The trichotomy theorem in difference fields

Thomas Scanlon

UC Berkeley

December 2011

Trichotomy theorem

Theorem

Every minimal type relative to the theory ACFA of difference closed fields is either modular or almost internal to a definable minimal fixed field. Moreover, the class of modular minimal types further decomposes into the trivial types and those which are almost internal to a generic of a modular group.

The main goal of today's lecture will be the explanation of the terms appearing in the trichotomy theorem. With the following lectures we shall see parts of three proofs.

The proof of the trichotomy theorem in all characteristics appears in a paper of Chatzidakis, Hrushovski and Peterzil. The first proof of the trichotomy theorem in characteristic zero is due to Chatzidakis and Hrushovski. Pillay and Ziegler gave a more geometric proof in characteristic zero using jet spaces. Each of these proofs yields further information beyond the raw statement of the trichotomy theorem.

Paradigmatic difference equations

- $\mathrm{Fix}(\sigma)$ is the definable set in one variable given by the equation $\sigma(x) = x$.
- Consider G, the definable set, again in one variable, defined by $\sigma(x) = x^2 \,\&\, x \neq 0$.
- Finally, we define T by $\sigma(x) = x^2 + 1$.

The fixed field

So that we may speak of the fixed field as an actual set, we shall work in a specific model $\mathbb{U} \models \mathrm{ACFA}$ of the theory of difference closed field and set $F := \mathrm{Fix}(\sigma)(\mathbb{U})$. Then as F is the subset of $\mathbb{U}$ fixed by the field automorphism σ, it is a subfield.

In fact, $\mathrm{Fix}(\sigma)$ is minimal (definition to follow) and F is pseudofinite.

By Ax's axiomatization of the theory of finite fields, to check that F is pseudofinite we should verify that

- F is perfect,
- F is quasifinite, and
- F is pseudoalgebraically closed.

Fix(σ) is perfect

- If $\mathrm{char}(F) = 0$, then there is nothing to check.
- Suppose now that $\mathrm{char}(F) = p > 0$ and let $a \in F$.
- As $\mathbb{U}$ is algebraically closed, there is some $\alpha \in \mathbb{U}$ with $\alpha^p = a$.
- Applying σ, we have $(\sigma(\alpha))^p = \sigma(\alpha^p) = \sigma(a) = a$. Hence, $\sigma(\alpha)$ is a root of $X^p - a = (X - \alpha)^p$ and must be equal to α. That is, $\alpha \in F$.

$\mathrm{Fix}(\sigma)$ is quasifininite

Strictly speaking, to say that a field is quasifinite includes the assertion that it is perfect. We shall concentrate on proving that $\mathrm{Gal}(F^{\mathrm{alg}}/F) = \widehat{\mathbb{Z}}$.

- From the Galois correspondence, it is clear that $\mathrm{Gal}(F^{\mathrm{alg}}/F)$ is topologically generated by σ (or to be more precise, the restriction of σ to F^{alg}).
- Let $N \in \mathbb{Z}_+$ be a positive integer. Consider $X := \mathbb{A}^N_{\mathbb{U}} = \mathrm{Spec}(\mathbb{U}[x_0, \ldots, x_{N-1}])$. We write X^σ as $\mathrm{Spec}(\mathbb{U}[y_0, \ldots, y_{N-1}])$. Let $\Gamma \subseteq X \times X^\sigma$ be defined by the equations $y_i = x_{i+1}$ for $i < N - 2$ and $x_0 = y_{N-1}$. Let $U \subseteq \Gamma$ be the dense constructible subset defined by the inequalities $x_0 \neq y_i$ for $i < N - 1$.
- By the axioms for ACFA there is a point $a = (a_0, \ldots, a_{N-1}) \in X(\mathbb{U})$ with $(a, \sigma(a)) \in U(\mathbb{U})$. From these equations, we see that the orbit of a_0 under σ has size exactly N. Hence, $[\mathrm{Fix}(\sigma^N)(\mathbb{U}) : \mathrm{Fix}(\sigma)(\mathbb{U})] \geq N$, while the opposite inequality holds on general grounds.

Fix(σ) is PAC

- Let V be an absolutely irreducible variety over F.
- Then $V_{\mathbb{U}}$ is an irreducible variety over $\mathbb{U}$, $(V_{\mathbb{U}})^{\sigma} = V_{\mathbb{U}}$, and the diagonal $\Delta_{V_{\mathbb{U}}} \subseteq (V_{\mathbb{U}}) \times (V_{\mathbb{U}})$ is an irreducible subvariety of the product which projects dominantly in both directions.
- Hence, there is a point $a \in V(\mathbb{U})$ with $(a, \sigma(a)) \in \Delta_V(\mathbb{U})$. That is, $\sigma(a) = a$ so that $a \in V(F)$.

Other fixed fields

If $\text{char}(\mathbb{U}) = p > 0$, then map $\tau : \mathbb{U} \to \mathbb{U}$ given by $x \mapsto x^p$ is a definable field automorphism which commutes with σ.

For what follows, if $\text{char}(\mathbb{U}) = 0$, then we define $\tau(x) := x$

For any pair of integers $m, n \in \mathbb{Z}$ the map $\rho := \sigma^n \tau^m$ is, thus, a definable field automorphism of $\mathbb{U}$ and the fixed field of ρ

$$\text{Fix}(\sigma^n \tau^m)(\mathbb{U}) := \{x \in \mathbb{U} \ : \ \sigma^n(x) = \tau^{-m}(x)\}$$

is a definable subfield of $\mathbb{U}$ which is pseudofinite provided that $n \neq 0$. Moreover, every definable field is definably isomorphic to $\text{Fix}(\sigma^n \tau^m)$ for some choice of n and m.

Induced structure on G defined by $\sigma(x) = x^2$ and $x \neq 0$

If $\mathrm{char}(\mathbb{U}) = 2$, then G is the set of nonzero elements of the fixed field of $\rho = \sigma\tau^{-1}$, and, thus, has the structure of a pseudofinite field (with zero removed).

