

RESEARCH ARTICLE

Hensel minimality I

Raf Cluckers^{1,2}, Immanuel Halupczok³ and Silvain Rideau-Kikuchi⁴

¹Univ. Lille, CNRS, UMR 8524, Laboratoire Paul Painlevé, F-59000 Lille, France; E-mail: Raf.Cluckers@univ-lille.fr.

²KU Leuven, Department of Mathematics, B-3001 Leuven, Belgium.

³Lehrstuhl für Algebra und Zahlentheorie, Mathematisches Institut, Heinrich-Heine-Universität Düsseldorf, Universitätsstr. 1, 40225 Düsseldorf, Germany; E-mail: math@karimmi.de.

⁴Université de Paris and Sorbonne Université, CNRS, IMJ-PRG, F-75013 Paris, France; E-mail: silvain.rideau@imj-prg.fr.

Received: 22 April 2021; **Revised:** 28 January 2022; **Accepted:** 19 February 2022

Keywords: Non-Archimedean geometry, tame geometry on Henselian valued fields, analogues to o-minimality, cell decomposition, quantifier elimination, Taylor approximation, Lipschitz continuity,

2020 Mathematics Subject Classification: *Primary* – 03C99; *Secondary* – 03C65, 12J20, 11D88, 03C98, 14E18, 41A58

Abstract

We present a framework for tame geometry on Henselian valued fields, which we call Hensel minimality. In the spirit of o-minimality, which is key to real geometry and several diophantine applications, we develop geometric results and applications for Hensel minimal structures that were previously known only under stronger, less axiomatic assumptions. We show the existence of t-stratifications in Hensel minimal structures and Taylor approximation results that are key to non-Archimedean versions of Pila–Wilkie point counting, Yomdin’s parameterization results and motivic integration. In this first paper, we work in equi-characteristic zero; in the sequel paper, we develop the mixed characteristic case and a diophantine application.

Contents

1	Introduction	1
2	Hensel minimality	6
3	Derivation, the Jacobian Property and Taylor approximation	23
4	Resplendency	33
5	Geometry in the Hensel minimal setting	41
6	Examples of Hensel minimal structures	57

1. Introduction

1.1. Background

Our goal is to understand what really lies behind the tameness of definable sets in valued fields and to describe this axiomatically, thereby providing a base for research in several directions of non-Archimedean geometry. In the present paper, we introduce the framework called ‘Hensel minimality’, which is very natural and readily verifiable from Ax–Kochen/Ershov type results. Moreover, it has strong consequences that were previously not expected to follow from such basic axioms.

Let us begin by describing some history of definable sets in valued fields and their uses for p -adic and motivic integration, geometry and diophantine applications. From the 1960s on, various results

were obtained, providing rather precise control over semi-algebraic definable sets in the fields \mathbb{Q}_p of p -adic numbers and $k((t))$ of formal Laurent series with coefficients in a field k of characteristic 0 by Ax, Kochen, Ershov, Cohen, Macintyre and others [1, 2, 3, 37, 25, 37, 53]. Based on those, Denef [26] solved a rationality question of Serre [66] which established a strong connection between number theory and the geometry of p -adic definable sets. This in turn became a motivation for the further development of the model theory of valued fields [60, 4] and motivic integration [28, 29] (an integration theory on arc spaces of varieties) and for the study of singularities and Berkovich spaces by Hrushovski and Loeser, among others [5, 30, 49]. The model theoretic approach also brought Fubini and Fourier to motivic integrals [21, 22, 47], which were key to applications to the Langlands program [16, 8]. Recently, semi-algebraic and subanalytic p -adic geometry has led to point counting results analogous to Pila–Wilkie point counting and Yomdin’s parameterization results [13, 14, 62, 72]. Applications to Lipschitz stratifications were also obtained for semi-algebraic and subanalytic sets from non-Archimedean geometry [43]. See [6] for a more extensive panorama of results. These historical results are all based on concrete settings like semi-algebraic and subanalytic geometry; we open them up by providing a natural axiomatic and more general framework.

In the field of real numbers, a successful framework exists for tame geometry since work by van den Dries, Pillay and others [31, 33, 52, 63]. Indeed, the tameness and geometry of real definable sets are very well captured by the notion of o-minimality, which is an axiomatic condition about definable sets. It became a central tool in real algebraic geometry and its generalizations on the one hand because of its beautiful naturality and on the other hand because of its vast consequences for the geometry of definable sets and its strong diophantine applications, for example to the André–Oort conjecture [61].

Soon after the first successes of o-minimality, the open quest for analogous frameworks in valued fields emerged as a central question in model theory. Several notions have been presented so far, each of which has some shortcomings or lack of generality. In analogy to o-minimality, there are P-minimality [46] and C-minimality [54, 45]; more motivated by applications are V-minimality [47] and b-minimality [20]; motivated by classification theory is dp-minimality [59, 51]. In the present paper, our new notion of Hensel minimality overcomes these shortcomings and provides a strong and natural framework for non-Archimedean tame geometry. Hensel minimality is easier to axiomatize and more natural than b-minimality and V-minimality; it applies more broadly than P-minimality, C-minimality, V-minimality or even dp-minimality; and it has stronger consequences than b-minimality, P-minimality and C-minimality.

In fact, we introduce two key notions of Hensel minimality: 1-h-minimality and ω -h-minimality. Most of our results are developed under the weaker assumption of 1-h-minimality; only some of the resplendency results from Section 4 need ω -h-minimality. As auxiliary notions, we define some intermediary variants of Hensel minimality, namely ℓ -h-minimality for integers $\ell > 1$. We use the term ‘Hensel minimality’ to talk about any of those variants in an implicit way. Most classically studied examples turn out to be ω -h-minimal, but for some interesting (somewhat less classical) examples, we were only able to prove 1-h-minimality (see Section 6).

In this paper, we work in equi-characteristic zero. The mixed characteristic case is derived from this in the sequel [17]. Indeed, the mixed characteristic version of Hensel minimality comes almost for free (including many of its consequences) from the equi-characteristic zero situation by using coarsened valuations. In Section 6, we briefly define the mixed characteristic variant of Hensel minimality for the sake of our examples.

We now sketch ideas, results and applications. In a way, o-minimality assumes that every unary definable set is controlled by a finite set, namely its set of boundary points; in this sense, Hensel minimality is very similar to o-minimality. To make this more precise, the definition of an o-minimal field R can be phrased as the following condition on definable subsets X of R : There exists a finite subset C of R such that, for any $x \in R$, whether x lies in X depends only on the tuple $(\text{sgn}(x - c))_{c \in C}$, where sgn stands for the sign, which can be $-$, 0 or $+$. The definition of Hensel minimality is similar, where the sign function is replaced by a suitable function adapted to the valuation, namely the leading

term map rv . However, whereas in the o-minimal world, C is automatically definable over the same parameters as X , in the valued field setting, we need to impose precise conditions on the parameters over which C can be defined. This is also where the differences between 1-h-minimality and ω -h-minimality arise.

A large part of this paper consists of proving our main geometric results in Hensel minimal structures, which are similar to those in o-minimal structures, in particular cell decomposition, dimension theory, the ‘Jacobian Property’ (which plays a key role in constructing motivic integration and can be considered an analogue of the Monotonicity Theorem from the o-minimal context, where sgn is replaced by rv), as well as higher-order and higher-dimension versions of the Jacobian Property, which state that definable functions have good approximations by their Taylor polynomials. Based on those properties, various recent results in the model theory of Henselian valued fields readily generalize to arbitrary Hensel minimal valued fields, like those on Lipschitz continuity [11] and t-stratifications (which were introduced in [43] and studied further in [42, 41, 44]). As an extra upshot, Hensel minimality intrinsically has resplendency properties in the spirit of resplendent quantifier elimination: that is, it is preserved by different kinds of expansions of the structure. In particular, it should be considered a notion of tameness ‘relative to the leading term structure RV ’.

As first applications, we show the existence of t-stratifications in arbitrary Hensel minimal structures, and we use our results in valued fields to deduce a uniform Taylor approximation result in power-bounded o-minimal real closed fields, which strengthens results from [43, 44]; here, the connection between valued fields and real closed fields is made using the notion of T -convexity from [35, 32]. Furthermore, our geometric results lay the ground for further generalizations of motivic integration from [21, 22, 47] and its use in, for example, [7, 9, 12, 15, 22, 39, 48, 50, 56, 58, 69, 70]. Concretely, we show how Hensel minimality relates to the axiomatic frameworks under which motivic integration is developed in [23] and [47]. Our results also lay the ground for C^r parameterizations and bounds on the number of rational points as in [13, 14], analogous to results by Yomdin [71] and Pila–Wilkie [62]. A first diophantine application to the counting of rational points of bounded height is given in [17].

1.2. The notion of Hensel minimality

We start by giving the definition of Hensel minimality; see Section 2.3 for more details.

Let K be a non-trivially valued field of equi-characteristic 0, considered an \mathcal{L} -structure for some language \mathcal{L} containing the language $\mathcal{L}_{\text{val}} = \{+, \cdot, \mathcal{O}_K\}$ of valued fields (where \mathcal{O}_K is a predicate for the valuation ring). We use multiplicative notation for valuations. We denote the value group by Γ_K^\times and the valuation map by

$$|\cdot|: K \rightarrow \Gamma_K := \Gamma_K^\times \cup \{0\};$$

see Section 2.1 for more detailed definitions.

The analogue to the sign map from the o-minimal context will be the ‘leading term map’ $\text{rv}: K \rightarrow \text{RV}$, defined as follows:

Definition 1.2.1 (Leading term structure RV_λ). Let $\lambda \leq 1$ be an element of Γ_K^\times , and set $I := \{x \in K \mid |x| < \lambda\}$. We define RV_λ^\times to be the quotient of multiplicative groups $K^\times / (1 + I)$, and we let

$$\text{rv}_\lambda: K \rightarrow \text{RV}_\lambda := \text{RV}_\lambda^\times \cup \{0\}$$

be the map extending the projection map $K^\times \rightarrow \text{RV}_\lambda^\times$ by sending 0 to 0. We abbreviate RV_1 and rv_1 by RV and rv , respectively. If several valued fields are around, we may also write RV_K and $\text{RV}_{K,\lambda}$.

We can now make precise in which sense a set $X \subset K$ can be controlled by a finite set $C \subset K$.

Definition 1.2.2 (Prepared sets; see Definition 2.3.1). Let $\lambda \leq 1$ be in Γ_K^\times , let C be a finite non-empty subset of K , and let $X \subset K$ be an arbitrary subset. We say that C λ -prepares X if whether some $x \in K$ lies in X depends only on the tuple $(\text{rv}_\lambda(x - c))_{c \in C}$. In other words, if $x, x' \in K$ satisfy

$$\text{rv}_\lambda(x - c) = \text{rv}_\lambda(x' - c) \text{ for each } c \in C, \tag{1.2.1}$$

then either x and x' both are elements of X , or none of them are elements of X .

A key ingredient of Hensel minimality is that any A -definable set $X \subset K$ (for $A \subset K$) can be 1-prepared by a finite A -definable set C . This, however, is not yet strong enough a condition: We need some precise control of parameters from RV and RV_λ . This is where we obtain different variants of Hensel minimality, the difference consisting only in some number ℓ of allowed parameters.

Definition 1.2.3 (ℓ -h-minimality; see Definition 2.3.3). Let $\ell \geq 0$ be either an integer or ω , and let \mathcal{T} be a theory of valued fields of equi-characteristic 0, in a language \mathcal{L} containing the language \mathcal{L}_{val} of valued fields. We say that \mathcal{T} is ℓ -h-minimal if every model $K \models \mathcal{T}$ has the following property:

(1.2.2) For every $\lambda \leq 1$ in Γ_K^\times , for every set $A \subset K$ and for every set $A' \subset \text{RV}_\lambda$ of cardinality $\#A' \leq \ell$, every $(A \cup \text{RV} \cup A')$ -definable set $X \subset K$ can be λ -prepared by a finite A -definable set $C \subset K$.

In this definition, the parameters from RV and RV_λ are imaginary elements (see Section 2.2).

The notion of ℓ -h-minimality seems most natural in the case $\ell = \omega$, and indeed, this case is especially interesting in view of the respendency results of Section 4. However, the case $\ell = 1$ plays an important role, too: most of the results in this paper follow from 1-h-minimality, and we obtain a rather different equivalent characterization of 1-h-minimality in terms of definable functions (see Theorem 2.9.1). This shows that 1-h-minimality is of a deeply geometric nature, even though its definition might appear somewhat contrived. Similar characterizations for integers $\ell \geq 2$ are developed in [68], involving properties of definable functions on ℓ -dimensional sets.

Many structures on valued fields whose model theory is known to behave well are Hensel minimal. In Section 6, we provide examples of Hensel minimal structures on valued fields of mixed and equi-characteristic 0. Some are new and less expected; others incorporate more classical structures. We show that the following classical structures are Hensel minimal:

1. The theory of Henselian valued fields of equi-characteristic 0 in the pure valued field language is ω -h-minimal.
2. The theory stays ω -h-minimal if we expand the valued field by analytic functions (forming an analytic structure as in [18] or [19]): For example, we can consider the theory of a field $K = k((t))$ with $\text{char } k = 0$, in a language containing all analytic functions $\mathcal{O}_K^n \rightarrow K$ for all n , where ‘analytic’ here means that the function is given by evaluating a t -adically converging power series in $x \in \mathcal{O}_K^n$ (see Section 3.4 of [19]).
3. If $\mathcal{T}_{\text{omin}}$ is the theory of a power-bounded o-minimal expansion of a real closed field, then the theory of $\mathcal{T}_{\text{omin}}$ -convex valued fields in the sense of [35] is 1-h-minimal. Recall that a $\mathcal{T}_{\text{omin}}$ -convex valued field is obtained by taking a (sufficiently big) model K of $\mathcal{T}_{\text{omin}}$ and turning it into a valued field using the convex closure of an elementary substructure $K_0 < K$ as the valuation ring; see Section 6.3 for details.
4. Any expansion of the structure by predicates on Cartesian powers of RV preserves 0-, 1- and ω -h-minimality (Theorem 4.1.19).
5. The notion of ω -h-minimality is preserved under coarsening of the valuation (see Corollary 4.2.4). In [17, Theorem 2.2.7], it is proved that 1-h-min is also preserved under coarsening of the valuation.

By combining (5) and (2), one recovers the examples of coarsened analytic structures from [64] and respently so by (4).

On the other hand, there are also some well-behaved structures on valued fields that are certainly not Hensel minimal: Since the very definition of Hensel minimality implies that every definable subset of

K is either finite or contains a ball, one cannot add a section of the residue field or the value group to the language, and neither can one add an automorphism of K .

As a converse to (1), we show that a valued field that is Hensel minimal is automatically Henselian; see Theorem 2.7.2. In Proposition 6.4.2, we compare Hensel minimality to V-minimality from [47].

1.3. Description of the main results

We now list some of the consequences of 1-h-minimality (often in simpler forms than the versions in the main part of the paper). The first result is the ‘Jacobian Property’, which plays a crucial role for example in motivic integration, both in the Cluckers–Loeser version [21, 22] and in the Hrushovski–Kazhdan version [47].

Theorem 1.3.1 (Jacobian Property; see Corollary 3.2.7). *Let $f : K \rightarrow K$ be definable (with parameters). Then there exists a finite set $C \subset K$ such that the following holds: for every fiber B of the map sending $x \in K$ to the tuple $(\text{rv}(x - c))_{c \in C}$, there exists a $\xi_B \in \text{RV}$ such that*

$$\text{rv} \left(\frac{f(x_1) - f(x_2)}{x_1 - x_2} \right) = \xi_B \tag{1.3.1}$$

for all $x_1, x_2 \in B, x_1 \neq x_2$.

Remark 1.3.2. This formulation of the Jacobian Property becomes exactly the o-minimal Monotonicity Theorem if one replaces all occurrences of rv by sgn and if one adds the condition that f should be continuous on the fibers B (which, in the valued field setting, follows automatically).

Corollary 3.2.7 also includes the following strengthenings of Theorem 1.3.1:

- The theorem still holds with rv replaced by rv_λ for any $\lambda \leq 1$ in Γ_K^\times (still only assuming 1-h-minimality), and one can moreover choose a single finite set C that works for all such λ .
- If f is definable with parameters from $A \cup \text{RV}$ with $A \subset K$, then C can taken to be A -definable. (Corollary 3.2.7 only speaks about \emptyset -definable sets f ; Remark 2.6.11 explains how to deduce this more general version.)

Another point of view of the Jacobian Property is that on each fiber B (using notation from the Theorem), f has a good approximation by its first-order Taylor series. One of the deepest results of this paper is a similar result for higher-order Taylor approximations:

Theorem 1.3.3 (Taylor approximations; see Theorem 3.2.2). *Let $f : K \rightarrow K$ be a definable function, and let $r \in \mathbb{N}$ be given. Then there exists a finite set C such that for every fiber B of the map sending $x \in K$ to the tuple $(\text{rv}(x - c))_{c \in C}$, f is $(r + 1)$ -fold differentiable on B and we have*

$$\left| f(x) - \sum_{i=0}^r \frac{f^{(i)}(x_0)}{i!} (x - x_0)^i \right| \leq |f^{(r+1)}(x_0) \cdot (x - x_0)^{r+1}| \tag{1.3.2}$$

for every $x_0, x \in B$.

As in Theorem 1.3.1, if f is $(A \cup \text{RV})$ -definable for $A \subset K$, then C can be taken A -definable (again using Remark 2.6.11).

Using that $\mathcal{T}_{\text{omin}}$ -convex valued fields are 1-h-minimal (see (3) above, on p. 4), this Taylor approximation result implies a uniform Taylor approximation result in power-bounded o-minimal real closed fields; see Corollary 6.3.7.

As in the o-minimal context, the preparation of unary sets and the Jacobian Property lend themselves well (by logical compactness) to obtain results about higher-dimensional objects (namely, definable subsets of K^n and definable functions on K^n). Indeed, this can be pursued, *mutatis mutandis*, in the style of how cell decomposition and dimension theory are built up from b-minimality in [20]. In particular, we obtain results about

- almost everywhere differentiability (Subsection 5.1),
- cell decomposition in two variants (Subsections 5.2 and 5.7),
- dimension theory (Subsection 5.3), and
- higher-dimensional versions of the Jacobian Property and Taylor approximations (Subsections 5.4 and 5.6).

For cell decomposition, Subsection 5.2 provides a new approach specific to Hensel minimality: Usually, the notion of cells in valued fields is a lot more technical than in o-minimal structures, partly due to the lack of (certain) Skolem functions. Item (4) above (on p. 4) allows us to add the missing Skolem functions to the language without destroying Hensel minimality, and there are also some tools enabling us to get back to the original language afterward (see Subsection 4.3). This allows us, in Subsection 5.2, to work with a notion of a cell, which is much less technical than most previous ones in valued fields.

One of our original goals was to deduce the Jacobian Property starting from abstract conditions on unary sets only. Previous proofs of the Jacobian Property either use piecewise analyticity arguments, even in the semi-algebraic case (like in [18]), or are restricted to rather specific situations (like [47] for V-minimal valued fields and [70] for power-bounded T -convex valued fields). Our new proof is more general and works for arbitrary Henselian valued fields, which are 1-h-minimal, and goes without analyticity arguments.

The remainder of this paper is organized as follows. In Section 2, after fixing notation and terminology, we develop many basic tools that are useful for proofs in Hensel minimal theories, and we obtain first elementary results like a key ingredient to dimension theory (Lemma 2.8.1) and a weak version of the Jacobian Property (Lemma 2.8.5). We also show that those two results are essentially equivalent to 1-h-minimality (Theorem 2.9.1).

The deepest results of this paper are contained in Section 3 about definable functions from K to K , namely almost everywhere differentiability and Taylor approximation.

Section 4 is devoted to understanding in which sense Hensel minimality is a notion ‘relative to RV’ and to proving that various modifications of the language preserve Hensel minimality.

The geometric results in K^n (like cell decomposition, dimension theory) are collected in Section 5. That section also contains our main application of t-stratifications, our higher-dimensional Taylor approximation results, and Cluckers–Loeser style motivic integration.

Finally, in Section 6, we show that many previously studied structures are Hensel minimal, and we give some new examples as well, namely coarsened valued fields as variants of classical analytic structures. Since some of those examples are of mixed characteristics, at the beginning of Section 6 we briefly define the mixed characteristic variant of Hensel minimality, which is treated in full in the sequel [17] to this paper.

As an application, we show in Section 6.3 how Hensel minimality results yield corresponding results in power-bounded real closed fields. At the end, we compare our notion to V-minimality from [47].

2. Hensel minimality

2.1. Notation and terminology for valued fields and balls

In the entire paper, K will denote a valued field of characteristic zero, with valuation ring $\mathcal{O}_K \subset K$ and maximal ideal $\mathcal{M}_K \subset \mathcal{O}_K$. In this paper, we only consider non-trivially valued fields: that is, $\mathcal{O}_K \neq K$. Moreover, apart from Section 6, K will always be of equi-characteristic 0, meaning that both K and the residue field $\mathcal{O}_K/\mathcal{M}_K$ have characteristic 0.

Note that we allow Krull-valuations (and thus valuations of arbitrary rank): that is, we allow \mathcal{O}_K to be an arbitrary (proper, by the non-triviality) subring of K such that for every element $x \in K^\times$, at least one of x or x^{-1} belongs to \mathcal{O}_K . The value group is then defined to be the quotient $\Gamma_K^\times := K^\times/\mathcal{O}_K^\times$ of multiplicative groups.

We denote the valuation map by $|\cdot|: K^\times \rightarrow \Gamma_K^\times$ and use multiplicative notation for the value group. We write Γ_K for the disjoint union $\Gamma_K^\times \cup \{0\}$, we extend the valuation map to $|\cdot|: K \rightarrow \Gamma_K$ by setting

$|0| := 0$, and we define the order on Γ_K in such a way that $\mathcal{O}_K = \{x \in K \mid |x| \leq 1\}$ and $|x| < |y|$ whenever $x/y \in \mathcal{M}_K$ for x and nonzero y in K .

For $x = (x_1, \dots, x_n) \in K^n$, we set $|x| := \max_i |x_i|$.

We use the (generalized) leading term structures RV_λ (for $\lambda \leq 1$ in Γ_K^\times) that have already been introduced in Definition 1.2.1, and we also denote the natural map $\text{RV}_\lambda \rightarrow \Gamma_K$ by $|\cdot|$. Note that RV_λ is a semi-group for multiplication.

Remark 2.1.1. Write RV^\times for $\text{RV} \setminus \{0\}$. Recall that one has a natural short exact sequence of multiplicative groups $(\mathcal{O}_K/\mathcal{M}_K)^\times \rightarrow \text{RV}^\times \rightarrow \Gamma_K^\times$. (So RV combines information from the residue field and value group.)

Example 2.1.2. In the case $K = k((t))$, the above short exact sequence naturally splits, giving an isomorphism $\text{RV}^\times \rightarrow (\mathcal{O}_K/\mathcal{M}_K)^\times \times \Gamma_K^\times$, which, for $a = \sum_{i=N}^\infty a_i t^i \in K^\times$ with $a_N \neq 0$, sends $\text{rv}(a)$ to (a_N, N) .

Remark 2.1.3. For $a, a' \in K$, one has $\text{rv}_\lambda(a) = \text{rv}_\lambda(a')$ if and only if either $a = a' = 0$ or $|a - a'| < \lambda \cdot |a|$.

We consider several kinds of balls:

Definition 2.1.4 (Balls).

1. We call a subset $B \subset K^n$ a *ball* if B is infinite, $B \neq K^n$, and for all $x, x' \in B$ and all $y \in K^n$ with $|x - y| \leq |x - x'|$, one has $y \in B$.
2. By an *open ball*, we mean a set of the form

$$B = B_{<\gamma}(a) := \{x \in K^n \mid |x - a| < \gamma\}$$

for some $a \in K^n$ and some $\gamma \in \Gamma_K^\times$.

3. By a *closed ball*, we mean a set of the form

$$B = B_{\leq\gamma}(a) := \{x \in K^n \mid |x - a| \leq \gamma\}$$

for some $a \in K^n$ and some $\gamma \in \Gamma_K^\times$.

4. For B as in (2) or (3), we call γ the *radius* of the open (respectively, closed) ball and denote it by $\text{rad}_{\text{op}}(B)$ (respectively, $\text{rad}_{\text{cl}}(B)$).

We call the valuation on K discrete if there is a uniformizing element ϖ in \mathcal{O}_K , namely satisfying $\varpi \mathcal{O}_K = \mathcal{M}_K$.

Remark 2.1.5. When Γ_K^\times is discrete, a ball B can be open and closed at the same time, but with $\text{rad}_{\text{cl}}(B) = |\varpi| \cdot \text{rad}_{\text{op}}(B) < \text{rad}_{\text{op}}(B)$.

Note also that for any $\lambda \leq 1$ in Γ_K^\times and any $\xi \in \text{RV}_\lambda \setminus \{0\}$, the preimage $\text{rv}_\lambda^{-1}(\xi)$ is an open ball satisfying $\text{rad}_{\text{op}}(\text{rv}_\lambda^{-1}(\xi)) = |\xi| \cdot \lambda$.

Definition 2.1.6 (λ -next balls). Fix $\lambda \leq 1$ in Γ_K^\times .

1. We say that a ball B is λ -next to an element $c \in K$ if

$$B = \{x \in K \mid \text{rv}_\lambda(x - c) = \xi\}$$

for some (nonzero) element ξ of RV_λ .

2. We say that a ball B is λ -next to a finite non-empty set $C \subset K$ if B equals $\bigcap_{c \in C} B_c$ with B_c a ball λ -next to c for each $c \in C$.

Remark 2.1.7. Using that the intersection of the finitely many balls B_c is either empty or equal to one of the B_c , one deduces that every ball B that is λ -next to C is in particular λ -next to one element $c \in C$.

Remark 2.1.8. Given a finite non-empty set $C \subset K$, the fibers of the map $x \mapsto (rv_\lambda(x - c))_{c \in C}$ are exactly the singletons consisting of one element of C and the balls λ -next to C . In particular, the balls λ -next to C form a partition of $K \setminus C$, and a subset X of K is λ -prepared by C (as in Definition 2.3.1) if and only if every ball B that is λ -next to C is either contained in X or disjoint from X .

Example 2.1.9. The balls 1-next to an element $c \in K$ are exactly the maximal balls in K not containing c . From this, one deduces that a ball 1-next to a finite set C is exactly a maximal ball disjoint from C . This means that a set $X \subset K$ is 1-prepared by C if and only if every ball disjoint from C is either contained in X or disjoint from X . Note again how closely ‘every definable set can be 1-prepared’ resembles o-minimality (where ‘balls disjoint from C ’ becomes ‘intervals disjoint from C ’).

Given a subset A of a Cartesian product $B \times C$ and an element $b \in B$, we write A_b for the fiber $\{c \in C \mid (b, c) \in A\}$. Also, for a function g on A , we write $g(b, \cdot)$ for the function on A_b sending c to $g(b, c)$.

2.2. Model theoretic notations and conventions

As already stated, in the entire paper, K is a non-trivially valued field of characteristic zero, and outside of Section 6, K is moreover of equi-characteristic 0. In the entire paper, we fix a language \mathcal{L} containing the language \mathcal{L}_{val} of valued fields, and we consider the valued field K as an \mathcal{L} -structure. More precisely, as ‘language of valued fields’, it suffices for us to take $\mathcal{L}_{\text{val}} = \{+, \cdot, \mathcal{O}_K\}$, where \mathcal{O}_K is a predicate for the valuation ring; in any case, we only care about which sets are definable (and not how they are definable). For that reason, we will often specify languages only up to interdefinability.

If not specified otherwise, ‘definable’ always refers to the fixed language \mathcal{L} . As usual, ‘ \mathcal{L} -definable’ means definable (in \mathcal{L}) without additional parameters, ‘ A -definable’ means $\mathcal{L}(A)$ -definable, and ‘definable’ means $\mathcal{L}(K)$ -definable (i.e., with arbitrary parameters).

Sometimes, we will also consider K a structure in other languages (e.g., \mathcal{L}'); in that case, we may specify the language as an index, writing for example $\text{Th}_{\mathcal{L}'}(K)$ for the theory of K considered an \mathcal{L}' -structure.

In almost the entire paper (more precisely, everywhere except in parts of Section 4.1), \mathcal{L} will be a one-sorted language. Nevertheless, we often work with imaginary sorts of K : that is, quotients K^n/\sim , where \sim is a \emptyset -definable equivalence relation. In particular, we consider imaginary definable sets and imaginary elements. As usual, this can be made formal either by working in the \mathcal{L}^{eq} -structure K^{eq} (see for example [67, Proposition 8.4.5]) or, equivalently, by ‘simulating’ imaginary objects in \mathcal{L} , namely as follows:

- By a ‘definable subset X of K^n/\sim ’, we really mean its preimage in K^n : that is, a definable set $\tilde{X} \subset K^n$ that is a union of \sim -equivalence classes.
- If, in a formula $\phi(x, \dots)$, the variable x runs over an imaginary sort K^n/\sim , this means that we really have an n -tuple \tilde{x} of variables (running over K^n) and that the truth value of $\phi(\tilde{x}, \dots)$ only depends on the equivalence class of \tilde{x} modulo \sim .
- If A is a set of potentially imaginary elements, then by ‘ A -definable’, we mean definable in the expansion of K by predicates for the equivalence classes (in K^n for some n) corresponding to the imaginary elements $a \in A$; the corresponding extension of the language \mathcal{L} by predicate symbols is denoted by $\mathcal{L}(A)$.
- A (potentially imaginary) element b is said to be in the definable closure of a set A (of potentially imaginary elements) if the equivalence class in K^n corresponding to b is $\mathcal{L}(A)$ -definable. If X is any imaginary sort (or even more generally an arbitrary set of imaginary elements), we write $\text{dcl}_X(A)$ for the set of elements from X that are in the definable closure of A . Being in the algebraic closure, with notation $\text{acl}_X(A)$, is defined accordingly.

The value group Γ_K is, of course, an imaginary sort. In general, RV_λ (for $\lambda \leq 1$ in Γ_K^\times) is, by itself, not an imaginary sort, since the equivalence relation used to define it may not be \emptyset -definable. However,

the disjoint union

$$RV_{\bullet} := \bigcup_{\lambda \leq 1} RV_{\lambda}$$

is an imaginary sort, and RV_{λ} is a definable subset of RV_{\bullet} ; in particular, it makes sense to use elements from RV_{λ} as parameters.

2.3. Hensel minimality

We first restate the definitions of prepared sets and Hensel minimality from the introduction. We use RV_{λ} and rv_{λ} from Definition 1.2.1.

Definition 2.3.1 (Prepared sets). Let K be a non-trivially valued field of equi-characteristic 0. Let $\lambda \leq 1$ be in Γ_K^{\times} , let C be a finite non-empty subset of K , and let $X \subset K$ be an arbitrary subset. We say that C λ -prepares X if there exists a set $\Xi \subset RV_{\lambda}^{\#C}$ such that X is the preimage of Ξ under the map $K \rightarrow RV_{\lambda}^{\#C}, x \mapsto (rv_{\lambda}(x - c))_{c \in C}$.

The condition that C is non-empty is essentially irrelevant, since one could always add 0 to C , but it will sometimes avoid pathologies.

Example 2.3.2. A subset of K is λ -prepared by the set $C = \{0\}$ if and only if it is the preimage under rv_{λ} of a subset of RV_{λ} .

Definition 2.3.3 (ℓ -h-minimality). Let $\ell \geq 0$ be either an integer or ω , and let \mathcal{T} be a (possibly non-complete) theory of valued fields of equi-characteristic 0, in a language \mathcal{L} containing the language \mathcal{L}_{val} of valued fields. We say that \mathcal{T} is ℓ -h-minimal if every model $K \models \mathcal{T}$ has the following property:

(2.3.1) For every $\lambda \leq 1$ in Γ_K^{\times} , for every set $A \subset K$ and for every set $A' \subset RV_{\lambda}$ of cardinality $\#A' \leq \ell$, every $(A \cup RV \cup A')$ -definable set $X \subset K$ can be λ -prepared by a finite A -definable set $C \subset K$.

Note that a theory \mathcal{T} is ℓ -h-minimal if and only if each of its completions is. Conversely, compactness arguments show that most of our results hold uniformly in a theory \mathcal{T} if and only they hold in every completion (see Proposition 2.6.2, for example). For this reason, we will often work with complete theories in this paper.

Remark 2.3.4. In the case $\ell = 0$, λ plays no role, and the definition simplifies to: Any $(A \cup RV)$ -definable set (for $A \subset K$) can be 1-prepared by a finite A -definable set.

Remark 2.3.5. In the case $\ell = \omega$, we can more generally allow X to use parameters $\xi_i \in RV_{\lambda_i}$ for different λ_i (using that we have surjections $RV_{\lambda} \rightarrow RV_{\lambda'}$ for $\lambda' > \lambda$); in that case, C is required to λ -prepare X for $\lambda := \min_i \lambda_i$.

Remark 2.3.6. The assumption that C is definable using only the parameters from K will enable us to simultaneously prepare families of sets parametrized by RV (and RV_{λ}). This plays a central role in making Hensel minimality independent of the structure on RV .

2.4. Basic model theoretic properties of Hensel minimality

Let us now verify that the notion of Hensel minimality (in all its variants) has some basic properties one would expect from any good model theoretic notion.

First of all, note that it is preserved under expansions of the structure by constants; more precisely:

Lemma 2.4.1 (Adding constants). *Let $\ell \geq 0$ be an integer or ω , suppose that $\text{Th}(K)$ is ℓ -h-minimal, and let A be any subset of $K \cup RV^{eq}$. Then $\text{Th}_{\mathcal{L}(A)}(K)$ is also ℓ -h-minimal.*

Here, by RV^{eq} , we mean imaginary sorts of the form RV^n/\sim for some n and some \emptyset -definable equivalence relations \sim . In particular, the lemma allows us to add constants from Γ_K . Note, however, that adding parameters from other sorts than K and RV^{eq} may destroy Hensel minimality.

The lemma should be clear in the case $A \subset K \cup \text{RV}$; we mainly give the following proof to show that parameters from RV^{eq} are not a problem either.

Proof of Lemma 2.4.1. We verify Definition 2.3.3: Let $K' \models \text{Th}_{\mathcal{L}(A)}(K)$, and let $X \subset K'$ be $\mathcal{L}(A \cup A' \cup \text{RV}_{K'} \cup A'')$ -definable for some $A' \subset K'$ and some $A'' \subset \text{RV}_{K',\lambda}$ satisfying $\#A'' \leq \ell$, with $\lambda \leq 1$ in $\Gamma_{K'}$. Choose $\tilde{A} \subset \text{RV}_{K'}$ such that every element of $A \cap \text{RV}_{K'}^{\text{eq}}$ is $\mathcal{L}(\tilde{A})$ -definable. Then X is $\mathcal{L}((A \cap K) \cup \tilde{A} \cup A' \cup \text{RV}_{K'} \cup A'')$ -definable, so by ℓ -h-minimality of $\text{Th}_{\mathcal{L}}(K')$, X can be λ -prepared by a finite $\mathcal{L}((A \cap K) \cup A')$ -definable set C . In particular, C is $\mathcal{L}(A \cup A')$ -definable, as desired. \square

With this lemma in mind, many results in this paper are formulated for \emptyset -definable sets; those results then automatically also hold for A -definable sets, when $A \subset K \cup \text{RV}^{\text{eq}}$, and using a compactness argument (given in Remark 2.6.3), one then often obtains family versions of the results.

As so often in model theory, it is sufficient to verify Hensel minimality in sufficiently saturated models. To see this, we first prove that preparation is a first-order property in the following sense:

Lemma 2.4.2 (Preparation is first order). *Let X_q and C_q be \emptyset -definable families of subsets of K , where q runs over a \emptyset -definable subset Q of an arbitrary possibly imaginary sort. Suppose that C_q is finite for every $q \in Q$. Then the set of pairs $(q, \lambda) \in Q \times \Gamma_K^\times$ with $\lambda \leq 1$ such that C_q λ -prepares X_q is \emptyset -definable.*

Proof. If $\phi(x, q)$ defines X_q and $\psi(z, q)$ defines C_q , then the above set of pairs (q, λ) is defined by the following \mathcal{L} -formula:

$$\forall x, x': \underbrace{(\forall z: (\psi(z, q) \rightarrow \text{rv}_\lambda(x - z) = \text{rv}_\lambda(x' - z))) \rightarrow (\phi(x, q) \leftrightarrow \phi(x', q)))}_{\text{i.e., } (\text{rv}_\lambda(x - c))_{c \in C_q} = (\text{rv}_\lambda(x' - c))_{c \in C_q}}$$

\square

Lemma 2.4.3 (Saturated models suffice). *Let $\ell \geq 0$ be either an integer or ω , and suppose that K is \aleph_0 -saturated. Then the theory $\text{Th}(K)$ of K is ℓ -h-minimal if and only if K satisfies Condition (2.3.1) from Definition 2.3.3.*

Proof. We need to show that if K satisfies (2.3.1), then so does any other model K' of $\text{Th}(K)$.

Suppose for contradiction that K' is a model not satisfying (2.3.1): that is, there exist a $\lambda \leq 1$ in $\Gamma_{K'}^\times$, tuples $a \in (K')^n$, $\zeta \in \text{RV}_{K'}^{n'}$, $\xi \in \text{RV}_{K',\lambda}^{n''}$ with n, n' arbitrary and $n'' \leq \ell$, and an (a, ζ, ξ) -definable set $X = \phi(K', a, \zeta, \xi) \subset K'$ such that no finite non-empty a -definable set $C \subset K$ λ -prepares X .

For fixed ϕ , the non-existence of C can be expressed by an infinite conjunction of \mathcal{L} -formulas in (λ, a, ζ, ξ) . Indeed, for every formula $\psi(z, a)$ that could potentially define C and for every integer $k \geq 1$, there is (by Lemma 2.4.2) an \mathcal{L} -formula $\chi_\psi(\lambda, a, \zeta, \xi)$ expressing ‘ $\psi(K', a)$ has cardinality k and $\psi(K', a)$ does not λ -prepare $\phi(K', a, \zeta, \xi)$.’

The fact that this partial type $\{\chi_\psi \mid \psi \text{ as above}\}$ is realized in K' implies that it is also realized in K , so that Condition (2.3.1) fails in K . \square

Whether a structure is o-minimal can also be characterized via its 1-types. For 0-h-minimality, we have a similar characterization:

Lemma 2.4.4 (0-h-minimality in terms of types). *Suppose that K is \aleph_0 -saturated. The theory $\text{Th}(K)$ is 0-h-minimal if and only if, for every parameter set $A \subset K$ and every ball $B \subset K \setminus \text{acl}_K(A)$, any two elements of B have the same type over $A \cup \text{RV}$.*

Remark 2.4.5. One could also formulate similar conditions for ℓ -h-minimality, but it would be more technical.

Proof of Lemma 2.4.4. ‘ \Rightarrow ’: Suppose for contradiction that B contains two elements x, x' with $\text{tp}(x/A \cup \text{RV}) \neq \text{tp}(x'/A \cup \text{RV})$. This means that there exists an $(A \cup \text{RV})$ -definable set X containing x but not x' . By 0-h-minimality, there exists a finite A -definable set C 1-preparing X . In particular, $C \subset \text{acl}_K(A)$ and hence $C \cap B = \emptyset$. However, by Example 2.1.9, this implies that B is either contained in X or disjoint from X , contradicting the properties of x and x' .

‘ \Leftarrow ’: Let X be $(A \cup \text{RV})$ -definable for some parameter set $A \subset K$ that we may assume to be finite, and suppose that no finite A -definable $C \subset K$ 1-prepares X . This means (using Example 2.1.9 again) that for every finite A -definable $C \subset K$, there exists an open ball $B \subset K$ that is disjoint from C and such that neither $B \subset X$ nor $B \subset K \setminus X$. Taking all those conditions on B together (for all finite A -definable C), we obtain a (finitely satisfiable) type in an imaginary variable running over the open balls in K . A realization of this type is a ball B that is disjoint from $\text{acl}_K(A)$ on the one hand but, on the other hand, contains elements x and x' that satisfy $x \in X$ and $x' \notin X$ and hence have different types over $A \cup \text{RV}$. \square

Yet another way to characterize o-minimality is: Every unary definable set is already quantifier-free definable in the language $\{<\}$. Using Lemma 4.1.10, one obtains a similar kind of characterization of 0-h-minimality and ω -h-minimality.

