

Computationally Categorical Fields via Fermat's Last Theorem

Russell Miller* & Hans Schoutens†

July 13, 2009

Abstract

We construct a computable, computably categorical field of infinite transcendence degree over \mathbb{Q} , using the Fermat polynomials and assorted results from algebraic geometry. We also show that this field has an intrinsically computable (infinite) transcendence basis.

1 Introduction

In computable model theory, we investigate the extent to which various model-theoretic constructions can be performed effectively. For instance, given two structures, model theorists naturally wish to consider isomorphisms between them. In computable model theory, we break this down into two problems. First we pose the *Isomorphism Problem*, in which we ask how difficult it is to determine in general whether two given structures are isomorphic to each other at all. If they are indeed isomorphic, then we ask about the difficulty of actually computing an isomorphism between them. This latter

*The first author was partially supported by grants numbered 67182-00-36, 68470-00 37, 69723-00 38, and 61467-00 39 from The City University of New York PSC-CUNY Research Award Program, and by grant #13397 from the Templeton Foundation. Both authors were supported by CUNY Collaborative Grant # 80209-04-12, and wish to acknowledge useful conversations with Gunther Cornelissen, Oleg Kudinov, Alexandra Shlapentokh, and Lucien Szpiro.

†The second author was partially supported by NSF grant # DMS-0500835, and by grants numbered 67195-00 36, 68093-00 37, 69135-00 38, and 61085-00 39 from The City University of New York PSC-CUNY Research Award Program.

question involves the notion of *computable categoricity*. When asking these questions, we usually assume that the structures themselves are computable, meaning that we can compute the functions and relations on them. If they are in fact isomorphic, then it is reasonable to ask whether there exists a computable isomorphism between them.

Fields were the first mathematical structures for which the notion of computable categoricity arose. The isomorphism problem for fields is addressed by Calvert in [2]. Long before that, though, in [8], Frohlich and Shepherdson had begun to consider the second question, by giving an example of (in their terminology; see their Corollary 5.51) two isomorphic, explicitly presented fields with no explicit isomorphism between them. This idea eventually grew into the following definition.

Definition 1.1. A computable structure \mathcal{A} is *computably categorical* if for every computable structure \mathcal{B} isomorphic to \mathcal{A} , there exists a computable isomorphism from \mathcal{A} onto \mathcal{B} .

From the point of view of our second question, such a structure is about as nice as it could be: no matter which two computable copies of \mathcal{A} we choose, there must exist a computable isomorphism between them. (An even nicer version, called *uniform computable categoricity* and examined in [4], requires not only that a computable isomorphism exist, but that we should be able to figure out a program for computing it, just given the ability to compute the structures \mathcal{A} and \mathcal{B} .) If the structure is not computably categorical, we may ask how strong an oracle is required to compute isomorphisms between two computable copies; for a consideration of this question for algebraic fields, for example, see [26].

Much research has been devoted to characterizing the computably categorical models of various theories, including work by Dzgoev, Goncharov, Khisamiev, Lempp, McCoy, Miller, Remmel, and Solomon. Some results are readily stated: we know that a computable linear order is computably categorical iff it has only finitely many pairs of consecutive elements, for example, and that a computable Boolean algebra is computably categorical iff it has finitely many atoms. On the other hand, the known structural characterization of computably categorical trees requires a description by recursion on the heights of finite trees. The question has been studied for a number of other theories as well, and results along these lines may be found in [9], [10], [11], [12], [13], [18], [20], [23], [29], and [30].

However, the original problem of computable categoricity for fields has defied all attempts at structural characterization. The most obvious conjecture would be that the transcendence degree of a field over its prime subfield should determine computable categoricity. For algebraically closed fields, this is indeed the case, as shown by Ershov in [5]: an ACF is computably categorical iff it has finite transcendence degree over its prime subfield. However, in the same work, Ershov built a field, algebraic over its prime subfield but not algebraically closed, which was not computably categorical. In this paper we refute the converse as well, by constructing a computably categorical field of infinite transcendence degree over the rationals \mathbb{Q} . Thus, neither implication in the original conjecture actually holds.

The counterexample F we build is readily described. It begins with an infinite, purely transcendental extension $\mathbb{Q}(x_0, x_1, \dots)$ of \mathbb{Q} . Then, for each i , we adjoin an element y_i such that (x_i, y_i) is a solution of the Fermat curve D_{p_i} defined by $X_i^{p_i} + Y_i^{p_i} = 1$, for an odd prime p_i . Thus each y_i is algebraic over x_i , but the infinite set $\{x_i : i \in \omega\}$ is still algebraically independent. Fermat's Last Theorem shows that each D_{p_i} has only the trivial solutions $(0, 1)$ and $(1, 0)$ in \mathbb{Q} , and the heart of our proof is a demonstration that there exists a computable sequence $p_0 < p_1 < \dots$ such that in F , every nontrivial solution of the equation of D_{p_i} generates the same subfield, namely $\mathbb{Q}(x_i, y_i)$. Therefore, mapping x_i and y_i to any nontrivial solution of D_{p_i} in a field \tilde{F} isomorphic to F will define an isomorphism. The algebraic geometry required is developed in Section 2, and the sequence $\langle p_i \rangle$ is chosen and the field presented in Section 3. In Section 4, we use this construction to give the first example of a computable field possessing an infinite, intrinsically computable transcendence basis.

We wish to make note here of unpublished work by Kudinov and Lvov. Working jointly, they addressed the same question, about computable categoricity for fields of infinite transcendence degree, and made significant progress on it. Like us, they combined techniques from algebraic geometry and computability theory, but their investigations were unfortunately cut short when Lvov passed away, and it has not been possible to reconstruct their work. As we understand it, their constructions did not make use of the Fermat polynomials – which suggests that there are alternative ways to approach this problem, awaiting discovery (or re-discovery) by an enterprising researcher. We salute the efforts of Kudinov and Lvov, and regret the demise of the latter.

We describe our principal conventions for this paper. A *computable field*

is a structure in the signature with addition and multiplication, whose domain is an initial segment of ω , and for which those two operations are computable and define a field. This conforms to the usual computable-model-theoretic definition of a computable structure. An introduction to such fields for non-logicians is given in [24] and continued in [25]. Classic references on computable fields include [5], [6], [7], [8], [19], [22], [28], and [32]. Our computability-theoretic notation is standard and can be found in [31].

The field \mathbb{Q} is known to be computably categorical, and so we will often just write \mathbb{Q} to denote a computable presentation of that field, without concern for the specifics of the presentation. Given a computable field F , its polynomial ring $F[X]$ may be viewed just as the set F^* of finite tuples of elements of F , with $\langle a_0, \dots, a_d \rangle$ identified with $\sum a_i X^i$. (For a perfect identification, ensure that if $a_d = 0$, then $d = 0$.) Iterating this process yields a computable presentation of the ring $F[X_1, \dots, X_n]$, uniformly in n .

Given a computable field F , we will treat any singly-generated field extension $F(x)$ as a computable field as well. To compute it, we will need to know whether x is algebraic over F or not, and if it is, we will need its minimal polynomial $p(X)$ over F . In the algebraic case, one views elements of $F(x)$ as F -linear combinations over the set $\{1, x, x^2, \dots, x^{d-1}\}$, where $d = \deg p$, with the obvious addition and multiplication (which requires knowledge of $p(X)$, of course). In the transcendental case, $F(x)$ is just the quotient field of the domain $F[X]$ given above, and this quotient field is computably presentable as the Cartesian product $F[x] \times (F[x] - \{0\})$ modulo a computable equivalence relation. We can iterate these extensions, even over infinitely many generators, as long as the minimal polynomial (or lack thereof) for each generator over the preceding ones is given effectively. Notice that the base field F is a computable subfield of each extension built this way.

