STRATIFICATIONS IN VALUED FIELDS

IMMANUEL HALUPCZOK

Abstract. In these notes, I present a new, strong notion of stratifications which describe singularities of sets in Henselian valued fields. In the first part, I give the definition, some examples, and I state the main result about their existence; in the second part, I explain how such a stratification in a valued field induces a classical Whitney stratification in $\mathbb{R}$.

1. Introduction

A classical tool to describe singularities of a subset X of $\mathbb{R}^n$ (or of $\mathbb{C}^n$) are Whitney stratifications. More precisely, if X is, say, a semi-algebraic subset of $\mathbb{R}^n$ (i.e., a set given by polynomial equations and inequations), then by a "Whitney stratification for X", in these notes we mean a partition of $\mathbb{R}^n$ into smooth submanifolds $S_0, \dots, S_n$ where S_d has dimension d, such that X is a union of some of the connected components of these sets, and such that certain regularity conditions relating the different sets S_d are satisfied. (Often, one only considers partitions $S_0, \dots, S_{\dim X}$ of X, but for the purpose of these notes, it is more handy to work with partitions of $\mathbb{R}^n$.) Roughly, these regularity conditions imply that S_0 contains all the points where X is most singular, S_1 contains the slightly less singular points, etc.; see Fig. 1 for an example and Definition 3.1 for a precise definition.

The main goal of these notes is to present "t-stratifications", a somewhat similar notion in Henselians valued fields, which has been introduced in [1]; I will first define and explain that notion and then show how it is related to Whitney stratifications. Let us fix some notation for the remainder of these notes.

The author was supported by the project SFB 878 of the Deutsche Forschungsgemeinschaft.

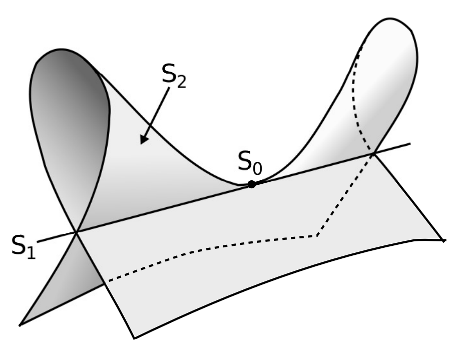

Figure 1. A surface X with a Whitney stratification (S_3 is the whole complement $\mathbb{R}^3 \setminus X$).

Notation 1.1. Let K be a Henselian valued field of equi-characteristic zero (i.e., both K and its residue field have characteristic zero). We set:

$$\begin{array}{rcl} \text{valued field} & K & \xrightarrow{v} \Gamma \cup \{\infty\} \quad \text{value group} \\ & \cup & \\ \text{valuation ring} & \mathcal{O}_K & \xrightarrow{\text{res}} k \quad \text{residue field} \\ & \cup & \\ \text{maximal ideal} & \mathcal{M}_K & \end{array}$$

A t-stratification is a partition $S_0 \,\dot{\cup}\, \dots \,\dot{\cup}\, S_n = K^n$ that describes singularities of a set $X \subseteq K^n$. Again, S_d is of dimension d and again, S_0 contains the worst singularities, and so on. However, the regularity conditions which ensure this are different and in a certain sense much stronger than the ones of Whitney stratifications. First of all, they imply that a t-stratification for a set X does not only see the singularities of X itself, but also those of its image in the residue field. (To be precise, some singularities might only become visible to the t-stratification if one works in the algebraic closure of K.)

Describing the singularities in the residue field is, by itself, not yet a strong property of t-stratificatoins. However, the notion of t-stratification is invariant under translation by elements of K^n and scaling by elements of $K^\times$. Thus, instead of intersecting X with the closed ball $\mathcal{O}_K^n$ and quotienting by the corresponding open ball $\mathcal{M}_K^n$ (which yields the image of X in the residue field), we can also intersect X with any other closed ball $B \subseteq K^n$ and take the corresponding quotient. The strength of t-stratifications lies in the fact that a single stratification simultaneously describes the singularities in all those quotients.

