MODEL THEORY OF DIFFERENTIAL FIELDS

MOSHE KAMENSKY

1. Introduction

Consider a finite system X of equations $p_i(x_1, \dots, x_k) = 0$, for $1 \le i \le n$, where p_i are polynomials over $\mathbb{Q}$. What can be said about the set S of solutions of X whose coordinates are all roots of unity? For example, in what cases is S infinite?

To make the question more precise, one needs in particular to specify where the solutions to X are taken to begin with. However, all roots of unity are contained in the algebraic closure of $\mathbb{Q}$, so we are free to choose any field that contains it, for example $\mathbb{C}$. In this case, the set $X(\mathbb{C})$ of complex valued solutions has a geometric structure (essentially a complex analytic manifold, though it may have some "corners"). Without choosing $\mathbb{C}$, we may still view X as a geometric object: it is an example of an *affine algebraic variety*. The set S we would like to study is the intersection $X(\mathbb{C}) \cap T$, where T is the set of all n-tuples of roots of unity.

If T was also an algebraic variety, i.e., if $T = Y(\mathbb{C})$ for some system of polynomial equations, it would be possible to study such questions via geometry: algebraic varieties admit a good notion of dimension, and a well developed theory of intersection, predicting the dimension and number of points in the intersection. However, no such system Y exists: it is easy to check that the set T does not (collectively) satisfy any non-trivial polynomial relation.

A fundamental idea then is to enlarge the class of possible equations, in a manner that will provide non-trivial information on the set T, while keeping the structure of such equations manageable. There are a number of useful choices for such structures, and in this course we will concentrate on a differential one.

In place of the field $\mathbb{C}$ we chose above, we consider for instance the field K of meromorphic functions on an open disc D in the complex plane. K is equipped with a natural additional structure: the derivative $'$ on meromorphic functions. Using this additional structure, we may form new equations, on top of the polynomial ones we had before. In particular, the set of all roots of unity is contained in the set of solutions of the equation $p(x) = 0$, where $p(x) = x'$ is now a *differential polynomial*. This follows from the fact that the set of solutions to this equation is $\mathbb{C}$, an algebraically closed field, but a more conceptual approach is replace the above equation by the equation $l(x) = 0$, where $l(x) = \frac{x'}{x}$. The point is that $l_K : K^* \to K$ is a group homomorphism from the multiplicative group to the additive one, and roots of unity are precisely the torsion points for the multiplicative structure, so must go to 0 in the additive one (which is torsion free).

As above, the full kernel of l is $\mathbb{C}^*$, so much "smaller" than K. With the theory of dimension that we will present, this will be one of the main examples of (the points of) a set of dimension 1. This example is in fact not very useful: after passing to the kernel, we no longer see the differential structure, and we are back to usual commutative algebra over $\mathbb{C}$. However, an analogous consideration for a different class of groups plays a role in one of the main applications of the model theory of differential fields to arithmetic, namely the proof (by Hrushovski) of the relative Mordell–Lang conjecture. We will explain elements of this proof, following the book [1], which is dedicated to it.

As another example, there is a classical function, the j-function, which is a holomorphic function on the upper-half plane $\mathbb{H} = \{z \in \mathbb{C} | \Im z > 0\}$. This function admits (and almost characterised by) the property that $j(z) = j(m(z))$ whenever $m : \mathbb{H} \to \mathbb{H}$ is an *integral Möbius transformation*, i.e., $m(z) = \frac{az+b}{cz+d}$, where $a, b, c, d \in \mathbb{Z}$ and $ad - bc = 1$. In other words, if $W_m = \{\langle z, m(z)\rangle | z \in \mathbb{H}\} \subseteq \mathbb{H} \times \mathbb{H}$ is the graph of m, then $j(W_m) \subseteq \Delta \subseteq \mathbb{A}^2$, so an algebraic subvariety W_m of $\mathbb{H}^2$ is mapped into a (proper) algebraic subvariety of $\mathbb{A}^2$. Since j is far from being a polynomial map, it is interesting to ask if there are other non-trivial algebraic relations among the values of j. It turns out that the answer is "no", even if one includes first and second derivatives of j (and also for generalisations of j). A crucial point in the proof (by Casale–Freitag–Nagloo, [2]) is that j satisfies a particular algebraic differential equation (of order 3). The proof proceeds by analysing the structure of this equation.

A fundamental observation is that the properties we are interested in are *algebraic* properties of the solutions, and thus one can expect to derive them from algebraic properties of the equations themselves. There is a number of approaches to formalising this idea, our basic notion will be that of a *differential field*:

derivation

Definition 1.0.1. Let A be a commutative ring. A *derivation* of A is an additive function $\partial : A \to A$ satisfying the Leibniz rule: $\partial(ab) = \partial(a)b + a\partial(b)$ for all $a, b \in A$.

differential ring

A *differential ring* is a pair $\langle A, \partial\rangle$ with A and ∂ as above.

Starting with a differential ring $A = \langle A, \partial\rangle$, it is possible to consider polynomial differential equations with coefficients in A, and their solutions (in a differential ring extending A). This way, algebraic properties of differential equations can be studied without reference to any analytic properties of their solutions. One expects to have a theory similar the theory of polynomial equations and their solutions, and such a theory, *differential algebra* and *differential algebraic geometry* indeed exists, but we will see that it is substantially more complicated than the algebraic situation. In particular, algebraically defined dimension is difficult to work with, there are no analogs of Noetherianity and primary decomposition, and so on.

Better tools are obtained via model theory. The relevant first order theory DCF belongs to the well-behaved class of ω-stable theories. Model theory provides good notions of rank for systems of equations (definable sets) in such a theory, and these turn out to be very useful in this setting. A central result is a detailed classification of definable sets of rank 1, which is the basis to the applications mentioned above. From a different point of view, DCF provides a non-degenerate example of an ω-stable theory, and examples of interesting model theoretic phenomena (for instance, distinction between Morley and Lascar ranks).

1.1. **More bibliography.** In addition to the references mentioned above, relevant information is contained in [9] and in [4]. For general model theory, [8] and [11] are useful.

1.2. **Tentative outline.**

- Review of first order logic (structures, models, formulas, theories, compactness) (1)
- The theory of fields, affine algebraic varieties (2)
- Quantifier elimination, model companions, ACF (3–4)
- More varieties, prolongations, DCF (5–6)
- Imaginaries (7)
- Morley rank, strong minimality, ω-stability (8–9)
- General properties of strongly minimal sets, orthogonality, Zilber trichotomy (abstract geometries?) (10–12)
- Abelian varieties, Isogenies, Manin kernels (13–14)
- Strong minimality of Fuchsian equations (CFN §3,4) (15–16)
- Zil'ber trichotomy in DCF (Zariski geometries/Jet spaces) (17–18)
- Geometric triviality of Fuchsian equations (CFN 5) (19)

2. Preliminaries

2.1. **Review of first order logic.** For completeness we recall the basic definitions. You might prefer to look in a basic logic book, or jump directly to Example 2.1.8.

Definition 2.1.1. A (1-sorted) *first order structure* is given by: **first order structure**

(1) A set M ("the universe")
(2) For every finite set J, a boolean sub-algebra D_J of $\mathcal{P}(M^J)$ ("definable subsets")

satisfying:

(1) If $X \in D_I$ and $Y \in D_J$ where I, J are disjoint, then $X \times Y \in D_{I\cup J}$
(2) For any function $t : I \to J$, for all $X \in D_J$, $t^*X = \{f \circ t | f \in X\} \in D_I$

Less formally, if $f \in M^J$ and $t : I \to J$, then $f \circ t \in M^I$ is the point obtained from f by picking the coordinates as dictated by t, and t^*X is the image of X under this map. In particular:

- If t is a permutation of J, then $f \circ t$ is a tuple obtained by permuting the coordinates
- If t is the inclusion of a subset I, t^* is the projection to the coordinates in I
- If J is a singleton (and t is the unique function from I), then t^* is the diagonal map from $M^J = M$ to M^I.

All other cases are determined as combinations of these ones. The definition for the general (multi-sorted) case is similar, and we will review it later. It is clear that given an arbitrary collection D of subsets of M^I, for various I, there is a smallest structure with universe M where all these subsets are definable. We will call it the structure generated by D.

Exercise 2.1.2. Assume $X \in D_I(M)$ and $Y \in D_J(M)$ for some structure M. Show that $\{f \in M^{I\cup J} | f_I \in X, f_J \in Y\} \in D_{I\cup J}(M)$, where f_I is the restriction of f to I. □

For X, Y definable in a structure M, we say that a function $f : X \to Y$ is definable if its graph $\Gamma_f = \{\langle x, y\rangle \in X \times Y | y = f(x)\}$ is definable.

Exercise 2.1.3. For each $t : I \to J$ and $X \in D_J$, the function $f \mapsto f \circ t$ is definable □

Exercise 2.1.4. If X, Y, Z are definable in M, and $f : X \to Z$, $g : Y \to Z$ are definable, then $X \times_Z Y = \{\langle x, y\rangle \in X \times Y | f(x) = g(y)\}$ is definable. □

The notion of a structure is important, but for us it will be more useful to have a more syntactic description, for a number of reasons. First, as will be seen below, it is a convenient and natural way to describe definable sets. More importantly, the syntax provides a way to relate "the same" definable subset of two different structures. Our variant of syntax is given as follows:

first order language

Definition 2.1.5. A (relational, 1-sorted) *first order language* is given by a set $\mathcal{F}_I$ of "formulas in the variables I", for every finite set I, along with:

(1) Functorial[1] maps $t_* : \mathcal{F}_I \to \mathcal{F}_J$ for every function $t : I \to J$.
(2) Operations $\to: \mathcal{F}_I \times \mathcal{F}_I \to \mathcal{F}_I$ and $\exists i : F_I \to \mathcal{F}_{I\setminus\{i\}}$ for all i.
(3) Prescribed elements $\mathbf{0} \in \mathcal{F}_\emptyset$ and $= \in \mathcal{F}_\mathbf{2}$

The sets I should be thought of as variables. If $I = \{x, y, z, \dots\}$, we write $\phi(x, y, z, \dots)$ for a typical element of $\mathcal{F}_I$, and call it "a formula in the free variables I". We write $\phi \to \psi$ in place of $\to (\phi, \psi)$, reading "ϕ implies ψ", etc. The operations t_* correspond to variable substitution.

We make the following abbreviations:

- $\neg\phi := \phi \to \mathbf{0}$ ("not ϕ"). We set $\mathbf{1} := \neg\mathbf{0}$

[1]This means that $(t \circ s)_* = t_* \circ s_*$ for all t, s, and t_* is the identity map whenever t is

- $\phi \vee \psi := (\neg\phi) \to \psi$, $\phi \wedge \psi := \neg((\neg\phi) \vee (\neg\psi))$
- $\forall x\phi := \neg\exists x(\neg\phi)$
- $\exists! x\phi := \exists x\phi \wedge \forall y, z(t_*\phi \wedge s_*\phi \to y = z)$, where $t, s : \{x\} \to \{y, z\}$ send x to y and to z, respectively.

We note that our definition of the syntax is somewhat more general than usual, but this will not make a substantial difference.

The relation between the syntax and the semantics is given by the following definition:

Definition 2.1.6. Let $\mathcal{F} = (\mathcal{F}_I)_I$ be a language. An $\mathcal{F}$-*structure* consists of a set M and an assignment $\phi \mapsto \phi(M) \subseteq M^I$ for each $\phi \in \mathcal{F}_I$, such that: **$\mathcal{F}$-structure**

- For each $t : I \to J$ and each $\phi \in F_I$, $(t_*\phi)(M) = \{f \in M^J | f \circ t \in \phi(M)\}$
- $\mathbf{0}(M) = \emptyset$, $=(M) = \{\langle m, m\rangle | m \in M\}$
- $(\phi \to \psi)(M) = \phi(M)^c \cup \psi(M)$
- $(\exists x\phi)(M) = t^*\phi(M)$, where $t : I \setminus \{x\} \to I$ is the inclusion map

Exercise 2.1.7. If $\mathcal{F}$ is a first order language and M is an $\mathcal{F}$-structure, then, with the collection of subsets $D_I = \{\phi(M) | \phi \in \mathcal{F}_I\}$ it is a first order structure.

Conversely, every first-order structure is an $\mathcal{F}$-structure, for a canonically defined $\mathcal{F}$. □

Given a first order language $\mathcal{F}$, and collection of subsets $R_I \subseteq \mathcal{F}_I$ for each I, it is easy to see that there is a smallest sub-language $\mathcal{F}' \subseteq \mathcal{F}$ containing the R_I. If $\mathcal{F}' = \mathcal{F}$, we say that $\mathcal{F}$ is generated by the R_I, and often describe only the R_I. Any $\mathcal{F}$ structure is determined by its restriction to R_I.

Furthermore, given a set R_I for each I, it is possible to construct a "free" language generated by the R_I, and in practice one restricts to languages of this form (the R_I are called a *signature*). The freeness implies that any assignment of subsets (of the correct arity) to the elements of R_I extends to and $\mathcal{F}$ structure. **signature**

Example 2.1.8. Let A be a commutative ring. We consider the first order language $\mathcal{F}$ generated by the formulas $R_I = \{p = 0 | p \in A[I]\}$, where $A[I]$ is the algebra of polynomials in the variables I, with coefficients in A. We will call this the language of (commutative) A-algebras.

Any commutative A-algebra B determines a structure for this language, by assigning to $p = 0$ the set of solutions of this equation in B. □

So far, we did not consider any substantial way of restricting the structures. For instance, in the last example, there are many more $\mathcal{F}$-structures than commutative A-algebras. This can be fixed by noting that the language considered above can be used to describe those structure that are commutative A-algebras.

To make this precise, we note that by definition, for an element $\phi \in \mathcal{F}_\emptyset$ and an $\mathcal{F}$-structure M, $\phi(M) \subseteq \mathbf{1} = \{\emptyset\} = M^\emptyset$. Such a ϕ is called a *sentence*, and we say that ϕ holds in M, or that M is a *model of* ϕ, if $\phi(M) = \mathbf{1}$. Similarly, if $\mathcal{T}$ is a set of sentences, we say that M is a *model of* $\mathcal{T}$ if it is a model of every element of $\mathcal{T}$. We denote by $\mathcal{Mod}(\mathcal{T})$ the class of models of $\mathcal{T}$, and say that a class of this form (for some $\mathcal{T}$) is an *elementary class*.

sentence **model of ϕ** **model of $\mathcal{T}$** **elementary class**

theory of C **logically follows** **theory** **axiomatises**

Conversely, if C is a class of $\mathcal{F}$-structures, the *theory of* C, denoted by $\mathcal{Th}(C)$, is the set of sentences that hold in all members of C. A sentence ϕ *logically follows* from a set of sentences $\mathcal{T}$ if $\phi \in \mathcal{Th}(\mathcal{Mod}(\mathcal{T}))$. A set of the form $\mathcal{Th}(C)$ for some class C is called a *theory*. Thus, a theory is a set of sentences closed under implication. Given a set of sentences $\mathcal{T}_0$, there is a smallest theory that contains it (namely, $\mathcal{T} = \mathcal{Th}(\mathcal{Mod}(\mathcal{T}_0))$), and we normally do not distinguish between $\mathcal{T}$ and $\mathcal{T}_0$ (one says that T_0 *axiomatises* $\mathcal{T}$).

consistent

A theory $\mathcal{T}$ is said to be *consistent* if it has a model, i.e., if $\mathcal{Mod}(\mathcal{T})$ is non-empty.

Example 2.1.9. The class of commutative A-algebras is elementary (in the language of A-algebras). Some examples of sentences that axiomatise the theory that shows it are:

(1) $\forall x \forall y \exists ! z (x + y - z = 0)$
(2) $\forall x \forall y \exists ! z (x * y - z = 0)$
(3) $\forall x, y (x - y = 0 \to x{=}y)$
(4) $\forall x, y, z, w (x + y - z = 0 \wedge y + x - w = 0 \to z{=}w)$
(5) ...

□

What is an example of a non-elementary class? A little experimenting shows that if A is finite (for example, $A = \mathbb{F}_p$), there is no theory axiomatising the class of finite A-algebras. To prove this, we recall:

Theorem 2.1.10 (The Compactness Theorem)**.** *If every finite subset of a theory $\mathcal{T}$ is consistent, then $\mathcal{T}$ is consistent.*

This theorem can be reformulated in many ways. For example:

Exercise 2.1.11. If a theory $\mathcal{T}$ implies a sentence ϕ, then a finite subset of $\mathcal{T}$ implies ϕ as well. □

As an application, we show that indeed the class of finite $\mathbb{F}_p$-algebras is not axiomatisable:

Corollary 2.1.12. *If $\mathcal{T}$ is a theory that has finite models of unbounded size, then it has an infinite model.*

Proof. We first note that for each $n \in \mathbb{N}$, there is a sentence ϕ_n whose models are structures of size at least n. Namely, ϕ_n is given by

$$\exists x_1, \dots, x_n (\bigwedge_{i<j\leq n} x_i \neq x_j)$$

Now, assume that $\mathcal{T}$ has arbitrary large models. Then every finite finite subset of $\mathcal{T}_1 = \mathcal{T} \cup \{\phi_n | n \in \mathbb{N}\}$ is consistent, since it has only finitely many ϕ_i. By compactness, $\mathcal{T}_1$ is also consistent, and each model of $\mathcal{T}_1$ is an infinite model of $\mathcal{T}$. □

End of lecture 1, Mar 22

2.2. **Expansion by constants.** There is a well defined notion of homomorphism for A-algebras. We discuss a variant of it for structures.

We fix a first-order language $\mathcal{F}$ and a subset $R_I \subseteq \mathcal{F}_I$ for each I, that freely generate $\mathcal{F}$, so that each R-structure extends uniquely to an $\mathcal{F}$-structure. We may say R-structures in place of $\mathcal{F}$-structures to stress when properties depend on R rather than $\mathcal{F}$. The elements of R_I are called *basic* (or *quantifier free*), and we assume they include $\mathbf{0}$ and $=$.

quantifier free

Example 2.2.1. We modify Example 2.1.8 so that $R_{x,y}$ includes $\exists! z(x+y-z = 0)$ and $\exists! z(xy - z = 0)$, in addition to all polynomial equations. This will be our typical example. □

Definition 2.2.2. Let M and N be two R-structures. A *homomorphism* from M to N is a function $f : M \to N$ such that for all $\phi \in R_I$ and all $x \in M^I$, if $x \in \phi(M)$ then $f \circ x \in \phi(N)$.

homomorphism

A subset M of a structure N such that the inclusion is a homomorphism is called a *substructure* of N. Note that a formula $\phi(x, y)$ defines (the graph of) a function on the set $\exists! y\phi(x, y)$, so if such a formula is included in R_I, a substructure is closed under the function defined by ϕ.

substructure

Exercise 2.2.3. Show that in the setting of Example 2.2.1, when M and N are two structures viewed as A-algebras, a homomorphism from M to N is the same as a map of A-algebras. Show also that such a homomorphism need not satisfy the condition in the definition for an arbitrary formula ϕ. □

For a commutative A, the class of A-algebras can be described as follows: We start with the (elementary) class C of all commutative rings, fix a structure A, and consider the class of pairs (M, f), where M is in C and $f : A \to M$ is a homomorphism. We may repeat the same procedure with an arbitrary class of structures C and an arbitrary structure A, and obtain the class of structure in C over A, denoted $C_{A/}$.

In the case of commutative rings, we saw that the class obtained in this manner is elementary. The same is true in general, with a similar construction:

Definition 2.2.4. Let $\mathcal{F}$ be a language, and A a set. By the *expansion by constants* of $\mathcal{F}$ we mean the language $\mathcal{F}^A$ given by ${\mathcal{F}^A}_I = \bigcup_{A_0 \subseteq A} \mathcal{F}_{I \cup A_0}$, where the union is over all finite subsets A_0 of A.

expansion by constants

If M is an $\mathcal{F}$-structure and $A \subseteq M$ (or, more generally, we are given an injective map from A to M), we view M as an $\mathcal{F}^A$-structure in the obvious way.

If $\mathcal{T}$ is a theory in $\mathcal{F}$ and A is an $\mathcal{F}$-structure, we let $\mathcal{T}_A$ be the theory obtained from $\mathcal{T}$ by adding all $\phi(a)$ for $\phi \in R_I$ such that $a \in \phi(A)$.

