Math 121. Transcendence bases

1. Introduction

Let L/k be an arbitrary extension of fields. A subset $S \subset L$ is called *algebraically independent* over k if for all finite sets of elements $a_1, \ldots, a_n \in S$, no nonzero $f \in k[X_1, \ldots, X_n]$ satisfies $f(a_1, \ldots, a_n) = 0$ in L. Equivalently, this says that the surjective k-algebra map $k[X_1, \ldots, X_n] \to k[a_1, \ldots, a_n]$ onto the k-subalgebra of L generated by the a_i's has vanishing kernel, or equivalently is an isomorphism onto that k-subalgebra. This isomorphism of domains then uniquely extends to an isomorphism

$$k(X_1, \ldots, X_n) \simeq k(a_1, \ldots, a_n)$$

of fields over k for all finite subsets $\{a_1, \ldots, a_n\} \subset S$. (If S is a subset of L that is not algebriacally independent over k then we say it is *algebraically dependent* over k.) As we shall see below, there are close analogies between the properties of algebraically independent subsets of a field extension L/k and linearly independent subsets of a vector space V over k.

We say that an algebraically independent subset $S \subset L$ over k is *maximal* if it is not a proper subset of another algebraically independent set $S' \subset L$. It is equivalent to say that the extension $L/k(S)$ over the subfield of L generated by S over k is an *algebraic* extension, or equivalently that for all $a \in L$ there is some nonzero $f \in k(S)[X]$ such that $f(a) = 0$. Indeed, if such an f exists (for a given $a \in L$) then it involves just a finite set of elements $s_1, \ldots, s_n \in S$ in its coefficients (relative to X), so by scaling through against a nonzero common denominator in $k[s_1, \ldots, s_n]$ we may arrange that

$$f \in k[s_1, \ldots, s_n][X] = k[s_1, \ldots, s_n, X].$$

Hence, the relation $f(a) = 0$ with $\deg_X(f) > 0$ shows that $S \cup \{a\}$ is *not* algebraically independent over k; i.e., S is maximal. Conversely, suppose S is maximal as an algebraically independent set in L over k. We want to show that L is algebraic over $k(S)$. For any $a \in L - S$ the maximality implies that $S \cup \{a\}$ is algebraically dependent, so for some $s_1, \ldots, s_n \in S$ there is a nonzero $f \in k[X_1, \ldots, X_n, X_{n+1}]$ such that $f(s_1, \ldots, s_n, a) = 0$. But f must involve X_{n+1}, as otherwise we would contradict the algebraic independence of $\{s_1, \ldots, s_n\} \subset S$ over k. Hence, we may write

$$f = h_d X_{n+1}^d + \cdots + h_1 X + h_0 \tag{1}$$

for $h_j \in k[X_1, \ldots, X_n]$ with $d > 0$ and $h_d \neq 0$. Then

$$0 = f(s_1, \ldots, s_n, a) = \sum h_j(s_1, \ldots, s_d) a^j,$$

and $h_d(s_1, \ldots, s_n) \neq 0$ in L because $h_d \neq 0$ and $\{s_1, \ldots, s_n\}$ is algebraically independent over k.

In class we stated the following result, the proof of which will occupy the rest of the handout.

Theorem 1.1. *Let L/k be a finitely generated extension, with $\{a_1, \ldots, a_n\}$ a finite subset of L that generates L over k. Then:*

1. *Every algebraically independent subset of $\{a_1, \ldots, a_n\}$ (relative to k) that is maximal as such is also maximal as an algebraically independent subset of L over k.*
2. *Every algebraically independent subset of L is finite, all such subsets live inside maximal ones, and the maximal ones all have the same size.*
3. *Every subfield $F \subset L$ is finitely generated over k. In particular, the subfield $k' \subset L$ consisting of all $a \in L$ algebraic over k is of finite degree over k.*

The common finite size of all maximal algebraically independent subsets of L over k is called the *transcendence degree* of L over k, and such maximal subsets are called *transdencence bases* of L over k. By design, if $B \subset L = k(a_1, \ldots, a_n)$ is a transcendence basis over k then every element of L not in B is algebraic over $k(B)$ due to the maximality condition, as are elements of B (elementary). Hence, each of the a_i's is $k(B)$-algebraic, and hence L is $k(B)$-finite. That is the general structure of L/k: a finite extension of $k(B) \simeq k(Y_1, \ldots, Y_d)$ where d is the transcendence degree of L over k.