2 Results from algebraic geometry

We now introduce some notation and review some results from algebraic geometry. Let k be a field of characteristic zero and fix an algebraic closure \bar{k} of k . A *variety* V over k is by definition an absolutely irreducible, separated, reduced scheme of finite type over k . Recall that V is called *absolutely irreducible* if the base change $V \times_k \bar{k}$, (or more correctly, the fiber product $V \times_{\text{Spec } k} \text{Spec } \bar{k}$) is irreducible. In fact, $V \times_k K$ is then irreducible for any extension K of k .

Suppose first that V is *affine*, that is to say, a closed subscheme of affine space \mathbb{A}_k^n for some n , or equivalently, a scheme of the form $V = \text{Spec } A$ with A an absolutely irreducible (meaning that $A \otimes_k \bar{k}$ is a domain), finitely generated k -algebra. For any extension field K of k , we define the K -rational points of V , denoted $V(K)$, as the collection of all $P \in \mathbb{A}_K^n(K) = K^n$ such that P lies on V (more precisely, in scheme-theoretic terms, this means that the morphism $\text{Spec } K \rightarrow \mathbb{A}_K^n$ induced by P factors through V , yielding a morphism $\text{Spec } K \rightarrow V$; or equivalently, that P is given by a maximal ideal \mathfrak{m}_P in A such that $A/\mathfrak{m}_P \cong K$). We define the *coordinate ring* $k[V]$ of V over k as the residue ring $k[x_1, \dots, x_n]/I_V$, where x_i are variables and I_V is the ideal of all polynomials vanishing on V (that is to say, the collection of all $f \in k[x_1, \dots, x_n]$ such that $f(P) = 0$ for all $P \in V(\bar{k})$). Hence, in the above notation, $A \cong k[V]$. Since we assumed V to be irreducible, I_V is a prime ideal and remains so when extended to $\bar{k}[x_1, \dots, x_n]$; we call such an ideal *absolutely prime*. Conversely, if I is an absolutely prime ideal in $k[x_1, \dots, x_n]$, the *variety defined by I* is the scheme $V := \text{Spec}(k[x_1, \dots, x_n]/I)$; for any extension $k \subseteq K$, its K -rational points are precisely those $P \in K^n$ such that $f(P) = 0$ for all $f \in I$. This establishes a one-one correspondence between (affine) subvarieties of \mathbb{A}_k^n and absolutely prime ideals in $k[x_1, \dots, x_n]$ (which in turn is in one-one correspondence with the class of all absolutely irreducible, finitely generated k -algebras).

We define the *function field* $k(V)$ of V over k as the field of fractions of $k[V]$. The function field of a variety is a so-called *birational invariant*, meaning that it only depends on the birational class of V . In particular, by resolution of singularities in characteristic zero, we may therefore assume, when studying the function field, that the variety has no singularities. The following quantities are all the same and are called the *dimension* of V :

- the transcendence degree of $k(V)$ over k ;
- the least number of *hypersurfaces* (=variety defined by a single, absolutely irreducible equation) H_1, \dots, H_d such that $V(\bar{k}) \cap H_1(\bar{k}) \cap \dots \cap H_d(\bar{k})$ is non-empty and finite;
- the combinatorial dimension of $V(\bar{k})$ viewed as a topological space via its Zariski topology;
- the Krull dimension of $k[V]$.

A point P on V is called *generic over k* if the field $k(P)$ generated by its coordinates is isomorphic to the function field of V , or equivalently, if the transcendence degree of $k(P)$ over k is equal to the dimension of V . Conversely, any point $P = (p_1, \dots, p_n) \in K^n$ in some extension field $K \supseteq k$ can be viewed as the generic point of an affine variety V over k , namely, let \mathfrak{m}_P be the maximal ideal in $K[x_1, \dots, x_n]$ generated by the linear polynomials $X_i - p_i$, let I_V be the contraction of this ideal to $k[x_1, \dots, x_n]$, and let V be the affine variety with coordinate ring $k[V] := k[x_1, \dots, x_n]/I_V$. One verifies that $k(V)$ is isomorphic with the field $k(P)$, that is to say, P is a generic point of V . If V is not affine, then its function field can still be defined as the function field $k(V_0)$ of any non-empty affine open subset V_0 of V . In particular, when studying the function field, we may take the variety to be projective as well. On occasion, we will use the following simple observation:

Lemma 2.1. *If V is a variety over k , then k is algebraically closed inside $k(V)$.*

Proof. By assumption $V \times_k \bar{k}$ is irreducible, showing that $k(V) \otimes_k \bar{k}$ is equal to the function field $\bar{k}(V)$ of V over \bar{k} . Towards a contradiction, assume $k \subsetneq l \subseteq k(V)$ is a finite extension of degree $d > 1$. Hence $l \otimes_k \bar{k} \cong \bar{k}^d$ is not a field, which contradicts that $l \otimes_k \bar{k} \subseteq k(V) \otimes_k \bar{k}$. \square

2.1 Curves

By a *curve* over k we will mean in these notes a non-singular (also called *smooth*), projective one-dimensional variety C over k . Recall that for any one-dimensional variety X defined over k , there exists a unique curve C to which X is birationally equivalent, and moreover, if X has no singularities, then it is isomorphic to an open subset of C . For instance, one can define a scheme structure on the set C_X of discrete valuations of $k(X)$ which are trivial on k , and then show that C_X is a curve birational to X (see for instance [17, Theorem 7.3.1] or [16, I. Corollary 6.11]); alternatively, we can take a completion \bar{X} of X (e.g., the Zariski closure of X viewed as a subvariety of some projective space over k), and then take its normalization. A one-dimensional variety over k which is birational to a curve C over k is sometimes called a *model* of C . Recall (see for instance [16, IV. Corollary 3.11]) that any curve admits an affine plane model (with at most nodes as singularities). By Lemma 2.1, a point P on C is generic (over k) if and only if $P \notin C(\bar{k})$.

The *genus* of a curve C can be defined as the number $1 - H_C(0)$, where H_C is the *Hilbert polynomial* of C , or equivalently, by Serre duality, as the vector space dimension over k of the global sections of the canonical sheaf ω_C of C (see for instance [16, §IV.1]). Note that the value of the genus does not change when making a base change. In case C is a plane curve (that is, a curve in \mathbb{P}_k^2), then its genus is calculated as $g(C) := (d-1)(d-2)/2$, where d is the degree of the defining equation of C (see for instance [16, I. Exercise 7.2]).

A morphism between curves $C \rightarrow D$ defined over k is either *constant* (that is to say, its image is a single point, necessarily k -rational), or otherwise *surjective*. Indeed, take non-empty affine open subsets C_0 and D_0 of C and D respectively so that $C \rightarrow D$ induces a k -algebra homomorphism $k[D_0] \rightarrow k[C_0]$. The map $C \rightarrow D$ is not dominant precisely when this homomorphism has a non-zero kernel, which then by necessity has to be a maximal ideal \mathfrak{m} of $k[D_0]$. Hence we have inclusions of fields $k \subseteq k[D_0]/\mathfrak{m} \subseteq k(C)$. By the Nullstellensatz, the former extension is algebraic, whence trivial by Lemma 2.1, showing that the image of $C \rightarrow D$ is the k -rational point on D determined by \mathfrak{m} . On the other hand, if $C \rightarrow D$ is dominant, then it must be surjective: since C and D are projective, $C \rightarrow D$ is proper whence in particular has a closed image. Since this image is also dense, it must be equal to D . Moreover, in the surjective case, we have an inclusion of function fields $k(D) \subseteq k(C)$, which is necessarily a finite extension, since both $k(C)$ and $k(D)$ have transcendence degree one over k . It follows that $C \rightarrow D$ is a *finite* morphism of *degree* $d := ([k(C) : k(D)])$; in particular, all the fibers of $C \rightarrow D$ are finite, of cardinality at most d . The following fact will be quite useful.