If $\mathrm{char}(\mathbb{U}) \neq 2$, then G is a subgroup of the multiplicative group having the property that every definable subset of G^n (for any $n \in \mathbb{Z}_+$) is a finite Boolean combination of cosets of subgroups. Let us prove this for quantifier-free definable sets.

Quantifier-free induced structure on G, $\mathrm{char}(\mathbb{U}) \neq 2$

- If $X \subseteq G^n$ is a quantifier-free definable set, then it is a finite Boolean combination of sets defined by difference equations of the form

$$f(x_1, \ldots, x_n; \sigma(x_1), \ldots, \sigma(x_n); \ldots; \sigma^\ell(x_1), \ldots, \sigma^\ell(x_n)) = 0$$

  for some polynomial f in $n(\ell + 1)$ variables.
- Using the fact that $\sigma(x) = x^2$ on G, we may replace such a difference equation with a strictly algebraic equation

$$f(x_1, \ldots, x_n; x_1^2, \ldots, x_n^2; \ldots; x_1^{2^\ell}, \ldots, x_n^{2^\ell}) = 0$$

- After having reduced to such algebraic equations, we may assume that X is the intersection of an irreducible algebraic variety in which X is Zariski dense.

Quantifier-free induced structure on G, char$(\mathbb{U}) \neq 2$, continued

- From our reduction, we conclude that for any m that $\overline{X}^{\sigma^m} = [2^m]_{\mathbb{G}_m{}^n}(\overline{X})$ where we have written $[2^m]_{\mathbb{G}_m{}^n}$ for the selfmap of algebraic groups given by $(x_1, \ldots, x_n) \mapsto (x_1^{2^m}, \ldots, x_n^{2^m})$.
- Let $d := \dim(X)$. Then $2^{m(n-d)} \deg(\overline{X}) = 2^{m(n-d)} \deg(\overline{X}^{\sigma^m}) = \deg([2^m]_{\mathbb{G}_m{}^n}^{-1}(\overline{X}^{\sigma^m}) = \deg(\bigcup_{\xi \in \mu_{2^m}^n} \xi\overline{X}) = (2^{mn}/\#\{\xi \in \mu_{2^m}^n \ : \ \xi\overline{X} = \overline{X}\})\deg(X)$.
- Hence, if we set $\mathrm{Stab}(\overline{X}) := \{\gamma \in \mathbb{G}_m{}^n \ : \ \gamma X = X\}$, we have $\#\mathrm{Stab}(\overline{X})(\mathbb{U}) \cap \mu_{2^m}^n = 2^{md}$. Thus, $\dim(\mathrm{Stab}(\overline{X})) = d$ and $\overline{X}$ is a coset of its stabilizer.

The structure on T defined by $\sigma(x) = x^2 + 1$

If $\mathrm{char}(\mathbb{U}) = 2$, then T is a definable principal homogenous space for the field $\mathrm{Fix}(\sigma\tau^{-1})$ under addition. However, there is no definable function $\mathrm{Fix}(\sigma\tau^{-1}) \to T$ defined without parameters as, for example, $\mathrm{Fix}(\sigma\tau^{-1})$ has definable elements (consider 0) but neither 0 nor 1 belongs to T.

Regardless of the characteristic of $\mathbb{U}$, since T is defined over $\mathrm{Fix}(\sigma)$, the restriction of σ to T gives a definable function $\sigma : T \to T$. If $\mathrm{char}(\mathbb{U}) \neq 2$, then there is no more structure.

Theorem

If $\mathrm{char}(\mathbb{U}) \neq 2$, *then every definable (with parameters) subset of* T^m *is definable (with parameters from* T*) in the structure* (T, σ).

It would be an interesting exercise to prove this theorem directly, but we shall do so only as a consequence of the general trichotomy theorem combined with a theorem of Ritt on polynomial decompositions.

- As we have seen, ACFA does not admit outright quantifier elimination, as, for example, the set $\{x \in \mathrm{Fix}(\sigma) \ : \ (\exists y) y^2 = x\}$ is not quantifier-free definable outside of characteristic two. Consequently, to describe the full induced structure on a definable set it does not suffice to consider only quantifier-free formulas.
- In characteristic zero, every modular minimal type is stable and stably embedded. Thus, to describe the induced structure it suffices to describe the induced structure over a small set of parameters over which the type is defined.
- In positive characteristic, there are modular minimal types which are not stably embedded. A simple example is given by the equation $\sigma(x) = x^p - x$.

Some notational conventions

- $\mathbb{U}$ is a fixed difference closed field, assumed to be saturated and much larger than any other difference field or parameter set we shall mention. All sets of parameters will be tacitly assumed to be small subsets of $\mathbb{U}$.
- Often, a tuple will be written without an explicit reference to its components or to its length. We may write expressions like $a \in A$ to mean that a is a finite tuple whose components are elements of the set A.
- We may write expressions such as $p \in S_x(A)$ to mean that p is a complete type in the variable x (which again may be properly a finite tuple of variables or a variable ranging over some interpretable set) over the parameter set A.

σ-degree

Definition

Given an inversive difference field K and a tuple a (from $\mathbb{U}$) we define $K(a)_\sigma$ to be the inversive difference field generated by a over K, which as a field, is $K(\{\sigma^j(a_i) \ : \ j \in \mathbb{Z}, j \leq n\})$ where $a = (a_1, \ldots, a_n)$.

Definition

We define $\deg_\sigma(a/K) := \operatorname{tr.deg}(K(a)_\sigma/K)$. If $\deg_\sigma(a/K)$ is infinite, then it is necessarily $\aleph_0$, but we shall simply write $\deg_\sigma(a/K) = \infty$. Over a general set B, we define $\deg_\sigma(a/B)$ to be $\deg_\sigma(a/K)$ where K is the inversive difference field generated by B.

Remark

Since the isomorphism type of $K(a)_\sigma$ over K, we may regard $\deg_\sigma(a/K)$ as a function of $\mathrm{qftp}(a/K)$, the quantifier-free type of a over K, or equivalently, of

$$I(a/K) := \{f \in K[x]_\sigma \ : \ f(a) = 0\}$$

the ideal of difference polynomials over K which vanish on a.

