Isomorphism and Classification for Countable Structures

Russell Miller [*]

July 19, 2018

Abstract

We introduce a topology on the space of all isomorphism types represented in a given class of countable models, and use this topology as an aid in classifying the isomorphism types. This mixes ideas from effective descriptive set theory and computable structure theory, extending concepts from the latter beyond computable structures to examine the isomorphism problem on arbitrary countable structures. We give examples using specific classes of fields and of trees, illustrating how the new concepts can yield classifications that reveal differences between seemingly similar classes. Finally, we use a computable homeomorphism to define a measure on the space of isomorphism types of algebraic fields, and examine the prevalence of relative computable categoricity under this measure.

1 Introduction

Computable structure theorists have often restricted their study to computable models of the particular theory or class in which they are interested. Attempts to measure the difficulty of determining whether two given structures are isomorphic, such as [4, 11], frequently consider only computable

1

structures. Indeed, sometimes the goal was to distinguish the computably presentable structures from the others. [5], [6], and [14] went further, defining computable transformations of one class of countable structures into another, a notion that we will adapt here to a broader context. The question of actually computing an isomorphism, in case the structures are isomorphic, was initially posed for computable structures, as the notion of computable categoricity, although its subsequent extension to all countable structures (as *relative computable categoricity*) was shown in [2, 3] to have a far simpler syntactic characterization than computable categoricity itself has. In [7, 11, 16]. questions of classification up to isomorphism are also restricted to computable structures.

Of course, the class of *all* structures isomorphic to a given structure is a proper class, and hence unwieldy for study. Here, however, given a class of countable structures, we will consider all those representatives of the class which have the domain ω. This is the natural context in which to study questions of Turing computability on countable structures, since n-ary functions and relations in the structure may then be viewed as subsets of ω^n, susceptible to analysis by computability theorists. One might request that subsets of ω be allowed as domains, but this essentially only matters when we wish to include finite structures in our classes; otherwise, given the domain itself, we may compute an isomorphism from the structure onto a structure whose domain is all of ω.

In this article we will address several classes of structures with domain ω. First we examine $\mathcal{ALG}_0$, the class of all algebraic field extensions of the field $\mathbb{Q}$ of rational numbers. From there we move on to algebraically closed extensions of $\mathbb{Q}$ (of arbitrary transcendence degree $\leq \omega$), which form a class usually viewed as simpler than the algebraic fields. Last we consider the class $\mathcal{FBT}$ of finite-branching trees, under the predecessor function. This class bears distinct similarities to $\mathcal{ALG}_0$, but in the context of all countable models, we will be able to draw distinctions between them which were not so clear when only computable models were considered. The goal is to develop and illustrate the usefulness of studying all countable models in the class. We have deliberately chosen fairly simple classes for this purpose, but nevertheless, several intriguing questions arose out of these investigations, which are discussed in the final three sections. It is hoped that subsequent work will extend this study to classes with more complex isomorphism problems.

2 Known Results on Computable Structures

Here we set up our notation and review some relevant background, mainly about computable structures. We give all specific definitions for the class of algebraic field extensions of $\mathbb{Q}$; those for algebraically closed fields and finite-branching trees are analogous. Recall that the *atomic diagram* of a structure $\mathcal{A}$ is the set of all atomic sentences true in $\mathcal{A}$, in the language of $\mathcal{A}$ extended by new constants for elements of the domain (which for us is always ω). For all structures on the domain ω in this signature, we use a single effective bijective Gödel coding to view the atomic diagram as a subset of ω.

Definition 2.1

$\mathcal{ALG}_0 = \{D \in 2^\omega : D \text{ is the atomic diagram of a field algebraic over } \mathbb{Q}\}$.

For $D \in \mathcal{ALG}_0$, F_D is the field with domain ω and atomic diagram D.

$\mathcal{I}(\mathcal{ALG}_0) = \{(D_0, D_1) \in (\mathcal{ALG}_0)^2 : F_{D_0} \cong F_{D_1}\}$.

$\mathcal{ALG}_0^c = \{e \in \omega : (\exists D \in \mathcal{ALG}_0) \ \chi_D = \varphi_e\}$.

For $e \in \mathcal{ALG}_0^c$, F_e is the field F_D for which $\chi_D = \varphi_e$.

$\mathcal{I}(\mathcal{ALG}_0^c) = \{(e_0, e_1) \in (\mathcal{ALG}_0^c)^2 : F_{e_0} \cong F_{e_1}\}$.

$\mathcal{ACF}_0 = \{D \in 2^\omega : D \text{ is the atomic diagram of a model of } \mathbf{ACF}_0\}$.

$\mathcal{FBT} = \{D : D \text{ is the atomic diagram of an infinite finite-branching tree}\}$.

The language for the trees in $\mathcal{FBT}$ is the language with equality and one unary function symbol P, the *predecessor function*. The root r of the tree is the unique node satisfying $P(r) = r$, while for every other x, $P(x)$ is the immediate predecessor of x in the tree. Of course, in this language, the greatest possible height for a tree is ω, and in fact all trees in $\mathcal{FBT}$ have height ω, by König's Lemma, since they are infinite but finite-branching. The tree is finite-branching if, for each node x, only finitely many y satisfy $P(y) = x$. Finite-branchingness for trees cannot be stated in the first-order languages we have given; nor can algebraicity for fields. The 0 in $\mathcal{ALG}_0$ and $\mathcal{ACF}_0$ is to specify the characteristic. We will not discuss fields of positive characteristic here, but they behave in essentially the same way, except that in $\mathcal{ALG}_p$ it would be helpful to allow finite structures.

It is well-known that $\mathcal{I}(\mathcal{ALG}_0^c)$ and the analogously defined set $\mathcal{I}(\mathcal{FBT}^c)$ are both Π_2^0-complete. The fact that they are Π_2^0 stems from the following characterizations, which essentially are consequences of König's Lemma.

Lemma 2.2 *Two algebraic fields of equal characteristic are isomorphic if and only if every finitely generated subfield of each one embeds into the other (equivalently, if every polynomial over the prime subfield with a root in either field also has a root in the other). Two finite-branching trees are isomorphic if and only if every finite subtree of each one embeds into the other.*

Moreover, these equivalence relations (on the relevant subsets of ω, namely $\mathcal{ALG}_0^c$ and $\mathcal{FBT}^c$) are computably reducible to each other, and also computably bireducible with the equivalence relation $=^e$ on ω defined by

$$i =^e j \iff W_i = W_j.$$

This means that there are computable functions f (which is total, with image $\subseteq \mathcal{ALG}_0^c$) and g (with domain $\supseteq \mathcal{ALG}_0^c$) such that

$$(\forall i, j) \; [i =^e j \iff F_{f(i)} \cong F_{f(j)}]$$
$$\& \quad (\forall i, j \in \mathcal{ALG}_0^c) \; [F_i \cong F_j \iff g(i) =^e g(j)],$$

and likewise for $\mathcal{FBT}^c$. These functions f and g are *reductions* between the equivalence relations involved, in the sense described in [17], for example. The function f here is defined by setting $f(i)$ to be the index for an atomic diagram of the field $\mathbb{Q}[\sqrt{p_n} \; : \; n \in W_i]$, where p_n is the n-th prime number. To define the function g, fix a listing $h_0, h_1, \ldots$ of all irreducible polynomials in $\mathbb{Q}[X]$, and let $W_{g(i)} = \{n \in \omega : h_n \text{ has a root in } F_i\}$. Likewise, to reduce $=^e$ to $\mathcal{I}(\mathcal{FBT}^c)$, let $f(i)$ be the index of the atomic diagram of the tree with a single infinite "spine" of length ω and with one other node at level $n + 1$ (above the spinal node at level n) just if $n \in W_i$; for the opposite reduction, use a computable list of all finite trees and consider the set of those which embed into the given tree.

3   Algebraic Fields

The reductions described above between $\mathcal{ALG}_0^c$ and $\mathcal{FBT}^c$ were computable reductions on equivalence relations on subsets of ω. We now switch to the notion of a computable reduction from one equivalence relation on one subset of Cantor space to another equivalence relation on another such subset, as has been done in [17] and other places. This can be seen as a generalization of [6, Defn. 8], where it is done for the relation of isomorphism on classes

such as $\mathcal{ALG}_0$ and $\mathcal{FBT}$. Here we will allow ourselves to work with arbitrary equivalence relations on Cantor space or Baire space, and our goal will be not merely reductions, but actual homeomorphisms. To begin with, therefore, we explain the topologies in question, with $\mathcal{ALG}_0$ as our first example.

The sets $\mathcal{ALG}_0$, $\mathcal{ACF}_0$, and $\mathcal{FBT}$ are all subsets of Cantor space, so each one becomes a topological space in its own right, under the subspace topology. The basic open sets in the topology on $\mathcal{ALG}_0$, for instance, are defined to be the sets of the form

$$\mathcal{U}_\sigma = \{D \in \mathcal{ALG}_0 : \sigma \subset D\},$$

with σ ranging over all initial segments of elements of $\mathcal{ALG}_0$ (or over all of $2^{<\omega}$, in which case many of the sets above are empty). Then, under the equivalence relation of isomorphism on $\mathcal{ALG}_0$, defined by $\mathcal{I}(\mathcal{ALG}_0)$, we get the quotient topology on the quotient space $\mathcal{ALG}_0/\cong$, sometimes known as the *isomorphism space* of $\mathcal{ALG}_0$. By definition this means that a subset $\mathcal{V}$ of $\mathcal{ALG}_0/\cong$ is open there if and only if its preimage under the quotient map

$$\mathcal{U} = \{D \in \mathcal{ALG}_0 : \text{the isomorphism type of } F_D \text{ lies in } \mathcal{V}\}$$

is open in the topology on $\mathcal{ALG}_0$. Since such an open set $\mathcal{U}$ must be the union of certain basic open sets $\mathcal{U}_\sigma$, we see that an open set $\mathcal{V}$ in $\mathcal{ALG}_0/\cong$ is defined by taking any set of initial segments σ of various $D \in \mathcal{ALG}_0$ and letting $\mathcal{V}$ contain all isomorphism types of fields F_D with $D \in \mathcal{U}_\sigma$, for all these strings σ. That is, $\mathcal{V}$ contains all isomorphism types such that one of these σ can be an initial segment of a field of that type. (The analogous statement holds for any other class of structures on the domain ω, of course.)

Now a string σ may contain enough information to ensure that all its extensions in $\mathcal{ALG}_0$ contain roots of one or of several polynomials: it might state, for example, that the square of some element, when added to the multiplicative identity, equals the additive identity, i.e., that $X^2 + 1$ has a root. It might also give incomplete information. For example, σ could specify that the field contains elements x and y with $x^8 = y^{12} = 2$, but without specifying the minimal polynomial of y over $\mathbb{Q}(x)$. Since $Y^{12} - 2$ factors as $(Y^3 - x^2)(Y^3 + x^2)(Y^6 + x^4)$ over $\mathbb{Q}(x)$, we cannot determine even the degree of the extension dictated by σ, let alone its isomorphism type. However, there are only finitely many possibilities, and in any particular field, some finite piece of the atomic diagram must specify the minimal polynomial of y over $\mathbb{Q}(x)$, so in practice we normally just choose a string sufficiently long to specify the isomorphism type of the relevant subfield.