Proposition 2.2. *Let C be a curve and V any variety, both defined over a field k . We have a functorial bijection of sets*

$$\mathrm{Mor}_k(C, V) \cong V(k(C)),$$

and under this bijection, the morphism $C \rightarrow V$ is constant if, and only if, the associated $k(C)$ -rational point of V is k -rational.

Proof. Let P be a $k(C)$ -rational point on V , which therefore corresponds to a morphism $\mathrm{Spec} k(C) \rightarrow V$. By the valuative criterion for properness (see for instance [16, I. Proposition 6.8 or II. Theorem 4.7] or [17, Proposition 7.2.3]), this extends to a morphism $C \rightarrow V$. Conversely given a morphism $v : C \rightarrow V$,

choose affine opens $C_0 \subseteq C$ and $V_0 \subseteq V$ such that $v(C_0) \subseteq V_0$. Hence we have an induced k -algebra homomorphism $k[V_0] \rightarrow k[C_0]$. Composing this with the inclusion $k[C_0] \subseteq k(C_0) = k(C)$ then yields a $k(C)$ -rational point $\text{Spec } k(C) \rightarrow V$.

Functoriality here means that we can view $\text{Mor}_k(\cdot, V)$ and $V(\cdot)$ as contravariant functors, and the above bijection is compatible with these functors in the following sense. Given a non-constant morphism $C \rightarrow D$, composition yields a map $\text{Mor}_k(D, V) \rightarrow \text{Mor}_k(C, V)$. Moreover, we have an extension $k(D) \subset k(C)$ of function fields, giving rise to an inclusion $V(k(D)) \subset V(k(C))$. One now easily checks that we have a commutative diagram

$$\begin{array}{ccc} \text{Mor}_k(D, V) & \xrightarrow{\cong} & V(k(D)) \\ \downarrow & & \downarrow \\ \text{Mor}_k(C, V) & \xrightarrow[\cong]{} & V(k(C)) \end{array}$$

Moreover, these bijections are also compatible when viewed as (covariant) functors in their second component, that is to say, given a non-constant morphism $V \rightarrow W$, composition yields $\text{Mor}_k(C, V) \rightarrow \text{Mor}_k(C, W)$, and we get a natural map $V(k(C)) \rightarrow W(k(C))$, making the diagram

$$\begin{array}{ccc} \text{Mor}_k(C, V) & \xrightarrow{\cong} & V(k(C)) \\ \downarrow & & \downarrow \\ \text{Mor}_k(C, W) & \xrightarrow[\cong]{} & W(k(C)) \end{array}$$

commute as well. □

Moreover, if V is actually a curve, then by our above discussion, under this isomorphism, the morphism $C \rightarrow V$ is constant if, and only if, the associated $k(C)$ -rational point of V is in fact k -rational.

Given two curves C and D over k , we say that C *covers* (or *dominates*) D if there exists a non-constant morphism $C \rightarrow D$ defined over some field extension $K \supseteq k$. Note that such a $C \rightarrow D$ is then automatically finite and surjective. By Proposition 2.2, therefore, C does not cover D if and only if $\text{Mor}_K(C, D) \cong D(K)$ for every $K \supseteq k$.

Lemma 2.3. *Let C and D_i be curves over k such that no D_i covers C . If F is the field generated by all the function fields $k(D_i)$, then $C(F) = C(k)$.*

Proof. Let F_i be the subfield of F generated by all $k(D_j)$ with $j \leq i$, that is to say, F_i is defined inductively as $F_{i-1}(D_i)$ (with $F_0 := k$). It suffices to show by induction on i that $C(F_i) = C(k)$. Suppose P is an F_i -rational point on C , and we want to show that P is k -rational. Since $P \in C(F_{i-1}(D_i))$, we get from Proposition 2.2 a morphism $D_i \rightarrow C$ defined over F_{i-1} . Since D_i does not cover C , this morphism must be constant, showing that P belongs to $C(F_{i-1})$ and by induction, the latter is just $C(k)$. \square

We will also use the following well-known inequality (see for instance [16, Corollary 2.4]):

Theorem 2.4 (Hurwitz's Formula). *Let k be an algebraically closed field of characteristic zero and let C and D be curves of genus $g(C)$ and $g(D)$ respectively. If $C \rightarrow D$ is a finite morphism of degree d , then $g(C) - 1 \geq d(g(D) - 1)$.*

Proof. Since the characteristic is zero, the morphism is separable. In this case we even have an equality

$$2(g(C) - 1) = 2d(g(D) - 1) + \deg(R),$$

where R is the (effective) *ramification divisor* of $C \rightarrow D$; see for instance [16, IV. Corollary 2.4]. \square

Proposition 2.5. *Let k be a field of characteristic zero and let C be a curve over k of genus $g \geq 2$. Then the function field $K := k(C)$ of C is generated by the coordinates of any K -rational point P of C which is not k -rational, that is to say, for any $P \in C(K) \setminus C(k)$, the natural inclusion $k(P) \subseteq K$ is an equality.*

Proof. Let C_0 be a (non-empty) affine open subset of C and let $A := k[C_0]$ be its coordinate ring. For instance, if C_0 is a plane curve with affine equation $f = 0$, then $A = k[x, y]/fk[x, y]$ and K is the field of fractions of A . We need to show that if $P := (a, b) \in K^2 \setminus k^2$ satisfies $f(a, b) = 0$, then $k(P) := k(a, b)$ is equal to K . In any case, since P is not defined over k , it is a generic point of C by Lemma 2.1. In particular, $k(P)$ and K are isomorphic over k (for instance, in the planar example, the k -algebra $k[a, b]$ is easily seen to be k -isomorphic to A). In particular, by our discussion of Proposition 2.2, the embedding $k(P) \subseteq K$ is finite of degree $d := (K : k(P))$ inducing a finite morphism $C \rightarrow C$. We need to show that $d = 1$.

Assume first that k is algebraically closed, so that we can apply Theorem 2.4 to this morphism, giving $g - 1 \geq d(g - 1)$. Since $g \geq 2$, we must have $d = 1$, as we wanted to show. For the general case, it follows from the algebraically closed case that there exists a finite extension l of k such that $l(P) = l(C)$. Let π be such that $l = k(\pi)$, so that in particular $l(P) = k(P, \pi)$ and $l(C) = K(\pi)$. However, since π is algebraic over k whence over the k -isomorphic fields $k(P)$ and K , the degrees $(l(P) : k(P))$ and $(l(C) : K)$ are the same. By transitivity, it follows that $(K : k(P)) = 1$, as we wanted to show. \square

Remark 2.6. *The proposition is false in positive characteristic p , precisely because of the purely inseparable extension given by the Frobenius Frob_p : if P is a K -rational point of C , then so is its Frobenius transform $\text{Frob}_p(P)$, which clearly generates a proper subfield.*

Remark 2.7. *By the argument in the proof, $g(C) \geq g(D)$ for any cover $C \rightarrow D$ over k (not necessarily algebraically closed). Moreover, if the genera are the same and at least two, then $C \rightarrow D$ must have degree one, hence is an isomorphism. In summary, given a cover $C \rightarrow D$ of curves of genus at least two, either $C \cong D$ (and the cover itself is an isomorphism) or $g(C) > g(D)$.*