Definition

For a definable set X defined over some set of parameters K we define

$$\deg_\sigma(X) := \sup\{\deg_\sigma(a/K) \ : \ a \in X(\mathbb{U})\}$$

With this definition we are implicitly asserting that if X is also defined over L, then $\deg_\sigma(X) = \sup\{\deg_\sigma(a/L) \ : \ a \in X(\mathbb{U})\}$.

Independence of σ-degree on field of definition

- We may reduce immediately to the case that $K \subseteq L$ and both L and K are inversive difference fields. We should then show that if there is some $a \in X(\mathbb{U})$ with $\deg_\sigma(a/K) \geq n$, then there is some $b \in X(\mathbb{U})$ with $\deg_\sigma(b/L) \geq n$.
- To simplify the presentation, I shall assume that a is a singleton leaving the reduction to this case as an exercise.
- The partial type $\{x \in X\} \cup \{f(x, \sigma(x), \ldots, \sigma^{n-1}(x)) \neq 0 \ : \ f \in L[x_0, \ldots, x_{n-1}] \ : \ f \not\equiv 0\}$ is consistent as otherwise (by taking products) there would be some $f \in L[x_0, \ldots, x_{n-1}]$ for which $\mathrm{ACFA} \vdash x \in X \rightarrow f(x, \sigma(x), \ldots, \sigma^{n-1}(x)) = 0$.
- Applying any automorphism ρ of $\mathbb{U}$ over K, we see that $\mathrm{ACFA} \vdash x \in X \rightarrow f^\rho(x, \ldots, \sigma^{n-1}(x)) = 0$. Taking conjunctions and using Noetherianity, we would find a an order $< n$ difference equation over the definable closure of K (which is contained in the algebraic closure of K in the usual algebraic sense) which is implied by $x \in X$. Taking a product of the finitely many conjugates of this equation over K, we contradict the hypothesis that $\deg_\sigma(a/K) \geq n$.

Definition

For an extension $K \subseteq L$ of inversive difference fields and a tuple a, we say that a is free from L over K, written $a \downarrow_K L$ if $K(a)_\sigma$ is algebraically independent from L over K.

- Visibly, the relation $a \downarrow_K L$ depends only on $\mathrm{qftp}(a/L)$.
- While we have defined independence algebraically, it is an instance of general non-forking independence.
- Since algebraic independence is insensitive to algebraic extensions, we could replace K and L with their algebraic closures without changing independence.
- If A, B and C are general sets then we define $A \downarrow_B C$ if for each finite tuple a from A we have $a \downarrow_K L$ where K is the inversive difference field generated by B and L is the inversive difference field generated by $B \cup C$.

Lascar rank

We define Lascar rank, or SU-rank (or even just U-rank) for a a tuple from $\mathbb{U}$ and $B \subseteq \mathbb{U}$ a small subset by transfinite recursion.

- $\mathrm{SU}(a/B) \geq 0$ always.
- $\mathrm{SU}(a/B) \geq \alpha + 1$ if there is some C so that $a \not\perp_B C$ and $\mathrm{SU}(a/C) \geq \alpha$.
- $\mathrm{SU}(a/B) \geq \lambda$ for λ a limit ordinal just in case $\mathrm{SU}(a/B) \geq \alpha$ for all $\alpha < \lambda$.

Definition

We say that $p(x) = \mathrm{tp}(a/B)$ is minimal if $\mathrm{SU}(a/B) = 1$.

Comparison of $\deg_\sigma$ and SU

Theorem

Let a be a tuple and K a small inversive difference field.

- $\mathrm{SU}(a/K) \leq \deg_\sigma(a/K)$
- $\mathrm{SU}(a/K) < \omega \iff \deg_\sigma(a/K) < \infty$
- $\mathrm{SU}(a/K) = 0 \iff \deg_\sigma(a/K) = 0 \iff K(a) \subseteq K^{\mathrm{alg}}$

Remark

Strictly speaking, the inequality $\mathrm{SU}(a/K) \leq \deg_\sigma(a/K)$ is not meaningful when either side is infinite as $\mathrm{SU}(a/K)$ is an ordinal and $\deg_\sigma(a/K)$ is a cardinal. We shall understand the inequality to mean only that $\deg_\sigma(a/K)$ is infinite if $\mathrm{SU}(a/K)$ is infinite.

Proof that $\mathrm{SU}(a/K) \leq \deg_\sigma(a/K)$

- Working by induction on n we show that for any such pair (a, K) if $\mathrm{SU}(a/K) \geq n$, then $\deg_\sigma(a/K) \geq n$.
- The base case of $n = 0$ is trivial.
- If $\mathrm{SU}(a/K) \geq n + 1$, then we may find an inversive difference field L with $a \not\downarrow_K L$ and $\mathrm{SU}(a/L) \geq n$.
- By induction, $\deg_\sigma(a/L) \geq n$.
- By definition of $\not\downarrow$, if $\mathrm{tr.deg}(K(a)_\sigma/K) < \infty$, then $\mathrm{tr.deg}(L(a)_\sigma/L) < \mathrm{tr.deg}(K(a)_\sigma/K)$. Thus, $\deg_\sigma(a/K) \geq n + 1$.

Proof that $\mathrm{SU}(a/K) = 0 \iff \deg_\sigma(a/K) = 0 \iff a \in K^{\mathrm{alg}}$

- From the definition of $\deg_\sigma(a/K) := \mathrm{tr.deg}(K(a)_\sigma/K)$ it is clear that $\deg_\sigma(a/K) = 0 \iff a \in K^{alg}$. (Note: K is an inversive difference field! If $a \in K^{\mathrm{alg}}$, then $\sigma^n(a) \in K^{\mathrm{alg}}$ for all n.)
- We have already shown that $\deg_\sigma(a/K) = 0 \implies \mathrm{SU}(a/K) = 0$.
- If $a \notin K^{\mathrm{alg}}$, then $K(a)_\sigma \not\perp_K K(a)_\sigma$ but $\mathrm{SU}(a/K(a)_\sigma) \geq 0$. Hence, $\mathrm{SU}(a/K) \geq 1$.