Beware that the finite generatedness in (3) is very specific to the case of subfields of fields: it breaks down completely in the setting of subalgebras of k-algebras. That is, there are examples (due to Nagata) of sufficiently large n so that there is a k-subalgebra $A \subset k[X_1, \ldots, X_n]$ that is *not* finitely generated as a k-algebra. (Google "Hilbert's 14th Problem" for more on this.)

2. Proof of Theorem

The proof of (1) is elementary, as follows. Letting $S \subset \{a_1, \ldots, a_n\}$ be maximal as a subset of $\{a_1, \ldots, a_n\}$ algebraically independent over k, we need to show that S is also maximal as an algebraically independent subset of L over k. By relabeling we may assume $S = \{a_1, \ldots, a_r\}$ for some $0 \leq r \leq n$. The maximality hypothesis gives that $S \cup \{a_i\}$ is algebraically dependent for all $r < i \leq n$. By the same considerations as with f in (1), since S is algebraically independent over k it follows that a_i is algebraic over $k(S)$ for each $i > r$. Hence, the field $L = k(S)(a_{r+1}, \ldots, a_n)$ is algebraic over $k(S)$, so for *all* $a \in L - S$ there is some nonzero $f \in k(S)[X]$ such that $f(a) = 0$. Clearing denominators from $k[S] - \{0\}$ as in the study of (1), we can modify the choice of f so that $f \in k[S][X] = k[X_1, \ldots, X_n, X]$. Then the condition $f(a) = 0$ expresses algebraic dependence of $S \cup \{a\}$, establishing the desired maximality of S as an algebraically independent subset of L over k. This completes the proof of (1).

Now we turn to the proof of (2), for which we may assume L/k is not algebraic (as otherwise there is nothing to do). This is a slightly tricky induction argument, modelled on the proof that dimension of a vector space is well-defined. Recall from (1) that we have built some maximal algebraically independent sets in L over k, namely algebraically independent subsets of $\{a_1, \ldots, a_n\}$ that are maximal as such.

We shall now prove the following general assertion, with L/k allowed to vary across *all* finitely generated extensions of fields: if $\{y_1, \ldots, y_r\}$ is *any* maximal algebraically independent subset of L over k then for *every* algebraically independent subset $S \subset L$ over k (not assumed to be maximal!) necessarily $\#S \leq r$ and S lies in a maximal one of size r. This will be shown by induction on $r \geq 0$, varying across all finitely generated extensions of fields L/k. The case $r = 0$ is trivial (why?), so we may assume $r > 0$ and that the result is known whenever there is a maximal algebraically independent subsets of size $< r$. We may also focus attention just on finite S (as S is finite with size $\leq r$ if and only if the same holds for all finite subsets of S), so enumerate S as $\{x_1, \ldots, x_m\}$. We may and do assume $m > 0$ (so each x_i is transcendental over k).

By maximality of $\{y_1, \ldots, y_r\}$, the collection $\{y_1, \ldots, y_r, x_1\}$ is algebraically dependent over k (this is trivial if $x_1 \in \{y_1, \ldots, y_r\}$). Thus, there is a nonzero $f \in k[Y_1, \ldots, Y_r, X_1]$ such that $f(y_1, \ldots, y_r, x_1) = 0$. This f *must* involve X_1 (as $\{y_1, \ldots, y_r\}$ is algebraically independent over k), and it *must* involve some Y_i (as x_1 is transcendental over k!), so by relabeling we may assume f involves Y_1.

Consider the subset $\Sigma = \{x_1, y_2, \ldots, y_r\}$ in L. Using f, which involves y_1, we see that y_1 is algebraic over $k(\Sigma) := k(x_1, y_2, \ldots, y_r)$, so $k(y_1, \ldots, y_r)$ is algebraic over $k(\Sigma)$, yet L is algebraic over $k(y_1, \ldots, y_r)$ and hence by transitivity of algebraicity we see that L is algebraic over $k(\Sigma)$. Thus, if Σ is algebraically *independent* over k then it must be maximal as such (why?). Let's rule out the possibility that Σ is *not* algebraically indendepent over k.