5

Since a single σ is finite, it can only specify finitely many extensions of finite degree to be contained in the field. So an open set $\mathcal{V}$ in $\mathcal{ALG}_0/\cong$ will contain every isomorphism type realizing the finite-degree extensions dictated by σ, as σ ranges over some set of strings. Of course, by the Primitive Element Theorem, every finite-degree algebraic extension of $\mathbb{Q}$ is generated by a single element, so it is natural to think of $\mathcal{V}$ as being defined by a particular subset S of $\mathbb{Q}[X]$, corresponding to the set of strings σ defining its preimage $\mathcal{U}$: $\mathcal{V}$ is the set of all isomorphism types containing a root of any polynomial in S. (For structures more generally, $\mathcal{V}$ will contain every isomorphism type realizing any of the finite partial quantifier-free n-types defined by these strings σ.) Conversely, for any given finite-degree extension K of $\mathbb{Q}$, we can choose a primitive generator y of K and its minimal polynomial $p(Y)$ over $\mathbb{Q}$. Take the open subset of $\mathcal{ALG}_0$ given by

$$\mathcal{U} = \bigcup_\sigma \mathcal{U}_\sigma,$$

where the union is over all σ which state that, for some $y < |\sigma|$, $p(y) = 0$. Thus the image $\mathcal{V}$ of $\mathcal{U}$ under the quotient map is an open set (since $\mathcal{U}$ is closed under isomorphism) and will contain exactly those algebraic fields extending K. This yields the following.

Lemma 3.1 *The topology on $\mathcal{ALG}_0/\cong$ has a basis consisting of those sets defined by the existence of a root for one irreducible polynomial from $\mathbb{Q}[Y]$. Equivalently, these are the sets of the form $\{[D] \in \mathcal{ALG}_0/\cong: K \subseteq F_D\}$, as K ranges over all finite-degree field extensions of $\mathbb{Q}$.*

It is important to notice that these sets do form a basis, not just a subbasis. To understand the potential difficulty, consider the sets $\mathcal{V}$ and $\mathcal{W}$ containing, respectively, those fields with an eighth root of 2 and those with a twelfth root of 2. As seen earlier, the intersection $\mathcal{V} \cap \mathcal{W}$, which is nonempty, cannot be defined by the existence of any single subfield; there are three fields K_0, K_1, and K_2, pairwise incomparable under embedding, such that the fields in $\mathcal{V} \cap \mathcal{W}$ are those containing at least one of these three as a subfield, but it is impossible to give two fields, let alone just one, which characterize $\mathcal{V} \cap \mathcal{W}$ in the same way. (The three fields correspond to the three factors of $Y^{12} - 2$ over $\mathbb{Q}(x)$, seen above.). However, if $\mathcal{V}_i$ contains all isomorphism types of fields extending K_i, then each $\mathcal{V}_i$ is an element of the basis in Lemma 3.1, and $\mathcal{V} \cap \mathcal{W} = \mathcal{V}_0 \cup \mathcal{V}_1 \cup \mathcal{V}_2$.

Corollary 3.2 *In the topology on $\mathcal{ALG}_0/\cong$, the only open set containing the isomorphism type of $\mathbb{Q}$ is the entire space $\mathcal{ALG}_0/\cong$, whereas the type of the algebraic closure $\overline{\mathbb{Q}}$ lies in every nonempty open set.* ∎

This corollary is not unique to $\mathcal{ALG}_0$; it would hold in any class $\mathfrak{C} \subseteq 2^\omega$ of structures, provided that $\mathfrak{C}$ contains both a least and a greatest structure under embedding. It can readily be related to [6, Prop. 4.1], re-expressed in terms of topology. As a topological space, $\mathcal{ALG}_0/\cong$ is not merely compact: every open cover has a subcover of size 1, since some element of the cover is an open set containing the type of $\mathbb{Q}$. However, the corollary does establish that this topology on $\mathcal{ALG}_0/\cong$ is not homeomorphic to any of the best-known topological spaces of size continuum. We discuss this further in Section 8.

To address this shortcoming, we now adjoin to the language $\mathcal{L}$ of fields, for each $n > 1$, the n-ary *root predicate* R_n, defined as follows:

$$\models_\mathbb{F} R_n(a_0, \ldots, a_{n-1}) \iff (\exists x \in \mathbb{F}) x^n + a_{n-1}x^{n-1} + \cdots + a_0 = 0.$$

Theorem 3.3 *In the language $\mathcal{L}^*$ with the predicates R_n, the class $\mathcal{ALG}_0^*/\cong$ of all algebraic fields of characteristic 0, modulo isomorphism, is computably homeomorphic to Cantor space 2^ω (under its usual topology). In particular, it is compact.*

In this context, a homeomorphism h from $\mathcal{ALG}_0^*/\cong$ onto 2^ω is *computable* if there is one Turing functional Γ such that, for every $D \in \mathcal{ALG}_0^*$, $\Gamma^D = h([D])$, the image of the equivalence class of D, and a second Turing functional Φ such that, for every $D \in \mathcal{ALG}_0^*$, $\Phi^{h([D])}$ lies in the class $[D]$. Every homeomorphism is continuous, hence S-computable for some $S \subseteq \omega$, but we assert here that for this homeomorphism, no oracle is required. (We refer readers to [24] for details of the relation between computability and continuity.)

Proof. Fix a computable copy E of the algebraically closed field $\overline{\mathbb{Q}}$, and let $f_0, f_1, \ldots$ enumerate all monic polynomials in $E[X]$ of degree > 1. For each $\sigma \in 2^{<\omega}$, we now define a subfield $F_\sigma \subseteq E$ and a polynomial f_σ from this list. The root node λ of $2^{<\omega}$ has $F_\lambda = \mathbb{Q}$, the prime subfield of E. For each σ, starting with λ, we now define f_σ recursively to be f_n for the least n such that:

- f_n lies in $F_\sigma[X]$ (within $E[X]$) and is irreducible of prime degree there; and

7

- for the least root x of f_n in E, and for every τ with $\tau^\frown 0 \subseteq \sigma$, the subfield $F_\sigma(x)$ of E contains no root of f_τ.

A theorem of Kronecker from [15] gives us splitting algorithms for every finitely generated subfield of E, uniformly in the generators, so these conditions are both decidable. (For details see [18, Thm. 2.1].) Since F_σ is finitely generated, there must exist some f_n satisfying these conditions. This defines f_σ, and we now set $F_{\sigma^\frown 0} = F_\sigma$ and $F_{\sigma^\frown 1} = F_\sigma(x)$, for the least root x of f_σ in E. (The roots of f_σ are elements of the domain ω of E, so by the "least root," we mean the smallest element of ω which is a root.) This completes the recursion, defining F_σ and f_σ for every σ.

We now define the functionals Γ and Φ, which will be inverses of each other when viewed as maps between $\mathcal{ALG}_0^*/\cong$ and 2^ω. Γ accepts as input the atomic diagram D of any field F_D in $\mathcal{ALG}_0^*$, in the language with the predicates R_n. Starting with the empty string σ_0 and the unique embedding g_0 of $F_{\sigma_0} = \mathbb{Q}$ into F_D, it defines each string $\sigma_{m+1} \in 2^{m+1}$ to extend σ_m by setting:

- $\sigma_{m+1}(m) = 1$ if F_D contains a root of f_{σ_m}, in which case g_{m+1} extends g_m by mapping the least root of f_{σ_m} in $F_{\sigma_{m+1}}$ to the least root in F_D of the g_m-image of f_{σ_m} in $F_D[X]$;

- $\sigma_{m+1}(m) = 0$ if not, in which case $g_{m+1} = g_m$.

So in both cases g_{m+1} embeds $F_{\sigma_{m+1}}$ into F_D. With the root predicate in the language, this procedure is effective: D allows Γ^D to determine the correct value for $\sigma_{m+1}(m)$, extending σ_m to length $m+1$. The output of Γ is the function $h = \cup_m \sigma_m$, and $g = \cup_m g_m$ is an isomorphism from the subfield $\cup_m F_{h \restriction m}$ of E onto F_D. Notice that for each m, the g-image of $f_{h \restriction m}$ in $F_D[X]$ has a root in F_D if and only if $h(m) = 1$.

Φ accepts any $h \in 2^\omega$ as oracle, and defines a field $K = \cup_m F_{h \restriction m}$, using the fields F_σ defined above. This is done in such a way that ω is the domain of K (even if h is the constant function 0.) Moreover, with h as oracle, Φ^h will eventually decide whether a given polynomial $X^d + a_{d-1}X^{d-1} + \cdots + a_0$ in $K[X]$ has a solution in K, because this polynomial (or factors of it) will eventually appear as polynomial(s) $f_{h \restriction m}$ for various m, and the existence of a root of $f_{h \restriction m}$ in K will be determined by the value of $h(m)$. Thus Φ^h decides the atomic diagram of the field K it builds, even in the expanded language with the root predicates. Moreover, applying Γ to the atomic diagram of this

field will clearly yield the same h we started with, and likewise, applying Φ to the function Γ^D will yield a field isomorphic to F_D. So the two computable functions are inverses, and we have the desired homeomorphism. ∎

Notice that this proof would not work if we did not have the root predicates in our language. As noted above, the structure $\overline{\mathbb{Q}}$ lies in every nonempty open set of $\mathcal{ALG}_0/\cong$, so clearly $\mathcal{ALG}_0/\cong$ and $\mathcal{ALG}_0^*/\cong$ cannot be homeomorphic. Having the root relations in the language allows the atomic diagram of $\mathbb{Q}$ to differentiate itself from those of other algebraic fields, as the failure of these relations to hold (on specific elements of $\mathbb{Q}$) expresses the absence from $\mathbb{Q}$ of roots of certain polynomials. There does exist a continuous bijection from $\mathcal{ALG}_0^*/\cong$ onto $\mathcal{ALG}_0/\cong$, sending each isomorphism type to itself; indeed, this bijection is computable, mapping the atomic diagram of any F_D^* to the atomic diagram of its reduct. (It follows that $\mathcal{ALG}_0/\cong$ is also compact, as we already knew.) However, the inverse map is not continuous.