Lemma: Additivity of $\deg_\sigma$ and SU-rank

It is clear from additivity of transcendence degree that for tuples a and b and inversive difference fields K one has
$\deg_\sigma(ab/K) = \deg_\sigma(a/K) + \deg_\sigma(b/K(a)_\sigma)$.
The corresponding additivity result for SU-rank goes under the name of the Lascar inequalities (which reduce to outright equalities when all of the quantities in question are finite) and holds in all generality for SU-rank.

Theorem

Given tuples a and b and a small parameter set C, one has

$$\mathrm{SU}(a/Cb) + \mathrm{SU}(b/C) \leq \mathrm{SU}(ab/C) \leq \mathrm{SU}(a/Cb) \oplus \mathrm{SU}(b/C)$$

where "$\oplus$" denotes Cantor's natural sum of ordinals.

Monotonicity lemmata about SU-rank

Lemma

- *For tuples a and b and any small parameter set C one has $\mathrm{SU}(b/C) \leq \mathrm{SU}(ab/C)$.*
- *For a tuple a and parameter sets $B \subseteq C$ one has $\mathrm{SU}(a/C) \leq \mathrm{SU}(a/B)$.*

- We argue by induction on α that if $\mathrm{SU}(b/C) \geq \alpha$, then $\mathrm{SU}(ab/C) \geq \alpha$ with the base case of $\alpha = 0$ being trivial and the limit case immediate.
- If $\mathrm{SU}(b/C) \geq \alpha + 1$ witnessed by some $D \supseteq C$ with $\mathrm{SU}(b/D) \geq \alpha$ and $b \not\perp_C D$, then by induction we have $\mathrm{SU}(ab/D) \geq \alpha$ while clearly $ab \not\perp_C D$. Hence, $\mathrm{SU}(ab/C) \geq \alpha + 1$ as required.
- We leave the second monotonicity assertion as an exercise.

Working by induction on α we show that if $\mathrm{SU}(b/C) \geq \alpha$, then $\mathrm{SU}(a/Cb) + \alpha \leq \mathrm{SU}(ab/C)$.

- The case of $\alpha = 0$ is the content of the lemma from the last slide and limit case is immediate.
- If $\mathrm{SU}(b/C) \geq \alpha + 1$ witnessed by $D \supseteq C$ with $b \not\downarrow_C D$ and $\mathrm{SU}(b/D) \geq \alpha$, then by induction we have $\mathrm{SU}(ab/D) \geq \mathrm{SU}(a/Db) + \alpha$ and visibly $ab \not\downarrow_C D$ so that $\mathrm{SU}(ab/C)$.

Now working by induction on α we show that if $\mathrm{SU}(ab/C) \geq \alpha$, then $\mathrm{SU}(a/Cb) \oplus \mathrm{SU}(b/C) \geq \alpha$.

- Again, the case of $\alpha = 0$ is trivial and the limit case is immediate.
- Suppose that $D \supseteq C$ witnesses $\mathrm{SU}(ab/C) \geq \alpha + 1$ in the sense that $ab \not\perp_C D$ and $\mathrm{SU}(ab/D) \geq \alpha$.
- By induction, $\mathrm{SU}(a/Db) \oplus \mathrm{SU}(b/D) \geq \alpha$.
- From the dependence, we conclude that either $a \not\perp_{Cb} Db$ or $b \not\perp_C D$. Thus, either the righthand side of our inequality is ∞ or $\mathrm{SU}(a/Db) < \mathrm{SU}(a/Cb)$ or $\mathrm{SU}(b/D) < \mathrm{SU}(b/C)$. In any case, we see that $\mathrm{SU}(a/Cb) \oplus \mathrm{SU}(b/C) \geq \alpha + 1$.

Proof that $\mathrm{SU}(a/K) \geq \omega \iff \deg_\sigma(a/K) = \infty$

- We already know that $\mathrm{SU}(a/K) \geq \omega \Longrightarrow \deg_\sigma(a/K) = \infty$.
- Writing $a = (a_1, \ldots, a_n)$ as an n-tuple of elements of $\mathbb{U}$, using additivity of $\deg_\sigma$, we see that $\deg_\sigma(a_i/K(a_1, \ldots, a_{i-1})_\sigma) = \infty$ for some $i \leq n$. Hence, we may assume that $n = 1$.
- Using the fact that $\mathrm{Fix}(\sigma^N)$ is an N-dimensional vector space over $\mathrm{Fix}(\sigma)$, one sees that $\mathrm{SU}(\mathrm{Fix}(\sigma^N)) \geq N$.
- It follows that upon setting $b_N := \sigma^N(a) - a$ one has $\mathrm{SU}(a/b_N) \geq N$.

Minimality and σ-degree

Recall that $\mathrm{tp}(a/K)$ is minimal if $\mathrm{SU}(a/K) = 1$.

- We have seen that $\deg_\sigma(a/K) = 1 \Longrightarrow \mathrm{SU}(a/K) = 1$.
- It may happen that $\mathrm{tp}(a/K)$ is minimal, but $\deg_\sigma(a/K) > 1$. Consider, for example, a satisfying $\sigma^2(a) = a^3$ with $a \notin K^{\mathrm{alg}}$.
- However, $\mathrm{SU}(a/K) = 1$ does imply that $\deg_\sigma(a/K) < \infty$. We shall use this fact to give a geometric criterion for minimality.

Definition

Let K be a difference field. By a pre-σ-variety over K we mean a pair (X, Γ) consisting of an algebraic variety X over K and a subvariety $\Gamma \subseteq X \times X^\sigma$. We say that (X, Γ) is a σ-variety if X and Γ are irreducible and the projection maps $\Gamma \to X$ and $\Gamma \to X^\sigma$ are dominant. We say that (X, Γ) is finitary if the projection maps restricted to Γ are finite.