2.5. Basic properties under weaker assumptions

Recall (from the beginning of Section 2.1) that K is a non-trivially valued field of equi-characteristic zero, considered a structure in a language $\mathcal{L} \supset \mathcal{L}_{\text{val}} = \{+, \cdot, \mathcal{O}_K\}$. In this subsection, we assume $\text{Th}(K)$ to be ‘Hensel minimal without control of parameters’, namely:

Assumption 2.5.1. For every $K' \equiv K$, every definable (with any parameters) subset of K' can be 1-prepared by a finite set C . (We do not impose definability conditions on C .)

Note that this assumption is preserved under adding constants to \mathcal{L} (even from arbitrary imaginary sorts), so below, every occurrence of ‘ \emptyset -definable’ can also be replaced by ‘ A -definable’.

Lemma 2.5.2 (\exists^∞ -elimination). *Under Assumption 2.5.1, every infinite definable set $X \subset K$ contains an (open) ball. In particular, if $\{X_q \mid q \in Q\}$ is a \emptyset -definable family of subsets X_q of K for some \emptyset -definable set Q in an arbitrary possibly imaginary sort, then the set $Q' := \{q \in Q \mid X_q \text{ is finite}\}$ is a \emptyset -definable set, and there exists a uniform bound $N \in \mathbb{N}$ on the cardinality of X_q for all $q \in Q'$.*

Proof. If X is infinite, it is not contained in the finite set C preparing it, which implies that it contains a ball. The definability of Q' then follows since the condition that a set contains an open ball can be expressed by a formula. The existence of the bound N then follows by compactness: If no bound would exist, then in a sufficiently saturated model, we would find a $q \in Q'$ with $\#X_q > N$ for every N . \square

Lemma 2.5.3 (Finite sets are RV-parametrized). *Under Assumption 2.5.1, let $C_q \subset K$ be a \emptyset -definable family of finite sets, where q runs over some \emptyset -definable set Q in an arbitrary possibly imaginary sort. Then there exists a \emptyset -definable family of injective maps $f_q: C_q \rightarrow \text{RV}^k$ (for some k).*

Proof of Lemma 2.5.3. Using Lemma 2.5.2, we can bound the cardinality $\#C_q$ and then assume that it is constant. We do an induction over $\#C_q$.

If C_q is always a singleton or empty, we can define f_q to always be constant. Otherwise, the lemma is obtained by repeatedly taking averages of the elements of C_q and subtracting. More precisely, setting $a_q := \frac{1}{\#C_q} \sum_{x \in C_q} x$, we get that the map $\hat{f}_q: C_q \rightarrow \text{RV}$, $x \mapsto \text{rv}(x - a_q)$ is not constant on C_q . Therefore, each fiber $\hat{f}_q^{-1}(\xi)$ of \hat{f} (for $\xi \in \hat{f}_q(C_q)$) has cardinality less than C_q , so by induction, we obtain a definable family of injective maps $g_{q,\xi}: \hat{f}_q^{-1}(\xi) \rightarrow \text{RV}^k$. Now set $f_q(x) := (\hat{f}_q(x), g_{q,\hat{f}_q(x)}(x))$. \square

The family of balls 1-next to some finite set $C \subset K$ can be parameterized by RV-variables; more precisely (and more generally), we have the following:

Lemma 2.5.4 (λ -next balls as fibers). *Under Assumption 2.5.1, let $\lambda \leq 1$ be an element of Γ_K^\times , let A be any set of possibly imaginary parameters containing λ , and let $C \subset K$ be a finite non-empty A -definable set. Then there exists an A -definable map $f: K \rightarrow \text{RV}^k \times \text{RV}_\lambda$ (for some k) such that each nonempty fiber of f is either a singleton contained in C or contained in a single ball λ -next to C . In the case $\lambda = 1$, we may even obtain that each fiber of f that is not a singleton is equal to a ball 1-next to C .*

Proof. Given $x \in K$, let $\mu(x) := \min\{|x - c| \mid c \in C\}$ be the minimal distance to elements of C , let $C(x) := \{c \in C \mid |x - c| = \mu(x)\}$ be the set $c \in C$ realizing that distance, and let $a(x) := \frac{1}{\#C(x)} \sum_{c \in C(x)} c$ be the average of those elements. Note that the map $a: K \rightarrow K$ has a finite image. Using Lemma 2.5.3, we find an injective map α from the image of a to RV^k . If $\lambda < 1$, we define

$$f(x) := (\alpha(a(x)), \text{rv}_\lambda(x - a(x))).$$

In the case $\lambda = 1$, to obtain the more precise statement, we define

$$f(x) := \begin{cases} (\alpha(a(x)), \text{rv}(x - a(x))) & \text{if } |x - a(x)| \geq \mu(x) \\ (\alpha(a(x)), \text{rv}(0)) & \text{if } |x - a(x)| < \mu(x). \end{cases}$$

It is now just a computation to check that the lemma holds. For the first part, suppose that $f(x_1) = f(x_2) = (\zeta, \xi) \in \text{RV}^k \times \text{RV}_\lambda$ for some $x_1, x_2 \in K$; our aim is to show that either $x_1 = x_2 \in C$ or they both lie in the same ball λ -next to C . Set $a_0 := a(x_1) = a(x_2)$.

If $\xi = 0$, then either $x_i = a_0$ or $|x_i - a_0| < \mu(x_i)$. The latter is only possible if $a_0 \notin C$, and it implies that x_i and a_0 lie in the same ball 1-next to C . This shows: Either $x_1 = x_2 \in C$ or $\lambda = 1$, and both lie in the same ball 1-next to C .

If $\xi \neq 0$, then the fact that $|x_i - a(x)| \leq \mu(x_i)$ shows that $\text{rv}_\lambda(x_i - a(x))$ determines $\text{rv}_\lambda(x_i - c)$ for each $c \in C$, so we are done with the first part of the lemma.

For the second part, pick x_1, x_2 in the same ball λ -next to C . Then one obtains $\mu(x_1) = \mu(x_2) =: \mu_0$, $C(x_1) = C(x_2) =: C_0$ and $a(x_1) = a(x_2) =: a_0$. If $|x_i - a_0| = \mu_0$, then $\text{rv}(x_i - a_0)$ is determined by $\text{rv}(x_i - c)$ for any $c \in C_0$, so $f(x_1) = f(x_2)$. Otherwise, if $|x_i - a_0| < \mu_0$, then $f(x_i) = (\alpha(a_0), 0)$ for $i = 1, 2$ by definition. \square

Remark 2.5.5. In Lemma 2.5.4, we can also find a map f with codomain RV_λ^{k+1} instead of $\text{RV}^k \times \text{RV}_\lambda$; indeed, in Lemma 2.5.3 and its proof, RV can be replaced by RV_λ everywhere.

2.6. Preparing families

In this subsection, we show that ℓ -h-minimality implies that we can prepare not only definable subsets of K (by finite sets C) but also various other kinds of definable objects that ‘live in $K \times \text{RV}^k \times \text{RV}_\lambda^\ell$ ’.

Definition 2.6.1 (Preparing families). Let $C \subset K$ be a non-empty set, and let $\lambda \leq 1$ be an element of Γ_K^\times .

1. Suppose that W is, up to permutation of coordinates, a subset of $K \times Q$, where Q is a Cartesian product of some (possibly imaginary) sorts. We say that C λ -prepares W if for every ball $B \subset K$ λ -next to C , the fiber $W_x \subset Q$ does not depend on x when x runs over B . (The terminology will only be applied when exactly one of the coordinates of W runs over K ; the fibers W_x always are over that coordinate.)

We synonymously also say that C λ -prepares the family $(W_\xi)_{\xi \in Q}$ of subsets of K . Indeed, the above condition is equivalent to C λ -preparing $W_\xi \subset K$ for each $\xi \in Q$.

2. If f is a definable function whose graph W lives in a Cartesian product as in (1), we say that C λ -prepares f if it λ -prepares W .

3. We say that C prepares a set

$$W \subset K \times Q \times \bigcup_{\lambda \in \Gamma_K^\times, \lambda \leq 1} \text{RV}_\lambda^k,$$

uniformly in λ , where Q is a product of sorts and $k \geq 0$ is an integer, if it λ -prepares $W \cap (K \times Q \times \text{RV}_\lambda^k)$ for each $\lambda \leq 1$ in Γ_K^\times . As in (1), we allow the coordinates of W to be in a different order, but making sure that this is unambiguous.

Proposition 2.6.2 (Preparing families). *Assume that $\text{Th}(K)$ is ℓ -h-minimal and that A is a subset of K . For any integer $k \geq 0$ and any $(A \cup \text{RV})$ -definable set*

$$W \subset K \times \text{RV}^k \times \bigcup_{\lambda \leq 1} \text{RV}_\lambda^\ell,$$

there exists a finite non-empty A -definable set C that prepares W uniformly in λ .

Proof. For each λ and each $\xi \in \text{RV}^k \times \text{RV}_\lambda^\ell$, let C_ξ be a finite A -definable set λ -preparing W_ξ . By a usual compactness argument (see Remark 2.6.3 below), we may suppose that $C := \bigcup_\xi C_\xi$ is finite and A -definable. It prepares each W_ξ for each ξ and hence also W . □

Remark 2.6.3. In the above proof, we used a compactness argument that we will be using (in variants) many times in this paper. We give some details once. First, recall that by Lemma 2.4.2, ‘preparing’ is a definable condition. In particular, the set

$$\Xi_\xi := \{\xi' \in \text{RV}^k \times \bigcup_{\lambda \leq 1} \text{RV}_\lambda^\ell \mid C_\xi \text{ } \lambda\text{-prepares } W_{\xi'}\}$$

is A -definable. Since $\xi \in \Xi_\xi$, the union of all Ξ_ξ is equal to $\text{RV}^k \times \bigcup_{\lambda \leq 1} \text{RV}_\lambda^\ell$, and then compactness implies that finitely many sets Ξ_{ξ_i} suffice to cover everything. Now $C := \bigcup_i C_{\xi_i}$ is a finite A -definable set that prepares every W_ξ .

Remark 2.6.4. A variant of the compactness argument shows that Proposition 2.6.2 holds even more uniformly, namely: Given an ℓ -h-minimal theory \mathcal{T} (possibly non-complete) and a formula ϕ defining a set

$$W_{K,a} := \phi(K, a) \subset K \times \text{RV}_K^k \times \bigcup_{\lambda \leq 1} \text{RV}_\lambda^\ell$$

for $K \models \mathcal{T}$ and $a \in K^\nu \times \text{RV}_K^{\nu'}$, there exists a formula ψ defining a set $C_{K,a} := \psi(K, a) \subset K$ such that for each model $K \models \mathcal{T}$ and each $a \in K^\nu \times \text{RV}_K^{\nu'}$, this set $C_{K,a}$ uniformly prepares $W_{K,a}$.

Remark 2.6.5. In Proposition 2.6.2, we can also replace RV^k by any Cartesian product Q of sorts from RV^{eq} . Indeed, in that case, just apply the original version of the proposition to the preimage of W under some quotient map $\text{RV}^k \rightarrow Q$. In particular, W can additionally use (arbitrarily many) Γ_K -coordinates, since Γ_K is a quotient of RV . (One can of course also allow RF -coordinates, given that RF can be considered a \emptyset -definable subset of RV .)

In almost all applications of Proposition 2.6.2, we will only need the following special case:

Corollary 2.6.6 (Preparing families). *Assume that $\text{Th}(K)$ is 0-h-minimal and that A is a subset of K . For any $k > 0$ and any $(A \cup \text{RV})$ -definable set*

$$W \subset K \times \text{RV}^k,$$

there exists a finite non-empty A -definable set $C \subset K$ 1-preparing W (in the sense of Definition 2.6.1).

Proof of Corollary 2.6.6. Clear. □

Recall that we set $\text{RV}_\bullet := \bigcup_{\lambda \leq 1} \text{RV}_\lambda$.

Corollary 2.6.7 (RV-unions stay finite).

1. Assume that $\text{Th}(K)$ is 1- h -minimal. For any $k \geq 0$ and any definable (with parameters) set $W \subset \text{RV}^k \times K$ such that the fiber $W_\xi \subset K$ is finite for each $\xi \in \text{RV}^k$, the union $\bigcup_\xi W_\xi$ is also finite.
2. Under the (weaker) assumption that $\text{Th}(K)$ is 0- h -minimal, the previous statement still holds if we assume $W \subset \text{RV}^k \times K$.

Proof. We suppose that $W \subset \text{RV}^k \times \text{RV}^{k'} \times K$ and proceed by induction on k' . If $k' \leq 1$, we obtain both claims of the corollary from Proposition 2.6.2 applied with $\ell = k'$, namely: We find a finite set C such that, for all λ and $\xi \in \text{RV}^k \times \text{RV}^{k'}$, W_ξ is λ -prepared by C . Since W_ξ is finite, we have $W_\xi \subset C$, and hence $\bigcup_\xi W_\xi \subset C$ is finite.

The case of $k' > 1$ now follows by induction on k' and the case $k' = 1$. □

Corollary 2.6.8 (Finite image in K).

1. Assume that $\text{Th}(K)$ is 1- h -minimal. Then the image of any definable (with parameters) function $f: \text{RV}^k \rightarrow K$ for any $k \geq 0$ is finite.
2. Under the (weaker) assumption that $\text{Th}(K)$ is 0- h -minimal, the previous statement still holds for functions $f: \text{RV}^k \rightarrow K$.

Proof. Apply Corollary 2.6.7 to the graph of f . □

Remark 2.6.9. Remark 2.6.5 also applies to Corollaries 2.6.6, 2.6.7 and 2.6.8. In particular, we can additionally allow (arbitrarily many) Γ_K -coordinates in W (in 2.6.6, 2.6.7) and the domain of f (in 2.6.8).

Corollary 2.6.10 (Removing RV-parameters). Assume that $\text{Th}(K)$ is 0- h -minimal. For any $A \subset K$ and any finite $(A \cup \text{RV}^{\text{eq}})$ -definable set $C \subset K$, there exists a finite A -definable set $C' \subset K$ containing C . In other words, $\text{acl}_K(A \cup \text{RV}^{\text{eq}}) = \text{acl}_K(A)$.

Proof. Add constants for A to the language. We have $C = W_{\xi_0}$ for some \emptyset -definable $W \subset K \times \text{RV}^k$ and some $\xi_0 \in \text{RV}^k$. We may assume that all fibers W_ξ have cardinality at most the cardinality of C . Let C' be their union, which is finite by Corollary 2.6.7. □

Remark 2.6.11. Many results in this paper are stated in the form:

- (★) For every \emptyset -definable object X of a certain kind, there exists a finite \emptyset -definable set $C \subset K$ that ‘prepares’ X in some sense (depending on the context).

By Lemma 2.4.1, we get for free that (★) holds more generally, namely if X is $(A \cup \text{RV})$ -definable, for $A \subset K$, we get an $(A \cup \text{RV})$ -definable C ‘preparing’ X . By applying Corollary 2.6.10, we then may even assume that C is A -definable. (It will always be the case that if C prepares X , then so does any set containing C .) Finally, using that the notions of preparation under consideration will be definable, we can apply compactness to get for free that this works uniformly in all models of a non-complete theory, in the same style as in Remark 2.6.4.

We end this subsection by noting that RV is stably embedded in a strong sense (namely, with the RV -parameters being in the definable closure of the original parameters):

Proposition 2.6.12 (Stable embeddedness of RV). Assume that $\text{Th}(K)$ is 0- h -minimal. Then RV is stably embedded in the following strong sense: Given any $A \subset K$, every A -definable set $X \subset \text{RV}^n$ is already $\text{dcl}_{\text{RV}}(A)$ -definable.

Proof. We may assume that A is finite; we do an induction on the cardinality of A .

Let $A = \hat{A} \cup \{a\}$. Then we have an \hat{A} -definable set $Y \subset K \times \text{RV}^n$ such that X is equal to the fiber $Y_a \subset \text{RV}^n$. By applying Corollary 2.6.6 to Y , we find a finite \hat{A} -definable set $C \subset K$ such that either $a \in C$ or, for every $a' \in K$ in the same ball 1-next to C as a , we have $Y_{a'} = Y_a$. Using Lemma 2.5.4, we find an \hat{A} -definable map $f: K \rightarrow \text{RV}^k$ (for some k) whose fibers are exactly the elements of C and the balls 1-next to C . In particular, the set $X = Y_a$ is definable using \hat{A} and $f(a)$ as parameters. Thus we

have $X = Z_{f(a)}$ for some \hat{A} -definable set $Z \subset \text{RV}^k \times \text{RV}^n$. By induction, Z is $\text{dcl}_{\text{RV}}(\hat{A})$ -definable, so X is $\text{dcl}_{\text{RV}}(\hat{A}) \cup \{f(a)\}$ -definable and hence $\text{dcl}_{\text{RV}}(A)$ -definable. \square

Remark 2.6.13. If $\text{Th}(K)$ is ω -h-minimal, there are also various variants of Proposition 2.6.12 involving RV_λ , with similar proofs. For example, building on Remark 2.5.5 instead of Lemma 2.5.4, one obtains that any A -definable subset of $\prod_i \text{RV}_{\lambda_i}$ (for $A \subset K$) is $\text{dcl}_{\text{RV}_\lambda}(A)$ -definable with $\lambda = \min_i \lambda_i$.

2.7. Henselianity of the valued field K

As an analogue of o -minimal fields being real closed, in this subsection, we prove that any equi-characteristic zero-valued field that is Hensel minimal (in any language containing \mathcal{L}_{val}) is Henselian. This is one reason we call our notion ‘Hensel minimality’.

A collection of balls is called nested if, for any two balls in the collection, one is contained in the other.

Lemma 2.7.1 (Definable spherical completeness). *Assuming 0-h-minimality, let $\{B_q \mid q \in Q\}$ be a definable family of nested balls in K , for some non-empty definable set Q in an arbitrary, possibly imaginary sort. Then the intersection $\bigcap_{q \in Q} B_q$ is non-empty.*

Proof. First we suppose that $Q \subset \Gamma_K^\times$ and that each B_q is an open ball of radius q . By Corollary 2.6.6 (and Remark 2.6.9), there exists a finite set C 1-preparing the family of balls B_q . Now one checks that at least one of the following two situations occurs. First: For each $q \in Q$, the intersection of C with B_q is non-empty. Second: The set Q has a minimum $q_0 \in Q$. (Indeed, suppose that the intersection of C with B_{q_0} is empty for some $q_0 \in Q$; then Q contains no $q < q_0$, since B_q would not be 1-prepared by C .) In both situations, the lemma follows.

Finally, we reduce the general case to the case that $Q \subset \Gamma_K^\times$ and that each B_q is an open ball of radius q . To this end, for $\gamma \in \Gamma_K^\times$, let B'_γ be the (necessarily unique) open ball of radius γ containing some B_q ($q \in Q$) if such a ball exists, and let it be the empty set otherwise. Then it is clear that the non-empty B'_γ form a nested definable family of open balls. Moreover, the intersection of the non-empty B'_γ equals the intersection of the B_q (since each B_q is equal to the intersection of all open balls containing B_q). \square

Theorem 2.7.2 (Hensel minimality implies Henselian). *Suppose that K is a valued field of equi-characteristic 0 with 0-h-minimal theory (in a language $\mathcal{L} \supset \mathcal{L}_{\text{val}}$). Then K is Henselian.*

If \mathcal{L} is the pure valued field language, Corollary 6.2.6 implies the converse. Combining, we have, for K of equi-characteristic 0: K is Henselian if and only if $\text{Th}_{\mathcal{L}_{\text{val}}}(K)$ is 0-h-minimal, if and only if $\text{Th}_{\mathcal{L}_{\text{val}}}(K)$ is ω -h-minimal.

Proof of Theorem 2.7.2. Let $P \in \mathcal{O}_K[X]$ be a polynomial such that $P(0) \in \mathcal{M}_K$ and $P'(0) \in \mathcal{O}_K^\times$; we need to prove that P has a root in \mathcal{M}_K . (Note that the uniqueness of such a root then follows automatically.) The idea is to use ‘Newton approximation’ as in the usual proof of Hensel’s lemma for complete discretely valued fields, but where complete and discretely valued is replaced by definably spherically complete.

To make this formal, we suppose that P has no root in \mathcal{M}_K and we set $B_x := B_{\leq |P(x)|}(x)$ for $x \in \mathcal{M}_K$. Note that the B_x form a definable family of balls, parameterized by $x \in \mathcal{M}_K$. We will prove that (a) all these balls form a chain under inclusion and (b) that an element in the intersection of all those balls (which is non-empty by Lemma 2.7.1) is, after all, a root of P .

(a) Let $x_1, x_2 \in \mathcal{M}_K$ be given, and set $\varepsilon := x_2 - x_1$. To see that the balls B_{x_1} and B_{x_2} are not disjoint, we verify that $|\varepsilon| \leq \max\{|P(x_1)|, |P(x_2)|\}$. Taylor expanding P around x_1 yields

$$|P(x_1 + \varepsilon) - P(x_1) - \varepsilon P'(x_1)| \leq |\varepsilon|^2, \tag{2.7.1}$$

which, together with $|P'(x_1)| = 1$, implies $|\varepsilon| \leq \max\{|P(x_1)|, |P(x_1 + \varepsilon)|\}$.

(b) Let x_1 be in the intersection $\bigcap_{x \in \mathcal{M}_K} B_x$, and suppose that $P(x_1) \neq 0$. Then equation (2.7.1) with $\varepsilon := -\frac{P(x_1)}{P'(x_1)}$ implies $|P(x_1 + \varepsilon)| \leq |\varepsilon|^2 < |\varepsilon|$ and hence $x_1 \notin B_{x_1 + \varepsilon}$, contradicting our choice of x_1 . \square

Remark 2.7.3. Lemma 2.7.1 implies a ‘definable Banach Fixed Point Theorem’ (exactly in the form of [43, Lemma 2.32], and with the same proof). The above proof of Theorem 2.7.2 can be considered applying that Fixed Point Theorem to the map $\mathcal{M}_K \rightarrow \mathcal{M}_K, x \mapsto x - P(x)/P'(x)$.

2.8. Definable functions

We continue assuming that K is an equi-characteristic 0 valued field, and we now assume that $\text{Th}(K)$ is 1-h-minimal (unless specified otherwise). Under those assumptions, we now prove the first basic properties of definable functions in one variable; in particular, we already obtain a weak version of the Jacobian Property (Lemma 2.8.5) and simultaneous domain and image preparation (Proposition 2.8.6).

The first result is key to dimension theory (although in our proofs of dimension theory, this will only be used indirectly, namely in the proof of Proposition 5.3.3).

Lemma 2.8.1 (Basic preservation of dimension). *Assume (as convened for the whole Section 2.8) that $\text{Th}(K)$ is 1-h-minimal. Let $f: K \rightarrow K$ be a definable function. Then there are only finitely many function values that are taken infinitely many times.*

Proof. We may assume f to be \emptyset -definable (say, after adding enough parameters from K to the language). Suppose f takes infinitely many values y infinitely many times. Then for each such y , $f^{-1}(y)$ contains a ball (by Lemma 2.5.2). Thus, letting $X \subset K$ be the set of points where f is locally constant, $f(X)$ is still infinite.

Let $W \subset K \times \Gamma_K^\times$ consist of those (x, λ) such that f is constant on $B_{<\lambda}(x)$. This set W is \emptyset -definable, so we find a finite set C 1-preparing W (by Corollary 2.6.6). By enlarging C , we may moreover assume that C also 1-prepares X .

From the fact that $f(X)$ is infinite, we can deduce that there exists a ball $B_0 \subset X$ 1-next to C such that $f(B_0)$ is still infinite. Indeed, letting $g: K \rightarrow \text{RV}^k$ be a map whose fibers are the singletons in C and the balls 1-next to C (using Lemma 2.5.4), if $f(g^{-1}(\xi))$ would be finite for every $\xi \in g(X)$, then so would be $f(X)$ (by Corollary 2.6.7).

Choose $x \in B_0$ and $\lambda_0 \in \Gamma_K^\times$ such that f is constant on $B_{<\lambda_0}(x)$. Since C 1-prepares W , f is constant on $B_{<\lambda_0}(x')$ for every $x' \in B_0$.

Set $\lambda_1 := \lambda_0/\text{rad}_{\text{op}}(B_0)$. Then the family F_1 of open balls of radius λ_0 contained in B_0 can be definably parametrized by a subset of RV_{λ_1} (using some parameters). Indeed, if we fix $c \in K$ such that B_0 is 1-next to c , then each member of F_1 is of the form $c + \text{rv}_{\lambda_1}^{-1}(\xi)$ for some $\xi \in \text{RV}_{\lambda_1}$. We define F_2 to be the family of $f(B)$, for B in F_1 . Then each family member of F_2 is a singleton, yet their union is infinite, contradicting Corollary 2.6.7. □

Using this, we obtain that definable functions are (in a strong sense) locally constant or injective:

Lemma 2.8.2 (Piecewise constant or injective). *For every \emptyset -definable map $f: K \rightarrow K$, there exists a finite \emptyset -definable set C such that for every ball B 1-next to C , f is either constant or injective on B .*

Proof. First, consider the set Y_∞ of $y \in K$ such that $f^{-1}(y)$ is infinite. By Lemma 2.8.1, this set Y_∞ is finite, so we can find a finite \emptyset -definable set C 1-preparing $f^{-1}(y)$ for each $y \in Y_\infty$. Indeed, for each $y \in Y_\infty$, one finds a finite y -definable set C_y 1-preparing $f^{-1}(y)$ (uniformly in y), and one lets C be the union of the sets C_y .

For each $y \in K \setminus Y_\infty$, the set $f^{-1}(y)$ is finite, so by Lemma 2.5.3, there exists a \emptyset -definable family of injective functions $g_y: f^{-1}(y) \rightarrow \text{RV}^k$ for some $k \geq 1$. For convenience, we set $g_y(x) := 0$ if $y \in Y_\infty$ and $x \in f^{-1}(y)$, so that we can define a function $h: K \rightarrow \text{RV}^k$ by $h(x) := g_{f(x)}(x)$. We then enlarge our above set C (using Corollary 2.6.6) so that it also 1-prepares (the graph of) h .

We claim that this set C is as desired, so let B be a ball 1-next to C . If $f(B) \cap Y_\infty \neq \emptyset$, then B is contained in one of the sets $f^{-1}(y)$, for $y \in Y_\infty$, and hence f is constant on B . Otherwise, we use that

h is constant on B to deduce that f is injective on B . Indeed, $f(x_1) = f(x_2) = y$ for some $y \in K \setminus Y_\infty$ implies $h(x_1) = g_y(x_1) = g_y(x_2) = h(x_2)$, so injectivity of h on $f^{-1}(y)$ implies $x_1 = x_2$. \square

The next lemma says that a definable function sends most open balls either to points or to open balls.

Lemma 2.8.3 (Images of balls are balls). *Let $f : K \rightarrow K$ be a \emptyset -definable function. There exists a \emptyset -definable finite set $C \subset K$ such that for every open ball B disjoint from C , $f(B)$ is either a point or an open ball.*

Proof. Define $W \subset K \times \Gamma_K^\times$ to consist of those (x, λ) for which the open ball $B_{<\lambda}(x)$ ‘has a good image’: that is, $f(B_{<\lambda}(x))$ is a singleton or an open ball. By Corollary 2.6.6, we find a finite \emptyset -definable subset C_0 of K such that for any ball B 1-next to C_0 , the fiber $W_x \subset \Gamma_K^\times$ does not depend on $x \in B$.

Fix a ball B_0 1-next to C_0 . We first prove that any open ball B strictly contained in B_0 has a good image. Suppose otherwise, namely that B_1 is an open ball strictly contained in B_0 with bad image. Then the fact that W_x does not depend on $x \in B_0$ implies that every translate of B_1 contained in B_0 also has a bad image. We can find an infinite definable (with parameters) family F_1 consisting of such translates of B_1 and parameterized by a subset of RV . Indeed, the sets $x_0 + rv^{-1}(\xi)$ form such a family for a suitable $x_0 \in K$ and when ξ runs over a suitable subset of RV .

Consider the family F_2 of the sets $f(B)$ for B in F_1 , use Corollary 2.6.6 to find a finite set $D \subset K$ 1-preparing the family F_2 , and let $g : K \rightarrow \text{RV}^k$ be a function whose fibers are the balls 1-next to D and the individual points of D (as obtained using Lemma 2.5.4). Since none of the balls B in F_1 are good, none of the sets $f(B)$ in F_2 are exactly a point or an open ball, so each $f(B)$ consists of several fibers of g ; in other words, $g \circ f : K \rightarrow \text{RV}^k$ is non-constant on every B in F_1 .

Now we get a contradiction by applying Corollary 2.6.6 to the graph of $g \circ f$. Indeed, any set C' 1-preparing that graph would have to contain at least one point in each ball B from F_1 (since $g \circ f$ is not constant on B), so C' cannot be finite. This finishes the proof that balls strictly contained in B_0 have a good image.

The only problematic open balls left (i.e., which are disjoint from C_0 and might have bad image) are the ones 1-next to C_0 . To get hold of those, we run a similar argument as above: We let F_1 be the family of balls 1-next to C_0 ; we find a finite set D 1-preparing $f(B)$ for each B in F_1 ; and we define $g : K \rightarrow \text{RV}^k$ as before, so that in particular if $g \circ f$ is constant on a ball B 1-next to C_0 , then that ball B has good image.

Now we find a finite set C' 1-preparing the graph of $g \circ f$ and we set $C := C_0 \cup C'$. In this way, among the balls 1-next to C_0 , all those that have bad image are not disjoint from C . Note that since this time, F_1 is \emptyset -definable, and hence so are D , g and C' . \square

The following is a key technical lemma, which serves later in the proof of the Jacobian Property. The main point of the statement is that a definable map cannot scale all small balls by one factor and all big balls by a different factor.

Lemma 2.8.4 (Preservation of scaling factor). *Let B be either \mathcal{O}_K or \mathcal{M}_K , and let $f : B \rightarrow K$ be a \emptyset -definable function. Suppose that there are α and β in Γ_K^\times with $\alpha < 1$ such that for every open ball $B' \subset B$ of radius α , $f(B')$ is contained in an open ball of radius β . Then the following hold:*

1. *The image $f(B)$ is contained in a finite union of closed balls of radius β/α .*
2. *If we moreover assume that $B = \mathcal{M}_K$ and that $f(B)$ is an open ball, then $\text{rad}_{\text{op}} f(B) \leq \beta/\alpha$.*

Proof. (1) The family of all open balls of radius α contained in B can be (definably) parametrized by a definable set $\Lambda \subset \text{RV}_\alpha$, namely, set $\Lambda := rv_\alpha(B)$, and for $\xi \in \Lambda$, let B_ξ be the open ball of radius α containing $rv_\alpha^{-1}(\xi)$ (which exists, since $\text{rad}_{\text{op}}(rv_\alpha^{-1}(\xi)) \leq \alpha$). By Proposition 2.6.2, there exists a finite set C such that for each $\xi \in \Lambda$, the set $f(B_\xi)$ is α -prepared by C . This implies that for every open ball B' that is α -next to C and every $\xi \in \Lambda$, one has

$$\text{either } B' \subset f(B_\xi), \quad \text{or } B' \cap f(B_\xi) = \emptyset.$$

Hence, if the radius of the open ball B' is larger than β , then $B' \cap f(B_\xi)$ is empty for all $\xi \in \Lambda$. (Indeed, by assumption, $f(B_\xi)$ is contained in an open ball of radius β , so $f(B_\xi)$ cannot contain the larger ball B' .)

Thus, we find that $f(B)$ is contained in the union of C with those open balls α -next to C that have a radius at most β . This union equals the union over $c \in C$ of the closed balls $B_{\leq \beta/\alpha}(c)$. This proves (1).

(2) If the value group is dense, then (2) follows from (1) using that the largest open balls contained in a finite union of closed balls of radius β/α have radius β/α .

If the value group is discrete (see just above Remark 2.1.5), we apply Part (1) to $g(x) := f(\varpi x) : \mathcal{O}_K \rightarrow K$, where $\varpi \in K$ is a uniformizing element. Since g sends open balls of radius $|\varpi|^{-1} \cdot \alpha$ to open balls of radius β , we obtain that $f(\mathcal{M}_K) = g(\mathcal{O}_K)$ is contained in a finite union of closed balls of radius $|\varpi| \cdot \beta/\alpha$, which is the same as a finite union of open balls of radius β/α . Now the claim follows from the assumption that $f(\mathcal{M}_K)$ is itself an open ball. \square

By combining various previous lemmas, we already obtain a weak form of the Jacobian Property.

Lemma 2.8.5 (Valuative Jacobian Property). *For every \emptyset -definable function $f : K \rightarrow K$, there exists a finite \emptyset -definable set $C \subset K$ such that for every ball B 1-next to C , we have the following: Either f is constant on B or there exists a $\mu_B \in \Gamma_K^\times$ such that*

1. *for every open ball $B' \subset B$, $f(B')$ is an open ball of radius $\mu_B \cdot \text{rad}_{\text{op}}(B')$;*
2. *for every $x_1, x_2 \in B$, we have $|f(x_1) - f(x_2)| = \mu_B \cdot |x_1 - x_2|$.*

Proof. Choose C \emptyset -definable such that for every ball B 1-next to C , we have:

- f is constant or injective on B (using Lemma 2.8.2);
- $f(B')$ is a point or an open ball for every open ball $B' \subset B$ (using Lemma 2.8.3).

Moreover, we assume (using Corollary 2.6.6) that C 1-prepares the graph of the function $r : K \times \Gamma_K^\times \rightarrow \Gamma_K$ defined by

$$r(x, \lambda) = \begin{cases} \text{rad}_{\text{op}}(f(B_{<\lambda}(x))) & \text{if } f(B_{<\lambda}(x)) \text{ is an open ball} \\ 0 & \text{otherwise.} \end{cases}$$

We claim that this C is as desired, so fix a ball B 1-next to C for the remainder of the proof. Note that for x running over B , the function $r(x, \cdot)$ is independent of x , so from now on we write $r(\lambda)$ instead of $r(x, \lambda)$.

If f is constant on B , there is nothing to do, so we may assume that f is injective on B . We claim that the lemma holds with $\mu_B := \text{rad}_{\text{op}}(f(B))/\text{rad}_{\text{op}}(B)$. (Note that indeed, $f(B)$ is an open ball.)

(1) Fix $\alpha \in \Gamma_K$ with $0 < \alpha < \text{rad}_{\text{op}}(B)$. Any open ball B' of radius α is sent to an open ball $f(B')$ of radius $r(\alpha)$. By applying Lemma 2.8.4 (to a suitably rescaled function), we deduce that the radius of the open ball $f(B)$ is at most $\text{rad}_{\text{op}}(B) \cdot r(\alpha)/\alpha$; this implies $r(\alpha)/\alpha \geq \mu_B$.

To get the other inequality, namely $r(\alpha)/\alpha \leq \mu_B$, we apply the same argument to the inverse function $f^{-1} : f(B) \rightarrow B$. This inverse is well-defined since f is injective on B , so it remains to verify that f^{-1} sends open balls $B'' \subset f(B)$ of radius $r(\alpha)$ to open balls of radius α . Indeed, choose any $x \in f^{-1}(B'')$; then $f(B_{<\alpha}(x)) = B''$, and hence $f^{-1}(B'') = B_{<\alpha}(x)$.

(2) For every $x_1, x_2 \in B$, (1) implies $|f(x_1) - f(x_2)| \leq \mu_B \cdot |x_1 - x_2|$. Indeed, applying (1) to a ball of the form $B_{<\alpha}(x_1)$ with $|x_1 - x_2| < \alpha \leq \text{rad}_{\text{op}}(B)$ yields $|f(x_1) - f(x_2)| < \mu_B \cdot \alpha$ for every $\alpha > |x_1 - x_2|$, and hence $|f(x_1) - f(x_2)| \leq \mu_B \cdot |x_1 - x_2|$. The same argument with f replaced by $f^{-1} : f(B) \rightarrow B$ yields the other inequality. \square

Using Lemma 2.8.5, we deduce that the domain and the image of a definable function can be prepared simultaneously and in a compatible way.

Proposition 2.8.6 (Domain and range preparation). *Let $f : K \rightarrow K$ be a \emptyset -definable function, and let $C_0 \subset K$ be a finite, \emptyset -definable set. Then there exist finite, \emptyset -definable sets $C, D \subset K$ with $C_0 \subset C$ such that $f(C) \subset D$, and for every $\lambda \leq 1$ in Γ_K^\times and every ball $B \subset K$ that is λ -next to C , the image $f(B)$ is either a singleton in D or a ball λ -next to D .*

The role of C_0 is to make it possible to combine this proposition with other preparation results for functions (notably Theorem 3.2.2). First apply the other results to get a set C_0 ; then apply Proposition 2.8.6 to enlarge C_0 to C and obtain D .

Note, however, that the proposition cannot be combined very well with itself: Given $f_1, f_2: K \rightarrow K$ as in the proposition, there are, in general, no C, D_1, D_2 such that (f_i, C, D_i) are as in the proposition for both $i = 1, 2$, as the following example shows.

Example 2.8.7. Fix $r \in K$ of negative valuation. For $x \in \mathcal{O}_K$, we set $f_1(x) = f_2(x) = f_1(x+r) = x$ and $f_2(x+r) = x+1$, with the exception that $f_2(0) = r$. Extend f_1 and f_2 by 0 outside of $\mathcal{O}_K \cup (\mathcal{O}_K + r)$. Then one successively deduces: $0 \in C \Rightarrow 0 \in D_1 \Rightarrow r \in C \Rightarrow 1 \in D_2 \Rightarrow 1 \in C \Rightarrow 1 \in D_1 \Rightarrow \dots$; this shows that C cannot be finite.

In Addendum 3 to the cell decomposition Theorem 5.2.4, we will state a version of simultaneous preparation of the domain and image that avoids this problem (and thus works for several functions simultaneously) by working piecewise.

Proof of Proposition 2.8.6. First, enlarge C_0 to a set C_1 using Lemma 2.8.5 so that on each ball B 1-next to C_1 , f sends open balls to open balls and the quotient

$$|f(x_1) - f(x_2)|/|x_1 - x_2| \tag{2.8.1}$$

is constant. Then let D be a set containing $f(C_1)$ and preparing the family $f(B)$, where B runs over the balls 1-next to C_1 . (This family is parametrized by RV-variables, by Lemma 2.5.4, so Corollary 2.6.6 applies.) Now denote by $X \subset K$ the locus where f is not locally constant, and let C be $(X \cap f^{-1}(D)) \cup C_1$. Then clearly C is finite. We claim that these C and D are as desired.

Note that we have $C \supset C_1$, so the only thing that is not clear from the construction is that balls λ -next to C are sent to elements of D or balls λ -next to D . We first deal with the case $\lambda = 1$.

Let B be a ball 1-next to C , and let B_1 be the ball 1-next to C_1 containing B . If $f(B_1)$ is a singleton, then it is an element of D , and we are done. If $f(B_1)$ is not a singleton, then $B_1 \subset X$ by the constancy of equation (2.8.1) on B_1 . If furthermore $f(B_1) \cap D = \emptyset$, then $B_1 \cap C = \emptyset$, so $B = B_1$ and $f(B_1)$ is a ball 1-next to D . Finally, suppose that $f(B_1) \cap D \neq \emptyset$. Then $B_1 \cap C \neq \emptyset$, so B is 1-next to some $c \in C \cap B_1$. Then using that equation (2.8.1) is constant on B_1 , we obtain that $f(B)$ is a ball 1-next to $f(c) \in D$.

Now suppose that $\lambda < 1$. Any ball B λ -next to C is contained in a ball B' 1-next to C . If f is constant on B' , we are done; otherwise, $f(B')$ is a ball 1-next to D , and we deduce that $f(B)$ is λ -next to D , using once more that equation (2.8.1) is constant on B' . □

2.9. An equivalent condition to 1-h-minimality

The conclusions of Lemmas 2.8.1 and 2.8.5 together are actually equivalent to 1-h-minimality. More precisely, we have the following equivalence, where Condition (T1) is slightly weaker than Lemma 2.8.5:

Theorem 2.9.1 (Criterion for 1-h-minimality). *Let \mathcal{L} be a language containing the language \mathcal{L}_{val} of valued fields, and let \mathcal{T} be an \mathcal{L} -theory whose models are valued fields of equi-characteristic 0. Then \mathcal{T} is 1-h-minimal if and only if, for every model K of \mathcal{T} , for every set $A \subset K \cup \text{RV}$ and for every $\mathcal{L}(A)$ -definable $f: K \rightarrow K$, we have the following:*

- (T1) *There exists a finite $\mathcal{L}(A)$ -definable $C \subset K$ such that for every ball $B \subset K$ 1-next to C , there exists a $\mu_B \in \Gamma_K$ such that for every $x_1, x_2 \in B$, we have $|f(x_1) - f(x_2)| = \mu_B \cdot |x_1 - x_2|$.*
- (T2) *The set $\{d \in K \mid f^{-1}(d) \text{ is infinite}\}$ is finite.*

Note that in the above, we intentionally allow C to use the parameters from RV in its definition, in contrast to the condition in the definition of 1-h-minimality.