Remark 2.8. *Let C be a curve of genus g at least two and let $K := k(C)$ be its function field. Any non-constant morphism $C \rightarrow C$ is necessarily an automorphism by our previous remark. Hence under Proposition 2.2, we have a one-one correspondence between the K -rational generic points of C and the automorphisms of C . In particular, $C(K) \setminus C(k)$ has cardinality at most $84(g - 1)$, as this is the maximum number of automorphisms of C (see for instance [16, IV Exercise 2.5]). In particular, if C has no non-trivial automorphism (which is the ‘generic’ case for $g \geq 3$), then C has a unique generic point (in any fixed function field).*

2.2 General collections of curves

By a *general collection of curves over k* , we mean a countable set \mathcal{C} of curves over k of genus at least two, such that no two distinct curves in \mathcal{C} are isomorphic over \bar{k} . The *function field $k(\mathcal{C})$* of \mathcal{C} is by definition the field generated by all the function fields of curves in \mathcal{C} . More precisely, take a universal field Ω (that is, an algebraically closed field containing k and of cardinality larger than any of the fields we use otherwise), and for each $C \in \mathcal{C}$, fix a subfield

$k_C \subseteq \Omega$ isomorphic to its function field $k(C)$. Then $k(\mathcal{C})$ is the subfield of Ω generated by all the k_C with $C \in \mathcal{C}$. Note that the field generated by the subfields k_{C_1}, \dots, k_{C_i} is isomorphic to the function field of the product variety $C_1 \times \dots \times C_i$, for C_1, \dots, C_i distinct curves in \mathcal{C} . Moreover, $k(\mathcal{C})$ is the union of the function fields of all such products. In particular, $k(\mathcal{C})$ is well-defined up to isomorphism.

Proposition 2.9. *Let \mathcal{C} be a general collection of curves over k and let $k(\mathcal{C})$ be its function field. Suppose all curves in \mathcal{C} have genus at most g and let D be an arbitrary curve of genus at least g . Then the function field $k(D)$ embeds in $k(\mathcal{C})$ if and only if $D \in \mathcal{C}$.*

Proof. Let us write k_D for the image of $k(D)$ in $k(\mathcal{C})$. Let C_1, C_2, \dots be an enumeration of \mathcal{C} , and let $k_i := k(C_1 \times \dots \times C_i) \subseteq k(\mathcal{C})$ be as above, so that in particular $k(\mathcal{C})$ is the union of all k_i . Since k_D is finitely generated, it lies in some k_i . Let i be minimal such. By Proposition 2.5 and minimality of i , the function field $k_{i-1}(D)$ of D over k_{i-1} is isomorphic to the field generated by k_{i-1} and k_D , and hence has transcendence degree one over k_{i-1} by Lemma 2.1. The extension $k_{i-1}(D) \rightarrow k_i = k_{i-1}(C_i)$ is finite, since both fields have transcendence degree one over k_{i-1} , and hence by the discussion of Proposition 2.2, determines a cover $C_i \rightarrow D$ over k_{i-1} . If $C_i \neq D$, then $g(C_i) > g(D) \geq g$ by Remark 2.7, contradicting the assumption that $g(C_i) \leq g$. \square

We say that a general collection of curves \mathcal{C} is *non-covering*, if there is no cover relation between any two distinct curves in \mathcal{C} . Immediately from Remark 2.7 we get:

Lemma 2.10. *The collection of all curves over k of a fixed genus $g \geq 2$ is non-covering.* \square

Note that the collection in Lemma 2.10 is in one-one correspondence with the set of k -rational points of the *moduli space* \mathcal{M}_g of all curves of genus g . Recall that the moduli space has dimension $3(g-1)$, for $g \geq 2$.

Theorem 2.11. *Let \mathcal{C} be a non-covering collection of curves over k , and let $k(\mathcal{C})$ be its function field, as in Proposition 2.9. For any curve C in \mathcal{C} , any embedding of its function field into $k(\mathcal{C})$ has image equal to k_C . In particular, if P is a $k(\mathcal{C})$ -rational point on C which is not k -rational, then P is k_C -rational and the natural inclusion $k(P) \subseteq k_C$ is an equality.*

Proof. Let K be the subfield of $k(\mathcal{C})$ generated by all the function fields k_D with $D \in \mathcal{C}$ different from C . Since \mathcal{C} is non-covering, we can apply Lemma 2.3 over the subfield $k_C \subseteq k(\mathcal{C})$. Since $k(\mathcal{C})$ is generated by K and k_C , Lemma 2.3 shows that P is k_C -rational. As P is not k -rational, Proposition 2.5 then shows that $k(P) = k_C$, as we wanted to show. To prove the first assertion, let $k(C) \rightarrow k(\mathcal{C})$ be an embedding, and let P be the corresponding generic point of C defined over $k(\mathcal{C})$. In particular, $k(P)$ is the image of the above embedding, and we just argued that $k(P) = k_C$. \square

Corollary 2.12. *Let \mathcal{C} be a non-covering collection of curves over k , and let $k(\mathcal{C})$ be its function field. Then the Galois group G of $k(\mathcal{C})$ over k is equal to the direct product of the automorphism groups of each curve in \mathcal{C} . In particular, if all curves in \mathcal{C} have genus at most g , then G has exponent at most $84(g-1)$.*

Proof. Let C be a curve in \mathcal{C} with function field $k_C \subseteq k(\mathcal{C})$, and let $g \in G$. Since $g(k_C) \cong k_C$, it must be equal to it by Theorem 2.11. In other words, the restriction of g to k_C belongs to the Galois group H_C of k_C over k . In particular, the restriction map induces a split group homomorphism $G \rightarrow H_C$. Since $k(\mathcal{C})$ is generated by all the k_C with $C \in \mathcal{C}$, it follows that G is isomorphic to the direct product of all H_C . Since $H_C \cong \text{Aut}_k(C)$ has order at most $84(g(C)-1)$, the result follows (see Remark 2.8). \square

3 Construction of a computably categorical field of infinite transcendence degree

In this section, Q is a (countable) computable field of characteristic zero. Given a collection $\mathcal{C} = \{C_0, C_1, \dots\}$ of curves C_i over Q , we say that \mathcal{C} admits *effective Mordell-Weil* (*effective M-W* for short) if $|C_i(Q)|$ is computable as a function of i . Of course, this depends not only on \mathcal{C} , but also on the ordering of the curves in \mathcal{C} . Normally we assume a fixed computable listing of these curves: some computable function $g(i)$ gives the defining equation of each C_i over Q . We allow the possibility that $|C_i(Q)| = \aleph_0$, but in practice we will deal with collections \mathcal{C} for which all $|C_i(Q)|$ are finite. In this case, with a computably enumerable ground field Q , for \mathcal{C} to admit effective M-W means that we can computably determine all of the finitely many solutions of each C_i , since to do so we only need to be able to compute their number and then

search for them. Note that by Faltings' positive solution of the Mordell-Weil conjecture, each curve of genus two has only finitely many solutions over a number field.

