrightarrow_{\text{fin}} \mathcal{B}$. Let σ describe a part of \mathcal{A} such that M makes a (correct) conjecture for \mathcal{A} on σ . Without loss of generality, suppose that all mentioned items of σ which are equivalent in \mathcal{A} are mentioned equivalent in σ . We have \mathcal{A}_σ (the finite structure coded by σ) is now a finite substructure of \mathcal{A} , so, by supposition, $\mathcal{A}_\sigma \hookrightarrow \mathcal{B}$. This shows that σ can be extended to an informant I for (an isomorphic copy of) \mathcal{B} . Since M already makes an output on σ , the output of M on I is $M(\sigma)$. This shows that $M(\sigma)$ is a correct conjecture for \mathcal{B} , which gives

$$\mathcal{A} \cong \mathcal{M}_{M(\sigma)} \cong \mathcal{B}.$$

Suppose now $\forall \mathcal{A}, \mathcal{B} \in \mathfrak{A} : \mathcal{A} \hookrightarrow_{\text{fin}} \mathcal{B} \Rightarrow \mathcal{A} \cong \mathcal{B}$. For each $\mathcal{A} \in \mathfrak{A}$, let $K(\mathcal{A})$ be a finite substructure of \mathcal{A} such that, for all $\mathcal{B} \in \mathfrak{A}$ with $\mathcal{A} \not\cong \mathcal{B}$, $K(\mathcal{A}) \not\hookrightarrow \mathcal{B}$ (which exists since \mathfrak{A} contains

only finitely many \cong -types). Then, for all $\mathcal{A}, \mathcal{B} \in \mathfrak{A}$ with $\mathcal{A} \not\cong \mathcal{B}$ we have that neither $K(\mathcal{A})$ is embeddable into $K(\mathcal{B})$ nor vice versa. This is equivalent to 3.

Using such a list of pairwise incomparable finite substructures, we can define a learner M such that, on input σ , $M(\sigma)$ is a conjecture for $\mathcal{A} \in \mathfrak{A}$ if $K(\mathcal{A}) \hookrightarrow \mathcal{A}_\sigma$ with an embedding which may not map elements from different equivalence classes of $K(\mathcal{A})$ into elements of different equivalence classes of \mathcal{A}_σ , unless σ explicitly contains the information that these equivalence classes are different. If no such \mathcal{A} exists (which would be necessarily unique), then $M(\sigma) = ?$. Clearly, this learner \mathbf{InfFin}_{\cong} -learns \mathfrak{A} . ■

Note that the above proof extends to any equivalence relation on structures in place of \cong .

Next we show that there is a class of two structures which is learnable up to bi-embeddability, but not up to isomorphism, showing Example 3 formally.

Theorem 12 *The following holds*

$$\mathbf{InfEx}_{\cong} \subset \mathbf{InfEx}_{\approx}.$$

Proof It is obvious that any family that is learnable up to isomorphism is also learnable up to bi-embeddability. This show that $\mathbf{InfEx}_{\cong} \subseteq \mathbf{InfEx}_{\approx}$. To see that $\mathbf{InfEx}_{\approx} \not\subseteq \mathbf{InfEx}_{\cong}$, let \mathcal{A} be a structure of type $[5 : \omega]$ and \mathcal{B} a structure of type $[5 : \omega, 2 : 1]$. We have that $\mathcal{A} \approx \mathcal{B}$ and thus, trivially, $\{\mathcal{A}, \mathcal{B}\} \in \mathbf{InfEx}_{\approx}$ by a learner that always conjectures the isomorphism type of \mathcal{A} . On the other hand, \mathcal{A} is clearly a limit of \mathcal{B} (see Definition 3), and so $\{\mathcal{A}, \mathcal{B}\}$ is not finitely separable. By Theorem 5, this means that $\{\mathcal{A}, \mathcal{B}\} \notin \mathbf{InfEx}_{\cong}$. ■

Finally we consider learning without negative information, that is, learning from text rather than from informant. We establish with Theorem 13 that these two settings are the same as long as only families of structures without infinite equivalence classes are considered; Theorem 14 then shows that, for families of structures with infinite equivalence classes, we get a separation of learning power.

Theorem 13 *Let \sim be any equivalence relation on \mathbb{E} . Let \mathfrak{A} be such that none of the elements of \mathfrak{A} has an infinite equivalence class. The following are equivalent.*

1. \mathfrak{A} is \mathbf{InfEx}_{\sim} -learnable.
2. \mathfrak{A} is \mathbf{TxtEx}_{\sim} -learnable.