2. Definition of t-stratifications

We will need the notion of a ball in K^n. The most natural norm in a valued field is the maximum norm, so a ball in K^n is just a product of one ball in each coordinate, all of which are of the same radius. More precisely, we define:

Definition 2.1. For a tuple $a = (a_1, \dots, a_n) \in K^n$, we set $v(a) := \min_i v(a_i)$. An *open ball* in K^n is a set of the form

$$B(a, > \lambda) := \{x \in K^n \mid v(x - a) > \lambda\}$$

for some $a \in K^n$ and $\lambda \in \Gamma$.

Before we go on, let us first fix the context in which we will be working. Whitney stratifications exist in different contexts: algebraic sets, semi-algebraic sets, analytic sets, and several others. We have similar choices for t-stratifications, the simplest one being to work with algebraic subsets of K^n. However, since we will use balls in the definition of a t-stratification and since balls are not algebraic sets, it is more natural to work with a larger class of sets, namely definable sets in the sense of model theory (to be more precise: sets definable in the language of rings, with a predicate for the valuation ring). These should be thought of as the right analogue in Henselian valued fields of the semi-algebraic sets in $\mathbb{R}^n$. Readers who are not interested in definable sets can just stay in the algebraic world everywhere, except at one place (in Definition 2.3), where we need to work with a definable map between balls; however, by simply ignoring the word "definable" there, one still gets a good approximation to the notion of t-stratification.

For readers who want to know it more precisely, here is a quick formal definition of definable sets.

Definition 2.2. The class of *definable sets* is the smallest class of subsets of K^n (for all n) which satisfies:

(1) algebraic sets (i.e., zero sets of polynomials) are definable;
(2) if x is an n-tuple of variables and $f, g \in K[x]$ are polynomials, then $\{x \in K^n \mid v(f(x)) \geq v(g(x))\}$ is definable;
(3) boolean combinations of definable sets are definable;
(4) if $X \subseteq K^n$ is definable, then so is $\pi(X)$, where $\pi\colon K^n \twoheadrightarrow K^{n-1}$ is a coordinate projection.

A map between definable sets is called *definable* if its graph is definable.

To get a rough idea of what this means: if K is algebraically closed, then (1)–(3) suffice to obtain all definable sets (this result, called quantifier elimination, was first proven in [3]); there are other results describing definable sets even more precisely, and one can also obtain similar results for general Henselian K; the most important one is probably the cell decomposition result by Denef–Pas [2]. This allows for example to define a good notion of dimension for definable sets.

Now let us come back to our t-stratifications and introduce the regularity condition. The central ingredient is a notion of being "roughly translation invariant in d dimensions" on a given ball B. More precisely, for suitable balls $B \subseteq K^n$, we will require that there exists a d-dimensional vector space $V \subseteq K^n$ such that all sets $S_i \cap B$ and also the set $X \cap B$ are "translation invariant in the directions V up to a small perturbation". Here, one has to be careful about the notion of "small perturbation"; the following definition makes this precise.

Definition 2.3. Suppose $B \subseteq K^n$ is a ball and $Y_1, \dots, Y_\ell$ are arbitrary subsets of K^n.

(1) We say that $(Y_i)_i$ is *V-translatable on B*, where $V \subseteq K^n$ is a vector space, if there exists a definable bijection $\phi\colon B \to B$ with the following properties:
   (a) For each i and for any two points $z, z' \in B$ with $z - z' \in V$, we have $z \in \phi(Y_i \cap B)$ iff $z' \in \phi(Y_i \cap B)$.
   (b) For any two points $y, y' \in B$, we have $v\big((\phi(y) - \phi(y')) - (y - y')\big) > v(y - y')$.
(2) We say that $(Y_i)_i$ is *d-translatable on B* for some $d \in \mathbb{N}$ if there exists a d-dimensional $V \subseteq K^n$ such that $(Y_i)_i$ is V-translatable on B.