- Sometimes the term σ-variety is reserved for the case where Γ is the graph of a regular map $f : X \to X^\sigma$.
- It makes sense to relax the requirement that X be a variety, allowing for example σ-schemes.
- We shall consider σ-varieties generically. As such, the finitary hypothesis may be relaxed to require only that the projections are generically finite.

$(X, \Gamma)^\sharp$

Definition

Given a σ-variety (X, Γ) over a difference field K and an extension difference field L, we define

$$(X, \Gamma)^\sharp(L) := \{a \in X(L) \ : \ (a, \sigma(a)) \in \Gamma(L)\}$$

The geometric axioms for ACFA assert that if K is difference closed and (X, Γ) is a σ-variety over K for which Γ is irreducible, then $(X, \Gamma)^\sharp(K)$ is Zariski-dense in X and indeed the set $\{(a, \sigma(a)) \ : \ a \in (X, \Gamma)^\sharp(K)\}$ is Zariski-dense in Γ.

Proposition

Every quantifier-free definable set D is in definable bijection (via a quantifier-free definable function) with a finite Boolean combination of sets of the form $(X,\Gamma)^\sharp$ for some σ-variety (X,Γ). If $\deg_\sigma(D) < \infty$, then (X,Γ) may be taken to be finitary.

- We may express D as the solutions to some quantifier-free formula $\phi(x) = \phi(x_1,\ldots,x_m)$ which we may further express as $\widetilde{\phi}(x,\sigma(x),\ldots,\sigma^n(x))$ for some quantifier-free formula in the language of rings.
- Since every quantifier-free formula in the language of rings is a finite Boolean combination of equations, we may assume that $\widetilde{\phi}$ is given by a polynomial equation, $F(y_0,y_1,\ldots,y_n) = 0$. Moreover, since σ is invertible, we may assume that both y_0 and y_n appear in F.
- The map $x \mapsto (x,\sigma(x),\ldots,\sigma^{n-1}(x))$ induces a bijection between D and $(\mathbb{A}^{nm},\Gamma)^\sharp$ where Γ is defined by $y_i' = y_{i+1}$ (for $i < n$) together with $F(y_0,\ldots,y_{n-1},y_{n-1}') = 0$.

If (X, Γ) is an absolutely irreducible σ-variety over K, then by a generic type in (X, Γ) we mean any complete (quantifier-free) type extending the partial type determined by

- $x \in (X, \Gamma)^\sharp$
- $(x, \sigma(x)) \notin Y$ for each proper closed subvariety (not necessarily absolutely irreducible) $Y \subseteq \Gamma$ defined over K.

If K is an inversive difference field and a is a tuple from $\mathbb{U}$ for which $\mathrm{tp}(a/K)$ is minimal, then $\deg_\sigma(a/K) < \infty$ so that $\mathrm{qftp}(a/K)$ is interdefinable with a generic type of a finitary σ-variety $(X, \Gamma)^\sharp$.

Question

For which σ-varieties (X, Γ) is a generic type of $(X, \Gamma)^\sharp$ minimal?

A geometric rank

We define a rank GU on σ-varieties with the usual conditions that $\mathrm{GU}(X, \Gamma) \geq 0$ as long as $X \neq \varnothing$ (we could define $\mathrm{GU}(\varnothing, \varnothing) := -1$) and $\mathrm{GU}(X, \Gamma) \geq \lambda$ a limit provided that $\mathrm{GU}(X, \Gamma) \geq \alpha$ for all $\alpha < \lambda$. For the successor case, we require the notion of a normal family of sub-σ-varieties.

- A morphism of σ-varieties $g : (Y, \Xi) \to (X, \Gamma)$ is given by a map of algebraic varieties $g : Y \to X$ for which the restriction of $(g \times g^\sigma)$ to Ξ maps to Γ. Note that g takes $(Y, \Xi)^\sharp$ to $(X, \Gamma)^\sharp$.
- A sub-σ-variety $(Y, \Xi) \subseteq (X, \Gamma)$ of a σ-variety is a σ-variety for which the inclusion map $\iota : Y \to X$ induces a morphism of σ-varieties.
- Observe that if $Y \subseteq X$ is any subvariety, then $(Y, (Y \times Y^\sigma) \cap \Gamma)$ is sub-pre-σ-variety of (X, Γ), but it is possible that the projection maps from $(Y \times Y^\sigma) \cap \Gamma$ are not dominant.
- In the special case that Γ is the graph of a regular function $f : X \to X^\sigma$, then (Y, Ξ) is a sub-σ-variety just in case the restriction of f to Y maps Y to Y^σ. In particular, the sub-σ-varieties of (X, Γ) are determined by certain subvarieties of X.
- Note that a point $\{a\} \subseteq X$ can be the underlying variety of a σ-variety of (X, Γ) only if $a \in (X, \Gamma)^\sharp(K)$.

Families of sub-σ-varieties

- By a family of sub-σ-varieties of a σ-variety (X, Γ) we mean a sub-σ-variety $(Y, \Upsilon) \subseteq (X, \Gamma) \times (B, \Xi)$ of the base change of (X, Γ) over some σ-variety (B, Ξ).
- For $b \in (B, \Xi)^\sharp$ the (projection to X) of the fibre (Y_b, Υ_b) is a sub-σ-variety of (X, Γ).
- We say that the family is normal if for distinct $b, b' \in B$ one has $(Y_b, \Upsilon_b) \neq (Y_{b'}, \Upsilon_{b'})$.
- We say that the family is generically covering if for a generic $a \in (X, \Gamma)^\sharp(\mathbb{U})$ there is some generic $b \in (X, \Xi)^\sharp(\mathbb{U})$ for which $a \in (Y_b, \Upsilon_b)^\sharp(\mathbb{U})$.
- For a normal family, we define dimension to be the dimension of B.

Definition of GU completed

We defined $\mathrm{GU}(X,\Gamma)$ for (X,Γ) a σ-variety by recursion.