Proof The direction “(2) \Rightarrow (1)” is immediate. Regarding “(1) \Rightarrow (2)”, let M be a strong informant locking \mathbf{InfEx}_{\sim} -learner for \mathfrak{A} (see Theorem 21). We now describe how \mathfrak{A} can be learned from text. Given an initial part of a text σ we reorder this information as follows. First comes all positive information regarding the “first” equivalence class, defined as the equivalence class of 0. Then comes the positive information regarding the second smallest equivalence class, the class containing 1 (unless 1 was already covered by the equivalence class of 0, then we take 2 and so on). Then comes all the negative information between the first and second class (this is not present in σ , just assumed). Then all positive information of the third class, followed by the negative information between the third class on the one hand and the first and second class on the other hand, and so on,

until all elements mentioned in σ are covered. The resulting partial informant we call $\bar{\sigma}$. We can now define a learner M' learning from text as $M'(\sigma) = M(\bar{\sigma})$.

For any given $\mathcal{A} \in \mathfrak{A}$, there is an informant I which presents the data in the form described above. Let k be the largest size for which there are infinitely many equivalence classes of that size in \mathcal{A} . Since M is strong informant locking, there is n such that $I[n]$ is a strong locking sequence for M on \mathcal{A} . For any Text T , there is now an n' such that $T[n']$ contains (a) all the positive information contained in $I[n]$; (b) all positive information about the equivalence classes of \mathcal{A} that are larger than k ; and (c) all positive information about all equivalence classes containing elements that are numerically smaller than any of the elements mentioned in (a) or (b). This gives that, for all $m > n'$, we have that $\overline{T[n']}$ is extensible to an isomorphic copy of \mathcal{A} . Thus, using that $I[n]$ is a strong locking sequence, for all $m > n'$ we have that $M'(T[m]) = M(I[n])$ as desired. ■

In contrast to this result, infinite equivalence classes quickly lead to differences between text- and informant-learning.

Theorem 14 *Let $\mathfrak{A} = \{[\omega : 1], [\omega : 2]\}$. Then we have the following.*

1. \mathfrak{A} is \mathbf{InfEx}_{\cong} -learnable.
2. \mathfrak{A} is not \mathbf{TxtEx}_{\approx} -learnable.

Proof Regarding (1), the learner conjectures $[\omega : 1]$ while no negative data was given, and $[\omega : 2]$ afterwards.

Regarding (2), suppose by contradiction that a \mathbf{TxtEx}_{\approx} -learner M for \mathfrak{A} exists. Using Theorem 23, there is a strong locking sequence σ of M on $[\omega : 1]$ with corresponding embedding f . We can extend σ to a sequence τ for $[\omega : 2]$ for which $M(\tau)$ is a correct conjecture for $[\omega : 2]$. Since we can extend f trivially to \mathcal{A}_τ , we get a contradiction. ■

5. Relation to Language Learning

In this section we compare \mathbf{InfEx}_{\cong} -learning to the cardinal learning criterion of formal languages, \mathbf{TxtEx} -learning. The formal definition of this criterion can be found for example in (Jain et al., 1999), we will here use the following characterization (see (Angluin, 1980)). A collection \mathcal{L} of formal languages is \mathbf{TxtEx} -learnable iff, for all $L \in \mathcal{L}$ there is a finite $D \subseteq L$ such that, for all $L' \in \mathcal{L}$ with $D \subseteq L' \subseteq L$ we have $L' = L$. Intuitively, D signals L as the minimal extrapolating target. We call this characterization *Angluin's tell-tale criterion* and the finite sets D are called *tell-tales*.

We are now interested in somehow mapping learning tasks for \mathbf{InfEx}_{\cong} -learning into learning tasks for \mathbf{TxtEx} -learning. The next theorem shows that there cannot be a mapping taking an (isomorphism class of a) structure to a single language in order to embed \mathbf{InfEx}_{\cong} -learning into \mathbf{TxtEx} -learning.

