Appendix A. Mastery material: Singular points

In the mastery material, we will define singular points of algebraic varieties. These are points where the variety is not smooth, for example the origin in the curves $y^2 = x^3$ or $y^2 = x^2(x+1)$.

To determine whether a point of a variety is singular, we use the tangent space of the variety at that point. This is a generalisation of the tangent line to a smooth curve (i.e. the line whose gradient is equal to the gradient of the curve). We say that the point is singular if the dimension of the tangent space at that point is equal to the dimension of the variety (for an irreducible variety – there is a slight complication for reducible varieties).

So first we have to define tangent spaces. We start with the case of affine algebraic sets, then generalise to quasi-projective varieties. We will finish with the proof that most points of a variety are not singular – more precisely, the singular points form a proper closed subset of the variety.

Tangent space of an affine algebraic set.

Let $V \subseteq \mathbb{A}^n$ be an affine algebraic set. We consider V equipped with an embedding $\mathbb{A}^n$ because we will use coordinates in our definition of the tangent space. There is also an intrinsic definition using local rings, which (unfortunately) we do not have time to cover.

Choose a point $x \in V$. For each polynomial $f \in k[X_1, \dots, X_n]$, let df_x denote the linear map $k^n \to k$ given by

$$df_x(a_1, \dots, a_n) = \sum_{i=1}^{n} \frac{\partial f}{\partial X_i}\bigg|_x a_i.$$

Informally: df_x sends a vector $\underline{a} \in k^n$ to the "directional derivative" of f at x along that vector. Thus $df_x(\underline{a}) = 0$ precisely for those directions in which f is stationary at x. Since the polynomials in $\mathbb{I}(V)$ are zero on V, we should expect polynomials in $\mathbb{I}(V)$ to be stationary along "tangent directions" to V. This motivates the following definition.

Definition. Let $V \subseteq \mathbb{A}^n$ be an affine algebraic set and let $x \in V$. The **tangent space** to V at x is

$$T_x V = \bigcap_{f \in \mathbb{I}(V)} \ker df_x \subseteq k^n.$$

This is sometimes called the **Zariski tangent space** when it is necessary to distinguish it from other kinds of tangent space.

In our definition of df_x, we used partial derivatives. Because we are only differentiating polynomials, these can be defined purely algebraically and therefore make sense over any field, even in positive characteristic where there is no analysis. However derivatives of polynomials can behave surprisingly in positive characteristic: over a field of characteristic p we have

$$\frac{\mathrm{d}}{\mathrm{d}X} X^p = pX^{p-1} = 0$$

so it is possible for a non-constant polynomial to have derivative equal to zero.

Similarly the informal motivation for the definition relied on our intuition from analysis about what happens over $\mathbb{C}$. Even over $\mathbb{C}$, our analytic intuition only works correctly if the variety is non-singular at x.

Let $f_1, \dots, f_n$ be a finite list of polynomials which generate $\mathbb{I}(V)$. It is easy to prove that the tangent space $T_x V$ can be calculated just by looking at this finite list of polynomials:

$$T_x V = \bigcap_{i=1}^{n} \ker d(f_i)_x.$$

Thus it is straightforward to calculate tangent spaces in practice, from a list of generators for $\mathbb{I}(V)$. There is just one thing to be careful of: the polynomials $f_1, \dots, f_n$ must *generate* the ideal $\mathbb{I}(V)$. It is not enough to take a list of functions which define V as an algebraic set but do not generate the ideal.

Example. As a very simple example, consider the line L in $\mathbb{A}^2$ defined by the polynomial $f(X, Y) = X$. At the point $(0, 0)$, we have

$$dX_0(a_1, a_2) = \left.\frac{\partial X}{\partial X}\right|_{(0,0)} a_1 + \left.\frac{\partial X}{\partial Y}\right|_{(0,0)} a_2 = 1.a_1 + 0.a_2 = a_1.$$

Since f generates $\mathbb{I}(L)$, we get

$$T_0 V = \ker dX_{(0,0)} = \{(a_1, a_2) \in k^2 : a_1 = 0\}.$$

This is what we should expect: the tangent space to a line is a line in the same direction.