Remark 2.9.2. The conditions given in Theorem 2.9.1 are very closely related to the tameness notion of Definition 2.1.6 of [13] and its variant from Section 2.1 of [14], one difference being that we do not assume the existence of an angular component map.

By Lemmas 2.8.5 and 2.8.1, (T1) and (T2) follow from 1-h-minimality, so in the remainder of Subsection 2.9, we assume that \mathcal{T} satisfies (T1) and (T2), our goal being to prove 1-h-minimality. We also assume throughout the subsection that K is a model of \mathcal{T} .

Remark 2.9.3. By applying (T1) to the characteristic function of an A -definable set $X \subset K$ (for $A \subset K \cup \text{RV}$), we find a finite A -definable set $C \subset K$ 1-preparing X . In particular, Assumption 2.5.1 is satisfied, so we may apply Lemmas 2.5.2 (\exists^∞ -elimination), 2.5.3 (existence of injective functions from finite sets to RV) and 2.5.4 (the balls 1-next to C are RV -parametrized).

Lemma 2.9.4. *Assume (as convened for the remainder of Subsection 2.9) that \mathcal{T} satisfies (T1) and (T2) from Theorem 2.9.1 and that K is a model. We have the following:*

1. Any definable (with parameters) function $f : \text{RV}^k \rightarrow K$ has a finite image.
2. If $C_\xi \subset K$ is a definable (with parameters) family of finite sets, parametrized by $\xi \in \text{RV}^k$, then the union $\bigcup_{\xi \in \text{RV}^k} C_\xi$ is still finite.

Proof. We first prove both claims for $k = 1$.

(1) The composition $f \circ \text{rv} : K \rightarrow K$ is locally constant everywhere except possibly at 0, so the claim follows from (T2).

(2) By Lemma 2.5.2, the cardinality of C_ξ is bounded independently of ξ , so we may assume that the cardinality is constant: say, equal to m . Let $\sigma_1, \dots, \sigma_m \in \mathbb{Z}[x_1, \dots, x_m]$ be the elementary symmetric functions in m variables, which we consider functions on the set of m -element subsets of K . By (1), for each i , the function $f_i(\xi) := \sigma_i(C_\xi)$ has a finite image. Since $\sigma_1(C), \dots, \sigma_m(C)$ together determine C , there are only finitely many different sets C_ξ , which implies the claim.

Now we deduce (2) for arbitrary k by induction: Given a definable family $C_{\xi, \xi'} \subset K$, for $\xi \in \text{RV}$ and $\xi' \in \text{RV}^k$, we first obtain that $D_{\xi'} := \bigcup_{\xi \in \text{RV}} C_{\xi, \xi'}$ is finite for every ξ' and then that the entire union is finite.

Finally, we obtain (1) for arbitrary k by applying (2) to the family of singletons $C_\xi := \{f(\xi)\}$. □

Lemma 2.9.5. *Given an A -definable function $f : K \rightarrow K$, with $A \subset K \cup \text{RV}$, we can find C as in Condition (T1) of Theorem 2.9.1 that is moreover $(A \cap K)$ -definable. In particular, $\text{Th}(K)$ is 0-h-minimal.*

Proof. Using (T1), we first find an A -definable set C . We consider it a member of an $(A \cap K)$ -definable family of sets C_ξ , parametrized by $\xi \in \text{RV}^k$. By Lemma 2.9.4, the union $C' := \bigcup_{\xi \in \text{RV}^k} C$ is still finite. Moreover, it is $(A \cap K)$ -definable, and since it contains C , it satisfies the requirements of (T1).

For the in-particular part, apply this improved (T1) to the characteristic function of an A -definable set $X \subset K$ (where $A \subset K \cup \text{RV}$), to find a finite $(A \cap K)$ -definable C 1-preparing X . □

Note that we can now use the results from Section 2.6 with $\ell = 0$; we will in particular use Corollary 2.6.6 several times to prepare RV -parametrized families of subsets of K .

The next lemma is a variant of Lemma 2.8.5 with different assumptions.

Lemma 2.9.6. *Given an A -definable function $f : K \rightarrow K$, for $A \subset K \cup \text{RV}$, we can find C that is $(A \cap K)$ -definable, such that Condition (T1) of Theorem 2.9.1 holds, and such that moreover, for $B \subset K$ 1-next to C , for μ_B as in (T1) and for $B' \subset B$ an open ball, the image $f(B')$ is either a singleton (when $\mu_B = 0$) or an open ball of radius $\text{rad}_{\text{op}}(B') \cdot \mu_B$.*

Proof. Let C be an $(A \cap K)$ -definable set satisfying (T1). By Lemma 2.5.4, the family of balls B 1-next to C can be parametrized using RV -parameters, so using Corollary 2.6.6, we can prepare the family of $f(B)$ using a finite $(A \cap K)$ -definable set D . Similarly, we find a finite $(A \cap K)$ -definable set C_0 preparing the family $f^{-1}(B_1)$, where B_1 runs over the balls 1-next to D . We claim that the set $C' := C \cup C_0$ does the job.

So let $B' \subset K$ be 1-next to C' , let B be the ball 1-next to C containing B' , and let $\mu_B \in \Gamma_K$ be as in (T1): that is, such that we have

$$|f(x_1) - f(x_2)| = \mu_B \cdot |x_1 - x_2| \tag{2.9.1}$$

for all $x_1, x_2 \in B$. We suppose that $\mu_B \neq 0$ (otherwise $f(B')$ is a singleton) and we have to show that for every open ball $B'' \subset B'$, $f(B'')$ is an open ball of radius $\text{rad}_{\text{op}}(B'') \cdot \mu_B$.

Let B''_1 be the smallest ball containing $f(B'')$. Using equation (2.9.1), one obtains that B''_1 is an open ball with radius $\text{rad}_{\text{op}} B''_1 = \text{rad}_{\text{op}}(B'') \cdot \mu_B$, so it remains to show that $f(B'')$ is equal to the entire ball B''_1 .

By our definition of C_0 , $f(B')$ (and hence also B''_1) is contained in a ball B_1 that is 1-next to D , and by definition of D (and since B_1 is certainly not disjoint from $f(B)$), B_1 is contained in $f(B)$. In particular, the entire ball B''_1 is contained in $f(B)$. However, using equation (2.9.1) again, we obtain that no element of $B \setminus B''$ can be sent into B''_1 , so we deduce $f(B'') = B''_1$, as desired. \square

Proof of Theorem 2.9.1, ‘ \Leftarrow ’. We are given $\lambda \leq 1$ in Γ_K^\times and an $(A \cup \{\xi\})$ -definable set $X \subset K$, with $A \subset K \cup \text{RV}$ and $\xi \in \text{RV}_\lambda$. We need to find a finite $(A \cap K)$ -definable set $C \subset K$ that λ -prepares X . We may assume that λ is A -definable.

By Lemma 2.9.4 (2), it is enough to find a C that is A -definable. Indeed, then we can get rid of the RV -parameters in the same way as in Lemma 2.9.5.

We consider X a member of an A -definable family X_y , where y runs over K , with $X = X_y$ for every $y \in Y := \text{rv}_\lambda^{-1}(\xi)$. (This uses $\lambda \in \text{dcl}_{\Gamma_K}(A)$.)

By Remark 2.9.3, we find, for each $y \in K$, a finite $(A \cup \{y\})$ -definable set \hat{C}_y that 1-prepares X_y . Using Lemma 2.5.3, we find an $(A \cup \{y\})$ -definable injective map $h_y: \hat{C}_y \rightarrow \text{RV}^k$. By compactness, we may assume that those \hat{C}_y and h_y form A -definable families parametrized by $y \in K$. This allows us to write the set $\bigcup_{y \in K} \{y\} \times \hat{C}_y \subset K^2$ as a disjoint union of graphs of functions $g_\eta: Y_\eta \subset K \rightarrow K$, which form an A -definable family parameterized by $\eta \in \text{RV}^k$, namely: $g_\eta(y) := h_y^{-1}(\eta)$, where the domain Y_η of g_η is the set of those $y \in K$ for which η is in the image of h_y . (For some η , Y_η may be empty.)

For each η , we find a finite $(A \cup \{\eta\})$ -definable set $D_\eta \subset K$ 1-preparing Y_η and $Z_\eta := \{y \in Y_\eta \mid g_\eta(y) \in X_y\}$ and also preparing g_η in the sense of Lemma 2.9.6 (by applying the lemma to g_η extended by 0 outside of Y_η). Using compactness again, we suppose that D_η is an A -definable family parametrized by η so that the union $D := \{0\} \cup \bigcup_\eta D_\eta$ is A -definable. Note that by Lemma 2.9.4 (2), D is finite.

If $D \cap Y \neq \emptyset$, we obtain a finite A -definable set $\bigcup_{y \in D} \hat{C}_y$ that λ -prepares (even 1-prepares) X , and hence we are done. So now assume that $D \cap Y = \emptyset$. In that case, we claim that the following set C λ -prepares X : By Lemma 2.5.4, the balls 1-next to D form an A -definable family $(B_\xi)_\xi$, where ξ runs over RV^k for some $k \geq 1$. We let (using Corollary 2.6.6) $C \subset K$ be a finite A -definable set that 1-prepares $g_\eta(B_\xi)$ for each (η, ξ) satisfying $B_\xi \subset Y_\eta$.

To prove that this C indeed λ -prepares X , we need to verify: For every $x \in X$ and every $x' \in K \setminus X$, there exists a $c \in C$ such that $|x - c| \cdot \lambda \leq |x - x'|$. Let $B_1 \subset K$ be the smallest (closed) ball containing x and x' .

Fix a $y_0 \in Y$. Since \hat{C}_{y_0} 1-prepares X , $\hat{C}_{y_0} \cap B_1$ is non-empty, so there exists an η such that $y_0 \in Y_\eta$ and $g_\eta(y_0) \in B_1$. We fix such an η for the remainder of the proof.

Recall that Y is a ball satisfying $Y \cap D = \emptyset$, and let B_0 the ball 1-next to D containing Y . Our choice of D in particular implies that g_η is defined on all of B_0 and that for μ_{B_0} as in Lemma 2.9.6, $g_\eta(Y)$ and $g_\eta(B_0)$ are open balls satisfying

$$\text{rad}_{\text{op}}(g_\eta(Y)) = \text{rad}_{\text{op}}(Y) \cdot \mu_{B_0} \quad \text{and} \quad \text{rad}_{\text{op}}(g_\eta(B_0)) = \text{rad}_{\text{op}}(B_0) \cdot \mu_{B_0}. \tag{2.9.2}$$

Moreover, since D also 1-prepares $\{y \in Y_\eta \mid g_\eta(y) \in X_y\}$, and since $X = X_y$ for all $y \in Y$, we obtain that $g_\eta(Y)$ is either contained in X or disjoint from X . By our choice of η , we have $g_\eta(Y) \cap B_1 \neq \emptyset$. However, B_1 is neither contained in X nor disjoint from X , so we deduce $B_1 \not\subset g_\eta(Y)$. This implies

$$\text{rad}_{\text{op}}(g_\eta(Y)) \leq \text{rad}_{\text{cl}}(B_1) = |x - x'|. \tag{2.9.3}$$

Since C 1-prepares $g_\eta(B_0)$, there exists a $c \in C$ such that

$$|g_\eta(y_0) - c| \leq \text{rad}_{\text{op}}(g_\eta(B_0)). \tag{2.9.4}$$

Finally, recall that Y is a fiber of rv_λ . Since $0 \notin B_0$ (which we ensured by putting 0 into D), we deduce $\lambda \cdot \text{rad}_{\text{op}}(B_0) \leq \text{rad}_{\text{op}}(Y)$. Putting this together with equations (2.9.2), (2.9.3) and (2.9.4), we obtain:

$$|g_\eta(y_0) - c| \cdot \lambda \leq \mu_{B_0} \cdot \text{rad}_{\text{op}}(B_0) \cdot \lambda \leq \mu_{B_0} \cdot \text{rad}_{\text{op}}(Y) \leq |x - x'|.$$

Now recall that $g_\eta(y_0) \in B_1$, so we obtain the desired result:

$$|x - c| \cdot \lambda \leq \max\{\underbrace{|x - g_\eta(y_0)|}_{\leq |x - x'|}, |g_\eta(y_0) - c|\} \cdot \lambda \leq |x - x'|. \quad \square$$

2.10. A criterion for preparability

We already proved that various kinds of objects ‘living over K ’ can be prepared by finite subsets of K in various senses. We finish Section 2 with a criterion simplifying such proofs (Lemma 2.10.3). Since we want to apply this to quite different notions of ‘being prepared’, we do this at an abstract level, using the following definition.

Definition 2.10.1 (Preparing bad balls). Let \mathcal{B} be a set of closed balls in K (the ‘bad balls’). We say that a finite set $C \subset K$ prepares \mathcal{B} if for every $B \in \mathcal{B}$, the intersection $C \cap B$ is non-empty.

Example 2.10.2. Given a definable subset $X \subset K$, we can let \mathcal{B} be the set of all closed balls B that are neither disjoint from X nor contained in X . Then a finite set $C \subset K$ 1-prepares X if and only if it prepares \mathcal{B} . Indeed, the implication ‘ \implies ’ is clear. For the other implication, suppose that B' is a ball 1-next to C containing both a point $x \in X$ and a point $x' \in K \setminus X$; then the smallest (closed) ball containing x and x' is disjoint from C but lies in \mathcal{B} .

Note that we can (and will) assume \mathcal{B} to be an imaginary definable set. As such, in the above example, \mathcal{B} is definable over the same parameters as X .

By a compactness argument (as in the proof of Lemma 2.4.4), a \emptyset -definable \mathcal{B} can be prepared by a finite \emptyset -definable $C \subset K$ if and only if, in a sufficiently saturated model K , no ball $B \in \mathcal{B}$ is disjoint from $\text{acl}_K(\emptyset)$. The following lemma provides an even weaker condition that needs to be checked if one wants to prove that \mathcal{B} can be prepared:

Lemma 2.10.3 (Criterion for preparability). *Suppose that $\text{Th}(K)$ is 0-h-minimal and that K is $|\mathcal{L}|^+$ -saturated. Let \mathcal{B} be a \emptyset -definable set of closed balls in K . We call an arbitrary ball $B' \subset K$ a ‘bad ball’ if it contains a ball $B \in \mathcal{B}$ as a subset.*

Suppose that for every $a \in K$, there is no open ball $B' \subset K \setminus \text{acl}_K(a)$ that is 1-next to a and bad. Then there exists a finite \emptyset -definable set $C \subset K$ preparing \mathcal{B} .

Remark 2.10.4. It might sound more natural to let the balls in \mathcal{B} be open instead of closed, but that would make the lemma false: We use, in the proof, that any open bad ball already contains a closed bad subball.

Proof of Lemma 2.10.3. Suppose that no finite \emptyset -definable set C prepares \mathcal{B} . Then, by saturation, we find a single ball $B \in \mathcal{B}$ that is disjoint from all finite \emptyset -definable C ; in other words, this ball satisfies $B \cap \text{acl}_K(\emptyset) = \emptyset$. We fix B for the remainder of the proof.

If B is 1-next to some $a \in \text{acl}_K(\emptyset)$, then we also have $B \cap \text{acl}_K(a) = \emptyset$, and we get a contradiction to the assumption that such a B should not be bad. Therefore, if for $a \in \text{acl}_K(\emptyset)$ we denote by B_a the ball 1-next to a containing B , none of the sets $B_a \setminus B$ is empty. By saturation, also the intersection $\bigcap_{a \in \text{acl}_K(\emptyset)} B_a \setminus B$ is non-empty. Let B' be the smallest ball containing B and any chosen element of that intersection. This ball has the following properties: it is closed, it is disjoint from $\text{acl}_K(\emptyset)$, and it is strictly bigger than B . Set $\gamma := \text{rad}_{\text{cl}}(B)$ and $\mu := \text{rad}_{\text{cl}}(B')$.

All elements of B' have the same type over RV ; indeed, if two elements of B' could be distinguished by a formula with parameters from RV , then any finite set C preparing the RV -parametrized family

of subsets of K defined by that formula would have to contain points of B' ; but then C cannot be \emptyset -definable, since $B' \cap \text{acl}_K(\emptyset) = \emptyset$.

We deduce that all the open balls $B_{<\mu}(b) \subset B'$ (for $b \in B'$) are bad. Indeed, for $b \in B$, we have $B_{<\mu}(b) \supset B_{\leq\gamma}(b) \in \mathcal{B}$, and since any other $b' \in B'$ has the same type as b over RV (and \mathcal{B} is \emptyset -definable), we also have $B_{\leq\gamma}(b') \in \mathcal{B}$, witnessing that $B_{<\mu}(b')$ is bad.

Now fix an arbitrary $a \in B'$. By saturation, there exists a $b' \in B'$ such that $B'' := B_{<\mu}(b')$ is disjoint from $\text{acl}_K(a)$. This however contradicts the assumption of the lemma, since B'' is bad and 1-next to a . \square

3. Derivation, the Jacobian Property and Taylor approximation

We continue assuming that K is a valued field of equi-characteristic 0, considered a structure in a fixed language $\mathcal{L} \supset \mathcal{L}_{\text{val}}$, and we assume that $\text{Th}(K)$ is 1-h-minimal. The goal of this section is to prove various preparation results for definable functions $f: K \rightarrow K$, starting with Theorem 3.1.4 that states that f is almost everywhere differentiable, and culminating in Theorem 3.2.2, which states that away from a finite set, f has good approximations by its Taylor polynomials.

3.1. Derivation, strict and classical

We define the derivative of functions $f: K \rightarrow K$ as usual:

Definition 3.1.1 (Classical derivative). We say that the (classical) derivative of a function $f: K \rightarrow K$ exists at $x \in K$ if there exists $a \in K$ such that for each ε in Γ_K^\times there exists δ in Γ_K^\times such that for all $y \in K$ with $|x - y| < \delta$ and $x \neq y$ one has

$$\left| \frac{f(x) - f(y)}{x - y} - a \right| < \varepsilon;$$

we then write $f'(x)$ for a .

In the context of totally disconnected fields, one sometimes needs to put a stronger condition on the existence of derivatives, namely:

Definition 3.1.2 (Strict derivative). We say that the strict derivative of a function $f: K \rightarrow K$ exists at $x \in K$ if there exists $a \in K$ such that for each ε in Γ_K^\times , there exists δ in Γ_K^\times such that for all $y_i \in K$ with $|x - y_i| < \delta$ for $i = 1, 2$ and $y_1 \neq y_2$ one has

$$\left| \frac{f(y_1) - f(y_2)}{y_1 - y_2} - a \right| < \varepsilon.$$

Remark 3.1.3. As usual, one easily verifies that the set of x where the classical derivative of a definable function f exists is definable over the same parameters as f , and similarly for the strict derivative. Moreover, the derivative of f is definable over the same parameters as f .

Theorem 3.1.4 (Existence of derivatives). *Let K be a valued field of equi-characteristic 0 such that $\text{Th}(K)$ is 1-h-minimal, and let $f: K \rightarrow K$ be a definable (with parameters) function. Then the strict derivative of f exists almost everywhere: that is, the set of $x \in K$ such that the strict derivative of f does not exist at x is finite.*

Most of the remainder of this subsection is devoted to the proof of this theorem, so from now on we fix a definable function $f: K \rightarrow K$. By Lemma 2.4.1, we may as well impose f to be \emptyset -definable.

We also fix a handy notation for difference quotients: Given unequal x_1, x_2 in K , we set

$$q_f(x_1, x_2) := \frac{f(x_1) - f(x_2)}{x_1 - x_2}.$$

For convenience, we put $q_f(x, x) := 0$ for any $x \in K$.

We start by proving an auxiliary lemma, which is a weak form of Theorem 3.1.4, stating that in some sense, one has finitely valued derivatives.

Lemma 3.1.5. *For $x \in K$, let A_x be the set of accumulation points of*

$$\lim_{\substack{y_1, y_2 \rightarrow x \\ y_1 \neq y_2}} q_f(y_1, y_2). \tag{3.1.1}$$

There is a finite \emptyset -definable set $C \subset K$ such that, for every $x \in K \setminus C$, A_x is finite and we moreover have

$$\lim_{\substack{y_1, y_2 \rightarrow x \\ y_1 \neq y_2}} \min_{a \in A_x} |q_f(y_1, y_2) - a| = 0. \tag{3.1.2}$$

Proof. Clearly, the set C of x where A_x is not as desired is \emptyset -definable (using Lemma 2.5.2 to express finiteness of A_x). We need to show that C is finite.

We may assume that K is sufficiently saturated, and we suppose that C is infinite. Then C contains a transcendental element x_0 (i.e., $x_0 \notin \text{acl}_K(\emptyset)$). We will prove that for transcendental $x_0 \in K$, there exists a finite set A satisfying equation (3.1.2). One easily checks that this implies that $A_{x_0} \subset A$ and that then A_{x_0} also satisfies equation (3.1.2), contradicting $x_0 \in C$. Constructing A is done in several steps.

Step 1: Using (once more) that K is sufficiently saturated, we find an entire open ball $B_1 := B_{<\mu}(x_0)$ that is disjoint from $\text{acl}_K(\emptyset)$. We fix these μ and B_1 once and for all and more generally define $B_\lambda := B_{<\lambda \cdot \mu}(x_0)$, for $\lambda \leq 1$.

For Steps 2–4, we fix $\lambda \leq 1$ in Γ_K^\times , and ζ will always be an element of RV satisfying $|\zeta| < \lambda \cdot \mu$ (so that $\zeta = \text{rv}(x - x')$ for some $x, x' \in B_\lambda$).

Step 2: For ζ as above, the subset of RV_λ defined by $Q_{x, \zeta} := \{\text{rv}_\lambda(q_f(x, x')) \mid x' \in x + \text{rv}^{-1}(\zeta)\}$ is independent of x when x runs over B_λ .

Proof. We λ -prepare the family $(Q_{x, \zeta})_{x, \zeta}$ using Proposition 2.6.2 (i.e., we consider $Q_{x, \zeta}$ a fiber of a definable set $Q \subset K \times \text{RV} \times \text{RV}_\lambda$ and λ -prepare Q). The finite \emptyset -definable set C' obtained in this way is disjoint from B_1 (since $B_1 \cap \text{acl}_K(\emptyset) = \emptyset$), so B_λ is contained in a ball λ -next to C' . This implies Step 2. □

Set $Q_\zeta := Q_{x, \zeta}$.

Step 3: $Q_\zeta \subset Q_{2\zeta}$ (where 2ζ is $\text{rv}(2) \cdot \zeta$).

Proof. Fix any $\xi \in Q_\zeta$. (We need to show that $\xi \in Q_{2\zeta}$.) Choose $x_1 \in B_\lambda$ witnessing $\xi \in Q_{x_0, \zeta}$: that is, such that $\text{rv}(x_0 - x_1) = \zeta$ and $\text{rv}_\lambda(q_f(x_0, x_1)) = \xi$; similarly, choose $x_2 \in B_\lambda$ such that $\text{rv}(x_1 - x_2) = \zeta$ and $\text{rv}_\lambda(q_f(x_1, x_2)) = \xi$. Then the following computation shows that $\text{rv}_\lambda(q_f(x_0, x_2)) = \xi$ (which implies $\xi \in Q_{2\zeta}$): Set $r_i := x_{i-1} - x_i$ and $s_i := f(x_{i-1}) - f(x_i)$ for $i = 1, 2$. We have

$$\begin{aligned} |q_f(x_0, x_2) - q_f(x_0, x_1)| &= \left| \frac{s_1 + s_2}{r_1 + r_2} - \frac{s_1}{r_1} \right| = \left| \left(\frac{s_2}{r_2} - \frac{s_1}{r_1} \right) \cdot \frac{r_2}{r_1 + r_2} \right| \\ &= \underbrace{|q_f(x_1, x_2) - q_f(x_0, x_1)|}_{< \lambda \cdot |q_f(x_0, x_1)|} \cdot \underbrace{\left| \frac{r_2}{r_1 + r_2} \right|}_{=1} \end{aligned}$$

and hence $\text{rv}_\lambda(q_f(x_0, x_2)) = \text{rv}_\lambda(q_f(x_0, x_1))$. □

Applying Step 3 repeatedly shows: $Q_\zeta \subset Q_{2^n \zeta}$ for every integer $n \geq 0$.

Step 4: Q_ζ is a singleton for every ζ .

Proof. Suppose otherwise: that is, $\xi_1, \xi_2 \in Q_\zeta$ with $\xi_1 \neq \xi_2$. The set

$$X := \{x \in B_\lambda \mid \text{rv}_\lambda(q_f(x_0, x)) = \xi_1\}$$

is definable, and hence it can be 1-prepared by a finite set D . (We do not care about the parameters needed to define X and D .) However, for each of the (disjoint) balls $B_n := \{x \in B_\lambda \mid \text{rv}(x_0 - x) = 2^n \zeta\}$ (where n runs over the non-negative integers), we have neither $X \cap B_n = \emptyset$ (since $\xi_1 \in Q_{2^n \zeta}$) nor $B_n \subset X$ (since $\xi_2 \in Q_{2^n \zeta}$); hence for D to 1-prepare X , we would need $D \cap B_n \neq \emptyset$ for every n , contradicting the finiteness of D . \square

Let us reformulate what we obtained until now in a slightly different way: Given any $\zeta \in \text{RV}$ and any $\lambda \in \Gamma_K$ satisfying $|\zeta|/\mu < \lambda \leq 1$, Step 4 states that the entire set $\tilde{Q}_{\zeta,\lambda} := \{q_f(x, x') \mid x, x' \in B_\lambda, \text{rv}(x - x') = \zeta\}$ is contained in a single ball $\tilde{B}_{\zeta,\lambda}$ λ -next to 0.

Step 5: Using Corollary 2.6.6, choose a finite set $A \subset K$ 1-preparing the family $(\tilde{Q}_{\zeta,\lambda})_{\zeta,\lambda}$, for $\zeta \in \text{RV}$ and $\lambda \in \Gamma_K^\times$ satisfying $|\zeta|/\mu < \lambda \leq 1$. (Again, we do not care about parameters.)

The last step consists of showing that the set A satisfies equation (3.1.2), as desired. More precisely, we show:

Step 6: There exists a constant $\kappa \in \Gamma_K$ such that for every $\lambda < 1$ and every $x, x' \in B_\lambda$ with $x \neq x'$, we have $\min_{a \in A} |q_f(x, x') - a| \leq \lambda \cdot \kappa$.

Proof. The constant is $\kappa := \max\{|a| \mid a \in A\}$. Let $x, x' \in B_\lambda$ be given. For $\zeta := \text{rv}(x - x')$, we have $q_f(x, x') \in \tilde{Q}_{\zeta,\lambda} \subset \tilde{B}_{\zeta,\lambda}$. Note that $\tilde{B}_{\zeta,\lambda}$ is an open ball of radius $\lambda \cdot |q_f(x, x')|$. Since A 1-prepares $\tilde{Q}_{\zeta,\lambda}$, there exists an $a \in A$ such that $|q_f(x, x') - a| \leq \text{rad}_{\text{op}}(\tilde{B}_{\zeta,\lambda}) = \lambda \cdot |q_f(x, x')|$. Since $\lambda < 1$, this in particular implies $|q_f(x, x')| = |a|$, so the right-hand side is $\lambda \cdot |a| \leq \lambda \cdot \kappa$, and we are done. \square

This finishes the proof of the lemma. \square

Using the lemma, we can now prove that the derivative of f exists almost everywhere.

Proof of Theorem 3.1.4. Let A_x (for $x \in K$) and C be as in Lemma 3.1.5. We need to show that, for almost all $x \in K \setminus C$, the set A_x is a singleton. Suppose otherwise: that is, A_x is not a singleton for infinitely many x . Then, as usual, we can find a ball B such that A_x consists of several elements for every $x \in B$.

We first shrink B in such a way that the cardinality $\#A_x$ is constant for $x \in B$, and then we further shrink it to make the map $x \mapsto A_x$ ‘approximately constant’ on B in the following sense: There is a $\mu \in \Gamma_K^\times$ such that for every $x, x' \in B$, the relation $a \sim_\mu a' : \iff |a - a'| < \mu$ defines a bijection between A_x and $A_{x'}$. This shrinking of B is possible as follows: By 1-preparing the (graph of the) function $B \rightarrow \Gamma_K^\times, x \mapsto \min_{a_1, a_2 \in A_x, a_1 \neq a_2} |a_1 - a_2|$, we find a subball (which we will now call B) on which this minimum is constant equal to some $\mu \in \Gamma_K^\times$. We then choose any $x_0 \in B$, and using the definition (3.1.1) of A_{x_0} , we replace B by an even smaller ball around x_0 such a way that for every $x_1, x_2 \in B$, we have

$$\min_{a \in A_{x_0}} \left| \frac{f(x_1) - f(x_2)}{x_1 - x_2} - a \right| < \mu.$$

This implies that \sim_μ defines bijections as desired.

Fix any $a_0 \in \bigcup_{x \in B} A_x$. We apply Lemma 2.8.5 (2) (the Valuative Jacobian Property) to the $\mathcal{L}(a_0)$ -definable function $\tilde{f}(x) := f(x) - a_0 x$. This allows us to further shrink B in such a way that the quotient $|\tilde{f}(x_1) - \tilde{f}(x_2)|/|x_1 - x_2|$ is constant for $x_1, x_2 \in B, x_1 \neq x_2$.

This now leads to a contradiction, as follows. Fix any $x \in B$. For every $a \in A_x$, there exist $x_1, x_2 \in B$ with $|\frac{f(x_1) - f(x_2)}{x_1 - x_2} - a| < \mu$ (by definition of A_x). Since

$$\frac{\tilde{f}(x_1) - \tilde{f}(x_2)}{x_1 - x_2} = \frac{f(x_1) - f(x_2)}{x_1 - x_2} - a_0,$$

this implies $|\frac{\tilde{f}(x_1) - \tilde{f}(x_2)}{x_1 - x_2}| < \mu$ if $a \sim_\mu a_0$ and $|\frac{\tilde{f}(x_1) - \tilde{f}(x_2)}{x_1 - x_2}| \geq \mu$ if $a \not\sim_\mu a_0$. Since A_x contains elements a of both kinds (with and without $a \sim_\mu a_0$), this contradicts $|\frac{\tilde{f}(x_1) - \tilde{f}(x_2)}{x_1 - x_2}|$ being constant. \square

Using the existence of derivatives, we can now reformulate the Valuative Jacobian Property (Lemma 2.8.5) in a nicer way:

Corollary 3.1.6 (Valuative Jacobian Property). *For every \emptyset -definable function $f: K \rightarrow K$, there exists a finite \emptyset -definable set $C \subset K$ such that for every ball B 1-next to C , f' exists on B , $|f'|$ is constant on B , and we have the following:*

1. *For every $x_1, x_2 \in B$, we have $|f(x_1) - f(x_2)| = |f'(x_1)| \cdot |x_1 - x_2|$. (In particular, f is constant on B if $f' = 0$ on B .)*
2. *If $f' \neq 0$ on B , then for every open ball $B' \subset B$, $f(B')$ is an open ball of radius $|f'(x)| \cdot \text{rad}_{\text{op}}(B')$ for any $x \in B$.*

Proof. The only difference between this and Lemma 2.8.5 is that the factor called μ_B in Lemma 2.8.5 is now claimed to be equal to $|f'(x)|$ for any $x \in B$. But indeed, by definition of the derivative, for $x \in B$ and x' sufficiently close to x , we have $|f'(x)| = \frac{|f(x') - f(x)|}{|x' - x|} = \mu_B$. □

3.2. The Jacobian Property and Taylor approximations

We now come to one of the central results of this paper: Every definable function $K \rightarrow K$ is, away from a finite set C , well approximated by its Taylor polynomials. Here follows the precise statement and various variants and corollaries. The proof will be given in the next subsection, and higher-dimensional variants will be deduced in Subsections 5.4 and 5.6.

Definition 3.2.1 (Taylor Polynomial). Let $f: X \subset K \rightarrow K$ be a function that is r -fold differentiable at x_0 for some $r \in \mathbb{N}$ and some $x_0 \in K$. Then we write

$$T_{f,x_0}^{<r+1}(x) := T_{f,x_0}^{\leq r}(x) := \sum_{i=0}^r \frac{f^{(i)}(x_0)}{i!} (x - x_0)^i \tag{3.2.1}$$

for the Taylor polynomial of degree r of f around x_0 . (Here, $f^{(i)}$ denotes the i th derivative of f .) Similarly, when $f: X \subset K^m \rightarrow K$ is r -fold differentiable at $x_0 \in X$, then we write

$$T_{f,x_0}^{<r+1}(x) := T_{f,x_0}^{\leq r}(x) := \sum_{i \in \mathbb{N}^m, |i| \leq r} \frac{f^{(i)}(x_0)}{i!} (x - x_0)^i \tag{3.2.2}$$

for the Taylor polynomial of degree r of f around x_0 (where equation (3.2.2) uses multi-index notation, with $|i| = \sum_{j=1}^m i_j$).

Theorem 3.2.2 (Taylor approximations). *Let $f: K \rightarrow K$ be a \emptyset -definable function, and let $r \in \mathbb{N}$ be given. Then there exists a finite \emptyset -definable set C such that for every open ball B 1-next to C , f is $(r + 1)$ -fold differentiable on B , $|f^{(r+1)}|$ is constant on B , and we have*

$$|f(x) - T_{f,x_0}^{\leq r}(x)| \leq |f^{(r+1)}(x_0)| \cdot |x - x_0|^{r+1} \tag{3.2.3}$$

for every $x_0, x \in B$.

Remark 3.2.3. As explained in Remark 2.6.11, here (in Theorem 3.2.2) and in the following (Corollaries 3.2.5, 3.2.6, 3.2.7), we can as well allow f to be $(A \cup \text{RV})$ -definable for $A \subset K$ and obtain an A -definable set C .

Remark 3.2.4. In fact, Theorem 3.2.2 also comes in a variant with equality instead of weak inequality in equation (3.2.3). Indeed, this variant follows from Corollary 3.2.6 for $r + 1$ and the non-Archimedean triangular property.

As in the reals, we can also express the error in terms of the r th derivative. We mention two different such variants of the theorem. (Both will also play a role in the proof of the theorem.)

Corollary 3.2.5 (Taylor approximations). *Let $f : K \rightarrow K$ be a \emptyset -definable function, and let $r \in \mathbb{N}$ be given. Then there exists a finite \emptyset -definable set C such that for every open ball B 1-next to C , f is r -fold differentiable on B , $|f^{(r)}|$ is constant on B , and we have*

$$|f(x) - T_{f,x_0}^{\leq r}(x)| \leq |f^{(r)}(x_0) \cdot (x - x_0)^{r+1}| / \text{rad}_{\text{op}}(B) \tag{3.2.4}$$

for every $x_0, x \in B$.

Proof. Apply (a) Theorem 3.2.2 to f , (b) Corollary 3.1.6 to $f^{(r)}$ and (c) Corollary 2.6.6 to $\text{rv} \circ f^{(r)}$, and let C be the union of the sets obtained in these ways. Then we have, for B 1-next to C and $x_0 \in B$,

$$|f^{(r+1)}(x_0)| \cdot \text{rad}_{\text{op}}(B) \stackrel{(b)}{=} \text{rad}_{\text{op}}(f^{(r)}(B)) \stackrel{(c)}{\leq} |f^{(r)}(x_0)|.$$

Using this, equation (3.2.3) implies equation (3.2.4). □

Corollary 3.2.6 (Taylor approximations). *Let $f : K \rightarrow K$ be a \emptyset -definable function, and let $r \in \mathbb{N}$ be given. Then there exists a finite \emptyset -definable set C such that for every open ball B 1-next to C , f is r -fold differentiable on B , $|f^{(r)}|$ is constant on B , and we either have*

$$|f(x) - T_{f,x_0}^{\leq r}(x)| < |f^{(r)}(x_0) \cdot (x - x_0)^r| \tag{3.2.5}$$

or $f(x) = T_{f,x_0}^{\leq r}(x)$ (if $|f^{(r)}(x_0)| = 0$) for every $x_0, x \in B$.

Proof. This follows directly from Corollary 3.2.5, using that $|x - x_0| < \text{rad}_{\text{op}}(B)$. □

Note that in the case $r = 1$, we in particular get back the RV_λ -Jacobian Property, studied in an analytic setting in [18, Theorem 6.3.7, Remark 6.3.16]. More precisely, the following version of the Jacobian Property is even uniform for all $\lambda \leq 1$.

Corollary 3.2.7 (Jacobian Property). *Let $f : K \rightarrow K$ be a \emptyset -definable function. Then there exists a finite \emptyset -definable set C such that for every $\lambda \leq 1$ in Γ_K^\times , for every ball B λ -next to C and every x_0 and x in B with $x \neq x_0$, we have:*

1. The derivative f' exists on B and $\text{rv}_\lambda \circ f'$ is constant on B .
2. $\text{rv}_\lambda(\frac{f(x) - f(x_0)}{x - x_0}) = \text{rv}_\lambda(f')$.
3. For every open ball $B' \subset B$, $f(B')$ is either a point or an open ball.

Definition 3.2.8 (Jacobian Property). If Conditions (1)–(3) of Corollary 3.2.7 hold, we say that $f|_B$ has the RV_λ -Jacobian Property (or just *Jacobian Property*, in the case $\lambda = 1$).

Proof of Corollary 3.2.7. Choose C using Corollary 3.2.5 (applied to f , with $r = 1$); this will ensure (2), as we shall see below.

To obtain that $\text{rv}_\lambda \circ f'$ is constant on balls λ -next to C , enlarge C using Proposition 2.6.2: Consider the graph of $\text{rv}_\lambda \circ f'$ a subset of $K \times \text{RV}_\lambda \subset K \times \text{RV}_\bullet$ and take the union of all of those as the set W in the proposition.

Finally, to obtain that $f(B')$ is either a point or a ball, enlarge C once more using Lemma 2.8.3.

It remains to verify (2), so fix $\lambda \leq 1$, fix a ball B λ -next to C , fix $x_0, x \in B$, and let B' denote the ball 1-next to C containing B . Then we have

$$|f(x) - f(x_0) - f'(x_0)(x - x_0)| \leq |f'(x_0) \cdot (x - x_0)^2| / \text{rad}_{\text{op}}(B'). \tag{3.2.6}$$

Since $\text{rad}_{\text{op}}(B) = \lambda \cdot \text{rad}_{\text{op}}(B')$, we have $|x - x_0| < \lambda \cdot \text{rad}_{\text{op}}(B')$, so after dividing both sides of equation (3.2.6) by $|x - x_0|$, we obtain

$$|\frac{f(x) - f(x_0)}{x - x_0} - f'(x_0)| < |f'(x_0)| \cdot \lambda,$$

which is equivalent to (2). □

3.3. Proof of Taylor approximation

Before we dive into the proof of Theorem 3.2.2, here is a somewhat technical lemma that will be needed for one special case.

Lemma 3.3.1 (Preparing equivalence relations). *Suppose that Γ_K is discrete. Fix an open ball $B \subset K$ that is disjoint from $\text{acl}(\emptyset)$. Suppose that \sim is an \emptyset -definable binary relation on K with the following properties: (i) the restriction of \sim to $B \times B$ is an equivalence relation with finitely many equivalence classes; (ii) every ball B' strictly contained in B (i.e., $B' \subsetneq B$) is contained in a single equivalence class. Then the entire ball B is a single equivalence class.*

Proof. We write $k := \mathcal{O}_K/\mathcal{M}_K$ for the residue field and $\text{res}: \mathcal{O}_K \rightarrow k$ for the residue map, and we choose a linear bijection $f: B \rightarrow \mathcal{O}_K$. (Note that since Γ_K is discrete, the open ball B is also closed.) Then \sim induces an equivalence relation \sim_k on k (satisfying $x \sim x' \iff \text{res}(f(x)) \sim_k \text{res}(f(x'))$) for $x, x' \in B$).

For $x \in K$, set $W_x := \{\text{rv}(x' - x) \mid x' \in K, x' \sim x\}$. Since B is disjoint from $\text{acl}_K(\emptyset)$, there exists a single set $W \subset \text{RV}$ such that $W_x = W$ for all $x \in B$ (by Corollary 2.6.6). This set W yields a set $T \subset k$ such that $z \sim_k z'$ if and only if $z - z' \in T$ (for $z, z' \in k$).

