

MATH 210C. ODE

As applications of connectivity arguments and the contraction mapping theorem in complete metric spaces (in fact, an infinite-dimensional complete normed vector space: continuous functions on a compact interval, endowed with the sup-norm), in this handout we prove and illustrate the classical local existence and uniqueness theorem for a wide class of first-order ordinary differential equations (ODE).

As a special case this includes linear ordinary differential equations of *any* order, and in this linear case we prove a remarkable *global* existence theorem (and give counterexamples in the non-linear case). We end by illustrating the surprisingly subtle problems associated with variation of auxiliary parameters and initial conditions in an ODE (a fundamental issue in applications to physical problems, where parameters and initial conditions must be allowed to have error); this effect of variation is a rather non-trivial problem that underlies the theory of integral curves which we address in a separate handout (and is absolutely basic in differential geometry).

You should at least skim over this handout (reading the examples, if not the proofs) in order to appreciate the elegance and power of the basic theory of ODE in terms of its applicability to rather interesting problems.

1. MOTIVATION

In classical analysis and physics, a fundamental topic is the study of systems of linear ordinary differential equations (ODE). The most basic example is as follows. Let $I \subseteq \mathbf{R}$ be a non-trivial interval (i.e., not a point but possibly half-open/closed, perhaps unbounded, etc.) and choose smooth functions f, a_0, \dots, a_{n-1} on I as well as constants $c_0, \dots, c_{n-1} \in \mathbf{R}$, with $n \geq 1$. We seek to find a smooth function $u : I \rightarrow \mathbf{R}$ satisfying

$$u^{(n)} + a_{n-1}u^{(n-1)} + \dots + a_1u' + a_0u = f$$

on I , subject to the “initial conditions” $u^{(j)}(t_0) = c_j$ for $0 \leq j \leq n-1$ at $t_0 \in I$. (For example, in physics one typically meets 2nd-order equations, in which case the initial conditions on the value and first derivative are akin to specifying position and velocity at some time.) We call this differential equation *linear* because the left side depends \mathbf{R} -linearly on u (as the operations of differentiation and multiplication by a smooth function are \mathbf{R} -linear self-maps of $C^\infty(I)$).

Actually, for more realistic examples we want more: we should let u be vector-valued rather than just \mathbf{R} -valued. Let V be a finite-dimensional vector space over \mathbf{R} (classically, \mathbf{R}^N). The *derivative* $u' : I \rightarrow V$ of a map $u : I \rightarrow V$ is defined in one of three equivalent ways: use the habitual “difference-quotient” definition (which only requires the function to have values in a normed \mathbf{R} -vector space), differentiate componentwise with respect to a choice of basis (the choice of which does not impact the derivative being formed; check!), or take $u'(t) = Du(t)(\partial|_t) \in T_{u(t)}(V) \simeq V$ for ∂ the standard “unit vector field” on the interval $I \subseteq \mathbf{R}$.

For C^∞ mappings $f : I \rightarrow V$ and $A_j : I \rightarrow \text{Hom}(V, V)$ ($0 \leq j \leq n-1$), as well as vectors $C_j \in V$ ($0 \leq j \leq n-1$), we seek a C^∞ mapping $u : I \rightarrow V$ such that

$$(1.1) \quad u^{(n)}(t) + (A_{n-1}(t))(u^{(n-1)}(t)) + \dots + (A_1(t))(u'(t)) + (A_0(t))(u(t)) = f(t)$$

for all $t \in I$, subject to the initial conditions $u^{(j)}(t_0) = C_j$ for $0 \leq j \leq n-1$. If $V = \mathbf{R}^N$ and we write $u(t) = (u_1(t), \dots, u_N(t))$ then (1.1) is really a system of N linked n th-order linear ODE's in the functions $u_i : I \rightarrow \mathbf{R}$, and the n initial conditions $u^{(j)}(t_0) = C_j$ in $V = \mathbf{R}^N$ are Nn conditions $u_i^{(j)}(t_0) = c_{ij}$ in \mathbf{R} , where $C_j = (c_{1j}, \dots, c_{Nj}) \in \mathbf{R}^N = V$. For most problems in physics and engineering we have $n \leq 2$ with $\dim V$ large.