However, if we were given the polynomial $g = X^2$, then $\mathbb{V}(g)$ is again L. We could try to calculate

$$dg_{(0,0)}(a_1, a_2) = \left.\frac{\partial X^2}{\partial X}\right|_{(0,0)} a_1 + \left.\frac{\partial X^2}{\partial Y}\right|_{(0,0)} a_2 = 0.a_1 + 0.a_2 = 0.$$

So $\ker dg_{(0,0)} = k^2$ which is too big. Thus using polynomials which do not generate the whole ideal of the variety may give the wrong answer for the tangent space.

Example. Consider the graph Γ of a polynomial function $g \in k[X]$. This is an affine algebraic set in $\mathbb{A}^2$ defined by the polynomial

$$f(X, Y) = Y - g(X).$$

Note that f is irreducible because it is monic of degree 1 in Y, so f generates $\mathbb{I}(\Gamma)$. We can calculate

$$df_{(x,g(x))}(a_1, a_2) = -g'(x).a_1 + 1.a_2$$

so

$$T_{(x,g(x))}\Gamma = \ker df_{(x,g(x))} = \{(a_1, a_2) \in k^2 : a_2 = g'(x).a_1\}.$$

Thus the tangent space to Γ at the point $(x, g(x))$ is a line with gradient $g'(x)$, as we expect.

Example. Consider the cuspidal cubic curve defined by the polynomial

$$g(X,Y) = Y^2 - X^3.$$

We have

$$\frac{\partial f}{\partial X} = -3X^2, \quad \frac{\partial f}{\partial Y} = 2Y.$$

Hence

$$\left.\frac{\partial f}{\partial X}\right|_{(0,0)} = \left.\frac{\partial f}{\partial Y}\right|_{(0,0)} = 0$$

and so $df_{(0,0)}$ is the zero map. Therefore $T_{(0,0)}V = k^2$.

We can't see this looking at a picture: this curve appears to have only the x-axis as a tangent line at the origin. This demonstates that we cannot rely on geometric intuition to calculate the tangent space at singular points: it is necessary to use the algebraic definition.

Definition of singular points.

Intuition suggests that, at non-singular points, the dimension of the tangent space should be equal to the dimension of the algebraic set. The above examples indicate that this breaks down at singular points. This motivates us to define a singular point to be a point $x \in V$ where $\dim T_x V \neq \dim V$.

However this simple definition only works correctly for irreducible algebraic sets. To see this: take a union of two irreducible components of different dimensions, say $V = V_1 \cup V_2$ where $\dim V_1 = 1$ and $\dim V_2 = 2$. By definition $\dim V = 2$. Being singular should be a "local" property: for a point $x \in V_1 \setminus (V_1 \cap V_2)$, whether x is a singular point of V should not care about V_2 – we ought to compare $\dim T_x V$ against $\dim V_1 = 1$, not against $\dim V = \dim V_2 = 2$.

In order to fix this and corerctly define singular points of reducible algebraic sets, we introduce a new definition:

Definition. Let V be a quasi-projective variety and let x be a point of V. The **local dimension** of V at x, written $\dim_x V$, is the maximum of the dimensions of those irreducible components of V which contain x.

(Taking the maximum of the dimensions of components fits with the way we defined the dimension of a reducible variety, but now we are ignoring components which do not contain x.)

Thus in our previous example $V = V_1 \cup V_2$, $\dim_x V = 2$ if $x \in V_2$ (including if $x \in V_1 \cap V_2$) while $\dim_x V = 1$ if $x \in V_1 \setminus (V_1 \cap V_2)$.

Now we can define singular points of a reducible algebraic set by using local dimension.

Definition. Let V be an affine algebraic set and let $x \in V$. Then x is a **singular point** of V if

$$\dim T_x V \neq \dim_x V.$$

(We may also express this as "V is **singular** at V.")

We will prove later that $\dim T_x V \geq \dim_x V$ always, so we could equivalently state this definition as: x is a singular point of V if $\dim T_x V > \dim_x V$.

If $V = \bigcup_{i=1}^r V_i$ is a union of irreducible components, then for any point x which lies in only one irreducible component V_i, we have

$$\dim_x V = \dim_x V_i$$

by definition. A little algebra also shows that $T_x V = T_x V_i$ and so V is singular at x if and only if V_i is singular at x.

On the other hand, if x lies in an intersection of two or more irreducible components of V, then it turns out that x is always a singular point of V. This is intuitively sensible, but requires too much algebra to prove in this course (specifically, it requires Nakayama's lemma).

Independence of embedding.