Exercise 2.2.5. The models of T_A are (naturally identified with) models M of $\mathcal{T}$ along with a homomorphism from A to M. $\square$

universal formula
universal sentence

We may now show another application of the compactness theorem. By a *universal formula* we mean one of the form $\forall x\phi(x,y)$, where $\phi(x,y) \in R_I$ (so $(x,y) \in I$ are tuples of any length). A *universal sentence* is a sentence which is a universal formula. For a theory $\mathcal{T}$, we denote by $\mathcal{T}_\forall$ the set of universal sentences implied by $\mathcal{T}$. Clearly, if M satisfies a given universal sentence ϕ, then any substructure satisfies ϕ. We claim that the converse also holds:

Proposition 2.2.6. *If $\mathcal{T}$ is any theory, then $\mathcal{Mod}(\mathcal{T}_\forall)$ is the class of substructures of models of $\mathcal{T}$ (in particular, this class is elementary)*

Proof. We already mentioned one direction, so we need to prove: if M is a model of $T_\forall$, then M is a substructure of a model of $\mathcal{T}$. We have seen above that such a model is the same as a model of T_M, so we need to show that T_M is consistent.

Assuming not, by compactness a finite subset T_0 of T_M is inconsistent, and by taking conjunction, we may assume that $T_0 = \{\phi_0 \wedge \phi_1\}$, where ϕ_0 is in the language $\mathcal{F}$ of $\mathcal{T}$, and $\phi_1 = \psi(m)$, with $\psi \in R$ and $m \in M$. So $\mathcal{T}$ implies $\neg\psi(m)$ for arbitrary m, but then $\mathcal{T}$ implies $\forall x\neg\psi(x)$, an element of $T_\forall$, contradicting that M was a model of it. $\square$

Example 2.2.7. Let $\mathbb{F}$ be the theory of fields, in the language of Example 2.2.1 (so that a structure is a set along with binary functions $+$ and $\cdot$). The standard axiomatisation of $\mathbb{F}$ involves a non-universal sentence $\forall x\exists y(x = 0 \vee xy = 1)$. One could ask if this axiom can be replaced by a universal one. According to the proposition, the answer is no: a substructure of a field need not be a field.

Can we describe $\mathbb{F}_\forall$? An element of this theory is the sentence $\forall x, y(xy = 0 \to x = 0 \vee y = 0)$, i.e., there are no zero-divisors. This is the full theory, again according to the last proposition, since every integral domain is the substructure of a field (its field of fractions).

We note the dependence of these notions on the collection of quantifier free formulas: If we include the set $\phi(x) = \exists y(xy = 1)$ of invertible elements as "quantifier free", fields are axiomatised by the universal sentence $\forall x(x = 0 \vee \phi(x))$. $\square$

2.3. **Quantifier free sets in the theory of fields.** We now completely specialize to the setup of Example 2.2.1, and the theory $\mathbb{F}$ of fields from the last example (possibly over a given ring A).

Our goal is to describe some of the structure of sets definable without quantifiers. This is (part of) the subject of *algebraic geometry*, and we will take a few small bites of it. We fix a field K, and denote by $S = K[x_1, \dots, x_n]$ the polynomial ring over K in n variables. Each $p \in S$ determines a function $p : K^n \to K$.

Zariski closed subset

Definition 2.3.1. A *Zariski closed subset* of K^n is the set of solutions of a finite system of equations $p_1 = \cdots = p_m = 0$ with $p_i \in S$

Example 2.3.2. For $K = \mathbb{R}$, the equation $x^2 + y^2 - 1 = 0$ shows that the unit circle is a Zariski closed subset of $\mathbb{R}^2$. □

For a subset $I_0 \subseteq S$, we set $Z(I_0) = \{a \in K^n | p(a) = 0 \forall p \in I_0\}$. To which extent is I_0 determined by $Z(I_0)$? We consider several cases:

- If $I_0 = p, q$, then for every $r \in S$, $Z(I_0) = Z(p, rp + q)$.
- For every $p \in S$, $Z(p) = Z(p^2)$
- For $K = \mathbb{R}$, $Z(x^2 + 1) = Z(1) = \emptyset$

In the last case, there is no easily expressible algebraic relation between $x^2 + 1$ and 1 that would account for the equality, but the equality depends on the field $\mathbb{R}$: if we consider solutions in other fields (for instance, $\mathbb{C}$), the sets of solutions are no longer the same. We will return to this case later.

For the first instance, the general phenomenon is that a system of polynomials defines the same subset as the ideal it generates, so to remove the ambiguity we may restrict attention to ideals. This does not resolve the ambiguity presented in the second example, since p is (usually) not in the ideal generated by p^2. We will describe an explicit description later, but for the moment we bypass the difficulty with the following definition: For any subset $Y \subseteq K^n$, we let $I(Y) = \{p \in S | p(y) = 0 \forall y \in Y\}$.

Clearly, for all $Y \subseteq K^n$ the set $I(Y)$ is an ideal in S, and directly from the definition it follows that $I_0 \subseteq I(Z(I_0))$ and $Y \subseteq Z(I(Y))$ for all $I_0 \subseteq S$ and $Y \subseteq K^n$. The examples above show that the inclusion can be strict, even when I_0 is an ideal, but we have the following:

Exercise 2.3.3. For all $I_0 \subseteq S$, $Z(I(Z(I_0))) = Z(I_0)$ and for all $Y \subseteq K^n$, $I(Z(I(Y))) = I(Y)$ □

A priori, a subset of the form $Z(I)$ is not Zariski closed, according to our definition, since we required it to be the zero set of a *finite* number of polynomials. Of course, if I is finitely generated, as an ideal, this is not an issue. It turns out that all ideals in S are finitely generated, i.e., S is a *Noetherian ring*. This is known as **Noetherian ring**

Theorem 2.3.4 (Hilbert's basis theorem). *If A is a Noetherian ring, then so is the polynomial ring $A[x]$.*

By induction, every polynomial ring in finitely many variables over A is Noetherian. Hence, the Zariski closed subsets are precisely the subsets of the form $Z(I)$ for some ideal I in the polynomial algebra. One could also ask:

Question 2.3.5. what ideals are of the form $I(Y)$, for some subset $Y \subseteq K^n$? □

and we will answer this question later. End of lecture 2, Mar 28

2.3.6. *Affine varieties.* One issue with our definition of "Zariski closed subset", is that it is only defined as a subset, rather than a standalone object. In some cases, we would like to consider them independently of the embedding into K^n. So we would like to view them as some kind of a "geometric space", where the geometry is determined by the algebra of functions on it.

k-algebra of functions Let k be a (commutative) ring, and X a set. By a *k-algebra of functions* on X we mean a k-sub-algebra of the algebra k^X of all functions from X to k (with pointwise operations). Given such an algebra S, every element $x \in X$ determines a k-algebra homomorphism $\phi_x : S \to k$, given by $\phi_x(s) = s(x)$.

affine variety **Definition 2.3.7.** An *affine variety* over a field k is a pair $\langle S, X\rangle$, where X is a set, and S is a k-algebra of functions on X. They are required to satisfy the following conditions:

(1) The k-algebra S is finitely generated
(2) The map $x \mapsto \phi_x$ (as above) is a bijection from X to the set $Hom_{k-alg}(S, k)$ of k-algebra homomorphisms from S to k.

Example 2.3.8. For any infinite field k and any natural number n, the pair $\mathbb{A}^n = \langle k[x_1, \dots, x_n], k^n\rangle$ is an affine variety, where we identify each element p of $S = k[x_1, \dots, x_n]$ with the function it defines on k^n. Indeed, S is visibly finitely generated, and k-algebra maps from S to k are in canonical correspondence with elements of $k^{\{x_1,\dots,x_n\}} = k^n$ (via the required map). □

Exercise 2.3.9. Where did we use that k is infinite, and why is it enough? □

We remark that the datum of an affine variety is completely determined by the algebra S, since X is identifies with $Hom_{k-alg}(S, k)$. However, we do not know, at this point, which algebras occur as the algebra of an affine variety. This is related to our previous question 2.3.5[2]:

Proposition 2.3.10. *Let $Z \subseteq K^n$ be a Zariski closed subset. Then the pair $\langle K[x_1, \dots, x_n]/I(Z), Z\rangle$ is an affine variety over K.*

Proof. We first remark that $S = K[x_1, \dots, x_n]/I(Z)$ may indeed be viewed as an algebra of functions on Z: we identify the class of a polynomial p with the function on Z that the restriction of p determines. By definition of $I(Z)$ this does not depend on the choice of p, and if p determines the zero function, it is in $I(Z)$, again by definition of this ideal.

Clearly S is finitely-generated. If $z, w \in Z$ determine the same homomorphism from S to k, then they also determine the same homomorphism from the algebra of polynomials, so $z = w$ by the previous case. It remains to see that an arbitrary homomorphism $\phi : S \to k$ corresponds to a point of Z. We already know it corresponds to a point x of K^n, and by definition, $s(x) = 0$ for all $s \in I(Z)$. Since Z was Zariski closed, this implies that $x \in Z$ by Ex. 2.3.3. □

[2]We assume for convenience that the base field is infinite, but the statements are easily modified for the general case

We would like to assert that every affine variety is of the above form, but this is only true up to isomorphism, so to state this precisely, we need to define what are maps between affine varieties. Intuitively, a map from $\langle S, X\rangle$ to $\langle T, Y\rangle$ should be a function of sets from X to Y that preserves the algebra of functions, i.e.:

Definition 2.3.11. A map from the affine k-variety $\langle S, X\rangle$ to the affine k-variety $\langle T, Y\rangle$ is a function $g : X \to Y$ such that $t \circ g \in S$ for all $t \in T$.

Notice that in this case, the function $t \mapsto t \circ g$ is a k-algebra homomorphism. While the definition is reasonable (and correct) there is a simpler description:

Exercise 2.3.12. In the situation of the definition, show that every k-algebra map from T to S corresponds to a unique map of affine varieties. □

Example 2.3.13. By way of a sanity check, setting $Y = K^0$ a singleton, and $S = k$ in the last exercise, we see that the points of X correspond to maps from Y, as expected. □

It is obvious that the composition of maps of varieties is again such a map, and as usual, an isomorphism of affine varieties is a map that has an inverse with respect to composition.

We may now formulate the converse to 2.3.10:

Proposition 2.3.14. *If $\langle K[x_1, \dots, x_n]/I, X\rangle$ is an affine variety, then $I = I(Z(I))$, and the variety is isomorphic to $\langle K[x_1, \dots, x_n]/I, Z(I)\rangle$. Every affine variety is isomorphic to one of this form.*

Proof. Let $Z = Z(I)$, and $J = I(Z)$. We need to show that $J \subseteq I$, so let $p \in J$. We need to show that the image s of p in $S = K[x_1, \dots, x_n]/I$ is 0, and since X is an affine variety, for that it suffices to show that $\phi(s) = 0$ for all $\phi : S \to k$. But each such ϕ is represented by some $z \in Z$, so $\phi(s) = 0$ since $J = I(Z)$.

The fact that X is isomorphic to Z follows, since they have the same algebra, and for the last part, if $\langle S, X\rangle$ is an affine variety, S is finite generated, so there is a surjective map to S from some polynomial ring, with some kernel I, which is of the form above by the first part. □

We extend the definition of Zariski closed subsets to arbitrary affine varieties: A Zariski closed subset of $\langle S, X\rangle$ is one of the form $Z = Z(I_0) = \{z \in X | s(z) = 0 \forall s \in I_0\}$. Since S is a quotient of a polynomial ring, and a quotient of a Noetherian ring is Noetherian, a finite number of elements of S suffice to define each Zariski closed subset. A map from $\langle S, X\rangle$ to $\langle T, Y\rangle$ where the algebra map $T \to S$ is surjective is called a *closed embedding* of X in Y. It identifies X with the closed subset $Z(I) \subseteq Y$, where I is the kernel of the algebra map. Note that not every embedding is closed: the subset $X = K^*$ of invertible elements in K has the structure of an affine variety, with algebra of functions given by the localisation $K[x, \frac{1}{x}] = K[x, y]/xy - 1$, **closed embedding** and the inclusion of X in K is the map of varieties that corresponds to the inclusion of $K[x]$ in $K[x, \frac{1}{x}]$, so is not closed.

2.3.15. *Dimension.* Can we prove that the circle, given by $x^2+y^2=1$ in $\mathbb{R}^2$, is not isomorphic to the sphere $x^2+y^2+z^2=1$ in $\mathbb{R}^3$? If p is a polynomial isomorphism from S^1 to S^2, it is in particular a diffeomorphism between them, viewed as smooth manifolds. Such a diffeomorphism does not exists, since the two manifolds have different dimensions.

We would like to make a similar argument, but without passing to the smooth category. In other words, we would like to define the dimension algebraically. The analog of smooth manifolds is difficult to define in this setting, and instead we will use a definition that is similar to the one for vector spaces.

Recall that the dimension of a finitely generated vector space V can be defined in two equivalent ways:

- It is the number of elements in each basis of V
- It is the length of the longest chain of subspaces $0 = V_0 \subset V_1 \subset \cdots \subset V_n = V$, where all inclusions are proper.

For algebraic varieties, the situation is similar but more complicated. Starting with the second approach, a direct generalization would be the longest chain of proper Zariski closed subsets.

Example 2.3.16. The subset $\{0,1,2\} \subseteq \mathbb{R}$ (given inside $\mathbb{R}$ as the zero set of $x(x-1)(x-2)$) is expected to be 0-dimensional, but the chain $\{0\} \subset \{0,1\} \subset \{0,1,2\}$ has length 3. □

Example 2.3.17. The subset defined by the equations $xz = yz = 0$ is the union of the plane $z = 0$ and the line $x = y = 0$ through it. Hence we expect its dimension to be 2, but the plane is a closed subset whose dimension should be 2 as well. □

The last two example are examples of reducible varieties:

reducible
irreducible variety

Definition 2.3.18. An affine variety is *reducible* if it is the union of finitely many proper closed subsets. Otherwise, it is called an *irreducible variety*.

If $\langle S, X\rangle$ is an affine variety, the Noetherian property of S implies that any descending chain of closed subsets of X is finite. In particular, X is a finite union of irreducible subvarieties, called its irreducible components. We expect the dimension of X to be equal to the maximal dimension of a component, and if $\langle S, X\rangle$ is irreducible, we expect each closed subset to be of lower dimension. We arrive at the following definition:

Krull dimension

Definition 2.3.19. The *Krull dimension* of a non-empty affine variety X is the maximal length of a chain $X_0 \subset \cdots \subset X_n$ of irreducible subvarieties X_i of X. The dimension of $\emptyset$ is $-\infty$.

Algebraically, the variety $\langle S, X\rangle$ is irreducible precisely when S is an integral domain. Hence, the dimension is the maximal length of a certain kind of chains of prime ideals (namely, those that correspond to proper subvarieties).

End of lecture 3, Mar 29

Exercise 2.3.20. Show that a finite, non-empty variety is irreducible if and only if it consists of one point. Conclude that a variety is 0-dimensional if and only if it is finite and non-empty. □

Example 2.3.21. For each $i \leq n$, the ideal generated by $x_{i+1}, \dots, x_n$ in $k[x_1, \dots, x_n]$ is prime, since the quotient is $k[x_1, \dots, x_i]$. Hence the dimension of k^n is at least n (geometrically, the corresponding subvarieties are the affine subspaces determined by the coordinate axes) □

To obtain more precise information about the dimension, we need a stronger tie between the algebra and the geometry. We will therefore work on the assumption that the correspondence between closed irreducible subsets and prime ideals is a bijection. We will later see that this assumption holds when the base field is algebraically closed. At the moment, we simply switch to algebra: the assumption is only needed to relate the algebraic results to the geometric notion of dimension. The algebraic counterpart of the definition of dimension is:

Definition 2.3.22. The *Krull dimension* of a (Noetherian) non-zero ring A is the maximal length of a chain of prime ideals $p_0 \subset \dots p_n \subset A$.

Krull dimension

One way to compute the dimension precisely, we may use the following result (by Emmy Noether):

Theorem 2.3.23 (Noether Normalization)**.** *If A is generated by n elements as an algebra over a field k, then there is a polynomial sub-algebra $B = k[y_1, \dots, y_m] \subseteq A$ such that A is finite over B, with $m < n$ if the n elements satisfy a relation.*

Recall that A is a *finite algebra* over B if it is finitely generated over B as a module. One significance of this condition (in the geometric case) to dimension is provided by the following:

finite algebra

Theorem 2.3.24. *If $A \subseteq B$ is a finite ring extension, the map restriction map $q \mapsto q \cap A$ from prime ideals in B to prime ideals in A is surjective and strongly inclusion preserving: if $q_1 \subset q_2$ then $q_1 \cap A \subset q_2 \cap A$*

Each point of an affine variety determines a maximal (hence prime) ideal, and if $q \subseteq A$ is of this form, the theorem asserts the fibre over that point is (at most) 0-dimensional. In fact, in the geometric case finite inclusions correspond to surjective, proper maps with finite fibres. In any case, the theorem implies that Krull dimensions of A and B are equal. Together with Noether normalization, it thus reduces the computation of Krull dimension to polynomial rings

Claim 2.3.25. *For every field k, the Krull dimension of $A = k[x_1, \dots, x_n]$ is n.*

Proof. We already saw one direction. In the other, let $0 = p_0 \subset \cdots \subset p_m \subset A$ be a strict chain of prime ideals. Let $f \in p_1$. By Noether normalization A/f is finite over $k[y_1, \ldots, y_r]$ for $r < n$, so the dimension of A/f is r by induction. But p_i/f are a chain of primes in A/f, so $m \le r < n$. □

Noether Normalization also provides a second way to determine the dimension of A, when A is a finitely generated integral domain over k. To explain it, we recall some definitions:

algebraically independent

Definition 2.3.26. Let A be a k-algebra. A subset $B \subseteq A$ is *algebraically independent* over k if $p(\bar{b}) \neq 0$ for every non-zero polynomial p over k.

In other words, B is independent over k if the tautological k-algebra map from $k[B]$ to A is injective. When A is a field extension of k, this is equivalent to the statement that the inclusion of B in A extends to a map of fields $k(B) \to A$.

transcendence degree

By Zorn's lemma, each algebra A admits a maximal independent subset. If B is such a subset, every element of A is algebraic over the sub-algebra generated by B (so if A is a field, this is an algebraic field extension). It is a basic fact (that we will prove later in greater generality) that when A is a field, all maximal independent subsets have the same cardinality, which is called the *transcendence degree* of A over k. In particular, the transcendence degree is preserved by algebraic extensions.

Corollary 2.3.27. *Let A be a finitely generated integral domain over a field k. Then the Krull dimension of A is equal to the transcendence degree of the fraction field $K(A)$ over k.*

Proof. Let $B \subseteq A$ be as in Noether normalization. Then $K(A)$ is a finite field extension of $K(B)$, so has the same transcendence degree. On the other hand, we saw that A and B have the same Krull dimension. So the statement reduces to the case of polynomial algebras, where it is obvious. □

2.4. **Differential algebra.**

We would like to repeat some of the above ideas with *differential* polynomial equations, instead of polynomials ones. To do that, we introduce a formal notion of derivative:

derivation

Definition 2.4.1. Let A be a commutative ring. A *derivation* of A is an additive function $\partial : A \to A$ satisfying the Leibniz rule: $\partial(ab) = \partial(a)b + a\partial(b)$ for all $a, b \in A$.

differential ring

A *differential ring* is a pair $\langle A, \partial \rangle$ with A and ∂ as above. A map of differential rings is a map of rings that commutes with the derivation. A is a domain, a field, etc., if it so as a ring.

If $p(x)$ is a polynomial over a commutative ring A, we considered the solutions to the equation $p(x) = 0$ in A (or extension of it). The analogous notion in the differential case is provided by differential polynomials:

differential polynomial

Definition 2.4.2. Let A be a ring, and I a set of variables. A *differential polynomial* in I over A is a polynomial over A in the variables $x^{(i)}$, where

$x \in I$ and i a natural number.