For our construction, we need a non-covering collection of curves over the ground field $Q = \mathbb{Q}$ admitting effective M-W. For instance, we might in light of Lemma 2.10 ask whether there are infinitely many curves of a fixed genus for which we can effectively determine their \mathbb{Q} -rational points? After all, a "generic" choice will produce a curve without any \mathbb{Q} -rational points, a trivial instance of effective M-W. Although we do not know the answer to this, we can prove:

Theorem 3.1. *There exists a non-covering collection \mathcal{C} of curves over \mathbb{Q} admitting effective M-W.*

Proof. Our collection will consist of *Fermat curves*, which is to say, plane curves D_N with affine equation $d_N(X, Y) = X^N + Y^N - 1$. By Wiles' positive solution of Fermat's Theorem, the set of Fermat curves D_N with $N \geq 3$ satisfies effective M-W, since the only \mathbb{Q} -rational points of such a curve are $(1, 0)$ and $(0, 1)$ (and the point at infinity $(1 : -1 : 0)$). Moreover, the genus of D_N is equal to $(N - 1)(N - 2)/2$. So it remains to find a non-covering subset. To this end, we will choose inductively a set of prime exponents p_0, p_1, \dots as follows. Let $p_0 = 5$, and choose p_{i+1} to be the least prime bigger than $(4(p_i - 1)(p_i - 2))^2$. Now D_{p_i} does not cover D_{p_j} for $i > j$ by Lemma 3.2 below, and for $i < j$ by Remark 2.7. \square

Lemma 3.2. *Let C be a curve of genus $g \geq 2$ and let D_p be the Fermat curve of exponent p , as above. If $p > 64g^2$, then there is no cover relation between C and D_p .*

Proof. Since D_p has genus $(p - 1)(p - 2)/2 > g$, it cannot be covered by C by Remark 2.7. By the well-known estimate for the Euler's totient function $\varphi(n) \geq \sqrt{n}$, we get $\varphi(n) > 8g$ for all $n > p$. In particular, p and $(64g^2)!$ are relatively prime, so that by the proof of [1, Lemma 9.3], neither can there be a covering $D_p \rightarrow C$. \square

Fixing the sequence of primes $\langle p_s \rangle_{s \geq 0}$ defined in Theorem 3.1, we begin with the purely transcendental extension $\mathbb{Q}(x_0, x_1, \dots)$, and for each s , adjoin an element y_s satisfying the p_s -th Fermat curve $d_{p_s}(x_s, y_s) = x_s^{p_s} + y_s^{p_s} - 1 = 0$. We write $F_0 = \mathbb{Q}$ and $F_{s+1} = F_s(x_s)[y_s]/(d_{p_s}(x_s, y_s))$, so F_{s+1} has transcendence degree 1 over F_s , and we may take each F_s to be computable

within F , uniformly in s . Clearly F has infinite transcendence degree over \mathbb{Q} .

Theorem 3.3. *The computable field F built above is computably categorical and has infinite transcendence degree.*

Proof. Let \tilde{F} be any computable field isomorphic to F , say via a noncomputable isomorphism $\psi : F \rightarrow \tilde{F}$. We define our computable isomorphism f on each F_s in turn, starting with the unique embedding of F_0 onto the subfield $\tilde{F}_0 = \tilde{\mathbb{Q}}$ of \tilde{F} . Given $f \upharpoonright F_{s-1}$, search for any nonzero elements \tilde{x} and \tilde{y} in \tilde{F} such that $d_{p_s}(\tilde{x}, \tilde{y}) = 0$ there. We must find them, since $F \cong \tilde{F}$, and (\tilde{x}, \tilde{y}) is then a generic point of D_p over \mathbb{Q} . So we define $f(x_s) = \tilde{x}$ and $f(y_s) = \tilde{y}$, then extend this f to the rest of F_s , which is generated over F_{s-1} by these elements. By Theorem 2.11, this gives a field embedding of F_s onto the image $\psi(F_s)$. \square

One could begin with any prime > 3 as p_0 , thereby building countably many non-isomorphic computable, computably categorical fields of infinite transcendence degree. (Our use of Proposition 2.5 requires that $p_0 \neq 3$, since the curves must all have genus ≥ 2 .) And by varying the choice of the subsequent primes p_1, p_2, \dots , one could get uncountably many relatively computably categorical fields of infinite transcendence degree, although of course only countably many of them would be computably presentable.

Next we adapt Theorem 3.3 to a more general setting.

Theorem 3.4. *Let k be any finitely generated field of characteristic 0, and \mathcal{C} any computably enumerable collection of curves over k , all of genus ≥ 2 , such that \mathcal{C} is non-covering and admits effective Mordell-Weil over k . (In particular, $C(k)$ is finite for all $C \in \mathcal{C}$.) Then the function field $k(\mathcal{C})$ is computably categorical.*

Saying that \mathcal{C} is computably enumerable (or *c.e.*) means that there is a computable function g such that for every s , $g(s) \in k[X, Y]$ is a polynomial defining a curve C_s , and C_0, C_1, \dots is a list of all curves in \mathcal{C} without repetitions. Below we write g_s for $g(s)$.

Proof. Let F and \tilde{F} be computable fields isomorphic to $k(\mathcal{C})$. Now the finitely generated field k is computably enumerable within F , and for a fixed (not necessarily computable) isomorphism $\alpha : F \rightarrow \tilde{F}$, the image $\tilde{k} = \alpha(k)$ will likewise be *c.e.* within \tilde{F} . Moreover, the isomorphism $f_0 = \alpha \upharpoonright k$ from the

subfield $F_0 = k$ onto $\tilde{F}_0 = \tilde{k}$ is computable: we need only know the images under α of the finitely many generators of k .

As in Theorem 3.3, the construction is straightforward. For each s , with effective Mordell-Weil, we can compute the (finite) number j of distinct solutions of g_s in k . Search until we have found all j solutions of g_s in k and one additional solution (x_s, y_s) of g_s in F . Then we search in \tilde{k} for j distinct solutions to the polynomial $\tilde{g}_s(X, Y) \in \tilde{k}[X, Y]$ whose coefficients are the images of those of g_s under f_0 , and in \tilde{F} for one additional solution $(\tilde{x}_s, \tilde{y}_s)$ to \tilde{g}_s . We must find such a pair $(\tilde{x}_s, \tilde{y}_s)$, and we set $f(x_s) = \tilde{x}_s$ and $f(y_s) = \tilde{y}_s$. It is then easy to compute f on all of F , since Theorem 2.11 shows F to be generated by $\{x_s, y_s : s \in \omega\}$. For any $z \in F$, we compute $f(z)$ just by searching for an n and an $h \in k(X_1, \dots, X_n, Y_1, \dots, Y_n)$ with $h(x_1, \dots, x_n, y_1, \dots, y_n) = z$, since then $f(z) = h(\tilde{x}_1, \dots, \tilde{x}_n, \tilde{y}_1, \dots, \tilde{y}_n)$.

To see that this f really is an isomorphism, we appeal to the following lemma. In Theorem 3.3, the corresponding fact is immediate (and appears explicitly in the discussion of the proof of Proposition 4.4 below). Now in this more general situation, we make sure that it is safe to choose an arbitrary solution $(\tilde{x}_s, \tilde{y}_s)$ of \tilde{g}_s in $(\tilde{F} - \tilde{k})$.