Note that 0-translatable is an empty condition, whereas n-translatable means that each of the sets Y_i is either disjoint from B or contains B. Note also that if $Y \cap B \neq \emptyset$ and Y is d-translatable on B, then Y has dimension at least d.

Remark 2.4. Condition (b) says: applying ϕ does not change the difference of two points y, y' too much. In particular, if ϕ is differentiable, then it implies that the derivative of ϕ is close to the identity map $K^n \to K^n$, which in turn implies that the tangent space of each set Y_i at any point of B is close to a space containing V. Neither the derivative of ϕ nor these tangent space need to really exist, but the statement about tangent spaces can be made precise anyway by introducing a suitable notion of "approximate tangent space".

Now we can define the notion of a t-stratification describing the singularities of a set $X \subseteq K^n$. A first attempt could be to require that for any point $x \in S_d$, there exists an open ball B containing x such that everything is d-translatable on B (everything means: X and all sets S_j). However, such a stratification would not see the singularities of the image of X in the residue field: to get that, we need that B is at least as big as the ball we are quotienting out when we pass to the residue field (i.e., $\mathcal{M}_K^n$). It turns out that the right notion of t-stratification is obtained by requiring d-translatability on every ball where we can possibly hope for it: any ball which is disjoint from $S_0 \cup \cdots \cup S_{d-1}$. Indeed, if B would not be disjoint from S_j for some $j < d$, then S_j would not be d-translatable on B since it has dimension $j < d$. Here is the complete definition.

Definition 2.5. Let $X \subseteq K^n$ be a subset. A *t-stratification* for X is a partition $K^n = S_0 \,\dot\cup\, \ldots \,\dot\cup\, S_n$ with the following properties; we write $S_{\le d}$ for the union $S_0 \cup \cdots \cup S_d$.

1. Each set $S_{\le d}$ is algebraic.
2. Each S_d is either empty or of dimension d.
3. For each open ball $B \subseteq K^n$, if $B \cap S_{\le d-1} = \emptyset$ (for some $d \ge 1$), then $(X, S_d, \ldots, S_n)$ is d-translatable on B.

Now we can state the main theorem of [1].

Theorem 2.6 ([1, Theorem 1.1])**.** *Let K be a Henselian valued field with* $\operatorname{char} K = \operatorname{char} k = 0$. *For any algebraic (or more generally definable) subset $X \subseteq K^n$, there exists a t-stratification $(S_i)_i$. Moreover, if X is defined over a subring $R \subseteq K$, then the S_i can also be chosen to be defined over R.*

Note that even if X is allowed to be definable, we can find a t-stratification such that the sets $S_{\le i}$ are algebraic. (And to make things precise: by a "definable set which is defined over a subring R", we mean that the polynomials appearing in (1) and (2) of Definition 2.2 have coefficients from R.)

Remark 2.7. Such a result can also be formulated "uniformly in K": using the language from algebraic geometry, given X defined over R, we can find $(S_i)_i$ defined over R such that for all K containing R (Henselian and of equi-characteristic 0), we have that $(S_i(K))_i$ is a t-stratification for $X(K)$.

Example 2.8. Consider the cusp curve $X = \{(x, y) \in K^2 \mid x^3 = y^2\}$ (see Figure 2). Here, we set $S_0 := \{(0, 0)\}$ (the singularity), $S_1 := X \setminus S_0$, and $S_2 := K^2 \setminus X$. To check whether this is a valid t-stratification for X, we have to verify the translatability condition on different balls B.

- By Definition 2.5, we have to check that (X, S_2) is 2-translatable on any ball $B \subseteq S_2$; but in that case, we have $B \cap X = \emptyset$ and $B \cap S_2 = S_2$, which indeed implies 2-translatability.
- Now suppose $B \cap X \ne \emptyset$. If $B \cap S_0 \ne \emptyset$, no condition is required, so suppose $(0, 0) \notin B$; then we have to check that X is 1-translatable on B (or, more precisely, that (X, S_1, S_2) is 1-translatable on B, but this then follows). The computation is not very difficult and left to the reader.