A straightforward computation shows: The fact that ' $z - z' \in T$ ' defines an equivalence relation on k implies that T is a \mathbb{Z} -submodule of k . The number of equivalence classes is equal to the cardinality of the \mathbb{Z} -module quotient k/T , so to finish the proof of the lemma,¹ suppose that this quotient has finite cardinality $N > 1$. Then for any $z \in k$, we have $Nz \in T$. However, this fails if we choose any $z' \in k \setminus T$, and set $z := \frac{1}{N}z'$. \square

We now come to the proof of Theorem 3.2.2. It goes by induction on r , where Corollaries 3.2.5 and 3.2.6 are used as intermediate steps. More precisely, Theorem 3.2.2 follows from the following four lemmas:

Lemma 3.3.2. *Theorem 3.2.2 holds for $r = 0$.*

Lemma 3.3.3. *For any $r \geq 1$, if Theorem 3.2.2 holds for $r - 1$ (for every 1-h-minimal theory) and Corollary 3.2.6 holds for all $r' < r$, then Corollary 3.2.6 holds for r .*

Lemma 3.3.4. *For any $r \geq 1$, if Corollary 3.2.6 holds for r (for every 1-h-minimal theory), then Corollary 3.2.5 holds for r .*

Lemma 3.3.5. *For any $r \geq 1$, if Corollary 3.2.5 holds for r (for every 1-h-minimal theory), then Theorem 3.2.2 holds for r .*

For all of the proofs, first note that the condition about sufficient differentiability of f on $K \setminus C$ appearing in the theorem and the corollaries is easily obtained using Theorem 3.1.4; this will not be further mentioned. Similarly, $|f^{(r)}|$ (or $|f^{(r+1)}|$) can easily be made constant on balls 1-next to C using Corollary 2.6.6.

Proof of Lemma 3.3.2. In the case $r = 0$, equation (3.2.3) becomes $|f(x) - f(x_0)| \leq |f'(x_0) \cdot (x - x_0)|$, so the theorem follows from Corollary 3.1.6. \square

The proofs of the other three lemmas are long and technical, but the main idea is simple (and the same in all three lemmas): Given a function f that we want to control on a ball B , we define $g(x) := f(x) - ax^r$ in such a way that the r th derivative of g is small on B . In this way, applying the inductive assumption to g yields a particularly strong bound, which will be good enough to obtain the desired bound on f .

The difficulty with this approach is that we need to control the parameters over which g is definable. A powerful ingredient for this is Lemma 2.10.3, which allows us to use points near B as parameters. Nevertheless, various additional technical tricks are needed (different for each of the lemmas) to make the proofs work.

¹The authors would like to thank W. Singhof for providing the argument at the end the proof of Lemma 3.3.1.

Proof of Lemma 3.3.3. Namely:

$$\underbrace{|f - T_{f,x_0}^{\leq r-1}| \leq |f^{(r)} \cdot (x - x_0)^r|}_{\text{equation(3.2.3) for } r-1} \quad \underbrace{|f - T_{f,x_0}^{\leq r}| < |f^{(r)} \cdot (x - x_0)^r|}_{\text{equation(3.2.5)}}$$

We assume that K is sufficiently saturated. Then it suffices to prove that on every open ball $B \subset K$ that is disjoint from $\text{acl}(\emptyset)$, either equation (3.2.5) holds or we have $f(x) = T_{f,x_0}^{\leq r}(x)$ (by the compactness argument from Lemma 2.4.4), so fix such a B , and suppose that $x_0, x \in B$ violate equation (3.2.5).

We may assume that $f^{(r)} \neq 0$ on B , since otherwise, Theorem 3.2.2 for $r - 1$ implies $f(x) = T_{f,x_0}^{\leq r-1}(x) = T_{f,x_0}^{\leq r}(x)$ for any $x_0, x \in B$, and we are done.

Case 1: There exists an open ball $B' := B_{<\delta}(x_0)$ containing x that is strictly smaller than B .

We fix the above δ for the remainder of the proof of Case 1. Also, fix any $a \in B$ and choose any open ball $B'' = B_{<\delta}(x'_0) \subset B$ of radius δ disjoint from $\text{acl}(a)$. (Such a B'' exists by saturation.) Since x_0 and x'_0 have the same type over Γ_K (by Lemma 2.4.4), there exists an $x' \in B''$ such that x'_0, x' violate equation (3.2.5).

We apply Theorem 3.2.2 for $r - 1$ to the function

$$g(x) := f(x) - \frac{f^{(r)}(a)}{r!} \cdot x^r.$$

Since g is a -definable and B'' is disjoint from $\text{acl}(a)$, we obtain

$$|g(x') - T_{g,x'_0}^{\leq r-1}(x')| \leq |g^{(r)}(x'_0)(x' - x'_0)^r|. \tag{3.3.1}$$

The definition of g has been chosen such that $g^{(r)}(x) = f^{(r)}(x) - f^{(r)}(a)$ for all $x \in B$, so using that $\text{rv}(f^{(r)})$ is constant on B (this uses Corollary 2.6.6 and that $B \cap \text{acl}_K(\emptyset) = \emptyset$), we deduce

$$|g^{(r)}(x)| < |f^{(r)}(x)| \tag{3.3.2}$$

for all $x \in B$. From this, we obtain the following (explanation of (\star) below):

$$\begin{aligned} |f(x') - T_{f,x'_0}^{\leq r}(x')| &\stackrel{(\star)}{=} |g(x') - T_{g,x'_0}^{\leq r}(x')| \\ &= |g(x') - T_{g,x'_0}^{\leq r-1}(x') + \frac{1}{r!}g^{(r)}(x'_0)(x' - x'_0)^r| \\ &\leq \max\{|g(x') - T_{g,x'_0}^{\leq r-1}(x')|, \frac{1}{r!}|g^{(r)}(x'_0)(x' - x'_0)^r|\} \\ &\stackrel{\text{equation(3.3.1)}}{\leq} |g^{(r)}(x'_0)(x' - x'_0)^r| \stackrel{\text{equation(3.3.2)}}{<} |f^{(r)}(x'_0) \cdot (x' - x'_0)^r|. \end{aligned} \tag{3.3.3}$$

(\star) : f and g differ by a polynomial of degree r , so their Taylor approximations differ by the same polynomial.

Thus equation (3.2.5) holds for x'_0, x' , which is a contradiction and hence finishes the proof of Case 1.

Case 2: The only open ball $B' \subset B$ containing both x_0 and x is already B itself.

Note that this can only happen if Γ_K is discrete (otherwise the radius of B' can be taken strictly between $|x_0 - x|$ and $\text{rad}_{\text{op}}(B)$), and it also means that $|x_0 - x| = \text{rad}_{\text{cl}}(B) =: \delta$.

By Case 1, we may assume that equation (3.2.5) holds on every proper subball of B . Moreover, we apply Corollary 3.2.6 inductively to derivatives of f to get, for $1 \leq i \leq r$:

$$|f^{(i)}(x) - T_{f^{(i)},x_0}^{\leq r-i}(x)| < |f^{(r)}| \cdot |x - x_0|^{r-i} \leq |f^{(r)}| \cdot \delta^{r-i} \tag{3.3.4}$$

for every $x_0, x \in B$. (Note that $|f^{(r)}|$ is constant on B ; here and in the sequel, we just write $|f^{(r)}|$ instead of $|f^{(r)}(x)|$ for $x \in B$.)

Step 2.1: Set

$$d(x_0, x) := T_{f, x_0}^{\leq r}(x) - f(x_0) = \sum_{i=1}^r \frac{f^{(i)}(x_0)}{i!} (x - x_0)^i$$

for $x_0, x \in B$. Then, for any $x_1, x_2, x_3 \in B$, we have

$$|d(x_1, x_2) + d(x_2, x_3) - d(x_1, x_3)| < |f^{(r)}| \cdot |x_2 - x_3| \cdot \delta^{r-1}. \tag{3.3.5}$$

Proof. This is a straightforward computation consisting of using equation (3.3.4) to approximate the derivatives appearing in $d(x_2, x_3)$ by their Taylor series around x_1 . The details are as follows. In the computation, ‘ \approx ’ means that the norm of the difference is less than $|f^{(r)}| \cdot |x_2 - x_3| \cdot \delta^{r-1}$:

$$\begin{aligned} d(x_2, x_3) &= \sum_{i=1}^r \frac{f^{(i)}(x_2)}{i!} (x_3 - x_2)^i \approx \sum_{i=1}^r \sum_{j=0}^{r-i} \frac{f^{(i+j)}(x_1)}{i! \cdot j!} (x_2 - x_1)^j (x_3 - x_2)^i \\ &= \sum_{k=1}^r \underbrace{\sum_{\substack{i+j=k, \\ i>0}} \frac{k!}{i! \cdot j!} (x_2 - x_1)^j (x_3 - x_2)^i}_{= ((x_2-x_1)+(x_3-x_2))^k - (x_2-x_1)^k (x_3-x_2)^0} \frac{f^{(i+j)}(x_1)}{k!} \\ &= \sum_{k=1}^r ((x_3 - x_1)^k - (x_2 - x_1)^k) \frac{f^{(k)}(x_1)}{k!} = d(x_1, x_3) - d(x_1, x_2). \end{aligned} \quad \square$$

For $x_1, x_2 \in K$, define

$$x_1 \sim x_2 : \iff |f(x_2) - f(x_1) - d(x_1, x_2)| < |f^{(r)}| \cdot \delta^r$$

(with some arbitrary convention if derivatives at $x_1 \notin B$ do not exist). This relation is definable using only the parameter $|f^{(r)}| \cdot \delta^r \in \Gamma$. Our aim is to verify the prerequisites of Lemma 3.3.1 (in the language where $|f^{(r)}| \cdot \delta^r$ has been added as a constant), but before that, let us verify that this then finishes the proof: The lemma then implies that all elements of B are \sim -equivalent. In particular, for our original $x_0, x \in B$ satisfying $|x_0 - x| = \delta$, we obtain

$$|f(x) - T_{f, x_0}^r(x)| = |f(x) - f(x_0) - d(x_0, x)| < |f^{(r)}| \cdot \delta^r, \tag{3.3.6}$$

that is, equation (3.2.5) holds, as desired.

Step 2.2: The restriction of \sim to $B \times B$ is an equivalence relation.

Proof. Fix $x_0 \in B$ and define, for $x \in B$,

$$\tilde{f}(x) := f(x) - d(x_0, x). \tag{3.3.7}$$

Then, using ‘ \approx ’ to mean that the difference has norm less than $|f^{(r)}| \cdot \delta^r$, we have:

$$\begin{aligned} \tilde{f}(x_2) - \tilde{f}(x_1) &= f(x_2) - f(x_1) - d(x_0, x_2) + d(x_0, x_1) \\ &\stackrel{\text{equation(3.3.5)}}{\approx} f(x_2) - f(x_1) - d(x_1, x_2). \end{aligned} \tag{3.3.8}$$

So $x_1 \sim x_2$ if and only if $\tilde{f}(x_2) \approx \tilde{f}(x_1)$, which is clearly an equivalence relation. □

Step 2.3: Each proper subball $B' \subsetneq B$ is contained in a single equivalence class of \sim .

Proof. This follows from our assumption that f satisfies equation (3.2.5) on B' (using a similar computation as in equation (3.3.6)). □

Step 2.4: \sim has only finitely many equivalence classes on B .

Proof. Consider \tilde{f} as in equation (3.3.7). If $x_1, x_2 \in B$ satisfy $|x_1 - x_2| \leq |\varpi| \cdot \delta$ (where $\varpi \in \mathcal{O}_K$ is a uniformizer), then in the ' \approx ' of equation (3.3.8), the difference is less than $|f^{(r)}| \cdot |x_1 - x_2| \cdot \delta^{r-1} \leq |f^{(r)}| \cdot |\varpi| \cdot \delta^r$, and the right-hand side of equation (3.3.8) satisfies

$$|f(x_2) - f(x_1) - d(x_1, x_2)| = |f(x_2) - T_{f,x_1}^r(x_2)| \stackrel{\text{Case 1}}{<} |f^{(r)}| \cdot |x_1 - x_2|^r.$$

Thus $|\tilde{f}(x_2) - \tilde{f}(x_1)| < |f^{(r)}| \cdot |\varpi| \cdot \delta^r$: that is, for any open ball $B' \subset B$ of (open) radius δ , $\tilde{f}(B')$ is contained in an open ball of (open) radius $|f^{(r)}| \cdot |\varpi| \cdot \delta^r$. By Lemma 2.8.4, $\tilde{f}(B)$ is therefore contained in the union of finitely many closed balls of (closed) radius $|f^{(r)}| \cdot |\varpi| \cdot \delta^r$. By the characterization of \sim from Step 2.2, each such closed ball corresponds exactly to one equivalence class of \sim , so we are done. □

These were all the prerequisites needed for Lemma 3.3.1, so this finishes the proof of Case 2 and hence of the entire Lemma. □

Proof of Lemma 3.3.4. Namely:

$$\underbrace{|f - T_{f,x_0}^{\leq r}| < |f^{(r)}| \cdot (x - x_0)^r}_{\text{equation (3.2.5)}} \quad \underbrace{|f - T_{f,x_0}^{\leq r}| \leq |f^{(r)}| \cdot (x - x_0)^{r+1}/\text{rad}_{\text{op}}(B)}_{\text{equation (3.2.4)}}$$

We assume that K is sufficiently saturated.

Step 1: The corollary follows if we can verify that for every $a \in K$, inequality (3.2.4) holds for every ball $B \subset K \setminus \text{acl}_K(a)$ that is 1-next to a .

Proof. An easy computation shows that given any open ball B , inequality (3.2.4) holds on B if and only if for every closed subball $B' \subset B$, and every $x_0, x \in B'$, we have the following corresponding strict inequality:

$$|f(x) - T_{f,x_0}^{\leq r}(x)| < |f^{(r)}(x_0) \cdot (x - x_0)^{r+1}|/\text{rad}_{\text{cl}}(B'). \tag{3.3.9}$$

In particular, for a and B as in the assumption of Step 1, every closed subball of B satisfies equation (3.3.9). This allows us to apply Lemma 2.10.3, where we take \mathcal{B} to be the set of closed balls on which equation (3.3.9) fails. The lemma yields a finite \emptyset -definable set $C \subset K$ intersecting each ball in \mathcal{B} and hence also intersecting (as desired) each open ball where equation (3.2.4) does not hold. □

For the remainder of the proof, let $a \in K$ be given, and let B be a ball that is disjoint from $\text{acl}_K(a)$ and 1-next to a . We may assume that $f^{(r)}$ is nowhere 0 on B , since otherwise, it would be 0 on all of B (since $B \cap \text{acl}_K(\emptyset) = \emptyset$) and Corollary 3.2.6 would yield $f(x) = T_{f,x_0}^{\leq r}(x)$ for $x \in B$.

Suppose that equation (3.2.4) fails on B . Choose $x, x_0 \in B$ witnessing this failure and fix, for the remainder of the proof,

$$\delta := |x - x_0| \tag{3.3.10}$$

and

$$\alpha := \frac{|f(x) - T_{f,x_0}^{\leq r}(x)|}{|f^{(r)}(x_0) \cdot (x - x_0)^{r+1}|/\text{rad}_{\text{op}}(B)}; \tag{3.3.11}$$

in other words, $\alpha > 1$ is the factor by which equation (3.2.4) fails for x_0, x . Moreover, set

$$\gamma := \min\{\delta \cdot \alpha, \text{rad}_{\text{op}}(B)\}$$

(so that the ball $B_{<\gamma}(x_0)$ is contained in B and is not much bigger than the smallest ball containing x_0 and x).

Step 2: There exists an open ball $B' \subset B$ of radius γ containing a (δ, α) -counter-example to equation (3.2.4) with the additional properties that B' is 1-next to some $a' \in K$ and $B' \cap \text{acl}_K(a, a') = \emptyset$. By a (δ, α) -counter-example to equation (3.2.4), we mean a pair $x_0, x \in B'$ with $|x - x_0| = \delta$ satisfying

$$|f(x) - T_{f,x_0}^{\leq r}(x)| \geq \alpha \cdot |f^{(r)}(x_0) \cdot (x - x_0)^{r+1}| / \text{rad}_{\text{op}}(B). \tag{3.3.12}$$

Proof. If $\gamma = \text{rad}_{\text{op}}(B)$, we can simply take $B' = B$ and $a' = a$, so now suppose that $\gamma < \text{rad}_{\text{op}}(B)$.

We claim that the fact that B contains a (δ, α) -counter-example already implies that each open ball $B_{<\gamma}(z) \subset B$ of radius γ contains a (δ, α) -counter-example. Indeed, the set

$$Z := \{z \in K \mid \exists x_0, x \in B_{<\gamma}(z) : (x_0, x) \text{ is a } (\delta, \alpha)\text{-counter-example}\}$$

is definable using only some value group parameters (namely α, δ and $\text{rad}_{\text{op}}(B)$), so since $B \cap \text{acl}_K(\emptyset) = \emptyset$, we have either $B \subset Z$ or $B \cap Z = \emptyset$. In particular, the existence of a single (δ, α) -counter-example in B already implies $B \subset Z$, which proves our claim.

Now fix any $a' \in B$ and fix an open ball B' with $\text{rad}_{\text{op}}(B') = \gamma$, which is 1-next to a' and such that $B' \cap \text{acl}_K(a, a') = \emptyset$; such a B' exists by saturation of K . Since $\gamma < \text{rad}_{\text{op}}(B)$, we have $B' \subset B$ and by the previous paragraph, B' contains a (δ, α) -counter-example. Hence a' and B' are as desired. \square

The remainder of the proof now consists of proving that there is no (δ, α) -counter-example on B' (so that we have a contradiction). More precisely, fix $x_0, x \in B'$ with $|x_0 - x| = \delta$; our goal is to prove that equation (3.3.12) does not hold.

Step 3: There exists a $d \in \text{acl}_K(a')$ such that the image $f^{(r)}(B')$ is either a ball 1-next to d or equal to the singleton $\{d\}$.

Proof. Apply Proposition 2.8.6 to compatibly prepare the domain and the image of $f^{(r)}$ using finite a' -definable sets C and D , where we additionally impose $a' \in C$. Since $B' \cap \text{acl}_K(a') = \emptyset$, B' is a ball 1-next to C , and the claim follows. \square

Step 4: We apply Corollary 3.2.6 to the function

$$g(x) := f(x) - \frac{d}{r!} \cdot x^r.$$

Since g is $\{d\}$ -definable and $B' \cap \text{acl}_K(d) = \emptyset$, we obtain

$$|f(x) - T_{f,x_0}^{\leq r}(x)| = |g(x) - T_{g,x_0}^{\leq r}(x)| < |g^{(r)}(x_0) \cdot (x - x_0)^r| \tag{3.3.13}$$

(where x_0 and x are as fixed above Step 3). The first equality follows from the fact that f and g differ by a polynomial of degree r , so their r th Taylor approximations differ by the same polynomial.

Step 5a: If $f^{(r)}$ is non-constant on B' , we obtain

$$|f(x) - T_{f,x_0}^{\leq r}(x)| < \text{rad}_{\text{op}}(f^{(r)}(B')) \cdot |x - x_0|^r. \tag{3.3.14}$$

Proof. By definition of g and by our choice of d (in Step 3), we have

$$|g^{(r)}(x_0)| = |f^{(r)}(x_0) - d| = \text{rad}_{\text{op}}(f^{(r)}(B')). \tag{3.3.15}$$

Combining this with equation (3.3.13) yields equation (3.3.14). \square

Step 5b: If $f^{(r)}$ is constant on B' , we obtain

$$|f(x) - T_{f,x_0}^{\leq r}(x)| = 0.$$

Proof. Instead of equation (3.3.15), we have $g^{(r)}(x_0) = f^{(r)}(x_0) - d = 0$; apply this in the same way as in Step 5a. □

In the constant case (as in Step 5b), we are already done for the lemma, since the right-hand side of equation (3.3.12) is non-zero, a contradiction. (Recall that we assume $f^{(r)} \neq 0$.)

The last ingredient for the non-constant case is the following:

Step 6: We have $\text{rad}_{\text{op}}(f^{(r)}(B')) \leq \alpha \cdot |f^{(r)}(x_0)| \cdot |x - x_0|/\text{rad}_{\text{op}}(B)$.

Proof. Using that $\text{rv}(f^{(r)})$ is constant on B , we obtain $\text{rad}_{\text{op}}(f^{(r)}(B)) \leq |f^{(r)}(x_0)|$. From this and by applying Lemma 2.8.5 to $f^{(r)}$, we deduce that $\text{rad}_{\text{op}}(f^{(r)}(B')) \leq |f^{(r)}(x_0)| \cdot \text{rad}_{\text{op}}(B')/\text{rad}_{\text{op}}(B)$. Combining this with $\text{rad}_{\text{op}}(B') = \gamma \leq \alpha \cdot |x - x_0|$ yields the claim. □

Now Steps 5a and 6 together imply that equation (3.3.12) fails, as desired, so we are done. □

Proof of Lemma 3.3.5. Namely:

$$\underbrace{|f - T_{f,x_0}^{\leq r}| \leq |f^{(r)} \cdot (x - x_0)^{r+1}|/\text{rad}_{\text{op}}(B)}_{\text{equation (3.2.4)}} \quad \underbrace{|f - T_{f,x_0}^{\leq r}| \leq |f^{(r+1)} \cdot (x - x_0)^{r+1}|}_{\text{equation (3.2.3)}}$$

Let \mathcal{B} be the set of all closed balls on which equation (3.2.3) does not hold. The strategy is to use Lemma 2.10.3 to find a finite \emptyset -definable C meeting every $B \in \mathcal{B}$. Note that then we are done, since if equation (3.2.3) fails for some $x, x_0 \in B'$, where B' is a ball 1-next to C , then equation (3.2.3) also fails on a ball from \mathcal{B} , namely the smallest (closed) ball containing x and x_0 .

So as needed for Lemma 2.10.3, let $a \in K$ be given, and let B be an open ball 1-next to a satisfying $B \cap \text{acl}_K(a) = \emptyset$. We need to verify that equation (3.2.3) holds on B .

By applying Proposition 2.8.6 to $f^{(r)}$ and to a set C_0 containing a , we find a $d \in \text{acl}_K(a)$ such that either $f^{(r)}(B) = \{d\}$ or $f^{(r)}(B)$ is a ball 1-next to d .

Now apply Corollary 3.2.5 to $g(x) := f(x) - \frac{d}{r!}x^r$. Since B is disjoint from the algebraic closure of the parameters used to define g , we obtain

$$|f(x) - T_{f,x_0}^{\leq r}(x)| = |g(x) - T_{g,x_0}^{\leq r}(x)| \leq |g^{(r)}(x_0) \cdot (x - x_0)^{r+1}|/\text{rad}_{\text{op}}(B) \tag{3.3.16}$$

for every $x, x_0 \in B$. As before, the first equality holds because f and g differ by a polynomial of degree r . To finish the proof, it now remains to bound the right-hand side of equation (3.3.16) by $|f^{(r+1)}(x_0) \cdot (x - x_0)^{r+1}|$.

If $f^{(r)}(B) = \{d\}$, then $g^{(r)}(x_0) = 0$, and we are done. Otherwise, $g^{(r)}(B)$ is a ball 1-next to 0, so $\text{rad}_{\text{op}}(g^{(r)}(B)) = |g^{(r)}(x_0)|$. Moreover, by Corollary 3.1.6 applied to $g^{(r)}$, we have $\text{rad}_{\text{op}}(g^{(r)}(B)) = \text{rad}_{\text{op}}(B) \cdot |g^{(r+1)}(x_0)| = \text{rad}_{\text{op}}(B) \cdot |f^{(r+1)}(x_0)|$. Combining these two equations yields the desired bound on the right-hand side of equation (3.3.16). □

4. Resplendency

The main goal of this section is to show that Hensel minimality behaves well with respect to expansions of the structure by predicates living only on Cartesian powers of RV , one of the main results (Theorem 4.1.19) being that if the \mathcal{L} -theory of a valued field K is 0-, 1- or ω -h-minimal, then so is its \mathcal{L}' -theory, where \mathcal{L}' is the language of such an RV -expansion of K .

The key ingredient to Theorem 4.1.19 is Proposition 4.1.7, which in some sense is even stronger: Any set $X \subset K$ definable in the bigger language \mathcal{L}' can already be prepared by a finite set definable in the smaller language \mathcal{L} . Using this, it often becomes possible, given a completely arbitrary set $Z \subset \text{RV}^k$, to ‘without loss assume that Z is definable’. This turns out to be pretty useful to get rid of some technicalities related to cell decomposition in valued fields; the preparations for this are done in Subsection 4.3.

Under the assumption of ω -h-minimality, preparation can also be generalized to more general leading term structures: One can define $\text{RV}_I := (K^\times/(1+I)) \cup \{0\}$ for arbitrary proper definable ideals $I \subset \mathcal{O}_K$,

and for most such I , any subset of K that is definable using parameters from RV_I can be ‘ I -prepared’ (see Theorem 4.2.3). Using this, we deduce that if the theory of a valued field is ω -h-minimal, then so is the theory of the field with any coarsened (not necessarily definable) valuation (Corollary 4.2.4). In fact, in [17, Theorem 2.2.7], it is shown that also 1-h-minimality is preserved under coarsenings of the valuation, and furthermore, some resplendency results in the mixed characteristic case are developed in [17].

Note that the proofs in this section need somewhat deeper methods from model theory than the remainder of the paper. In particular, the emphasis shifts from a geometric description of definable sets to questions revolving around the extension of automorphisms.

4.1. Resplendency for a fixed ideal

As convened earlier, K is a valued field of equi-characteristic zero, considered a structure in a language \mathcal{L} containing \mathcal{L}_{val} . For this entire subsection, we fix a proper definable (with parameters) ideal $I \subset \mathcal{O}_K$. We start by defining the I -version of preparation.

Definition 4.1.1 (I -preparing sets).

1. We define RV_I to be the disjoint union of the quotient $K^\times/(1+I)$ with $\{0\}$, and we write rv_I for the map $K \rightarrow \text{RV}_I$ that extends the projection map $K^\times \rightarrow K^\times/(1+I)$ by sending 0 to 0.
2. We say that a ball $B \subset K$ is I -next to c for some $c \in K$ if $B = c + \text{rv}_I^{-1}(\xi)$ for some (nonzero) element ξ of RV_I . We say that B is I -next to C for some finite (non-empty) set $C \subset K$ if B equals $\bigcap_{c \in C} B_c$ with B_c a ball I -next to c for each $c \in C$.
3. Let C be a finite non-empty subset of K . We say that a set $X \subset K$ is I -prepared by C if belonging to X depends only on the tuple $(\text{rv}_I(x - c))_{c \in C}$, or, equivalently, if every ball I -next to C is either contained in X or disjoint from X .

Remark 4.1.2. If $I = B_{<\lambda}(0)$ for some $\lambda \leq 1$ in Γ_K , then of course we have $\text{RV}_I = \text{RV}_\lambda$, I -next means λ -next, and I -prepared means λ -prepared.

By considering I a member of a \emptyset -definable family of proper ideals of \mathcal{O}_K , we see that RV_I is a definable subset of an imaginary sort. In particular, it makes sense to work using parameters from RV_I .

Definition 4.1.3 (Having I -preparation). We say that K has I -preparation if for every set $A \subset K$ and every $(A \cup \text{RV}_I)$ -definable subset $X \subset K$, there exists a finite A -definable set $C \subset K$ such that X is I -prepared by C .

Remark 4.1.4. Note how this is related to Hensel minimality: $\text{Th}(K)$ is 0-h-minimal iff every $K' \equiv K$ has $\mathcal{M}_{K'}$ -preparation, and $\text{Th}(K)$ is ω -h-minimal iff every $K' \equiv K$ has $B_{<\lambda}(0)$ -preparation for every $\lambda \leq 1$ in $\Gamma_{K'}$.

By ‘resplendent I -preparation’, we mean that one can I -prepare sets that are definable in arbitrary expansions of \mathcal{L} by predicates on RV_I . More precisely:

Definition 4.1.5 (Resplendent I -preparation).

1. An RV_I -expansion of K is an expansion obtained by adding (any) predicates that live on Cartesian powers of the (imaginary) definable set RV_I .
2. We say that K has resplendent I -preparation if for every set $A \subset K$, for every RV_I -expansion of K with language $\mathcal{L}' \supset \mathcal{L}$, and for every $\mathcal{L}'(A)$ -definable subset $X \subset K$, there exists a finite $\mathcal{L}(A)$ -definable set $C \subset K$ such that X is I -prepared by C .

Note that we intentionally require C to be definable in the smaller language \mathcal{L} .

For the remainder of this subsection, we will assume that I is definable without parameters. In this case, it also makes sense to introduce the notions of I -preparation for the theory of K :

Definition 4.1.6 (I -preparation for theories). Suppose that I is \emptyset -definable (as convened for the remainder of this subsection).

1. Given $K' \equiv K$, we write $I_{K'}$ for the ideal of $\mathcal{O}_{K'}$ defined by the formula that defines I in K .
2. We say that $\text{Th}(K)$ has I -preparation if every model $K' \equiv K$ has $I_{K'}$ -preparation.
3. We say that $\text{Th}(K)$ has *resplendent I -preparation* if every model $K' \equiv K$ has resplendent $I_{K'}$ -preparation.

(In particular, for $\text{Th}(K)$, having \mathcal{M}_K -preparation is the same as 0-h-minimality.)

Since adding parameters from RV_I is a specific kind of RV_I -expansion, resplendent I -preparation clearly implies I -preparation. The central result of this subsection is the converse given in the following proposition:

Proposition 4.1.7 (Resplendency). *Suppose that I is \emptyset -definable. Then the following are equivalent:*

- (i) $\text{Th}(K)$ has I -preparation.
- (ii) $\text{Th}(K)$ has *resplendent I -preparation*.

Note that the proposition in particular implies that 0-h-minimal theories have resplendent \mathcal{M}_K -preparation; in Theorem 4.2.3, we will see a strong version of this for ω -h-minimality.

The proof of (i) \Rightarrow (ii) requires a number of lemmas that we will prove now. It will be sufficient to consider models K that are sufficiently saturated, so from now on, we fix such a sufficiently saturated K : We assume K to be κ -saturated for some $\kappa > |\mathcal{L}|$. As usual, we call a set *small* if its cardinality is less than κ .

Convention 4.1.8. For the remainder of Subsection 4.1, we consider \mathcal{L} a genuine two-sorted language, with sorts K and RV_I .

Let us first rephrase preparation in terms of definability in a certain sublanguage, namely:

Definition 4.1.9 ($\mathcal{L}_{\mathbb{Q},I}$). Let $\mathcal{L}_{\mathbb{Q},I}$ be the language $\{0, +, -\} \cup \{s \cdot \mid s \in \mathbb{Q}\}$ of \mathbb{Q} -vector spaces on the valued field K (where ‘ $s \cdot$ ’ denotes multiplication by s), together with the sort RV_I and the map rv_I .

Lemma 4.1.10 (Preparation in terms of $\mathcal{L}_{\mathbb{Q},I}$). *For every (not necessarily definable) $X \subset K$ and every \mathbb{Q} -vector space $A \subset K$, the following are equivalent:*

1. *There exists a finite set $C \subset A$ such that X is I -prepared by C .*
2. *There exists an RV_I -expansion of K with language $\mathcal{L}'_{\mathbb{Q},I} \supset \mathcal{L}_{\mathbb{Q},I}$ such that X is defined by a quantifier free $\mathcal{L}'_{\mathbb{Q},I}(A)$ -formula.*
3. *There exists an RV_I -expansion of K with language $\mathcal{L}'_{\mathbb{Q},I} \supset \mathcal{L}_{\mathbb{Q},I}$ such that X is defined by a field quantifier free $\mathcal{L}'_{\mathbb{Q},I}(A)$ -formula.*

Proof. (2) \Rightarrow (3) is clear.

(1) \Rightarrow (2): Let us assume that (1) holds, and let $C = \{c_1, \dots, c_k\}$, $f(x) := (\text{rv}_I(x - c_1), \dots, \text{rv}_I(x - c_k))$ and $Y := f(X) \subset \text{RV}_I^k$. Then, because X is I -prepared by C , $X = f^{-1}(Y)$, and this is quantifier free definable in $\mathcal{L}_{\mathbb{Q},I}(A) \cup \{Y\}$.

(3) \Rightarrow (1): Every field quantifier free $\mathcal{L}'_{\mathbb{Q},I}(A)$ -formula in a single variable x is equivalent to one of the form $\phi(\text{rv}_I(m_1x + c_1), \dots, \text{rv}_I(m_\ell x + c_\ell))$, where ϕ is a formula living entirely in the sort RV_I , $m_i \neq 0$ are rational numbers and c_i are elements of A . Since $\text{rv}_I(x + c_i/m_i)$ determines $m_i \text{rv}_I(x + c_i/m_i) = \text{rv}_I(m_i x + c_i)$, X is I -prepared by $C = \{-c_1/m_1, \dots, -c_\ell/m_\ell\} \subset A$. \square

We now recall some general model theoretic notions.

Notation 4.1.11. In the remainder of this subsection, we use the following conventions common in model theory:

- Given a set A and a tuple of variables x , we write A^x for the Cartesian power of A corresponding to the length of the tuple of variables x .
- Given sets A and B , we sometimes write AB for their union, and we freely interpret tuples as sets.

Definition 4.1.12 (Partial (elementary) isomorphisms). Suppose that $\tilde{\mathcal{L}}$ is an arbitrary language, M and N are $\tilde{\mathcal{L}}$ -structures, $A \subset M$, $B \subset N$ and $f: A \rightarrow B$ a bijection.

- We say that f is a *partial $\tilde{\mathcal{L}}$ -isomorphism* if for every quantifier free $\tilde{\mathcal{L}}$ -formula $\phi(x)$ and $a \in A^x$, $M \models \phi(a)$ if and only if $N \models \phi(f(a))$.
- We say that f is a *partial elementary $\tilde{\mathcal{L}}$ -isomorphism* if for every $\tilde{\mathcal{L}}$ -formula $\phi(x)$ and $a \in A^x$, $M \models \phi(a)$ if and only if $N \models \phi(f(a))$.

Note that any partial $\tilde{\mathcal{L}}$ -isomorphism has a unique extension to the $\tilde{\mathcal{L}}$ -structure generated by its domain. We implicitly identify f and this extension.

In the following, for a subset $A \subset K$, $\langle A \rangle_{\mathbb{Q}}$ denotes the \mathbb{Q} -sub-vector space of K generated by A .

Remark 4.1.13. For any set $A \subset K$, the $\mathcal{L}_{\mathbb{Q},I}$ -substructure of K generated by A consists of $\langle A \rangle_{\mathbb{Q}}$ together with its image $\text{rv}_I(\langle A \rangle_{\mathbb{Q}})$ in RV_I . In particular, since $\mathcal{L}_{\mathbb{Q},I}$ has no language on RV_I (except for the maps rv_I), we have: To obtain that a sort-preserving map $f: A_1 \rightarrow A_2$ (for some $A_1, A_2 \subset K \cup \text{RV}_I$) is a partial $\mathcal{L}_{\mathbb{Q},I}$ -automorphism, it suffices to verify that the restriction $f|_{A_1 \cap K}$ is a partial $\mathcal{L}_{\mathbb{Q},I}$ -automorphism and that on $\tilde{A}_1 := A_1 \cap \text{rv}_I(A_1 \cap K)$, the map induced by $f|_{A_1 \cap K}$ agrees with $f|_{\tilde{A}_1}$.

Lemma 4.1.14 (Preparation in terms of partial isomorphisms). *Let $A \subset K$ be a small \mathbb{Q} -sub-vector space. The following are equivalent:*

- (i) Any $\mathcal{L}(A \cup \text{RV}_I)$ -definable set $X \subset K$ can be I -prepared by some finite set $C \subset A$.
- (ii) For all $A_2 \subset K$, $c_1, c_2 \in K$ and all (potentially large) $B_1, B_2 \subset \text{RV}_I$ with $\text{rv}_I(\langle A, c_1 \rangle_{\mathbb{Q}}) \subset B_1$, if $f: AB_1c_1 \rightarrow A_2B_2c_2$ is a partial $\mathcal{L}_{\mathbb{Q},I}$ -isomorphism sending c_1 to c_2 whose restriction $f|_{AB_1}$ is a partial elementary \mathcal{L} -isomorphism, then the entire f is a partial elementary \mathcal{L} -isomorphism.
- (iii) For all $c_1, c_2 \in K$ and all (potentially large) $B \subset \text{RV}_I$ containing $\text{rv}_I(\langle A, c_1 \rangle_{\mathbb{Q}})$, any partial $\mathcal{L}_{\mathbb{Q},I}(A \cup B)$ -isomorphism $f: \{c_1\} \rightarrow \{c_2\}$, is a partial elementary $\mathcal{L}(A \cup B)$ -isomorphism.

Remark 4.1.15. If (ii) holds for $B_1 = B_2 = \text{RV}_I$, then it also holds in general, since by Remark 4.1.13, the partial $\mathcal{L}_{\mathbb{Q},I}$ -isomorphism $f: AB_1c_1 \rightarrow A_2B_2c_2$ extends to $f: Ac_1 \cup \text{RV}_I \rightarrow A_2c_2 \cup \text{RV}_I$. Analogously, we may assume $B = \text{RV}_I$ in (iii).

Proof of Lemma 4.1.14. (i) \Rightarrow (iii): Let f be as in (iii). We have to check that for every $\mathcal{L}(A \cup B)$ -definable set $X \subset K$, $c_1 \in X$ if and only if $c_2 \in X$. By (i), there exists a finite $C \subset A$ such that X is I -prepared by C . Since f is an $\mathcal{L}_{\mathbb{Q},I}(A \cup B)$ -isomorphism and B contains $\text{rv}_I(\langle A, c_1 \rangle_{\mathbb{Q}})$, for all $a \in C$ and all $r \geq 1$, we have

$$\text{rv}_I(c_2 - a) = \text{rv}_I(f(c_1) - f(a)) = f(\text{rv}_I(c_1 - a)) = \text{rv}_I(c_1 - a).$$

Since X is I -prepared by C , it follows that $c_1 \in X$ if and only if $c_2 \in X$.

(iii) \Rightarrow (ii): Let f be as in (ii). Since f is \mathcal{L} -elementary if and only its restriction to every finite domain is, we may assume B_i small. Using the assumption that $f|_{AB_1}$ is \mathcal{L} -elementary, we can extend $(f|_{AB_1})^{-1}$ \mathcal{L} -elementarily to some g defined at c_2 . Let $c'_1 := g(c_2)$. Then $g \circ f: \{c_1\} \rightarrow \{c'_1\}$ is a partial $\mathcal{L}_{\mathbb{Q},I}(A \cup B_1)$ -isomorphism. Since $\text{rv}_I(\langle A, c_1 \rangle_{\mathbb{Q}}) \subset B_1$, it follows by (iii) that $g \circ f$ is an elementary \mathcal{L} -isomorphism. As g is also \mathcal{L} -elementary, so is f .

(ii) \Rightarrow (i): Let X be as in (i), and let $B \subset \text{RV}_I$ be a finite subset such that X is $\mathcal{L}(A \cup B)$ -definable.

Consider any $c_1, c_2 \in K$ that have the same $\text{qf-}\mathcal{L}'_{\mathbb{Q},I}$ -type over $A \cup B$, where $\mathcal{L}'_{\mathbb{Q},I} \supset \mathcal{L}_{\mathbb{Q},I}$ is the language of the expansion of K by the full $\mathcal{L}(A)$ -induced structure on RV_I . Then the map $f: c_1 \rightarrow c_2$ is an $\mathcal{L}'_{\mathbb{Q},I}(A \cup B)$ -isomorphism and extends to $f: AB_1c_1 \rightarrow A_2B_2c_2$, where $B_i := B \cup \text{rv}_I(\langle A, c_i \rangle_{\mathbb{Q}})$. By definition of $\mathcal{L}'_{\mathbb{Q},I}$, the restriction $f|_{AB_1}$ is \mathcal{L} -elementary, so by (ii), also the entire f is \mathcal{L} -elementary. Since moreover f is the identity on $A \cup B$, this implies that c_1 and c_2 have the same \mathcal{L} -type over $A \cup B$.