Lemma 3.5. *For every s and any pairs $(x, y), (x', y') \in (F_{s+1} - F_s)^2$ with $g_s(x, y) = g_s(x', y') = 0$, every automorphism ψ of F_s which fixes k pointwise extends to a unique automorphism of F_{s+1} which maps x to x' and y to y' .*

Proof. By Lemma 2.1, F_s is algebraically closed within $F_s(C_s) = F_{s+1}$, so each of x and x' must be transcendental over F_s . Therefore ψ extends to a partial automorphism ψ' of F_{s+1} by mapping x to x' . Next, by absolute irreducibility, $g_s(x, Y)$ is irreducible in $F_s(x)[Y]$, and likewise for x' , so

$$F_s(x, y) \cong F_s(x)[Y]/(g_s(x, Y)) \cong F_s(x')[Y]/(g_s(x', Y)) \cong F_s(x', y'),$$

with the middle isomorphism being induced by ψ' on the quotients of the polynomial rings. (It is important here that ψ' fixes the coefficients of $g_s(X, Y)$.) But by Theorem 2.11, x and y together generate F_{s+1} over F_s , as do x' and y' , which makes it clear both that we have an automorphism of F_{s+1} , and that it is unique. \square

Our claim that the $f = \cup_s f_s$ constructed above is an isomorphism will follow from $(f^{-1} \circ \alpha)$ being an automorphism of F , where α was the given isomorphism from F to \tilde{F} . Since $f_0 = \alpha \upharpoonright k$, we know that $(f_0^{-1} \circ \alpha)$ is the

identity on $k = F_0$. For each s , $(f^{-1} \circ \alpha) \upharpoonright F_{s+1}$ is just the unique extension of $\psi = (f^{-1} \circ \alpha) \upharpoonright F_s$ given by the lemma, with $x = \alpha^{-1}(\tilde{x}_s)$, $y = \alpha^{-1}(\tilde{y}_s)$, $x' = x_s$, and $y' = y_s$. \square

We have a criterion for computable enumerability of certain collections \mathcal{C} of curves.

Theorem 3.6. *Let $g \geq 2$, let Q be a computable field, and let \mathcal{C} be an infinite subset of the Q -rational points of the moduli space \mathcal{M}_g . Suppose for each Q -rational point of the moduli space \mathcal{M}_g , we can effectively compute a defining set of equations over Q of an affine model of the curve of genus g determined by this Q -rational point. Then relative to a computable representation of the field extension $Q \subseteq Q(\mathcal{C})$, the subset \mathcal{C} is computably enumerable inside $\mathcal{M}_g(Q)$.*

Proof. Let C be a Q -rational point of \mathcal{M}_g , viewed as a curve over Q of genus g , and let $f_1, \dots, f_s \in Q[x_1, \dots, x_m]$ be the defining equations of an affine model of C (meaning that the function field of C is the field of fractions of $Q[x_1, \dots, x_m]/(f_1, \dots, f_s)Q[x_1, \dots, x_m]$). We now search for a solution $P \in Q(\mathcal{C})^m$ of the system of equations $f_1 = \dots = f_s = 0$ which is not defined over the computable subfield Q . If such a solution exists, then $C \in \mathcal{C}$ by Proposition 2.9. \square

Remark 3.7. *For instance, if $g = 2$, any curve C of genus 2 has an affine model C_f given by an equation $y^2 = f(x)$ with f a polynomial of degree 6 without double roots (so that in particular C_f is non-singular whence isomorphic to an open subset of C). Moreover, two such models C_{f_1} and C_{f_2} are birational to the same curve if and only if f_1 and f_2 are equal up to a fractional linear transformation, that is to say, if and only if*

$$f_1(x) = (cx + d)^6 \cdot f_2\left(\frac{ax + b}{cx + d}\right)$$

for some $a, b, c, d \in Q$ with $ad - bc \neq 0$ (see for instance [3, §1.1]). Put differently, the subset \mathbb{H} in the Hilbert scheme \mathbb{P}_Q^7 of degree six polynomials without a double root admits a natural action of the group of fractional linear transformations $\Gamma := \text{PGL}(2, Q)$, and the geometric quotient \mathbb{H}/Γ is the moduli space \mathcal{M}_2 .

4 An Intrinsically Computable Basis

The field F built in Theorem 3.3 has uncountably many automorphisms, as one sees by defining, for any $S \subseteq \omega$, an automorphism extending the map

$$x_i \mapsto \begin{cases} x_i, & \text{if } i \in S; \\ y_i, & \text{if } i \notin S. \end{cases}$$

Therefore, between F and any computable field $\tilde{F} \cong F$, there exist 2^ω -many isomorphisms. The point of Theorem 3.3 was that at least one must be computable, but the image of the transcendence basis $B = \{x_i : i \in \omega\}$ can have arbitrarily high Turing degree as well: its image under the above automorphism is Turing-equivalent to S . In the language of computable model theory, this says that B is far from being *intrinsically computable*.

Definition 4.1. Let \mathcal{M} be any computable structure, and R an n -ary relation on \mathcal{M} (generally not in the signature of \mathcal{M}). The *degree spectrum of R on \mathcal{M}* is the set

$$\text{DgSp}_{\mathcal{M}}(R) = \{\text{deg } f(R) : \mathcal{B} \cong \mathcal{M} \text{ via } f \text{ \& } \mathcal{B} \text{ is a computable structure}\}.$$

If $\text{DgSp}_{\mathcal{M}}(R) = \{\mathbf{0}\}$, then R is *intrinsically computable*.

So the degree spectrum measures the amount of complexity that can be added to (or withheld from) R under isomorphisms onto other computable structures. If that amount is bounded above (or below), then we think of R as being intrinsically no more complex than (or no less complex than) that bound. Degree spectra of relations have been studied widely in computable model theory; see [15] for a survey of results. With our field F , the degree spectrum of the basis B , viewed as a unary relation on F , contains every Turing degree, simply by the argument above using automorphisms from F to itself.

In this section we show that we can go to the opposite extreme: an infinite transcendence basis for a computable field can be intrinsically computable. This would be trivial if the basis were finite, of course, but for the infinite case, we believe this is the first proof of Theorem 4.2.

Theorem 4.2. *There exists a computable field F , of infinite transcendence degree over its prime field \mathbb{Q} , with an intrinsically computable transcendence basis.*

F will be the computable field defined in Theorem 3.3, and the key to the proof is knowing the exact number of solutions in F to each of the Fermat polynomials $d_{p_i}(X, Y)$ used to identify transcendence basis elements. Therefore, we need the following essential result of Tzermias and Leopoldt.

Theorem 4.3 (Tzermias [33]; Leopoldt [21]). *Over an algebraically closed field K of characteristic 0, the automorphism group of the projective curve $X^p + Y^p + Z^p$ is the semidirect product of the symmetric group S_3 and the group $(\mu(p))^2$, where $\mu(p)$ is the multiplicative group of p -th roots of unity in K . \square*

In particular, S_3 acts by permuting the coordinates in $\mathbb{P}^2(K)$, and $(\alpha, \beta) \in (\mu(p))^2$ maps $(x : y : z)$ to $(\alpha x : \beta y : z)$. Clearly each of these operations maps solutions of $X^p + Y^p + Z^p = 0$ to other solutions; the content of the theorem is that there are no other automorphisms.