Assume such failure, so by applying (1) to $k(\Sigma)$ and its finite generating set Σ, we would obtain a subset $\Sigma_0 \subset \Sigma$ with size $r_0 < r$ such that Σ_0 is a maximal algebraically indepdent subset of $k(\Sigma)$. Hence, $k(\Sigma)$ is a *finite* extension of $k(\Sigma_0)$. But $L/k(\Sigma)$ is algebraic too, so then L is algebraic over $k(\Sigma_0)$. Then the algebraically independent subset $\Sigma_0 \subset L$ over k with size $r_0 < r$ could be fed into the inductive hypothesis (!), forcing *all* algebraically independent subsets of L over k to have size $\leq r_0$, contradicting that $\{y_1, \ldots, y_r\}$ is algebraically independent in L over k with size $r > r_0$.

We conclude that Σ is indeed algebraically independent over k. In other words, by replacing $\{y_1, \ldots, y_r\}$ with Σ, we have passed to the case that $y_1 = x_1$. Now comes the clever trick: consider L as an extension of the field $k(x_1)$! The subset $\{y_2, \ldots, y_r\}$ of L is algebraically independent over $k(x_1)$ because if $h \in k(x_1)[Y_2, \ldots, Y_r]$ is nonzero with $h(y_2, \ldots, y_r) = 0$ then scaling through by a suitable common deniminator in $k[x_1] - \{0\}$ in its $k(x_1)$-coefficients brings us to the case $h \in k[x_1][Y_2, \ldots, Y_r]$, so we would obtain a nonzero $H \in k[X_1, Y_2, \ldots, Y_r]$ such that $H(x_1, y_2, \ldots, y_r) = 0$. But that contradicts the algebraic independence of $\{x_1, y_2, \ldots, y_r\} = \Sigma$ over k. Hence, $L/k(x_1)$ has a maximal algebraically independent subset $\{y_2, \ldots, y_r\}$ of size $r - 1 < r$, so by induction (!) it follows that *every* algebraically independent subset of L over $k(x_1)$ has size at most $r - 1$ and lies inside an algebraically independent subset of L over $k(x_1)$ with size $r - 1$.

By the exact same reasoning as just given with $\{x_1, y_2, \ldots, y_r\} = \Sigma$, the subset $\{x_2, \ldots, x_m\}$ in L with size $m-1$ is algebraically independent over $k(x_1)$ too. Aha, but then $m-1 \leq r-1$, so $m \leq r$, and likewise the subset $\{x_2, \ldots, x_m\}$ lies inside an algebraically independent subset $\{x_2, \ldots, x_r\}$ of L over $k(x_1)$. The equality

$$k[X_1, \ldots, X_r] = k[X_1][X_2, \ldots, X_r]$$

then shows that $\{x_1, x_2, \ldots, x_r\}$ is algebraically independent over k. (Indeed, if $h \in k[X_1, \ldots, X_r]$ is nonzero and $h(x_1, \ldots, x_r) = 0$ then h must involve some of $X_2, \ldots, X_r$ or else we contradict the transcendence of x_1 over k, and hence h instead contradicts the algebraic independence of $\{x_2, \ldots, x_r\}$ over $k(x_1)$.) Since we have already shown that *every* (finite) algebraically independent subset of L over k has size at most r, it follows thar the algebraically independent subset $\{x_1, \ldots, x_r\}$ of L over k is maximal as such. This completes the induction on r.

With our induction finished, we conclude that any algebraically independent subset of L over k is finite and lies in a maximal one, that the size of any single maximal one is an upper bound on the size of all of them. Hence, the sizes of any two maximal ones must be the *same* (as each size is an upper bound on the other, due to what we just proved by induction). This finally completes the proof of (2). In particular, the terminology *transcendence degree* and *transcendence basis* for a general finitely generated extension of fields L/k now makes sense.