We just proved that the $\mathcal{L}(A \cup B)$ -type of any element $c \in K$ is implied by its $\text{qf-}\mathcal{L}'_{\mathbb{Q},I}(A \cup B)$ -type. By a classical compactness argument (cf. the proof of [67, Theorem 3.2.5]), it follows that any $\mathcal{L}(A \cup B)$ -formula in one valued field variable is equivalent to a quantifier-free $\mathcal{L}'_{\mathbb{Q},I}(A \cup B)$ -formula. In particular, this applies to our set X . Since $\mathcal{L}'_{\mathbb{Q},I}(B)$ is the language of an RV_I -expansion of K , our claim follows from Lemma 4.1.10. □

Lemma 4.1.16 (Back and forth over RV_I). *Suppose that K has I -preparation. Then the set of partial elementary \mathcal{L} -isomorphisms $f : \text{RV}_I \cup A_1 \rightarrow \text{RV}_I \cup A_2$ (where A_1, A_2 run over all small subsets of K) has the back and forth.*

Recall that ‘having the back and forth’ means: for any such f and any $c_1 \in K \setminus A_1$, f can be extended to c_1 while staying in that set of maps, and similarly for f^{-1} .

Proof. Let f be as above. Since partial elementary isomorphisms can always be extended to the algebraic closure of their domain, we may assume that $\text{acl}_K(A_i) = A_i$. By I -preparation, statement (i) of Lemma 4.1.14 now holds for $A = A_i$. Pick any $c_1 \in K$, set $B_1 := \text{rv}_I(\langle A_1, c_1 \rangle_{\mathbb{Q}})$, and let $c_2 \models f_* \text{qftp}_{\mathcal{L}_{\mathbb{Q},I}}(c_1/A_1 B_1)$. (Such a c_2 exists, since K is $|A_1 B_1|^+$ -saturated.) Let g extend f by sending c_1 to c_2 . By construction, $g|_{A_1 B_1 c_1}$ is a partial $\mathcal{L}_{\mathbb{Q},I}$ -isomorphism. This implies that the entire g is a partial $\mathcal{L}_{\mathbb{Q},I}$ -isomorphism (cf. Remark 4.1.13). Note also that $g|_{A_1 \text{RV}_I} = f$ is a partial elementary \mathcal{L} -isomorphism. Hence, by Lemma 4.1.14 (ii), the entire g is \mathcal{L} -elementary. \square

We are now ready for the proof of the central result of this subsection:

Proof of Proposition 4.1.7. Set $\mathcal{T} := \text{Th}(K)$. Recall that the implication (ii) to (i) is trivial, so let us assume (i) (namely, \mathcal{T} has I -preparation) and prove (ii) (namely, \mathcal{T} has resplendent I -preparation): Suppose that K is an arbitrary model of \mathcal{T} , let $\mathcal{L}' \supset \mathcal{L}$ be the language of an RV_I -expansion of K , let A be a subset of K , and let $X \subset K$ be an $\mathcal{L}'(A)$ -definable set; we need to find a finite $\mathcal{L}(A)$ -definable set $C \subset K$ that I -prepares X . Since the condition that C I -prepares X is first order, we may replace K with a sufficiently saturated elementary extension. For the remainder of the proof, we fix such a K .

We may moreover assume that $A = \text{acl}_{\mathcal{L},K}(A)$. Indeed, if we find a finite, $\mathcal{L}(\text{acl}_{\mathcal{L},K}(A))$ -definable set C that I -prepares X , then the union of the (finitely many) conjugates of C over A is an $\mathcal{L}(A)$ -definable finite set I -preparing X . Therefore, we may apply Lemma 4.1.14 (and Remark 4.1.15), meaning that it suffices to show that every partial $\mathcal{L}_{\mathbb{Q},I}(A \cup \text{RV}_I)$ -isomorphism $f : \{c_1\} \rightarrow \{c_2\}$ is an elementary $\mathcal{L}'(A \cup \text{RV}_I)$ -isomorphism.

By (i) and Lemma 4.1.14 applied in \mathcal{L} , such an f is a partial elementary $\mathcal{L}(A \cup \text{RV}_I)$ -isomorphism, and since f is the identity on RV_I , it is a partial $\mathcal{L}'(A \cup \text{RV}_I)$ -isomorphism. It remains to show that f preserves all \mathcal{L}' -formulas and not just the quantifier-free ones.

Let \mathcal{F} be the class of partial elementary $\mathcal{L}(\text{RV}_I)$ -isomorphisms with small domains $A \subset K$. By Lemma 4.1.16, \mathcal{F} has the back and forth. Moreover, any $g \in \mathcal{F}$ is also a partial \mathcal{L}' -isomorphism, as it fixes RV_I . Recall that if a class of partial $\tilde{\mathcal{L}}$ -isomorphisms has the back and forth (for any given language $\tilde{\mathcal{L}}$), by an easy induction on the structure of formulas, those partial isomorphisms are automatically $\tilde{\mathcal{L}}$ -elementary. Thus any $g \in \mathcal{F}$, and in particular f , is a partial elementary \mathcal{L}' -isomorphism, which is what we had to show. \square

We now mention some easy consequences of Proposition 4.1.7 and its proof.

Corollary 4.1.17 (RV-expansions preserve acl). *Suppose that $\text{Th}(K)$ has I -preparation for some 0-definable proper ideal $I \subset \mathcal{O}_K$. Then for any RV_I -expansion of K with language $\mathcal{L}' \supset \mathcal{L}$ and any $A \subset K$, we have $\text{acl}_{\mathcal{L}',K}(A) = \text{acl}_{\mathcal{L},K}(A)$. In particular:*

1. *If $\text{Th}(K)$ is 0- h -minimal, then for any RV -expansion of K with language $\mathcal{L}' \supset \mathcal{L}$ and any $A \subset K$, we have $\text{acl}_{\mathcal{L}',K}(A) = \text{acl}_{\mathcal{L},K}(A)$.*
2. *If $\text{Th}(K)$ is ω - h -minimal, then for any $\lambda \in \Gamma_K^\times$, any RV_λ -expansion of K with language $\mathcal{L}' \supset \mathcal{L}$ and any $A \subset K$, we have $\text{acl}_{\mathcal{L}',K}(A) = \text{acl}_{\mathcal{L},K}(A)$.*

Proof. Any $b \in \text{acl}_{\mathcal{L}',K}(A)$ is an element of a finite $\mathcal{L}'(A)$ -definable set $X \subset K$. By Proposition 4.1.7, X can be I -prepared by a finite $\mathcal{L}(A)$ -definable set C . This implies $X \subset C$ and hence $b \in \text{acl}_{\mathcal{L},K}(A)$. \square

Remark 4.1.18. We already saw some stable embeddedness results in Proposition 2.6.12 and Remark 2.6.13. For similar reasons, I -preparation implies stable embeddedness of RV_I . Alternatively, one may deduce the stable embeddedness from Lemma 4.1.16 and [10, Appendix, Lemma 1].²

²This uses the existence of fully saturated models; getting rid of those is left as an exercise to the reader.

We conclude this subsection with the result that various notions of Hensel minimality automatically pass to RV-expansions of the structure.

Theorem 4.1.19 (RV-expansions preserve Hensel minimality). *Let K be a valued field of equi-characteristic 0 in a language \mathcal{L} containing the language \mathcal{L}_{val} of valued fields. Fix $\ell \in \{0, 1, \omega\}$, and suppose that $\text{Th}_{\mathcal{L}}(K)$ is ℓ -h-minimal. Moreover, fix an RV-expansion of K with language $\mathcal{L}' \supset \mathcal{L}$ (i.e., an expansion by predicates on Cartesian powers of RV). Then $\text{Th}_{\mathcal{L}'}(K)$ is also ℓ -h-minimal.*

Whereas the cases 0 and ω follow easily from Proposition 4.1.7, the only proof we have for $\ell = 1$ makes a detour through the criterion given in Theorem 2.9.1 (which is generalized for $\ell \geq 2$ in [68]). In [68], Theorem 4.1.19 is generalized to hold for all ℓ .

Proof of Theorem 4.1.19. The case $\ell = 0$ follows directly from Proposition 4.1.7: Since $\text{Th}_{\mathcal{L}}(K)$ has \mathcal{M}_K -preparation, it even has resplendent \mathcal{M}_K -preparation, so every $\mathcal{L}'(A)$ -definable set $X \subset K$ (for $A \subset K$) can be prepared by a finite $\mathcal{L}(A)$ -definable (and hence in particular $\mathcal{L}'(A)$ -definable) subset of K .

For the case $\ell = \omega$, we use a similar argument, where \mathcal{M}_K is replaced by $I = B_{<\lambda}(0)$ for arbitrary $\lambda \leq 1$ in Γ_K^\times . Note that indeed, the given RV-expansion of K can also be considered an RV_λ -expansion, by pulling back all the new predicates along the canonical map $\text{RV}_\lambda \rightarrow \text{RV}$.

Finally, in the case $\ell = 1$, we use the 1-h-minimality criterion given in Theorem 2.9.1: Let $f : K \rightarrow K$ be $\mathcal{L}'(A)$ -definable for $A \subset K \cup \text{RV}$. We need to prove the two conditions (T1) and (T2) stated in the theorem.

First of all, note that we already know that $\text{Th}_{\mathcal{L}'}(K)$ is 0-h-minimal.

By Corollary 4.1.17 (applied with $I = \mathcal{M}_K$), for every $x \in K$, we have $f(x) \in \text{acl}_{\mathcal{L},K}(A \cup \{x\})$. Let C_x be an $\mathcal{L}(A)$ -definable family of finite sets such that $f(x) \in C_x$ for all $x \in K$. Lemma 2.5.3 yields an $\mathcal{L}(A)$ -definable family of injective maps $g_x : C_x \rightarrow \text{RV}^k$. Using this, we define maps $h : K \rightarrow \text{RV}^k, x \mapsto g_x(f(x))$ and $\tilde{f} : K \times \text{RV}^k \rightarrow K$ with $\tilde{f}(x, \xi) = g_x^{-1}(\xi)$ if $\xi \in g_x(C_x)$ and $\tilde{f}(x, \xi) = 0$ otherwise. Note that we obtain $f(x) = \tilde{f}(x, h(x))$ for all $x \in K$. Also note that h is $\mathcal{L}'(A)$ -definable and \tilde{f} is $\mathcal{L}(A)$ -definable.

Applying Corollary 2.6.6 (in \mathcal{L}') to the graph of h yields a finite $\mathcal{L}'(A)$ -definable set C' such that h is constant on every ball B 1-next to C' . Moreover, using Lemma 2.8.5 in \mathcal{L} , we find an $\mathcal{L}(A)$ -definable family of sets C_ξ preparing $x \mapsto \tilde{f}(x, \xi)$ in the sense of that lemma for every $\xi \in \text{RV}^k$. Now let C be the union of C' and all those C_ξ . That union is still finite (using Corollary 2.6.7 in \mathcal{L}), and it prepares f in the way Condition (T1) requires it. Indeed, on each fixed ball B 1-next to C , we have $f(x) = \tilde{f}(x, \xi)$ for one fixed ξ , our application of Lemma 2.8.5 ensured that $x \mapsto \tilde{f}(x, \xi)$ is prepared in the desired sense.

Condition (T2) now is also clear: For each ξ , $\tilde{f}(\cdot, \xi)$ has only finitely many infinite fibers (by Lemma 2.8.1 in \mathcal{L}). Taking the union of all those sets (for all ξ) still yields a finite set. □

4.2. Changing the ideal

We now switch back to considering \mathcal{L} a single-sorted language. In this subsection, we prove that ω -h-minimality implies I -preparation for most ideals I , where ‘most’ means ‘ Γ_K -open’ in the following sense:

Definition 4.2.1 (Notions of openness). By an *open ball ideal* in \mathcal{O}_K , we mean an ideal of the form $B_{<\lambda}(0)$ for some $\lambda \leq 1$ in Γ_K . By a Γ_K -*open ideal* in \mathcal{O}_K , we mean an ideal that is equal to the union of all open ball ideals it contains.

In other words, an ideal is Γ_K -open if its image in Γ_K is open with respect to the interval topology on Γ_K . Note that none of the two above openness notions coincides with being topologically open in the valued field topology.

Remark 4.2.2. If the value group is discrete, then every proper ideal is Γ_K -open; otherwise, an ideal is Γ_K -open if and only if it is not a closed ball. In [68], the condition of Γ_K -openness on I in Theorem 4.2.3 is removed.

Theorem 4.2.3 (*I*-preparation). *Let K be a valued field of equi-characteristic 0, considered a structure in a language $\mathcal{L} \supset \mathcal{L}_{\text{val}}$. Suppose that $\text{Th}(K)$ is ω -h-minimal. Then K has resplendent *I*-preparation (see Definition 4.1.5) for every proper Γ_K -open definable (with parameters) ideal I of \mathcal{O}_K .*

Proof of Theorem 4.2.3. Let K and I be as in the theorem, let $A \subset K$, fix an RV_I -expansion of K in some language $\mathcal{L}' \supset \mathcal{L}$, and let $X \subset K$ be $\mathcal{L}'(A)$ -definable. We need to find a finite $\mathcal{L}(A)$ -definable set $C \subset K$ such that X is *I*-prepared by C . By passing to an elementary extension, we may assume that K is sufficiently saturated as an \mathcal{L}' -structure.

Fix $\lambda \in \Gamma_K^\times$ such that $B_{<\lambda}(0)$ is contained in I . Then up to interdefinability, the \mathcal{L}' -structure on K is an RV_λ -expansion of the \mathcal{L} -structure on K , namely by the preimage of each of the \mathcal{L}' -definable sets in RV_I^k under the map $\text{RV}_\lambda^k \rightarrow \text{RV}_I^k$.

By Lemma 2.4.1, adding a λ as a constant to the language preserves ω -h-minimality: that is, $\mathcal{T} := \text{Th}_{\mathcal{L}(\lambda)}(K)$ is still ω -h-minimal. In that language, $B_{<\lambda}(0)$ is \emptyset -definable, so Proposition 4.1.7 implies that \mathcal{T} has resplendent $B_{<\lambda}(0)$ -preparation. In particular, there exists a finite $\mathcal{L}(A, \lambda)$ -definable set C such that no ball λ -next to C ‘intersects X properly’: that is, every such ball is either contained in X or disjoint from X . Using Corollary 2.6.7, we may even assume that C is $\mathcal{L}(A)$ -definable.

The above set C might depend on the choice of λ , so let us denote it by C_λ instead. (Contrary to what the notation might suggest, this dependence is not definable, in general.) Using a similar compactness argument as in Proposition 2.6.2, we now make C_λ independent of λ , as follows: The condition that X intersects no ball λ -next to C_λ properly can be expressed by a first-order formula in λ , and there are only a bounded number of choices for C_λ . It follows that if no finite $\mathcal{L}(A)$ -definable C works for all λ , then by our saturation assumption on K , we find a λ such that for every C , some ball λ -next to C intersects X properly, which is a contradiction. Let now C be a set that works for every λ .

To finish the proof, it remains to prove that every ball *I*-next to C is either contained in X or disjoint from X . Let $a, a' \in K$ be such that $\text{rv}_I(a - c) = \text{rv}_I(a' - c)$ for all $c \in C$. Since I is the union of the open ball ideals $B_{<\lambda}(0)$ it contains, $\frac{a-c}{a'-c} \in 1 + I$ implies $\frac{a-c}{a'-c} \in B_{<\lambda}(1)$ for one such λ . We may moreover choose a single λ such that this holds for all $c \in C$. Then $\bigcap_{c \in C} \text{rv}_\lambda^{-1}(\text{rv}_\lambda(a - c))$ is a ball λ -next to C containing both a and a' . It follows that $a \in X$ if and only if $a' \in X$. □

From this, we can now deduce that ω -h-minimality is preserved under coarsening of the valuation:

Corollary 4.2.4 (Coarsening the valuation). *Suppose that $\text{Th}_{\mathcal{L}}(K)$ is ω -h-minimal and that $|\cdot|_c : K \rightarrow \Gamma_{K,c}$ is a non-trivial coarsening of the valuation $|\cdot|$ on K with valuation ring $\mathcal{O}_{K,c}$ (non-trivial meaning $\mathcal{O}_{K,c} \neq K$). Let K_c be the expansion of K by a predicate for $\mathcal{O}_{K,c}$, considered a valued field with the valuation $|\cdot|_c$, and denote by $\mathcal{L}_c \supset \mathcal{L}$ the corresponding language. Then $\text{Th}_{\mathcal{L}_c}(K_c)$ is also ω -h-minimal.*

Proof. Since $\mathcal{O}_{K,c}$ is an rv -pullback (i.e., a preimage in K of a subset of RV), Theorem 4.1.19 implies that $\text{Th}_{\mathcal{L}_c}(K)$ is ω -h-minimal for the valuation $|\cdot|$.

Let K' be \mathcal{L}_c -elementary equivalent to K ; we need to verify that K' has $B_{<\lambda}(0)$ -preparation for every $\lambda \leq 1$ in $\Gamma_{K',c}^\times$. This follows from Theorem 4.2.3 since $B_{<\lambda}(0)$ is $\Gamma_{K'}$ -open, which in turn holds because $B_{<\lambda}(0)$ is the preimage under $K' \rightarrow \Gamma_{K'}$ of the set of those $\mu \in \Gamma_{K'}$ which get sent to an element less than λ by the map $\Gamma_{K'} \rightarrow \Gamma_{K',c}$. □

Remark 4.2.5. In [17], it is shown that also 1-h-minimality is preserved under coarsening of the valuation, and, in [68], this is further generalized to all $\ell \geq 1$.

4.3. Algebraic Skolem functions

As usual, K is a valued field of equi-characteristic 0, in a language $\mathcal{L} \supset \mathcal{L}_{\text{val}}$.

The mere statement of cell decomposition in valued fields is usually very technical, one reason being that one cannot definably pick the centers of the cells. This problem is easier to deal with if one expands the structure by certain Skolem functions. Usually, one would not want to modify the structure in such a

way. However, in this subsection, we provide tools that make it possible, in many situations, to assume the existence of such Skolem functions (without losing power or generality).

The first thing to note is that properly adding the required Skolem functions preserves Hensel minimality; this follows from Theorem 4.1.19 and the observation that adding ‘algebraic’ Skolem functions on RV is enough (Lemma 4.3.1). This by itself is useful whenever one only wants to prove that every definable set (or function) has some good properties. A similar approach has been recently followed in [14].

Sometimes, one wants to use cell decomposition to prove that every definable object yields some other kind of definable object. If one proves such a result in a structure expanded by Skolem functions, one needs to be able to get rid of the Skolem functions again afterward. There is probably no general recipe for this, but Lemma 4.3.4 is a useful tool; we apply it, for example, in the proof of Theorem 5.4.10.

Lemma 4.3.1. *Suppose that $\text{Th}(K)$ is 0-h-minimal and that for every set $A' \subset \text{RV}$, we have $\text{acl}_{\text{RV}}(A') = \text{dcl}_{\text{RV}}(A')$. Then for every set $A \subset K$, we have $\text{acl}_K(A) = \text{dcl}_K(A)$.*

Proof. Let $C \subset K$ be a finite A -definable set; we need to prove that C is contained in $\text{dcl}_K(A)$. By Lemma 2.5.3, there exists an A -definable bijection $f: C \rightarrow C' \subset \text{RV}^k$. By Proposition 2.6.12, C' is definable with parameters from $\text{dcl}_{\text{RV}}(A)$, so $C' \subset \text{acl}_{\text{RV}}(\text{dcl}_{\text{RV}}(A))$. By assumption, this implies $C' \subset \text{dcl}_{\text{RV}}(A)$, which, using f , implies $C \subset \text{dcl}_K(A)$. \square

Remark 4.3.2. The above property ‘ $\text{acl}_K(A) = \text{dcl}_K(A)$ ’ holds in all models if and only if algebraic Skolem functions exist: that is, for all integers $n \geq 0$ and $m > 0$ and every \emptyset -definable set $X \subset K^{n+1}$ with the property that the coordinate projection map $p: X \rightarrow K^n$ has fibers of cardinality precisely m , there is a \emptyset -definable function $f: K^n \rightarrow K$ whose graph is a subset of X .

Proposition 4.3.3 (Obtaining $\text{acl} = \text{dcl}$). *Given a valued field K (in a language \mathcal{L}) whose theory is ℓ -h-minimal, for $\ell \in \{0, 1, \omega\}$, there exists an RV-expansion of K in a language $\mathcal{L}^{\text{as}} \supset \mathcal{L}$ such that $\text{Th}_{\mathcal{L}^{\text{as}}}(K)$ is ℓ -h-minimal and such that for every model $K' \models \text{Th}_{\mathcal{L}^{\text{as}}}(K)$ and for every subset $A \subset K'$, we have $\text{acl}_{\mathcal{L}^{\text{as}}, K'}(A) = \text{dcl}_{\mathcal{L}^{\text{as}}, K'}(A)$.*

Proof. Let us be lazy and simply expand K by a predicate for each subset of RV_K^n (for every n); let \mathcal{L}^{as} be the corresponding language. Then $\text{Th}_{\mathcal{L}^{\text{as}}}(K)$ is still ℓ -h-minimal by Theorem 4.1.19. Since in this expansion, we have algebraic Skolem functions in RV_K^n , we in particular have $\text{acl}_{\text{RV}}(A') = \text{dcl}_{\text{RV}}(A')$ for every $A' \subset \text{RV}_{K'}$ and for every model $K' \equiv_{\mathcal{L}^{\text{as}}} K$. Now Lemma 4.3.1 implies $\text{acl}_K(A) = \text{dcl}_K(A)$ for every $A \subset K'$. \square

One may use Proposition 4.3.3 to ensure the condition that acl equals dcl ; the next lemma can be useful to get back to the original language.

Lemma 4.3.4 (Undoing algebraic skolemization). *Suppose that $\text{Th}_{\mathcal{L}}(K)$ is 0-h-minimal and that we have an RV-expansion of K in some language $\mathcal{L}' \supset \mathcal{L}$. Let $\chi': K^n \rightarrow \text{RV}^{k'}$ be an \mathcal{L}' -definable map (for some $k' \geq 0$). Then there exists an \mathcal{L} -definable map $\chi: K^n \rightarrow \text{RV}^k$ (for some $k \geq 0$) such that χ' factors over χ : that is, $\chi' = g \circ \chi$ for some function $g: \text{RV}^k \rightarrow \text{RV}^{k'}$ (which is automatically \mathcal{L}' -definable).*

Proof. We do an induction over n . The case $n = 0$ is trivial, so assume $n \geq 1$.

Fix $a \in K$, and set $\chi'_a(\tilde{a}) := \chi'(a, \tilde{a})$, for $\tilde{a} \in K^{n-1}$. By induction, we find an $\mathcal{L}(a)$ -definable map $\chi_a: K^{n-1} \rightarrow \text{RV}^k$ such that $\chi'_a = g_a \circ \chi_a$ for some $\mathcal{L}'(a)$ -definable g_a . By compactness, we may assume that χ_a and g_a are definable uniformly in a , so that in particular, we obtain an \mathcal{L}' -definable set $W = \{(a, \zeta', g_a(\zeta')) \mid a \in K, \zeta' \in \text{RV}^k\}$. Using Corollary 2.6.6, we find a finite \mathcal{L}' -definable set 1-preparing W , and by Corollary 4.1.17, that set is contained in a finite \mathcal{L} -definable set C .

Finally, let $f: K \rightarrow \text{RV}^\ell$ be an \mathcal{L} -definable map as provided by Lemma 2.5.4, namely such that each nonempty fiber of f is either a singleton in C or a ball 1-next to C . We claim that the map $\chi(a, \tilde{a}) := (f(a), \chi_a(\tilde{a}))$ is as desired. Indeed, it is clearly \mathcal{L} -definable, and since C 1-prepares W , the function g_a is determined by $f(a)$, so that $f(a)$ and $\chi_a(\tilde{a})$ together determine $\chi'(a, \tilde{a}) = \chi'_a(\tilde{a}) = g_a(\chi_a(\tilde{a}))$. \square

5. Geometry in the Hensel minimal setting

In this section, we deduce geometric results in K^n under the assumption of 1-h-minimality. Some of those are already known under stronger assumptions, like C^k properties of definable functions (Subsection 5.1), cell decomposition (Subsection 5.2) and dimension theory (Subsection 5.3). Since some proofs are very similar to those in many other papers, we will be somewhat succinct.

As highlights, we then prove a version of the Jacobian Property in many variables (Subsection 5.4) that allows us to get general t-stratifications for definable sets in 1-h-minimal structures (Subsection 5.5). We also provide a higher-dimension version of the Taylor-approximation result (Subsection 5.6), the aim being to lay the ground for the first axiomatic approach in the non-Archimedean context of analogues of results by Yomdin–Gromov and Pila–Wilkie on parameterizations and point counting.

The version of cell decomposition presented in Subsection 5.2 uses a simplified notion of cells by temporarily adding certain Skolem functions to the language, as detailed in Subsection 4.3. Since that approach might not work for all potential applications, we conclude Section 5 by linking with more classical viewpoints on cells (Subsection 5.7).

5.1. Continuity and differentiability

We start by proving that 1-h-minimality implies that definable functions are almost everywhere continuous and even C^k . Note that all this follows directly from Theorem 5.6.1; moreover, the notion of ‘almost everywhere’ makes more sense once dimension theory is developed (see Subsection 5.3). For those reasons, for now, the statements will be a bit weaker and the proofs a bit sketchy.

Theorem 5.1.1 (Continuity). *Assume that $\text{Th}(K)$ is 1-h-minimal. For every definable function $f: X \subset K^n \rightarrow K$, the set U of those $u \in X$ such that f is continuous on a neighborhood of u is dense in X .*

For the moment, we only give the proof in the case where X is open. The general case is proved in a joint induction with cell decomposition (more precisely with Addenda 2 and 4 of Theorem 5.2.4), so we postpone that proof until Subsection 5.2.

The proof given here uses a different approach than the one used for similar results in [13], building on the following two lemmas that might be of independent interest.

Lemma 5.1.2 (Continuity on small boxes). *Assume that $\text{Th}(K)$ is 1-h-minimal. If $B \subset K^n$ is a product of balls such that all elements of B have the same type over RV and $f: B \rightarrow K$ is a \emptyset -definable function, then f is uniformly continuous on B .*

Proof. Set $B =: \hat{B} \times B' \subset K^{n-1} \times K$. Using Lemma 2.8.5, we deduce that for every fixed $\hat{a} \in \hat{B}$, the map $a' \mapsto f(\hat{a}, a')$ is continuous on B' . Moreover, this continuity is uniform in both, \hat{a} and a' , since all elements $(\hat{a}, a') \in B$ have the same type over RV . (The type $\text{tp}(\hat{a}a'/\text{RV})$ knows which δ works for which ϵ , where $\epsilon, \delta > 0$ are from the definition of continuity at (\hat{a}, a') .) After applying the same argument to all other coordinates, we deduce that f as a whole is uniformly continuous in B : For every $\epsilon > 0$, there exists a $\delta > 0$ such that if $a_1, a_2 \in B$ with $|a_1 - a_2| < \delta$ differ only in one coordinate, then $|f(a_1) - f(a_2)| < \epsilon$. If a_1 and a_2 differ in several coordinates, apply this repeatedly. \square

Lemma 5.1.3 (Balls with constant type). *Assume that K is $|\mathcal{L}|^+$ -saturated and that acl_K equals dcl_K . If $X \subset K^n$ is a \emptyset -definable set with non-empty interior, then X contains a ball B such that all elements of B have the same type over RV (i.e., $\text{tp}(a/\text{RV}) = \text{tp}(a'/\text{RV})$ for all $a, a' \in B$).*

Proof of Lemma 5.1.3. We may assume that X is of the form $\hat{B} \times B'$ for some balls $\hat{B} \subset K^{n-1}, B' \subset K$, and by induction, we may assume that all elements of \hat{B} have the same type over RV .

Since any \emptyset -definable function $f: K^{n-1} \rightarrow K$ is continuous on \hat{B} by Lemma 5.1.2, given such an f , we can first shrink \hat{B} so that $B' \setminus f(\hat{B})$ still contains a ball and then shrink B' so that $\hat{B} \times B'$ is disjoint from the graph of f . After possibly further shrinking \hat{B} , we then obtain

$$\text{rv}(a'_1 - f(\hat{a}_1)) = \text{rv}(a'_2 - f(\hat{a}_2)) \tag{5.1.1}$$

for any $(\hat{a}_i, a'_i) \in \hat{B} \times B'$. By saturation, this shrinking of \hat{B} can be done for all \emptyset -definable f simultaneously. In particular, B' is now disjoint from $\text{acl}_K(\hat{a}_1)$ for every $\hat{a}_1 \in \hat{B}$, and one deduces that all elements of $\hat{B} \times B'$ have the same type over RV , namely as follows: Given $(\hat{a}_i, a'_i) \in \hat{B} \times B'$ for $i = 1, 2$, we find some $a''_1 \in B'$ such that $\text{tp}(\hat{a}_1 a''_1 / \text{RV}) = \text{tp}(\hat{a}_2 a'_2 / \text{RV})$. This implies $\text{rv}(a''_1 - f(\hat{a}_1)) = \text{rv}(a'_2 - f(\hat{a}_2))$ for every \emptyset -definable f and hence (using equation (5.1.1) and Lemma 2.4.4) $\text{tp}(\hat{a}_1 a''_1 / \text{RV}) = \text{tp}(\hat{a}_1 a'_1 / \text{RV})$. \square

Proof of Theorem 5.1.1 when X is open. We may suppose that f is \emptyset -definable, that K is $|\mathcal{L}|^+$ -saturated and that acl_K equals dcl_K (by Proposition 4.3.3). Suppose that the set U from the theorem is not dense in X . Choose a ball $B \subset X \setminus U$ as provided by Lemma 5.1.3. By Lemma 5.1.2, f is continuous on B , which is a contradiction to $B \not\subset U$. \square

We now come to the differentiability of definable functions. The definition of C^k is the usual one:

Definition 5.1.4 (C^k). Let $U \subset K^m$ be open, and let $f: U \rightarrow K^n$ be a map. We say that f is C^0 if it is continuous and that it is C^k for $k \geq 1$ if there exists a C^{k-1} -map $\text{Jac } f: U \rightarrow K^{n \times m}$ (the Jacobian of f) such that for every $u \in U$, we have $\lim_{x \rightarrow u} \frac{|f(x) - f(u) - (\text{Jac } f)(u)(x-u)|}{|x-u|} = 0$. In the case $n = 1$, we also write $\text{grad } f$ instead of $\text{Jac } f$ and call it the gradient of f .

Theorem 5.1.5 (C^k). Assume that $\text{Th}(K)$ is 1- h -minimal and fix $k \geq 0$. For every definable function $f: K^n \rightarrow K$, the set U of those $u \in K^n$ such that f is C^k on a neighborhood of u is dense in K^n .

Remark 5.1.6. Clearly, U is open and definable over the same parameters as f , so f is C^k on a definable dense open subset of K^n .

Remark 5.1.7. In Subsection 5.3, we will introduce a notion of dimension, and we will see that U being dense implies that $K^n \setminus U$ has dimension less than n . Thus, definable functions are almost everywhere C^k in this rather strong sense.

Proof of Theorem 5.1.5. ³ The case $k = 0$ is Theorem 5.1.1. The cases $k \geq 2$ follow from the case $k = 1$ by induction, so now assume $k = 1$. By using the same strategy as in the proof of Theorem 5.1.1, it suffices to prove the following: If $B \subset K^n$ is a product of balls such that all elements of B have the same type over RV and $f: B \rightarrow K$ is a \emptyset -definable function, then f is C^1 on B .

So let such a B be given. To prove that f is C^1 on B , we use three ingredients:

First, we use that all second partial derivatives of f have constant valuation on B , say, bounded by λ . Second, for each $i \leq n$, we apply Theorem 3.2.2 with $r = 1$ to the function f considered a function in the i th coordinate and with all other coordinates fixed. Using that all elements of B have the same type over RV , we deduce that for $x, x' \in B$ differing only in one coordinate, we have

$$|f(x) - f(x') - \nabla f(x') \cdot (x - x')| \leq \lambda |x - x'|^2. \tag{5.1.2}$$

Third, we apply Theorem 3.2.2 with $r = 0$ in a similar way to each of the first partial derivatives of f .

Now the same computation as in the proof of Theorem 5.6.1 implies that f is C^1 on B . The idea of this computation is the following: Given arbitrary x and x' in B , we change one coordinate at a time, apply equation (5.1.2) in each step and use the third ingredient to bound the difference between the value of ∇f at the starting point and its value at the intermediate points. In the end, we obtain the same bound as in equation (5.1.2) for arbitrary $x, x' \in B$. \square

5.2. Cell decomposition

In this subsection, we present a version of cell decomposition results that are simpler than usual by imposing the following condition.

³We would like to thank Yatir Halevi for pointing out that a previous argument for Theorem 5.1.5 was wrong.

Assumption 5.2.1. We assume in this subsection that we have algebraic Skolem functions in K ; or, equivalently, that for every $K' \equiv K$ and every $A \subset K'$, we have $\text{acl}_{K'}(A) = \text{dcl}_{K'}(A)$. We will abbreviate this by ‘acl equals dcl (in $\text{Th}(K)$)’.

The tools provided in Subsection 4.3 often make it possible to reduce to the case of languages where this assumption holds, as we will see in later subsections of Section 5. In fact, Assumption 5.2.1 does more than just simplify the arguments: it also allows one to formulate stronger results like piecewise Lipschitz continuity results as in Theorem 5.2.8. This is similar to [14], where this condition on acl is furthermore used to obtain parametrization results with finitely many maps (in the Pila-Wilkie sense [62]), but not under axiomatic assumptions.

Close variants of the results of this subsection appear in [14]; indeed, the tameness condition of [14] is very similar to the 1-h-minimality criteria from Theorem 2.9.1; see Remark 2.9.2.

In the following, for $m \leq n$, we denote the projection $K^n \rightarrow K^m$ to the first m coordinates by $\pi_{\leq m}$ or $\pi_{< m+1}$.

Definition 5.2.2 (Cells, twisted boxes). Fix any parameter set $A \subset K^{\text{eq}}$, and consider a non-empty A -definable set $X \subset K^n$ for some n , and, for $i = 1, \dots, n$, values j_i in $\{0, 1\}$ and A -definable functions $c_i : \pi_{< i}(X) \rightarrow K$. Then X is called an A -definable cell with center tuple $c = (c_i)_{i=1}^n$ and of cell-type $j = (j_i)_{i=1}^n$ if it is of the form

$$X = \{x \in K^n \mid (\text{rv}(x_i - c_i(x_{< i})))_{i=1}^n \in R\}$$

for a (necessarily A -definable) set

$$R \subset \prod_{i=1}^n (j_i \cdot \text{RV}^\times),$$

where $x_{< i} = \pi_{< i}(x)$ and where $0 \cdot \text{RV}^\times = \{0\} \subset \text{RV}$, and $1 \cdot \text{RV}^\times = \text{RV}^\times \subset \text{RV}$.

If X is such a cell, then, for any $r \in R$, the subset

$$\{x \in K^n \mid \text{rv}(x_i - c_i(x_{< i})))_{i=1}^n = r\}$$

of X is called a *twisted box* of the cell X . We also call X itself a twisted box if it is a cell consisting of a single twisted box (i.e., if R is a singleton).

Remark 5.2.3. Given a cell X as above, let $\pi : K^n \rightarrow K^d$ be the projection to those coordinates i for which $j_i = 1$. Then $\pi(X)$ is a cell of cell-type $(1, \dots, 1)$, and the restriction of π to X is injective. If the components c_i of the center tuple of X are continuous, then $\pi|_X$ is even a homeomorphism.

There are many variants of cell decomposition results preparing different kinds of data. We first state the simplest version and then formulate other versions as addenda:

Theorem 5.2.4 (Cell decomposition). *Suppose that acl equals dcl in $\text{Th}(K)$ (see Assumption 5.2.1) and that $\text{Th}(K)$ is 1-h-minimal. Consider a \emptyset -definable set $X \subset K^n$ for some n . Then there exists a cell decomposition of X , more precisely, a partition of X into finitely many \emptyset -definable cells A_ℓ .*

Remark 5.2.5. Given finitely many \emptyset -definable sets $X_i \subset K^n$, we can find a cell decomposition of K^n such that each X_i is a union of cells, namely by applying Theorem 5.2.4 to suitable intersections of the sets X_i and $K^n \setminus X_i$.

Addendum 1 (Preparation of RV -sets). On top of the assumptions from Theorem 5.2.4, let also a \emptyset -definable set $P \subset X \times \text{RV}^k$ be given for some k . We consider P the function sending $x \in X$ to the fiber $P_x := \{\xi \in \text{RV}^k \mid (x, \xi) \in P\}$.

Then the cells A_ℓ from Theorem 5.2.4 can be taken such that moreover P (seen as a function) is constant on each twisted box of each cell A_ℓ .

Addendum 2 (Continuous functions). On top of the assumptions from Theorem 5.2.4 (with Addendum 1, if desired), suppose that finitely many \emptyset -definable functions $f_j: X \rightarrow K$ are given. Then the A_ℓ can be taken such that moreover, the restriction $f_j|_{A_\ell}$ of each function f_j to each cell A_ℓ is continuous.

In the one-dimensional case, we can prepare the domain and image of the functions in a compatible way:

Addendum 3 (Compatible preparation of domain and image). Under the assumptions of Addendum 2, if the ambient dimension n is equal to 1, we may moreover impose that for each ℓ and each j , $f_j|_{A_\ell}$ is either constant or injective, $f_j(A_\ell)$ is a \emptyset -definable cell and f_j maps each twisted box of A_ℓ onto a twisted box of $f_j(A_\ell)$.

Addendum 4 (Continuous centers). In Theorem 5.2.4 (with any of Addenda 1, 2, if desired), we may assume that for each cell A_ℓ , each component $c_i: \pi_{<i}(A_\ell) \rightarrow K$ of its center tuple is continuous. In particular, each twisted box of each cell A_ℓ of cell-type $(1, \dots, 1)$ is an open subset of K^n .

Remark 5.2.6. In Addenda 2 and 4, one can replace continuous by C^k , provided that one fixes a reasonable definition of a function being C^k on non-interior points of its domain.

To formulate the next addendum, we need the notion of Lipschitz continuity, which is defined in the usual way:

Definition 5.2.7 (Lipschitz continuity). For a valued field K and an element λ in its value group Γ_K^\times , a function $f: X \subset K^n \rightarrow K^m$ is called Lipschitz continuous with Lipschitz constant λ if for all x and x' in X one has

$$|f(x) - f(x')| \leq \lambda|x - x'|,$$

where the norm of tuples is, as usual, the sup-norm. We call such f shortly λ -Lipschitz.

Call f locally λ -Lipschitz, if for each $x \in X$ there is an open neighborhood U of x such that the restriction of f to U is λ -Lipschitz.

Addendum 5 (1-Lipschitz centers). Theorem 5.2.4 (with any of Addenda 1, 2, 4, if desired), is also valid in the following variant: Instead of imposing that A_ℓ itself is a cell in the sense of Definition 5.2.2, we only impose that $\sigma_\ell(A_\ell)$ is a cell for some coordinate permutation $\sigma_\ell: K^n \rightarrow K^n$. In that version, we may moreover impose that $\sigma_\ell(A_\ell)$ is of cell-type $(1, \dots, 1, 0, \dots, 0)$ and that each component c_i of the center tuple $(c_i)_i$ of $\sigma_\ell(A_\ell)$ is 1-Lipschitz.

Closely related to that addendum, we also have the following reformulation of the piecewise continuity result of [14] in the Hensel minimal setting.

Theorem 5.2.8 (Piecewise Lipschitz continuity). *Suppose that acl equals dcl in $\text{Th}(K)$ (see Assumption 5.2.1) and that $\text{Th}(K)$ is 1-h-minimal. Consider a \emptyset -definable set $X \subset K^n$ for some n and a \emptyset -definable function $f: X \rightarrow K$. Suppose that f is locally 1-Lipschitz. Then there exists a finite partition of X into \emptyset -definable sets A_ℓ such that the restriction of f to A_ℓ is 1-Lipschitz, for each ℓ .*

All of the above results in this subsection have already been proved under related assumptions. For the convenience of the reader, we nevertheless provide short, self-contained versions of the proofs, with the exception of Addendum 5 and Theorem 5.2.8 that will not be used in this paper and where we will only explain how to adapt the proof from [14]. However, we do provide a proof of the weak version of Addendum 5 stated as Proposition 5.4.11, since that proposition is used to deduce the existence of t -stratifications.

Proof of Theorem 5.2.4 with Addendum 1. It suffices to find a cell decomposition adapted to a set $P \subset K^n \times \text{RV}^k$.