It is important to notice that, in order to make $\mathcal{ALG}_0$ homeomorphic to 2^ω, we had to be careful not to add too *much* information to the language either. It would have been very natural to extend each field F in $\mathcal{ALG}_0$ to its *jump* F', as defined in [22, 23] by Montalbán, Soskov, and Soskova. The root predicates R_n all belong to the signature of the jump of a field, and so we would still have a computable function Γ from the class $\mathcal{ALG}_0^{(1)}/\cong$ of jumps of fields in $\mathcal{ALG}_0$, modulo isomorphism, onto Cantor space; it would be exactly the same map as from $\mathcal{ALG}_0^*$ onto 2^ω. However, in $\mathcal{ALG}_0^{(1)}/\cong$, the isomorphism type of $\mathbb{Q}$ forms an open set all by itself, because the language of the jump of a field contains a predicate for the following computable existential $L_{\omega_1\omega}$ statement about the field:

$$(\exists p \in \mathbb{Q}[X])(\exists x)[p \text{ is irreducible of degree} > 1 \ \& \ p(x) = 0].$$

A field $F \in \mathcal{ALG}_0$ is isomorphic to $\mathbb{Q}$ just if the predicate for this statement is false in F'. Thus the topological space $\mathcal{ALG}_0^{(1)}/\cong$ contains a singleton open set, whereas 2^ω does not, and so there can be no homeomorphism between them. In particular, the inverse of the map Γ is not computable, nor even continuous. As an insight into this failure, notice that, with the constant function 0 as h, an inverse function Φ^h would have to compute a copy of the jump of the field $\mathbb{Q}$, and this copy would have to compute the degree $\mathbf{0}'$, as all jumps of structures do; but it is clearly impossible for a computable functional Φ, with the computable oracle h, to output an atomic diagram which decides $\mathbf{0}'$. Of course, a continuous inverse function might require an

9

oracle S, i.e., might only be given by an S-computable functional Φ^S, but a similar argument with a field whose spectrum is the upper cone above $\deg(S)$ shows this also to be impossible.

Proposition 3.4 *Let $\mathfrak{C} \subseteq 2^\omega$ be a class of structures, closed under isomorphism. Then, for each structure $\mathcal{D}$ with atomic diagram $D \in \mathfrak{C}$, the singleton of the isomorphism type of $\mathcal{D}$ forms an open set in $\mathfrak{C}/\cong$ if and only if there is some (finitary) Σ_1 formula $\exists \vec{x}\alpha$ in the language which is a Scott sentence for $\mathcal{D}$ within the class $\mathfrak{C}$. (This means that, among structures in $\mathfrak{C}$, α holds just in those structures isomorphic to $\mathcal{D}$.)*

To clarify: in the example with $\mathbb{Q}$ in $\mathcal{ALG}_0^{(1)}$, the Scott sentence is actually a quantifier-free sentence α in the language of jumps of fields. In the more general context, however, finitary Σ_1 formulas may be necessary.

Proof. First, if there is such a Scott sentence $\exists \vec{x}\alpha(x_1, \ldots, x_n)$ in the language, then for every $\mathcal{C} \cong \mathcal{D}$ with $\Delta(\mathcal{C}) \in \mathfrak{C}$, we can choose $(k_1, \ldots, k_n) \in \omega^n$ realizing α in $\mathcal{C}$ and let $\sigma_{\mathcal{C}}$ be an initial segment of $\Delta(\mathcal{C})$ sufficiently long to state that $\models_{\mathcal{C}} \alpha(k_1, \ldots, k_n)$. Then $\mathcal{U}_{\sigma_{\mathcal{C}}} \cap \mathfrak{C}$ is an open set of $\mathfrak{C}$ containing only diagrams of structures in which $\exists \vec{x}\alpha$ holds, including $\mathcal{D}$ itself, and the union of all these $\mathcal{U}_{\sigma_{\mathcal{C}}}$ (over all $\mathcal{C} \cong \mathcal{D}$ in $\mathfrak{C}$) is an open set whose image in $\mathfrak{C}/\cong$ is the singleton of the class of $\mathcal{D}$. Thus this singleton is an open set in $\mathfrak{C}/\cong$.

Conversely, suppose that the isomorphism type of $\mathcal{D}$ forms a singleton open set in $\mathfrak{C}/\cong$. Then its preimage $\mathcal{U}$ in $\mathfrak{C}$ is open and contains the atomic diagram D of $\mathcal{D}$, but contains no atomic diagrams of any structures $\not\cong \mathcal{D}$. Let $\sigma \subset D$ be sufficiently long that $\mathcal{U}_\sigma \cap \mathfrak{C}$ is contained in this open preimage, and let α be the conjunction of all atomic statements in $\sigma^{-1}(1)$ and all negations of atomic statements in $\sigma^{-1}(0)$. Replacing each element j of ω used in α by the variable x_j yields a finitary Σ_1 formula

$$\exists x_0 \cdots \exists x_{|\sigma|}\ \alpha(x_0, \ldots, x_{|\sigma|})$$

which holds in $\mathcal{D}$. Whenever $\mathcal{C}$ is a structure (with diagram in $\mathfrak{C}$) realizing this Σ_1 formula, there must be some $\mathcal{B} \cong \mathcal{C}$ such that the atomic diagram of $\mathcal{B}$ begins with σ: just use a permutation of ω mapping the elements of $\mathcal{C}$ realizing α onto $(0, \ldots, |\sigma|)$. Since $\mathfrak{C}$ is closed under isomorphism, the diagram of $\mathcal{B}$ lies in $\mathcal{U}_\sigma \cap \mathfrak{C}$, hence lies in $\mathcal{U}$, forcing $\mathcal{B} \cong \mathcal{D}$. Thus $\exists \vec{x}\alpha$ is indeed a Scott sentence for $\mathcal{D}$ within $\mathfrak{C}$. ∎

We will discuss singleton open sets further in the next section.

4   Algebraically Closed Fields

We postponed consideration of algebraically closed fields until after algebraic fields because, although the former is a simpler class, the latter is more representative of the usual situation. An algebraically closed field is determined up to isomorphism by its characteristic and its transcendence degree; the field is countable just if the latter is $\leq \omega$. So there are only countably many isomorphism classes of countable algebraically closed fields. We will restrict our discussion to those of characteristic 0.

With only the usual language of fields, the goal is to count the number of elements in a transcendence basis of an ACF F. Recognizing transcendentals is not a decidable process, so we first consider the easier situation in which the language is enriched by n-ary *dependence predicates* D_n, for every $n > 0$, which hold of an n-tuple from F if and only if that tuple is algebraically dependent over the prime subfield $\mathbb{Q}$. Now a field F (with domain ω) has a natural transcendence basis, containing exactly those elements x which are transcendental over all preceding elements of the domain ω of F:

$$(\forall p \in \mathbb{Q}[Y_0, \dots, Y_x]) \ [p(0, 1, \dots, x - 1, x) = 0 \implies p(0, 1, \dots, x - 1, Y_x) = 0.$$

This basis is not invariant under automorphisms, of course, but its elements are recognizable from the atomic diagram (in the language with the predicates D_n). So it is readily seen that for the class $\mathcal{ACF}_0^D$ of algebraically closed fields of characteristic 0, with domain ω, in the language with the dependence predicates, the space $\mathcal{ACF}_0^D/\cong$ is homeomorphic to $2^\omega/E_{\mathrm{card}}$, where the equivalence relation E_{card} (the *cardinality relation*) is defined by:

$$A \ E_{\mathrm{card}} \ B \iff |A| = |B|.$$

In one direction, a field F is mapped to the transcendence basis defined above; while in the reverse direction, each time we find a new element of the oracle set A, we adjoin a new transcendental to the field we are building, while always taking one further step to make the field algebraically closed. Notice that this does allow computation of a transcendence basis of the field (namely, the set of elements added in response to the discovery of new elements of A), from which in turn one can compute the dependence relations on the field.

The cardinality relation is not among the usual Borel equivalence relations on 2^ω. It has only countably many equivalence classes, among which one is a singleton (the class of $\emptyset$), one has size continuum, and all others are countable.

E_{card} is not *smooth*: the equality relation on 2^ω has no Borel reduction to E_{card}, as E_{card} has too few classes. In the topology on $2^\omega/E_{\text{card}}$, the basic open sets are those defined by having at least n elements, and these are the only nonempty open sets in the space:

$$\mathcal{U}_n = \{A \in 2^\omega : |A| \geq n\}.$$

Therefore, the E_{card}-class of infinite sets belongs to every nonempty open set, whereas the E_{card}-class of $\emptyset$ belongs to no open set except the entire space. The same holds in $\mathcal{ACF}_0^D/\cong$: the prime model $\overline{\mathbb{Q}}$ belongs to no open set except the whole space, whereas the algebraically closed field with transcendence degree ω belongs to every nonempty open set. The open sets are just those closed upwards under transcendence degree.

Of course, by adding more definable relation symbols to the language, one could change this topology. The transcendence degree itself is $L_{\omega_1\omega}$-definable, after all. One specific method is to add n-ary predicates B_n, which hold of an n-tuple $\vec{a}$ if and only if $\vec{a}$ is a transcendence basis for F. (It is helpful to include B_0, which states $\forall x D_1(x)$; this is essentially 0-ary, and holds only in transcendence degree 0.) In the space $\mathcal{ACF}_0^*/\cong$ of algebraically closed fields in this language, the isomorphism type of a field of finite transcendence degree will itself form an open set, but the space is not quite discrete, as the only open sets containing the type of infinite transcendence degree are those sets which, for some d, contain all types of degree $\geq d$. Thus this space is still compact. To make the singleton of the infinite-degree type open as well (and thus make the space noncompact), one must add a predicate C defining it, such as $(\forall d)(\exists \vec{x})\neg D_d(\vec{x})$. Once again this is only $L_{\omega_1\omega}$-definable.

Several principles are illustrated here. First, this process recalls Proposition 3.4: in a space $\mathfrak{D}/\cong$, the singleton set containing just the isomorphism type of a structure $\mathcal{A}$ from $\mathfrak{D}$ is an open set if and only if $\mathcal{A}$ has a Scott sentence which is (finitary) Σ_1 in the language of $\mathfrak{D}$. Such a sentence says that among structures in $\mathfrak{D}$, being isomorphic to $\mathcal{A}$ is equivalent to containing a finite collection of elements satisfying a certain configuration, and so a finite piece of the atomic diagram of $\mathcal{A}$ is sufficient to guarantee isomorphism to $\mathcal{A}$. Each predicate B_n above yielded a Scott sentence for the algebraically closed field of transcendence degree n (among algebraically closed fields of characteristic q), but the infinite-degree field required a bit more work.

Second, once all the Scott sentences (including for infinite degree) were adjoined to the language, the space became the discrete countable topological

space, hence noncompact. Consider the general situation of a class $\mathfrak{D}$ of countable structures and a class $\mathfrak{D}^*$ of the same structures in a language extended by predicates which are $L_{\omega_1\omega}$-definable in the original language. Thus $\mathfrak{D}/\cong$ and $\mathfrak{D}^*/\cong$ have an obvious bijective correspondence as sets, and the map from $\mathfrak{D}^*/\cong$ onto $\mathfrak{D}/\cong$ is clearly computable, since the atomic diagram of a reduct can easily be computed from the atomic diagram of the structure. (If the languages were countable but noncomputable, then this reduction might not be computable either, yet it is still continuous, requiring only an oracle for the languages.) However, the inverse map, from the reduct $\mathcal{A}$ to the original $\mathcal{A}^*$, may well not be computable, and often is not even continuous, as seen in the examples above, where the topological spaces $\mathfrak{D}^*/\cong$ and $\mathfrak{D}/\cong$ are not homeomorphic. Therefore, although compactness of the space in the larger language implies compactness of the space in the smaller language (as the continuous image of a compact space is always compact), the converse fails. The class $\mathcal{ACF}_0$ and its expansions provide a ready counterexample.

The situation above resembles the well-known case of the *jump map*, taking a set $A \in 2^\omega$ to its jump A'. This map is not continuous – its points of continuity are precisely the 1-generic sets – but its inverse is readily seen to be a continuous map on the domain $\{A' : A \in 2^\omega\}$, and in fact extends to a continuous map on all of 2^ω.