The “ideal theorem” is that (1.1) with its initial conditions should exist and be unique, but we want more! Indeed, once one has an existence and uniqueness theorem in hand, it is very natural to ask: how does the solution depend on the initial conditions $u^{(j)}(t_0) = C_j$? That is, as we vary the C_j 's, does the solution exhibit smooth dependence on these initial conditions? And what if the A_j 's depend continuously (or smoothly) on some auxiliary parameters not present in the initial conditions (such as some friction constants or other input from the surrounding physical setup)? That is, if $A_j = A_j(t, z)$ for z in an auxiliary space, does the solution $t \mapsto u(t, z)$ for each z depend “as nicely” on z as do the A_j 's? How about dependence on the *combined* data (t, z, C_1, \dots, C_n) ?

It has been long recognized that it is hopeless to try to study differential equations by explicitly exhibiting solutions. Though clever tricks (such as integrating factors and separation of variables) do find some solutions in special situations, one usually needs a general theory to predict the “dimension” of the space of solutions and so to tell us whether there may be more solutions remaining to be found. Moreover, we really want to understand *properties* of solutions: under what conditions are they unique, and if so do they exist for all time? If so, what can be said about the long-term behavior (growth, decay, oscillation, etc.) of the solution? Just as it is unwise to study the properties of solutions to polynomial equations by means of explicit formulas, in the study of solutions to differential equations we cannot expect to get very far with explicit formulas. (Though in the rare cases that one can find an explicit formula for some or all solutions it can be helpful.)

The first order of business in analyzing these questions is to bring the situation into more manageable form by expressing the given n th-order problem (1.1) as a first-order differential equation. This immensely simplifies the notational complexity and thereby helps us to focus more clearly on the essential structure at hand. Here is an illustration of the classical trick known to all engineers for reducing linear initial-value problems to the first-order case; the idea is to introduce an *auxiliary vector space*.

Example 1.1. Consider the equation

$$(1.2) \quad y'' + P(x)y' + Q(x)y = R(x)$$

on an interval I , with $P, Q, R \in C^\infty(I)$. Say we impose the conditions $y(t_0) = c_0$ and $y'(t_0) = c_1$ for some $c_0, c_1 \in \mathbf{R}$. We shall rephrase this as a first-order equation with values in \mathbf{R}^2 . The idea is to study the derivative of the vector $u(x) = (y(x), y'(x)) \in \mathbf{R}^2$. Define

$$A(x) = \begin{pmatrix} 0 & 1 \\ -Q(x) & -P(x) \end{pmatrix}, \quad f(x) = \begin{pmatrix} 0 \\ R(x) \end{pmatrix}, \quad c = \begin{pmatrix} c_0 \\ c_1 \end{pmatrix}.$$

For a mapping $u : I \rightarrow V = \mathbf{R}^2$ given by $u(x) = (u_0(x), u_1(x))$, consider the first-order initial-value problem in \mathbf{R}^2 :

$$(1.3) \quad u'(x) = (A(x))(u(x)) + f(x), \quad u(t_0) = c.$$

Unwinding what this says on components, one gets the equivalent statements $u'_0 = u_1$ and u_0 satisfies the original differential equation (1.2) with $u_0(t_0) = c_0$ and $u'_0(t_0) = c_1$. Hence, our second-order problem for an \mathbf{R} -valued function u_0 with 2 initial conditions (on u_0 and u'_0 at t_0) has been reformulated as a first-order problem (1.3) for an \mathbf{R}^2 -valued mapping u with one initial condition (at t_0).

To apply this trick in the general setup of (1.1), let W be the n -fold direct sum V^n , and let $A : I \rightarrow \text{Hom}(W, W)$ be defined by

$$A(t) : (v_0, \dots, v_{n-1}) \mapsto (v_1, v_2, \dots, v_{n-1}, -((A_{n-1}(t))(v_{n-1}) + \dots + (A_0(t))(v_0))) \in V^n = W,$$

so A is clearly smooth. Also, let $F : I \rightarrow W = V^n$ be the smooth mapping $F(t) = (0, \dots, 0, f(t))$, and let $C = (C_0, \dots, C_{n-1}) \in W$. A mapping $u : I \rightarrow W$ is given by

$$u(t) = (u_0(t), u_1(t), \dots, u_{n-1}(t))$$

for mappings $u_j : I \rightarrow V$, and $u : I \rightarrow W$ is smooth if and only if all $u_j : I \rightarrow V$ are smooth (why?). By direct calculation, the first-order equation

$$(1.4) \quad u'(t) = (A(t))(u(t)) + F(t)$$

with initial condition $u(t_0) = C$ is the *same* as requiring two things to hold:

- $u_j = u_0^{(j)}$ for $1 \leq j \leq n-1$ (so u_1, \dots, u_{n-1} are redundant data when we know u_0),
- u_0 satisfies the initial-value problem

$$(1.5) \quad u_0^{(n)}(t) + (A_{n-1}(t))(u_0^{(n-1)}(t)) + \dots + (A_0(t))(u_0(t)) = f(t), \quad u_0^{(j)}(t_0) = C_j \quad (0 \leq j \leq n-1)$$

that is (1.1) by another name.