Case $n = 1$: In this case, Corollary 2.6.6 provides a cell decomposition, namely: Let $C \subset K$ 1-prepare P . Using the assumption that acl equals dcl , each element $c \in C$ is \emptyset -definable, so we obtain a cell

decomposition as follows: For each ball B 1-next to C , choose (in a definable way) a $c(B) \in C$ such that B is 1-next to $c(B)$. Now the cell decomposition consists of two cells with center c for each $c \in C$, namely (i) $\{c\}$ and (ii) the union of all those B 1-next to C for which $c(B)$ equals c .

Case $n > 1$: We apply the case $n = 1$ to each fiber $P_a = \{(b, \xi) \in K \times \text{RV}^k \mid (a, b, \xi) \in P\}$, where a runs over K^{n-1} . (Note that by compactness this works uniformly so that in particular we get definable cell centers $K^{n-1} \rightarrow K$.) Then we finish by applying induction to a set $P' \subset K^{n-1} \times \text{RV}^{k'}$ ‘describing’ the fibers: For each 0-cell $\{c\} \subset K$ of the fiber at $a \in K^{n-1}$, P'_a encodes the set $P_{a,c} \subset \text{RV}^k$; for each 1-cell $X \subset K$ of the fiber at $a \in K^{n-1}$, P'_a encodes (a) the set denoted by R in Definition 5.2.2 and (b) for each $\xi \in R$, the fiber $P_{a,b} \subset \text{RV}^k$, where $b \in K$ is an arbitrary element of the twisted box corresponding to ξ (i.e., $\text{rv}(b - c) = \xi$, where c is the center of X). \square

Proof of Theorem 5.1.1 and Addenda 2, 4. For $n = 0$, all three results are trivial. We now assume that all three results are already known for $n - 1$, and we deduce them for n . Concerning Theorem 5.1.1, note that we may as well assume that acl equals dcl (by Proposition 4.3.3).

Addendum 4: First, find a cell decomposition with possibly non-continuous centers. By inductively applying Addendum 2 to each component of the center tuple of each cell, we may refine the cell decomposition to get continuous centers.

Theorem 5.1.1: We may suppose that X has an empty interior, since the proof given at the end of Subsection 5.1 applies to the interior of X . Choose a cell decomposition of X with continuous centers. No cell $A_\ell \subset X$ is of cell-type $(1, \dots, 1)$ (since such cells have non-empty interior, by the ‘in particular’ part of Addendum 4). Thus the homeomorphism from Remark 5.2.3 allows us to reduce the problem on A_ℓ to one of lower ambient dimension. Apply induction.

Addendum 2: The above proof of Theorem 5.1.1 also yields a finite partition of X such that f is continuous on each piece. Apply this to each of the given functions f_j , and then choose a cell decomposition respecting all pieces from all those partitions. \square

Proof of Addendum 3. Using Lemma 2.8.1 and our assumption $\text{acl} = \text{dcl}$, we find a partition of X such that on each piece, f_j is continuous and either constant or injective for each j ; assume without loss that X is a single such piece. In a similar way, assume without loss that X is a cell and that each f_j has the Jacobian Property (Definition 3.2.8) on each twisted box of X . Since constant functions pose no problem, we assume that all f_j are injective.

Next, choose a finite \emptyset -definable set $\tilde{C} \subset K$ 1-preparing $f_j(B)$ for every j and every twisted box B of X (using Corollary 2.6.6), and set $C := \bigcup_j f_j^{-1}(\tilde{C})$. After a further finite partition of X , we may assume that either (i) X consists of a single twisted box or that (ii) C is empty.

In Case (i), choose any $c \in C \subset X$, and decompose X as $\{c\} \cup (X \setminus \{c\})$, both of which are \emptyset -definable cells. Then the desired properties follow from the Jacobian Property, namely the image of both cells are cells with center $f_j(c)$.

In Case (ii), definably choose, for each twisted box B of X and each j , an element $\tilde{c}_{j,B} \in \tilde{C}$ in such a way that $f_j(B)$ is 1-next to $\tilde{c}_{j,B}$. By a further finite partition of X , we may assume that $\tilde{c}_{j,B}$ does not depend on B . Then $f_j(X)$ is a cell with center $\tilde{c}_{j,B}$, and we are done. \square

As explained above, we will not give detailed proofs of Theorem 5.2.8 and its related result from Addendum 5 to Theorem 5.2.4, and we do not use these results in this paper. We nevertheless specify where this is worked out, under very closely related assumptions.

Proof of Theorem 5.2.8 and Addendum 5 to Theorem 5.2.4. Under our Assumption 5.2.1 that acl equals dcl , but assuming a notion of tameness with an angular component map $\overline{\text{ac}}$ (instead of 1-h-minimality with rv), both results are proved in [14], and the proof readily adapts. (By Theorem 2.9.1 and Remark 2.9.2, 1-h-minimality and the tameness notion are very closely related.) \square

(The proof of Addendum 5 and its variant in [14] essentially come from [13], apart from the improvement made possible by the assumption $\text{acl}_K = \text{dcl}_K$.)

5.3. Dimension theory

Under the assumption of 1-h-minimality, there is a good notion of the dimension of definable subsets of K^n . It can be defined in various equivalent ways; here is one possible definition.

Definition 5.3.1 (Dimension). We define the dimension of a non-empty definable set $X \subset K^n$ as the maximal integer m such that there is a K -linear function $\ell : K^n \rightarrow K^m$ such that $\ell(X)$ has a non-empty interior in K^m . If X is empty, we set $\dim X := -\infty$.

Remark 5.3.2. In Proposition 5.3.4 (1), we will see that one could equivalently only consider coordinate projections $\ell : K^n \rightarrow K^m$ (instead of arbitrary linear maps).

In [20], there already exists a proof that dimension behaves well if one imposes a suitable condition on $\text{Th}(K)$ called b -minimality. By the following proposition, 1-h-minimality implies b -minimality, so for several results about dimension, we could simply refer to [20]. For the convenience of the reader, we will nevertheless give a self-contained proof of those results.

Proposition 5.3.3 (b -minimality). *Assume that $\text{Th}(K)$ is 1-h-minimal. Then the two sorted structure on (K, RV) obtained from K by adding the sort RV , and the map rv is b -minimal in the sense of Definition 2.1 of [20], with K as the main sort. More specifically, the structure (K, RV) is b -minimal with centers and preserves all balls in the sense of Definitions 5.1 and 6.2 of [20].*

Proof. The axioms of Definition 2.1 of [20] clearly hold, and Definition 5.1 of [20] follows from the Jacobian Property as formulated in Corollary 3.2.7. \square

The definition of dimension given in [20, Definition 4.1] is different than ours, but the results from [20, Section 4] imply that the definitions are equivalent: If $X \subset K^n$ is a finite union of cells, then the dimension of X in our sense equals the dimension of X in the sense of [20], namely the maximum of the dimensions of the cells, where the dimension of a cell of cell-type $(j_i)_{i=1}^n$ is $\sum_i j_i$.

The following proposition summarizes the good properties of dimension; in particular, we have definability of dimension, as in o -minimal structures. Property (6) is new.

Proposition 5.3.4 (Dimension theory). *Assume that $\text{Th}(K)$ is 1-h-minimal. Let $X \subset K^n$, $Y \subset K^n$ and $Z \subset K^m$ be non-empty definable sets, and let $f : X \rightarrow Z$ be a definable function. Then the following properties hold:*

1. *For any $d \leq n$, we have $\dim X \geq d$ if and only if there exists a projection $\pi : K^n \rightarrow K^d$ to a subset of the coordinates such that $\pi(X)$ has a non-empty interior. In particular, $\dim X = 0$ if and only if X is finite.*
2. $\dim(X \cup Y) = \max\{\dim X, \dim Y\}$.
3. *For any $d \leq n$, the set of $z \in Z$ such that $\dim f^{-1}(z) = d$ is definable over the same parameters as f .*
4. *If all fibers of f have dimension d , then $\dim X = d + \dim Z$.*
5. *There exists an $x \in X$ such that the local dimension of X at x is equal to the dimension of X : that is, such that for every open ball $B \subset K^n$ around x , we have $\dim(X \cap B) = \dim X$.*
6. *One has $\dim(\overline{X} \setminus X) < \dim X$, where \overline{X} is the topological closure of X , for the valuation topology.*

Although most likely property (6) can be proved in a similar way as Theorem (1.8) of [33], we postpone that proof until the end of Subsection 5.5, where we have t -stratifications at our disposal, which make the proof much simpler. We do, however, right away prove the ‘easy’ case of Property (6), namely when $\dim X$ is equal to the ambient dimension n .

As announced, we present a proof of the proposition that does not rely on b -minimality. Readers willing to take b -minimality as a black box can jump directly to the paragraph ‘Property (5). . .’ (on p. 48) in the proof of the proposition. Indeed, (1)–(4) are proved in [20, Section 4], except for the ‘in particular’ part of (1), which is easy to deduce from the fact that infinite definable subsets of K have non-empty interior.

Our own proof of (1)–(4) consists of showing that dimension agrees with the acl-dimension (which we quickly recall below). In particular, the key is to show that acl has the exchange property (so that the acl-dimension makes sense).

Lemma 5.3.5. *If $\text{Th}(K)$ is 1-h-minimal, then acl_K has the exchange property: that is, for any $a \in K^n$ and any $b, c \in K$ satisfying $c \in \text{acl}_K(a, b) \setminus \text{acl}_K(a)$, we have $b \in \text{acl}_K(a, c)$.*

Proof. Let a, b, c be given as in the lemma; we need to show that $b \in \text{acl}_K(a, c)$. We may assume that a is the empty tuple, otherwise adding it to the language. Since RV-expansions do not change acl_K (by Corollary 4.1.17), we may moreover assume that acl equals dcl (using Proposition 4.3.3). Thus our assumptions yield $c \in \text{dcl}_K(b)$, which implies that there exists a \emptyset -definable function $f: K \rightarrow K$ satisfying $f(b) = c$. If $f^{-1}(c)$ is finite, we deduce that $b \in \text{acl}_K(c)$, and we are done. Otherwise, c lies in the (\emptyset -definable) set Y of those $y \in K$ for which $f^{-1}(y)$ is infinite. By Lemma 2.8.1, Y is finite, so this contradicts $c \notin \text{acl}_K(\emptyset)$. □

Since acl has the exchange property, it defines a pre-geometry and yields a notion $\text{dim}^{\text{acl}}(X)$ of dimension of a definable set $X \subset K^n$; we quickly recall how it is defined. First, note that we may assume that K is \aleph_0 -saturated (otherwise replacing it by an elementary extension). One calls a tuple $b = (b_1, \dots, b_m) \in K^m$ algebraically independent over a set $A \subset K$ if $b_i \notin \text{acl}_K(A, b_1, \dots, b_{i-1})$ for every $i \leq m$. Now, given an A -definable set $X \subset K^n$ for some finite $A \subset K$, we define $\text{dim}^{\text{acl}}(X)$ as the maximal m for which there exists a coordinate projection $\pi: K^n \rightarrow K^m$ such that $\pi(X)$ contains a tuple that is algebraically independent over A . The theory of pre-geometries shows that dim^{acl} is well-defined (i.e., does not depend on the specific finite set A) and well-behaved; in particular, it satisfies Properties (2) and (4) of Proposition 5.3.4, and using \exists^∞ -elimination (Lemma 2.5.2), it also satisfies (3).

To see that our dimension agrees with dim^{acl} , we need the following lemma:

Lemma 5.3.6. *Suppose that K is \aleph_0 -saturated and that $A \subset K$ is finite. An A -definable set $X \subset K^m$ has a non-empty interior if and only if $\text{dim}^{\text{acl}} X = m$: that is, if and only if it contains a tuple that is algebraically independent over A . More precisely, if $a \in X$ is algebraically independent over A , then a lies in the interior of X .*

Proof. For the implication from left to right, suppose without loss that X is a ball. Then we find the desired algebraically independent tuple $a \in X$ by choosing one coordinate after the other, each one outside of the algebraic closure of A and the previously chosen coordinates; this is possible by \aleph_0 -saturation.

For the other direction (including the ‘more precisely’ part), suppose that $a \in X$ is algebraically independent over A , and write it as $a = (\hat{a}, a') \in K^{m-1} \times K$. On the one hand, \hat{a} is algebraically independent over $A \cup \{a'\}$, so by induction, it lies in the interior of the fiber $X_{a'} \subset K^{m-1}$. Fix $\lambda \in \Gamma_K^\times$ such that $B_{<\lambda}(\hat{a}) \subset X_{a'}$. On the other hand, we have $a' \notin \text{acl}_K(A \cup \{\hat{a}\})$. By saturation, we find a ball $B \subset K \setminus \text{acl}(A \cup \{\hat{a}\})$ containing a' , and by 0-h-minimality in the form of Lemma 2.4.4, all elements of B have the same type over $A \cup \{\hat{a}\} \cup \text{RV}^{\text{eq}}$ as a' . In particular, we have $B_{<\lambda}(\hat{a}) \subset X_{a''}$ for every $a'' \in B$, which shows that X contains $B \times B_{<\lambda}(\hat{a})$. □

Proof of Proposition 5.3.4, except for (6) when $\text{dim} X < n$. We claim that for every definable set $X \subset K^n$, we have $\text{dim}^{\text{acl}}(X) = \text{dim}(X)$.

To prove this claim, we suppose that K is \aleph_0 -saturated. Set $m := \text{dim}^{\text{acl}}(X)$. Then there exists a coordinate projection $\ell: K^n \rightarrow K^m$ such that $\ell(X)$ contains an algebraically independent tuple (over the parameters needed to define X). By Lemma 5.3.6, this tuple lies in the interior of $\ell(X)$, which shows that $\text{dim}(X) \geq m$. For the other inequality, suppose for contradiction that we have a linear map $\ell: K^n \rightarrow K^{m'}$ for some $m' > m$ such that $\ell(X)$ has a non-empty interior. Then by Lemma 5.3.6, $\ell(X)$ contains an algebraically independent tuple (over the parameters of $\ell(X)$), so $\text{dim}^{\text{acl}}(\ell(X)) = m'$, but this contradicts that taking the image under a definable map cannot increase dim^{acl} .

This already proves (2)–(4), and we also get (1) for free, using once more that a definable set in K^d has a non-empty interior if and only if its dimension is d . (The ‘in particular’ part of (1) holds trivially for \dim^{acl} .)

Property (5) is proved in [40] in a much more general context; here is a much shorter proof in the present setting: We may assume that acl equals dcl in $\text{Th}(K)$ (using Proposition 4.3.3) so that we can apply cell decomposition (Theorem 5.2.4) to X ; we also use Addendum 4 to get continuous centers.

Choose a cell $A_\ell \subset X$ of maximal dimension (of cell-type $(j_i)_{i=1}^n$, with $\sum_i j_i = \dim X =: d$). Then for any $x \in A_\ell$, the local dimension of X at x is d . Indeed, the projection $\pi(A_\ell) \subset K^d$ to the coordinates $\{i \leq n \mid j_i = 1\}$ is a cell of cell-type $(1, \dots, 1)$ with continuous center, and hence open. Thus for every sufficiently small ball $B \subset K^n$ around x , we have $\pi(X \cap B) = \pi(B)$, witnessing $\dim(X \cap B) \geq d$.

To prove Property (6) in the case $\dim X = n$, we again first expand the structure so that acl equals dcl and then find a cell decomposition of $\bar{X} \setminus X$. Since every n -dimensional cell has a non-empty interior, no such cells can be contained in $\bar{X} \setminus X$. This implies $\dim(\bar{X} \setminus X) < n$. □

5.4. Jacobian Properties in many variables

There are different ways to generalize the Jacobian Property (Definition 3.2.8) to functions f in several variables. The one presented in this subsection (which we now call the Supremum Jacobian Property) has been introduced in [43] and is used to obtain t-stratifications. (To be precise, [43, Definition 2.19] is a bit weaker than Definition 5.4.5 below.)

First of all, we need a specific higher-dimensional version of the rv-map.

Notation 5.4.1 (\lesssim). Given $\lambda, \mu \in \Gamma_K$, we define $\lambda \lesssim \mu$ as $\lambda < \mu \vee \lambda = \mu = 0$.

Definition 5.4.2 (Higher-dimensional RV). For every $n \geq 1$, we define $\text{RV}^{(n)}$ as the quotient K^n / \sim , where $x \sim x' \iff |x - x'| \lesssim |x|$. We write $\text{rv}^{(n)}$ for the canonical map $K^n \rightarrow \text{RV}^{(n)}$. (For matrices $M \in K^{n \times m}$, we will use the more suggestive notation $\text{rv}^{(n \times m)}(M)$ instead of $\text{rv}^{(n \cdot m)}(M)$.)

(Recall that for $x \in K^n$, $|x|$ denotes the maximum norm of x .)

Remark 5.4.3. Note that $\text{RV}^{(1)}$ is just the usual RV. For $n \geq 2$, $\text{RV}^{(n)}$ is not the same as RV^n , but $\text{rv}^{(n)}$ factors over coordinate-wise rv so that we have a natural surjection $\text{RV}^n \rightarrow \text{RV}^{(n)}$. Moreover, the maximum norm on K^n factors over $\text{RV}^{(n)}$.

Remark 5.4.4. As explained in [43, Section 2.2], $\text{rv}^{(n)}$ interacts well with $\text{GL}_n(\mathcal{O}_K)$; in particular, given $M \in \text{GL}_n(\mathcal{O}_K)$ and $x \in K^n$, $\text{rv}^{(n)}(Mx)$ is determined by $\text{rv}^{(n \times n)}(M)$ and $\text{rv}^{(n)}(x)$.

Definition 5.4.5 (Sup-Jac-prop, sup-preparation). For $X \subset K^n$ open and $f: X \rightarrow K$, we say that f has the *Supremum Jacobian Property* (*sup-Jac-prop* for short) on X if f is C^1 on X , $\text{rv}^{(n)}(\text{grad } f)$ is constant on X , and for every x_0 and x in X we have:

$$|f(x) - f(x_0) - ((\text{grad } f)(x_0)) \cdot (x - x_0)| \lesssim |\text{grad } f| \cdot |x - x_0|. \tag{5.4.1}$$

As usual, we consider $(\text{grad } f)(x_0)$ a matrix with a single row, which we multiply with the column vector $x - x_0$ in the usual way. We say that a map $\chi: X \rightarrow \text{RV}^k$ *sup-prepares* f (for some $k \geq 0$) if each n -dimensional fiber $F \subset K^n$ of χ is open and f has the sup-Jac-prop on each such F .

Remark 5.4.6. One easily checks that the validity of equation (5.4.1) does not depend on the precise value of $(\text{grad } f)(x_0)$, but only on $\text{rv}^{(n)}((\text{grad } f)(x_0))$, so it does not play a role whether we evaluate $\text{grad } f$ at x_0 or at any other point of X .

Remark 5.4.7. In the case $n = 1$, equation (5.4.1) is equivalent to $\text{rv}(f(x) - f(x_0)) = \text{rv}(f'(x_0)) \cdot \text{rv}(x - x_0)$ (which is exactly the main condition of the one-dimensional Jacobian Property; see Definition 3.2.8). For $n \geq 2$ however, equation (5.4.1) does not always determine $\text{rv}(f(x) - f(x_0))$. Indeed, if for example $x - x_0$ is orthogonal to $(\text{grad } f)(x_0)$, then equation (5.4.1) only imposes an upper bound on $|f(x) - f(x_0)|$.

Remark 5.4.8. One cannot expect to be able to sup-prepare definable functions in the stronger sense that $\text{rv}(f(x) - f(x_0))$ is equal to $\text{rv}(((\text{grad } f)(x_0)) \cdot (x - x_0))$ within fibers of χ (which would correspond to replacing the right-hand side of equation (5.4.1) by $|((\text{grad } f)(x_0)) \cdot (x - x_0)|$). Indeed, consider for example $f(x, y) = y - x^2$, fix any $(x_0, y_0) \in K^2$ and any $\epsilon \in K^\times$, and set $(x, y) := (x_0 + \epsilon, y_0 + 2x_0\epsilon)$. Then $((\text{grad } f)(x_0, y_0)) \cdot ((x, y) - (x_0, y_0)) = 0$ but $f(x, y) - f(x_0, y_0) \neq 0$. For any χ potentially preparing f , we can choose x_0, y_0 and ϵ such that (x_0, y_0) and (x, y) lie in the same fiber.

The following lemma states that the sup-Jac-prop is preserved by certain transformations.

Lemma 5.4.9 (Preservation of sup-Jac-prop). *Let $X \subset K^m$ and $Y \subset K^n$ be open subsets, and let $\alpha: X \rightarrow Y$ be a C^1 -map. Suppose that $\text{rv}^{(n \times m)}(\text{Jac } \alpha)$ is constant on X and that for every $x_1, x_2 \in X$, we have*

$$|\alpha(x_2) - \alpha(x_1)| = |\text{Jac } \alpha| \cdot |x_2 - x_1| \tag{5.4.2}$$

and

$$\text{rv}^{(n)}(\alpha(x_2) - \alpha(x_1)) = \text{rv}^{(n)}((\text{Jac } \alpha)(x_1) \cdot (x_2 - x_1)). \tag{5.4.3}$$

Finally, suppose that $f: Y \rightarrow K$ is a C^1 -map such that $f \circ \alpha$ has the sup-Jac-prop (on X). Then f satisfies equation (5.4.1) for all $x_0, x \in \alpha(X)$.

Proof. Let $x_1, x_2 \in X$ be given, and set $y_i := \alpha(x_i)$ and $z_i := f(y_i)$. In the following, gradients and Jacobians will always be computed at x_1 or y_1 ; we will omit those points from the notation.

What we need to show is:

$$|z_2 - z_1 - (\text{grad } f) \cdot (y_2 - y_1)| \leq |\text{grad } f| \cdot |y_2 - y_1|. \tag{5.4.4}$$

By assumption, we have

$$|z_2 - z_1 - (\text{grad}(f \circ \alpha)) \cdot (x_2 - x_1)| \leq |\text{grad}(f \circ \alpha)| \cdot |x_2 - x_1|. \tag{5.4.5}$$

Applying $\text{grad}(f \circ \alpha) = (\text{grad } f) \cdot (\text{Jac } \alpha)$ to the right-hand side of equation (5.4.5) gives

$$|\text{grad}(f \circ \alpha)| \cdot |x_2 - x_1| \leq |\text{grad } f| \cdot |\text{Jac } \alpha| \cdot |x_2 - x_1| \stackrel{\text{equation (5.4.2)}}{=} |\text{grad } f| \cdot |y_2 - y_1|.$$

On the left-hand side of equation (5.4.5), we do the following:

$$(\text{grad}(f \circ \alpha)) \cdot (x_2 - x_1) = (\text{grad } f) \cdot (\text{Jac } \alpha) \cdot (x_2 - x_1) \approx (\text{grad } f) \cdot (y_2 - y_1),$$

where in the ‘ \approx ’, we use equation (5.4.3) to get an error e with $e \leq |\text{grad } f| \cdot |y_1 - y_2|$ (and this is what ‘ \approx ’ means here). Putting things together yields equation (5.4.4), as desired. \square

The main result of this subsection is that every definable function on K^n can be sup-prepared:

Theorem 5.4.10 (Sup-preparation). *Suppose that $\text{Th}(K)$ is 1-h-minimal. For every \emptyset -definable function $f: K^n \rightarrow K$, there exists a \emptyset -definable map $\chi: K^n \rightarrow \text{RV}^k$ (for some $k \geq 0$) sup-preparing f (in the sense of Definition 5.4.5).*

The proof needs some kind of cell decomposition with 1-Lipschitz centers, as for example provided by Theorem 5.2.4, Addendum 5. To keep this paper more self-contained (since we did not give the proof of Addendum 5 in full detail), we will instead prove and use the following weaker version of the addendum; more precisely, this proposition is proved in a joint induction with Theorem 5.4.10.

Proposition 5.4.11 (Twisted boxes with 1-Lipschitz centers). *Assume 1-h-minimality and that acl equals dcl (in the sense of Assumption 5.2.1). Then, for every \emptyset -definable set $X \subset K^n$, there exists a \emptyset -definable map $\chi: X \rightarrow \text{RV}^{k'}$ such that each nonempty fiber F of χ is, up to permutation of coordinates, an*

RV-definable twisted box of cell-type $(1, \dots, 1, 0, \dots, 0)$ with 1-Lipschitz center (i.e., each component $c_i: \pi_{<i}(F) \rightarrow K$ of the center tuple is 1-Lipschitz).

Here, by ‘twisted box of cell-type...’, we mean a cell with those properties consisting of a single twisted box; see Definition 5.2.2.

A key step in the proof of the proposition consists of swapping two coordinates to make the derivative of some center smaller. This uses the following lemma, which has a similar but simpler proof than Cases 1 and 2 of the proof of Proposition 2.4 of [11]. By a ‘genuine box’ in K^n , we mean a Cartesian product of balls in K (as opposed to the more general notion of ‘twisted box’).

Lemma 5.4.12. *Assuming only that K is a valued field in a language containing \mathcal{L}_{val} , suppose that*

$$X = \{(x_1, x_2) \in K^2 \mid \text{rv}(x_1 - c_1) = \xi_1, \text{rv}(x_2 - c_2(x_1)) = \xi_2\}$$

is an A-definable twisted box for some set of parameters $A \subset K$ and that c_2 has the Jacobian Property. Then either X is a genuine box (i.e., a Cartesian product of two balls) or we have $X = \{(x_1, x_2) \in K \times Y \mid \text{rv}(x_1 - c_2^{-1}(x_2)) = \xi_3\}$ for some suitable constant $\xi_3 \in \text{dcl}_{\text{RV}}(A)$ and where Y is the projection of X to the second coordinate.

Proof. If c_2 is constant, there is nothing to do. So, suppose that c_2 is not constant. Then the image $c_2(\pi_{\leq 1}(X))$ is an open ball of radius $\rho := |\xi_1| \cdot |c'_2|$ (where c'_2 is the derivative of c_2). If $\rho \leq |\xi_2|$, then whether $\text{rv}(x_2 - c_2(x_1)) = \xi_2$ holds does not depend on the choice on $x_1 \in \pi_{\leq 1}(X)$ and X is a genuine box. Otherwise, if $\rho > |\xi_2|$, then c_2^{-1} is defined on the whole of Y and the Jacobian property implies that $\text{rv}(x_2 - c_2(x_1)) = \xi_2$ if and only if $\text{rv}(x_1 - c_2^{-1}(x_2)) = -\text{rv}(c'_2)^{-1} \cdot \xi_2$. □

Proof of Proposition 5.4.11 in the case $n = 1$. Prepare X by a finite set C , and let χ be the map given by Lemma 2.5.4. □

Proof of Proposition 5.4.11, assuming Proposition 5.4.11 and Theorem 5.4.10 for $n - 1$. By an ‘RV-partition’ of X , we mean a partition into fibers of an RV-definable map $X \rightarrow \text{RV}^k$. Note that if we are already given an RV-partition of X , it suffices to prove the proposition for each fiber individually. (Then put everything together using compactness.)

We say that X is a ‘thick graph’ (of the map c_n) if it is of the form $\{(y, x_n) \in Y \times K \mid \text{rv}(x_n - c_n(y)) = \xi\}$ for some $Y \subset K^{n-1}$, some $c_n: Y \rightarrow K$ and some $\xi \in \text{RV}$. (Note that we allow $\xi = 0$, which means that X is just the graph of c_n .)

Note that it suffices to obtain the claim in the last coordinate: that is, to RV-partition X into sets that are, up to permutation of coordinates, thick graphs of 1-Lipschitz functions $c_n: Y \rightarrow K$. After that, the proposition follows by applying induction to Y .

Using cell decomposition, we reduce to the case where X is a thick graph of a function $c_n: Y \rightarrow K$ and Y is a twisted box. In particular, X is a twisted box. We may assume that Y either is open or has an empty interior (by treating the interior separately).

Step 1: If Y has an empty interior, we apply induction to Y to reduce to the case that Y is a twisted box with 1-Lipschitz centers, and we translate the centers away so that X lives in a subspace of K^n where some of the coordinates are 0; then apply induction once more to finish. Note that translating the variables of a 1-Lipschitz function by 1-Lipschitz functions yields a 1-Lipschitz function, and hence the original X does have the required properties.

So suppose from now on that Y is open. By partitioning Y , we may assume that c_n is C^1 and that $|\partial c_n / \partial x_i|$ is constant on Y for each i . (Lower-dimensional pieces are treated as in Step 1.)

Step 2: Assume $|\text{grad } c_n| \leq 1$. Using the $n - 1$ case of Theorem 5.4.10, we may assume that c_n has the sup-Jac-prop. This, together with $|\text{grad } c_n| \leq 1$, implies that c_n is 1-Lipschitz, and hence we are done.

Step 3: So now suppose $|\text{grad } c_n| > 1$. We do an induction on the number of partial derivatives of c_n satisfying $|\partial c_n / \partial x_i| > 1$. We suppose without loss that $|\partial c_n / \partial x_{n-1}| = |\text{grad } c_n|$. Let Z be the projection of Y to the first $n - 2$ coordinates. By further partitioning, we reduce to the case where, for each individual

$a \in Z$, the function $c_n(a, \cdot)$ has the Jacobian Property (using Corollary 3.2.7 and compactness) and has an open ball as a domain. (Again, lower-dimensional pieces are treated as in Step 1.)

Note that for each $a \in Z$, the fiber $X_a \subset K^2$ is a twisted box. By further partitioning Z , we may assume that either all of them or none of them are genuine boxes.

Step 3.a: If all fibers are genuine boxes: by induction on n , we may assume that the projection \tilde{X} of X to the coordinates $1, \dots, n-2, n$ is a thick graph of a 1-Lipschitz function $\tilde{c}_n: Z \rightarrow K$. (This involves permuting the coordinates $1, \dots, n-2, n$.) Then X is a thick graph of $c_n(z, x_{n-1}) := \tilde{c}_n(z)$, which is 1-Lipschitz, so we are done.

Step 3.b: If no fiber is a genuine box, we apply the map $\sigma: K^n \rightarrow K^n$ swapping the coordinates $n-1$ and n . By compactness and Lemma 5.4.12, $\sigma(X)$ is the thick graph of the function $c_{n,\text{new}}$ sending $(x_1, \dots, x_{n-2}, c_n(x_{n-1}))$ to x_{n-1} . Since $|\partial c_{n,\text{new}}/\partial x_i| \leq |\partial c_n/\partial x_i|$ for $i \leq n-2$ and $|\partial c_{n,\text{new}}/\partial x_n| = 1/|\partial c_n/\partial x_{n-1}| < 1$, $c_{n,\text{new}}$ has fewer partial derivatives bigger than 1, so we can finish by the induction from Step 3. □

Proof of Theorem 5.4.10 in the case $n = 1$. This follows directly from Corollary 3.2.7 and Lemma 2.5.4: The corollary yields a finite \emptyset -definable set C such that equation (5.4.1) holds on every ball 1-next to C , and Lemma 2.5.4 then yields a map $\chi: K \rightarrow \text{RV}^k$ whose 1-dimensional fibers are exactly those balls. □

Proof of Theorem 5.4.10, assuming Proposition 5.4.11 for n and Theorem 5.4.10 for $n - 1$. The proof consists of three parts.

Part 1: Some preliminaries:

Claim 5.4.13. *It suffices to prove the theorem under the assumption that acl equals dcl .*

Proof. Let $\mathcal{L}^{\text{as}} \supset \mathcal{L}$ be as given by Proposition 4.3.3: that is, $\text{Th}_{\mathcal{L}^{\text{as}}}(K)$ is still 1-h-minimal, and in $\text{Th}_{\mathcal{L}^{\text{as}}}(K)$, we have acl equals dcl . Assuming that Theorem 5.4.10 holds for this language, we find an \mathcal{L}^{as} -definable $\chi': K^n \rightarrow \text{RV}^{k'}$ sup-preparing f . Since the \mathcal{L}^{as} -structure on K is an RV-expansion of the \mathcal{L} -structure, Lemma 4.3.4 provides an \mathcal{L} -definable $\chi: K^n \rightarrow \text{RV}^k$ such that each fiber F of χ is contained in a fiber of χ' ; in particular, equation (5.4.1) holds whenever $(\text{grad}(f|_F))(x_0)$ is defined. It remains to refine χ in such a way that each of its n -dimensional fibers is open. We do this by splitting each n -dimensional fiber F into its interior $\overset{\circ}{F}$ and the remainder. Then indeed $\overset{\circ}{F}$ is open, and $F \setminus \overset{\circ}{F}$ has dimension less than n , by Proposition 5.3.4 (6) applied to $K^n \setminus F$. (Note that we only use the case of Proposition 5.3.4 (6) that we already proved right after the proposition.) □

So for the remainder of the proof, we assume that acl equals dcl (so that we can apply cell decomposition and Proposition 5.4.11).

Recall that we inductively assume that the theorem holds up to dimension $n - 1$. From this, we deduce the following for functions defined on some n -dimensional neighborhoods of lower-dimensional subsets of K^n .

Claim 5.4.14. *Given any $(n - 1)$ -dimensional \emptyset -definable $Z \subset K^n$ and any \emptyset -definable C^1 -function f to K defined on an open neighborhood of Z , there exists a \emptyset -definable map $\chi: Z \rightarrow \text{RV}^k$ (for some $k \geq 0$) such that if $F \subset Z$ is an $(n - 1)$ -dimensional fiber of χ , then equation (5.4.1) holds for every pair $x_0, x \in F$.*

Proof. We can (and will repeatedly) partition Z into fibers of a \emptyset -definable map $Z \rightarrow \text{RV}^k$. (If the claim holds for each fiber of such a partition, we then obtain the desired map $Z \rightarrow \text{RV}^{k'}$ using compactness.) By partitioning Z into the twisted boxes of a cell decomposition, we may assume that $\text{rv}^{(n)}(\text{grad } f)$ is constant. By Proposition 5.4.11, we may assume that Z is a twisted box of cell-type $(1, \dots, 1, 0)$ with 1-Lipschitz center.

Let \hat{Z} be the projection of Z to the first $n - 1$ coordinates so that Z is the graph of a 1-Lipschitz function $c: \hat{Z} \rightarrow K$. Apply Theorem 5.4.10 to c and $f \circ \alpha$, where $\alpha: \hat{Z} \rightarrow Z, x \mapsto (x, c(x))$, and partition \hat{Z} and Z accordingly: that is, so that after the partition, c and $f \circ \alpha$ have the sup-Jac-prop on \hat{Z} .

Using that c is 1-Lipschitz, one obtains that the map $\alpha: \hat{Z} \rightarrow Z, x \mapsto (x, c(x))$ satisfies the assumptions of Lemma 5.4.9. Therefore, the fact that $f \circ \alpha$ satisfies equation (5.4.1) on \hat{Z} implies that f satisfies equation (5.4.1) on Z , as desired. \square

Let now a \emptyset -definable function $f: K^n \rightarrow K$ be given (with $n \geq 2$); we need to find a \emptyset -definable map $K^n \rightarrow \text{RV}^k$ sup-preparing f . We more generally allow the domain of f to be any \emptyset -definable set $X \subset K^n$. As in the proof of Claim 5.4.14, if we have a \emptyset -definable map $\chi: X \rightarrow \text{RV}^k$, it suffices to sup-prepare the restrictions of f to each fiber of χ . Moreover, fibers of dimension less than n can always be neglected. This argument will be applied repeatedly.

We write elements of K^n as (x, y) , with $x \in K^{n-1}$ and $y \in K$.

Part 2: Reducing to the case where X is of the form $\bar{X} \times B$ for some $\bar{X} \subset K^{n-1}$ and some ball $B \subset K$, f is C^1 , $\text{rv}^{(n)}(\text{grad } f)$ is constant, and

(SJP1) for each fixed $x \in \bar{X}$, the function $f(x, \cdot)$ has the sup-Jac-prop.

We start by partitioning K^n as follows:

- (a) By repeatedly applying the case $n = 1$ (and using compactness), we may assume that f has the sup-Jac-prop fiberwise: for every fixed $x \in K^{n-1}$ and every coordinate permutation $\sigma: K^n \rightarrow K^n$, the map $y \mapsto f(\sigma(x, y))$ has the sup-Jac-prop.
- (b) We moreover assume that f is C^1 (using Theorem 5.1.5) and that $\text{rv}^{(n)}(\text{grad } f)$ is constant.

By applying Proposition 5.4.11 and permuting coordinates, we may assume that X is a twisted box with 1-Lipschitz center:

$$X = \{(x, y) \in K^{n-1} \times K \mid x \in \bar{X}, \text{rv}(y - c(x)) = \rho\} \tag{5.4.6}$$

for some $\rho \in \text{RV}^\times$, some definable 1-Lipschitz $c: \bar{X} \rightarrow K$ and where $\bar{X} := \pi_{\leq n-1}(X) \subset K^{n-1}$ is the projection to the first $n - 1$ coordinates. We may assume that c is C^1 , and applying the Theorem inductively to c allows us to moreover assume that c has the sup-Jac-prop.

Set $X' := \bar{X} \times \text{rv}^{-1}(\rho) \subset K^{n-1} \times K$. The bijection $\alpha: X' \rightarrow X, (x, y) \mapsto (x, y + c(x))$ satisfies the assumptions of Lemma 5.4.9, so to prove that f has the sup-Jac-prop on some set $F \subset X$, it suffices to verify that $f' := f \circ \alpha$ has the sup-Jac-prop on $\alpha^{-1}(F)$. In other words, it remains to sup-prepare $f': X' \rightarrow K$.

By (b), f' is C^1 and $\text{rv}^{(n)}(\text{grad } f') = \text{rv}^{(n)}((\text{grad } f)(\text{Jac } \alpha))$ is constant (using Remark 5.4.4); by (a), $f'(x, \cdot)$ has the sup-Jac-prop for each $x \in \bar{X}$; thus we are done with Part 2.

Part 3: Finishing under the assumptions obtained in Part 2.

Recall that $X = \bar{X} \times B \subset K^{n-1} \times K$.

(SJP2) By induction, we find a map $\chi: X \rightarrow \text{RV}^k$ such that for each fixed $y \in B$, $g(\cdot, y)$ is sup-prepared by $\chi(\cdot, y)$.

We choose a cell decomposition of X with continuous centers, and we refine χ in such a way that the fibers of χ are exactly the twisted boxes of the cells. Given such a cell A_ℓ , let $\bar{A}_\ell := \pi_{\leq n-1}(A_\ell)$ be its projection, and let $c_\ell: \bar{A}_\ell \rightarrow K$ be the last component of its center tuple.

Let Z_ℓ be the intersection of the graph of c_ℓ with X . We apply Claim 5.4.14 to Z_ℓ and f , yielding a map $\chi_\ell: Z_\ell \rightarrow \text{RV}^k$, we extend χ_ℓ by 0 to the whole graph of c_ℓ , and we replace χ by the refinement $(x, y) \mapsto (\chi(x, y), (\chi_\ell(x, c_\ell(x))))_\ell$. In this way, we achieve the following:

(SJP3) Given any $(x_1, y_1), (x_2, y_2)$ in the same n -dimensional fiber of χ , and given any ℓ , the pair of points $(x_i, c_\ell(x_i))$ ($i = 1, 2$) satisfies equation (5.4.1), provided that both of those points $(x_i, c_\ell(x_i))$ lie in X .

Note that since this refinement of χ only depends on x , each fiber F of χ is still of the form

$$F = \{(x, y) \in K^{n-1} \times K \mid x \in \bar{F}, \text{rv}(y - c_\ell(x)) = \xi\} \subset \bar{X} \times B \tag{5.4.7}$$

for some ℓ , for some $\xi \in \text{RV}$ and where $\bar{F} = \pi_{\leq n-1}(F)$. Using one last refinement of χ (also depending only on x), we may assume that \bar{F} either is open or has dimension less than $n - 1$, so that if F is n -dimensional, it is open. To finish the proof of the theorem, we will prove that f has the sup-Jac-prop on each such n -dimensional fiber F .

We already know that f is C^1 on F and that $\text{rv}^{(n)}(\text{grad } f)$ is constant on F , so it remains to verify equation (5.4.1); thus let $(x_1, y_1), (x_2, y_2) \in F$ be given.

Recall (Remark 5.4.6) that in equation (5.4.1), it does not matter at which point of F we evaluate the gradient $\text{grad } f$. Using this, an easy computation shows that equation (5.4.1) can be verified in several steps, jumping to certain intermediate points $(x_3, y_3) \in F$ first, namely: If equation (5.4.1) holds for $(x_1, y_1), (x_3, y_3)$ and also for $(x_3, y_3), (x_2, y_2)$, and if moreover $|(x_1, y_1) - (x_3, y_3)| \leq |(x_1, y_1) - (x_2, y_2)|$, then equation (5.4.1) follows for $(x_1, y_1), (x_2, y_2)$. In a similar way, we can also jump through several intermediate points. Note also that the intermediate points can be arbitrary points of X (and do not need to lie in F), since $\text{rv}^{(n)}(\text{grad } f)$ is constant on all of X .