We identify $x^{(0)}$ with x, and sometimes write x', x'', etc. in place of $x^{(1)}$, $x^{(2)}$, etc. Note that the definition does not require a derivation on A. As usual, the collection $A\{I\}$ of differential polynomials is an A-algebra, and the differential structure is provided by the fact that each derivation on A extends canonically to $A\{I\}$:

Exercise 2.4.3. Let $A = \langle A, \partial \rangle$ be a differential ring and I a set. There is a unique derivation ∂ on $A\{I\}$ determined by the requirements that $\partial(x^{(i)}) = x^{(i+1)}$ for all $x \in I$ and $i \in \mathbb{N}$, and the map $A \to A\{I\}$ is a map of differential rings. It classifies functions from I to differential ring extensions of A: the restriction $Hom_{\langle A,\partial \rangle}(A\{I\}, B) \to B^I$ is a bijection for every differential ring B over A (in other words, it is the free differential A-algebra on I) □

As in the algebraic situation, we will mostly be interested in sets of solutions in fields:

Definition 2.4.4. Let $\langle K, \partial \rangle$ be a differential field. A *Kolchin closed subset* of K^n is the set of solutions of a finite system of equations $p_1 = \cdots = p_m = 0$, where each $p_i \in K\{x_1, \ldots, x_n\}$. **Kolchin closed subset**

End of lecture 4, Apr 4

As with usual polynomials, we would like to remove some of the ambiguity by passing to ideals. Since the derivative of 0 is 0, we require it to be closed under the derivative:

Definition 2.4.5. A *differential ideal* in a differential ring $\langle A, \partial \rangle$ is an ideal $I \subseteq A$ such that $\partial(x) \in I$ for all $x \in I$. **differential ideal**

If $I \subseteq A$ is a differential ideal, A/I is a differential ring, uniquely determined by the property that the quotient map is a map of differential rings.

As in the algebraic case, differential polynomial equations for a language of a universal theory:

Definition 2.4.6. Let $\langle A, \partial \rangle$ be a differential ring. The language of differential fields over A consists of differential polynomial equations over A as basic (quantifier free) relations. The theory DF_A consists of the axioms for differential fields of characteristic 0 over A[3]

Where does one obtain examples of differential rings? We will soon see some constructions, but the natural source is geometry:

Example 2.4.7. The ring of smooth functions on the open interval $(0, 1)$ is a differential ring with the usual derivative (and similarly for other intervals). This example is far from being a field, it is not even a domain. □

[3]Strictly speaking, setting the characteristic is not necessary, but the theory behaves very differently in positive characteristic

Example 2.4.8. To obtain a differential integral domain, replace the real interval by a complex analytic domain X (open connected subset of $\mathbb{C}$), and consider holomorphic functions on X. The fraction field is the field of meromorphic functions on X. □

Both examples can be generalised by considering a smooth (or analytic) manifold X along with a vector field on it. We will return to this later. An easy special case:

Example 2.4.9. The field $\mathbb{C}(t)$ of rational functions, with the derivative determined by sending t to 1.

In fact, it is easy to see that for every $f \in \mathbb{C}(t)$, there is a unique derivation ∂ on $\mathbb{C}(t)$ determined by the property that $\partial(t) = f$. For example, we could set $t' = t$. Analytically, this would mean that t actually represents an exponential function $t = e^s$: algebraically, there is no difference between the coordinate t and the exponential function e^s, both are algebraically transcendental over $\mathbb{C}$. But differentially, they satisfy distinct equations. □

Example 2.4.10. Each ring can be viewed as a differential ring with the 0 derivation. This means, informally, that objects of algebraic geometry can be viewed as objects of the differential world. □

Exercise 2.4.11. Let A be a commutative ring, and let $A[\epsilon] = A[x]/x^2$ (we denote by ϵ the image of x), with $\pi : A[\epsilon] \to A$ the A-algebra map that sends ϵ to 0.

(1) If $\partial : A \to A$ is a derivation, the function $t_\partial : A \to A[\epsilon]$ given by $t(a) = a + \partial(a)\epsilon$ is a map of rings. Conversely, any map of rings $t : A \to A[\epsilon]$ such $\pi \circ t$ is the identity is of the form $t = t_\partial$ for a unique ∂ (we will later discuss the geometric meaning of this)
(2) If A is a differential domain, there is a unique derivation on $K(A)$ that extends the one on A (if you know what is an étale map, show that any derivation extends uniquely along them as well)
(3) The subset $C_A = \{c \in A | \partial(c) = 0\}$ is sub-ring of A, which is a subfield if A is a field. It is called the *subring of constants* of A

subring of constants

□

Suppose that $L \subseteq K$ is an extension of differential fields, and $a \in K$. What can be said of a from the point of view of L? The element a corresponds to an L-algebra map $L\{x\} \to K$, sending x to a, and its kernel is a prime differential ideal, the ideal of differential polynomials satisfied by a. Hence, we would like to understand prime differential ideals in $L\{x\}$.

Let us first recall the analogous situation in the algebraic setting. The polynomial algebra $k[x]$ in one variable is a *principal ideal domain*: Every (prime) ideal is generated by one element. The element generating an ideal I can be found as an element of minimal degree in I. In particular, it is a unique factorisation domain: every irreducible element is prime (recall that an element a is reducible if it is a product of two non-invertible elements, and prime if the ideal it generates is prime; it follows immediately that a prime element is irreducible, but in general the converse is false).

In the differential setting, the definition of "irreducible" remains the same, but the ideal in the definition of a being prime is replaced by the *differential* ideal $\langle a\rangle$ generated by a. With these definitions, the analogous facts to the above are false for $k\{x\}$: the polynomial $p(x) = (x'')^2 - 2x'$ is irreducible, but its derivative (which is in $\langle p\rangle$) is $2x''x^{(3)} - 2x'' = 2x''(x^{(3)} - 1)$, and neither of the factors is in $\langle p\rangle$, so the ideal is not prime.

We can still achieve a description that somewhat resembles the algebraic one. To each irreducible differential polynomial p, we will assign a prime ideal $I(p)$ containing it. We will also define a well-founded pre-order $\ll$, and will show that each prime differential ideal I is of the form $I(p)$ for $p \in I$ minimal with respect to $\ll$. We start with the definition of $\ll$:

Definition 2.4.12. The *order* $\mathrm{ord}(p)$ of non-zero differential polynomial p is the highest i such that $x^{(i)}$ appears in p. **order**

For any $p, q \in K\{x\}$, we say p is simpler than q, writing $p \ll q$, if $\mathrm{ord}(p) < \mathrm{ord}(q)$, or $\mathrm{ord}(p) = \mathrm{ord}(q) = n$, and $\deg_{x^{(n)}}(p) \leq \deg_{x^{(n)}}(q)$.

To proceed, we make the following computation, that will be used (and reinterpreted) also later: suppose that $p = p(t) = \sum_{i=0}^{m} a_i t^i$ is a (regular) polynomial over a differential ring $\langle A, \partial\rangle$, and $b \in A$. Then

$$\partial(p(b)) = \partial(\sum_{i=0}^{m} a_i b^i) = \sum_{i=0}^{m} \partial(a_i) b^i + (\sum_{i=1}^{m} i a_i b^{i-1})\partial(b) =$$
$$p^{\partial}(b) + \frac{\partial p}{\partial t}(b) \cdot \partial(b) \quad (2.1)$$

where p^{∂} is the polynomial obtained from p by applying ∂ to the coefficients. Assuming now that $p(x)$ is a differential polynomial of order n, we view it as a regular polynomial $p = p_0(x^{(n)})$ with p_0 over $k\{x\}$ (with coefficients of order at most $n-1$), and applying the above calculation we obtain

$$\partial(p) = {p_0}^{\partial}(x^{(n)}) + \frac{\partial p}{\partial x^{(n)}} \cdot x^{(n+1)} \quad (2.2)$$

the expression $\frac{\partial p}{\partial x^{(n)}}$ is called the *separant* of p, denoted s_p. In the example above, $\partial(p)$ was reducible because $s_p = 2x''$ divides the first summand (which was also $2x''$). Note that s_p is strictly simpler than p, so assuming p is minimal within a given prime ideal I, we expect the other factor to lie in I. Hence we define **separant**

$$I(p) = \{q \in k\{x\} | \exists m \in \mathbb{N}\ s_p^m q \in \langle p\rangle\} \quad (2.3)$$

This is often called the *saturation* of $\langle p\rangle$ with respect to s_p. It turns out that s_p accounts for all the non-primeness of $\langle p\rangle$:

Proposition 2.4.13. *For each irreducible $p \in k\{x\}$, $I(p)$ is a prime differential ideal containing p. Conversely, if I is a prime differential ideal in $k\{x\}$, then $I = I(p)$ for $p \in I$ minimal with respect to $\ll$.*

End of lecture 5, Apr 5

2.4.14. To prove Proposition 2.4.13, we mimic the algebraic case, with suitable modifications. We set $A = k\{x\}$, and let $A_i \subseteq A$ be the sub-ring of polynomials of order less than i (so $A_0 = k$). We fix $p \in A$ of order n (not necessarily irreducible), so $p(x) = a(x){x^{(n)}}^d + p_1(x)$, where $a = a(x)$ is of order lower than p, and p_1 is simpler than p. We denote by s the separant of p.

Let $B_i = A_{i+1}/A_i$. This is a module over A_i (and therefore over A_j, for $j < i$), and since $\partial(A_i) \subseteq A_{i+1}$, we have an induced map $\partial_i : B_i \to B_{i+1}$, which is A_i-linear. The polynomial p has a non-zero image $\bar{p}$ in B_n, and applying the above observations along with formula (2.2), we see that for each $m > 0$, $\partial^m(p) = sx^{(n+m)}$ up to lower order terms. Similar arguments allow us to work up to simpler terms.

We now have the following two variants of division with remainder:

Lemma 2.4.15. *With notation as in 2.4.14, for any $q \in A$ we have*

(1) *$l, m \in \mathbb{N}$ and $r \in A$ with $r \ll p$ and $a^l s^m q - r \in \langle p \rangle$*
(2) *$m \in \mathbb{N}$ and $r \in A$ with $\mathrm{ord}(r) \leq n$ and $s^m q - r \in \langle p \rangle$*

Proof. Let j be the order of q. If $j = n$, we may take $r = q$ and $m = 0$ for the second version. For the first, we may divide by a and use regular polynomial division (for polynomials in the variable $x^{(n)}$) to find r of lower degree so that $a^l q - r$ is divisible by p (in this case, $m = 0$).

If $j > n$, we remarked above that up to simpler terms, $\partial^{j-n}(p) = sx^{(j)}$. Writing up to simpler terms $q = b{x^{(j)}}^e$ for some b of order lower than j and $e \in \mathbb{N}$, we see that $q_1 = b(\partial^{j-n}p)^e - s^e q$ is simpler than q, so by induction, there are l_1, m_1, r_1 as in the statement (with $l_1 = 0$ for the second version), so that $a^{l_1} s^{m_1} q_1 - r_1 \in \langle p \rangle$. Since $b(\partial^{j-n}p)^e \in \langle p \rangle$, we are done. □

We next note that for q has order n, q could only be in $I(p)$ for the obvious reason:

Lemma 2.4.16. *Notation as in 2.4.14. If $q \in I(p)$ is of order at most n, then $s^m q \in (p)$ for some m. If p is irreducible, we may take $m = 0$.*

In particular, it follows that in this situation, the order of q is exactly n.

Proof. By assumption, $s^m q \in \langle p \rangle$ for some m, so $s^m q \in (p, \ldots, \partial^k(p))$ for some k, which we may take minimal. Assume $k > 0$. Localising at s, we have $q \in (p, p_1, \ldots, p_k)$, where $p_i = x^{(n+i)} + r_i$ with r_i of lower order. This is true in particular for $i = k$, and since $x^{(n+k)}$ does not appear in q, we may plug $x^{(n+k)} = -r_k$ (note that the differential structure no longer plays a role), so $q \in (p, \ldots, p_{k-1})$. This contradicts the minimality of k and proves the first statement.

For the second, note that (p) is prime (the polynomial ring is a UFD), and $s \notin (p)$ since the degree is too low. Since $s^m q \in (p)$ we must have $q \in (p)$. □

proof of 2.4.13. It is obvious that for all p, $I(p)$ is an ideal. To show that it is a differential ideal, let $g \in I(p)$, so that $s^m g \in \langle p \rangle$ for some m. Then $s^{m+1} g' = (s^{m+1} g)' - (m+1) s^m g$, and both terms on the right are in $\langle p \rangle$, so $g' \in I(p)$.

To show that $I(p)$ is prime when p is irreducible, assume $fg \in I(p)$, so that $s^m fg \in \langle p \rangle$ for some m. Applying long division 2.4.15 in the second version (possibly modifying m), we may assume that f, g have order at most n. But then p divides $s^m fg$ by Lemma 2.4.16. Since it is irreducible, it must divide one of the factors f, g (it cannot divide s, since it is simpler). Hence f or g are in $I(p)$.

Conversely, assume that I is a prime (differential) ideal, and let p be simplest non-zero in I (p need not be unique). Note that p is irreducible, since a factor would be in I and simpler. If $g \in I(p)$, then $s^m g \in \langle p \rangle \subseteq I$ for some m. Since I is prime, s^m or g is in I. It cannot be s^m since it is simpler than p. This shows that $I(p) \subseteq I$.

Assume that $q \in I$. Applying long division, we find $a^l s^m q - r \in \langle p \rangle \subseteq I$ for some l, m and some r simpler than p, so that $r \in I$ as well. Since p was assumed to be simplest non-zero in I, we have $r = 0$. Hence $a^l s^m q \in \langle p \rangle$, so $a^l q \in I(p)$. We already saw that $I(p)$ is prime, so $a^l \in I(p)$ or $q \in I(p)$. The former is impossible by simplicity. □

The direct analogue of Hilbert's basis theorem fails: there is no ascending chain condition on differential ideals in $k\{x\}$. For example, the sequence $I_0 = 0$ and $I_{k+1} = \langle I_k \cup \{(x^{(k)})^2\} \rangle$ is a strictly increasing chain (this is not trivial). However, we are interested in the geometry of the sets determined by these ideals, and we have, for example, $Z(I_1) = Z(x)$. Recall that the *radical* of an ideal I is the ideal $\sqrt{I} = \{b | \exists k \in \mathbb{N}\ b^k \in I\}$. It is clear that $Z(I) = Z(\sqrt{I})$. Furthermore, **radical**

Exercise 2.4.17. If I is a differential ideal (in any differential ring), then $\sqrt{I}$ is also a differential ideal. □

An ideal I is called a *radical ideal* if $I = \sqrt{I}$ (this is an algebraic notion). A differential ideal I is *well-mixed* if $ab \in I$ implies $ab' \in I$. Of course, if I is prime then it is well-mixed, but also **radical ideal** **well-mixed**

Exercise 2.4.18. Any radical differential ideal is well-mixed □

It turns out that the analogue of the Hilbert basis theorem is true for *radical* differential ideals. Let us say that a differential ring A is *differentially Noetherian* if every strictly ascending chain of radical differential ideals stabilizes. **differentially Noetherian**

Theorem 2.4.19 (Ritt–Raudenbush basis theorem)**.** *If A is a differentially Noetherian differential ring (in characteristic* 0*), then so is $A\{x\}$.*

We skip the proof, see [9] or [6]. We recall again that geometrically, this means that every strictly descending chain of Kolchin closed subsets stabilises, and that every Kolchin closed subset can be given by a *finite* number of equations.

End of lecture 6, Apr 11

3. The theory of differentially closed fields

3.1. **Quantifier Elimination.**

We saw above that we have some understanding of *quantifier free* definable subsets in fields, and even in differential fields. However, model theoretic notions are normally described via the collection of *all* definable sets, including quantifiers. These could be vastly more complicated:

Example 3.1.1. Consider the field $\mathbb{Q}$ of rational numbers. Julia Robinson showed (in her thesis, see [10] for example) that the subset of integers is definable: there is a formula $\phi(x)$ such that $\phi(q)$ holds for a rational q if and only if q is an integer (of course, ϕ must have quantifiers, but it turns out that just a few quantifiers suffice). Once this is known, the definable sets are essentially the same as the ones definable in $\mathbb{Z}$, the structure studied by Gödel in his incompleteness theorem. The proof of this theorem shows that this structure is extremely rich and "unmanageable". In particular, there is a definable bijection between any two Cartesian powers, so no reasonable theory of dimension for definable sets can exist. □

constructible

The situation is rather different with the field $\mathbb{C}$ of complex numbers. We will call a subset of $\mathbb{C}^n$ *constructible* if it can be defined by a quantifier-free formula (in the language of fields). In other words, it is a finite boolean combination of Zariski closed subsets. We will prove below:

Proposition 3.1.2. *Every definable subset of* $\mathbb{C}^n$ *is constructible*

A similar kind of statement is usually required to have any hope of understanding model-theoretically a structure, so there is some machinery in place to prove them. First, we note that this statement is really about the *theory* $\mathcal{T}$ of $\mathbb{C}$: for every formula $\phi(x)$ there is a quantifier free formula ψ with $\forall x(\phi \leftrightarrow \psi) \in \mathcal{T}$. This has a name:

quantifier elimination

Definition 3.1.3. A theory $\mathcal{T}$ admit *quantifier elimination* if for every formula $\phi(x)$ there is a quantifier free formula ψ with $\forall x(\phi \leftrightarrow \psi) \in \mathcal{T}$.

One tool to prove quantifier elimination is the following:

Proposition 3.1.4. *Let* $\mathcal{T}$ *be a theory, and let* $\phi(x)$ *be formula. Then* ϕ *is equivalent (with respect to* $\mathcal{T}$*) to a quantifier-free formula if and only if for all models* M *of* $\mathcal{T}$ *and any* $a \in \phi(M)$*,* $\mathcal{T}_a$ *implies* $\phi(a)$

Proof. We may assume ϕ is consistent with $\mathcal{T}$. If ϕ is equivalent to a quantifier-free formula, we may assume it itself is already quantifier-free, and then $\phi(a) \in \mathcal{T}_a$.

Conversely, let Γ be the set of quantifier-free formulas implied by ϕ. We claim that Γ implies ϕ (with respect to $\mathcal{T}$). If not, let N be a model with an element $a \in N$ such that $a \in \neg\phi(N)$ but $a \in \psi(M)$ hold for all $\psi \in \Gamma$.

We claim that $\mathcal{T}_a$ is consistent with $\phi(a)$: otherwise, $\Gamma \cup \{\phi\}$ is inconsistent with $\mathcal{T}$, so by compactness, ϕ is inconsistent with some $\psi \in \Gamma$. But then ϕ implies both ψ and $\neg\psi$, contradicting its consistency.

Since $\mathcal{T}_a$ is consistent with $\phi(a)$, there is a model M of $\mathcal{T}_a$ for which $a \in \phi(M)$. By assumption, this means that $\mathcal{T}_a$ implies $\phi(a)$, contradicting the existence of N.

We proved that Γ implies ϕ. Again by compactness, some $\psi \in \Gamma$ implies it. Since by definition $\phi \to \psi$, we are done. $\square$

The above result provides a way of checking that a particular formula is quantifier-free, but for an arbitrary formula it might still be difficult to check this condition. If we wish to prove that *all* formulas are equivalent to quantifier-free ones, we have the following observation:

Exercise 3.1.5. If $\mathcal{T}$ is a theory, and for each quantifier-free formula $\phi(x, y)$, where y is one variable, the formula $\exists y(\phi(x, y))$ is equivalent to a quantifier-free one, then $\mathcal{T}$ admits quantifier-elimination. $\square$

We may now prove the statement we started with, in slightly greater generality:

Proposition 3.1.6. *Let K be an algebraically closed field, and let $\mathcal{T}$ be its theory (in the language of fields). Then $\mathcal{T}$ admit quantifier-elimination.*

Proof. By the last exercise, it suffice to show that each formula of the form $\exists y\phi(x, y)$, where ϕ is quantifier-free, is equivalent to a quantifier-free one. By the criterion 3.1.4, we need to show that if L is a model of $\mathcal{T}$ and a is a tuple in L for which $\phi(a, y)$ has a solution in L, then $\phi(a, y)$ also has a solution in any other field E satisfying $\mathcal{T}$ and containing a.