In other words, our original n th-order linear V -valued problem (1.5) with n initial conditions in V is equivalent to the *first-order* linear W -valued problem (1.4) with *one* initial condition. This is particularly striking in the classical case $V = \mathbf{R}$: an n th-order linear ODE for an \mathbf{R} -valued function can be recast as first-order linear ODE provided we admit *vector-valued* problems (for $V = \mathbf{R}$ we have $W = \mathbf{R}^n$).

For these reasons, in the theory of linear ODE's the "most general" problem is the vector-valued problem

$$u'(t) = (A(t))(u(t)) + f(t), \quad u(t_0) = v_0$$

for smooth $f : I \rightarrow V$ and $A : I \rightarrow \text{Hom}(V, V)$ and a point $v_0 \in V$, where V is permitted to be an *arbitrary* finite-dimensional vector space. (We could write $V = \mathbf{R}^N$, but things are cleaner if we leave V as an abstract finite-dimensional vector space.) Note in particular that the classical notion of "system of n th-order linear ODE's" (expressing derivatives of each of several \mathbf{R} -valued functions as linear combinations of all of the functions, with coefficients that depend smoothly on t) is just a single first-order vector-valued ODE in disguise. Hence, we shall work exclusively in the language of *first-order vector-valued ODE's*. Observe also that we may replace "smooth" with C^p in everything that went before (for $0 \leq p \leq \infty$), with the caveat that we cannot expect the solution to have class of differentiability any better than C^{p+1} .

2. THE LOCAL EXISTENCE AND UNIQUENESS THEOREM

For applications in differential geometry and beyond, it is important to allow for the possibility of non-linear ODE's (e.g, non-linear expressions such as u^2 or $(u')^2$ showing up in the equation for \mathbf{R} -valued u). The engineer's trick of introducing auxiliary functions whose first derivatives are set equal to other functions still reduces problems to the first-order case (via an auxiliary vector space).

What do we really mean by "non-linear first-order ODE"? Returning to the linear case considered above, we can rewrite the equation $u'(t) = (A(t))(u(t)) + f(t)$ as follows:

$$u'(t) = \phi(t, u(t))$$

where $\phi : I \times V \rightarrow V$ is the smooth mapping $\phi(t, v) = (A(t))(v) + f(t)$. (This is a C^p mapping if A and f are C^p in t .)

Generalizing, we are interested in C^p solutions to the "initial-value problem"

$$(2.1) \quad u'(t) = \phi(t, u(t)), \quad u(t_0) = v_0$$

where $\phi : I \times U \rightarrow V$ is a C^p mapping for an open set $U \subseteq V$ and $v_0 \in U$ is a point. Of course, it is implicit that a solution $u : I \rightarrow V$ has image contained in U so that the expression $\phi(t, u(t))$ makes sense for all $t \in I$. Such a solution u is certainly of class C^0 , and in general if it is of class C^r with $r \leq p$ then by (2.1) we see u' is a composite of C^r mappings and hence is C^r . That is, u is C^{r+1} . Thus, by induction from the case $r = 0$ we see that a solution u is necessarily of class C^{p+1} .

The key point of the above generalization is that $\phi(t, \cdot) : U \rightarrow V$ need not be the restriction of an affine-linear map $V \rightarrow V$ (depending on t). Such non-linearity is crucial in the study of vector flow on manifolds, as we shall see.