We use the notation from equation (5.4.7) and distinguish three cases:

Case 1: $|c_\ell(x_1) - y_1| > |y_2 - y_1|$. Then we have $(x_1, y_2) \in F$, so we can jump from (x_1, y_1) to (x_1, y_2) by (SJP1) and from (x_1, y_2) to (x_2, y_2) by (SJP2).

Case 2: $|c_\ell(x_2) - y_2| > |y_2 - y_1|$: analogous to Case 1.

Case 3: $|c_\ell(x_i) - y_i| \leq |y_2 - y_1|$ for $i = 1, 2$: From $y_1, y_2 \in B$, we deduce $c_\ell(x_i) \in B$, so that $(x_i, c_\ell(x_i))$ lies in X . Moreover, we have $|(x_i, y_i) - (x_i, c_\ell(x_i))| \leq |(x_1, y_1) - (x_2, y_2)|$. This means that we can jump from (x_1, y_1) to $(x_1, c_\ell(x_1))$ by (SJP1), then to $(x_2, c_\ell(x_2))$ by (SJP3), and then to (x_2, y_2) by (SJP1) again. □

5.5. *t*-stratifications

In [43], a notion of stratifications in valued fields has been introduced, called ‘*t*-stratifications’. Intuitively, given a definable set $X \subset K^n$, a *t*-stratification captures, for every ball $B \subset K^n$, the dimension of the space of directions in which $X \cap B$ is ‘roughly translation invariant’. This strengthens classical notions of stratifications (like Whitney or Verdier stratifications), which capture rough translation invariance only locally.

The existence proof of *t*-stratifications given in [43] is carried out under some axiomatic assumptions, namely [43, Hypothesis 2.21]. Those assumptions hold in valued fields with analytic structure (in the sense of [18]) by [43, Proposition 5.12] and in power-bounded *T*-convex structures by [41]. We will now show that the assumptions hold in any 1-h-minimal theory of equi-characteristic 0, hence implying that *t*-stratifications exist in this generality. By the examples of 1-h-minimal theories given in Section 6, this generalizes both of the above results.

In this entire section, let K be an equi-characteristic 0 valued field with 1-h-minimal theory.

We quickly recall the necessary definitions related to *t*-stratifications. First, here is the precise notion of ‘roughly translation invariant’:

Definition 5.5.1 (Risometries, translatability). Let $B \subset K^n$ be a ball.

1. A bijection $f : B \rightarrow B$ is a *risometry* if for every $x_1, x_2 \in B$, we have $\text{rv}^{(n)}(f(x_1) - f(x_2)) = \text{rv}^{(n)}(x_1 - x_2)$.
2. A map $\chi : K^n \rightarrow Q$ (for some arbitrary set Q) is *d-translatable* on B , for some $0 \leq d \leq n$, if there exists a definable (with parameters) risometry $f : B \rightarrow B$ and a d -dimensional sub-vector space $V \subset K^n$ such that for every $x, x' \in B$ satisfying $x - x' \in V$, we have $\chi(f(x)) = \chi(f(x'))$.
3. A subset $X \subset K^n$ is called *d-translatable* if its characteristic function $1_X : K^n \rightarrow \{0, 1\}$ is *d-translatable*.

Definition 5.5.2 (*t*-stratifications). Let $\chi : K^n \rightarrow Q$ be a map (for some arbitrary set Q), and let A be a set of parameters. An *A-definable t-stratification reflecting χ* is a partition of K^n into A -definable sets S_0, \dots, S_n with the following properties:

1. $\dim S_d \leq d$.
2. Set $\chi'(x) := (\chi(x), d(x)) \in Q \times \{0, \dots, n\}$, where $d(x)$ is defined such that $x \in S_d$ for every $x \in K^n$. For each $d \leq n$ and each open or closed ball $B \subset S_d \cup \dots \cup S_n$, this map χ' is d -translatable on B .

Theorem 5.5.3 (t-stratifications). *Let K be a valued field of equi-characteristic 0 with 1-h-minimal theory, and let $\chi: K^n \rightarrow Q$ be a \emptyset -definable map, where Q is a sort of RV^{eq} . Then there exists a \emptyset -definable t-stratification $(S_i)_{i \leq n}$ reflecting χ .*

Proof. According to [43, Theorem 4.12], the existence of t-stratifications follows from [43, Hypothesis 2.21], so the theorem follows from the following lemma. □

Lemma 5.5.4. *Hypothesis 2.21 of [43] holds in 1-h-minimal theories.*

Proof. The hypothesis consists of the following four parts.

1. RV is stably embedded:
This is Proposition 2.6.12.
2. Definable maps from RV to K have finite image:
This is (a special case of) Corollary 2.6.8.
3. For every $A \subset K \cup \text{RV}^{\text{eq}}$, every A -definable set $X \subset K$ can be 1-prepared by a finite A -definable set $C \subset K$:
This is clear from the definition of 1-h-minimality and Lemma 2.4.1.
4. The theory has the Jacobian Property in the sense of [43, Theorem 2.19], namely: For every $A \subset K \cup \text{RV}^{\text{eq}}$, every A -definable function $f: K^n \rightarrow K$ can be sup-prepared (in the sense of Definition 5.4.5) by an A -definable map $\xi: K^n \rightarrow Q$, where Q is a sort of RV^{eq} :
Add A as constants to the language. Then (4) is just Theorem 5.4.10. □

Note that the proof we gave here also simplifies the proofs from [43] (in the case of fields with analytic structure) and [41] (in the case of T -convex structures): In those papers, the proof of (4) was done using a complicated inductive argument using the existence of t-stratifications in a lower dimension. This has been replaced by the more direct proof of our Theorem 5.4.10.

We end this subsection with the promised proof of the missing part of Proposition 5.3.4, namely that for definable sets $X \subset K^n$, the frontier $\bar{X} \setminus X$ has lower dimension than X :

Proof of Proposition 5.3.4 (6). Choose a t-stratification reflecting the Cartesian product $\chi(x) := (1_X(x), 1_Y(x))$ of the characteristic functions of X and the frontier $Y := \bar{X} \setminus X$, and set $d := \dim Y$. For dimension reasons, Y contains at least one point $y \in S_d$. (Note that the definition of t-stratification implies $Y \cap S_j = \emptyset$ for $j > \dim Y$; see [43, Lemma 3.10].) Assuming $\dim X \leq d$, we will show that Y cannot be contained in the topological closure of X .

Since $S_{\leq d-1} := S_0 \cup \dots \cup S_{d-1}$ is closed (by [43, Lemma 3.17 (a)]), there exists a ball $B \subset S_d \cup \dots \cup S_n$ containing y . Let $f: B \rightarrow B$ be a isometry and $V \subset K^n$ be a vector space witnessing d -translatability of χ on this B , as in Definition 5.5.1. Since f is a homeomorphism (and $y \notin X$), to obtain $y \notin \bar{X}$, it suffices to show that $X' := f^{-1}(X \cap B)$ is closed in B . Let $\pi: K^n \rightarrow K^n/V$ be the canonical map. The definition of d -translatability implies that $X' = B \cap \pi^{-1}(\pi(X'))$. Now the assumption $\dim X \leq d$ implies $\dim \pi(X') = 0$, so indeed X' is closed in B . □

5.6. Taylor approximation on boxes disjoint from a lower-dimensional set

We prove a higher-dimensional version of the Taylor approximation Theorem 3.2.2. By a *box* in K^n , we mean a Cartesian product of balls in K .

Theorem 5.6.1 (Taylor approximations on boxes). *Given a \emptyset -definable function $f: K^n \rightarrow K$, there exists a \emptyset -definable set $C \subset K^n$ of dimension at most $n - 1$ such that for any box $B \subset K^n \setminus C$, f is $(r + 1)$ -*

fold differentiable on B , for each $i \in \mathbb{N}^n$ with $|i| = r + 1$ one has that $|f^{(i)}|$ is constant on B , and we have

$$|f(x) - T_{f,x'}^{\leq r}(x)| \leq \max_{|i|=r+1} |f^{(i)}(x')(x - x')^i| \tag{5.6.1}$$

for every $x, x' \in B$. (Here, i runs over n -tuples and we use the usual multi-index notation.)

One may search for a variant of equation (5.6.1) that not only holds on boxes disjoint from C but even holds globally on each open fiber of a map $K^n \rightarrow \text{RV}^k$, as in Theorem 5.4.10; however, we do not yet know how to formulate such a variant in general.

Proof of Theorem 5.6.1. We do an induction over n , the case $n = 1$ being Theorem 3.2.2. Applying the induction hypothesis fiberwise (and using compactness) allows us to find a $C \subset K^n$ (\emptyset -definable, of dimension less than n) such that for every box $B = \hat{B} \times B_n \subset K^{n-1} \times K$ disjoint from C , for $x_n \in B_n$ and for every $\hat{x}, \hat{x}' \in \hat{B}$, we have

$$|f(x) - \sum_{|\hat{i}| \leq r} \frac{f^{(\hat{i},0)}(\hat{x}', x_n)}{\hat{i}!} (\hat{x} - \hat{x}')^{\hat{i}}| \leq \max_{|\hat{i}|=r+1} |f^{(\hat{i},0)}(\hat{x}', x_n) (\hat{x} - \hat{x}')^{\hat{i}}|. \tag{5.6.2}$$

For each $\hat{i} \in \mathbb{N}^{n-1}$ with $|\hat{i}| \leq r$, we moreover apply the $n = 1$ case of the theorem to $f^{(\hat{i},0)}(\hat{x}, \cdot)$ (for each fixed \hat{x}) and $r' := r - |\hat{i}|$, so that for $\hat{x}' \in \hat{B}$ and $x_n, x'_n \in B'$, we have:

$$|f^{(\hat{i},0)}(\hat{x}', x_n) - \sum_{i_n=0}^{r'} \frac{f^{(\hat{i},i_n)}(\hat{x}', x'_n)}{i_n!} (x_n - x'_n)^{i_n}| \leq |f^{(\hat{i},r'+1)}(\hat{x}', x'_n) (x_n - x'_n)^{r'+1}|. \tag{5.6.3}$$

Using equation (5.6.3) to estimate the $f^{(\hat{i},0)}(\hat{x}', x_n)$ from the left-hand side of equation (5.6.2) yields equation (5.6.1), as desired, where we also use that we may assume that $|f^{(\hat{i},0)}(\hat{x}', \cdot)|$ is constant for each \hat{x}' and each \hat{i} so that in the right-hand side of equation (5.6.2), we can replace $|f^{(\hat{i},0)}(\hat{x}', x_n)|$ by $|f^{(\hat{i},0)}(\hat{x}', x'_n)|$. □

5.7. Classical cells

We end this section by recalling the older, more classical notion of cells, which has to be used in the absence of the condition that acl equals dcl in $\text{Th}(K)$, and by stating the corresponding classical cell decomposition results under the assumption of 1-h-minimality.

Definition 5.7.1 (Reparameterized cells). Consider integers $n, k \geq 0$, a \emptyset -definable set $X \subset K^n$ and a \emptyset -definable function

$$\sigma : X \rightarrow \text{RV}^k.$$

Then (X, σ) is called a \emptyset -definable reparameterized cell (reparameterized by σ) if, for each $\xi \in \text{RV}^k$, the set $\sigma^{-1}(\xi)$, when non-empty, is a ξ -definable cell with some center tuple c_ξ (see Definition 5.2.2), such that moreover c_ξ depends definably on ξ and such that the cell-type of $\sigma^{-1}(\xi)$ is independent of ξ . If (X, σ) is such a reparameterized cell, then by ‘a twisted box of X ’, we mean a twisted box of $\sigma^{-1}(\xi)$ for some ξ as in Definition 5.2.2, and similarly, by the center tuple and the cell-type of (X, σ) , we mean the definable family of the center tuples of the cells $\sigma^{-1}(\xi)$ (with family parameter ξ), and the cell-type of the $\sigma^{-1}(\xi)$, respectively.

Remark 5.7.2. In the above definition, one can always modify σ in such a way that afterward, each $\sigma^{-1}(\xi)$ is either empty or a single twisted box. In a different direction, if the language \mathcal{L} has an angular component map $\overline{\text{ac}}$ sending K to its residue field, then one can take σ from Definition 5.7.1 to be residue

field valued (instead of RV-valued). Either of those additional assumptions on σ can in particular be imposed on the cells appearing in the following theorem.

Theorem 5.7.3 (Reparameterized cell decomposition). *Suppose that $\text{Th}(K)$ is 1-h-minimal. Consider $n, k \geq 0$ and \emptyset -definable sets $X \subset K^n$ and $P \subset X \times \text{RV}^k$. We consider P the function sending $x \in X$ to the fiber $P_x := \{\xi \in \text{RV}^k \mid (x, \xi) \in P\}$. Then there exists a finite decomposition of X into \emptyset -definable reparameterized cells (A_i, σ_i) such that moreover P (considered a function) is constant on each twisted box of each A_i .*

The other addenda of Theorem 5.2.4 can be adapted in a similar way. In particular, for the analogue of Addendum 3 of 5.2.4, given finitely many \emptyset -definable $f_j : X \rightarrow K$ and assuming $n = 1$, we can moreover assume that there are \emptyset -definable reparameterized cells (B_{ij}, τ_{ij}) such that $f_j(A_i) = B_{ij}$ and such that any twisted box of A_i is mapped by f_j onto a twisted box of B_{ij} .

For the analogue of Addendum 5 of 5.2.4, up to allowing a well-chosen coordinate permutation for each for each i , we can moreover for each ξ obtain that the center of $\sigma_i^{-1}(\xi)$ is 1-Lipschitz.

Proof. The proof is analogous to the proof of Theorem 5.2.4 and its addenda. □

Note that the above version of Addendum 5 is weaker than the original Addendum 5 of Theorem 5.2.4, since instead of obtaining finitely many 1-Lipschitz centers, we now only obtain that for each ξ separately, the center of $\sigma_i^{-1}(\xi)$ is 1-Lipschitz; this corresponds to an infinite partition.

A similar phenomenon also arises with Theorem 5.2.8 in the absence of the condition that $\text{acl} = \text{dcl}$, as in Theorem 2.1.7 of [13]: f being locally 1-Lipschitz implies globally 1-Lipschitz only after some infinite partition of the domain.

5.8. Motivic integration

In this subsection, we show that the Cluckers–Loeser style of motivic integration from [21] can be generalized to the 1-h-minimal setting under some natural extra assumptions about the induced structure on VG and RF. It is beyond the scope of this paper to recall what motivic integration is. Instead, we rely on [23], which generalizes an essential part of [21] to more characteristics. In [23], some axiomatic assumptions are given under which motivic integration works; below, we just verify those assumptions. More precisely, we check the conditions of [23, Definition 3.9] (in the variant from Remark 5.8.1 below) for the theory $\text{Th}_{\mathcal{L}}(K)$ where K is a Henselian, discretely valued field of equi-characteristic zero with an \mathcal{L} -structure that is 1-h-minimal and satisfies some natural conditions; motivic integration is developed uniformly for such K in [23] (as well as in mixed characteristic assuming bounds on the ramification). Note that the framework with motivic additive characters and motivic Fourier transforms from [22] can be generalized similarly to the 1-h-minimal setting under the same extra assumptions; as this would take more time to explain, we leave such adaptation to some future work.

Remark 5.8.1. We indicate a small adaptation of the framework of [23] in the equi-characteristic zero case. If one only considers equi-characteristic zero Henselian valued fields, then the language $\mathcal{L}_{\text{high}}$ used in [23] can be replaced by the Denef–Pas language (i.e., one omits the higher-order angular component maps $\overline{\text{ac}}_n$ for $n > 1$ from $\mathcal{L}_{\text{high}}$, one omits the symbol π for the uniformizer and one keeps only the sorts VF, RF and VG). One can adapt all definitions and results from [23] correspondingly, as long as one only considers equi-characteristic zero Henselian valued fields. Let us call this variant (with the Denef–Pas language) the non-high variant of [23].

Theorem 5.8.2. *Let K be an equi-characteristic zero, Henselian, discretely valued field, equipped with \mathcal{L} -structure for some language $\mathcal{L} \supset \mathcal{L}_{\text{val}}$. Suppose that $\text{Th}(K)$ is 1-h-minimal and that we additionally have the following:*

1. The language \mathcal{L} contains an angular component map $\overline{ac}: K \rightarrow k := \mathcal{O}_K/\mathcal{M}_K$.
2. The induced structure on Γ_K^\times is that of the pure ordered abelian group: that is, every definable subset of $(\Gamma_K^\times)^\ell$ (for every ℓ) is definable in the language $\{+, <\}$ on Γ_K^\times .
3. The value group Γ_K^\times and the residue field k are orthogonal: that is, every definable subset of $k^m \times (\Gamma_K^\times)^\ell$ (for every m and ℓ) is a finite union of sets of the form $Y \times Z$, with $Y \subset k^m$ and $Z \subset (\Gamma_K^\times)^\ell$.

Then the \mathcal{L} -theory $\mathcal{T} := \text{Th}(K)$ satisfies the assumptions from the non-high variant of [23, Definition 3.9] (see Remark 5.8.1).

Remark 5.8.3. In Items (1) and (3) of Theorem 5.8.2, the residue field k is naturally considered a subset of RV_K , consisting of the elements of norm 1 inside RV_K , and the zero element.

Proof of Theorem 5.8.2. According to [23, Definition 3.9] and with the terminology of [23, Section 3.2], we have to verify that \mathcal{T} is split, is finitely b -minimal, has the Jacobian Property, and has at least one model that is a $(0, 0, 0)$ -field, all in their non-high variants from Remark 5.8.1 above.

By Proposition 5.3.3, \mathcal{T} is b -minimal, and the additional condition ‘finitely’ follows from Lemma 2.8.1. By Corollary 3.2.7 (with $\lambda = 1$), \mathcal{T} has the Jacobian Property, and the given field K is a model that is a $(0, 0, 0)$ -field.

That \mathcal{T} is split ([23, Definition 3.6]) means that the induced structure on Γ_K^\times is the pure ordered abelian group (which is one of our assumptions) and that k and Γ_K^\times are ‘orthogonal with control of parameters’ in the following sense: For every set $A \subset K \cup \text{RV}^{\text{eq}}$ of parameters, every A -definable subset of $k^m \times (\Gamma_K^\times)^\ell$ is a finite union of sets of the form $Y \times Z$, where $Y \subset k^m$ and $Z \subset (\Gamma_K^\times)^\ell$ are A -definable.

It turns out that using the order on Γ_K^\times , this strong form of orthogonality follows from the normal one, namely as follows: Let an A -definable set $X \subset k^m \times (\Gamma_K^\times)^\ell$ be given, and let $(\Gamma_K^\times)^\ell = \bigcup_i Z_i$ be the partition defined by the condition that $z, z' \in (\Gamma_K^\times)^\ell$ lie in the same set Z_i if and only if the fibers X_z and $X_{z'}$ (in k^m) are equal. By (normal) orthogonality of k and Γ_K^\times , this partition is finite. Using the order on Γ_K^\times , we can definably distinguish between the sets Z_i , so each Z_i is A -definable. Denote by $Y_i \subset k^m$ the fiber corresponding to Z_i (so that $X_z = Y_i$ for every $z \in Z_i$). Those are also A -definable, and as desired, X is the union of the sets $Y_i \times Z_i$. □

6. Examples of Hensel minimal structures

In this section, we provide many examples of Hensel minimal valued fields of equi-characteristic 0, with various kinds of languages. Note that the structures from [64, 65] for which a proper subfield of K is definable (e.g., as the fixed field of an automorphism) are not covered by our framework. In some cases, we prove ω -h-minimality; in others, we only obtain 1-h-minimality (namely for power-bounded \mathcal{T} -convex structures). We also provide examples of valued fields of mixed characteristics.

So far, we have introduced Hensel minimality only in equi-characteristic 0. A mixed characteristic analogue is developed in the sequel [17]. For the sake of the examples below, we quickly give a definition of mixed characteristic Hensel minimality here.

6.1. Mixed characteristic

As usual, we fix a theory \mathcal{T} of valued fields in a language \mathcal{L} containing the language \mathcal{L}_{val} of valued fields. We do require the characteristic of models to be 0, but we allow the models to have arbitrary residue field characteristic.

Notation 6.1.1 (Equi-characteristic 0 coarsening). Given a model $K \models \mathcal{T}$, we write $\mathcal{O}_{K,\text{ecc}}$ for the smallest subring of K containing \mathcal{O}_K and \mathbb{Q} , and we let $|\cdot|_{\text{ecc}}: K \rightarrow \Gamma_{K,\text{ecc}}$ be the corresponding valuation. (Thus, $|\cdot|_{\text{ecc}}$ is the finest coarsening of $|\cdot|$ that has equi-characteristic 0; note that $|\cdot|_{\text{ecc}}$ can be a trivial valuation on K .) If $|\cdot|_{\text{ecc}}$ is a non-trivial valuation (i.e., $\mathcal{O}_{K,\text{ecc}} \neq K$), then we also use the following notation: $\text{rv}_{\text{ecc}}: K \rightarrow \text{RV}_{\text{ecc}}$ is the leading term map with respect to $|\cdot|_{\text{ecc}}$; given $\lambda \in \Gamma_{K,\text{ecc}}$,

$\text{rv}_\lambda: K \rightarrow \text{RV}_\lambda$ is the leading term map with respect to λ ; and \mathcal{L}_{ecc} is the extension of \mathcal{L} by a predicate symbol for $\mathcal{O}_{K,\text{ecc}}$.

Definition 6.1.2 (Hensel minimality in mixed characteristic). Let \mathcal{T} be a theory of valued fields of characteristic 0 (and arbitrary residue field characteristic) in a language \mathcal{L} containing the language \mathcal{L}_{val} of valued fields. Let $\ell \geq 0$ be either an integer or ω . We say that \mathcal{T} is ℓ -h^{ecc}-minimal if for every model $K \models \mathcal{T}$ the following holds: If the valuation $|\cdot|_{\text{ecc}}$ on K is non-trivial, then the \mathcal{L}_{ecc} -theory of K , when considered a valued field with the valuation $|\cdot|_{\text{ecc}}$, is ℓ -h-minimal in the sense of Definition 2.3.3.

Remark 6.1.3. For theories of valued fields of equi-characteristic 0, ℓ -h^{ecc}-minimality is trivially equivalent to ℓ -h-minimality, for each ℓ (since $\mathcal{O}_{K,\text{ecc}} = \mathcal{O}_K$).

A natural adaptation of the definition of h-minimality to mixed characteristics would consist of (suitably) replacing λ -preparation by $|m|\lambda$ -preparation for some nonzero integer m . In [17], we show in the case of 1-h-minimality that this adaptation is equivalent to the above notion of 1-h^{ecc}-minimality. Moreover, we prove that many preparation results from this paper can be adapted to mixed characteristics in a similar manner, replacing λ -preparation by $|m|\lambda$ -preparation.

6.2. Valued fields with or without analytic structure

In this subsection, we prove ω -h-minimality of arbitrary Henselian valued fields K of equi-characteristic 0 with analytic structure in the sense of [18]. If K has mixed characteristic and analytic structure, we will show that it is ω -h^{ecc}-minimal, which means that its equi-characteristic 0 coarsenings are ω -h-minimal. As is so often the case, obtaining results in the positive equi-characteristic case seems completely out of reach at present.

The pure valued field language is a special case of an analytic structure on K . Nevertheless, we will treat this case separately in this section, avoiding the machinery of analytic structures and instead building only on classical quantifier elimination results. Note that before, the only known proofs of the Jacobian Property (Corollary 3.2.7) either went via analytic structures (as in [21, 18]) or were restricted to algebraically closed valued fields (as in [47]).

The proofs of ω -h-minimality (i) in the pure field language and (ii) in fields with analytic structure are very similar, the main differences being that, in case (ii), we require additional input from [18]. We therefore formulate both proofs simultaneously, tagging differences with (i) and (ii).

We fix the following language and structure for the remainder of this subsection. Note that in the mixed characteristic case, the Definition 6.1.2 of ω -h^{ecc}-minimality requires us to work with two different valuations simultaneously (although in Case (i), just understanding the coarser valuation is enough for the proof of ω -h-minimality).

- (i) Let $\mathcal{L} := \mathcal{L}_{\text{val}} \cup \{0, 1, -\} = \{+, -, 0, 1, \cdot, \mathcal{O}_K\}$ be the pure valued field language together with $0, 1, -$ with their natural meaning, and let K be a Henselian valued field of equi-characteristic 0, considered an \mathcal{L} -structure.
- (ii) Let $\mathcal{A} = (A_{m,n})_{m,n}$ be a separated Weierstrass system in the sense of [18]. Let K be a characteristic zero valued field (possibly of positive residue field characteristic) with a separated analytic \mathcal{A} -structure, as in [18, Definitions 4.1.5 and 4.1.6]. Note that by [18, Proposition 4.5.10 (i)], any such K is Henselian.

We denote the valuation ring of K by $\mathcal{O}_{K,\text{fine}}$, and we fix an equi-characteristic 0 coarsening $\mathcal{O}_K \supset \mathcal{O}_{K,\text{fine}}$. (If $\mathcal{O}_{K,\text{fine}}$ itself is of equi-characteristic 0, one can as well choose $\mathcal{O}_K = \mathcal{O}_{K,\text{fine}}$.) We write $\mathcal{M}_{K,\text{fine}}, \mathcal{M}_K$ for the corresponding maximal ideals and $|\cdot|_{\text{fine}}: K \rightarrow \Gamma_{K,\text{fine}}, |\cdot|: K \rightarrow \Gamma_K$ for the corresponding valuations.

We let, still in case (ii), \mathcal{L} be the extension of the language from (i) by a function symbol for field division (extended by zero on zero), by a predicate symbol for $\mathcal{O}_{K,\text{fine}}$ and by one function symbol for each f in $A_{m,n}$, interpreted as a function $\mathcal{O}_{K,\text{fine}}^m \times \mathcal{M}_{K,\text{fine}}^n \rightarrow K$ via the analytic \mathcal{A} -structure on K (and extended by 0 outside of its domain).

Theorem 6.2.1 (Fields with analytic structure). *Let $(K, |\cdot|)$ be an equi-characteristic 0 valued field in a language \mathcal{L} as above in (i) or (ii). (Recall that in particular, in Case (ii) there is a valuation ring $\mathcal{O}_{K, \text{fine}}$ of K that may differ from \mathcal{O}_K .) Then the \mathcal{L} -theory of K with the valuation $|\cdot|$ is ω -h-minimal.*

In one word, the idea of the proof of Theorem 6.2.1 goes via quantifier elimination of valued field quantifiers, in a language to which the sort RV_λ and the map rv_λ have been added for some given $\lambda \leq 1$ in Γ_K^\times . The conditions of ω -h-minimality are then easily checked. Similar quantifier elimination results have been proved but not yet with such an RV_λ , so we take care to give sufficient details.

The first technical ingredient, inspired by [34], is the following:

Lemma 6.2.2. *Let $F \leq K$ be a subfield that is moreover an \mathcal{L} -substructure of K , let $\lambda \leq 1$ be an element of Γ_K^\times , and let $\tau(x)$ be an $\mathcal{L}(F)$ -term, with x a single variable. Then there exists a finite set $C \in K$ consisting only of algebraic elements over F such that $\text{rv}_\lambda(\tau(x))$ only depends on $(\text{rv}_\lambda(x - c))_{c \in C}$.*

Moreover, if we restrict (in our desired property for C) to the x that are solutions of a given degree d polynomial equation $P \in F[X]$, (i.e., if we want to obtain the implication $(\text{rv}_\lambda(x - c))_{c \in C} = (\text{rv}_\lambda(x' - c))_{c \in C} \Rightarrow \text{rv}_\lambda(\tau(x)) = \text{rv}_\lambda(\tau(x'))$ only under the assumption $P(x) = P(x') = 0$), then the elements of C can all be assumed to have degree strictly less than d over F .

Proof in Case (i). Note that our $\mathcal{L}(F)$ -term τ is simply a polynomial in $F[x]$, so the first part of the lemma is immediate from [38, Proposition 3.6] and its proof, namely using for C the set of all roots of all derivatives of τ (including the roots of τ itself). Note that the proof of [38, Proposition 3.6] yields that the ‘Swiss cheeses’ U_i appearing in the statement of the proposition are 1-prepared by C .

For the second part, choose a polynomial Q of degree less than $\text{deg}(P)$ such that the polynomial τ is congruent to Q modulo P . Then for every $x \in K$ that is a zero of P , we have $\tau(x) = Q(x)$, so we may as well replace τ by Q . Then choosing C as in the proof of the first part yields the claim. \square

Proof in Case (ii). We apply [18] to K with the valuation $|\cdot|_{\text{fine}}$. The main idea is to use [18, Theorem 5.5.3] to reduce to Case (i). As a preparation, note that it is sufficient to obtain the conclusion of the lemma for $x \in \mathcal{O}_K$. Indeed, the x outside of \mathcal{O}_K can be treated by applying the lemma separately to the $\mathcal{L}(F)$ -term $\tau(1/x)$.

By [18, Theorem 5.5.3 and Remark 5.5.4], there is a cover of the valuation ring $\mathcal{O}_{K_{\text{alg}}}$ of the algebraic closure of K by finitely many F -annuli \mathcal{U}_i (cf. [18, Definition 5.1.1]), and there is a finite F -definable set $S \subset K$ such that, on each $\mathcal{U}_i \setminus S$, we have

$$\tau(x) = G_i(x)/H_i(x) \cdot E_i(x),$$

where $G_i, H_i \in F[x]$ are polynomials and E_i is an $\mathcal{L}(F)$ -term that is a strong unit on \mathcal{U}_i (cf. [18, Definition 5.1.4]). Indeed, that these data are defined over F follows from [18, Remark 5.5.4] with $K' = F$.

Let $P_S \in F[x]$ be the polynomial whose set of roots is S , and let $P_j \in F[x]$ be the collection of polynomials appearing in the definition of the F -annuli \mathcal{U}_i . (In particular, each \mathcal{U}_i is defined by a boolean combination of inequalities between the valuations of the P_j .) Since $|\cdot|_{\text{fine}}$ factors over rv_λ , whether an element $x \in Z$ lies in \mathcal{U}_i is determined by $(\text{rv}_\lambda(P_j(x)))_j$. Moreover, by [18, Lemma 6.3.12] and [19, Remark A.1.12], $\text{rv}_\lambda(E_i)$ only depends on $(\text{rv}_\lambda(P_j(x)))_j$. Thus to prove the lemma for τ , it suffices to prove it for the polynomials G_i, H_i, P_j and P_S (considered functions on K). But this has already been done in Case (i). \square

The second ingredient is a quantifier elimination result. We fix $\lambda \leq 1$ in Γ_K^\times and consider the following expansion of K in a language $\mathcal{L}_{\text{qe}, \lambda} \supset \mathcal{L}$: We add RV_λ as a new sort, together with the map rv_λ and with the \emptyset -induced structure on RV_λ : that is, one predicate for each \emptyset -definable subset of RV_λ^n , for every n .

Proposition 6.2.3. *The $\mathcal{L}_{\text{qe}, \lambda}$ -theory of K eliminates field quantifiers.*

The particularity of Proposition 6.2.3 is not only that it has rv_λ and RV_λ on top of the analytic structure, but also that there are two valuation rings at play, namely $\mathcal{O}_{K, \text{fine}}$ (which may be of mixed characteristic) and \mathcal{O}_K (which is of equi-characteristic zero).

Proof of Proposition 6.2.3. First note that in the case $\lambda = 1$, this result is known: in Case (i), this is [38, Proposition 4.3] (or also [4, Theorem B]), and in Case (ii), it is [64, Theorem 3.10]. This already implies a partial result for arbitrary $\lambda \leq 1$, namely:

(6.2.1) Every $\mathcal{L}_{\text{qe},\lambda}$ -formula $\phi(x)$ having only K -variables is equivalent to a field quantifier free formula.

Indeed, $\phi(x)$ is equivalent to an $\mathcal{L}_{\text{qe},1}(\lambda)$ -formula $\psi(x)$ without K -quantifiers, and each RV-quantifier of $\psi(x)$ can easily be replaced by some RV_λ -quantifiers.

To prove the general case, we need the following variants of results from [38].

Fix a non-zero polynomial $P \in K[x]$ and an element $a_0 \in K$. Let $c_i \in K$ be the coefficients of P when developed around a_0 : that is, $P(x) = \sum_i c_i(x - a_0)^i$.

Given $b \in K$, we say (cf. [38, Definition 3.1]) that P has a λ -collision at b around a_0 if $|P(b)| < \lambda \max_i |c_i(b - a_0)^i|$. Note that whether P has a λ -collision at b around a_0 only depends on $\text{rv}_\lambda(c_i)$ (for all i) and $\text{rv}_\lambda(b - a_0)$. (This will be useful later.)

Claim 6.2.4. *Suppose that the above $P \in K[x]$ has a λ -collision at b around a_0 . Then there exists an integer $n \geq 0$ with $n < \deg P$ and an element $b' \in K$ ‘close to b ’ such that $P^{(n)}(b') = 0$. Here, ‘close to b ’ means $\text{rv}_\lambda(b' - a_0) = \text{rv}_\lambda(b - a_0)$ if $n = 0$ and $\text{rv}(b' - a_0) = \text{rv}(b - a_0)$ if $n \geq 1$, and $P^{(n)}$ stands for the n th derivative, with $P^{(0)} = P$.*

Proof. Without loss, we may assume that $a_0 = 0$, $b = 1$ and $\max_i |c_i| = 1$. Let $Q = \text{res}(P)$, the reduction of P modulo \mathcal{M}_K . We have $Q(1) = 0$. Let n be such that 1 is a root of multiplicity 1 of $Q^{(n)}$. Then Hensel’s Lemma yields a root b' of $P^{(n)}$ such that $|b' - 1| \leq |P^{(n)}(1)|$. Since $|P^{(n)}(1)| \leq 1$, this implies $\text{rv}(b') = \text{rv}(b)$. In the case $n = 0$, that 1 is not a root of Q' implies $|P'(1)| = 1$. Together with $|P(1)| \leq \lambda$, Hensel’s Lemma implies $|b' - b| < \lambda$ and hence $\text{rv}_\lambda(b') = \text{rv}_\lambda(b)$. □

Claim 6.2.5. *Suppose that the above $P \in K[x]$ has no common zero with any of its proper derivatives and fix $\xi \in \text{RV}_\lambda^\times$. The following are equivalent:*

1. *there exists a root $b \in K$ of P with $\text{rv}_\lambda(b - a_0) = \xi$;*
2. *there exists $b \in K$ such that*
 - (2a) *$\text{rv}_\lambda(b - a_0) = \xi$, P has a λ -collision at b around a_0 , and*
 - (2b) *for every root $a \in K$ of every proper derivative of P , P has a λ -collision at b around a .*

Proof. If b is a root of P , then P has a λ -collision at b around any $a \neq b$. So (1) implies (2). Let us now assume that we have b such that (2) holds. If (1) does not hold, then by Claim 6.2.4 (and (2a)), there exists a root a of some proper derivative of P with $\text{rv}(a - a_0) = \text{rv}(b - a_0)$. Pick the closest such a to b . By Claim 6.2.4, around a this time (and using (2b)), there exists a root c of some $P^{(m)}$ with $\text{rv}(c - a) = \text{rv}(b - a)$; in particular, $|b - c| < |b - a|$, a contradiction to our choice of a . □

We now come back to the actual proof of field quantifier elimination. We abbreviate ‘field quantifier free’ by ‘fqf’. It suffices to prove the following: Suppose that $K' \equiv K$ is $|K|^{+-}$ -saturated, that $A \subset K \cup \text{RV}_{K,\lambda}$ and $A' \subset K' \cup \text{RV}_{K',\lambda}$ are substructures and that $\alpha: A \rightarrow A'$ is an fqf-elementary bijection: that is, it preserves the validity of fqf formulas. Then for any $a \in K$, there exists an a' such that α extends to an fqf-elementary map sending a to a' .

For α to be fqf-elementary, it suffices that it is an isomorphism of substructures and that $\alpha|_{A \cap \text{RV}_{K,\lambda}}$ is fqf-elementary. Indeed, suppose that ϕ is an fqf $\mathcal{L}_{\text{qe},\lambda}(A)$ -sentence that holds in K . Then without loss, $\phi = \psi((\text{rv}_\lambda(\tau_i))_i)$ for some $\mathcal{L}_{\text{qe},\lambda}(A \cap K)$ -terms τ_i and for $\psi(y)$ an fqf $\mathcal{L}_{\text{qe},\lambda}(A \cap \text{RV}_{K,\lambda})$ -formula. Let $\xi_i \in A \cap \text{RV}_{K,\lambda}$ be the interpretation of $\text{rv}_\lambda(\tau_i)$. Then ϕ follows from $\bigwedge_i \tau_i = \xi_i$ (which is quantifier free) and $\psi((\xi_i)_i)$ (which is an fqf $\mathcal{L}_{\text{qe},\lambda}(A \cap \text{RV}_{K,\lambda})$ -sentence).

We may assume $\text{RV}_{K,\lambda} \subset A$, since $\alpha|_{A \cap \text{RV}_{K,\lambda}}$ extends to an fqf-elementary map on $\text{RV}_{K,\lambda}$ and the union of this extension with the original α is an isomorphism of substructures. In particular, when further extending α , we now only need to make sure that it remains an isomorphism of substructures. In terms of formulas, this means (by a usual compactness argument) that given a quantifier free $\mathcal{L}_{\text{qe},\lambda}(A)$ -

formula $\phi(x)$ with x a valued field variable, we need to check that $K \models \exists x \phi(x)$ implies $K' \models \exists x \phi^\alpha(x)$ (where ‘ ϕ^α ’ is the $\mathcal{L}_{\text{qe},\lambda}(A)$ -formula obtained from ϕ by applying α to the parameters from A).

We may assume that $F := K \cap A$ is a subfield. In Case (ii), this is automatic, since $\mathcal{L}_{\text{qe},\lambda}$ contains field division. In Case (i), the ring homomorphism $\alpha_{K \cap A}$ uniquely extends to the fraction field F of $K \cap A$, and extending α in this way yields an isomorphism of substructures, since for $\frac{b}{b'} \in F$ ($b, b' \in K \cap A$), we have $\text{rv}_\lambda(\frac{b}{b'}) = \frac{\text{rv}_\lambda(b)}{\text{rv}_\lambda(b')}$.

From now on, we identify A with its image $\alpha(A)$. Fix $a \in K$, and fix a quantifier free $\mathcal{L}_{\text{qe},\lambda}(A)$ -formula $\phi(x)$ such that $K \models \phi(a)$ holds. We need to show that $K' \models \exists x: \phi(x)$. To do so, we will successively reduce to simpler formulas until we can get rid of the K -quantifier $\exists x$.

Let $P \in F[x]$ be the minimal polynomial of a over F , and set $d := \deg P$. If a is transcendental over F , we set $P := 0$ and $d := \infty$. By induction on d , we may assume:

(6.2.2) F contains all roots b in K of polynomials over F of degree strictly less than d .

As before, we can assume that the above formula ϕ is of the form $\phi(x) = \psi((\text{rv}_\lambda(\tau_i(x))))_i$ for some $\mathcal{L}(F)$ -terms τ_i . By Lemma 6.2.2 (and (6.2.2)), $\phi(x)$ is equivalent, in the structure K , to a formula of the form

$$\phi'(x) = \psi'((\text{rv}_\lambda(x - c_j))_j) \wedge P(x) = 0$$

for some $c_j \in F$, where $\psi'(y) = \psi((\eta_i(y))_i)$ for suitable $\mathcal{L}_{\text{qe},\lambda}(F)$ -definable functions η_i . We claim that this equivalence also holds in K' so that we can without loss replace ϕ by ϕ' .

To prove the claim, note that the equivalence $\phi \leftrightarrow \phi'$ follows from an $\mathcal{L}_{\text{qe},\lambda}(F)$ -sentence χ , namely $\chi = \bigwedge_i \forall x \in K: \tau_i(x) = \eta_i((\text{rv}_\lambda(x - c_j))_j)$. Since χ only uses K -parameters, we already know (by (6.2.1)) that it is equivalent to a fqf $\mathcal{L}_{\text{qe},\lambda}(F)$ -sentence χ' (modulo only the $\mathcal{L}_{\text{qe},\lambda}$ -theory of K : that is, without using the specific embedding of F into K). Thus the truth of χ' is preserved by α so that we obtain the desired equivalence in K' .⁴

Next, note that we can get rid of all the c_j appearing in ϕ' except for the one closest to a : that is, denoting that closest c_j by c , we can replace ϕ' by

$$\phi''(x) = \psi''(\text{rv}_\lambda(x - c)) \wedge P(x) = 0.$$

Indeed, one can easily choose ψ'' in such a way that $K \models \phi''(a)$ and $\text{Th}_{\mathcal{L}_{\text{qe},\lambda}}(K)$ implies $\phi'' \rightarrow \phi'$.