Example 2.9. The surface in Figure 1 is a standard example for Whitney stratifications illustrating that it is not enough to require the strata to be smooth.

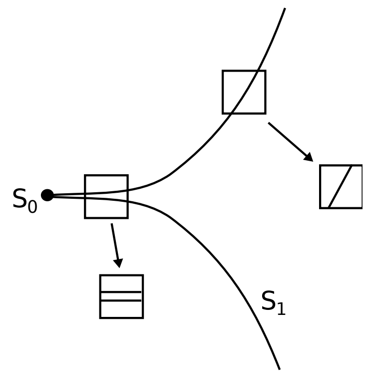

FIGURE 2. This is a picture of the cusp in $\mathbb{R}^2$, but nevertheless, it illustrates well what happens in K^2: in any square containing the singularity $(0,0)$, the curve is not close to something translation invariant; however, in the little squares which have been drawn, it is almost straight, and indeed it can be made translation invariant by a "small perturbation" in the sense of Definition 2.3.

(Otherwise, one could take $S_0 = \emptyset$ and S_1 to be the whole line.) The corresponding surface also exists in valued fields and we get exactly the corresponding t-stratification. Indeed, one can check:

- Away from S_1, the surface is 2-translatable.
- On a ball intersecting S_1 but not S_0, the surface is V-translatable where V is the direction of the line S_1.
- The surface is not translatable on any ball containing S_0, since the intersection with a plane Y perpendicular to S_1 has a little loop if Y does not contain S_0, whereas it is just a cusp curve if Y does contain S_0.

Example 2.10. Consider the hyperbola $X = \{(x,y) \in K^2 \mid xy = 1\}$. Since this does not seem to have a singularity, a first guess would be to set $S_1 := X$ and $S_2 := K^2 \setminus X$. However, it turns out that X is not 1-translatable on any ball B which contains $\mathcal{O}_K^2$, and the reason is a singularity which we overlooked. To see it, first choose any element $r \in K^\times$ of positive valuation and scale everything down by a factor r; we get a new set $X' = \{(x,y) \in K^2 \mid xy = r^2\}$. Obviousely, a t-stratification for X should also scale down to a t-stratification of X'. Now the image of X' in the residue field is $\{(x,y) \in k^2 \mid xy = 0\}$ (since $v(r^2) > 0$), and here, we see the missing singularity at $(0,0)$. To obtain a valid t-stratification, it suffices to remove the point $(0,0)$ from S_2 and to put it into S_0 instead.

This example also illustrates one of the stranger properties of t-stratifications: instead of adding $(0,0)$ to S_0, any other point of $\mathcal{O}_K^2$ would also have worked, so t-stratifications are pretty uncanonical. One could get rid of this uncanonicity using some abstract nonsense, but on the other hand, one strength of the theorem about existence of t-stratifications is that in higher dimensions (i.e., if a whole stratum S_d is uncanonical for $d \geq 1$), we can choose those non-canonical points in a uniform way.

3. Connection to Whitney stratifications

Up to now, the relation between t-stratifications and Whitney stratifications seems rather vague. In this section, we will make it precise. We start by recalling the definition of a Whitney stratification.

The regularity condition in the definition of a Whitney stratification describes the tangent space T_yS_j of a stratum S_j at a point $y \in S_j$ near a point $x \in S_d$ for $d < j$. Roughly, T_yS_j should be close to a space T containing (a) the tangent space T_xS_d and (b) the line connecting y to the closest point of S_d. Here is a formal definition.

Definition 3.1. Let $X \subseteq \mathbb{R}^n$ be a subset. A *Whitney stratification for X* is a partition $K^n = S_0 \dot\cup \dots \dot\cup S_n$ satisfying the following properties; we write $S_{\le d}$ for the union $S_0 \cup \dots \cup S_d$.