- $\mathrm{GU}(\varnothing,\varnothing) = -1$
- $\mathrm{GU}(X,\Gamma) \geq 0$ provided X is nonempty.
- $\mathrm{GU}(X,\Gamma) \geq \alpha + 1$ if there is some positive dimensional, normal generically covering family $(Y,\Upsilon) \subseteq (X,\Gamma) \times (B,\Xi)$ of sub-σ-varieties of (X,Γ) for which $\mathrm{GU}(Y_b,\Upsilon_b) \geq \alpha$ for generic $b \in (B,\Xi)^\sharp(\mathbb{U})$.
- $\mathrm{GU}(X,\Gamma) \geq \lambda$ for λ a limit ordinal just in case $\mathrm{GU}(X,\Gamma) \geq \alpha$ for all $\alpha < \lambda$.

Proposition

If $\mathrm{qftp}(a/K)$ *is a generic type of the* σ*-variety* (X,Γ)*, then* $\mathrm{SU}(a/K) = \mathrm{GU}(X,\Gamma)$*.*

Definition

For a tuple a and inversive difference field K, we define the canonical base of $\mathrm{tp}(a/K)$, $\mathrm{Cb}(a/K)$, to be the smallest inversive difference field over which $I_\sigma(a/K)$ is defined.

- In stable theories, one defines canonical bases only for stationary types. In a general simple theory, the canonical base makes sense only for Lascar strong types.
- If $\mathrm{qftp}(a/K)$ is a generic type of a σ-variety (X, Γ), then $\mathrm{Cb}(a/K)$ has the same algebraic closure as the field of definition of (X, Γ).
- We sometimes wish to treat a given inversive difference field L as a common base. In that case, we define $\mathrm{Cb}_L(a/K)$ to be the compositum of L and $\mathrm{Cb}(a/K)$.

A geometric interpretation of the canonical base

- Suppose now that $\mathrm{qftp}(a/K)$ is a generic type of the σ-variety (X, Γ).
- Let k be the prime field and set $\widetilde{X} := \mathrm{loc}(a/k) = V(I(a/K))$ and $\widetilde{\Gamma} := \mathrm{loc}((a, \sigma(a))/k)$.
- Then (X, Γ) is a σ-subvariety of $(\widetilde{X}_K, \widetilde{\Gamma}_K)$.
- Realize (X, Γ) as the base change to K of a generic fibre (Z_b, Υ_b) of some family $(Z, \Upsilon) \to (B, \Xi)$ of sub-σ-varieties of $(\widetilde{X}, \widetilde{\Gamma})$.
- Then $\mathrm{Cb}(a/K)$ is interalgebraic with the function field $k(B)$ for (B, Ξ) irreducible of minimal possible rank.

Definition

Let $p(x)$ be a not necessarily complete type over some small inversive difference field K. We say that p is modular if for any finite tuple a of realizations of p and any small inversive difference field L extending K we have $\mathrm{Cb}_K(a/L) \subseteq K(a)_\sigma^{\mathrm{alg}}$.

- What is called modularity here is one of several related notions, such as weak normality, one-basedness, matroid modularity, linearity, etc.
- Modularity has a clean geometric interpretation for types of finite σ-degree (an, in fact, every modular type has finite σ-degree). If p is a generic type of the finitary difference variety (X, Γ), then p is modular just in case whenever $Y \subseteq (X, \Gamma)^n \times (B, \Xi)$ is a generically covering normal family of sub-σ-varieties of some power of (X, Γ), then $\dim(Y) \leq n \dim(X)$.

Examples of modular types: $\sigma(x) = x^2$

Let K be a small inversive difference field with $\text{char}(K) \neq 2$ and $a \in \mathbb{U} \smallsetminus K^{\text{alg}}$ be some element satisfying $\sigma(x) = x^2$. Then $\text{tp}(a/K)$ is modular.

- $\text{qftp}(a/K)$ is a generic type of $(\mathbb{G}_m, x \mapsto x^2)$
- We say that every irreducible σ-subvariety $(Y, \Gamma) \subseteq (\mathbb{G}_m, x \mapsto x^2)^n$ is a translate of an algebraic subgroup of ${\mathbb{G}_m}^n$.
- Thus, if $b = (b_1, \ldots, b_n)$ is an n-tuple of realizations of $\text{qftp}(a/K)$, L is an inversive difference field extending K and $Y =: \text{loc}(b/L) \subseteq {\mathbb{G}_m}^n$, then $Y = bG$ for some algebraic torus $G \subseteq {\mathbb{G}_m}^n$.
- Then Y is defined over $K(b) \subseteq K(b)_\sigma$. Hence, $I_\sigma(b/L)$ is defined over $K(b)^{\text{alg}}$ so that $\text{Cb}_K(b/L) \subseteq K(b)^{\text{alg}}$.

A trivial example of a modular type

If $\text{char}(K) \neq 2$, then the induced structure on the set T defined by $\sigma(x) = x^2 + 1$ is definable in the structure (T, σ). In this structure, one sees easily that all families of definable come from varying a coordinate so that from a single point in the set one can compute a field of definition.

Definition

A definable set X is internal to the definable set Y if there is a definable function f (possibly defined over parameters not appearing in the definitions of X and Y) from some Cartesian power of X onto Y. We say that Y is almost-internal to X if there is a definable set $Z \subseteq X^n \times Y$ (again, we allow new parameters for the definition of Z) so that the projection of Z to Y is all of Y and for each $a \in X^n$ the fibre Z_a is finite.

In many cases of interest, the internalizing function $f : X^n \to Y$ requires new parameters for its definition. (Consider, for example, X defined by $\sigma(x) = x + 1$, $Y = \mathrm{Fix}(\sigma)$ and K the prime field.)