Theorem 15 *Let a mapping Θ be given which takes a structure and returns a language such that $\forall \mathcal{A}, \mathcal{B} \in \mathbb{E} : \Theta(\mathcal{A}) = \Theta(\mathcal{B}) \Leftrightarrow \mathcal{A} \cong \mathcal{B}$. Then there is a class of structures \mathfrak{A} such that*

$$\mathfrak{A} \notin [\mathbf{InfEx}_{\cong}] \text{ and } \{\Theta(\mathcal{A}) \mid \mathcal{A} \in \mathfrak{A}\} \in [\mathbf{TxtEx}].$$

Proof Let \mathfrak{A} consist of the \cong -closure of $[5 : \omega, 2 : 1]$ and $[5 : \omega]$. We know that $\mathfrak{A} \notin [\mathbf{InfEx}_{\cong}]$. We have $\{\Theta(\mathcal{A}) \mid \mathcal{A} \in \mathfrak{A}\}$ contains only two languages, which is trivially \mathbf{TxtEx} -learnable by Angluin’s tell-tale condition. \blacksquare

In order to bypass this phenomenon, instead of associating only *one* language with any isomorphism type of structures, we can associate an *infinite* set of languages. The next theorem shows that this way we can derive an embedding. We will use the definition of a *finite permutation*, which is any permutation π of ω such that, for all but finitely many $x \in \omega$, $\pi(x) = x$.

Theorem 16 *For any structure $\mathcal{A} \in \mathbb{E}$ we define a set of languages as follows. Let $f : \omega \cup \{\omega\} \rightarrow \omega \cup \{\omega\}$ represent the isomorphism type of \mathcal{A} , that is, $[f] = [\mathcal{A}]_{\cong}$. Let $g_{\mathcal{A}}$ be any computable function (where we allow ω as a special symbol output) such that, for all $a \in \omega_+ \cup \{\omega\}$, $f(a) = |\{i \in \omega \mid g_{\mathcal{A}}(i) = a\}|$.³ For any function $h : \omega \rightarrow \omega \cup \{\omega\}$ we let $L(h) = \{\langle i, j \rangle \mid j < h(i)\}$. We let $\mathcal{L}_{\mathcal{A}} = \{L(g_{\mathcal{A}} \circ \pi) \mid \pi \text{ finite permutation}\}$.*

Then, for all $\mathfrak{A} \subseteq \mathbb{E}$ closed under \cong , the following are equivalent.

1. \mathfrak{A} is \mathbf{InfEx}_{\cong} -learnable.
2. $(\bigcup_{\mathcal{A} \in \mathfrak{A}} \mathcal{L}_{\mathcal{A}})$ is \mathbf{TxtEx} -learnable.

Proof Regarding “ \Rightarrow ”, let an \mathbf{InfEx}_{\cong} -learner M' for \mathfrak{A} be given. Now we construct a learner M which is given some sequence of data σ . From this sequence the learner constructs a sequence σ' for M so that σ' encodes the finite structure on the elements mentioned in σ , where $\langle i, j \rangle$ is equivalent to $\langle i', j' \rangle$ iff $i = i'$. Let $\mathcal{A} \in \mathbb{E}$ be the structure conjectured by M' on σ' . In some canonical listing of all finite permutations, find the minimal finite permutation π such that σ is consistent with $L(g_{\mathcal{A}} \circ \pi)$ and conjecture this language.

Let now $\mathcal{A} \in \mathfrak{A}$ and $L \in \mathcal{L}_{\mathcal{A}}$ and a text T for L be given. Let n_0 be large enough such that M' on σ' is converged on the structure \mathcal{A} . Let π be the minimal finite permutation such that $L = L(g_{\mathcal{A}} \circ \pi)$. It now suffices to show that, for all $\pi' < \pi$ in the canonical listing, $L(g_{\mathcal{A}} \circ \pi')$ is inconsistent with some finite part of L . Let $\pi' < \pi$ be given, and let i be minimal such that $g_{\mathcal{A}}(\pi(i)) > g_{\mathcal{A}}(\pi'(i))$; such an i has to exist, since either we have for all $g_{\mathcal{A}}(\pi(i)) = g_{\mathcal{A}}(\pi'(i))$, in which case π was not chosen minimal as required, or there is a difference, in which case a difference has to be found both ways, since we only consider permutations. This shows that $\langle i, g_{\mathcal{A}}(\pi'(i)) \rangle \in L \setminus L(g_{\mathcal{A}} \circ \pi')$ as desired.