Note that $\phi(a, y)$ is boolean combination of polynomial equations in y (a single variable), with coefficients in the subfield L_0 generated by a. In fact, we may assume that $\phi(a, y)$ has the form $p_1 = 0 \wedge \cdots \wedge p_k = 0 \wedge q \neq 0$ for some polynomials p_i and q. Since $L_0[x]$ is a pid, this system is implied by one equation $p = 0$ with p a non-unit polynomial over L_0, $p_a(x) = x^n + \sum a_i x^i$. Since K is algebraically closed, $\mathcal{T}$ includes the sentence $\forall y \exists x(p_y(x) = 0)$, so it is true in E. Speciallising to $y = a$, we see that p_a has a solution in E. $\square$

Tracing through the proof, we see explicitly that we proved the follow stronger fact:

Corollary 3.1.7. *There is a first order theory ACF in the language of fields, whose models are precisely the algebraically closed fields. This theory admits quantifier-elimination.*

Since $\mathbb{C}$ is an algebraically closed field, we have $ACF \subseteq \mathcal{Th}(\mathbb{C})$. Is there anything else we can say in a first order manner about $\mathbb{C}$? The following result is often obtained via the Löwenheim–Skolem theorems, using the uniqueness of models in a fixed uncountable cardinality. For each ideal I in $\mathbb{Z}$, write $ACF_I = ACF \cup \{\underline{n} = 0 | n \in I\} \cup \underline{m} \neq 0 m \notin I$, where $\underline{n} = 1 + \dots + 1$ ($|n|$ times). We write ACF_p in place of $ACF_{(p)}$.

Corollary 3.1.8. *The theories ACF_p are complete when p is prime. In particular, $ACF_0 = \mathcal{Th}(\mathbb{C})$.*

Proof. Let ψ be a sentence. According the quantifier-elimination in ACF, it is equivalent there to a quantifier-free one, so we may assume it is quantifier-free. Every such sentence is a statement about the field structure of the prime field, which is uniquely determined by p. □

We next would like to finish describing the relation between algebra and geometry for affine varieties. To do that, we use the following definition, which will also be important later:

existentially closed

Definition 3.1.9. Let M be a model of a theory $\mathcal{T}$. We say that M is *existentially closed* if any quantifier-free formula $\phi(x)$ in $\mathcal{T}_M$ that is consistent with $\mathcal{T}_M$ has a point $m \in \phi(M)$.

Exercise 3.1.10. In the definition, it doesn't matter if x is a variable or a tuple of variables. □

Exercise 3.1.11. If $\mathcal{T}$ admit quantifier-elimination, then every model of $\mathcal{T}$ is existentially closed. □

For example, $\mathbb{C}$ is existentially closed (with respect to the theory of algebraically closed fields, or of fields). In the case of one variable, it says that if a polynomial has a root in some field extension, then it has a root in $\mathbb{C}$. In the case of several variables, it can be thought of as a form of Hilbert's Nullstellensatz:

Corollary 3.1.12. *Let K be an algebraically closed field, I a prime ideal in $A = K[x_1, \dots, x_n]$ and $g \in A \setminus I$. Then $Z(I) \setminus Z(g) \subseteq K^n$ is non-empty.*

Proof. Let L be the fraction field field of the domain A/I. The canonical map $p : A \to L$ determines a point of $Z(I)$ as in Proposition 2.3.10, and since the kernel of p is precisely I, this point is not in $Z(g)$.

Let $p_1, \dots, p_k$ be generators for I. We just showed that $p_1 = 0 \wedge \dots \wedge p_k = 0 \wedge g \neq 0$ is satisfied in an extension of K. Since K is existentially closed, it is satisfied in K. □

A standard way to state the Nullstellensatz is as follows:

Corollary 3.1.13 (Hilbert's Nullstellensatz)**.** *If K is algebraically closed, and I, J are two distinct radical ideals in $K[x_1, \dots, x_n]$, then $Z(I) \neq Z(J)$. For any ideal I, $I(Z(I)) = \sqrt{I}$. Any finitely generated reduced K algebra is affine (i.e., the ring of functions on an affine variety).*

We outline the argument: We may assume there is $g \in J \setminus I$. Any radical ideal is the intersection of the prime ideals that contain it (this is essentially the algebraic counterpart of irreducible components), so there is a prime ideal I_1 containing I and not g. We have $Z(I_1) \subseteq Z(I)$ and $Z(J) \subseteq Z(g)$, so to show that $Z(I)$ is different from $Z(J)$, it suffices to find a point in $Z(I_1)$ which is not in $Z(g)$. This is provided by 3.1.12. The second statement follows since both sides are radical ideals that define the same Zariski closed subset, and the last statement is the translation via Proposition 2.3.10. End of lecture 7, Apr 12

3.2. **Model companions.** We saw above that the theory of algebraically closed fields enjoys good model-theoretic properties, not shared by other fields. However, we still would like study fields (or integral domains) in general. We note that the theory D is precisely the universal part of ACF: This follows from Prop. 2.2.6, since integral domains are precisely the substructures of (algebraically closed) fields. In general, many of the natural theories we would like to study are universal, and we would like to do that by looking for interesting theories with the given universal part. Thus, it makes sense to make the following definition.

Definition 3.2.1. The theories $\mathcal{T}_1$ and $\mathcal{T}_2$ are *companions* if they have the same universal part. **companions**

In particular, each theory has a unique universal companion, its universal part. Among all companions of a given theory, we would like to find one with better properties. In light of the example of D and ACF, we might aim for the following:

Definition 3.2.2. A *model completion* of a theory $\mathcal{T}_0$ is a companion theory $\mathcal{T}$ that admits quantifier elimination. **model completion**

Since the notion of a companion depends only on the universal part, we will normally assume that $\mathcal{T}_0$ is universal. Combining the definition with Prop. 2.2.6, we see that $\mathcal{T}$ is a model completion of a universal theory $\mathcal{T}_0$ if and only if:

(1) Every model of $\mathcal{T}$ is a model of $\mathcal{T}_0$
(2) Every model of $\mathcal{T}_0$ can be extended to a model of $\mathcal{T}$
(3) $\mathcal{T}$ admits quantifier elimination

We have seen that ACF is a model completion of D, but at the moment, we do not know that it is determined uniquely by this property.

It turns out that it is more natural and more convenient to consider a slightly weaker condition than quantifier elimination. To explain it, consider a formula $\phi(x)$, and an inclusion of models $M \subseteq N$. The two models determine sets $\phi(M)$ and $\phi(N)$, but in general we know little about the relation between them: an element $m \in M$ could be in $\phi(M)$ but not in $\phi(N)$ or vice versa. The situation is better if ϕ is quantifier free: in this case, by definition, $\phi(M) = \phi(N) \cap M$. For the inclusion $M \subseteq N$ as models of $\mathcal{T}$ to be meaningful, we would like the same to hold arbitrary formulas:

Definition 3.2.3. An inclusion $M \subseteq N$ of structures is an *elementary inclusion*, denoted $M \preccurlyeq N$, if for every formula ϕ we have $\phi(M) = \phi(N) \cap M$. **elementary inclusion**

elementary embedding Similarly, an *elementary embedding* is a homomorphism $f : M \to N$ such that $f(M) \preccurlyeq N$.

Applying the definition to the case where ϕ is a sentence, we see that each structure has the same theory as any of its elementary substructures. The following observation will be useful:

Proposition 3.2.4. *Let M be a structure, and let $\mathcal{T}_{(M)}$ be the theory in the language with constants for M, consisting of all sentences that hold in M. Then the elementary extensions of M are precisely the models of $\mathcal{T}_{(M)}$. In particular, every infinite structure has a proper elementary extension.*

The proof consists of expanding the definitions, and is left as an exercise.

As observed above, the condition on ϕ in the definition is satisfied automatically for quantifier-free formulas ϕ. In particular, if $\mathcal{T}$ admits quantifier-elimination, then every inclusion of its models is elementary. A similar observation applies for existential formulas: If $\phi(x)$ is existential and $M \subseteq N$, then $\phi(M) \subseteq \phi(N) \cap M$. In this case, we also have the converse:

Proposition 3.2.5. *Let $\mathcal{T}$ be a theory and ϕ a formula. The following are equivalent:*

(1) *ϕ is (equivalent to) an existential formula*
(2) *For each inclusion $M \subseteq N$ of models of $\mathcal{T}$, $\phi(M) \subseteq \phi(N)$*

The proof is similar to that of Prop. 3.1.4:

Proof. Let $\Gamma = \{\psi | \mathcal{T} \models \psi \to \phi, \psi \text{ is existential}\}$ If ϕ is not existential, we may (by compactness) find a model M and an element $m \in M$ satisfying ϕ but no element of Γ.

We claim that $\mathcal{T}_1 = \mathcal{T}_M \cup \{\neg\phi(m)\}$ is consistent. Otherwise, there is a quantifier-free formula $\theta(x, y)$ and an element $m' \in M$, such that $\theta(m, m')$ holds in M and is inconsistent with $\neg\phi(m)$. Therefore, $\mathcal{T} \models \theta(m, m') \to \phi(m)$, hence also $\mathcal{T} \models \exists y\theta(m, y) \to \phi(m)$, so that $\psi = \exists y\theta \in \Gamma$. But $\psi(m)$ holds in M, contradicting the choice of m. This proves that $\mathcal{T}_1$ is consistent.

Let N be a model of $\mathcal{T}_1$. Then N is a model of $\mathcal{T}$ containing M. By assumption, $m \in \phi(N)$, a contradiction. □

Applying the proposition in a contrapositive manner, we obtain:

Corollary 3.2.6. *Let $\mathcal{T}$ be a theory and ϕ a formula. The following are equivalent:*

(1) *ϕ is (equivalent to) a universal formula*
(2) *For each inclusion $M \subseteq N$ of models of $\mathcal{T}$, $\phi(N) \cap M \subseteq \phi(M)$*

Hence, $\phi(M) = \phi(N) \cap M$ for all inclusions $M \subseteq N$ of models of $\mathcal{T}$ if and only if ϕ is both existential and universal.

Of course, if a formula is quantifier-free, it is both universal and existential, and in this case we recover the observation we began with. However, as the following example shows, this is not an equivalence:

Example 3.2.7. In the theory $\mathcal{T}$ of the real field $\mathbb{R}$, the formula $\exists y(y^2 = x)$ is clearly existential, but the same set is defined by the universal formula $x = 0 \vee \forall y(y^2 \neq -x)$. □

We are now ready to define our main condition.

Proposition 3.2.8. *The following conditions on the theory $\mathcal{T}$ are equivalent.*

(1) *A model of $\mathcal{T}_\forall$ is a model of $\mathcal{T}$ if and only if it is existentially closed (i.e., $\mathcal{T}$ is precisely the theory of existentially closed models of $\mathcal{T}_\forall$)*
(2) *Every model of $\mathcal{T}$ is existentially closed*
(3) *Every formula is equivalent with respect to $\mathcal{T}$ to an existential one*
(4) *For every model M of $\mathcal{T}$, the theory $\mathcal{T}_M$ is complete*
(5) *Every embedding of models of $\mathcal{T}$ is elementary*

Definition 3.2.9. A theory satisfying the equivalent conditions of Prop. 3.2.8 is said to be *model complete*. If $\mathcal{T}'$ is a companion of $\mathcal{T}$, $\mathcal{T}$ is called the *model companion* of $\mathcal{T}'$. **model complete** **model companion**

The word "the" in the last sentence can be justified by the first item in the Proposition: it implies, in particular, that a theory has at most one model companion.

Remark 3.2.10. It follows from the remarks above that a theory with quantifier-elimination is model complete. It turns out ([7]) that the theory of the real field is also model complete, but as above, it does not admit quantifier-elimination. In fact, it turns out that algebraically closed fields are the only infinite fields with quantifier-elimination, but there is a number of interesting theories of fields that are model complete. □

For the proof of Prop. 3.2.8, most of the implications are straightforward: (2) follows tautologically from (1). To prove (3), it suffices to show that each universal formula is existential. If $\phi = \forall x\psi(x, y)$ is such a formula, it suffices to show (by Prop. 3.2.5) that for each $m \in M$ and any extension $M \subseteq N$, $m \in \phi(N)$. If this is not the case, there is $n \in N$ satisfying $\neg\psi(n, m)$. Since M is existentially closed, such an n also exists in M, contradicting the assumption $m \in \phi(M)$.

The equivalence of (3) with (4) is similar to Prop. 3.1.4, and (5) follows from (3) since each formula is also universal in this case. Also, (5) trivially implies (2), so it only remains to show that every existentially closed model of $\mathcal{T}_\forall$ is a model of $\mathcal{T}$ (assuming (5)).

Let M be an existentially closed model of $\mathcal{T}_\forall$. We first note that M is a model of $\mathcal{T}_{\forall\exists}$, the subset of $\mathcal{T}$ consisting of sentences (implied by sentences) of the form $\forall x\exists y\phi(x, y)$. Indeed, given such an element of $\mathcal{T}$, let N be a model of $\mathcal{T}$ containing M, and let $m \in M$. We need to show that $M \models \exists y\phi(m, y)$. Since $m \in N$ and the sentence holds in N, there is $n \in N$ satisfying $\phi(m, n)$. Since M is existentially closed, there is such n in M as well. Thus, to complete the proof, it suffices to show that, given (5), we have $\mathcal{T} = \mathcal{T}_{\forall\exists}$. This is attained by the following.

Proposition 3.2.11. *Let $\mathcal{T}$ be a theory. The following are equivalent:*

(1) $\mathcal{T} = \mathcal{T}_{\forall\exists}$ *(up to equivalence)*
(2) *The class of models of $\mathcal{T}$ is closed under unions of chains: Given a non-empty collection $\mathcal{C}$ of models of $\mathcal{T}$ ordered by inclusion, the union $\bigcup\mathcal{C}$ is also a model (the union is viewed as a structure with the unique interpretation that makes the inclusions of elements of $\mathcal{C}$ in it homomorphisms)*

Example 3.2.12. Let $\mathcal{T}$ be the theory of $\mathbb{Z}$ as an ordered set. Then $\mathcal{T}$ includes the axiom $\forall x\exists y\forall z(x < y \wedge y \leq z \vee z \leq x)$ stating that every element has a successor. Can this axiom be replaced by some $\forall\exists$ axioms? According to the proposition, the answer is *no*: The structures $M_i = \frac{1}{2^i}\mathbb{Z}$ form a chain of models of $\mathcal{T}$ (each of them is isomorphic to $\mathbb{Z}$), but the union fails to satisfy the above axiom, so is not a model of $\mathcal{T}$. □

End of lecture 8, Apr 25

Given this last proposition, we may finish as follows:

Proof of 3.2.8. As explained above, it remains to prove that under the assumption (5) of 3.2.8, we have $\mathcal{T} = \mathcal{T}_{\forall\exists}$. To do this, it is enough to prove that the equivalent condition in Prop. 3.2.11 holds. Given a chain $\mathcal{C}$ of models of $\mathcal{T}$, assumption (5) implies that the chain is elementary. Hence, we are done by the following general result. □

Proposition 3.2.13. *The union of a chain of elementary embeddings is an elementary extension of each element of the chain.*

Proof. Let $\mathcal{C}$ be a non-empty set of structures, linearly ordered by inclusion, and assume all embeddings in $\mathcal{C}$ are elementary. Let N be the union, and let $M \in \mathcal{C}$. We show that the condition for elementary inclusion of M in N is satisfied, by induction on the construction of the formula.

The condition is trivial for quantifier-free formulas and closed under boolean operations, so we assume that for some $m \in M$, $\exists y\phi(m, y)$ holds in N and show it holds in M. Let $n \in N$ be such that $N \models \phi(m, n)$. Then there is $M_1 \in \mathcal{C}$ containing M and n. By induction, $\phi(m, n)$ holds in M_1, so $\exists y\phi(m, y)$ holds there. Since the embedding of M in M_1 is elementary, it also holds in M. □

It remains to prove Prop. 3.2.11. The main step is

Lemma 3.2.14. *Let M be a model of $\mathcal{T}_{\forall\exists}$. Then there are $M \subseteq N \subseteq M_1$ where N is a model of $\mathcal{T}$ and $M \subseteq M_1$ elementary.*

Proof. Let $\mathcal{T}' = \mathcal{T}_{(M)}$, as in Prop. 3.2.4, and let $\mathcal{T}_1 = \mathcal{T} \cup \mathcal{T}'_{\forall}$. A model N of $\mathcal{T}_1$ is a model of $\mathcal{T}$, but also satisfies $\mathcal{T}'_{\forall}$, so contains M, and can be embedded in a model M_1 of $\mathcal{T}'$, which is an elementary extension of M according to Prop. 3.2.4. Hence, it suffices to show that $\mathcal{T}_1$ is consistent.

If not, $\mathcal{T}$ implies $\neg\phi(m)$ for some universal ϕ. Since m does not appear in the language of $\mathcal{T}$, it follows that $\mathcal{T} \models \forall y\neg\phi(y)$. This is a sentence in $\mathcal{T}_{\forall\exists}$, contradicting the choice of M. □

Exercise 3.2.15. Use the lemma to prove Prop. 3.2.11 □

Remark 3.2.16. It follows from Prop. 3.2.11 that if a universal theory $\mathcal{T}_0$ has a model companion $\mathcal{T}$, then $\mathcal{T} = \mathcal{T}_{\forall\exists}$. It turns out that any universal theory has unique biggest $\forall\exists$ companion $\mathcal{T}$. It follows that if a model companion for $\mathcal{T}_0$ exists, then it must be this $\mathcal{T}$, and this happens precisely if all models of $\mathcal{T}$ are existentially closed. □

3.3. **Differentially closed fields.** In this section we prove that the (universal) theory DD of differential domains of characteristic 0 admits a model completion. The following axiomatisation is due to Lenore Blum:

Definition 3.3.1. The theory DCF asserts that a model K is a differential field, and for any non-zero differential polynomials $f(x), g(x)$ in one variable over K, with $\mathrm{ord}(f) > \mathrm{ord}(g)$, there is $a \in K$ with $f(a) = 0$ and $g(a) \neq 0$. A model of DCF is called a *differentially closed field*.

differentially closed field

This is clearly first order. We note, applying the definition in the case $\mathrm{ord}(f) = 0$, that any such differential field is algebraically closed. In particular, none of the "geometric" examples is a model. In fact, it might not be clear that the theory has a model at all, so we first prove:

Proposition 3.3.2. *Every differential domain embeds in a differentially closed field*

Proof. The fraction field of a differential domain is (uniquely) a differential field, so we start with a differential field k. Take differential polynomials f, g over k as in the definition. At least one of the factors of f will have the same order as f, so we may assume that f is irreducible. We saw in Prop. 2.4.13 that $I(f)$ is a prime differential ideal containing f, and since $\mathrm{ord}(g) < \mathrm{ord}(f)$, $g \notin I(f)$ by Lemma 2.4.16. Dividing by $I(f)$ we thus obtain a differential field extension L of k where the problem with the given f, g has a solution.

We now proceed as in the construction of an algebraic closure: fixing a well-order on the set of all pairs (f, g) over k, and proceeding by it, we obtain a differential field k_1 where all instances of the problem for pairs of polynomials over k can be solved. Repeating this for k_1 we again obtain a chain, whose union is a model of DCF containing k. □

This shows that DCF is consistent, and also that DCF is a companion of DD. We aim for:

Theorem 3.3.3. *The theory DCF eliminates quantifiers. Hence, it is the model completion of DD.*

End of lecture 9, Apr 26

Before proving the theorem, it is convenient to introduce an important notion that will be used again later. We say that a set is *κ-small*, for a cardinal κ, if its cardinality is less than κ.

κ-small

Definition 3.3.4. A set of formulas Γ is said to be *finitely satisfiable* in the

finitely satisfiable

structure M if every finite subset of Γ is satisfied in M.

κ-saturated For a cardinal κ, A structure M is *κ-saturated* if for each κ-small subset $A \subseteq M$, every finitely satisfiable set of formulas Γ over A is satisfied in M.

Example 3.3.5. Let E be a κ-saturated differentially closed field, and let k be a κ-small differential subfield. Let $f(x)$ be a differential polynomial of order n over k. Then there is an element $c \in E$ such that $f(c) = 0$, and $g(c) \neq 0$ for *all* differential polynomials g of order less than n over k. □

We will need the following basic fact:

Proposition 3.3.6. *Let M be a structure. For every cardinal κ, M has an elementary extension which is κ-saturated.*

We only sketch the proof: Fixing a finitely satisfiable Γ over a subset $A \subseteq M$ of cardinality less than Γ, we may use Prop. 3.2.4 to find an elementary extension of M of the same cardinality as M where Γ is satisfied. We proceed in this way for all subsets A and Γ, creating a chain of elementary extensions. The cardinality bound shows that the chain will not be "too long", and Prop. 3.2.13 shows that we remain with an elementary extension when we take unions.