The first serious theorem in the theory of differential equations is the *local* existence/uniqueness theorem for equations of the form (2.1):

Theorem 2.1. *Let $\phi : I \times U \rightarrow V$ be a C^p mapping with $p \geq 1$, and choose $v_0 \in U$ and $t_0 \in I$. There exists a connected open subset $J \subseteq I$ around t_0 and a differentiable mapping $u : J \rightarrow U$ satisfying (2.1) (so u is C^{p+1}). Moreover, such a solution is uniquely determined on any J where it exists.*

Before we prove the theorem, we make some observations and look at some examples. The existence aspect of the theorem is local at t_0 , but the uniqueness aspect is more global: on any J around t_0 where a solution to the initial-value problem exists, it is uniquely determined. In particular, if $J_1, J_2 \subseteq I$ are two connected open neighborhoods of t_0 in I on which a solution exists, the solutions must agree on the interval neighborhood $J_1 \cap J_2$ (by uniqueness!) and hence they “glue” to give a solution on $J_1 \cup J_2$. In this way, it follows from the uniqueness aspect that there exists a maximal connected open subset $J_{\max} \subseteq I$ around t_0 (depending on u, ϕ , and v_0) on which a solution exists (containing all other such connected opens). However, in practice it can be hard to determine J_{\max} ! (We give some examples to illustrate this below.) Also, even once the local existence and uniqueness theorem is proved, for applications in geometry we need to know more: does the solution u exhibit C^p -dependence on the initial condition v_0 , and if ϕ depends “nicely” (continuously, or better) on “auxiliary parameters” then does u exhibit just as nice dependence on these parameters? (See Example 3.6.) Affirmative answers will be given at a later time.

Remark 2.2. Although consideration of $\phi(t, u(t))$ permits equations with rather general “non-linear” dependence on u , we are still requiring that u' only intervene linearly in the ODE. Allowing non-linear dependence on higher derivatives is an important topic in advanced differential geometry, but it is not one we shall need to confront in our study of elementary differential geometry.

A fundamental dichotomy between the first-order *linear* ODE’s (i.e., the case when $\phi(t, \cdot)$ is the restriction of an affine-linear self-map $v \mapsto A(t)v + f(t)$ of V for each $t \in V$) and the general case is that in Theorem 3.1 we will prove a *global* existence theorem in the linear case (for which we may and do always take $U = V$): the initial-value problem will have a solution across the *entire* interval I (i.e., $J_{\max} = I$ in such cases). Nothing of the sort holds in the general non-linear case, even when $U = V$:

Example 2.3. If we allow non-linear intervention of derivatives then uniqueness fails. Consider the \mathbf{R} -valued problem $(u')^2 = t^3$ on $(-1, 1)$. This has exactly two solutions (not just one!) on $(0, 1)$ with any prescribed initial condition at $t_0 = 1/2$. The solutions involve $\pm\sqrt{t}$ and so continuously extend to $t = 0$ without differentiability there. In other words, even without “blow-up” in finite time, solutions to initial-value ODE’s that are non-linear in the derivative may not exist across the entire interval even though they admit limiting values at the “bad” point. The non-uniqueness of this example shows that the linearity in u' is crucial in the uniqueness aspect of Theorem 2.1.

Since this example does not fit the paradigm $u'(t) = \phi(t, u(t))$ in the local existence and uniqueness theorem, you may find this sort of non-linear example to be unsatisfying. The next example avoids this objection.

Example 2.4. Even for first-order initial-value problems of the form $u'(t) = \phi(t, u(t))$, for which there is uniqueness (given $u(t_0)$), non-linearity in the second variable of ϕ can lead to the possibility that the solution does not exist across the entire interval (in contrast with what we have said earlier in the linear case, and will prove in Theorem 3.1). For example, consider

$$u' = 1 + u^2, \quad u(0) = 0.$$

The unique local solution near $t_0 = 0$ is $u(t) = \tan(t)$, and as a solution it extends to $(-\pi/2, \pi/2)$ with blow-up (in absolute value) as $t \rightarrow \pm\pi/2$. It is not at all obvious from the shape of this particular initial-value problem that the solution fails to propagate for all time. Hence, we see that the problem of determining $J_{\max} \subseteq I$ in any particular case can be tricky. The blow-up aspect of this example is not a quirk: we will prove shortly that failure of the image of u near an endpoint τ of J in I to be contained in a compact subset of U as $t \rightarrow \tau$ is the only obstruction to having $J_{\max} = I$ for initial-value problems of the form in Theorem 2.1.

Let us now turn to the proof of Theorem 2.1.

Proof. We will prove the local existence result and weaker version of uniqueness: any two solutions agree *near* t_0 in I , where “near” may depend on the two solutions. This weaker form is amenable to shrinking I around t_0 . At the end we will return to the proof of global uniqueness (without shrinking I) via connectivity considerations.