Regarding “ \Leftarrow ”, we construct a learner M' for \mathfrak{A} as follows. Given an initial part σ' of an informant, we assign each element mentioned in σ' in order of (first) appearance to some $i \in \omega$: any element x known to be equivalent to a previously assigned element y is assigned to the same i as y was assigned; any element not thus assigned is assigned to the smallest i not yet used as an assignment. Let $g(i)$ be the number of elements assigned to i . Then M' checks whether there is a minimal superset of the finite set $\{\langle i, j \rangle \mid j < g(i)\}$ in $(\bigcup_{\mathcal{A} \in \mathfrak{A}} \mathcal{L}_{\mathcal{A}})$. If not, the conjecture is ?. Otherwise, let $L(g_{\mathcal{A}} \circ \pi)$ be the minimal superset and M' conjectures (a canonical index for) \mathcal{A} .

Let now $\mathcal{A} \in \mathfrak{A}$ be given and let I be an informant for \mathcal{A} . Note that, once M' conjectures an index for \mathcal{A} , M' has converged to \mathcal{A} . Otherwise a later conjecture for B would imply the existence of a finite permutation π such that $L(g_B \circ \pi)$ is consistent with known data; this would have been consistent already at the earlier point when the conjecture was \mathcal{A} , which would be a contradiction.

3. Intuitively, each equivalence class of \mathcal{A} is associated with some $i \in \omega$, and $g_{\mathcal{A}}(i)$ is the size of this equivalence class.

Let g be as produced in the limit by the construction of M' on I . Then there is a permutation π of ω (not necessarily finite) such that $g = g_A \circ \pi$. Using Angluin's tell-tale condition, let a tell-tale D for $L(g_A)$ be given. This implies that once at least the set $\{\langle \pi^{-1}(i), j \rangle \mid \langle i, j \rangle \in D\}$ has been used in the construction of M' , the output of M' will be a conjecture for \mathcal{A} as desired. ■

Acknowledgments

The authors are grateful to the referees for their thorough comments. Fokina was supported by the Austrian Science Fund, project P 27527. San Mauro was supported by the Austrian Science Fund FWF, project M 2461. Kötzing was supported by the German Research Foundation (DFG) under grant number KO 4635/1-1.

References

- U. Andrews, S. Lempp, J. S. Miller, K. M. Ng, L. San Mauro, and A. Sorbi. Universal computably enumerable equivalence relations. *The Journal of Symbolic Logic*, 79:60–88, 2014.
- D. Angluin. Inductive inference of formal languages from positive data. *Information and Control*, 45:117–135, 1980.
- C. J. Ash and J. Knight. *Computable Structures and the Hyperarithmetical Hierarchy*, volume 144 of *Studies in Logic and the Foundations of Mathematics*. North-Holland Publishing Co., Amsterdam, 2000.
- N. Bazhenov, E. Fokina, D. Rossegger, and L. San Mauro. Degrees of bi-embeddable categoricity of equivalence structures. *Archive for Mathematical Logic*, 2018. To appear.
- W. Calvert, D. Cenzer, V. Harizanov, and A. Morozov. Effective categoricity of equivalence structures. *Annals of Pure and Applied Logic*, 141:61–78, 2006.
- R. G. Downey, A. G. Melnikov, and K. M. Ng. A Friedberg enumeration of equivalence structures. *Journal of Mathematical Logic*, 17:1750008, 2017.
- E. B. Fokina, S.-D. Friedman, V. Harizanov, J. F. Knight, C. McCoy, and A. Montalbán. Isomorphism relations on computable structures. *The Journal of Symbolic Logic*, 77:122–132, 2012.
- S.-D. Friedman and L. M. Ros. Analytic equivalence relations and bi-embeddability. *The Journal of Symbolic Logic*, 76:243–266, 2011.
- C. Glymour. Inductive inference in the limit. *Erkenntnis*, 22:23–31, 1985.
- E. Gold. Language identification in the limit. *Information and Control*, 10:447–474, 1967.
- V. S. Harizanov and F. Stephan. On the learnability of vector spaces. *Journal of Computer and System Sciences*, 73:109–122, 2007.
- S. Jain, D. Osherson, J. Royer, and A. Sharma. *Systems that Learn: An Introduction to Learning Theory*. MIT Press, Cambridge, Massachusetts, second edition, 1999.