We now go back to the proof of the theorem.

Proof of Theorem 3.3.3. As in the proof of Prop. 3.1.6, we need to show that if L is a differentially closed field, $a \in L$ a tuple, and $\phi(x, y)$ is a quantifier-free formula with y a single variable, such that $\phi(a, y)$ is satisfied in L, then it is also satisfied in any other differentially closed field E containing the differential field k generated by a (note that k is countable). Since the satisfaction of ϕ is a first-order statement, we may replace E by an elementary extension, and so we may assume that E is κ-saturated for some uncountable κ.

Again as in 3.1.6, this reduces to the case when ϕ is of the form $p_1 = 0 \wedge \ldots \wedge p_r = 0 \wedge q \neq 0$, where this time, p_i and q are differential polynomials over k. Assume b satisfies $\phi(a, y)$ in L, and let I be the ideal of elements $k\{x\}$ that vanish on b. Then I is a prime differential ideal, so by Prop. 2.4.13, we have $I = I(f)$ for some irreducible f of some order n and degree d. By Example 3.3.5, we may find $c \in E$ such that $f(c) = 0$ and $g(c) \neq 0$ for each differential polynomial g over k of lesser order.

We claim that $g(c) \neq 0$ for g of order n and degree less than d. For that, we note that the choice of c implies that $c, c', \ldots, c^{(n-1)}$ are algebraically independent. It follows that as a (regular) polynomial of $c^{(n)}$, f is irreducible.

We now showed that f is the simplest differential polynomial satisfying $f(c) = 0$. In particular, c is not a zero of the separant s of f. By Lemma 2.4.15 (first version), there are number l, m and a differential polynomial g simpler than f, such that $a^l(c)s^m(c)q(c) - g(c) = 0$. Since $g(c) \neq 0$, $q(c) \neq 0$ as well. Likewise, for each j there is i such that $s^i(c)p_j(c) = 0$, so we must have $p_j(c) = 0$. Thus c satisfies ϕ. □

Corollary 3.3.7. *DCF is complete*

Proof. Every model of DCF has $\mathbb{Q}$ as a substructure, and the validity of each quantifier-free sentence is determined there. By quantifier elimination, this determines the whole theory. □

End of lecture 10, May 2

We now would like to provide another, more geometric, set of axioms for DCF. One advantage of it will be that it refers only to order one equations. Recall that to axiomatise the model companion, we need to effectively decide what systems of equations are solvable (in a potential extension differential field).

Example 3.3.8. The system of equations $x^2+y^2=1, x'=x, y'=y$ admits no solution, in any differential field (or ring). Indeed, if $\langle u,v\rangle$ solve the algebraic equation, differentiating it reveals the differential identity $uu'+vv'=0$. If, in addition, $u'=u$ and $v'=v$, we obtain $u^2+v^2=0$, incompatible with the original equation. □

We note that the equation $uu'+vv'=0$ in the above example can be expressed geometrically as saying that the vector $\langle u',v'\rangle$ is orthogonal to the vector $\langle u,v\rangle$ and thus is tangent to the variety given by $x^2+y^2=1$. We would like to explain that this is the situation in general. To do that, we review the notion of a tangent bundle.

3.3.9. *Tangent bundles.* Let B be an open ball around 0 in $\mathbb{R}^n$. Let A be the algebra of smooth functions on B. Each vector v in $\mathbb{R}^n$ (which we think of as a tangent vector at 0) determines an operation $\partial_v : A \to \mathbb{R}$, the directional derivative in the direction v. The operation is a derivation at 0: it is $\mathbb{R}$-linear, and for all $f,g\in A$, $\partial_v(fg)=\partial_v(f)g(0)+f(0)\partial_v(g)$. If B_1 is smaller ball around 0, and A_1 the corresponding algebra of smooth functions, this operation is compatible with the restriction. Furthermore, any operation satisfying these properties has the form ∂_v for a unique v (this requires an argument. For example, if $n=1$, every smooth function f satisfying $f(0)=0$ is of the form $xg(x)$, for a smooth function g). Hence, we may canonically identify the space of tangent vectors at 0 with the set of such derivations.

If M is a smooth manifold and $a\in M$, we may try to define a tangent vector to a as represented by a smooth curve $\gamma:(-u,u)\to M$ with $\gamma(0)=a$. Each such curve determines a derivation $\partial_\gamma : A\to\mathbb{R}$ at a, given by $\partial_\gamma(f)=(f\circ\gamma)'(0)$. Using a chart at a, we may translate the situation to an open ball around 0, where this derivation is identified with the direction derivative ∂_v in the direction $v=\gamma'(0)$, and the associated derivation the same as above. Thus, we may define the tangent space at a to be the space of derivations $\partial : A\to\mathbb{R}$ at a (note that any smooth function in a neighbourhood of a extends to a smooth function on A).

We would like to repeat same definition algebraically. We recall that an affine algebraic variety over k was given, essentially, by a (finitely generated, reduced) k-algebra A, and the points of the space associated to A were k-algebra maps $x : A \to k$.

Definition 3.3.10. Let A be a k-algebra. A tangent vector at the point $x : A \to k$ (to the affine variety X corresponding to A) is a k-linear map $\partial : A \to k$ which is a derivation with respect to x: $\partial(ab) = \partial(a)x(b) + x(a)\partial(b)$. The collection of tangent vectors at x forms a vector space over k, called the *tangent space* at x, denoted $T_x(X)$.

tangent space

Given an affine variety X with algebra A, we would like to produce an affine variety TX whose points consist of pairs $\langle x, v\rangle$, where x is a point of X, and v is a tangent vector at x. To do that, it is convenient to recall the description of tangent vectors given in Exercise 2.4.11: A tangent vector $\partial : A \to k$ at $x : A \to k$ determines a map $v : A \to k[\epsilon] = k[t]/t^2$, whose composition with $\epsilon \mapsto 0$ is x. Conversely, every such map determines a tangent vector, obtained by reading the coefficient of ϵ. Thus, the set of all tangent vectors of X can be identified with the set $X(k[\epsilon]) = Hom_{k-alg}(A, k[\epsilon])$, and we are looking for an affine variety TX with an identification $TX(k) = X(k[\epsilon])$.

More generally, we may consider the set $X(B)$ of B-valued points of X for some k-algebra B, i.e., the set of maps $A \to B$. We had seen that if B is the algebra of functions of some affine variety Y, this set corresponds to the set of (algebraic) maps from Y to X. The set $X(B[\epsilon])$ then corresponds to a family of tangent vectors to X, parametrised by Y. Such a family should correspond to a family of points of TX, i.e., to an element $TX(B)$. So we would like the equality $TX(B) = X(B[\epsilon])$ to hold for all k-algebras B. It turns out that this condition suffices to determine TX uniquely.

In fact, this whole discussion applies equally well when A is an arbitrary k-algebra, not necessarily related to an affine variety, except that the geometric intuition appears to be lost. However, all our geometric notions are defined algebraically, so may apply them to general such algebras. We will thus imagine a "space" $\mathrm{spec}(A)$ attached to an arbitrary k-algebra A, called *affine scheme* associated to A (in fact, A could be any commutative ring, but it will be more convenient to restrict to the case of algebras over a field k). When A is reduced and of finite type over a field k, this will coincide with the previous interpretation.

affine scheme

Definition 3.3.11. Let $X = \mathrm{spec}(A)$ be the affine scheme associated to a k-algebra A. The *tangent bundle* to X is an affine scheme TX with an identification $X(B[\epsilon]) \xrightarrow{\sim} TX(B)$, for all k-algebras B, functorial in B (i.e., commuting with maps $f : B \to B'$)

tangent bundle

In algebraic terms, $TX = \mathrm{spec}(TA)$, where TA is a k-algebra along with a bijection $Hom(TA, C) \xrightarrow{\sim} Hom(A, C[\epsilon])$, for each k-algebra C, functorially in C. We will prove the existence of such a TA below. The fact that it is determined uniquely by the defining condition is a special case of the Yoneda Lemma: apply the bijection to the identity map, in the case $C = TA$.

If $Y \subseteq X$ is a closed subscheme (corresponding, by definition, to a surjective k-algebra map $\pi : A \to B$, where $X = \mathrm{spec}(A)$ and $Y = \mathrm{spec}(B)$), the tangent bundle of Y is a subscheme of TX: for all k-algebras C, $TY(C) = Y(C[\epsilon]) \subseteq X(C[\epsilon]) = TX(C)$. In terms of derivation, for a point $y \in Y(C)$, the space T_yY consists of those derivations $v : A \to C$ that come from B: $v = v_0 \circ \pi$ for some derivation $v_0 : B \to C$ with respect to y. If I is the kernel of π, this is equivalent to $v(f) = 0$ for all $f \in I$. Thus, v is tangent to the "level-sets" $f = 0$ for $f \in I$, as expected.

It follows that we can explicitly compute the tangent bundle of any affine variety: First, we compute the tangent bundle of $\mathbb{A}^n$, the affine space:

Example 3.3.12. The affine space $\mathbb{A}^n$ satisfies $\mathbb{A}^n(R) = R^n$ for any ring R. In particular, $\mathbb{A}^n(R[\epsilon]) = (R[\epsilon])^n$, which we may identify with R^{2n}. Thus, $T\mathbb{A}^n = \mathbb{A}^{2n}$ (geometrically, the tangent space at each point of $\mathbb{A}^n$ may be canonically identified with the tangent space $\mathbb{A}^n$ at 0) □

Now, by definition, each affine variety can be embedded as a closed subset in some affine space, given by some equations $f_1 = \cdots = f_k = 0$. We compute the solutions in $R[\epsilon]$, and use the identification above.

Example 3.3.13. Let $\mathbf{X}$ be the affine subvariety of $\mathbb{A}^2$ given by $x^2+y^2 = 1$. A general point of the tangent space is given by $(x_0 + x_1\epsilon)^2 + (y_0 + y_1\epsilon)^2 = 1$. Expanding, we get the equations $x_0^2 + y_0^2 = 1$ and $2x_0x_1 + 2y_0y_1 = 0$ (recall that $\epsilon^2 = 0$), as expected. □

The above recipe essentially proves existence, but in an inconvenient form that requires an embedding and additional choices. We outline the proof of a more general fact, that implies our statement and also a twisted version we need below. To see this, we note that for a k-algebra R, we may write $R[\epsilon] = R \otimes_k E$, where $E = k[\epsilon]$ is a k-algebra which is of dimension 2 as a vector space over k. Given a k-algebra A_0, we are thus looking for another k-algebra A_E with a functorial identification $Hom_{k-alg}(A_E, R) = Hom_{k-alg}(A_0, R \otimes_k E)$. Since $Hom_{k-alg}(A_0, R \otimes_k E) = Hom_E(A, R \otimes_k E)$, where $A = E \otimes_k A_0$, our goal is achieved by the following general statement.

Theorem 3.3.14 (Weil restriction of scalars, affine case)**.** *Let k be a field, and let E be a k-algebra which is finite-dimensional as a vector space over k. Then for every E-algebra A there is a (necessarily unique) k-algebra A_E with a functorial identification $Hom_{k-alg}(A_E, R) = Hom_{E-alg}(A, R \otimes_k E)$ for all k-algebras R.*

Sketch of proof. Let U, W be k-vector spaces. Since E is finite dimensional, we have on the level of linear maps,

$$Hom_k(Hom_k(E, U), W) = Hom_k(U, E \otimes_k W).$$

If, in addition, U is an E-module, the subspace $Hom_E(U, E \otimes_k W)$ of the right hand side corresponds to the subspace $Hom_k(Hom^E(E, U), W)$, where

$Hom^E(E,U)$ is the largest quotient of $Hom_k(E,U)$ identifying the two E-module structures (i.e., the quotient of $Hom_k(E,U)$ by the subspace generated by all differences $t \circ e - et$, for all $e \in E$). Note that if $U = E \otimes_k U_0$ for some k-vector space U_0, the map $Hom_k(E,U_0) \to Hom^E(E,U)$ is injective, hence an isomorphism.

Applying this to $U = A$ and $W = R$, we see that we are looking to classify those linear maps from $V = Hom^E(E,A)$ that correspond to those maps on the right that respect the ring structure.

Let $V_0 = Hom_k(E,k)$. By the remark above, V_0 is canonically idetified with $Hom^E(E,E) \subseteq V$. Let $V_1 = Hom^E(E, A \otimes_E A)$. The product $m : A \otimes_E A \to A$ on A is an E-module map, so induces a map $m_1 : V_1 \to V$ (by composition), while the product $n : E \otimes E \to E$ induces a map $n_1 : V_1 \to V \otimes_k V$ (again since E is finite-dimensional).

Let $t : V \to R$ be a k-linear map, and let $t' : A \to E \otimes_k R$ be the corresponding E-linear map as above. It is easy to check that

(1) $t'(1) = 1$ if and only if $t(\phi) = \phi(1) \in k \subseteq R$ for all $\phi \in V_0$
(2) t' respects the product if and only if $m_R \circ (t \otimes_k t) \circ n_1 = t \circ m_1$, where $m_R : R \otimes_k R \to R$ is the product on R.

Now, a general k-linear map $t : V \to R$ corresponds to a k-algebra map $\tilde{t} : Sym_k(V) \to R$, so we define A_E as $Sym_k(V)$ divided by the relations above, namely by the (ideal generated by the) elements $n_1(v) - m_1(v)$ and $v_0 - v_0(1)$ for $v \in V_1$ and $v_0 \in V_0$. □

End of lecture 11, May 3

Exercise 3.3.15. Compute explicitly A_E in the following cases:

(1) $E = k \times k$
(2) E is a finite Galois field extension of k
(3) $E = k[\epsilon]$ with the obvious k-algebra structure: If you are familiar with Kahler differentials, show explicitly that A_E is isomorphic to $Sym_A(\Omega A)$ in this case.
(4) E is a finite purely inseparable field extension of k.

□

When K is a differential field, we will need a twisted version of the tangent bundle: Given an affine variety $\mathbf{X}$ over K, with algebra of functions A, we are looking for an affine variety $\tau_K\mathbf{X}$ with an identification of $\tau\mathbf{X}(R)$ with points $A \to R[\epsilon]$ that restrict to the map $K \to K[\epsilon]$ determined by the derivative on K (for any algebra R over K), i.e., which classifies derivations from A to R that extend the one on K. Geometrically, if $K = \mathbb{R}(t)$, with $t' = 1$ interpreted as time, the variety $\mathbf{X}$ itself changes with time, and as a result, the tangent vector to a smooth curve into $\mathbf{X}$ will not lie in the tangent bundle, but rather in a shifted version.

The construction of $\tau\mathbf{X}$ is similar to the tangent bundle, taking Weil restriction with respect to $E = K[\epsilon]$ with the K-algebra structure given by the derivation. In terms of formulas, the computation is similar, but applying the map $K \to K[\epsilon]$ to the coefficients.

Let K be a differential field, $\mathbf{X}$ an affine variety over K with A its algebra of functions, and $\tau\mathbf{X} = \tau_K\mathbf{X}$ the corresponding twisted tangent bundle. If $x : A \to K$ is a K-point of $\mathbf{X}$, we may compose with the differential structure $K \to K[\epsilon]$, to obtain a $K[\epsilon]$-point (commuting with the map $K \to K[\epsilon]$ in the base), so a point $\delta(a) \in \tau\mathbf{X}(K)$ (and similarly for differential extensions of K). This point is called the *prolongation* of a. If $a \in \mathbb{A}^n(K)$, then $\delta(a) = \langle a, a'\rangle$, where a' is the derivative of K applied to each coordinate. In particular, the map $x \mapsto \delta(x)$ is definable (without quantifiers) in the theory of differential fields. **prolongation**

Let $\mathbf{X}$ be an affine variety over the differential field K, and let $\mathbf{Z}$ be an affine subvariety of $\tau\mathbf{X}$. We set $\mathbf{Z}^{\#} = \{x \in \mathbf{X} | \delta(x) \in \mathbf{Z}\}$ (viewed either as a Kolchin closed or definable subset of $\mathbf{X}$). In terms of points, $\mathbf{Z}^{\#}$ can be described as follows. Let A and C be the K-algebras of functions on $\mathbf{X}, \mathbf{Z}$, respectively. The inclusion of $\mathbf{Z}$ in $\tau\mathbf{X}$ corresponds to a map of rings $A \to C[\epsilon]$ (restricting to the map $K \to K[\epsilon]$ corresponding to the derivation on K). If L is a differential ring extending K, an L-point of $\mathbf{Z}$ is an L-point $x : A \to L$ in the algebraic sense, such that the induced map $d_L \circ x : A \to L[\epsilon]$ factors via $C[\epsilon]$.

Example 3.3.16. Let $\mathbf{X} = \mathbb{A}^1_K$, the 1-dimensional affine space over $K = \mathbb{C}(t)$ (with $t' = 1$). Since $\mathbb{A}^1$ is defined over $\mathbb{Q}$, $\tau\mathbf{X}$ is the usual tangent bundle, which we identified as $\mathbb{A}^2$. Algebraically, the algebra A above is $K[x]$, and the algebra for the tangent bundle is $K[x_0, x_1]$. Let $\mathbf{Z} \subseteq \mathbb{A}^2$ be given by $x_1 = x_0$, so $C = K[x_0, x_1]/(x_1 - x_0)$. The canonical map $A \to C[\epsilon]$ induced by the inclusion sends x to $x_0 + x_1\epsilon$ (and a to $a + a'\epsilon$ for $a \in K$). If L is the field of meromorphic functions on $\mathbb{C}$, with the usual derivation and $K \to L$ the obvious inclusion, a map from A to L is determines by the image $f \in L$ of x. This image is in $\mathbf{Z}^{\#}$ precisely if $x_0 + x_1\epsilon$ can be sent to $f + f'\epsilon$. Since $x_0 = x_1$ in C, this means that $f' = f$, as expected. $\square$

Example 3.3.17. Again let $K = \mathbb{C}(t)$ with $t' = 1$, and let $\mathbf{X}$ be given by $x^2 + ty^2 = 1$. To compute $\tau\mathbf{X}$, we work as before, but also apply the map $d : K \to K[\epsilon]$ to the coefficients:

$$1 = (x_0 + x_1\epsilon)^2 + (t + t'\epsilon)(y_0 + y_1\epsilon)^2 = x_0^2 + ty_0^2 + \epsilon(y_0^2 + 2x_0x_1 + 2ty_0y_1)$$

So the equation for $\langle x_1, y_1\rangle$ is no longer homogenous (when $y_0 \neq 0$). Geometrically, we may think of $\mathbf{X}$ as an ellipse changing with the time t, always going through the points $\langle \pm 1, 0\rangle$, but becoming closer to the x-axis as time flows in the positive direction. As a result, a point moving on this ellipse through the point of intersection with the y-axis will have a non-zero velocity component in the y-direction, so will not be tangent to the ellipse, but instead will satisfy the relation defining $\tau\mathbf{X}$. $\square$

A map $f : \mathbf{Z} \to \mathbf{X}$ of affine varieties (or schemes) is called *dominant* if the corresponding map of algebras is injective. Geometrically, the image of f is dense in $\mathbf{X}$. **dominant**

Theorem 3.3.18. *Let E be a differential field. Then E is existentially closed if and only if the following two conditions hold:*

1. *E is algebraically closed*
2. *For every irreducible affine variety $\mathbf{X}$ over E and every irreducible affine subvariety $\mathbf{Z} \subseteq \tau\mathbf{X}$ projecting dominantly on $\mathbf{X}$, $\mathbf{Z}^{\#}(E) \neq \emptyset$*

Zariski closure

locus

Before the proof, we point out some terminology: If $\mathbf{X}$ is an affine variety, and $S\mathbf{X}(A)$ is *any* subset (for some K-algebra A), the Noetherian property implies that there is a minimum Zariski closed subset $\mathbf{S}$ of $\mathbf{X}$ over K containing S. $\mathbf{S}$ is called the *Zariski closure* of S. When $S = \{s\}$ is a singleton, it is also called the *locus* of s (over K).

The proof of the theorem relies upon the following observation, which is of independent interest.

Proposition 3.3.19. *Let K be a (char. 0) subfield of a field L, and let A be an L-algebra. Then any derivation $\partial : K \to A$ extends to L. If L is algebraic over K, the extension is unique.*

Geometrically, if we are given an A-point of $\operatorname{spec}(L)$, and a vector field along its image in $\operatorname{spec}(K)$, we may lift the vector field to the point (uniquely, in the algebraic case).