- (Almost-)internality might be defined at the level of (quantifier-free) types. We would have that $\mathrm{tp}(a/K)$ is internal to $\mathrm{tp}(b/K)$ if there is some L extending K and $b_1, \ldots, b_n$ realizing $\mathrm{tp}(b/K)$ for which $a \downarrow_K L$, $b_i \downarrow_K L$, and a is definable from $L(b_1, \ldots, b_n)$. For almost internality we would ask only that $a \in L(b_1, \ldots, b_n)_\sigma^{\mathrm{alg}}$.
- One could refine the definition of internality for quantifier-free definable sets/quantifier-free types to ask that the internalizing function also be quantifier-free definable. In general, even for quantifier-free definable X and Y this is a proper refinement.
- Almost-internality is closely related to the notion of non-orthogonality: We say that $\mathrm{tp}(a/K)$ and $\mathrm{tp}(b/K)$ are orthogonal, written $\mathrm{tp}(a/K) \perp \mathrm{tp}(b/K)$, if for any L extending K for which $a \downarrow_K L$ and $b \downarrow_K L$ we have $a \downarrow_L b$. If $\mathrm{tp}(b/K)$ is minimal, then it is non-orthogonal to some type $p(x)$ just in case it is almost internal to p.

Recalling the trichotomy theorem

We have now defined all of the terms of the first part of the trichotomy theorem.

Theorem

Every non-modular minimal type relative to ACFA *is almost internal to a minimal fixed field* $\mathrm{Fix}(\rho)$.

For the second cut in the trichotomy theorem we need a precise definition of the notion of triviality and a description of modular groups.

Triviality

Definition

We say that the type $p(x) \in S_x(A)$ is trivial if for any finite sequence $a_1, \ldots, a_n$ of realizations of p if $a_i \downarrow_A a_j$ for all $i < j$, then $a_i \downarrow_A \{a_j \ : \ i \neq j\}$ for all $i < n$. That is, a type is trivial if all dependencies on realizations of p are essentially binary.

Proposition

A trivial minimal type is modular.

Theorem

Every non-trivial modular minimal type is non-orthogonal to a modular group.

- If a definable group G is modular, then every (relatively) quantifier-free definable subset of any Cartesian power is a finite Boolean combination of cosets of definable subgroups.
- This theorem may be proven as a special case of the group configuration theorem asserting under appropriate stability theoretic hypotheses that the presence of a definable group may be detected solely from dependence-theoretic data.

- Working over parameters, the non-triviality of a minimal type p may be detected by the existence of a_1, a_2 and a_3 realizing p which are pairwise independent but for which a_3 is algebraic over $\{a_1, a_2\}$.
- We may regard a_2 as a code for a many-to-many correspondence $f_{a_2} : p \to p$.
- Composing such correspondences as a_2 varies through the realizations of p we obtain an ostensibly two-parameter family of p-curves, but by modularity this family must actually be one dimensional.
- It follows that the resulting family of maps is closed under composition and thus forms a group non-orthogonal to p.

Some complications in the proof of the modular dichotomy theorem

- Since the correspondence f_{a_2} is not actually a function, one must first reduce to the case that it is a function by replacing the elements a_1, a_2 and a_3. For example, one begins by replacing a_3 with its finite set of conjugates over a_1 and a_2.
- The proof outlined above uses stability in some crucial ways. The theory ACFA is not stable, but working with quantifier-free formulas, one has enough stability to push the proof through.

Canonical base property

Theorem

If K is any inversive difference field of characteristic zero, a and c are tuples for which $\deg_\sigma(a/K) < \infty$ and $K(c)_\sigma^{\mathrm{alg}} = \mathrm{Cb}_K(a/K(c)_\sigma))^{\mathrm{alg}}$, then $\mathrm{qftp}(c/K(a)_\sigma)$ is internal to $\mathrm{Fix}(\sigma)$.

Dichotomy theorem as a consequence of the canonical base property

Theorem

If K is any inversive difference field of characteristic zero and p is a minimal type over K, then either p is modular or p is almost internal to $\mathrm{Fix}(\sigma)$.

- If p were non-modular, then there would exist some tuple a of realizations of p and tuple b so that $K(b)_\sigma^{\mathrm{alg}} = \mathrm{Cb}_K(a/K(b)_{\mathrm{alg}})$ but $b \notin K(a)_\sigma^{\mathrm{alg}}$.
- By the canonical base property theorem, $\mathrm{tp}(b/K)$ is internal to $\mathrm{Fix}(\sigma)$.
- Since b is not algebraic over K, $\mathrm{tp}(b/K)$ is non-orthogonal to p. Hence, p is non-orthogonal to (and, thus, almost internal to) $\mathrm{Fix}(\sigma)$.

Geometric reformulation of canonical base property

- As the hypotheses and conclusions are insensitive to replacements of a and c by interdefinable tuples, we may assume that $\mathrm{qftp}(a/K)$ is a generic type of a finitary σ-variety (X, Γ) and that $\mathrm{qftp}(a/K(c)_\sigma)$ is a generic type of a sub-σ-variety (Y_c, Υ_c) where $(Y, \Upsilon) \subseteq (X, \Gamma) \times (C, \Xi)$ is a normal, irreducible family of sub-σ-varieties for which $\mathrm{qftp}(c/K)$ is a generic type of (C, Ξ).
- The theorem asserts that in reversing the roles of X and C, so that now (Y, Υ) gives a family of sub-σ-varieties of (C, Ξ), the fibre above a may be identified with a quotient of the $\mathrm{Fix}(\sigma)$-points of some algebraic variety.
- In fact, we shall realize this fibre as a certain definable subset of a Grassmannian.

Definition

For X an algebraic variety over a field K, $a \in X(K)$ a rational point on X, and $n \in \mathbb{Z}_+$ a positive integer, we define

$$\mathrm{Jet}_n X_a(K) := \mathrm{Hom}(\mathfrak{m}_{X,a}/\mathfrak{m}_{X,a}^{n+1}, K)$$

to be the n^{th} jet space of X at a.

- $\mathrm{Jet}_n X_a$ is represented by a linear space $\mathrm{Jet}_n X \to X$ over X and the association from X to this bundle is functorial.
- There are other natural interpretations of $\mathrm{Jet}_n X_a(K)$, as, for example, the germ of the sheaf of order at most n linear differential operators on $\mathscr{O}_X$ at a.
- $\mathrm{Jet}_1 X \to X$ is the Zariski tangent bundle.