It remains to prove (3): if L/k is finitely generated then *every* subfield $F \subset L$ is finitely generated over k. Observe that every subset of F algebraically independent over k may be viewed as such inside L and hence is *finite*, with size at most the transcendence degree of L over k. Thus, we may choose an algebraically independent subset $\{x_1, \ldots, x_d\}$ of F over k that is of maximal size as such, so for *every* $a \in F$ necessarily a is algebraic over $k(x_1, \ldots, x_d) = k(X_1, \ldots, X_d)$. Indeed, such algebraicity is clear if $a = x_i$ for some i, and otherwise the subset $\{x_1, \ldots, x_d, a\}$ of F is *algebraically dependent* over k due to the maximality hypothesis on $\{x_1, \ldots, x_d\}$, where a nonzero $h \in k[X_1, \ldots, X_d, X_{d+1}]$ satisfying

$$h(x_1, \ldots, x_d, a) = 0$$

must involve X_{d+1} (otherwise we would contradict algebraic independence of $\{x_1, \ldots, x_d\}$ over k). By writing h in $k(X_1, \ldots, X_d)[X_{d+1}]$ with positive X_{d+1}-degree, we see that a is algebraic over $k(x_1, \ldots, x_d)$ as claimed. Hence, we have shown that the extension of fields $F/k(x_1, \ldots, x_d)$ is *algebraic*.

Our aim is to show that F is finitely generated over k, so it suffices to show that it is of finite degree over $k(x_1, \ldots, x_d)$. It is harmless now to rename $k(x_1, \ldots, x_d)$ as the ground field k (!), so in other words we have reduced to proving (3) in the special case that F is *algebraic* over k, Hence, the proof of the Theorem is reduced to the following lemma of independent interest:

Lemma 2.1. *Let L/k be a finitely generated extension of fields. The subextension $k' \subset L$ of all $a \in L$ algebraic over k is of finite degree over k.*

Proof. We begin by induction on the new invariant we have created, the transcendence degree of L/k (allowing ourselves to vary over *all* L/k!). The case of transcendence degree 0 is trivial, as in such cases the finitely generated extension L/k is an algebraic extension (why?), forcing it to be of finite degree (so all subextensions are k-finite, as desired).

Next suppose that the transcendence degree is $n > 0$ and that the result is known for all finitely generated field extensions with transcendence degree $< n$. Let $\{x_1, \ldots, x_n\}$ be a transcendence basis of L/k, so the finitely generated algebraic extension $L/k(x_1, \ldots, x_n)$ is finite and L has transdencence degree $n-1$ over the subfield $K := k(x_1)$ (why?). By the inductive hypothesis, the subfield $K' \subset L$ of elements $a \in L$ algebraic over K is of *finite* degree over $K = k(x_1)$, and visibly $k' \subset K'$. Thus, K'/k is a finitely generated extension with transcendence degree 1 (why?) and we have just seen that $k' \subset K'$, so k' sits in relation to K'/k as it does to L/k. In other words, we may rename K' as L to reduce to the case that L/k has transcendence degree 1.

Now pick $x \in L$ transcendental over k, so the finitely generated extension $L/k(x)$ is algebraic and hence of *finite* degree. We want to prove that k' is k-finite, and do so we will induct on $[L : k(x)]$! When this degree is 1 then $L = k(x) \simeq k(X)$, in which $k' = k$ (why?). Thus, we may assume $[L : k(x)] > 1$. We may also assume $k' \neq k$ (or else there is nothing to prove), so choose $a \in k' - k$. Note that x is transcendental over the finite extension $k(a)$ of k (as otherwise it would be algebraic over $k(a)$ and hence algebraic over k, an absurdity). Hence, L is a finitely generated extension of $k(a)$ in which x remains a transcendence basis, and

$$[L : k(a)(x)] = [L : k(x)]/[k(a)(x) : k(x)]$$

with $k(a)(X)$ a finite *nontrivial* extension of $k(X)$ (why?), so $[k(a)(x) : k(x)] > 1$. Thus, $[L : k(a)(x)] < [L : k(x)]$, so by our new inductive process applied to $L/k(a)$ we conclude that the subfield of elements of L algebraic over $k(a)$ has finite degree over $k(a)$. But this subfield of L is the same as k' (why?), so we conclude that $[k' : k(a)]$ is finite. Hence, $[k' : k]$ is finite as desired. ■