1. X is a union of some of the connected components of the sets S_d.
2. Each set $S_{\le d}$ is algebraic.
3. For each d, S_d is either empty or smooth of dimension d.
4. *Whitney's condition (b):* suppose that for some d, j with $d < j$ we have one sequence of points $x_\mu \in S_d$ and one sequence $y_\mu \in S_j$, both of which converge to the same point $x \in S_d$ for $\mu \to \infty$. Suppose moreover that the limit spaces $T := \lim_{\mu\to\infty} T_{y_\mu}S_j$ and $T' := \lim_{\mu\to\infty} \mathbb{R}\cdot(x_\mu - y_\mu)$ both exist (these limits are computed in the corresponding Grassmanians). We require that under these assumptions, we have $T' \subseteq T$.

Before the definition, I claimed that T is also required to contain T_xS_d (this is called *Whitney's condition (a)*); indeed, this can be deduced from Whitney's condition (b) by choosing a sequence x_μ which converges to x more slowly than y_μ, and which approaches x from any direction of T_xS_d.

Remark 3.2. For a t-stratification, we have d-translatability on a neighbourhood of any point $x \in S_d$; intuitively, this is closely related to the condition $T_xS_d \subseteq T$ of Whitney stratifications: by Remark 2.4, the approximate tangent space of S_j at y contains the approximate tangent space S_d at x; so one just needs to remove the words "approximate" and instead insert "for $y \to x$" to obtain Whitney's condition (a).

However, up to now we saw no obvious relation between the d-translatability condition of a t-stratifications and the full condition (b) of Whitney; indeed, to get such a relation, we will first need to prove an additional property of t-stratifications.

3.1. t-stratifications induce Whitney stratifications.

Now we can formulate a precise relation between Whitney stratifications and t-stratifications, namely that for suitable valued fields K, a t-stratification of K^n induces a Whitney stratification of $\mathbb{R}^n$. I only proved this under a rather strong assumption on K coming from model theory: K is an $\aleph_1$-saturated elementary extension of $\mathbb{R}$; however, it is plausible that the only assumption which is really needed (and which follows from the model theoretic one) is that K is a real closed field strictly containing $\mathbb{R}$.

On such a field K, there is a natural valuation, obtained by defining the valuation ring to be the convex closure of $\mathbb{R}$: $\mathcal{O}_K = \bigcup_{r\in\mathbb{R}_{>0}}(-r, r)$. With this definition, one easily checks that the maximal ideal conists of all "infinitesimal" elements: $\mathcal{M}_K = \bigcap_{r\in\mathbb{R}_{>0}}(-r, r)$, that the residue field is isomorphic to $\mathbb{R}$, and that K is Henselian.

Theorem 3.3 ([1, Theorem 6.11]). *Let K be as above and let $X(\mathbb{R}) \subseteq \mathbb{R}^n$ be a semi-algebraic set. Suppose that $(S_i(K))_i$ is a t-stratification for the corresponding semi-algebraic set $X(K) \subseteq K^n$, and suppose that it is defined over $\mathbb{R} \subset K$. (Such t-stratifications exist by Theorem 2.6, since X is defined over $\mathbb{R}$.) Then $(S_i(\mathbb{R}))_i$ is a partition of $\mathbb{R}^n$ which is a Whitney stratification for $X(\mathbb{R})$, except that the strata might be only C^1 instead of smooth.*

We obtain smooth strata if we additionally assume that $(S_i(\tilde{K}))_i$ is a t-stratification for $X(\tilde{K})$, where $\tilde{K}$ is the algebraic closure of K. (This can be assumed by Remark 2.7.)

The reason for this slightly strange additional condition to get smooth strata is that we did not put any smoothness condition in the definition of t-stratifications. However, a variant of Theorem 3.3 also works for $\mathbb{R}$ replaced by $\mathbb{C}$ (also yielding a Whitney stratification whose strata are C^1), but there, C^1 already implies smooth.