5 Finite-Branching Trees and Graphs

For us, the language of trees has equality and a single unary function symbol P, the *predecessor function*. We define $P(r) = r$ for the root r of the tree, for convenience and also to enable the root to be identified from the atomic diagram. We stipulate that for every x in $\mathcal{T}$, there must exist an n with $P^{n+1}(x) = P^n(x)$, meaning that every node is only finitely many levels above the root. (The *level* of x is the least n for which this holds.) A tree $\mathcal{T}$ is *finite-branching* if, for every $x \in \mathcal{T}$, the inverse image $P^{-1}(x)$ is a finite set (possibly empty). These definitions use $L_{\omega_1\omega}$ formulas, of course; they cannot be expressed finitarily. The class of all infinite finite-branching trees with domain ω is denoted by $\mathcal{FBT}$.

We enrich the class to $\mathcal{FBT}^*$ by adjoining to this language, for each $n > 0$,

the unary *branching predicate* B_n, defined as follows:

$$\models_{\mathcal{T}} B_n(a) \iff (\exists x_0, \ldots, x_{n-1} \in \mathcal{T}) \left[\bigwedge_{i<n} P(x_i) = a \ \& \ \bigwedge_{j<k<n} x_j \neq x_k \right].$$

Since these trees are all finite-branching, it would be equivalent to have predicates saying that a has exactly n immediate successors.

Theorem 5.1 *The isomorphism space $\mathcal{FBT}^*/\cong$ of infinite finite-branching trees, in the language with these predicates B_n, is homeomorphic to Baire space ω^ω. Thus it is not compact and not homeomorphic to $\mathcal{ALG}_0^*/\cong$.*

Proof. We write $T|_m$ for the subtree of T containing all nodes at levels $\leq m$. Fix a computable list $\mathcal{L} = \{S_0, S_1, \ldots\}$ of all finite trees, without repetitions. Given the atomic diagram of a tree T in $\mathcal{FBT}$, the predicates R_n allow us to determine, for each m, exactly which tree $S_{f(m)}$ in $\mathcal{L}$ is isomorphic to the finite subtree $T|_m$. (Without the branching predicates, we would not be able to determine this.) So our functional Γ outputs a function $h = \Gamma^{\Delta(T)}$ in ω^ω, defined by setting $h(m)$ to be the size of the set

$$\{j < f(m+1) : S_j \text{ is a tree of height } m+1 \text{ with } S_j|_m \cong T|_m\}.$$

The intuition is that we list out and number all trees in $\mathcal{L}$ of height exactly $m+1$ which extend $T|_m$ to the next level. These are all the possibilities for T_{m+1}. We define $h(m)$ so that the tree isomorphic to $T|_{m+1}$ is the $h(m)$-th tree on the list (allowing $h(m) = 0$).

The inverse Φ is also defined readily: given $h \in \omega^\omega$, Φ^h builds a tree T by choosing the $h(0)$-th possibility among all trees on $\mathcal{L}$ of height 1, then extending it to the $h(1)$-st possibility among all trees of height 2 on $\mathcal{L}$ which extend the tree of height 1 already built, and so on. Since Φ^h knows from this the exact isomorphism type of the first m levels of T, it can compute the branching predicates for this tree and include them in the atomic diagram. Clearly this procedure is inverse to the Γ above. ∎

The main point of Theorem 5.1 is the remark at the end: that $\mathcal{FBT}^*$ and $\mathcal{ALG}_0^*$, modulo isomorphism, are not homeomorphic, and hence do not have the same classification. In light of the substantial similarities between the classes $\mathcal{ALG}_0^c$ and $\mathcal{FBT}^c$ of computable structures (including the analogous

use of the root and branching predicates) this could be a surprise. Considering all countable structures sheds light on the situation in a way that we did not get from the computable structures alone.

The situation becomes a little more complicated when we extend $\mathcal{FBT}$ to allow finite trees: now it is possible that a tree T simply ends at some finite level. The simplest device for allowing finite structures is to allow the atomic diagram to include or exclude the statement $n = n$ for each $n \in \omega$; the domain of the structure, which now consists of those n satisfying $n = n$, is thus decidable from the atomic diagram. Of course, the finiteness of the structure is not decidable. With branching predicates, however, finiteness is Σ_1, since the branching predicates will tell us when we have reached the top level of a finite tree T. Without giving all the details, we state here that this class of structures, under isomorphism, is homeomorphic to the topological space which includes all of ω^ω and also includes, for each $\sigma \in \omega^{<\omega}$, one path h_σ which extends σ but is isolated above σ. (Equivalently, it is homeomorphic to the subclass $\{h \in \omega^\omega : (\forall n)[h(n) = 0 \implies h(n+1) = 0]\}$, under the subspace topology inherited from ω^ω.)

The class $\mathcal{FBT}$ is closely related to the class of finite-valence infinite connected pointed graphs, i.e., connected symmetric irreflexive graphs for which each node has only finitely many neighbors, with one extra constant in the language (which makes the graph *pointed*). Indeed, this class turns out also to be classified effectively via a computable homeomorphism onto Baire space, once the valence function is added to the language. The *valence function* tells how many neighbors each node has; equivalently, one can add unary predicates $V_n(x)$ saying that x has n neighbors. The homeomorphism is similar to that for $\mathcal{FBT}^*$: $f(0)$ describes the isomorphism type of the constant node and its neighbors; $f(1)$ describes the isomorphism type of their neighbors, i.e., of the subgraph of radius 2 around the constant, given the subgraph of radius 1 described by $f(0)$; and so on.

When one abandons the constant, leaving the class of connected finite-valence graphs, the isomorphism problem becomes Σ_3^0 instead of Π_2^0. This is a class of structures for possible further study. In general, as the isomorphism problem grows more difficult, we suspect that it will be necessary to use various of the well-known basic equivalence relations on 2^ω or on ω^ω, such as the relation E_0 of having finite symmetric difference, or the relation E_{perm} of having the same columns (when viewed as subsets of ω^2) up to permutation. This means that a homeomorphism onto the space $2^\omega/E_0$ or $\omega^\omega/E_{\text{perm}}$, for instance, would fall under the head of an effective classification, as would

$2^\omega/E_{\mathrm{card}}$, which we already used to classify $\mathcal{ACF}_0^D$. Certain other useful equivalence relations will be discussed in Section 8.

6   Measure and Category for Field Properties

In classes of structures with a sufficiently nice classification, we can consider various properties of isomorphism types and inquire into the frequency with which those properties occur. For an example, we return now to the class $\mathcal{ALG}_0^*$ of algebraic fields of characteristic 0, in the language with all root predicates R_n. By Theorem 3.3, we have a computable homeomorphism Φ from Cantor space 2^ω onto the space $\mathcal{ALG}_0^*/\cong$ of all isomorphism types of such fields, allowing the transfer of the standard notions of Lebesgue measure and Baire category from 2^ω onto $\mathcal{ALG}_0^*/\cong$. The Lebesgue measure of a class of isomorphism types of fields is simply the Lebesgue measure of the corresponding set of reals in Cantor space, and the class of fields is nowhere dense, meager, etc., just if the corresponding set of reals is. Later, in Section 7, we will discuss whether Lebesgue measure is the best measure to use for this purpose. First, however, we illustrate our goals by offering potential uses of measure and category.

It is clear that the class of all isomorphism types of fields normal over $\mathbb{Q}$ has measure 0, and is meager. Infinitely many irreducible polynomials f in $\mathbb{Q}[X]$ have non-cyclic Galois groups, and for each of these, there is positive probability that f will have at least one root but not all of its roots in a given field. (More precisely, one can give infinitely many polynomials whose Galois groups are all S_3, the symmetric group on the three roots in $\overline{\mathbb{Q}}$, and such that the splitting field of each is linearly disjoint from that of the others. Each of these polynomials has the same positive probability of having exactly one root in a given field, and linear disjointness ensures that these probabilities are all independent: the number of roots of one such polynomial in F is independent of the number of roots of any of the others. We will not go further into the details here; the reader may refer to [12, Prop. 2.4].)

A more involved investigation is required when we consider the property of *relative computable categoricity* from computable structure theory. A countable structure $\mathcal{A}$ is relatively computably categorical if, for every two copies $\mathcal{B}$ and $\mathcal{C}$ of $\mathcal{A}$ with domain ω, there is an isomorphism from $\mathcal{B}$ onto $\mathcal{C}$ which is computable from the atomic diagrams of $\mathcal{B}$ and $\mathcal{C}$. An equivalent condition, discovered in [2, 3] and expanded in [19], is for $\mathcal{A}$ to have a

Scott family of (finitary) Σ_1 formulas which is e-reducible to the (finitary) Σ_1-theory of an expansion of $\mathcal{A}$ by finitely many constants. This in turn is equivalent to a uniform version of categoricity. Useful sources for details include [1] and [13], which we will assume here as background.

Theorem 6.1 *The class*

$$\{[F] \in \mathcal{ALG}_0^*/\cong: F \text{ is relatively computably categorical}\}$$

of isomorphism types of relatively computably categorical fields in $\mathcal{ALG}_0^*$ (in the language with all root predicates R_n) has measure 1 and is comeager within $\mathcal{ALG}_0^*/\cong$. Indeed, there is a single Turing functional Θ such that

$$\{[F] \in \mathcal{ALG}_0^*/\cong: (\forall K \cong F)\ \Theta^{F \oplus K} : F \to K \text{ is an isomorphism}]\}$$

is comeager of measure 1 there.

Proof. We prove the stronger statement, writing $f = \Theta^{F \oplus K}$ for simplicity. (We have now come to identify a field F_D with its atomic diagram D, writing F for both and trusting the reader to understand which is intended.) Let f_0 be the empty map. On input $x \in F$, we compute $f_{x+1}(x)$ using recursion on x, knowing $f_x(0), \ldots, f_x(x-1)$. From the atomic diagram F, we can find the minimal polynomial $p(0, \ldots, x-1, X)$ of each x in F over the subfield F_x generated by $\mathrm{dom}(f_x) = \{0, \ldots, x-1\}$, and then find all roots $x_1 = x, x_2, \ldots, x_k$ of $p(0, \ldots, x-1, X)$ in F. (The root predicates compute k, so we know when we have them all.) Likewise, we can then find all k roots $y_1, \ldots, y_k$ of $p(f_x(0), \ldots, f_x(x-1), Y)$ in K; there must be exactly k of them, since $F \cong K$ (and since, by induction, $f_x = f \upharpoonright x$ extends to some isomorphism). If $k = 1$, we define $f_{x+1}(x) = y_1$ immediately. Otherwise, for each $i \leq k$, we use the root predicates in F and K to search for a polynomial $q \in \mathbb{Q}[X_0, \ldots, X_x, Z]$ such that

$$(\exists a \in F)\ q(0, \ldots, x, a) = 0 \iff (\forall b \in K)\ q(f_x(0), \ldots, f_x(x-1), y_i, b) \neq 0.$$

Whenever we find such a q, we say that it has *ruled out* y_i as an image for x. If we reach a stage where all but one of $y_1, \ldots, y_k$ has been ruled out thus, then we define $f_{x+1}(x)$ to be the remaining y_i. If this never happens, then the program never halts.