Proof. For both parts, we may assume L is finitely generated over K as a field, hence by one element, $L = K(t)$. If t is transcendental, we may define $\partial(t) = 0$. Otherwise, let p be the minimal polynomial of t over K. Then we must have $0 = p^{\delta}(t) + p'(t)\partial(t)$ by (2.2), which implies uniqueness. By minimality of p (and char 0), $p'(t) \neq 0$ hence is invertible since it is in the field L. Defining $\partial(t) = \frac{-p^{\delta}(t)}{p'(t)} \in A$ solves the problem. □

Proof of Theorem 3.3.18. Assume that E is existentially closed. Then we need to show that each of the two properties holds in some differential field extension of E. The first follows immediately from Prop. 3.3.19. For the second, let $\mathbf{X}$ and $\mathbf{U}$ be as in the condition. The dominant projection corresponds to an injective map $A \to B$ of E-algebras, which are integral domains since both are irreducible, hence to a map of fraction fields $K \to L$. Further, the inclusion of $\mathbf{Z}$ in $\tau\mathbf{X}$ corresponds to a derivation $\partial : K \to L$, as discussed above. By Prop. 3.3.19, ∂ extends to L. The inclusion of A in L corresponds to the L-point $B \to L$ of $\mathbf{Z}$.

In the other direction, assume that E satisfies the two conditions, and let ϕ be some quantifier-free formula satisfied in an extension L of E. By adding more variables, we may assume that $\phi(x) = \psi(x, x')$, where ψ is in the language of fields. If $a \in L$ satisfies ϕ, let $\mathbf{X}$ be the locus of a over E, and $\mathbf{U}$ the locus of a, a'. Then $\mathbf{U} \subseteq \tau\mathbf{X}$, projects dominantly on $\mathbf{X}$, and both are irreducible. Then $\mathbf{U}$ and ψ agree on some affine subvariety $\mathbf{Z} \subseteq \mathbf{U}$, still projecting dominantly on $\mathbf{X}$. By the axioms, $\mathbf{Z}$ has an E-point, which is then also a solution of ϕ. □

End of lecture 12, May 9

Corollary 3.3.20. *The theory DD of differential domains admits a model companion. The geometric conditions in Theorem 3.3.18 are equivalent to Blum's axioms for DCF*

Proof. We need that the geometric conditions are first order. The only issue is the definability of irreducibility. This is classical, but can also be deduced from [3]. $\square$

Corollary 3.3.21. *Let* $\mathbf{X}$ *be an irreducible variety over a differentially closed field* K*, and let* $s : \mathbf{X} \to \tau\mathbf{X}$ *be a regular section (i.e., an algebraic map such that* $\pi \circ s$ *is the identity on* $\mathbf{X}$*). Then there is* $a \in \mathbf{X}(K)$ *such that* $a' = s(a)$*.*

Proof. Apply the axioms to $\mathbf{X}$ and $\mathbf{Z}$ the graph of s. $\square$

We note that in both the theorem and the corollary, there are actually Zariski-densly many solutions, since we may replace $\mathbf{X}$ by an affine open subset.

Example 3.3.22. Let A be an $n \times n$ matrix over a differential field K. The set of solutions to the system $y' = Ay$ is easily seen to be a vector space over C_K, of dimension at most n.

Assume that K is differentially closed. Then the dimension is precisely n. To see that, consider $\mathbf{X} = GL_n = \{B \in End(K^n) | \det(B) \in K^x\}$. This is an irreducible affine variety, and to show our claim, we need to show there is a point B of $\mathbf{X}(K)$, such that $B' = AB$. This follows from Corollary 3.3.21, with $s(Y) = AY$. $\square$

4. Structure of definable sets

We review some structure theory for definable sets in general theories. Good references for this include [8, §4.3, 6.2, 6.3] and [11, §6].

4.1. **Working with parameters.**

Let $\mathcal{T}_0$ be a theory with model completion $\mathcal{T}$. We would like to work in the following manner: given a model A of $\mathcal{T}_0$, embed it in a model M of $\mathcal{T}$, work there (taking advantage of quantifier elimination and other good properties of $\mathcal{T}$), then go back to the set of parameters A that we started with. The fact that we may embed A in some model of $\mathcal{T}$ is part of the definition of model companion. However, we may have more than one model of $\mathcal{T}_0$, and we may like to embed all of them in the same model of $\mathcal{T}$ (possibly with some relation between them).

Descending to A we did not discuss at all. For that, we first have the following definition.

Definition 4.1.1. Let $C \subseteq M^n$ be a subset definable with parameters, and let $A \subseteq M$. We say that C is A-*definable*, or *definable over* A if $C = \phi(a, x)^M$ for some formula ϕ and some $a \in A$.

A-definable
definable over A

An element $b \in M$ is A-definable if the singleton $\{b\}$ is A-definable. The set of all $b \in M$ definable over A is called the *definable closure* of A, denoted $\mathrm{dcl}(A)$.

definable closure

definable function

By a *definable function* we mean a definable set which is the graph of a function (note that this is a first order property).

Exercise 4.1.2. For all $A \subseteq M$, $\mathrm{dcl}(A) = \{f(a) | f \text{ definable}\}$ □

Example 4.1.3. Let M be an algebraically closed field, and let $A \subseteq M$. Every element in the subfield K generated by A is definable over A, by the above exercise, since the field operations are definable. In characteristic $p > 0$, the inverse $x \mapsto x^{\frac{1}{p}}$ is a definable function, so the perfect closure K_p of K is in the definable closure as well.

So $K_p \subseteq \mathrm{dcl}(A)$. Is there anything else? Let $\sigma : M \to M$ be an automorphism fixing A (pointwise). If $\phi(x, a)$ is any formula with $a \in A$, then $m \in \phi(x, a)^M$ if and only if $\sigma(m) \in \phi(x, \sigma(a))^M = \phi(x, a)^M$, so σ preserves ϕ^M (as a subset). In particular, such a σ will fix any element of $\mathrm{dcl}(A)$. However, any $m \in M$ outside K_p is moved by some element of $Aut(M/A)$. Hence $K_p \subseteq \mathrm{dcl}(A)$. □

The method of automorphisms worked well in the last example, will it work always?

Example 4.1.4. The field $\mathbb{R}$ (the real numbers) admits no non-trivial automorphisms: Each such automorphism would have to preserve the set of non-negative reals, since it coincides with those reals that have a square. Hence, it is a topological automorphism, and is therefore determined by its restriction to the dense subset $\mathbb{Q}$. On this subset it must be the identity, since $\mathbb{Q}$ is the prime field.

So the whole of $\mathbb{R}$ is fixed by $Aut(\mathbb{R}/A)$, where $A = \mathbb{Q}$, but the whole $\mathbb{R}$ cannot be the definable closure of A, since that definable closure is countable. □

The point is that the definable closure (and more generally, definability over a set) does not depend on the enclosing model, while the automorphism group does, and so we might choose the wrong model for it. Can we always choose a right one?

κ-homogeneous

strongly κ-homogeneous

Definition 4.1.5. Let κ be a cardinal. A model M is called *κ-homogeneous* if for every κ-small $A \subseteq M$ and every $b \in M$, each elementary map $i : A \to M$ extends to b. It is called *strongly κ-homogeneous* if every such elementary map i extends to an automorphism of M.

If κ is omitted, we take $\kappa = |M|$.

Strongly κ-homogeneous is the notion we are after: it implies that for κ-small A, the automorphism group $G = Aut(M/A)$ acts transitively on the realisations (in M) of each type over A.

Exercise 4.1.6. Let κ be a cardinal.

(1) If M is strongly κ-saturated, then it is κ-homogeneous.
(2) If M is κ-saturated, then it is κ-homogeneous.

(3) If M is homogeneous, then it is strongly-homogeneous (note we may assume that as a set, $M = \kappa$ for this case).

□

It follows from the last exercise that a saturated model (i.e., κ-saturated for κ its own cardinality) is strongly homogeneous. If we have such a model $\mathbb{U}$, of cardinality κ, and we are only interested in structures of smaller size, we may embed any such structure A in $\mathbb{U}$ (by saturation), along with any additional parameters or sets we are interested in, and the automorphism group $Aut(\mathbb{U}/A)$ will act transitively on all types over A. In other words, we may work entirely within $\mathbb{U}$.

In general, saturated models need not exist: their existence is independent of ZFC. This issue may be overcome in several ways:

(1) Use a slightly weaker (but more complicated) type of model called *special models* that enjoys almost the same benefits, and whose existence does not depend on set theory.
(2) Use saturated models, and rely on general absoluteness results that allow to deduce the final outcome independently of set theory.
(3) Use κ-saturated models (of arbitrary cardinality) and give more complicated arguments
(4) Actually, for the theories we are interested in (ACF, DCF), saturated models always exist

Either way, we will work with a fixed universal domain $\mathbb{U}$, satisfying the properties mentioned above. End of lecture 13, May 10

4.1.7. *Type spaces.* We now go back to the question of parameters for definable sets. It will be convenient to adopt a topological language coming from Stone duality:

Let $\mathbf{X}$ be a definable set, and let $B \subseteq U$ be a set of parameters. The collection $\mathbb{D}_{\mathbf{X}}(B)$ of B-definable subsets of $\mathbf{X}$ forms a boolean algebra, and the compactness theorem implies that the set of complete types over B that include $\mathbf{X}$ can be identified with the *Stone space* of this algebra. We denote this space $\mathcal{S}t_{\mathbf{X}}(B)$, and call it the *type-space* of $\mathbf{X}$ over B. In case $\mathbf{X}(M) = M^n$, this is also denoted $\mathcal{S}t_n(B)$. Recall that this is a compact Hausdorff topological space, with a basis of clopen sets that admit a canonical bijection with the elements of $\mathbb{D}_{\mathbf{X}}(B)$: For each $\phi \in \mathbb{D}_{\mathbf{X}}(B)$, the set $\bar{\phi} = \{p \in \mathcal{S}t_{\mathbf{X}}(B) | \phi \in p\}$ is open (hence also closed). Alternatively, $\mathbb{D}_{\mathbf{X}}(B)$ is the algebra of continuous functions from $\mathcal{S}t_{\mathbf{X}}(B)$ to $\mathbb{F}_2$ (with the discrete topology). In some cases, $\mathbf{X}$ does not play an explicit role, so we will omit it from the notation (i.e., $\mathbb{D}(A)$ denotes $\mathbb{D}_{\mathbf{X}}(A)$ for some fixed definable $\mathbf{X}$) **type-space**

If $\mathbb{U}$ is a universal domain, and $B \subseteq \mathbb{U}$ is small, we have a surjective map $\mathbf{X}(\mathbb{U}) \to \mathcal{S}t_{\mathbf{X}}(B)$, $x \mapsto \mathrm{tp}(x/B)$, and the group $Aut(\mathbb{U}/B)$ acts transitively on the set of realisations $p(\mathbb{U})$ for all $p \in \mathcal{S}t_{\mathbf{X}}(B)$.

This topological language simplifies definability arguments. For example:

Proposition 4.1.8. *Let* $\mathbf{X}$ *be* U*-definable, and let* $A \subseteq \mathbb{U}$ *(small). Then* $\mathbf{X}$ *is* A*-definable if and only if* $\mathbf{X}(\mathbb{U})$ *is invariant under* $G = Aut(\mathbb{U}/A)$*.*

In particular, $\mathrm{dcl}(A) = \mathbb{U}^G$ *(the fixed points under* G*).*

Proof. Assume that $\mathbf{X}$ is defined over $B \supseteq A$. The inclusion $\mathbb{D}(A) \subseteq \mathbb{D}(B)$ induces a surjective continuous restriction map $r : \mathcal{S}t(B) \to \mathcal{S}t(A)$. If $p \in r(\bar{\mathbf{X}})$, there is a type q over B extending p and including $\mathbf{X}$. Since $Aut(\mathbb{U}/A)$ acts transitively on $p(\mathbb{U})$ and preserves $\mathbf{X}(\mathbb{U})$, every other pre-image q' of p also includes $\mathbf{X}$. It follows that $\bar{\mathbf{X}} = r^{-1}(r(\bar{\mathbf{X}}))$ so $r(\bar{\mathbf{X}}^c) = r(\bar{\mathbf{X}})^c$, hence $r(\mathbf{X})$ is a clopen subset of $\mathcal{S}t(A)$, so corresponds to an element of $\mathbb{D}(A)$. □

4.2. **Morley rank.** We would like to define abstractly a notion of dimension (or *rank*) for definable sets. Our main notion will follow an intuition similar to that of the Krull dimension for affine varieties (Def. 2.3.19). However, we do not have a notion of closed sets in general, so it will have to be modified.

Definition 4.2.1. The relation $\mathcal{MR}(\mathbf{X}) \geq \alpha$ is defined inductively for $\mathbb{U}$-definable sets $\mathbf{X}$ and α an ordinal or -1 (where 0 is an immediate successor of -1) as follows:

1. $\mathcal{MR}(\mathbf{X}) \geq -1$ for all $\mathbf{X}$
2. If α is a limit ordinal, $\mathcal{MR}(\mathbf{X}) \geq \alpha$ if $\mathcal{MR}(\mathbf{X}) \geq \beta$ for all $\beta < \alpha$
3. For all α, $\mathcal{MR}(\mathbf{X}) \geq \alpha + 1$ if $\mathbf{X} \supset \coprod_{i \in \mathbb{N}} \mathbf{X}_i$ for some $\mathbb{U}$-definable $\mathbf{X}_i$ with $\mathcal{MR}(\mathbf{X}_i) \geq \alpha$

If, for some α, $\mathcal{MR}(\mathbf{X}) \geq \alpha$ but not $\mathcal{MR}(\mathbf{X}) \geq \alpha + 1$, we say $\mathbf{X}$ has Morley rank α. Otherwise, the Morley rank of $\mathbf{X}$ does not exist, or is ∞ (so $\mathcal{MR}(\mathbf{X}) \geq \mathcal{MR}(\mathbf{Y})$ for all $\mathbf{Y}$, in this case)

Clearly, $\mathbf{X}$ has Morley rank both α and β, then $\alpha = \beta$, so we may refer to *the* Morley rank of $\mathbf{X}$ (if it exists). We note that $\mathcal{MR}(\mathbf{X}) \geq 0$ if and only if $\mathbf{X}$ is non-empty, and $\mathcal{MR}(\mathbf{X}) = 0$ if and only if it is non-empty and finite.

Proposition 4.2.2. *For all* $\mathbb{U}$*-definable* $\mathbf{X}$ *and* $\mathbf{Y}$*,*

1. $\mathcal{MR}(\mathbf{X}) = -1$ *if and only if* $\mathbf{X} = \emptyset$, $\mathcal{MR}(\mathbf{X}) = 0$ *if and only if* $\mathbf{X}$ *is finite and non-empty.*
2. *If* $\mathbf{X} \subseteq \mathbf{Y}$*, then* $\mathcal{MR}(\mathbf{X}) \leq \mathcal{MR}(\mathbf{Y})$
3. $\mathcal{MR}(\mathbf{X} \cup \mathbf{Y}) = \max(\mathcal{MR}(\mathbf{X}), \mathcal{MR}(\mathbf{Y}))$
4. *If* $\mathcal{MR}(\mathbf{X}) = \alpha < \infty$*, then for each* $\beta \leq \alpha$*, there is a definable* $\mathbf{Z} \subseteq \mathbf{X}$ *with* $\mathcal{MR}(\mathbf{Z}) = \beta$*.*

Proof. By ordinal induction □

Proposition 4.2.3. *In* ACF_p $(p \geq 0)$*, each definable set has Morley rank, which is equal to its Krull dimension.*

Proof. By induction on the dimension n of $\mathbf{X}$. If $\mathbf{Y} = \bar{\mathbf{X}} \setminus \mathbf{X}$, then $\mathbf{Y}$ has Krull dimension less than n. Hence, by the previous proposition and induction, we may assume $\mathbf{X}$ is closed. By Noether Normalization, there is a finite surjective map $p : \mathbf{X} \to \mathbb{A}^n$. For $\mathbb{A}^n$, we have $\mathbb{A}^{n-1}$ embedded of Krull dimension $n-1$, and so are its translates. These translate can be lifted to disjoint subsets of $\mathbf{X}$ of Krull dimension $n-1$, so by induction, $\mathcal{MR}(\mathbf{X}) \geq n$. On the other hand, if $\mathbf{Y} \subset \mathbf{X}$ has $\mathcal{MR}(\mathbf{Y}) \geq n$, then by induction, the Krull dimension of $\mathbf{Y}$ is at least n, so there cannot be infinitely many disjoint such. □

End of lecture 14, May 16

Every affine variety has a finite number of irreducible components, which are (essentially) uniquely determined. This decomposition depends on the notion of closed sets, and cannot be recovered abstractly. However, we may recover the components of highest dimension, at least up to subsets of smaller dimension. To make this precise, let $\mathbf{Z}$ be a definable set with Morley rank $\beta < \infty$, let $\alpha \leq \beta$ be an ordinal, let I_α be the ideal in $\mathbb{D}_{\mathbf{Z}}(U)$ consisting of Morley rank less than α, and let $\mathbb{D}^\alpha_{\mathbf{Z}}(U)$ be the quotient algebra, with $[\mathbf{X}]_\alpha$ the image of $\mathbf{X}$ in it. Note that the Morley rank of an element there is well defined (and is at least α or -1)

Recall that an *atom* in a boolean algebra is a minimal non-zero element there, and the algebra is an *atomic boolean algebra* if each non-zero element has an atom below it.

atom

atomic boolean algebra

Proposition 4.2.4. *Let $\mathcal{MR}(\mathbf{Z}) = \beta < \infty$, and let $\alpha \leq \beta$. The boolean algebras $\mathbb{D}^\alpha_{\mathbf{Z}}(U)$ are atomic. An element $[\mathbf{X}]_\alpha$ is an atom if and only if it has Morley rank α and is not the union of two such sets. An element has Morley rank α if and only if it is a finite union of such atoms, which are then uniquely determined.*

Proof. We show that every element of rank α is a finite union of elements as in the statement. If b is such an element and is not an atom, then $b = b_0 \vee b_1$, with b_i disjoint, and of rank α. Continuing this way, we produce a sequence b_w where w is a finite sequence of $0, 1$, $b_u < b_w$ if w is a prefix of u, each of rank α, and a sequence c_w of complements to the b_w, also of rank α. Then the c_w show that b has rank at least $\alpha + 1$. □

Definition 4.2.5. For $\mathbf{X}$ of rank α, the number of components of $[\mathbf{X}]_\alpha$ in the above decomposition is called the *Morley degree* of $\mathbf{X}$, denoted $\mathcal{MD}(\mathbf{X})$ (a positive natural number).

Morley degree

Example 4.2.6. The degree of a finite non-empty set (i.e., a set of rank 0) is the number of elements □

Example 4.2.7. A definable set $\mathbf{X}$ is of Morley rank 1 and Morley degree 1 if and only if it is infinite, and each definable subset of it is finite or co-finite. Every set of Morley rank 1 is a union of d such, where d is the degree, uniquely determined up to finite sets. □

Example 4.2.8. In ACF, we saw that Morley rank coincides with Krull dimension. Taking Zariski closure only adds sets of lower dimension, so to understand Morley degree, we may assume $\mathbf{X}$ is Zariski closed. It will then be a finite union of irreducible components, each of Morley degree 1 (here, we allow extending parameters even when using the algebraic language: the components might be defined over a larger field than the original set). The degree of $\mathbf{X}$ will thus be the number of irreducible components of the same dimension as $\mathbf{X}$. $\square$

Exercise 4.2.9. If $\mathbf{X}$ and $\mathbf{Y}$ are disjoint and have the same Morley rank, then $\mathcal{MD}(\mathbf{X} \cup \mathbf{Y}) = \mathcal{MD}(\mathbf{X}) + \mathcal{MD}(\mathbf{Y})$. $\square$

Example 4.2.10. Let $\mathcal{T}$ be the theory of the reals $\mathbb{R}$, as an ordered field, and let $\mathbf{X}$ be the universe (i.e., the set defined by $x = x$). In $\mathcal{T}$ there is a definable bijection between $\mathbf{X}$ and the positive reals $\mathbf{X}_+$: For example, the map given by $t \mapsto 1 - \frac{1}{t}$ on $(0, 1)$ and $t \mapsto t - 1$ on $[1, \infty)$ is a bijection. Similarly, it is definably isomorphic to the negative reals $\mathbf{X}_-$. If all these sets have Morley rank α and Morley degree $d > 0$, we would have by the previous exercise that $d = d + d = 2d$. Hence, $\mathbf{X}$ does not have a Morley rank. $\square$

End of lecture 15, May 17

4.2.11. *Ranks for types.* The definitions Morley rank and degree can be extended to infinite intersections P of (realizations in $\mathbb{U}$ of) definable sets, by setting $\mathcal{MR}(P) = \min\{\mathcal{MR}(\mathbf{X}) | P \subseteq \mathbf{X}\}$ and $\mathcal{MD}(P) = \min\{\mathcal{MD}(\mathbf{X}) | P \subseteq \mathbf{X}, \mathcal{MR}(\mathbf{X}) = \mathcal{MR}(P)\}$. In particular, these definitions apply when P is a complete type p (over some small set A). We then have, by Prop. 4.2.4, that $\mathcal{MR}(\mathbf{X}) = \max\{\mathcal{MR}(p) | p \in \mathcal{S}t_{\mathbf{X}}(B)\}$, and each type contains a formula of the same rank and degree, uniquely determined up to smaller rank, and each such formula topologically isolates it from the rest of the corresponding type space.