The philosophy behind the proof of this theorem is that if we have a statement in $\mathbb{R}$ which speaks about a limit for some $\epsilon \to 0$, then this can be translated to an equivalent statement in K, where ϵ is replaced by an "infinitesimal" element of K, i.e., an element of positive valuation. A standard example (which already works when K is any field extension of $\mathbb{R}$, with the natural valuation defined as above) is that the derivative of a polynomial $f(x) \in \mathbb{R}[x]$ at $a \in \mathbb{R}$ is equal to $\operatorname{res}(\frac{f(a+h)-f(a)}{h})$ for any $h \in K$ of positive valuation. If K is an $\aleph_1$-saturated elementary extension of $\mathbb{R}$, then this philosophy can be formulated as a precise theorem, which allows to translate a pretty large class of statements. In particular, it can be applied to the intutitive similarity between translatability and Whitney's condition (a) from Remark 3.2, yielding an actual implication. In a similar way, one obtains that each $S_d(\mathbb{R})$ is C^1 and all the other properties of Whitney stratifications, except for (the full version of) Whitney's condition (b). In the last part of these notes, I will present a result about t-stratifications in general which shows that they automatically satisfy an analogue of Whitney's condition (b) in the valued field setting, thus filling the gap.

4. An additional property of t-stratifications

Let now K again be an arbitrary Henselian valued field K with $\operatorname{char} K = \operatorname{char} k = 0$, and let us suppose that $(S_i)_i$ is a t-stratification for some set $X \subseteq K^n$. Consider an open ball $B \subseteq K^n$ which is disjoint from $S_0 \cup \dots \cup S_{d-1}$ and an open sub-ball $B' \subseteq B$ which additionally is disjoint from S_d. By definition of t-stratification, we have V-translatability on B for some d-dimensional V and V'-translatability on B' for some V' of dimension at least $d+1$. It is not hard to check that we can choose V' to contain V, but apart from that, we do not (yet) know anything about V'. Under the assumption that B' is sufficiently close to S_d, we can prove such an additional statement, which corresponds exactly to the missing part of Whitney's condition (b).

Theorem 4.1. *Let K, X, $(S_i)_i$, d, and B be as above. Then there exists a $\lambda \in \Gamma$ with the following property. Suppose that $B' \subseteq B$ is an open subball which is disjoint from S_d and suppose that we have points $x \in S_d \cap B$ and $y \in B'$ with $v(x-y) \geq \lambda$. Then $(X, S_{d+1}, \dots, S_n)$ is $K \cdot (x-y)$-translatable on B'.*

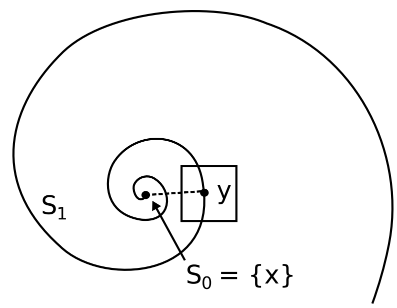

FIGURE 3. An evil example: no matter how close y gets to x, the logarithmic spiral restricted to the square around y is far from being translatable in the direction $x - y$.

This result really consists of two statements: first, for any x, we get a statement about all y which are sufficiently close, and second, this "sufficiently close" is, to some extend, uniform in x.

A counter-example to the first statement would be (a valued field version of) a logarithmic spiral (see Figure 3): here, no matter how close to x the point y is, S_1 will not be $K \cdot (x - y)$-translatable on any ball containing y. The idea to prove that such a counter-example does not exist is that otherwise, we could find a straight line L through x such that the intersection $L \cap S_1$ would be an infinite discrete set, and for such a set, there would be no t-stratification.

Once we have a bound λ which works for one single $x \in S_d \cap B$, we may simply apply an "almost-translation" sending x to any other point $x' \in S_d \cap B$ to verify that the same λ also works for x'.

REFERENCES

[1] I. HALUPCZOK, *Non-archimedean Whitney stratifications*, 2011. Preprint.
[2] J. PAS, *Uniform p-adic cell decomposition and local zeta functions*, J. Reine Angew. Math., 399 (1989), pp. 137–172.
[3] A. ROBINSON, *Complete theories*, North-Holland Publishing Co., Amsterdam, 1956.