It is clear that, if f actually is total, then it will be an isomorphism from F onto K. Of course, there will be cases where f is not total. For

17

example, whenever F is a proper normal extension of $\mathbb{Q}$, this f will not be total. (This includes many of the best-known fields in $\mathcal{ALG}_0$, such as $\overline{\mathbb{Q}}$ and all splitting fields properly extending $\mathbb{Q}$.) Also, if $f = \Theta^{F \oplus K}$ is non-total for some particular F and K, then $\Theta^{\widetilde{F} \oplus \widetilde{K}}$ will fail to be total whenever $F \cong \widetilde{F}$ and $K \cong \widetilde{K}$, since the image of x in $\widetilde{F}$ will always be a stumbling block.

However, in order for $f(x)$ to diverge, two of the roots in K – say y_1 and y_2, without loss of generality – must both fail ever to be ruled out by the procedure above. Now, for every finite initial segment $\sigma \in 2^{<\omega}$ of F and every conjugate x_2 of $x_1 = x$ in F over $\mathbb{Q}(0, \ldots, x-1)$, there is always a way to extend σ so that some q does distinguish between the two of them. It follows that the class of fields F for which (for some K, equivalently for all $K \cong F$) two such y_1 and y_2 exist for this x is a nowhere dense class. Therefore, the class of isomorphism types in $\mathcal{ALG}_0$ for which Θ fails to work is a meager class. Moreover, for any given $x = x_1$ and for each of its conjugates $x_i \neq x$ in F, the probability that no q ever distinguishes between them is 0, since Lemma 6.2 will give an infinite collection of polynomials, each of which has a one-half chance to distinguish between them, and for which these one-half chances are all independent. This completes the proof of the theorem. ∎

In the foregoing proof, Lemma 6.2 establishes the intuitively clear fact that, with probability 1, two conjugates y_1 and y_2 in an algebraic field F can always be distinguished by polynomials of the form $f(y_1, Z)$ and $f(y_2, Z)$: for some f, one of these will have a root in the field and the other will not. We will next prove the more surprising Theorem 6.4, stating that, even without the root relations in the language, measure-1-many of the fields in $\mathcal{ALG}_0$ remain relatively computably categorical. This will require Lemma 6.2 in exact detail.

Lemma 6.2 *Let* $\alpha_0, \ldots, \alpha_n \in \overline{\mathbb{Q}}$ *be algebraic numbers conjugate over* $\mathbb{Q}$. *Then, for every finite algebraic field extension* $E \supseteq \mathbb{Q}$ *with* $\alpha, \beta \in E$, *and for all distinct* $i \leq n$ *and* $j \leq n$, *there exists an infinite set* $D = \{q_0, q_1, \ldots\} \subseteq \mathbb{Q}$ *of rational numbers such that for every* k, *both of the following hold:*

$$\sqrt{\alpha_i + q_k} \notin E(\sqrt{\alpha_i + q_l}, \sqrt{\alpha_j + q_l} \; : \; l \neq k)(\sqrt{\alpha_j + q_k});$$

$$\sqrt{\alpha_j + q_k} \notin E(\sqrt{\alpha_i + q_l}, \sqrt{\alpha_j + q_l} \; : \; l \neq k)(\sqrt{\alpha_i + q_k}).$$

Moreover, there is a procedure, uniform in $\alpha_0, \ldots, \alpha_n$, i, j, *and the generators of* E, *for deciding such a set* D.

That is, adjoining $\sqrt{\alpha_i + q_k}$ to the field generated by all the other square roots (from other elements $q_l \in D$, with $l \neq k$) will not cause $\sqrt{\alpha_j + q_k}$ to appear in that same field, nor vice versa. The same therefore holds for adjoining $\sqrt{\alpha_i + q_k}$ to any subfield of that field.

Proof. We start by setting E_0 to be the normal closure of E, the smallest extension of E (clearly also of finite degree) which is normal over $\mathbb{Q}$. We first use E_0 to find q_0, then define E_1 to be the normal closure of $E_0(\sqrt{\alpha_i + q_0}, \sqrt{\alpha_j + q_0})$ and repeat the process with E_1 to get q_1, and so on, using the following step recursively.

We now appeal to the following well-known fact, which appears, for instance, as [9, Lemma 11.6].

Lemma 6.3 *Let L be a finite separable extension of a field K, and let $f \in L(Z)[T]$ be irreducible as a polynomial in T over $L(Z)$. Then there exists an irreducible $p \in K(Z)[T]$ such that for every $q \in K$, if $p(q, T)$ is irreducible in $K[T]$, then $f(q, T)$ is irreducible in $L[T]$.*

We apply this lemma with E_k as L and with $\mathbb{Q}$ as K. Recall that $\mathbb{Q}$ is a *Hilbertian field*: for every $p \in \mathbb{Q}(Z)[T]$ irreducible over $\mathbb{Q}(Z)$, there exist infinitely many $q \in \mathbb{Q}$ for which $p(q, T)$ is irreducible in $\mathbb{Q}[T]$. Therefore, the lemma implies that for every irreducible $f \in E_k(Z)[T]$, some $q \in \mathbb{Q}$ makes $f(q, T)$ irreducible in $E_k[T]$. Now the polynomial $T^2 - \frac{\alpha_i + Z}{\alpha_j + Z}$ is irreducible in the polynomial ring $E_k(Z)[T]$ provided that $\frac{\alpha_i + Z}{\alpha_j + Z}$ is not a square in $E_k(Z)$, which holds since $\alpha_i \neq \alpha_j$. Therefore, there exists some $q \in \mathbb{Q}$ for which $T^2 - \frac{\alpha_i + q}{\alpha_j + q}$ is irreducible in $E_k[T]$. In particular, $(\alpha_i + q)$ and $(\alpha_j + q)$ are not both squares in E_k. However, if $(\alpha_i + q)$ were a square, say $\alpha_i + q = z^2$, then we could apply an automorphism h of $\overline{\mathbb{Q}}$ mapping α_i to its $\mathbb{Q}$-conjugate α_j. Since E_k is normal over $\mathbb{Q}$, it would also contain $h(z)$, so $\alpha_j + q = h(\alpha_i + q) = (h(z))^2$ would also have been a square in E_k, which is impossible. Therefore $(\alpha_i + q)$ is not a square in E_k, nor is $(\alpha_j + q)$, by the same reasoning.

Now suppose $\sqrt{\alpha_j + q} \in E_k(\sqrt{\alpha_i + q})$. Then we have $x, y \in E_k$ with $(x + y\sqrt{\alpha_i + q})^2 = \alpha_j + q$. However, then

$$x^2 + 2xy\sqrt{\alpha_i + q} + y^2(\alpha_i + q) = \alpha_j + q \in E_k,$$

so the coefficient $2xy$ must be zero. But $y = 0$ would force $\alpha_j + q$ to be a square in E_k, while $x = 0$ would force $\frac{\alpha_i + q}{\alpha_j + q} = \frac{1}{y^2}$ to be a square in E_k.

Since both of these are impossible, we see that $\sqrt{\alpha_j + q} \notin E_k(\sqrt{\alpha_i + q})$ and likewise $\sqrt{\alpha_i + q} \notin E_k(\sqrt{\alpha_j + q})$.

This argument shows that infinitely many numbers $q \in \mathbb{Q}$ exist for which $\sqrt{\alpha_j + q} \notin E_k(\sqrt{\alpha_i + q})$ and $\sqrt{\alpha_i + q} \notin E_k(\sqrt{\alpha_j + q})$. Knowing the generators of E_k, we have a splitting algorithm for E_k, by Kronecker's Theorem, and so we may identify such a q when we find one. Choose q_k to be some such q which (as an element of the domain ω of $\overline{\mathbb{Q}}$) is greater than k; this will make our set D decidable. Setting E_{k+1} to be the normal closure of $E_k(\sqrt{\alpha_i + q_k}, \sqrt{\alpha_j + q_k})$, we proceed by recursion on k, and thus build the set D.

Now suppose for a contradiction that for some k,

$$\sqrt{\alpha_j + q_k} \in E(\sqrt{\alpha_i + q_l}, \sqrt{\alpha_j + q_l} \; : \; l \neq k)(\sqrt{\alpha_i + q_k}).$$

Take any k for which this holds, and fix the least p such that

$$\sqrt{\alpha_j + q_k} \in E(\sqrt{\alpha_i + q_l}, \sqrt{\alpha_j + q_l} \; : \; l \leq p \; \& \; l \neq k)(\sqrt{\alpha_i + q_k}).$$

By our construction, $\sqrt{\alpha_j + q_k} \notin E_k(\sqrt{\alpha_i + q_k})$, so clearly $p > k$, and (by the minimality of p) either the adjoinment of $\sqrt{\alpha_i + q_p}$ to the field

$$\widetilde{E} = E(\sqrt{\alpha_i + q_l}, \sqrt{\alpha_j + q_l} \; : \; l < p \; \& \; l \neq k)(\sqrt{\alpha_i + q_k})$$

or the subsequent adjoinment of $\sqrt{\alpha_j + q_p}$ to $\widetilde{E}(\sqrt{\alpha_i + q_p})$ caused $\sqrt{\alpha_j + q_k}$ to enter the field. But each of these two extensions was of degree 2, and so that extension must also be generated by $\sqrt{\alpha_j + q_k}$. Therefore either

$$\sqrt{\alpha_i + q_p} \in \widetilde{E}(\sqrt{\alpha_j + q_k}) \subseteq E_p$$

or

$$\sqrt{\alpha_j + q_p} \in \widetilde{E}(\sqrt{\alpha_i + q_p}, \; \sqrt{\alpha_j + q_k}) \subseteq E_p(\sqrt{\alpha_i + q_p}),$$

both of which contradict our construction at stage p. This proves one of the two conditions required by the lemma, and by symmetry on i and j, the other also holds. ∎

Theorem 6.4 *The class*

$$\{[F^*] \in \mathcal{ALG}_0^*/\cong: F \text{ is relatively computably categorical in } \mathcal{ALG}_0\}$$

of isomorphism types of relatively computably categorical fields in $\mathcal{ALG}_0$ (in the language without the root relations R_n) has measure 1 and is comeager within $\mathcal{ALG}_0^*/\cong$.

20

To be clear: the measure here is still taken in the isomorphism space $\mathcal{ALG}_0^*/\cong$ with the root relations, as this is the only isomorphism space in which this measure makes sense. The theorem states that measure-1-many of these fields remain relatively computably categorical even when the isomorphisms are forbidden to use information about the root relations from the atomic diagrams. However, we lose the uniformity of Theorem 6.1.

Proof. Fix any $\epsilon > 0$. We will enumerate, effectively, a family Σ of formulas which, for all $[F]$ in a subset of $\mathcal{ALG}_0^*/\cong$ of measure $> 1 - \epsilon$, forms a Scott family for F. This will show that all these $[F]$ are relatively computably categorical, since having a c.e. Scott family implies relative computable categoricity. Perhaps the most striking thing about this family is that it will be built with no reference to any specific $[F]$ at all, apart from its use of a fixed computable presentation of $\overline{\mathbb{Q}}$, which we denote simply as $\overline{\mathbb{Q}}$, and a fixed computable enumeration $p_0, p_1, \ldots$ of the irreducible monic polynomials in $\mathbb{Q}[X]$, with the coefficients viewed as elements of $\overline{\mathbb{Q}}$.