We have defined the tangent space only for affine algebraic sets, depending on the embedding into $\mathbb{A}^n$. Our definition of the tangent space gives not just an abstract vector space, but a subspace of k^n (where n is the dimension of the affine space into which we embed V). We can canonically identify k^n with the tangent space of $\mathbb{A}^n$ at any point, so one way of looking at this is to say: we have $V \subseteq \mathbb{A}^n$ and we have identified $T_x V$ as a subspace of $T_x \mathbb{A}^n$.

There is an intrinsic way to define the tangent space which does not depend on an embedding into affine space, using a bit more algebra. We will omit this here, and just state:

Lemma A.1. Let $V \subseteq \mathbb{A}^m$ and $W \subseteq \mathbb{A}^n$ be affine algebraic sets. Let $\varphi\colon V \to W$ be an isomorphism.

For any $x \in V$, if $y = \varphi(x) \in W$, then φ induces an isomorphism

$$d\varphi_x \colon T_x V \to T_y W.$$

In particular, $\dim T_x V = \dim T_y W$.

Outline proof. Choose polynomials $f_1, \dots, f_n$ such that $\varphi = (f_1, \dots, f_n)$. Define a linear map $k^m \to k^n$ by the matrix

$$\left(\frac{\partial f_i}{\partial X_j}\bigg|_x \right).$$

We define $d\varphi_x$ to be the restriction of this map to $T_x V$. (Recall that $T_x V$ is a subspace of k^m and $T_y W$ is a subspace of k^n.)

Using the chain rule for partial derivatives, one can check that:

(i) $d\varphi_x$ maps $T_x V$ into $T_y W$.
(ii) $d\varphi_x$ is independent of the choice of polynomials representing φ.
(iii) Because φ is an isomorphism of algebraic sets, $d\varphi_x$ is an isomorphism of vector spaces. □

Corollary A.2. The dimension of $T_x V$ is independent of the embedding of V into $\mathbb{A}^n$. Hence, whether a point of V is singular or not is independent of the embedding

In Lemma A.1 and Corollary A.2, we don't actually need an isomorphism between V and W themselves, just an isomorphism between open subsets (in other words, a birational map $V \dashrightarrow W$).

Lemma A.3. Let $V \subseteq \mathbb{A}^n$ and $W \subseteq \mathbb{A}^m$ be affine algebraic sets. Let $U_1 \subseteq V$ and $U_2 \subseteq W$ be open subsets, and let $\varphi\colon U_1 \to U_2$ be an isomorphism of quasi-projective varieties.

For any $x \in U_1$, if $y = \varphi(x) \in U_2$, then φ induces an isomorphism

$$d\varphi\colon T_xV \to T_yW.$$

In Lemma A.3, we talk about T_xV and T_yW rather than T_xU_1 and T_yU_2, even though the isomorphism is between U_1 and U_2. We have to do this because U_1 and U_2 need not be affine algebraic sets, so we have not yet defined their tangent spaces. Nevertheless, since U_1 is open in V, we intuitively expect $T_xU_1 = T_xV$ for $x \in U_1$ (the tangent space at x only depends on what happens near x, and "near x" there is no difference between U_1 and V). Hence the lemma makes sense.

The proof of Lemma A.3 is much the same as Lemma A.1 – we just write down fractions of polynomials which represent φ instead of polynomials. Because these fractions will have non-zero denominators at x, the calculus still works out fine.

Singular points on quasi-projective varieties.

Let V be a quasi-projective variety. By **??**, for any point $x \in V$ we can find an open subset $U \subseteq V$ which contains x and such that there is an isomorphism $\varphi\colon U \to W$ where W is an affine algebraic set.

We define the "dimension of the tangent space to V at x" to be the dimension of $T_{\varphi(x)W}$. Of course there are many choices for U, W and φ, but Lemma A.3 guarantees that $\dim T_{\varphi(x)}W$ will be independent of the choice. However, we do not get a specific vector space independent of choices which we can call T_xV. Of course, knowing the dimension is enough to determine a vector space up to isomorphism, but that is not as good as having a concrete vector space. There is a way of defining a concrete vector space which is "the tangent space" to V at x using local rings, but we will not do that here.

Knowing the dimension of tangent spaces is enough to define singular points. Just as in the case of an affine algebraic set, we define a **singular point** of V to be a point $x \in V$ where $\dim T_xV \neq \dim_x V$.