Fix a norm on V . Let $B = \overline{B}_{2r}(v_0)$ be a compact ball around v_0 that is contained in U with $0 < r < 1$. For the purpose of local existence and local uniqueness, we may also replace I with a compact subinterval that is a neighborhood of t_0 in the given interval I , so we can assume I is compact, say with length c . Let $M > 0$ be an upper bound on $\|\phi(t, v)\|$ for $(t, v) \in I \times B$. Also, let $L > 0$ be an upper bound on the “operator norm” $\|D\phi(t, v)\|$ for $(t, v) \in I \times B$; here, $D\phi(t, v) : \mathbf{R} \times V \rightarrow V$ is the total derivative of ϕ at $(t, v) \in I \times U$. Such upper bounds exist because ϕ is C^1 and $I \times B$ is compact.

Our interest in L is that it serves as a “Lipschitz constant” for $\phi(t, \cdot) : B \rightarrow V$ for each $t \in I$. That is, for points $v, v' \in B$,

$$\|\phi(t, v) - \phi(t, v')\| \leq L\|v - v'\|.$$

This inequality is due to the Fundamental Theorem of Calculus: letting $g(x) = \phi(t, xv + (1-x)v')$ be the C^1 restriction of $\phi(t, \cdot)$ to the line segment in B joining v and v' (for a fixed t), we have

$$\phi(t, v) - \phi(t, v') = g(1) - g(0) = \int_0^1 g'(y)dy = \int_0^1 D\phi(t, yv + (1-y)v')(v - v')dy,$$

so

$$\|\phi(t, v) - \phi(t, v')\| \leq \int_0^1 \|D\phi(t, yv + (1-y)v')(v - v')\|dy \leq \int_0^1 L\|v - v'\|dy = L\|v - v'\|.$$

Since the solution u that we seek to construct near t_0 has to be continuous, the Fundamental Theorem of Calculus allows us to rephrase the ODE with initial condition near t_0 in I as an “integral equation”

$$u(t) = v_0 + \int_{t_0}^t \phi(y, u(y))dy$$

for a continuous mapping

$$u : I \cap [t_0 - a, t_0 + a] \rightarrow B = \overline{B}_{2r}(v_0) \subseteq U$$

for small $a > 0$. How small do we need to take a to get existence? We shall require $a \leq \min(1/2L, r/M)$.

Let $W = C(I \cap [t_0 - a, t_0 + a], V)$ be the space of continuous maps from $I \cap [t_0 - a, t_0 + a]$ to V , endowed with the sup norm. This is a complete metric space. Let $X \subseteq W$ be the subset of such maps f with image in the closed subset $B \subseteq V$ and with $f(t_0) = v_0$, so X is closed in W and hence is also a complete metric space. For any $f \in X$, define $T(f) : I \cap [t_0 - a, t_0 + a] \rightarrow V$ by

$$(T(f))(t) = v_0 + \int_{t_0}^t \phi(y, f(y)) dy;$$

this makes sense because $f(y) \in B \subseteq U$ for all $y \in I \cap [t_0 - a, t_0 + a]$ (as $f \in X$). By the continuity of ϕ and f , it follows that $T(f)$ is continuous. Note that $(T(f))(t_0) = v_0$ and $\|\phi(y, f(y)) - \phi(y, v_0)\| \leq L\|f(y) - v_0\| \leq 2rL$. Since

$$(T(f))(t) - v_0 = \int_{t_0}^t (\phi(t, f(y)) - \phi(y, v_0)) dy + \int_{t_0}^t \phi(y, v_0) dy$$

and $|t - t_0| \leq a$, we therefore get

$$\|(T(f))(t) - v_0\| \leq 2arL + Ma \leq 2r.$$

This shows $(T(f))(t) \in B$ for all $t \in I \cap [t_0 - a, t_0 + a]$, so $T(f) \in X$.

We conclude that the “integral operator” $f \mapsto T(f)$ is a self-map of the complete metric space X . The usefulness of this is that it is a contraction mapping: for $f, g \in X$,

$$\|T(f) - T(g)\|_{\sup} \leq \sup_t \left\| \int_{t_0}^t (\phi(y, f(y)) - \phi(y, g(y))) dy \right\| \leq a \sup_y \|\phi(y, f(y)) - \phi(y, g(y))\|,$$

where $t, y \in I \cap [t_0 - a, t_0 + a]$, and this is at most $aM\|f - g\|_{\sup}$ due to the choice of M . Since $aM \leq r < 1$, the contraction property is verified. Thus, there exists a *unique* fixed point $f_0 \in X$ for our integral operator. Such a fixed point gives a solution to the initial-value problem on $I \cap (t_0 - a, t_0 + a)$, and so settles the local existence result.