Consider the n-th polynomial $p_n(X)$. Let d be its degree. If $d = 1$, so that $p_i = X - \frac{m}{k}$ for some $m \in \mathbb{Z}$ and nonzero $k \in \omega$, then we add to Σ a formula saying that $kX = m$. Otherwise, we now apply Lemma 6.2 with a vengeance, with the splitting field of p_n over $\mathbb{Q}$ as our E, and with the roots $\alpha_1, \ldots, \alpha_d$ of p_n in $\overline{\mathbb{Q}}$. Let $\delta = \frac{\epsilon}{2^{n+1}}$, and compute an N so large that both of the following hold:

- The probability that, out of $100N$ independent coin flips, at least $40N$ will be heads, is $> (1 - \delta)^{\frac{1}{2d}}$; and

- The probability that, out of $100N$ independent flips of two coins each, at most $35N$ will yield two heads, is $> (1 - \delta)^{\frac{1}{d(d-1)}}$.

Fix elements $q_0, \ldots, q_{100N-1}$ from the set D given by Lemma 6.2. For each subset $S \subseteq \{q_0, \ldots, q_{100N-1}\}$ with $|S| = 40N$, we add to Σ the fomula

$$p_n(X) = 0 \ \& \ \bigwedge_{q \in S} \exists Y_q \ Y^2 = X + q,$$

saying that X is one of the α_i and that, for all $q \in S$, the field contains an element $\sqrt{\alpha_i + q}$.

The first item ensures that, with probability $> \sqrt{1 - \delta}$, every α_i in F will satisfy $(\exists x) \ x^2 = \alpha_i + q_k$ for at least $(40N)$-many values $j < 100N$. Therefore, with this probability, every root of p_n in F (if there are any) will

indeed realize one of the formulas now in Σ. On the other hand, the second item ensures that the probability is $> ((1 - \delta)^{\frac{1}{d(d-1)}})^{\binom{d}{2}} = \sqrt{1 - \delta}$ that, for every two distinct roots α_i, α_j of p_n in F, no more than $(35N)$ values of $k < 100N$ have the property that

$$(\exists y \in F)\ y^2 = \alpha_i + q_k \quad \& \quad (\exists z \in F)\ z^2 = \alpha_j + q_k.$$

Therefore, with probability $> 1 - \delta = 1 - \frac{\epsilon}{2^{n+1}}$, Σ behaves exactly as a Scott family should, as far as the roots of p_n are concerned. (No other formula in Σ can be realized by any root of p_n, since all such formulas either specify a rational value for X, or else require X to be a root of some p_m with $m \neq n$.)

Finally, to make a true Scott family, we adjoin to Σ all conjunctions of the form

$$\bigwedge_{i \leq m} \alpha_{n_i}(X_{n_i}),$$

for all strictly increasing tuples $(n_0, \ldots, n_m) \in \omega^{<\omega}$, where α_{n_i} can be any of the formulas added to Σ for p_{n_i} above. Scott families, after all, must contain formulas realizing every n-tuple from a structure, even though, in this case, the single-variable formulas actually suffice, being specific enough that, when required to identify a new element with respect to the finitely many already considered, they can actually identify it uniquely without even considering the preceding elements.

The probability that our Σ fails to work for the roots of a single p_n is $< \frac{\epsilon}{2^{n+1}}$, by the analysis above. Therefore, for at least $(1 - \epsilon)$-many of the isomorphism types in $\mathcal{ALG}_0^* / \cong$, Σ is a computably enumerable Scott family, and thus all these types are relatively computably categorical. Moreover, Σ does not use the root predicates anywhere, and therefore these isomorphism types remain relatively computably categorical even in the original language of fields, without the root predicates. Since ϵ was arbitrary, the theorem follows. ∎

It is well-known that normal algebraic field extensions of $\mathbb{Q}$ are also relatively computably categorical, even without the root predicates, although the specific procedure is different from the functionals in Theorems 6.1 and 6.4. With normal extensions, one simply finds the minimal polynomial p of x over $\mathbb{Q}(0, 1, \ldots, x - 1)$ and defines $f_{x+1}(x)$ to be the first root of $p(f_x(0), \ldots, f_x(x - 1), Y)$ that appears in the target field K, knowing that by normality, this map must extend to an isomorphism. Of course, normal

extensions are a meager class of measure 0 within the class of all algebraic extensions of $\mathbb{Q}$; our point is simply that separate procedures may work in special cases. In the forthcoming work [8], it is shown that the procedures given in Theorems 6.1 and 6.4 succeed for all *random fields*, i.e., for all F such that the $h \in 2^\omega$ corresponding to $[F]$ is Martin-Löf random, or even just Schnorr random.

7 Different Measures

In the preceding section, we transferred the Lebesgue measure on Cantor space to the space $\mathcal{ALG}_0^*/\cong$, using our homeomorphism from Theorem 3.3. This seems like a natural choice, yet in this section we present another measure on $\mathcal{ALG}_0^*/\cong$, which we consider to be, if anything, a more natural choice. The results of Section 6 hold no matter which of the two measures one uses.

We were careful, in the proof of Theorem 3.3, to consider only polynomials f_n of prime degree. This avoided one possible trap regarding the measure: it does not depend on the particular enumeration $f_0, f_1, \ldots$ of the monic polynomials in $E[X]$ there. Had we not required prime degree, the following situation could have arisen.

(a) If $f_0(X) = X^2 - 2$, then the measure of the class of all isomorphism types of fields containing a square root of 2 is $\frac{1}{2}$.

(b) However, if instead $f_0(X) = X^4 - 2$ and $f_1(X) = X^2 - 2$, then Lebesgue measure would dictate that half of all fields contain a fourth root of 2, and that half of the remaining fields contain a square root of 2 (but no fourth root). In this case, the measure of the class of all isomorphism types of fields containing a square root of 2 is $\frac{3}{4}$.

By requiring prime degree, we avoided this trap to some extent. Nevertheless, certain peculiarities of the measure remain. For example, the fields containing an 181-st root of 2 form a class of measure $\frac{1}{2}$, since 181 is prime, whereas the fields containing a 180-th root of 2 form a class of far smaller measure, since such a root is built up by a long series of prime-degree extensions. This is not inconsistent in any way, but it seems a little strange, and there is a way of avoiding it.

Definition 7.1 The *Haar-compatible measure H* of a basic open set $\mathcal{U}_\sigma$ in Cantor space is defined as follows, using the same notation as in the proof

of Theorem 3.3: within a computable copy E of $\overline{\mathbb{Q}}$, we choose the same polynomials f_σ and build the same subfields F_σ of E. For the empty string λ, of course we define $H(\mathcal{U}_\lambda) = 1$. Given $H(\mathcal{U}_\sigma)$, we set

$$H(\mathcal{U}_{\sigma\frown 1}) = \frac{1}{d} \cdot H(\mathcal{U}_\sigma) \quad \& \quad H(\mathcal{U}_{\sigma\frown 0}) = \frac{d-1}{d} \cdot H(\mathcal{U}_\sigma),$$

where d is the degree of the polynomial f_σ. The Haar-compatible measure of an arbitrary subset $\mathcal{V} \subseteq 2^\omega$ is then the infimum of the sums $\Sigma_i H(\mathcal{U}_{\sigma_i})$, over all countable sequences $\langle \sigma_i \rangle_{i \in \omega}$ for which $\mathcal{V} \subseteq \bigcup_i \mathcal{U}_{\sigma_i}$.

So the set $\mathcal{U}$ of fields containing an 11-th root of 2 has $H(\mathcal{U}) = \frac{1}{11}$, as opposed to its Lebesgue measure $\frac{1}{2}$. Of course, $H(\mathcal{U}_{\sigma\frown 0}) + H(\mathcal{U}_{\sigma\frown 1}) = H(\mathcal{U}_\sigma)$, so this really is a measure. The next lemma justifies the name "Haar-compatible measure."

Lemma 7.2 *For every finite Galois extension K of $\mathbb{Q}$, the set $\mathcal{V}$ of all fields in $\mathcal{ALG}_0^*/\cong$ containing K has $H(\mathcal{V})$ equal to the usual Haar measure $\frac{1}{[K:\mathbb{Q}]}$ of the (pointwise) stabilizer of K within the Galois group $Gal(\overline{\mathbb{Q}}/\mathbb{Q})$.*

Proof. Let h be any element of Cantor space, with corresponding field $F_h = \cup_{\sigma \subset h} F_\sigma \in \mathcal{ALG}_0^*$. Assuming $K \neq \mathbb{Q}$, there will be at least one $\sigma \subset h$, for which $F_\sigma \cap K$ is a proper subfield of $F_{\sigma\frown 1} \cap K$. These σ are the nodes at which h makes decisions about whether to include K (if $h(|\sigma|) = 1$) or not (if $h(|\sigma|) = 0$). At the first such σ_1, with $p_1 = \deg(f_{\sigma_1})$, we will have $H(\mathcal{U}_{\sigma_1\frown 0}) = \frac{p_1 - 1}{p_1} \cdot H(\mathcal{U}_{\sigma_1})$, and, since K is normal over $\mathbb{Q}$, none of these fields will lie in $\mathcal{V}$. (Notice that if K were not normal, it might still be possible to have an $h \supset \sigma_1$ such that F_h contained a subfield isomorphic to K.) For the remaining $\frac{1}{p_1} \cdot H(\mathcal{U}_{\sigma_1}) = H(\mathcal{U}_{\sigma_1\frown 1})$-many fields, either $F_{\sigma_1\frown 1} \supseteq K$ and these fields all lie in $\mathcal{V}$; or else we continue up through Cantor space until we find another σ_2 with the same property. After we reach σ_n (where n is the number of prime factors of $[K : \mathbb{Q}]$, counted by multiplicity), we must have $K \subseteq F_{\sigma_n\frown 1}$, and multiplying all the numbers to this point proves the lemma, since the degrees of all the extensions $[F_{\sigma_i\frown 1} : F_{\sigma_i}]$ must have product $[K : \mathbb{Q}]$. ∎

Lemma 7.2 emphatically does require that K be normal over $\mathbb{Q}$. Our Haar-compatible measure for finite non-normal field extensions is more complicated, as is the Lebesgue measure. For example, consider the class $\mathcal{V}$ of

fields K containing a fourth root of 2 (equivalently, a copy of $\mathbb{Q}(\sqrt[4]{2})$). If our sequence of polynomials in the proof of Theorem 3.3 begins with:

$$f_0 = X^2 - 2; \quad f_1 = X^2 - \sqrt{2}; \quad f_2 = X^2 + \sqrt{2},$$

then $H(\mathcal{V}) = \frac{3}{8}$, as $\mathcal{V}$ corresponds to the union of $\mathcal{U}_{11}$ and $\mathcal{U}_{101}$. However, if instead we used a sequence beginning with:

$$g_0 = X^2 + 1; \quad g_1 = f_0; \quad g_2 = f_1; \quad g_3 = f_2,$$

then $H(\mathcal{V}) = \frac{5}{16}$ instead, as now $\mathcal{V}$ corresponds to the union $\mathcal{U}_{111} \cup \mathcal{U}_{011} \cup \mathcal{U}_{0101}$. (The point is that the string $\sigma = 110$ rules out any square root of $-\sqrt{2}$, since the presence of i in the field means that we must have either no fourth roots of 2, or else all four of them.) The Haar-compatible measure is not to blame, insofar as $H(\mathcal{V})$ is also the Lebesgue measure of $\mathcal{V}$ in both of these cases. For non-normal field extensions, the measure still depends on the sequence of polynomials used. We conjecture that various questions arising in this and the following section are more readily addressed if one considers only the class of normal algebraic field extensions of $\mathbb{Q}$, in place of the class $\mathcal{ALG}_0$.