Singular locus of a variety.

Let V be a quasi-projective variety. The set of singular points of V is called the **singular locus** of V and denoted $\operatorname{Sing} V$.

The singular locus turns out to be a proper closed subset of V. We will prove this only for irreducible affine algebraic sets. (To generalise to irreducible quasi-projective varieties, simply use a cover by affine open subsets. Generalising to reducible varieties requires the fact that every point which lies in the intersection of two or more irreducible components is singular; we remarked that this requires more algebra to prove.)

Theorem A.4. Let V be an irreducible affine algebraic set. Then $\operatorname{Sing} V$ is a proper closed subset of V.

A key intermediate step in the proof of Theorem A.4, which is also interesting in its own right, is the fact that $\dim T_x V \geq \dim_x V$ for every point $x \in V$.

The singular locus of a hypersurface.

We begin by proving Theorem A.4 for a hypersurface.

Let $V \subseteq \mathbb{A}^n$ be a hypersurface and let f be a polynomial which generates $\mathbb{I}(V)$. Since f generates $\mathbb{I}(V)$, the tangent space $T_x V$ is just $\ker df_x$. In other words $T_x V$ is the kernel of a linear map $k^n \to k$, and so

$$\dim T_x V = n - 1 \text{ if } df_x \text{ is not the zero map;}$$
$$\dim T_x V = n \text{ if } df_x \text{ is the zero map.}$$

For any point $x \in V$, we have $\dim_x V = \dim V = n - 1$. Hence

$$\operatorname{Sing} V = \{x \in V : df_x = 0\}.$$

Going back to the definition of df_x, we can write this as

$$\operatorname{Sing} V = \left\{ x \in V : \left. \frac{\partial f}{\partial X_i} \right|_x = 0 \text{ for } i = 1, \dots, n \right\}.$$

This may look like the definition of singular points which you have seen before for curves in $\mathbb{A}^2$.

For each i, $\partial f / \partial X_i$ is a polynomial. Therefore:

Lemma A.5. For any hypersurface $V \subseteq \mathbb{A}^n$, $\operatorname{Sing} V$ is a closed subset of V.

Now we want to show that for a hypersurface V, $\operatorname{Sing} V \neq V$. This is a little harder in positive characteristic than in characteristic zero, so first we prove a lemma on derivatives in positive characteristic. Over a field of characteristic p, X^{ip} has derivative zero for any positive integer i. We prove that these span all the polynomials with zero derivative.

Lemma A.6. Let k be a field of characteristic $p > 0$. Let $f \in k[X]$ be a polynomial. If $\frac{\mathrm{d}f}{\mathrm{d}X} = 0$, then for every term of f, the exponent of X is a multiple of p, that is,

$$f = \sum_{i=0}^{d} a_{ip} X^{ip}.$$

Proof. Consider a term $a_j X^j$ in f. This term differentiates to $j a_j X^{j-1}$. No other term of f differentiates to a scalar multiple of X^{j-1}, so this term can never cancel with another term in $\mathrm{d}f/\mathrm{d}X$.

Hence if $\mathrm{d}f/\mathrm{d}X = 0$, then $j a_j = 0$ (in k) for every j. If j is not a multiple of p, then j is invertible in k so this forces $a_j = 0$. Thus only terms where j is a multiple of p can appear in f. □

Proposition A.7. If V is a non-empty hypersurface, then $\operatorname{Sing} V$ is strictly contained in V.

Proof. Assume for contradiction that $\operatorname{Sing} V = V$. Then $\partial f/\partial X_1, \dots, \partial f/\partial X_n$ are all zero on V.

Since f generates $\mathbb{I}(V)$, this implies that f divides $\partial f/\partial X_i$ for each i. But $\partial f/\partial X_i$ has strictly smaller X_i-degree than f. This forces $\partial f/\partial X_i = 0$ for each i (as a polynomial in $k[X_1, \dots, X_n]$).

Over a field of characteristic zero, this implies that f is constant. But then V would be empty, contradicting the hypothesis.

Over a field of characteristic $p > 0$, by Lemma A.6, the fact that $\partial f/\partial X_i = 0$ implies that very term of f must have its X_i-exponent being a multiple of p. Since this holds for all i, each term of f is a p-th power (the constant in the term must be a p-th power because k is algebraically closed).