Proposition 4.4. *The function field $\mathbb{Q}(D_p)$ of the Fermat curve D_p of prime exponent $p > 3$ contains exactly eight solutions (given below) to its defining equation $d_p(X, Y) = 0$.*

Proof. Since \mathbb{Q} contains no non-trivial p -th roots of unity, the automorphism group of D_p is isomorphic to S_3 , and corresponds by Remark 2.8 to the six generic points on D_p . By Wiles' positive solution to Fermat's Last Theorem, $D_p(\mathbb{Q})$ consists of three points, including the point at infinity $(1 : -1 : 0)$. Excluding the point at infinity gives a total of eight affine solutions to the equation $d_p = 0$. \square

We now describe this proof more explicitly for readers with less background in algebraic geometry. The function field $\mathbb{Q}(D_p)$ must contain at least one nontrivial solution (x, y) to $X^p + Y^p = 1$. Now with $p > 3$, D_p has genus ≥ 2 , so if $(x', y') \in (\mathbb{Q}(D_p))^2$ is any nontrivial solution, then x' and y' generate $\mathbb{Q}(D_p)$, by Proposition 2.5 (as do x and y also). Hence x' is transcendental over \mathbb{Q} , forcing $\mathbb{Q}(x) \cong \mathbb{Q}(x')$, which in turn yields a chain whose composition is an automorphism of $\mathbb{Q}(D_p)$:

$$\mathbb{Q}(x, y) \cong \mathbb{Q}(x)[Y]/(d_p(x, Y)) \cong \mathbb{Q}(x')[Y]/(d_p(x', Y)) \cong \mathbb{Q}(x', y').$$

Thus $(x' : y' : -1)$ gives an automorphism of the projective curve $X^p + Y^p + Z^p$, using Remark 2.8, and so Theorem 4.3 shows that the only projective solutions to $X^p + Y^p + Z^p = 0$ in the function field are $(\alpha x : \beta y : -1)$ and

permutations thereof, with α and β ranging over the p -th roots of unity. But all those roots except 1 lie outside of \mathbb{Q} , since p is odd. Therefore, any solution to $X^p + Y^p + Z^p = 0$ must be of the form $(x : y : -1)$ or a permutation of this, which translates to an affine solution $(\lambda x, \lambda y, -\lambda)$ or a permutation. For our purposes the affine solution must end with -1 , in order to yield $X^p + Y^p = 1$. Clearly $(x, y, -1)$ works, and choosing $\lambda = -\frac{1}{x}$ gives $(-1, -\frac{y}{x}, \frac{1}{x})$, while choosing $\lambda = -\frac{1}{y}$ gives $(-\frac{x}{y}, -1, \frac{1}{y})$, which both can be permuted to the appropriate form. However, no other value for λ can give the necessary coordinate -1 .

Corollary 4.5. *The six nontrivial solutions of $X^p + Y^p = 1$ in $\mathbb{Q}(D_p)$ are (x, y) , $(-\frac{y}{x}, \frac{1}{x})$, $(-\frac{x}{y}, \frac{1}{y})$, and the transpositions of these, where (x, y) is any one nontrivial solution. \square*

So in Theorem 3.3, when we built the computable isomorphism f from F onto an arbitrary computable copy \tilde{F} , there were actually six possible images for each x_s in \tilde{F} , which we identified by finding nontrivial solutions to $d_{p_s} = 0$. We simply chose $f(x_s)$ to be the first one we recognized, since Theorem 2.11 shows that each of the six is a valid choice. For Theorem 4.2, however, we need a transcendence basis each element of which has only one possible image in the target field.

Proof of Theorem 4.2. Let F be the computable field defined in Theorem 3.3, presented as

$$F = \mathbb{Q}(x_0, x_1, \dots)[y_0, y_1, \dots] / (d_{p_i}(x_i, y_i) : i \in \omega),$$

where $\langle p_i \rangle_{i \in \omega}$ was the computable sequence of primes chosen there. In the nice presentation F constructed there, the transcendence basis $B = \{x_i : i \in \omega\}$ was computable, of course, but not intrinsically computable, as noted above. However, we can build an intrinsically computable basis A from it.

The nontrivial solutions of each $d_{p_i}(X, Y) = 0$ in F were given by Proposition 4.4: $\langle x_i, y_i \rangle$, $\langle y_i, x_i \rangle$, $\langle \frac{1}{y_i}, \frac{-x_i}{y_i} \rangle$, $\langle \frac{-x_i}{y_i}, \frac{1}{y_i} \rangle$, $\langle \frac{1}{x_i}, \frac{-y_i}{x_i} \rangle$, $\langle \frac{-y_i}{x_i}, \frac{1}{x_i} \rangle$, and no others. For each i , let z_i be the sum of these six numbers:

$$z_i = x_i + y_i + \frac{1}{y_i} - \frac{x_i}{y_i} + \frac{1}{x_i} - \frac{y_i}{x_i} = \left(\frac{1 - x_i}{1 - x_i^{p_i}} \right) y_i^{p_i-1} + \left(\frac{x_i - 1}{x_i} \right) y_i + \frac{x_i^2 + 1}{x_i},$$

where we have used the algebraic relation $y_i^{p_i} = 1 - x_i^{p_i}$ to rewrite the expression as a polynomial in y_i over $\mathbb{Q}(x_i)$. Thus z_i clearly lies in $\mathbb{Q}(x_i)[y_i]$ but

not in \mathbb{Q} , as y_i is algebraic of degree p_i over $\mathbb{Q}(x_i)$, and so z_i is algebraically interdependent with x_i in F over \mathbb{Q} . So we may replace each x_i by z_i in the basis B , giving a new transcendence basis $A = \{z_i : i \in \omega\}$ for F .

Now each z_i is defined by a simple existential formula $\psi_i(z)$: it is the sum of six distinct elements $x \in F$, each satisfying

$$x \neq 0 \ \& \ x \neq 1 \ \& \ (\exists y \in F)[x^{p_i} + y^{p_i} = 1].$$

By Proposition 4.4, this formula $\psi_i(z)$ actually defines z_i in F . So the transcendence basis A is defined by

$$(\exists i \in \omega)\psi_i(z).$$

(This definition is in fact a computable infinitary Σ_1 formula, quantifying over ω as well as over E , which is acceptable for us, though anathema to model theorists.) Therefore, the image $f(A)$ under any isomorphism f (not necessarily computable) from F onto any computable field \tilde{F} must also be existentially defined in \tilde{F} , hence c.e. Now we invoke a simple lemma to show that $f(A)$ is computable in \tilde{F} .

Lemma 4.6. *In a computable field K , if a transcendence basis A is computably enumerable, then A is computable.*

Proof. K is algebraic over the purely transcendental extension $Q(A)$, where Q is the prime subfield of K . Given any $t \in K$, we find an $n \in \omega$ and a nonzero polynomial $p(T) \in Q(a_0, \dots, a_n)[T]$ with $p(t) = 0$, by enumerating $A = \{a_0, a_1, \dots\}$ and searching through such polynomials (for all n simultaneously). But then t is algebraic over $Q(a_0, \dots, a_n)$, and so $t \in A$ iff $t \in \{a_0, \dots, a_n\}$. \square

This completes the proof of Theorem 4.2. \square

We remark that Proposition 4.4 was essential to this proof of Theorem 4.2. In a field as described more generally in Theorem 3.4, one might not be able to compute the number of solutions to C_i in its function field, and therefore it could be impossible to state the definitions $\psi_i(z)$ uniformly in i .

We also remark that this same proof shows the basis A to be *relatively intrinsically computable*: its image $f(A)$ under any isomorphism f is computable relative to the Turing degree of the field $f(F)$. Likewise, this field F , and also those described by Theorem 3.4, are *relatively computably categorical*: if \tilde{F} is a field isomorphic to F but of arbitrary Turing degree, then there is an isomorphism between these fields which is computable in the degree of \tilde{F} .

5 Questions

Several basic questions arise from the algebraic geometry presented here. For instance, we do not know whether cover relations exist among any of the Fermat curves D_p and D_q with $3 < p < q$. D_3 is a special case, of course, since its genus equals 1, and in Theorem 3.1 we chose specific primes such that no cover relations exist among those curves. But the general case remains unclear.

Also, while Theorems 3.3 and 3.4 give examples of computably categorical fields of infinite transcendence degree, they leave open the larger question of determining a structural criterion for computable categoricity of fields. If anything, they make this question appear more difficult, since now we know that infinite transcendence degree is not sufficient to rule out computable categoricity.