8   Topologies on 2^ω

As promised above, we return here to the topology of $\mathcal{ALG}_0/\cong$, which arose in Section 3. There we noted that it is not homeomorphic to Cantor space. (By *Cantor space* we really mean not just the set 2^ω, but rather this set with its usual topology, where the basic open sets are defined by $\mathcal{U}_\sigma = \{h \in 2^\omega : \sigma \subset h\}$ for all $\sigma \in 2^{<\omega}$.) Here we consider other possibilities.

The *Scott topology* on 2^ω is a sort of positive version of Cantor space. Here the basic open sets are defined to be the sets

$$\mathcal{W}_F = \{h \in 2^\omega : F \subseteq h^{-1}(1)\}$$

for all finite subsets F of ω. If $F = \{3, 10\}$, then $\mathcal{W}_F$ contains all $h \in 2^\omega$ with $h(3) = h(10) = 1$, or equivalently, all subsets of ω which contain both 3 and 10. However, sets such as $\{h \in 2^\omega : h(3) = 0\}$, which were open in Cantor space, are not open in the Scott topology. We will write $\mathfrak{S}$ for the space 2^ω under the Scott topology.

The Scott topology seems much closer in nature to $\mathcal{ALG}_0/\cong$, being defined by positive information only. Both contain an element which lies in

25

every nonempty open set (for $\mathfrak{S}$, it is the constant function $h = 1$) and an element which sits in no proper open subset of the space (the other constant function $h = 0$). In fact, there is a very natural equivalence relation on 2^ω whose quotient has the Scott topology. Earlier we used $=^e$ to denote a binary relation on ω. Now we revise it to denote the related notion on 2^ω:

$$A =^e B \iff (\forall x)[(\exists y \; \langle x, y \rangle \in A) \iff (\exists z \langle x, z \rangle \in B)].$$

This relation is best understood by considering the image $\pi_1(A) = \{x \in \omega : \exists y \; \langle x, y \rangle \in A\}$ of A under projection π_1 onto the first coordinate. The set $\pi_1(A)$ is said to be *enumerated by* A, and A and B satisfy $=^e$ if and only if they enumerate the same set. Notice that a set is C-computably enumerable just if there is a C-computable set enumerating it: the terminology captures our intuition about enumeration of sets. Likewise, it gives an equivalent definition of enumeration reducibility $\leq_e$: an *enumeration reduction from D to C* is simply a Turing functional Γ such that, for every enumeration B of D, Γ^B is the characteristic function of an enumeration of C.

The homeomorphism from $2^\omega /=^e$ onto $\mathfrak{S}$ is simple: one maps each $A \in 2^\omega$ to $\pi_1(A)$ in $\mathfrak{S}$. The preimage of each $\mathcal{W}_F$ is open, so this is continuous, and it respects $=^e$, so it gives rise to a continuous function from $2^\omega /=^e$ onto $\mathfrak{S}$, which is quickly seen to be bijective. The inverse is even easier: map $C \in \mathfrak{S}$ to the $=^e$-class of the set $\{\langle x, 0 \rangle : x \in C\}$. Again, this is easily seen to be continuous, so we have a homeomorphism. Indeed, the natural bases for the two spaces both consist of sets defined by the requirement of including a certain finite subset F of ω.

We now introduce another natural and closely related equivalence relation $=^f$ on 2^ω, which aids in effective classification of several classes of countable structures not considered in this article, such as equivalence structures and torsion-free abelian groups of rank 1. Recall that the x-th *column* of $A \subseteq \omega$ is defined to be $A^{[x]} = \{y \in \omega : \langle x, y \rangle \in A\}$. The relation $=^f$ is defined by:

$$A =^f B \iff (\forall x)[A^{[x]} \; E_{\mathrm{card}} \; B^{[x]}]$$
$$\iff (\forall x) \; |\{y : \langle x, y \rangle \in A\}| = |\{y : \langle x, y \rangle \in B\}|.$$

So we think of each column $A^{[x]}$ of A as "counting" something – perhaps the number of times a fixed element of a torsion-free abelian group is divisible by the prime p_x, for example. Instead of enumerating a set, A thus approximates a function g_A, by computable approximations from below. Now $A =^f B$ just if $g_A = g_B$ as functions. Notice, however, that this is not quite the same

notion as the limitwise monotonic approximation of functions used in many other topics in computable structure theory, since the function g_A maps ω into $(\omega + 1)$, taking on the value ω whenever the set $A^{[x]}$ is infinite.

Basic open sets in $2^\omega/=^f$ may be thought of as given by finite pieces of functions. For instance, using the function g_0 with $g_0(3) = 4$, $g_0(10) = 2$, and $g_0(x) = 0$ for all other x, we get a basic open set containing all A for which $g_A(3) \geq 4$ and $g_A(10) \geq 2$. Notice that g_0 must have finite support (i.e., $g_0^{-1}(0)$ must be cofinite) and must take on only finite values: setting $g_0(x) = \omega$ is not allowed. (A finite initial segment of A can ensure $g_A(3) \geq 4$, but no finite initial segment can ensure $g_A(x) = \omega$.) The space $2^\omega/=^f$ is similar in many respects to the Scott topology, but the two are not the same.

Theorem 8.1 *The spaces $2^\omega/=^e$ (i.e., the Scott topology) and $2^\omega/=^f$ are not homeomorphic.*

Proof. We use the notion of a *principal open set*. In a topological space, an open set $\mathcal{V}$ is *principal* if there exists some point x in the space for which $\mathcal{V}$ is the smallest open set containing x. (Such an x may be said to *generate* this $\mathcal{V}$.) This property is clearly preserved under homeomorphisms.

Now in $2^\omega/=^e$ and $2^\omega/=^f$, principal open sets do exist, and in fact they are just the basic open sets described above. Intuitively, each one is generated by a given finite subset of ω. Formally, for each finite $F \subset \omega$, the open set

$$\mathcal{V}_F = \{[A] \in 2^\omega/=^e \colon F \subseteq \pi_1(A)\}$$

is principal, being generated by $[F]$, and all principal open sets are of this form, including the entire space, which is generated by the singleton $=^e$-class $[\emptyset]$. (In the notation for the Scott topology $\mathfrak{S}$, this $\mathcal{V}_F$ was the set $\mathcal{W}_F$.)

Likewise, in $2^\omega/=^f$, principal open sets are of the form

$$\mathcal{X}_g = \{[A] : (\forall x)g_A(x) \geq g(x)\},$$

where g ranges over all total functions from ω into ω with finite support. $\mathcal{X}_g$ is generated by the $=^f$-class $[\{\langle x, y \rangle : y < g(x)\}]$ of a set whose x-th column always has cardinality $g(x)$.

Now we can explain why the two spaces are not homeomorphic. Let $\mathcal{V} = \mathcal{V}_F$ be any principal open set in $2^\omega/=^e$. Then the principal open supersets of $\mathcal{V}$ in this space are precisely those of the form $\mathcal{V}_G$ with $G \subseteq F$, and so there are exactly $2^{|F|}$ principal open sets containing $\mathcal{V}_F$, including

27

both $\mathcal{V}_F$ itself and the entire space. However, in $2^\omega/=^f$, the principal open supersets of each $\mathcal{X}_g$ are those $\mathcal{X}_h$ satisfying $(\forall x\ h(x) \le g(x))$. In particular, the principal open set $\mathcal{X}_{g_2}$ for the function $g_2(0) = 2$ (with support $\{2\}$) has exactly three principal open supersets: itself, the entire space, and $\mathcal{X}_{g_1}$, where $g_1(0) = 1$. Thus no homeomorphism can map $\mathcal{X}_g$ to any principal open set in $2^\omega/=^e$. ∎

Theorem 8.1 is of interest on its own, but another reason for proving it was to prepare the way for proving the answer to our initial question from this section.

Theorem 8.2 *The Scott topology and $\mathcal{ALG}_0/\cong$ are not homeomorphic; nor are $2^\omega/=^f$ and $\mathcal{ALG}_0/\cong$.*
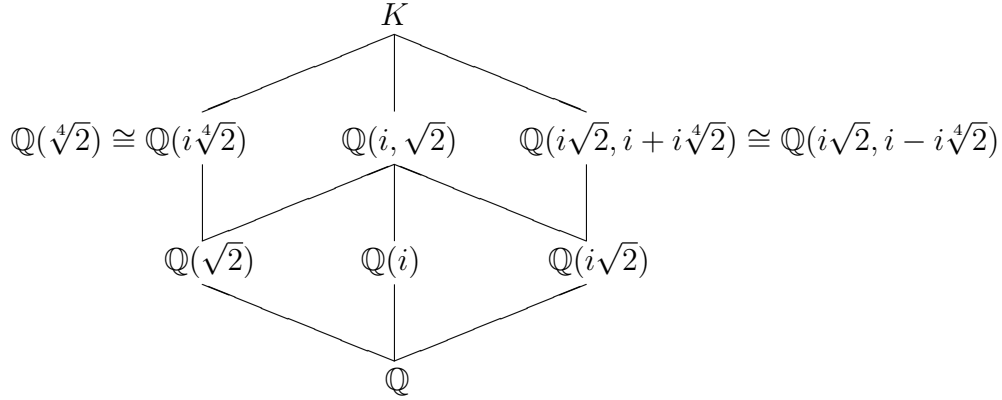
Proof. We use the characterization from Theorem 8.1 of the principal open sets in $2^\omega/=^e$ and $2^\omega/=^f$. In $\mathcal{ALG}_0/\cong$, by much the same reasoning, the principal open sets are the basic open sets from Lemma 3.1. As an example, consider the subfield $\mathbb{Q}(\theta)$ of $\overline{\mathbb{Q}}$ generated by a single primitive fifth root θ of unity. The minimal polynomial of θ is the cyclotomic polynomial $X^4 + X^3 + X^2 + X + 1$, and its conjugates are θ^2, θ^3, and θ^4, all of which lie in $\mathbb{Q}(\theta)$. Thus the Galois group of this normal extension is $\mathbb{Z}/(4)$. Since this Galois group has a proper nontrivial subgroup, there is a subfield of degree 2 within $\mathbb{Q}(\theta)$, generated by an element $\sqrt{z}$ for some $z \in \mathbb{Q}$. Now the subfield $\mathbb{Q}(\sqrt{z})$ generates a principal open set (more formally, the $\cong$-class $[\mathbb{Q}(\sqrt{z})]$ generates it), which is a proper superset of the principal open set generated by $[\mathbb{Q}(\theta)]$. $[\mathbb{Q}]$ itself generates another principal open superset, namely the entire space, and since $\mathbb{Q}(\theta)$ has no other subfields, the principal open set generated by $[\mathbb{Q}(\theta)]$ has exactly three principal open supersets. As before, this shows that $2^\omega/=^e$ cannot be homeomorphic to $\mathcal{ALG}_0/\cong$.