Jets determine subvarieties

Proposition

Let X be an algebraic variety over a field K, $a \in X(K)$ a K-rational points, and Y and Z two irreducible subvarieties of X both passing through a. Then $X = Y$ if and only if for every n, $\mathrm{Jet}_n Y_a(K) = \mathrm{Jet}_n Z_a(K)$ as subspaces of $\mathrm{Jet}_n X_a(K)$.

- Passing to a dense open affine neighborhood of a, we may assume that X is affine. (For me, "variety" means separated, reduced scheme of finite type over a field.)
- For any given n, the image of $\mathrm{Jet}_n Y_a(K)$ in $\mathrm{Jet}_n X_a(K)$ is $\{\psi : \mathfrak{m}_{X,a}/\mathfrak{m}_{X,a}^{n+1} \to K \ : \ (\forall g \in I(Y))\psi(g) = 0\}$ (and likewise for Z).
- If Y and Z differ, then, possibly at the cost of exchanging Y and Z, we can find $f \in I(Y) \smallsetminus I(Z)$.
- Since $\mathscr{O}_{X,a}$ is noetherian, there is some n for which $f \notin \mathfrak{m}_{X,a}^{n+1}$. Hence, we can find some $\psi : \mathfrak{m}_{X,a}/\mathfrak{m}_{X,a}^{n+1} \to K$ which vanishes on the image of $I(Z)$ but not on f.
- Then $\psi \in \mathrm{Jet}_n Z_a(K) \smallsetminus \mathrm{Jet}_n Y_a(K)$ contradicting our hypothesis.

Jets of étale maps

Proposition

If $f : X \to Y$ is a étale map of algebraic varieties over a field K, then for any point $a \in X(K)$ and any $n \in \mathbb{Z}_+$, the map $\mathrm{Jet}_n(f) : \mathrm{Jet}_n X_a(K) \to \mathrm{Jet}_n Y_{f(a)}(K)$ is an isomorphism of K-vector spaces.

Proof.

f induces an isomorphism between $\widehat{\mathscr{O}}_{Y,f(a)} = \varprojlim \mathscr{O}_{Y,f(a)}/fm^{n+1}_{Y,f(a)}$ and $\widehat{\mathscr{O}}_{X,a}$. Truncating and taking duals we conclude the proof. □

Suppose now that (X, Γ) is a finitary σ-variety over some difference field K. Provided that the projection maps $\pi_1 : \Gamma \to X$ and $\pi_2 : \Gamma \to X^\sigma$ are separable (which is automatic in characteristic zero), then for a sufficiently general point $a \in X(L)$ (where L might be a proper extension of K) the projections are étale in a Zariski neighborhood of $(a, \sigma(a))$. Hence, for any $n \in \mathbb{Z}_+$ we obtain an invertible linear map

$$\gamma_{n,a} := \mathrm{Jet}_n(\pi_2)_{(a,\sigma(a))} \circ \mathrm{Jet}_n(\pi_1)^{-1}_{(a,\sigma(a))} : \mathrm{Jet}_n X_a \to \mathrm{Jet}_n X^\sigma_{\sigma(a)}$$

Note that $(\mathrm{Jet}_n X_a, \gamma_{n,a})$ is a finitary σ-variety. We define $\mathrm{Jet}_n(X, \Gamma)_a := (\mathrm{Jet}_n X_a, \gamma_{n,a})^\sharp$.

Proposition

Let $n \in \mathbb{Z}_+$ be a positive integers and (X, Γ) be a finitary σ-variety over a difference closed field $\mathbb{U}$ for which both projections $\pi_1 : \Gamma \to X$ and $\pi_2 : \Gamma \to X^\sigma$ are separable. Then for a sufficiently general point $a \in (X, \Gamma)^\sharp(\mathbb{U})$ the σ-jet space $(\mathrm{Jet}_n X_a, \gamma_{n,a})^\sharp(\mathbb{U})$ is a finite dimensional $\mathrm{Fix}(\sigma)$-vector space which is Zariski dense in $\mathrm{Jet}_n X_a$.

- Since $\gamma_{n,a}$ is invertible, Zariski density is an immediate consequence of the geometric axioms for ACFA.
- That $\mathrm{Jet}_n(X, \Gamma)_a$ is a $\mathrm{Fix}(\sigma)$-vector space follows from the fact that $\gamma_{n,a}$ and σ are $\mathrm{Fix}(\sigma)$-linear.
- The dimension of this space is equal to $\dim J_n X_a$ from our general result on the σ-degree of sharp-sets.

Proof of canonical base property

- Let (X,Γ) be a finitary σ-variety over $\mathbb{U} \models \mathrm{ACFA}_0$, $a \in (X,\Gamma)^\sharp(\mathbb{U})$ a sharp point and $(Y,\Xi) \subseteq (X,\Gamma) \times (B,\Xi)$ an irreducible, normal family of sub-σ-varieties for which Y_b is irreducible and $a \in (Y_b,\Upsilon_b)^\sharp(\mathbb{U})$ for $b \in (B,\Xi)^\sharp(\mathbb{U})$.
- By compactness, for n large enough, we have that if $\mathrm{Jet}_n Y_b = \mathrm{Jet}_n Y_c$, then $b = c$.
- As $\mathrm{Jet}_n(Y_b,\Upsilon_b)_a$ is Zariski dense in $\mathrm{Jet}_n Y_b$, we see that $b = c$ if and only if $\mathrm{Jet}_n(Y_b,\Upsilon_b)_a(\mathbb{U}) = \mathrm{Jet}_n(Y_c,\Upsilon_c)_a(\mathbb{U})$ for $b, c \in (B,\Xi)^\sharp(\mathbb{U})$.
- Hence, we have a definable embedding of $(B,\Xi)^\sharp$ into a Grassmannian of subspaces of $\mathrm{Jet}_n(X,\Gamma)_a$ establishing that $(B,\Xi)^\sharp$ is $\mathrm{Fix}(\sigma)$ internal.