Specifically, we remarked in Lemma 2.10 that there is no cover relation among any collection of curves of any single fixed genus ≥ 2 . Whether this gives rise to other computably categorical fields of infinite transcendence degree depends on whether one can produce a computable collection of such curves which admits effective Mordell-Weil. It seems plausible that this can be done, but the proof remains elusive, and introduces the larger question of computing the number of solutions in \mathbb{Q} of arbitrary curves satisfying Mordell-Weil. Of course, the exact number of \mathbb{Q} -rational points on a curve given by a polynomial $q \in \mathbb{Q}[X, Y]$ defines a limitwise monotonic function $\lim_s g(q, s)$ (that is, with g computable and with $g(q, s) \leq g(q, s + 1)$ for all q and s), and when we restrict to curves of genus ≥ 2 , this function is total, i.e., the limit is always finite. However, it remains unknown whether the exact number is computable from q or not.

We regard this as a natural and challenging question. It seems related to Hilbert's Tenth Problem for \mathbb{Q} , which demands a decision procedure for the existence of \mathbb{Q} -rational solutions to polynomials in arbitrarily many variables over \mathbb{Q} . However, it is not immediately clear whether such a decision procedure would allow one to compute the exact number of rational solutions (even given that this number is finite). Conversely, even if we had effective Mordell-Weil for all curves of genus ≥ 2 , we would still not have an obvious algorithm for Hilbert's Tenth Problem for polynomials in more than two variables.

References

- [1] M.H. Baker, E. González-Jiménez, J. González & B. Poonen; Finiteness results for modular curves of genus at least 2. *Amer. J. Math.* **127** (2005) 6, 1325–1387.
- [2] W. Calvert; The isomorphism problem for classes of computable fields, *Archive for Mathematical Logic* **43** (2004), 327–336.
- [3] J. Cassels & E. Flynn; *Prolegomena to a middlebrow arithmetic of curves of genus 2*, London Mathematical Society Lecture Note Series, vol. 230. (Cambridge: Cambridge University Press, 1996.)
- [4] R.G. Downey, D.R. Hirschfeldt, & B. Khoussainov; Uniformity in Computable Structure Theory, *Algebra and Logic* **42** (2003), 318–332.
- [5] Yu.L. Ershov; Theorie der Numerierungen, *Zeits. Math. Logik Grund. Math.* **23** (1977), 289–371.
- [6] Yu.L. Ershov & S.S. Goncharov; Constructive fields, Section 2.5 in *Constructive Models* (New York: Kluwer Academic/Plenum Press, 2000).
- [7] M.D. Fried & M. Jarden; *Field Arithmetic* (Berlin: Springer-Verlag, 1986).
- [8] A. Frohlich & J.C. Shepherdson; Effective procedures in field theory, *Phil. Trans. Royal Soc. London, Series A* **248** (1956) 950, 407–432.
- [9] S.S. Goncharov; Autostability and computable families of constructivizations, *Algebra and Logic* **14** (1975), 647–680 (Russian), 392–409 (English translation).
- [10] S.S. Goncharov; Nonequivalent constructivizations, *Proc. Math. Inst. Sib. Branch Acad. Sci.* (Novosibirsk: Nauka, 1982).
- [11] S.S. Goncharov; Autostable models and algorithmic dimensions, *Handbook of Recursive Mathematics*, vol. 1 (Amsterdam: Elsevier, 1998), 261–287.
- [12] S.S. Goncharov & V.D. Dzgoev; Autostability of models, *Algebra and Logic* **19** (1980), 45–58 (Russian), 28–37 (English translation).

- [13] S.S. Goncharov, S. Lempp & R. Solomon; The computable dimension of ordered abelian groups, *Advances in Mathematics* **175** (2003) 1, 102–143.
- [14] A. Grothendieck; Revêtements étales et groupe fondamental (SGA 1), Séminaire de géométrie algébrique du Bois Marie 1960-61, *Documents Mathématiques (Paris)* 3, Société Mathématique de France, Paris, 2003. Updated and annotated reprint of the original [Berlin: Springer, Lecture Notes in Mathematics 224, 1971].
- [15] V.S. Harizanov; Pure computable model theory, *Handbook of Recursive Mathematics*, vol. 1 (Amsterdam: Elsevier, 1998), 3–114.
- [16] R. Hartshorne; *Algebraic Geometry* (New York: Springer-Verlag, 1977).
- [17] G. Kempf; *Algebraic Varieties*, London Mathematical Society Lecture Note Series, vol. 172. (Cambridge: Cambridge University Press, 1993).
- [18] B. Khoussainov & R.A. Shore; Effective Model Theory: The Number of Models and their Complexity, *Models And Computability: Invited Papers from Logic Colloquium '97*, ed. S.B. Cooper & J.K. Truss, LMSLNS **259** (Cambridge: Cambridge University Press, 1999), 193–240.
- [19] L. Kronecker; Grundzüge einer arithmetischen Theorie der algebraischen Größen, *J. f. Math.* **92** (1882), 1–122.
- [20] S. Lempp, C. McCoy, R. Miller, & R. Solomon; Computable categoricity of trees of finite height, *Journal of Symbolic Logic* **70** (2005), 151–215.
- [21] H.-W. Leopoldt; Über die Automorphismengruppen des Fermatkörpers, *Journal of Number Theory* **56** (1996) 2, 256–282.
- [22] G. Metakides & A. Nerode; Effective content of field theory, *Annals of Mathematical Logic* **17** (1979), 289–320.
- [23] R.G. Miller; The computable dimension of trees of infinite height, *Journal of Symbolic Logic* **70** (2005), 111–141.
- [24] R.G. Miller; Computable fields and Galois theory, *Notices of the AMS* **55** (August 2008) 7, 798–807.

- [25] R.G. Miller; Computability and differential fields: a tutorial, to appear in *Differential Algebra and Related Topics: Proceedings of the Second International Workshop*, eds. L. Guo & W. Sit. Also available at qcpages.qc.cuny.edu/~rmiller/research.html.
- [26] R.G. Miller; \mathbf{d} -Computable categoricity for algebraic fields, to appear in the *Journal of Symbolic Logic*.
- [27] B. Poonen; Unramified covers of Galois covers of low genus curves, *Math Res. Lett.* **12** (2005) 4, 475–481.
- [28] M. Rabin; Computable algebra, general theory, and theory of computable fields, *Transactions of the American Mathematical Society* **95** (1960), 341–360.
- [29] J.B. Remmel; Recursively categorical linear orderings, *Proceedings of the American Mathematical Society* **83** (1981), 387–391.
- [30] J.B. Remmel; Recursive isomorphism types of recursive Boolean algebras, *Journal of Symbolic Logic* **46** (1981), 572–594.
- [31] R.I. Soare; *Recursively Enumerable Sets and Degrees* (New York: Springer-Verlag, 1987).
- [32] V. Stoltenberg-Hansen & J.V. Tucker; Computable Rings and Fields, in *Handbook of Computability Theory*, ed. E.R. Griffor (Amsterdam: Elsevier, 1999), 363–447.
- [33] P. Tzermias; The group of automorphisms of the Fermat curve, *Journal of Number Theory* **53** (1995) 1, 173–178.

DEPARTMENT OF MATHEMATICS
 QUEENS COLLEGE – C.U.N.Y.
 65-30 KISSENA BLVD.
 FLUSHING, 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

DEPARTMENT OF MATHEMATICS
NEW YORK CITY COLLEGE OF TECHNOLOGY
300 JAY STREET
BROOKLYN, NEW YORK 11201 U.S.A.
PH.D. PROGRAM IN MATHEMATICS
C.U.N.Y. GRADUATE CENTER
365 FIFTH AVENUE
NEW YORK, NEW YORK 10016 U.S.A.
E-mail: hschoutens@citytech.cuny.edu