But the binomial expansion implies that

$$(a + b)^p = a^p + b^p$$

over a field of characteristic p. So if every term of f is a p-th power, then f itself is a p-th power. But then the ideal generated by f is not a radical ideal. Via the Nullstellensatz, this contradicts the assumption that f generates $\mathbb{I}(V)$. ☐

The singular locus of an irreducible variety.

Lemma A.8. Let $V \subseteq \mathbb{A}^n$ be an affine algebraic set. For any integer d, the set

$$\Sigma_d(V) = \{x \in V : \dim T_x V > d\}$$

is a closed subset of V.

Proof. Choose polynomials $f_1, \dots, f_m$ which generate $\mathbb{I}(V)$. Recall that

$$T_x V = \bigcap_{i=1}^{m} \ker d(f_i)_x.$$

In other words, $T_x V$ is the kernel of the matrix

$$M_x = \left(\frac{\partial f_i}{\partial X_j} \bigg|_x \right)_{ij}$$

which represents a linear map $k^n \to k^m$

By the rank–nullity theorem, $\dim T_x V$ is equal to $n - \operatorname{rk} M_x$. Hence

$$\Sigma_d(V) = \{x \in V : \operatorname{rk} M_x < n - d\}.$$

By linear algebra, $\operatorname{rk} M_x < n - d$ is equivalent to: every $(n-d) \times (n-d)$ submatrix of M_x has determinant zero.

The determinant of a submatrix of M_x is a polynomial, hence this gives us polynomial equations defining $\Sigma_d(V)$. ☐

Lemma A.9. Let V be an irreducible affine algebraic set. Then the non-singular points of V are dense in V.

Proof. By Proposition 11.5 from lectures, V is birational to a hypersurface $H \subseteq \mathbb{A}^{d+1}$. By Lemma 16.1 we can find non-empty open sets $U \subseteq V$ and $J \subseteq H$ such that there is an isomorphism $\varphi \colon U \to J$.

By Lemma A.5 and Proposition A.7, the non-singular points of H form a non-empty open subset $H_{ns} \subseteq H$. Since H is irreducible, H_{ns} must intersect J.

Since φ is continuous, $A = \varphi^{-1}(H_{ns} \cap J)$ is an open subset of U. Since $H_{ns} \cap J \neq \emptyset$ and since φ is surjective onto J, A is non-empty. Since V is irreducible, we conclude that A is dense in V.

For any $x \in A$, let $y = \varphi(x) \in H_{ns} \cap J$. Lemma A.3 tells us that T_xV is isomorphic to T_yH. Since $y \in H_{ns}$, we have $\dim T_yH = \dim H$. Thus

$$\dim T_xV = \dim T_yH = \dim H = \dim V$$

so V is non-singular at x. □

Lemma A.10. Let V be an irreducible affine algebraic set. For every $x \in V$, $\dim T_xV \geq \dim_x V$.

Proof. Since V is irreducible, $\dim_x V = \dim V$ for every $x \in V$ so we can work with $\dim V$ instead of $\dim_x V$. (The lemma is true for reducible V as well, but we do not have the tools to prove it when $\dim_x V$ is not constant.)

Let $d = \dim V$ and consider the set $\Sigma_{d-1}(V)$ as in Lemma A.8. By Lemma A.8, $\Sigma_{d-1}(V)$ is closed. Every non-singular point of V is in $\Sigma_{d-1}(V)$, so Lemma A.9 implies that $\Sigma_{d-1}(V)$ is dense in V.

Since $\Sigma_{d-1}(V)$ is closed and dense in V, we conclude that $\Sigma_{d-1}(V) = V$. □

Theorem A.11. Let V be an irreducible affine algebraic set. Then $\operatorname{Sing} V$ is a proper closed subset of V.

Proof. By Lemma A.10, $x \in V$ is a singular point if and only if $\dim T_xV > \dim_x V$. Again since V is irreducible, we can replace $\dim_x V$ by $\dim V$.

Thus

$$\operatorname{Sing} V = \Sigma_d(V)$$

where $d = \dim V$. So Lemma A.8 tells us that $\operatorname{Sing} V$ is closed in V. Lemma A.9 implies that $\operatorname{Sing} V$ is properly contained in V. □

Problem sheet.

There is a problem sheet (sheet 6) on singular points available on Blackboard, which I encourage you to do. I will provide solutions to it near the end of the holidays.