- E. Martin and D. Osherson. *Elements of scientific inquiry*. MIT Press, 1998.
- W. Merkle and F. Stephan. Trees and learning. *Journal of Computer and System Science*, 68: 134–156, 2004.
- D. Osherson, M. Stob, and S. Weinstein. *Systems that Learn: An Introduction to Learning Theory for Cognitive and Computer Scientists*. MIT Press, Cambridge, Mass., 1986.
- Robert I. Soare. *Recursively Enumerable Sets and Degrees*. Perspectives in Mathematical Logic, Omega Series. Springer-Verlag, Heidelberg, 1987.
- F. Stephan and Y. Ventsov. Learning algebraic structures from text. *Theoretical Computer Science*, 268:221 – 273, 2001.

Appendix A. Locking Sequences

Locking sequences are a powerful tool for many theorems in the area of learning theory. We first explore the concept of *weak locking sequences* which corresponds to standard locking sequences in learning languages (and the proof is standard, see (Jain et al., 1999)). Then we consider a strong variant where locking happens for a much larger class of possible extensions.

A.1. Weak Locking Sequences

We start by considering the classic locking sequences and give a straightforward generalization to arbitrary equivalence relations on the structures.

Definition 17 *Let M be a learner and \mathcal{A} a structure. We say that a sequence σ describing a finite part of \mathcal{A} is a weak locking sequence of M on \mathcal{A} iff, for every $\tau \supseteq \sigma$ describing a finite part of \mathcal{A} , we have $M(\sigma) = M(\tau)$. We distinguish between weak informant locking sequences which consist of negative and positive data (and the extensions τ are allowed positive and negative data), and weak text locking sequence which consist of positive data only (and the extensions τ are also only allowed positive data).*

Theorem 18 *Let \sim be any equivalence relation on \mathbb{E} . Suppose M \mathbf{InfEx}_{\sim} -learns a structure $\mathcal{A} \in \mathbb{E}$. Let any sequence σ_0 be given which describes a finite part of \mathcal{A} . Then there is a finite sequence $\sigma \supseteq \sigma_0$ such that σ is a weak informant locking sequence of M on \mathcal{A} . Furthermore, $\mathcal{M}_{M(\sigma)} \sim \mathcal{A}$.*

Proof Let σ_0 be given and suppose no such σ exists. Thus we can fix, for any $\sigma \supseteq \sigma_0$ describing a finite part of \mathcal{A} , $\text{ext}(\sigma)$ as an extension of σ which gives a mind change. Let I_0 be an informant for \mathcal{A} .

We define inductively

$$\begin{aligned} \forall i : \sigma_{2i+1} &= \sigma_{2i} \diamond I_0(i); \\ \forall i : \sigma_{2i+2} &= \text{ext}(\sigma_{2i+1}). \end{aligned}$$

Finally, let $I = \bigcup_i \sigma_i$. Then I is an informant for \mathcal{A} on which M does not converge, a contradiction. The “furthermore” clause of the statement follows since M needs to converge to a \sim -correct conjecture on any informant. \blacksquare

The first application of the locking sequences theorem is to prove a normal form for learners, which we will define next. This normal form can be very convenient for proofs.

Definition 19 *Let M be a learner and \mathcal{A} a structure. We call M locking on \mathcal{A} iff, for all informants I for \mathcal{A} , there is n such that $I[n]$ is a weak locking sequence for M on \mathcal{A} . Let \sim be any equivalence relation on \mathbb{E} . We call M locking iff M is locking for all $\mathcal{A} \in \mathbf{InfEx}_{\sim}$. Again we distinguish between informant locking and text locking.*

Theorem 20 *Let \sim be any equivalence relation on \mathbb{E} and let M be a \mathbf{InfEx}_{\sim} -learner. Then there is an informant locking \mathbf{InfEx}_{\sim} -learner M' which learns at least all structures learned by M .*

Proof For an input informant I we define inductively σ_0 as the empty sequence and, for all n , σ_{n+1} as σ_n concatenated with any information τ of $I[n]$ such that $M(\sigma_n) \neq M(\sigma_n \diamond \tau)$. Note that, for all \mathcal{A} which are \mathbf{InfEx}_{\sim} -learned by M , we get that $(\sigma_n)_n$ converges. Further note that σ_n can be computed from $I[n]$.

We now define a learner M' such that $M'(I[n]) = M(\sigma_n)$. Clearly, M' \mathbf{InfEx}_{\sim} -learns any structure \mathbf{InfEx}_{\sim} -learned by M . Furthermore, M' is strongly locking, since $(\sigma_n)_n$ converges. ■

By the analogous proof we get the analogous theorem for texts.