It follows that Morley rank can be defined directly in terms of types: the rank of p is (inductively) the smallest α for which p is isolated in the space $\mathcal{S}t^{\alpha}$ where the types of smaller rank were removed (this is a purely topological description, called the Cantor–Bendixon rank).

4.2.12. *ω-stability.* Morley rank is intuitively appealing, but might be difficult to verify directly. We have the following alternative characterization.

Proposition 4.2.13. *Let $\mathbf{X}$ be a definable set in a countable theory $\mathcal{T}$. The following are equivalent:*

1. $\mathcal{MR}(\mathbf{X}) < \infty$
2. *For all structures A, $|\mathcal{S}t_{\mathbf{X}}(A)| = |\mathcal{T}_A|$ ($= |A|$ if A is infinite)*
3. *For all countable A, $|\mathcal{S}t_{\mathbf{X}}(A)| < 2^{\aleph_0}$*
4. *Each of the above for every $\mathbf{Z} \subseteq \mathbf{X}^n$ in place of $\mathbf{X}$*

Before the proof, we make the following observation: The Morley ranks of definable sets that do have a Morley rank is bounded, say be α (since there are more ordinals than elements of $\mathbb{U}$). Hence, if $\mathbf{X}$ has no Morley rank, we may find $\mathbf{Y} \subseteq \mathbf{X}$ such that both $\mathbf{Y}$ and $\mathbf{X} \setminus \mathbf{Y}$ have no Morley rank.

Proof.

(1) $\Longrightarrow$ (2)**:** Since $\mathbf{X}$ has Morley rank, so does each type in $\mathcal{S}t_{\mathbf{X}}(A)$. If p is such a type, of rank α, we say that it is isolated by a formula of the same rank, with the different types corresponding to different formulas. This defines an injective map from $\mathcal{S}t_{\mathbf{X}}(A)$ to the set of formulas.

(2) $\Longrightarrow$ (3)**:** Obvious

(3) $\Longrightarrow$ (1)**:** If $\mathbf{X}$ has no Morley rank, it contains two disjoint $\mathbb{U}$-definable subsets $\mathbf{X}_0, \mathbf{X}_1$ that also have no Morley rank. Replacing $\mathbf{X}$ with $\mathbf{X}_0$, we may continue, building a complete infinite binary tree. The branches of this tree extend to distinct types over the countable set A required to define the sets along the way, contradicting the assumption.

(3) $\Longrightarrow$ (4)**:** It suffices to prove the case $\mathbf{Z} = \mathbf{X}^n$, which follows by considering the maps $\mathcal{S}t_{\mathbf{X}^n} \to \mathcal{S}t_{\mathbf{X}^{n-1}}$ and counting (note that the type of $\langle a, b\rangle$ over A is "the same" as the type of a over A, b) $\square$

We remark that the above can also be viewed as a purely topological statement: every compact Hausdorff space with no isolated points must be of size at least continuum.

The property in (2) is central in classification theory:

Definition 4.2.14. For a cardinal κ, a theory is *κ-stable* if for each structure A of cardinality κ and each infinite definable set $\mathbf{X}$, $|\mathcal{S}t_{\mathbf{X}}(A)| = \kappa$. The theory is *stable* if it is κ-stable for some κ. **κ-stable** **stable**

So Prop. 4.2.13 asserts that if a theory is ω-stable, then it is κ-stable for each infinite κ, and that is equivalent to the existence of Morley rank for all definable sets.

It is not completely obvious how to prove directly that if $\mathcal{MR}(\mathbf{X})$ exists, so does $\mathcal{MR}(\mathbf{X}^2)$. In fact, the equality $\mathcal{MR}(\mathbf{X}\times\mathbf{Y}) = \mathcal{MR}(\mathbf{X}) + \mathcal{MR}(\mathbf{Y})$ one might expect need not hold. This is one of the reasons it is convenient to introduce another rank:

Definition 4.2.15. Let p be a type over some set A, with Morley rank $\alpha < \infty$

1. If $B \supseteq A$ is an extension, and q is a type over B extending p, q is a *forking extension* of p if $\mathcal{MR}(q) < \mathcal{MR}(p)$. We also say that q *forks over* A. **forking extension** **forks over**
2. The *Lascar rank* $U(p)$ of p is defined inductively: $U(p) \geq \alpha + 1$ if $U(q) \geq \alpha$ for some forking extension q of p (and by continuity for limit ordinals). **Lascar rank**

Remark 4.2.16. Our definition of "forking extension" is correct when Morley rank exists. However, forking can be meaningfully defined in much more general situations $\square$

4.3. **Ranks in DCF.** We now study the manifestation of ranks in the context of DCF. We closely follow [9, §5.5]. For $f \in K\{\bar{x}\}$ we write ϕ_f for the formula $f = 0$. By quantifier elimination, a type $p(\bar{x})$ over k is determined uniquely by the prime differential ideal $I_p = \{f \in k\{\bar{x}\} | \phi_f \in p\}$. Conversely, any prime differential ideal is of the form I_p for a (unique) type p. Both correspond to irreducible Kolchin closed subsets on $\mathbb{A}^n$.

Proposition 4.3.1. *The theory DCF is ω-stable*

Proof. It suffices to prove that the set of types in one variable over a countable differential field is countable. We just mentioned that each such type p is determined by a prime differential ideal I_p, and we had seen that each such ideal is of the form $I(f)$ for some $f \in I_p$. Hence, there are countably many of them. □

To facilitate rank computations, it is convenient to introduce two additional, more algebraic ranks.

Kolchin dimension

Definition 4.3.2. Let $\mathbf{X}$ be a Kolchin closed set. The *Kolchin dimension* $\mathcal{K}(\mathbf{X})$ of $\mathbf{X}$ is its dimension for the Noetherian topology of Kolchin closed subsets, i.e., $\mathcal{K}(\mathbf{X}) = \sup\{\mathcal{K}(\mathbf{Y}) + 1 | \mathbf{Y} \subset \mathbf{X} \text{ closed irreducible}\}$ for irreducible $\mathbf{X}$.

differential order

If $\mathbf{X} \subset \mathbb{A}^1$ is irreducible, it is given by a single polynomial f of minimal order and degree. The *differential order* $\mathcal{O}(\mathbf{X})$ of $\mathbf{X}$ is the order of f.

In either case, the rank of a type is defined as the rank of the corresponding Kolchin closed set. These ranks may be viewed as analogs of Krull dimension and transcendence degree of the function field, respectively. The last point is reinforced by the following.

Proposition 4.3.3. *If $a \in L$ is a generic solution of $f(x) = 0$, where f is a differential polynomial over K of order n, then the transcendence degree of $K\{a\} \subseteq L$ over K is n*

Proof. By assumption, $a, a', \dots, a^{(n-1)}$ are algebraically independent, and higher derivatives depend on them by f and its derivatives, using minimality of f and equation (2.2). □

End of lecture 16, May 23

We would like to relate the four ranks we defined, and also to compute them in some cases. As a first observation, we note

Proposition 4.3.4. *For all types p over a differential field k, $\mathcal{U}(p) \leq \mathcal{MR}(p) \leq \mathcal{K}(p)$. If p is a 1-type, $\mathcal{K}(p) \leq \mathcal{O}(p)$.*

Proof. We mentioned already that $\mathcal{U}(p) \leq \mathcal{MR}(p)$ in an arbitrary theory. We show $\mathcal{MR}(p) \leq \mathcal{K}(p)$ by induction on $\alpha = \mathcal{MR}(p)$. Let $\mathbf{X}$ be the Kolchin closed set corresponding to p. To show $\mathcal{K}(\mathbf{X}) \geq \alpha$, it suffices to find, for each $\beta < \alpha$, a closed irreducible proper subset $\mathbf{Y}$ of $\mathbf{X}$ with $\mathcal{K}(\mathbf{Y}) \geq \beta$. By induction, we may replay $\mathcal{K}(\mathbf{Y})$ by $\mathcal{MR}(\mathbf{Y})$. Since $\mathcal{MR}(\mathbf{X}) = \alpha$ we may find infinitely many pairwise disjoint such $\mathbf{Y}$.

Finally, if p is a 1-type, we need to show that if $I_p = I(f) \subset I_q = I(g)$ is a proper extension of differential prime ideals, then the order of f is bigger than that of g. This is true since $f \in I(g)$, so by Lemma 2.4.16, if they have the same order, f is divisible by g, contradicting irreducibility of f. □

Corollary 4.3.5. *The following are equivalent for a type $p(x)$ over k:*

(1) *p is algebraic, i.e., $p(\mathbb{U})$ is finite*
(2) $\mathcal{U}(p) = 0$
(3) $\mathcal{O}(p) = 0$

Proof. Using the dimension inequality, the only thing to prove is that if p is algebraic, then $\mathcal{O}(p) = 0$. Let $a_1, \dots, a_n$ be all realisations of p, and let $I_p = I(f)$. If f has order bigger than 0, Blum's axioms for DCF imply that the formula ϕ_f has solutions distinct from the a_i. □

Corollary 4.3.6. *If $\mathcal{K}(p) = 1$ or p is a 1-type and $\mathcal{O}(p) = 1$, then $\mathcal{U}(p) = 1$.*

Proof. The upper bound is given by Prop. 4.3.4, and the lower by Cor. 4.3.5. □

Example 4.3.7. If $\mathbf{X}$ is a linear differential equation of order n, then $\mathcal{U}(\mathbf{X}) = \mathcal{O}(\mathbf{X}) = n$: The linear equation itself shows that $\mathcal{O}(\mathbf{X}) \leq n$. On the other, $\mathbf{X}$ is a linear space over $\mathbf{C}$ (the constants) of dimension n by Example 3.3.22. Choosing a basis, we find a definable bijection with $\mathbf{C}^n$ (note that $\mathcal{U}$ rank is preserved under definable bijections and when adding parameters). Since $\mathbf{C}$ is a pure algebraically closed field, $\mathcal{U}$-rank agrees with Krull dimension there, so $\mathcal{U}(\mathbf{X}) = \mathcal{U}(\mathbf{C}^n) = n$ □

Corollary 4.3.8. *The Morley and Lascar ranks of the universe $\mathbf{A}$ (i.e., of $x = x$) is ω.*

Proof. By Example 4.3.7, $\mathcal{U}(\mathbf{A}) \geq \omega$. On the other hand, if $\mathbf{Y} \subseteq \mathbf{A}$, then either $\mathbf{Y}$ or $\mathbf{A} \setminus \mathbf{Y}$ has finite differential order, hence finite Morley rank. Therefore $\mathcal{MR}(\mathbf{A}) < \omega + 1$. □

We had not seen examples where any of the ranks differ, so we may imagine that as in the case of ACF, they are all equal. It turns out that each of the three inequalities in Prop. 4.3.4 may be strict.

For $\mathcal{U}$-rank vs. Morley rank, the question may be asked abstractly: is there any type p in any theory, where $\mathcal{U}(p) < \mathcal{MR}(p) < \infty$? This question is easy to answer:

Exercise 4.3.9. Let $\mathcal{T}$ be the theory in a language with unary predicate symbols P_i ($i \in \mathbb{N}$), stating for each i that $P_{i+1} \subseteq P_i$ and $P_i \setminus P_{i+1}$ is infinite. Then each P_i has Morley rank 2, hence the unique type p containing all P_i has Morley rank 2, but $\mathcal{U}$-rank 1. □

The point is that the collection $\mathbf{X}_i$ of subsets witnessing higher Morley rank might be "of different shapes", rather than coming from a uniform family.

It could still be the case that $\mathcal{U}$-rank and Morley rank agree in DCF. This question was open for a while, but finally it was shown that this is not the case in [5].

Let $g(x)$ be a rational function, and consider the equation ϕ_g given by $x'' = g(x)x'$. If $h(x)$ is a rational anti-derivative of g, this equation is satisfied by any solution of $x' = h(x) + c$, where $c' = 0$. This equation has order 1, so its $\mathcal{U}$-rank is 1 as well, and varying c provides an infinite family of such. Hence, all ranks of ϕ_g are 2 in this case.

On the other hand, consider $\phi_{\frac{1}{x}}$, so the equation definable set $\mathbf{X}$ given by $xx'' = x'$. Any constant solves this equation, so $\mathbf{C} \subseteq \mathbf{X}$. Since this is a proper Kolchin-closed subset (and $\mathbf{X}$ is irreducible), we have $\mathcal{K}(\mathbf{X}) = 2$. We show outline a proof that $\mathcal{MR}(\mathbf{X}) = 1$ by showing that any $\mathbf{X}_i \subset \mathbf{X}$ of Morley rank 1 that could witness a higher Morley rank is in fact contained in the constants. Recall that each such potential subsets is given as the zero-set of an ideal $I(g)$, where g has order 1. Hence we are done once we show:

Lemma 4.3.10. *If g is an order 1 irreducible differential polynomial such that $xx'' - x' \in I(g)$, then $x' \in I(g)$*

Sketch of proof. We may write $g(x) = \sum_{i=0}^{n} a_i(x)(x')^i$, with a_i a polynomial. Differentiating, we get

$$\partial(g(x)) = \sum_i a_i^\partial(x)(x')^i + \sum_i \delta(a_i(x))(x')^{i+1} + x'' \sum_i i a_i(x)(x')^{i-1} \tag{4.1}$$

Where we denote by $\delta(h(x))$ the usual derivative of a polynomial $h(x)$ with respect to x (and, as before, by h^∂ the polynomial obtained by differentiating the coefficients). Since we are working modulo the equation $xx'' = x'$, multiplying the above by x yields

$$x\partial(g(x)) = x\sum_i a_i^\partial(x')^i + x\sum_i \delta(a_i)(x')^{i+1} + \sum_i i a_i(x)(x')^i \tag{4.2}$$

Since $I(g)$ contains the differential ideal generated by g, this differential polynomial is in $I(g)$. Since it is of order 1 and g is irreducible, we have by Lemma 2.4.16 that it is divisible by g. Hence the top coefficient a_n of g divides the top coefficient $x\delta(a_n)$. Comparing degrees, there must a be an element m of the ground differential field k such that $ma_n = x\delta(a_n)$. This equation can only be solved by a polynomial a_n if m is a natural number, and then $a_n = cx^m$ for some c, which we may assume to be 1. Considering the free coefficient, further calculations show that m must be 0, and that $g = x'$. □

The (completely algebraic) fact that differential order might be higher than Kolchin dimension is witnessed by an important classical equation, the *Painlevé equation* $x'' = 6x^2 + t$, where $t \in K = \mathbb{C}(t)$ satisfies $t' = 1$. It follows directly from the following fact:

Fact 4.3.11 (Kolchin)**.** *Every solution of the Painlevé equation $x'' = 6x^2 + t$ is of differential order 0 or 2*

This fact immediately yields the result, since a generic element of an infinite proper closed subset would have differential order 1. The proof of the fact is another calculation, which can be found in [9, §5.5] or the references there. The fact that this equation has Kolchin dimension (and therefore Morley rank) 1 is interesting because it makes it amenable to rank 1 analysis that we will see next. End of lecture 17, May 24

5. Definable sets of rank 1 in DCF

We now concentrate on definable sets of Morley rank 1 (hence also U-rank 1). Since each such set is a finite union sets of degree 1, we concentrate on the latter, which have their own name:

Definition 5.0.1. An infinite definable set $\mathbf{X}$ is *strongly minimal* if each definable subset of $\mathbf{X}$, with parameters, is finite or co-finite. **strongly minimal**

Our goal will be to give an explicit description of such sets occurring in DCF. This may seem like an ambitious goal, but not nearly as ambitious as classifying all strongly minimal sets in general, which we discuss first.

5.1. **Strongly minimal sets.** To get a feeling, let us consider a few examples of strongly minimal sets:

Example 5.1.1. The theory of infinite sets without any additional structure □

Exercise 5.1.2. Let $\mathcal{T}$ be the theory of $\langle \mathbb{Z}, s \rangle$, where $s : \mathbb{Z} \to \mathbb{Z}$ is the successor functions. Show that $\mathcal{T}$ eliminates quantifiers, and hence that it is strongly minimal. □

Example 5.1.3. Let F be a ring (not necessarily commutative). An F-module V can be viewed as a first-order structure whose signature includes 0, the addition, and for each element $a \in F$ a unary function symbol $\underline{a}$ for a function from V to itself. The theory reflects the relations among the various $\underline{a}$ that come from F, in addition to stating that V is an abelian group.

If V has any non-trivial definable subgroup, it cannot be strongly minimal. In particular, the kernel and image of each $\underline{a}$ must be either 0 or V, so that for $a \neq 0$, the map induced by a is invertible. Hence, V can only be strongly minimal if (it is infinite and) F is a division ring. It turns out that when F is a division ring (and V is infinite), one does obtain a strongly minimal set. □

Example 5.1.4. A variant of Example 5.1.3 is obtained by "forgetting 0" in V, i.e., the structure whose basic relations are affine subspaces of V. □

Example 5.1.5. Another variant of Example 5.1.3 is obtained by considering the associated projective space $P(V)$, the set of 1-dimensional subspaces (with the induced structure) □

Example 5.1.6. We saw that an algebraically closed field is strongly minimal □

Our goal in this section is to understand to which extent these examples are exhaustive. More precisely, we would like to answer the following questions:

(1) When should two examples be considered "essentially the same"?
(2) Can we distinguish different examples combinatorially?
(3) Are there examples that are essentially different than the ones we considered?

We begin with the second question, starting with the following definition:

algebraic
algebraic closure

Definition 5.1.7. Let $A \subseteq \mathbb{U}$ be a set of parameters. An element $b \in \mathbb{U}$ is *algebraic* over A if $\mathrm{tp}(b/A)$ has finitely many realisations. The *algebraic closure* $\mathrm{acl}(A)$ of A is the set of all elements algebraic over A.

As usual, the definition does not really depend on $\mathbb{U}$. The following can be compared with Prop. 4.1.8.

Exercise 5.1.8. Let $A \subseteq \mathbb{U}$ and $b \in \mathbb{U}$. The following are equivalent:

(1) b is algebraic over A
(2) There is a formula $\phi(x, \bar{a})$ with $\bar{a} \in A$, such that $\phi(b, \bar{a})$ holds, and $\phi^{\mathbb{U}}$ is finite.
(3) The orbit of b under $Aut(\mathbb{U}/A)$ is finite.
(4) b belongs to every model $M \subseteq \mathbb{U}$ that contains A.

□

We would like to abstract the combinatorial properties of the algebraic closure operator on subsets of $\mathbb{U}$. This is provided by the following definition:

finitary closure operator

Definition 5.1.9. A *finitary closure operator* on a set D is a map $\mathrm{cl} : \mathcal{P}(D) \to \mathcal{P}(D)$ satisfying for all $X \subseteq D$:

(1) $X \subseteq \mathrm{cl}(X)$ and $\mathrm{cl}(X) = \mathrm{cl}(\mathrm{cl}(X))$
(2) For all $Y \subseteq X$, $\mathrm{cl}(Y) \subseteq \mathrm{cl}(X)$
(3) $\mathrm{cl}(X) = \bigcup_{X_0 \subseteq X} \mathrm{cl}(X_0)$, where the union is over all finite $X_0 \subseteq X$ (this is the "finitary" part).