For $=^f$, a similar numerical argument will not work. Instead, we consider the principal open supersets of a given principal open set as a partial order, under $\subset$, and show that certain finite partial orders possible in $\mathcal{ALG}_0/\cong$ are impossible in $2^\omega/=^f$. In the latter space, when $\mathcal{X}_g$ is a principal open set, its principal open supersets are those of the form $\mathcal{X}_h$ with $h \le g$ on all inputs. In particular, let $\omega - g^{-1}(0) = \{x_1 < \cdots < x_n\}$ and, for each $i < n$, let

$$g_i(x) = \begin{cases} g(x), & \text{if } x = x_i; \\ 0, & \text{if not.} \end{cases}$$

28

Each g_i gives rise to a proper chain $\mathcal{X}_{i,0} \supset \mathcal{X}_{i,1} \supset \cdots \supset \mathcal{X}_{i,g(x_i)}$ of principal open supersets of $\mathcal{X}_g$, with $\mathcal{X}_{i,0}$ always being the entire space. The principal open supersets of $\mathcal{X}_g$ are just the $\Pi_i(1 + g(x_i))$-many intersections of these finitely many sets, and those which are maximal among the principal open proper subsets of $2^\omega/\cong$ are precisely the sets $\mathcal{X}_{i,1}$, for $i = 1, \ldots, n$.

Now we wish to produce a field K, algebraic of finite degree over K, for which the partial order of the subfields of K is not of the form above. Our specific example is the splitting field K of the polynomial $X^4 - 2$, for which we present here the lattice of subfields (up to isomorphism) under inclusion:

$$
\begin{array}{ccccc}
 & & K & & \\
 & & | & & \\
\mathbb{Q}(\sqrt[4]{2}) \cong \mathbb{Q}(i\sqrt[4]{2}) & & \mathbb{Q}(i, \sqrt{2}) & & \mathbb{Q}(i\sqrt{2}, i + i\sqrt[4]{2}) \cong \mathbb{Q}(i\sqrt{2}, i - i\sqrt[4]{2}) \\
 & & & & \\
 & \mathbb{Q}(\sqrt{2}) & \mathbb{Q}(i) & \mathbb{Q}(i\sqrt{2}) & \\
 & & | & & \\
 & & \mathbb{Q} & &
\end{array}
$$

One sees from this lattice that the principal open set $\mathcal{U}_{[K]}$ generated by $[K]$ in $\mathcal{ALG}_0/\cong$ has eight principal open supersets, including itself and the entire space. Those generated by $[\mathbb{Q}(\sqrt{2})]$, $[\mathbb{Q}(i)]$, and $[\mathbb{Q}(i\sqrt{2})]$ are maximal proper. In order for an $\mathcal{X}_g$ as above to give rise to the same lattice, g would need to have exactly three elements in its support (that is, $n = 3$), and to have $(1 + g(x_1)) \cdot (1 + g(x_2)) \cdot (1 + g(x_3)) = 8$. Moreover, in order for the lattice not to have a chain of length > 3, we would need all $g(x_i) \leq 2$, making $g(x_i) = 1$ for each i. However, in the lattice for such an $\mathcal{X}_g$, each maximal proper subset contains itself, $\mathcal{X}_g$, and exactly two of the other principal open sets: for example, if $h(x_1) = 1$ and $h = 0$ everywhere else, then $\mathcal{X}_h$ contains $\mathcal{X}_f$ for exactly those f with $h \leq f \leq g$, and such an f must have $f(x_1) = 1$, $f(x_2) \in \{0, 1\}$, $f(x_3) \in \{0, 1\}$, and $f = 0$ elsewhere. Thus, such an $\mathcal{X}_g$ cannot be mapped onto $\mathcal{U}_{[K]}$ by any homeomorphism, as no $\mathcal{X}_h$ could be mapped onto $\mathcal{U}_{[\mathbb{Q}(i)]}$. This proves the theorem. ∎

The broader conclusion of the foregoing proof is that the question of the actual structure of the principal open sets of $\mathcal{ALG}_0/\cong$ is the same question as the structure of the finite subfields of $\overline{\mathbb{Q}}$ under inclusion, and thus related to

the study of $\mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$, a notoriously difficult topic. We consider this to be a fair validation of the claim in Section 3 that $\mathcal{ALG}_0/\cong$ is not a recognizable topological space, and therefore should not be regarded as a useful effective classification of $\mathcal{ALG}_0$ up to isomorphism. Adding the root predicates to the language really is necessary. On the other hand, we would be quite willing to accept spaces such as $2^\omega/=^e$, or modulo $=^f$, as giving effective classifications of other classes $\mathcal{C}$ of structures, provided that a homeomorphism (preferably computable) exists from that space onto $\mathcal{C}/\cong$. Indeed, descriptive set theory provides quite a number of standard Borel equivalence relations on Cantor space and on Baire space, noncomputable but of relatively low complexity; these include the relations E_0, E_1, E_2, E_3, E_{set}, and Z_0 (see [17] or many other places for their definitions), and also the relation E_{perm}, computably bireducible with E_{set}, which holds of A and B just if the columns of A are exactly the columns of B, up to a permutation of columns. (E_{set} merely requires that every column of each should also appear as a column in the other, disregarding multiplicities.) Beyond that, it is possible to combine these notions: for example, one might mix $=^e$ with E_0 by declaring A and B to be equivalent whenever $\pi_1(A)$ E_0 $\pi_1(B)$ holds. Another variant, the equivalence relation $E_{\mathrm{card}}^\forall$ defined by

$$A\ E_{\mathrm{card}}^\forall\ B \iff (\omega - \pi_1(A))\ E_{\mathrm{card}}\ (\omega - \pi_1(B)),$$

is quickly seen to have $2^\omega/E_{\mathrm{card}}^\forall$ homeomorphic to $\mathcal{ACF}_0/\cong$ (just in the language of fields, with no dependence relations).

Finally, in addition to Cantor space, Baire space and all their quotients, there are the topological spaces $\mathbb{R}$ and $[0,1]$ under the usual topologies. We have no idea whether any isomorphism space $\mathcal{C}/\cong$ might be homeomorphic to either of these: it seems unlikely, but for now it stands as a good question.

References

[1] C.J. Ash & J.F. Knight; *Computable Structures and the Hyperarithmetical Hierarchy* (Amsterdam: Elsevier, 2000).

[2] C.J. Ash, J.F. Knight, M.S. Manasse, & T.A. Slaman; Generic copies of countable structures, *Annals of Pure and Applied Logic* **42** (1989), 195–205.

[3] J. Chisholm; On intrisically 1-computable trees, unpublished MS.

[4] W. Calvert; The isomorphism problem for classes of computable fields, *Archive for Mathematical Logic* **43** (2004), 327–336.

[5] W. Calvert, D. Cummins, J.F. Knight, & S. Miller; Comparing classes of finite structures, *Algebra and Logic* **43** (2004), 365–373.

[6] W. Calvert & J.F. Knight; Classification from a computable viewpoint, *Bulletin of Symbolic Logic* **12** (2006), 191–218.

[7] R.G. Downey, A.G. Melnikov, & K.M. Ng; A Friedberg enumeration of equivalence structures, *Journal of Mathematical Logic* **17** (2017) 2.

[8] J.N.Y. Franklin & R. Miller; Measuring computable categoricity for classes of countable structures, to appear.

[9] M.D. Fried & M. Jarden; *Field Arithmetic* (Berlin: Springer, 1986).

[10] H. Friedman & L. Stanley; A Borel reducibility for classes of countable structures. *Journal of Symbolic Logic* **54** (1989), 894–914.

[11] S.S. Goncharov & J.F. Knight; Computable structure and non-structure theorems, *Algebra and Logic* **41** (2002) 6, 351–373.

[12] D. Hirschfeldt, K. Kramer, R. Miller, & A. Shlapentokh, Categoricity properties for computable algebraic fields, *Transactions of the American Mathematical Society* **367** (2015) 6, 3955–3980.

[13] J.F. Knight; Degrees coded in jumps of orderings, *Journal of Symbolic Logic* **51** (1986), 1034-1042.

[14] J.F. Knight, S. Miller, & M. Vanden Boom; Turing computable embeddings. *Journal of Symbolic Logic* **72** 3 (2007), 901–918.

[15] L. Kronecker; Grundzüge einer arithmetischen Theorie der algebraischen Größen, *J. f. Math.* **92** (1882), 1–122.

[16] K. Lange, R. Miller, & R.M. Steiner; Effective classification of computable structures, to appear in *The Notre Dame Journal of Formal Logic* **59** (2018) 1.

[17] R. Miller; Computable reducibility for Cantor space, to appear in the *Proceedings of the SouthEastern Logic Symposium* (World Scientific).

[18] R. Miller; Is it easier to factor a polynomial or to find a root? *Transactions of the American Mathematical Society*, **362** (2010) 10, 5261–5281.

[19] R. Miller; Revisiting uniform computable categoricity: for the sixtieth birthday of Prof. Rod Downey, in *Computability and Complexity: Essays Dedicated to Rodney G. Downey on the Occasion of His Sixtieth Birthday*, eds. A. Day, M. Fellows, N. Greenberg, B. Khoussainov, A. Melnikov, & F. Rosamond, **LNCS 10010** (Berlin: Springer-Verlag, 2017), 254–270.

[20] R. Miller; Computable transformations of structures, in *Unveiling Dynamics and Complexity: 13th Conference on Computability in Europe, CiE 2017*, eds. J. Kari, F. Manea, & I. Petre, *Lecture Notes in Computer Science* **10307** (Berlin: Springer-Verlag, 2017), 88–97.

[21] R. Miller & K.M. Ng; Finitary reducibility on equivalence relations, *Journal of Symbolic Logic* **81** (2016) 4, 1225–1254.

[22] A. Montalbán; Notes on the jump of a structure, *Mathematical Theory and Computational Practice* 2009, 372–378.

[23] A.A. Soskova & I.N. Soskov; A jump inversion theorem for the degree spectra, *Journal of Logic and Computation* **19** (2009) 1, 199–215.

[24] K. Weihrauch; *Computable Analysis* (Springer-Verlag, 2000).

DEPARTMENT OF MATHEMATICS
    QUEENS COLLEGE – C.U.N.Y.
      65-30 KISSENA BLVD.
        QUEENS, NEW YORK 11367 U.S.A.
PH.D. PROGRAMS IN MATHEMATICS & COMPUTER SCIENCE
    C.U.N.Y. GRADUATE CENTER
      365 FIFTH AVENUE
        NEW YORK, NEW YORK 10016 U.S.A.
  E-mail: Russell.Miller@qc.cuny.edu