Now $\exists x: \phi''(x)$ is equivalent to $\exists \xi \in \text{RV}_\lambda: \psi''(\xi) \wedge (\exists x: \text{rv}_\lambda(x - c) = \xi \wedge P(x) = 0)$, and it remains to get rid of the $\exists x$ in that formula. If P is the zero polynomial, then the $\exists x$ part is trivially true, and we are done. Otherwise, by Claim 6.2.5, the existence of such an x is equivalent to the existence of an x with $\text{rv}_\lambda(x - c) = \xi$ such that P has a λ -collision at x around certain points b_j from F (namely around c and around the roots of the derivatives of P , which are in F by (6.2.2)). For fixed P and b_i , the existence of such a collision is determined by $\text{rv}_\lambda(x - b_j)$, so it remains to eliminate the $\exists x$ from an $\mathcal{L}_{\text{qe},\lambda}(B)$ -formula of the form $\exists x: \psi'''((\text{rv}_\lambda(x - b_j))_j)$ (with ψ''' fqf expressing that the collisions exist). This can then be further simplified to a formula of the form $\exists x: \bigwedge_j \text{rv}_\lambda(x - b_j) = \xi_j$ (where we take $\xi_j := \text{rv}_\lambda(a - b_j)$). This formula now expresses that the intersection of certain balls is non-empty, a condition that can easily be seen to only depend on $\text{rv}_\lambda(b_j - b_{j'})$ (see [38, Proposition 4.1] for details). Thus we are done. \square

Proof of Theorem 6.2.1. We need to show that for every $K' \equiv K$ and every $A \subset K'$, every $(A \cup \text{RV}_{K',\lambda})$ -definable set $X = \phi(K') \subset K'$ can be λ -prepared by a finite A -definable set C .

By Proposition 6.2.3, we may assume that ϕ contains no field quantifiers so that it suffices to λ -prepare the graph of functions of the form $\text{rv}_\lambda(\tau(x))$, where τ is an $\mathcal{L}_{\text{qe},\lambda}(A)$ -term. (Here, ‘preparing a graph’ is in the sense of Definition 2.6.1.)

⁴In Case (i), we do not need to invoke another field quantifier elimination result. As in Lemma 6.2.2, we assume that each τ_i has degree smaller than P ; and we choose, as c_j , all the roots of derivatives of all τ_i , including τ_i itself. Then, by [38, Proposition 3.6], for every x and i , there exists a j such that if we write $\tau_i(x) = \sum_k a_k(x - c_j)^k$, the sum $\sum_k \text{rv}_\lambda(a_k) \text{rv}_\lambda(x - c_j)^k$ is well-defined and hence equal to $\text{rv}_\lambda(\tau_i(x))$. Using that the well-definedness of the sum is an fqf condition, we can define $\eta_i: (\text{rv}_\lambda(x - c_j))_j \mapsto \text{rv}_\lambda(\tau_i(x))$ without field quantifiers so that the equality $\eta_i(\text{rv}_\lambda(x - c_j))_j = \text{rv}_\lambda(\tau_i(x))$ is preserved by f .

By the first half of Lemma 6.2.2, such a graph can be prepared by a finite set C of elements that are algebraic over the field $F \leq K'$ generated by A . Since such a set C is contained in a finite A -definable set C' , we are done. \square

Corollary 6.2.6 (Equi-characteristic 0 examples). *Let K be a Henselian valued field of equi-characteristic 0 in a language \mathcal{L} containing the pure valued field language $\mathcal{L}_{\text{val}} = \{+, \cdot, \mathcal{O}_K\}$. Then in each of the following cases, $\text{Th}_{\mathcal{L}}(K)$ is ω -h-minimal.*

1. \mathcal{L} is equal to \mathcal{L}_{val} .
2. The \mathcal{L} -structure K is an expansion of the \mathcal{L}_{val} -structure by an analytic \mathcal{A} -structure in the sense of [18], for some separated Weierstrass system \mathcal{A} .

Proof. These are just examples of Theorem 6.2.1, namely with $\mathcal{O}_K = \mathcal{O}_{K, \text{fine}}$. \square

Corollary 6.2.7 (Mixed characteristic examples). *Let K be a Henselian valued field of mixed characteristic in a language \mathcal{L} . Then in each of the following cases, $\text{Th}(K)$ is ω -h^{ecc}-minimal (as in Definition 6.1.2).*

1. \mathcal{L} is the pure valued field language $\mathcal{L}_{\text{val}} = \{+, \cdot, \mathcal{O}_K\}$.
2. K is a finite field extension of \mathbb{Q}_p and \mathcal{L} is the sub-analytic language from [34] (which is a variant on the language from [27]).
3. The \mathcal{L} -structure K is an expansion of the \mathcal{L}_{val} -structure by an analytic \mathcal{A} -structure in the sense of [18], for some separated Weierstrass system \mathcal{A} .

Proof. Given $K' \equiv_{\mathcal{L}} K$, let $|\cdot|_{\text{ecc}}$ be the finest equi-characteristic 0 coarsening of the valuation $|\cdot|$, and let \mathcal{L}_{ecc} be the extension of \mathcal{L} by a predicate symbol for the valuation ring $\mathcal{O}_{K, \text{ecc}}$ corresponding to $|\cdot|_{\text{ecc}}$. Suppose that $|\cdot|_{\text{ecc}}$ is non-trivial. Under those assumptions, we need to show that $\text{Th}_{\mathcal{L}_{\text{ecc}}}(K'_{\text{ecc}})$ is ω -h-minimal, where K'_{ecc} is the field K' considered a valued field with the valuation $|\cdot|_{\text{ecc}}$.

(1) If \mathcal{L} is the pure valued field language, we consider \mathcal{L}_{ecc} an extension of $\mathcal{L}_{\text{val, ecc}} := \{+, \cdot, \mathcal{O}_{K', \text{ecc}}\}$ (which is also the pure valued field language) by a predicate symbol for $\mathcal{O}_{K'}$. By Theorem 6.2.1, $\text{Th}_{\mathcal{L}_{\text{val, ecc}}}(K')$ is ω -h-minimal. Since the map $K' \rightarrow \Gamma_{K'}$ factors over RV_{ecc} (where RV_{ecc} denotes the leading term structure with respect to $|\cdot|_{\text{ecc}}$), the \mathcal{L}_{ecc} -structure on K' is an RV_{ecc} -expansion of the $\mathcal{L}_{\text{val, ecc}}$ -structure, so Theorem 4.1.19 implies that $\text{Th}_{\mathcal{L}_{\text{ecc}}}(K')$ is also ω -h-minimal.

(2) This language \mathcal{L} defines an analytic structure on K (by [18, Section 4.4] (2)) and hence also on K' . Hence, it suffices to prove (3).

(3) The language \mathcal{L}_{ecc} has the shape of the language called \mathcal{L} in the above Case (ii) of Section 6.2, so by Theorem 6.2.1, $\text{Th}_{\mathcal{L}_{\text{ecc}}}(K')$ is ω -h-minimal. \square

6.3. $\mathcal{T}_{\text{omin}}$ -convex valued fields

Fix a language $\mathcal{L}_{\text{omin}}$ containing the language of ordered rings, and fix a complete o-minimal $\mathcal{L}_{\text{omin}}$ -theory $\mathcal{T}_{\text{omin}}$ containing the theory RCF of real closed fields. Given a pair of models $K_0 < K$ of $\mathcal{T}_{\text{omin}}$, we can turn K into a valued field by using the convex closure of K_0 in K as the valuation ring \mathcal{O}_K . We suppose that $\mathcal{O}_K \neq K$ and we let \mathcal{L} be the extension of $\mathcal{L}_{\text{omin}}$ by a predicate symbol for \mathcal{O}_K . In [35, 32] van den Dries–Lewenberg obtained various results about the model theory of such valued fields K as \mathcal{L} -structures. In particular, the theory $\mathcal{T} := \text{Th}_{\mathcal{L}}(K)$ only depends on $\mathcal{T}_{\text{omin}}$ (and not on the choice of K and K_0 , provided that $\mathcal{O}_K \neq K$) [35, Corollary 3.13]. (Van den Dries–Lewenberg call such a ring \mathcal{O}_K a ‘ $\mathcal{T}_{\text{omin}}$ -convex subring of K ’. Accordingly, and following other subsequent literature, we call K a ‘ $\mathcal{T}_{\text{omin}}$ -convex valued field’.)

We will prove that this theory \mathcal{T} is 1-h-minimal, under the assumption that no fast-growing functions are definable in $\mathcal{T}_{\text{omin}}$. In \mathbb{R} , ‘no fast-growing functions’ means that every definable function is eventually bounded by a function of the form $x \mapsto x^n$. In arbitrary real closed fields, the right generalization is power-boundedness; see [55]:

Definition 6.3.1 (Power-bounded). A power function in K is an $\mathcal{L}_{\text{omin}}$ -definable function $g: K^\times \rightarrow K^\times$ that is an endomorphism of the multiplicative group K^\times . We call the $\mathcal{L}_{\text{omin}}$ -structure K (and its theory $\mathcal{T}_{\text{omin}}$) power-pounded if for every $\mathcal{L}(K)$ -definable function $f: K \rightarrow K$, there exists a power function g such that $|f(x)| \leq g(x)$ for all sufficiently large $x \in K$.

From now on, we will assume that $\mathcal{T}_{\text{omin}}$ is power-bounded.

The proof that $\text{Th}(K)$ is 1-h-minimal is essentially contained in the existing literature: Using the criteria given in Theorem 2.9.1, this can be deduced from [41, Theorems 2.1 and 2.9]. However, Theorem 2.9 is a lot deeper than what we really need, so we give a more direct proof below (mainly following the ideas from [41]).

Lemma 6.3.2. *The theory of K (as an \mathcal{L} -structure) is 0-h-minimal.*

Proof. We assume that K is sufficiently saturated and use the criterion given by Lemma 2.4.4: Given a parameter set $A \subset K$ and a ball $B \subset K \setminus \text{acl}_K(A)$, we need to verify that all elements of B have the same type over $A \cup \text{RV}$. We may assume $A = \text{acl}_K(A)$.

Since $B \cap A = \emptyset$ and since $\mathcal{L}_{\text{omin}}$ -types over A correspond to cuts in A , all elements of B have the same $\mathcal{L}_{\text{omin}}$ -type over A . Thus [70, Lemma 3.15] applies and tells us that for any $x, x' \in B$, there exists an automorphism of K fixing A and RV but sending x to x' . This shows that x and x' have the same type over $A \cup \text{RV}$. \square

The following lemma states that \mathcal{L} -definable functions are piecewise $\mathcal{L}_{\text{omin}}$ -definable. This is already stated in [32, Lemma 2.6], but we shall use a variant from [70]:

Lemma 6.3.3 ([70, Lemma 3.3]). *Let $f: K \rightarrow K$ be an $\mathcal{L}(A)$ -definable function for some $A \subset K \cup \text{RV}$. Then there exists a partition of K into finitely many $\mathcal{L}(A)$ -definable sets X_i and $\mathcal{L}_{\text{omin}}(A \cap K)$ -definable functions $g_i: K \rightarrow K$ such that $f|_{X_i} = g_i|_{X_i}$ for each i .*

It might be a bit unclear from the formulation of that lemma in [70] whether it is intended that parameters from RV are allowed. In any case, the proof given in [70] goes through with parameters from RV .

Theorem 6.3.4 ($\mathcal{T}_{\text{omin}}$ -convex examples). *Let $\mathcal{T}_{\text{omin}}$ be a power-bounded o-minimal theory containing the theory of real closed fields in a language $\mathcal{L}_{\text{omin}}$ containing the language of rings. Let \mathcal{T} be the theory of $\mathcal{T}_{\text{omin}}$ -convex valued fields, in the language \mathcal{L} that is obtained by extending $\mathcal{L}_{\text{omin}}$ by a predicate symbol for the valuation ring (as explained at the beginning of this subsection). Then \mathcal{T} is 1-h-minimal.*

Proof. We use the criteria from Theorem 2.9.1 to prove 1-h-minimality, so let an A -definable map $f: K \rightarrow K$ be given for some $A \subset K \cup \text{RV}$. Let $X_i \subset K$ and $g_i: K \rightarrow K$ be as obtained from Lemma 6.3.3.

Condition (T2) of Theorem 2.9.1 holds for each g_i by o-minimality (namely, the set $\{d \in K \mid g_i^{-1}(d) \text{ is infinite}\}$ is finite), so it also holds for f .

Condition (T1), too, follows for f if we can prove it for each g_i . Indeed, take the union of the sets C obtained for all the g_i , and further enlarge C so that it 1-prepares each X_i . (This is possible, since by Lemma 6.3.2, we already have 0-h-minimality.) So it remains to prove Condition (T1) for an $\mathcal{L}_{\text{omin}}(A \cap K)$ -definable function g_i .

By o-minimality, we find a finite $(A \cap K)$ -definable $C \subset K$ such that g_i is continuously differentiable on $K \setminus C$. Further enlarge C (using 0-h-minimality and Corollary 2.6.6) so that it 1-prepares the map $K \rightarrow \Gamma_K, x \mapsto |g'_i(x)|$.

Let $B \subset K$ be a ball 1-next to C . We claim that Condition (T1) is satisfied with $\mu_B := |g'_i(x)|$ for any $x \in B$. Indeed, let $x_1, x_2 \in B$ be given, with $x_1 \neq x_2$. By the Mean Value Theorem for o-minimal fields, there exists an x_3 in-between (and hence also in B) such that $g_i(x_1) - g_i(x_2) = g'_i(x_3) \cdot (x_1 - x_2)$. Taking valuations on both sides implies $|g_i(x_1) - g_i(x_2)| = \mu_B \cdot |x_1 - x_2|$, as desired. \square

Remark 6.3.5. The assumption that $\mathcal{T}_{\text{omin}}$ is power-bounded is necessary to obtain 1-h-minimality of \mathcal{T} . Indeed, in the presence of an exponential map, we can define $K \rightarrow \text{RV}, x \mapsto \text{rv}(e^x)$, whose fibers are

exactly the translates of the maximal ideal $B_{<1}(0)$ and that hence cannot be 1-prepared in the sense of Corollary 2.6.6.

Remark 6.3.6. We were not able to prove that power-bounded $\mathcal{T}_{\text{omin}}$ -convex valued fields are ω -h-minimal. This is one of the main reasons we only assume 1-h-minimality in most of the paper.

Using methods from non-standard analysis, results in a $\mathcal{T}_{\text{omin}}$ -convex valued field K can often be translated into results about K as an $\mathcal{L}_{\text{omin}}$ -structure. We finish this subsection by stating what our Taylor approximation result (Theorem 3.2.2) becomes under such a translation, namely a version of Taylor approximation that has some uniformity even when one approaches a bad point.

Corollary 6.3.7. *Let $\mathcal{T}_{\text{omin}}$ be a power-bounded o -minimal theory containing the theory of real closed fields, in a language $\mathcal{L}_{\text{omin}}$ containing the language of rings. Let $K \models \mathcal{T}_{\text{omin}}$ be a model, let $f : K \rightarrow K$ be an $\mathcal{L}_{\text{omin}}$ -definable function, and let $r \in \mathbb{N}$ be given. Then there exists a finite $\mathcal{L}_{\text{omin}}$ -definable set $C \subset K$ and a constant $c \in K_{>0}$ such that for every pair $x_0, x \in K$ satisfying*

$$c \cdot |x - x_0| < \min_{a \in C} |x_0 - a|, \tag{6.3.1}$$

we have

$$|f(x) - T_{f, x_0}^{\leq r}(x)| \leq c \cdot |f^{(r+1)}(x_0) \cdot (x - x_0)^{r+1}| \tag{6.3.2}$$

(where $|\cdot|$ denotes the usual absolute value and $T_{f, x_0}^{\leq r}$ is the Taylor polynomial of f around x_0 of degree r ; see Definition 3.2.1).

More generally, if $(f_q)_{q \in K^m}$ is an $\mathcal{L}_{\text{omin}}$ -definable family of functions $K \rightarrow K$, then we obtain an $\mathcal{L}_{\text{omin}}$ -definable family of sets $(C_q)_{q \in K^m}$ and a constant $c \in K_{>0}$ that is independent of q such that the above holds for every $q \in K^m$.

Remark 6.3.8. Note that this result would be false without the assumption on power-boundedness. Indeed, one can check that it fails near 0 for the function $x \mapsto e^{1/x}$. On the other hand, it should be rather easy to obtain for sub-analytic functions, so this is another instance (along with the Jacobian Property) of a generalization of a result from an analytic setting to an axiomatic one.

In the proof, we use the following lemma:

Lemma 6.3.9. *For any $A \subset K$, every finite $\mathcal{L}(A)$ -definable set $C \subset K$ is already $\mathcal{L}_{\text{omin}}(A)$ -definable.*

Proof. Using the order, we reduce to the case where $C = \{a\}$ is a singleton. $\mathcal{L}(A)$ -definability means that $a = f(0)$ for some $\mathcal{L}(A)$ -definable function $f : K \rightarrow K$. By Lemma 6.3.3, $f(0) = g(0)$ for some $\mathcal{L}_{\text{omin}}(A)$ -definable $g : K \rightarrow K$; this implies that C is $\mathcal{L}_{\text{omin}}(A)$ -definable. □

Proof of Corollary 6.3.7. Fix a $|K|^+$ -saturated elementary extension $K' \succ K$, and let $\mathcal{O}_{K'}$ be the convex closure of K in K' . By Theorem 6.3.4, K' is 1-h-minimal as an \mathcal{L} -structure, for \mathcal{L} as in the theorem. (Note that the saturation assumption implies $\mathcal{O}_{K'} \neq K'$.) In the following, we denote the valuation on K' by $|\cdot|_v$, to distinguish it from the absolute value $|\cdot|$. We suppose that a family of functions $(f_q)_{q \in K^m}$ is given as in the corollary, and by abuse of notation, we also write $(f_q)_{q \in (K')^m}$ for the corresponding family in K' (defined by the same formula).

Suppose that the corollary fails: that is, for every $c \in K_{>0}$ and every formula ψ that could potentially define the family $(C_q)_{q \in K^m}$, there exists a $q \in K^m$ and a pair x_0, x for which the implication from equation (6.3.1) to equation (6.3.2) fails. By our saturation assumption, there exist $q \in (K')^m$ and $x_0, x \in K'$ such that the implication fails for every $c \in K_{>0}$ and every ψ . Using that $\mathcal{O}_{K'}$ is the convex closure of K , this failure for every $c \in K_{>0}$ is equivalent to the conjunction

$$|x - x_0|_v < \min_{a \in C_q} |x_0 - a|_v \tag{6.3.3}$$

and

$$|f(x) - T_{f_q, x_0}^{\leq r}(x)|_v > |f_q^{(r+1)}(x_0) \cdot (x - x_0)^{r+1}|_v. \tag{6.3.4}$$

So we obtained: For every $\mathcal{L}_{\text{omin}}(q)$ -definable set C_q , there exist $x_0, x \in K'$ in the same ball 1-next to C_q (by equation (6.3.3)) such that equation (6.3.4) holds. This contradicts Theorem 3.2.2: *a priori*, that theorem only provides a finite $\mathcal{L}(q)$ -definable set C , but $\mathcal{L}_{\text{omin}}(q)$ -definability of that set then follows using Lemma 6.3.9. \square

One can expect that similar results in higher dimensions can be obtained and that they may lead to finer versions of the preparation results in power-bounded real closed fields from [36, 57], which are used to deduce the existence of Mostowski’s Lipschitz stratifications.

6.4. Comparison to V-minimality

In [47], Hrushovski–Kazhdan introduced the notion of V -minimal theories, which at first sight has the same goal as Hensel minimality: namely, to provide a powerful axiomatic framework for geometry in valued fields. However, the relation between V -minimality and Hensel minimality is similar to the relation between strong minimality and o -minimality: By working in a strongly minimal theory of (algebraically closed) fields, one obtains many useful results about geometry in real closed fields, but one cannot treat genuinely o -minimal languages like \mathbb{R}_{exp} . In a similar way, working in a V -minimal theory of (algebraically closed) valued fields does provide many useful insights about Henselian valued fields (as explained in [47, Section 12]), but there are examples of Hensel minimal theories that cannot be treated in this way.

Concretely, since a V -minimal theory has no more structure on RV than a pure valued field, every definable function $K \rightarrow K$ ultimately grows like $x \mapsto x^r$ for some rational number r , and this remains true if we expand the language by predicates on RV (following [47, Section 12]). In contrast, Section 6.3 provides examples of 1- h -minimal structures without this property: for example, the T -convex structure obtained from the (power-bounded o -minimal) expansion of \mathbb{R} by one function $x \mapsto x^r$ for every real number r .

On the other hand, if we restrict to the context for which V -minimality has been designed, then it agrees with Hensel minimality. Moreover, in this case, 0- h -minimality and 1- h -minimality agree. Let us recall the definition of V -minimality.

Definition 6.4.1 (V -minimality; [47, Section 3.4]). Fix a language $\mathcal{L} \supset \mathcal{L}_{\text{val}}$ and a complete theory \mathcal{T} containing the theory $\text{ACVF}_{0,0}$ of algebraically closed fields of equi-characteristic 0. The theory \mathcal{T} is called V -minimal if for every model $K \models \mathcal{T}$, we have the following:

0. Every definable (with parameters) subset of K is a finite boolean combination of points, open balls and closed balls.
1. Every $\mathcal{L}(K)$ -definable subset of RV^n is already $\mathcal{L}_{\text{val}}(K)$ -definable (where \mathcal{L}_{val} is the pure valued field language).
2. Every definable (with parameters) family of nested closed balls in K has a non-empty intersection.
3. For every $A \subset K$, if $X \subset K$ is an A -definable set that is the union of finitely many disjoint closed balls B_i , then $\text{acl}_K(A) \cap B_i \neq \emptyset$ for every i .

Proposition 6.4.2 (V -minimality). *Suppose that \mathcal{T} is a complete theory containing $\text{ACVF}_{0,0}$, in a language $\mathcal{L} \supset \mathcal{L}_{\text{val}}$, and suppose moreover that every $\mathcal{L}(K)$ -definable subset of RV^n is already $\mathcal{L}_{\text{val}}(K)$ -definable. Then the following are equivalent:*

- (i) \mathcal{T} is V -minimal.
- (ii) \mathcal{T} is 0- h -minimal.
- (iii) \mathcal{T} is 1- h -minimal.

Proof of Proposition 6.4.2. (iii) \Rightarrow (ii) is trivial.

(i) \Rightarrow (iii): We use the criteria from Theorem 2.9.1, so let $f: K \rightarrow K$ be A -definable, for some $A \subset K \cup \text{RV}$.

First note that by the Remark just above [47, Lemma 3.30], adding parameters from $K \cup \text{RV}$ to the language preserves V -minimality, so using compactness, the results from [47] hold uniformly in families parametrized by K or RV .

By dimension theory (for example [47, Lemma 3.55]), f has only finitely many infinite fibers: that is, Condition (T2) from Theorem 2.9.1 holds. By applying [47, Corollary 4.3] to all fibers $f^{-1}(b)$ of f (where b runs over K), we find an A -definable map $\rho: K \rightarrow \text{RV}^k$ (for some $k \geq 0$) such that for each $\xi \in \text{RV}^k$, the restriction $f|_{\rho^{-1}(\xi)}$ is either constant or injective. Apply [47, Corollary 5.9] to each injective restriction $f|_{\rho^{-1}(\xi)}$ and refine the map ρ accordingly: that is, such that afterwards, f is ‘nice’ on each open ball contained in one fiber of ρ in the sense of [47, Definition 5.8]. Finally, by applying [47, Corollary 4.3] to the graph of ρ , we find a finite A -definable set C 1-preparing ρ (namely, the image of the map c provided by the corollary). Then f is nice on each ball 1-next to C , and this implies Condition (T1) from Theorem 2.9.1.

(ii) \Rightarrow (i): We prove the conditions from Definition 6.4.1:

(0) Let $X \subset K$ be definable, and let $C \subset K$ be a finite set preparing X . Then X can be written as a union of the form $\bigcup_{c \in C} X_c$, where $X_c = \{c + x \mid \text{rv}(x) \in Z_c\}$ for suitable definable sets $Z_c \subset \text{RV}$. Using the assumption that definable subsets of RV are already definable in the language \mathcal{L}_{val} , we obtain that each X_c is a finite boolean combination of points, open balls and closed balls. This then also follows for X :

(1) holds by assumption;

(2) is a special case of Lemma 2.7.1;

(3) by 0-h-minimality, there exists a finite A -definable set C preparing X . This set C cannot be disjoint from any B_i , since for any $c \in C \setminus B_i$, the ball 1-next to c containing B_i is strictly bigger than B_i . \square

6.5. Some open questions

We finish the paper with a few questions; more questions are stated in the sequel [17]. Probably the most immediate question in this context is:

Question 6.5.1. Does 0-h-minimality imply 1-h-minimality, and does 1-h-minimality imply ω -h-minimality? More generally: For which $\ell < \ell'$ does ℓ -h-minimality imply ℓ' -h-minimality?

If 1-h-minimality is not equivalent to ω -h-minimality, there are still the following questions, motivated by the results from Sections 6.3 and 6.4.

Question 6.5.2. Are the $\mathcal{T}_{\text{omin}}$ -convex structures from Section 6.3 (with $\mathcal{T}_{\text{omin}}$ power-bounded) ω -h-minimal? Does V -minimality imply ω -h-minimality?

Acknowledgements. The authors would like to thank Y. Halevi, J. Koenigsmann, F. Loeser, A. Macintyre, W. Singhof and F. Vermeulen for interesting discussions on the topics of the paper, and L. van den Dries for suggesting the name of Hensel minimality.

Funding statement. The author R. C. was partially supported by the European Research Council under the European Community’s Seventh Framework Programme (FP7/2007-2013) with ERC Grant Agreement nr. 615722 MOTMELSUM and by KU Leuven IF C14/17/083, and thanks the Labex CEMPI (ANR-11-LABX-0007-01). The author I. H. was partially supported by *SFB 878: Groups, Geometry and Actions*, by the research training group *GRK 2240: Algebra-Geometric Methods in Algebra, Arithmetic and Topology*, and by individual research grant No. 426488848, all three funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation). Part of the work has been done while I. H. was affiliated with the University of Leeds. The author S. R. was partially supported by *GeoMod AAPG2019* (ANR-DFG), Geometric and Combinatorial Configurations in Model Theory.

Conflicts of Interest. None.

References

- [1] J. Ax and S. Kochen, *Diophantine problems over local fields. I*, *Amer. J. Math.* **87** (1965), 605–630.
- [2] J. Ax and S. Kochen, *Diophantine problems over local fields. II.: A complete set of axioms for p -adic number theory*, *Amer. J. Math.* **87** (1965), 631–648.
- [3] J. Ax and S. Kochen, *Diophantine problems over local fields. III. decidable fields*, *Amer. J. Math.* **83** (1966), 437–456.
- [4] Ş. Basarab, *Relative elimination of quantifiers for Henselian valued fields*, *Ann. Pure Appl. Logic* **53** (1991), no. 1, 51–74.
- [5] V. Batyrev and A. Moreau, *The arc space of horospherical varieties and motivic integration*, *Compos. Math.* **149** (2013), no. 8, 1327–1352. MR 3103067
- [6] L. Bélair, *Panorama of p -adic model theory*, *Annales des sciences mathématiques du Québec* **36** (2012), 43–75.
- [7] M. Bilu, *Motivic Euler products and motivic height zeta functions*, *Mem. Amer. Math. Soc.* (2018). arxiv.org/abs/1802.06836.
- [8] W. Casselman, J. E. Cely, and T. Hales, *The spherical Hecke algebra, partition functions, and motivic integration*, *Trans. Amer. Math. Soc.* **371** (2019), no. 9, 6169–6212.
- [9] A. Chambert-Loir and F. Loeser, *Motivic height zeta functions*, *Amer. J. Math.* **138** (2016), no. 1, 1–59. MR 3462879
- [10] Z. Chatzidakis and E. Hrushovski, *Model theory of difference fields*, *Trans. Amer. Math. Soc.* **351** (1999), no. 8, 2997–3071.
- [11] R. Cluckers, G. Comte, and F. Loeser, *Lipschütz continuity properties for p -adic semi-algebraic and subanalytic functions*, *GAGA (Geom. Funct. Anal.)* **20** (2010), 68–87.
- [12] R. Cluckers, G. Comte, and F. Loeser, *Local metric properties and regular stratifications of p -adic definable sets*, *Commentarii Mathematici Helvetici* **87** (2012), no. 4, 963–1009.
- [13] R. Cluckers, G. Comte, and F. Loeser, *Non-archimedean Yomdin-Gromov parametrizations and points of bounded height*, *Forum of Mathematics, Pi* **3** (2015), no. e5, 60 pages, doi:[10.1017/fmp.2015.4](https://doi.org/10.1017/fmp.2015.4).
- [14] R. Cluckers, A. Forey, and F. Loeser, *Uniform Yomdin-Gromov parametrizations and points of bounded height in valued fields*, *Algebra Number Theory* **14** (2020), no. 6, 1423–1456.
- [15] R. Cluckers, J. Gordon, and I. Halupczok, *Uniform analysis on local fields and applications to orbital integrals*, *Trans. Amer. Math. Soc. Ser. B* **5** (2018), 125–166, doi.org/[10.1090/btran/25](https://doi.org/10.1090/btran/25).
- [16] R. Cluckers, T. Hales, and F. Loeser, *Transfer principle for the fundamental lemma, On the stabilization of the trace formula*, *Stab. Trace Formula Shimura Var. Arith. Appl.*, vol. **1**, Int. Press, Somerville, MA, 2011, pp. 309–347.
- [17] R. Cluckers, I. Halupczok, S. Rideau and F. Vermeulen, *Hensel minimality II: Mixed characteristic and a diophantine application*, [arXiv](https://arxiv.org/abs/2012.08001) (2021).
- [18] R. Cluckers and L. Lipshitz, *Fields with analytic structure*, *J. Eur. Math. Soc. (JEMS)* **13** (2011), 1147–1223.
- [19] R. Cluckers and L. Lipshitz, *Strictly convergent analytic structures*, *J. Eur. Math. Soc. (JEMS)* **19** (2017), no. 1, 107–149.
- [20] R. Cluckers and F. Loeser, *b -minimality*, *J. Math. Log.* **7** (2007), no. 2, 195–227.
- [21] R. Cluckers and F. Loeser, *Constructible motivic functions and motivic integration*, *Inventiones Mathematicae* **173** (2008), no. 1, 23–121.
- [22] R. Cluckers and F. Loeser, *Constructible exponential functions, motivic Fourier transform and transfer principle*, *Annals of Mathematics* **171** (2010), 1011–1065.
- [23] R. Cluckers and F. Loeser, *Motivic integration in all residue field characteristics for Henselian discretely valued fields of characteristic zero*, *J. Reine Angew. Math.* **701** (2015), 1–31, DOI [10.1515/crelle-2013-0025](https://doi.org/10.1515/crelle-2013-0025).
- [24] R. Cluckers and F. Loeser, *Motivic mellin transformation*, **preprint** (2020).
- [25] P. J. Cohen, *Decision procedures for real and p -adic fields*, *Comm. Pure Appl. Math.* **22** (1969), 131–151.
- [26] J. Denef, *The rationality of the Poincaré series associated to the p -adic points on a variety*, *Inventiones Mathematicae* **77** (1984), 1–23.
- [27] J. Denef and L. van den Dries, *p -adic and real subanalytic sets*, *Annals of Mathematics* **128** (1988), no. 1, 79–138.
- [28] J. Denef and F. Loeser, *Germes of arcs on singular algebraic varieties and motivic integration*, *Inventiones Mathematicae* **135** (1999), 201–232.
- [29] J. Denef and F. Loeser, *Definable sets, motives and p -adic integrals*, *Journal of the American Mathematical Society* **14** (2001), no. 2, 429–469.
- [30] J. Denef and F. Loeser, *Geometry on arc spaces of algebraic varieties*, European Congress of Mathematics, Vol. **I** (Barcelona, 2000), Progr. Math., vol. 201, Birkhäuser, Basel, 2001, pp. 327–348.
- [31] L. van den Dries, *Remarks on Tarski’s problem concerning $(\mathbb{R}, +, \cdot, \exp)$* , *Logic colloquium ’82 (Florence, 1982)*, Stud. Logic Found. Math., vol. **112**, North-Holland, Amsterdam, 1984, pp. 97–121. MR 762106
- [32] L. van den Dries, *T -convexity and tame extensions. II*, *J. Symbolic Logic* **62** (1997), no. 1, 14–34. MR 1450511 (98h:03048)
- [33] L. van den Dries, *Tame topology and o -minimal structures*, Lecture note series, vol. **248**, Cambridge University Press, 1998.
- [34] L. van den Dries, D. Haskell, and D. Macpherson, *One-dimensional p -adic subanalytic sets*, *Journal of the London Mathematical Society* **59** (1999), no. 1, 1–20.
- [35] L. van den Dries and A. H. Lewenberg, *T -convexity and tame extensions*, *J. Symbolic Logic* **60** (1995), no. 1, 74–102. MR 1324502 (96a:03048)
- [36] L. van den Dries and P. Speissegger, *O -minimal preparation theorems, Model theory and applications*, Quad. Mat., vol. **11**, Aracne, Rome, 2002, pp. 87–116.
- [37] J. Eršov, *On the elementary theory of maximal normed fields*, *Dokl. Akad. Nauk SSSR* **165** (1965), 21–23.
- [38] J. Flennner, *Relative decidability and definability in henselian valued fields*, *J. Symbolic Logic* **76** (2011), no. 4, 1240–1260.

- [39] A. Forey, *Motivic local density*, *Math. Z.* **287** (2017), no. 1–2, 361–403. MR 3694680
- [40] A. Fornasiero and I. Halupczok, *Dimension in topological structures: topological closure and local property*, *Groups and Model Theory*, *Contemp. Math.*, vol. **576**, Amer. Math. Soc., Providence, RI, 2012, pp. 89–94.
- [41] E. García Ramírez, *Non-archimedean stratifications in power bounded t -convex fields*, *Notre Dame J. Form. Log.* **61** (2020), no. 3, 441–465.
- [42] E. García Ramírez, *Non-archimedean stratifications of tangent cones*, *MLQ Math. Log. Q.* **63** (2017), no. 3–4, 299–312. MR 3724385
- [43] I. Halupczok, *Non-Archimedean Whitney stratifications*, *Proc. Lond. Math. Soc. (3)* **109** (2014), no. 5, 1304–1362. MR 3283619
- [44] I. Halupczok and Y. Yin, *Lipschitz stratifications in power-bounded o -minimal fields*, *J. Eur. Math. Soc. (JEMS)* **20** (2018), no. 11, 2717–2767. MR 3861807
- [45] D. Haskell and D. Macpherson, *Cell decompositions of C -minimal structures*, *Ann. Pure Appl. Logic* **66** (1994), no. 2, 113–162.
- [46] D. Haskell and D. Macpherson, *A version of o -minimality for the p -adics*, *J. Symbolic Logic* **62** (1997), no. 4, 1075–1092.
- [47] E. Hrushovski and D. Kazhdan, *Integration in valued fields*, *Algebraic geometry and number theory*, *Progr. Math.*, vol. **253**, Birkhäuser Boston, Boston, MA, 2006, pp. 261–405.
- [48] E. Hrushovski and D. Kazhdan, *Motivic Poisson summation*, *Mosc. Math. J.* **9** (2009), no. 3, 569–623.
- [49] E. Hrushovski and F. Loeser, *Non-archimedean tame topology and stably dominated types*, *Annals of Mathematics Studies*, vol. **192**, Princeton University Press, Princeton, NJ, 2016.
- [50] E. Hrushovski, B. Martin, and S. Rideau, *Definable equivalence relations and zeta functions of groups, with an appendix by Raf Cluckers*, *J. Eur. Math. Soc. (JEMS)* **20** (2018), no. 10, 2467–2537.
- [51] Franziska Jahnke, Pierre Simon, and Erik Walsberg, *Dp -minimal valued fields*, *J. Symb. Log.* **82** (2017), no. 1, 151–165. MR 3631280
- [52] J. F. Knight, A. Pillay, and C. Steinhorn, *Definable sets in ordered structures. II*, *Trans. Amer. Math. Soc.* **295** (1986), no. 2, 593–605. MR 833698
- [53] A. Macintyre, *On definable subsets of p -adic fields*, *Journal of Symbolic Logic* **41** (1976), 605–610.
- [54] D. Macpherson and C. Steinhorn, *On variants of o -minimality*, *Annals of Pure and Applied Logic* **79** (1996), no. 2, 165–209.
- [55] C. Miller, *A growth dichotomy for o -minimal expansions of ordered fields*, *Logic: from foundations to applications* (Staffordshire, 1993), Oxford Sci. Publ., Oxford Univ. Press, New York, 1996, pp. 385–399. MR 1428013 (98a:03052)
- [56] K. H. Nguyen, *Uniform rationality of the Poincaré series of definable, analytic equivalence relations on local fields*, *Trans. Amer. Math. Soc. Ser. B* **6** (2019), 274–296.
- [57] N. Nguyen and G. Valette, *Lipschitz stratifications in o -minimal structures*, *Ann. Sci. École Norm. Sup. (4)* **49** (2016), no. 2, 399–421.
- [58] J. Nicaise, *Geometric invariants for non-archimedean semialgebraic sets*, *Algebraic geometry: Salt Lake City 2015*, *Proc. Sympos. Pure Math.*, vol. **97**, Amer. Math. Soc., Providence, RI, 2018, pp. 389–403. MR 3821179
- [59] A. Onshuus and A. Usvyatsov, *On dp -minimality, strong dependence and weight*, *J. Symbolic Logic* **76** (2011), no. 3, 737–758. MR 2849244
- [60] J. Pas, *Uniform p -adic cell decomposition and local zeta functions*, *Journal für die reine und angewandte Mathematik* **399** (1989), 137–172.
- [61] J. Pila, *O -minimality and the André-Oort conjecture for \mathbb{C}^n* , *Ann. of Math. (2)* **173** (2011), no. 3, 1779–1840.
- [62] J. Pila and A. J. Wilkie, *The rational points of a definable set*, *Duke Math. J.* **133** (2006), no. 3, 591–616.
- [63] A. Pillay and C. Steinhorn, *Definable sets in ordered structures. I*, *Trans. Amer. Math. Soc.* **295** (1986), no. 2, 565–592. MR 833697
- [64] S. Rideau, *Some properties of analytic difference valued fields*, *Journal of the Institute of Mathematics of Jussieu* **16** (2017), no. 3, 447–499, DOI: <http://doi.org/10.1017/S1474748015000183>.
- [65] T. Scanlon, *Valuation theory and its applications volume II*, *Fields Institute Communications Series*, Quantifier elimination for the relative Frobenius, pp. 323 – 352, AMS, Providence, (2003).
- [66] J.-P. Serre, *Quelques applications du théorème de densité de Chebotarev*, *Publ. Math. Inst. Hautes Études Sci.* (1981), no. **54**, 323–401.
- [67] K. Tent and M. Ziegler, *A course in model theory*, *Lecture Notes in Logic*, vol. **40**, Association for Symbolic Logic, La Jolla, CA; Cambridge University Press, Cambridge, 2012. MR 2908005
- [68] F. Vermeulen, *Hensel minimality: Geometric criteria for ℓ - h -minimality*, (2022), arxiv.org/abs/2204.04158.
- [69] Y. Yin, *Integration in algebraically closed valued fields*, *Ann. Pure Appl. Logic* **162** (2011), no. 5, 384–408.
- [70] Y. Yin, *Generalized Euler characteristic in power-bounded T -convex valued fields*, *Compos. Math.* **153** (2017), no. 12, 2591–2642. MR 3705299
- [71] Y. Yomdin, *C^k -resolution of semialgebraic mappings. Addendum to: 'Volume growth and entropy'*, *Israel J. Math.* **57** (1987), no. 3, 301–317. MR 889980
- [72] Y. Yomdin, *Volume growth and entropy*, *Israel J. Math.* **57** (1987), 285–300.