Theorem 21 *Let \sim be any equivalence relation on \mathbb{E} . Let M be a \mathbf{TextEx}_{\sim} -learner. Then there is a text locking \mathbf{TextEx}_{\sim} -learner M' which learns at least all structures learned by M .*

A.2. Strong Locking Sequences

The above theorems about locking sequences are already rather powerful and used extensively in the case of language learning. However, since a learner needs to learn *any* ω -presentation of structure, we can get even stronger locking than the ones given by Theorem 18. Here we do not only lock on extensions from the chosen concept, but also on other sequences which can be considered equivalent. This was already proven by (Martin and Osherson, 1998).

Recall that, for any σ , σ encodes a finite equivalence relation on the elements mentioned either positively or negatively by using the transitive closure of all positively mentioned pairs and assuming all other relations to be negative. This finite equivalence relation we call \mathcal{A}_{σ} .

Definition 22 *Let M be a learner and \mathcal{A} a structure. We say that a sequence σ describing a finite part of \mathcal{A} is a strong locking sequence of M on \mathcal{A} iff there is an embedding f embedding \mathcal{A}_{σ} into \mathcal{A} such that, for any τ extending σ and any g embedding from \mathcal{A}_{τ} to \mathcal{A} which extends f , $M(\sigma) = M(\tau)$. We distinguish between strong informant locking sequences which consist of negative and positive data (and the extensions τ are allowed positive and negative data), and strong text locking sequence which consist of positive data only (and the extensions τ are also only allowed positive data).*

Theorem 23 (Martin and Osherson (1998)) *Let \mathfrak{A} be a set of equivalence structures and let $\mathcal{A} \in \mathfrak{A}$. Suppose M \mathbf{InfEx}_{\approx} -learns \mathfrak{A} and let any sequence σ_0 be given which describes a finite part of \mathcal{A} . Then there is a strong informant locking sequence $\sigma \supseteq \sigma_0$ of M on \mathcal{A} .*

Proof Suppose no such σ, f exist. Thus, for any given σ, f with $\sigma \supseteq \sigma_0$, let σ', f' be first extensions τ, g found in a dovetailing search which lead to a mind change. Furthermore, for any $\sigma \supseteq \sigma_0$ and any embedding f of \mathcal{A}_{σ} into \mathcal{A} , we let $\text{ext}(\sigma, f, i)$ denote a pair (τ, g) such that τ extends σ and τ labels at least all pairs from $\{0, \dots, i\}^2$, the image of g contains $\{0, \dots, i\}$ and g is an embedding of \mathcal{A}_{τ} into \mathcal{A} ; we denote τ by $\text{ext}_{\text{seq}}(\sigma, f, i)$ and g by $\text{ext}_{\text{emb}}(\sigma, f, i)$.

Let f_0 be any embedding of \mathcal{A}_{σ_0} into \mathcal{A} . We define inductively

$$\begin{aligned} \forall i : \sigma_{i+1} &= \text{ext}_{\text{seq}}(\sigma'_i, f'_i, i); \\ \forall i : f_i &= \text{ext}_{\text{emb}}(\sigma'_i, f'_i, i). \end{aligned}$$

Finally, let $I = \bigcup_i \sigma_i$, \mathcal{B} the structure described by I , and $f = \bigcup_i f_i$. Then f witnesses $\mathcal{A} \cong \mathcal{B}$ and I witnesses that M does not learn \mathcal{B} , a contradiction. ■

We can now have the analogous definition about locking learners as for weak locking sequences.

Definition 24 Let M be a learner and \mathcal{A} a structure. We call M strong-locking on \mathcal{A} iff, for all informants I for \mathcal{A} , there is n such that $I[n]$ is a strong locking sequence for M on \mathcal{A} . Let \sim be any equivalence relation on \mathbb{E} . We call M strong locking iff M is locking for all $\mathcal{A} \in \mathbf{InfEx}_{\sim}$. Again we distinguish between strong informant locking and strong text locking.

Theorem 25 Let \sim be any equivalence relation on \mathbb{E} and let M be a \mathbf{InfEx}_{\sim} -learner. Then there is a strong informant locking \mathbf{InfEx}_{\sim} -learner M' which learns at least all structures learned by M .

Proof Analogous to the proof of Theorem 20. ■