A subset of the form $\mathrm{cl}(X)$ is said to be *closed*.

pregeometry
matroid
Steinitz exchange property

A *pregeometry* (or *matroid*) is a pair $\langle D, \mathrm{cl}\rangle$ where cl is a finitary closure operator on the set D, satisfying in addition the *Steinitz exchange property*: For $a, b \in D$ and $X \subseteq D$, if $a \in \mathrm{cl}(X \cup \{b\}) \setminus \mathrm{cl}(X)$, then $b \in \mathrm{cl}(X \cup \{a\})$

It is easy to check that acl is a finitary closure operator on $\mathbb{U}$, for any first order structure $\mathbb{U}$. For strongly minimal sets we further have:

Proposition 5.1.10. *If* $\mathbb{U}$ *is a strongly minimal structure,* $\langle \mathbb{U}, \mathrm{acl}\rangle$ *is a pregeometry.*

Proof. Let a, b, X be as in the assumptions of the exchange property. We may assume $X = \emptyset$. Since a is algebraic over b, there is a formula $\phi(x, b)$ satisfied by finitely many elements, including a. We may assume that all fibres $\phi(x, b')$ have the same finite number of elements.

We claim $\phi(a, y)$ is algebraic (and then we are done). Otherwise, $\phi(a, c)$ is satisfied for all but k elements $c \in \mathbb{U}$, for some natural k. Then the formula

$$\psi(x) = \exists y_1 \dots y_k \forall y (\bigwedge_i y \neq y_i \to \phi(x, y))$$

is satisfied by a, algebraic, and defined without parameters, contradicting the assumption that $a \notin \mathrm{acl}(\emptyset)$ $\square$

We mentioned in Cor. 2.3.27 that the dimension of an algebraic variety can be computed as the transcendence degree of its fraction field. We will now see that an analogous statement holds for definable sets in a strongly minimal theory (in fact, the former is a special case of the latter, once Morley rank is identified with Krull dimension). To do that, we first note that the we may repeat the definition of transcendence degree for an arbitrary pregeometry.

Definition 5.1.11. Let $\langle D, \mathrm{cl}\rangle$ be a pregeometry. $A \subseteq D$ is *independent* if cl is injective on $\mathcal{P}(A)$. A is a *basis* of D if it is independent and $\mathrm{cl}(A) = D$. It follows from Steinitz exchange that all bases of D have the same cardinality, which is called the *dimension* of D, denoted $\dim(D)$.

independent
basis
dimension

Note that by Zorn's lemma, each pregeometry has a basis, hence a well defined dimension. More generally, if $A \subseteq B = \mathrm{cl}(B) \subseteq D$ for some pregeometry $\langle D, \mathrm{cl}\rangle$, and we let $\mathrm{cl}_A(C) = \mathrm{cl}(A \cup C)$ for subsets $C \subseteq B$, we obtain a pregeometry on B. We say that C is a independent, a basis for B, etc over A, if they are such for this pregeometry. Likewise, we denote $\dim_A(C)$ the dimension of C in this geometry (note that dimension and independence are independent of B).

We now have:

Proposition 5.1.12. *Let $\bar{b} \in \mathbb{U}$ in a strongly minimal theory $\mathcal{T}$. Then $\mathcal{MR}(\mathrm{tp}(\bar{b}/A)) = \dim_A(\bar{b})$ for all $A \subseteq \mathbb{U}$.*

Proof. We may assume $A = \emptyset$. Let $\bar{b} = \langle b_1, \dots, b_n\rangle$. We use induction on n. If b_n is algebraic over the previous b_i, both sides do not change by its removal, so the result follows by induction. Hence, we assume the b_i are independent, and prove that the Morley rank is n. By induction $\mathcal{MR}(b_1, \dots, b_{n-1}) = n-1$, and since b_n is not algebraic over them, the Morley rank is at least n. The result follows from the following exercise. $\square$

Exercise 5.1.13. Let $\mathcal{T}$ be strongly minimal. Then $\mathcal{MR}(\mathcal{T}^n) = n$. $\square$

We now return to the task of classifying our examples in combinatorial terms. We need the following definition.

Definition 5.1.14. Let $\langle D, \mathrm{cl}\rangle$ be a pregeometry.

(1) D is a *geometry* if $\mathrm{cl}(\emptyset) = \emptyset$ and points are closed.
(2) D is a *trivial pregeometry* (or geometry) if $\mathrm{cl}(A) = \bigcup_{a \in A} \mathrm{cl}(\{a\})$ for all $A \subseteq D$.

geometry
trivial pregeometry

(3) D is a *modular* if $\dim(A \cup B) + \dim(A \cap B) = \dim(A) + \dim(B)$ for all closed $A, B \subseteq D$. **modular**

locally modular

(4) D is *locally modular* if cl_a is modular for some $a \in D$

End of lecture 18, Jun 6

Exercise 5.1.15. Each pregeometry has a canonical geometry associated to it: For $\langle D, \mathrm{cl}\rangle$ a pregeometry, let $PD = D'/\sim$, where $D' = D \setminus \mathrm{cl}(0)$, and $a \sim b$ if $\mathrm{cl}(a) = \mathrm{cl}(b)$. Then there is an induced closure operator cl' on PD making it a geometry. Each $A \subseteq PD$ is cl'-independent if and only if $\pi^{-1}(A)$ is cl-independent. $\square$

Exercise 5.1.16. Each trivial pregeometry is modular (hint: using the previous exercise, we may assume we have a geometry) $\square$

Exercise 5.1.17. The examples above are classified as follows:

(1) Ex. 5.1.1 is a trivial geometry
(2) Ex. 5.1.2 is a trivial pregeometry which is not a geometry.
(3) Ex. 5.1.3 is a non-trivial modular pregeometry (and not a geometry)
(4) Ex. 5.1.4 is a non-trivial locally modular pregeometry (which is not modular)
(5) Ex. 5.1.5 is a non-trivial locally modular geometry. It is the canonical geometry associated to Ex. 5.1.3.

$\square$

Example 5.1.18. Ex. 5.1.6 is non-locally modular: Let $A = \mathrm{acl}(K(a, b))$, $B = \mathrm{acl}(K(c, ac + b))$, where $a, b, c \in \mathbb{U}$ are independent over $K \subseteq \mathbb{U}$. Note that by quantifier elimination, algebraic closure here coincides with the algebraic notion. Hence, A and B have dimension 2 over K, $A \cup B$ has dimension 3 and $A \cap B$ dimension 0. $\square$

The distinction between examples 5.1.3, 5.1.4 and 5.1.5 is not substantial, and we will view them all as locally-modular pregeometries.

Since there is only one trivial geometry (on a give set), no further classification is possible in terms of the geometry for this class. However, the fact that the geometry is trivial has strong consequences for definable sets. For example, a strongly minimal group cannot be trivial: If a and b are independent elements, then ab belongs to the closure of $\{a, b\}$ but not to the closure of a (otherwise $b = a^{-1}(ab)$ is there as well, contradicting independence). We will come back to this later.

We would like to understand if there are any substantially different examples than the ones we already discussed. To do that, we need to explain more precisely what "substantially different" means, so we make a minor digression.

5.2. Imaginaries.

We quickly review definable equivalence relations, see [11, §8.4] for more details. A definable subset $\mathbf{Y} \subseteq \mathbf{X} \times \mathbf{X}$ is a *definable equivalence relation* if it satisfies the usual (first order) axioms of an equivalence relation. A *quotient* for such a $\mathbf{Y}$ is a surjective definable map $\pi : \mathbf{X} \to \bar{\mathbf{X}}$

definable equivalence relation
quotient

whose kernel $\pi(x_1) = \pi(x_2)$ is $\mathbf{Y}$. We note that these are all first order conditions, so depend only on the theory.

As with sets, any two quotients for the same equivalence relation are uniquely definably isomorphic. However, quotients for definable equivalence relations need not exist. To verify that, it is convenient to reformulate the condition. If $\mathbf{Z}$ is a set definable with parameters, we say it is defined with a *canonical parameter* if it can be defined in the form $\phi(z, c)$, with c the unique parameter defining $\mathbf{Z}$ **canonical parameter**

Proposition 5.2.1. *Let $\mathcal{T}$ be a complete theory with a universal domain $\mathbb{U}$. The following are equivalent:*

(1) *Each definable equivalence relation in $\mathcal{T}$ admits a quotient.*
(2) *Each definable set with parameters is definable with a canonical parameter*
(3) *Each definable set $\mathbf{Z}$ with parameters is definable with a parameter fixed by all $\sigma \in Aut(\mathbb{U})$ which preserve $\mathbf{Z}(\mathbb{U})$ (as a set).*

Proof.

(1) $\implies$ (2): If $\mathbf{Z}$ is defined by $\phi(z, c)$, apply the assumption to the definable equivalence relation $\forall z(\phi(z, u) \leftrightarrow phi(z, v))$. If b is the image of c under the quotient π, $\forall y(\pi(y) = b \to \phi(z, y))$ is a definition of $\mathbf{Z}$ with a canonical parameter.

(2) $\implies$ (1): Apply the assumption to the equivalence classes (i.e., when $\mathbf{Z}$ has the form $\phi(z, c)$, with ϕ defining the given equivalence relation)

(3) $\implies$ (2): Will be proved below, after Ex. 5.2.2

(2) $\implies$ (3): If $\mathbf{Z}$ is defined by $\phi(z, c)$, then $\sigma(\mathbf{Z}(\mathbb{U}))$ is defined by $\phi(z, \sigma(c))$. Hence, if σ preserves $\mathbf{Z}(\mathbb{U})$, then $\phi(z, c)$ is equivalent to $\phi(z, \sigma(c))$. If c is canonical, we must have $\sigma(c) = c$ □

Any theory $\mathcal{T}$ has a canonical extension $\mathcal{T}^{eq}$ the eliminates imaginaries. It is constructed by adding, for each definable equivalence relation $\mathbf{Y}$ on $\mathbf{X}$, a new sort $\mathbf{Q}_{\mathbf{Y}}$ and a new function symbol $\pi_{\mathbf{Y}} : \mathbf{X} \to \mathbf{Q}_{\mathbf{Y}}$, with the theory stating that $\pi_{\mathbf{Y}}$ is a quotient for $\mathbf{Y}$. It is easy to see that this theory admits elimination of imaginaries. Since quotients of equivalence relations exist for usual sets, each model $\mathbb{M}$ of $\mathcal{T}$ has an extension to a model $\mathbb{M}^{eq}$ of $\mathcal{T}^{eq}$, and every elementary map $\sigma : \mathbb{M} \to \mathbb{N}$ between models of $\mathcal{T}$ has a unique extension $\sigma^{eq} : \mathbb{M}^{eq} \to \mathbb{N}^{eq}$. Conversely, the restriction of a model of $\mathcal{T}^{eq}$ to the sorts in $\mathcal{T}$ is a model of $\mathcal{T}$ (in other words, the categories of models are equivalent). In particular, the automorphism groups of $\mathbb{M}$ and $\mathbb{M}^{eq}$ coincide. Many model theoretic properties are shared by $\mathbb{M}$ and $\mathbb{M}^{eq}$. For example, $\mathbb{M}$ is saturated if and only if $\mathbb{M}^{eq}$ is.

If a definable equivalence relation $\mathbf{Y} \subseteq \mathbf{X} \times \mathbf{X}$ admits two quotients $\pi_i : \mathbf{X} \to \mathbf{Z}_i$, there is a unique definable map $\mathbf{Z}_1 \to \mathbf{Z}_2$ commuting with the projections, as for sets. In particular, if $\mathbf{Z}$ is a quotient for $\mathbf{Y}$ in $\mathcal{T}$, each element $a \in \mathbf{Q}_{\mathbf{Y}}(\mathbb{U})$ (often called an *imaginary element*) is inter-definable with an element $b \in \mathbf{Z}(\mathbb{U})$: $\operatorname{dcl}^{eq}(b) = \operatorname{dcl}^{eq}(a)$, where dcl^{eq} is the definable **imaginary element** closure computed in $\mathcal{T}^{eq}$. In particular, if $\mathcal{T}$ eliminates imaginaries, each imaginary element is inter-definable with a "real" one. The converse is also true:

Exercise 5.2.2. A theory $\mathcal{T}$ eliminates imaginaries if and only if for each model $\mathbb{M}$, each $a \in \mathbb{M}^{eq}$ is inter-definable with some $b \in \mathbb{M}$ ("each model" can be replaced by "some model" if $\mathcal{T}$ is complete) □

One advantage of this criterion is that it enables working with automorphisms of a saturated model. For instance, to prove (3) $\implies$ (2) of Prop. 5.2.1, let $\phi(x, b)$ be a definition of $\mathbf{Z}$, and assume that b is fixed by all automorphisms preserving $\mathbf{Z}(\mathbb{U})$. Let a be a canonical parameter for $\mathbf{Z}$ in $\mathcal{T}^{eq}$ (which exists by the direction already proven). If $\sigma(a) = a$ then $\sigma(\mathbf{Z}(\mathbb{U})) = \sigma(\mathbf{Z}(\mathbb{U}))$, so $\sigma(b) = b$. Hence, a is inter-definable with b. By compactness, there is a definable injective function f in $\mathcal{T}^{eq}$, such that $f(a) = b$ and for all c, $f(c)$ and c define the same subset. Since all such c are canonical, so are the $f(c)$, so b is a canonical parameter.

As another application, we obtain an example where EI fails:

Example 5.2.3. Let $\mathcal{T}$ be the theory of an infinite set, with universal domain $\mathbb{U}$. The equivalence relation on $\mathbb{U}^2$ that identifies $\langle a, b\rangle$ with $\langle b, a\rangle$ is definable. Its quotient can be thought of as the set of non-empty subsets of $\mathbb{U}$ of cardinality at most 2. We claim that this quotient is not definable: If it was, the set $C = \{a, b\}$ (for different $a, b \in \mathbb{U}$) would have to be definable over the elements fixed by all automorphisms preserving C. However, we may find an automorphism switching a and b, and moving every other element, so C would have to be definable over the empty set, which is clearly impossible. □

The last example can be generalized. If G is a finite group acting by definable maps on some definable set $\mathbf{X}$, the relation of being in the same orbit is definable, but its quotient $\mathbf{X}/G$ need not be. The last example occurs in this way with $G = C_2$ and $\mathbf{X}$ the set of pairs. More generally, the symmetric group S_n acts on $\mathbf{Y}^n$ for each definable $\mathbf{Y}$ (in any theory) by permuting the coordinates. The quotient, which can be identified with the set non-empty subsets of $\mathbf{Y}$ of size at most n, need not be definable. If it is (for all n), we say that $\mathcal{T}$ *eliminates finite imaginaries*, or *codes finite sets*. Note that if $\mathcal{T}$ codes finite sets, then so does any expansion that does not introduce new sorts. In particular, every theory of fields (possibly with additional structure) eliminates finite imaginaries, by the following observation:

eliminates finite imaginaries
codes finite sets

Proposition 5.2.4. *The (incomplete) theory of fields codes finite sets*

Proof. We need to show that each set $\{a_1, \ldots, a_n\}$ of field elements is inter-definable with a tuple. We may take the coefficients of the unique monic polynomials of degree n whose roots are the a_i □

Using this notion, we may reduce proving EI to the following weaker notion:

Definition 5.2.5. A theory $\mathcal{T}$ admits *weak elimination of imaginaries* if each imaginary element $b \in \mathbb{U}^{eq}$ is definable over a real tuple in its algebraic closure.

weak elimination of imaginaries

Exercise 5.2.6. $\mathcal{T}$ eliminates imaginaries if and only if it admits weak EI and codes finite sets □

Weak EI has the following slightly surprising sufficient condition:

Proposition 5.2.7. *If $\mathcal{T}$ is strongly minimal with* $\mathrm{acl}(0)$ *infinite, then it admits weak EI.*

Proof. Let b be an imaginary element, and let A be the set of real elements in its algebraic closure. We have $\pi(c_1, \dots, c_n) = b$ for some real c_i, and we need to show they can be taken from A. Assume by induction that this is true for $i < k$ and the formula $\phi(y)$ given by $\exists y_{k+1} \dots y_n(\pi(c_1, \dots, c_{k-1}, y, y_{k+1}, \dots, y_n) = b)$ is satisfiable. By strong minimality, it is finite or confinite. In the first case we may choose any solution: it will be algebraic over b since the c_i are. Otherwise, we may choose c_k from $\mathrm{acl}(0)$, since it is infinite by assumption. □

Corollary 5.2.8. *The theory ACF_p eliminates imaginaries for each p.*

The theory DCF is not strongly minimal, but we may deduce

Corollary 5.2.9. *DCF eliminates imaginaries*

We skip the proof, which consists, essentially, of reducing to the case of ACF_0.

End of lecture 19, Jun 7

We say that a theory $\mathcal{T}$ *interprets* a group, field, etc. if such an object is definable in $\mathcal{T}^{eq}$ (in general, a theory $\mathcal{T}$ interprets a theory $\mathcal{T}_0$ if there is a "model" of $\mathcal{T}_0$ in $\mathcal{T}^{eq}$). We saw above that a definable group cannot have a trivial geometry. For the same reason, such a structure cannot interpret a group.

interprets

References

[1] Elisabeth Bouscaren, ed. *Model theory and algebraic geometry.* Lecture Notes in Mathematics 1696. An introduction to E. Hrushovski's proof of the geometric Mordell-Lang conjecture. Berlin: Springer-Verlag, 1998. ISBN: 3-540-64863-1 (cit. on p. 2).

[2] Guy Casale, James Freitag, and Joel Nagloo. "Ax-Lindemann-Weierstrass with derivatives and the genus 0 Fuchsian groups". In: (2018). DOI: `10.48550/ARXIV.1811.06583`. arXiv: `1811.06583` (cit. on p. 2).

[3] Lou van den Dries and K. Schmidt. "Bounds in the theory of polynomial rings over fields. A nonstandard approach". In: *Invent. Math.* 76.1 (1984), pp. 77–91. ISSN: 0020-9910 (cit. on p. 35).

[4] Deirdre Haskell, Anand Pillay, and Charles Steinhorn, eds. *Model theory, algebra, and geometry*. Mathematical Sciences Research Institute Publications 39. Cambridge: Cambridge University Press, 2000. ISBN: 0-521-78068-3. URL: http://www.msri.org/communications/books/Book39/contents.html (cit. on p. 3).

[5] Ehud Hrushovski and Thomas Scanlon. "Lascar and Morley ranks differ in differentially closed fields". In: *J. Symbolic Logic* 64.3 (1999), pp. 1280–1284. ISSN: 0022-4812 (cit. on p. 43).

[6] I. Kaplansky. *An Introduction to Differential Algebra*. 2nd ed. Paris: Hermann, 1976 (cit. on p. 20).

[7] Angus Macintyre, Kenneth McKenna, and Lou van den Dries. "Elimination of quantifiers in algebraic structures". In: *Advances in Mathematics* 47.1 (1983), pp. 74–87. ISSN: 0001-8708. DOI: https://doi.org/10.1016/0001-8708(83)90055-5. URL: https://www.sciencedirect.com/science/article/pii/0001870883900555 (cit. on p. 25).

[8] David Marker. *Model theory: An introduction*. Graduate Texts in Mathematics 217. New York: Springer-Verlag, 2002. ISBN: 0-387-98760-6 (cit. on pp. 3, 35).

[9] David Marker, Margit Messmer, and Anand Pillay. *Model theory of fields*. 2nd ed. Lecture Notes in Logic 5. La Jolla, CA: Association for Symbolic Logic, 2006. ISBN: 1-56881-282-5 (cit. on pp. 3, 20, 41, 45).

[10] J. Robinson and S. Feferman. *The Collected Works of Julia Robinson*. Collected works series. American Mathematical Society, 1996. ISBN: 9780821805756. URL: https://books.google.co.il/books?id=%5C_33D84OENIAC (cit. on p. 20).

[11] Katrin Tent and Martin Ziegler. *A Course in Model Theory*. Lecture Notes in Logic. Cambridge University Press, 2012. DOI: 10.1017/CBO9781139015417 (cit. on pp. 3, 35, 48).

Department of Math, Ben-Gurion University, Be'er-Sheva, Israel
Email address: mailto:kamensky.bgu@gmail.com
URL: https://www.math.bgu.ac.il/~kamenskm