Since the value of a could have been taken to be arbitrarily small (for a given L, M, r), we also get a local uniqueness result: if $u_1, u_2 : J \rightarrow U$ are two solutions, then they coincide near t_0 . Indeed, we may take a as above so small that $I \cap [t_0 - a, t_0 + a] \subseteq J$ and both u_1 and u_2 map $I \cap [t_0 - a, t_0 + a]$ into B . (Here we use continuity of u_1 and u_2 , as $u_1(t_0) = u_2(t_0) = v_0 \in \text{int}_V(B)$.) For such small a , u_1 and u_2 lie in the complete metric space X as above and they are each fixed points for the same contraction operator. Hence, $u_1 = u_2$ on $I \cap [t_0 - a, t_0 + a]$. That is, u_1 and u_2 agree on J near t_0 . This completes the proof of local uniqueness near t_0 .

Finally, we must prove “global uniqueness”: if $J \subseteq I$ is an open connected set containing t_0 and u_1 and u_2 are solutions on all of J , then we want $u_1 = u_2$ on J . By local uniqueness, they agree on an open around t_0 in J . We now argue for uniqueness to the right of t_0 , and the same method will apply to the left. Pick any $t \in J$ with $t \geq t_0$. We want $u_1(t) = u_2(t)$. The case $t = t_0$ is trivial (by the initial condition!), so we may assume $t > t_0$. (In particular, t_0 is not a right endpoint of J .) Let $S \subseteq [t_0, t]$ be the subset of those $\tau \in [t_0, t]$ such that u_1 and u_2 coincide on $[t_0, \tau]$. For example, $t_0 \in S$. Also local uniqueness at t_0 implies that $[t_0, t_0 + \varepsilon] \subseteq S$ for some small $\varepsilon > 0$. It is clear that S is a subinterval of $[t_0, t]$ and that it is closed (as u_1 and u_2 are continuous), so if $\rho = \sup S \in (t_0, t]$ then $S = [t_0, \rho]$. Thus, the problem is to prove $\rho = t$.

We assume $\rho < t$ and we seek a contradiction. Let $v = u_1(\rho) = u_2(\rho)$. Hence, near the point ρ that lies on the *interior* of J in \mathbf{R} , we see that u_1 and u_2 are solutions to the same initial-value problem

$$u'(\tau) = \phi(\tau, u(\tau)), \quad u(\rho) = v$$

on an open interval around ρ contained in J . By local uniqueness applied to *this* new problem (with initial condition at ρ), it follows that u_1 and u_2 coincide near ρ in J . But ρ is not a right endpoint of J , so we get points to the right of ρ lying in S . This contradicts the definition of ρ , so the hypothesis $\rho < t$ is false; that is, $\rho = t$. Hence, $u_1(t) = u_2(t)$. Since $t \geq t_0$ in J was arbitrary, this proves equality of u_1 and u_2 on J to the right of t_0 . ■

Example 2.3 shows that when u' shows up non-linearly in the ODE, we can fail to have a solution across the entire interval I even in the absence of unboundedness problems. However, in Example 2.4 we had difficulties extending the solution across the entire interval due to unboundedness problems. Let us show that in the general case of initial-value problems of the form $u'(t) = \phi(t, u(t))$ with $u(t_0) = v_0$, “unboundedness” is the only obstruction to the local solution propagating across the entire domain. By “unboundedness” we really mean that the local solution approaches the boundary of U , which is to say that it fails to remain in a compact subset of U as we evolve the solution over time.

Corollary 2.5. *With notation as in Theorem 2.1, let u be a solution to (2.1) on some connected open subset $J \subseteq I$ around t_0 . Let $K \subseteq U$ be a compact subset. If $\tau_0 \in I$ is an endpoint of the closure of J in I such that u has image in K at all points near τ_0 in $J - \{\tau_0\}$, then u extends to a solution of (2.1) around τ_0 . That is, J_{\max} contains a neighborhood of τ_0 in I .*

The condition on $u(t)$ as $t \in J$ tends to τ_0 is trivially necessary if there is to be an extension of the solution around τ_0 ; the interesting feature is that it is sufficient, and more specifically that we do not need to assume anything stronger such as existence of a limit for $u(t)$ as $t \in J$ approaches τ_0 : containment in a compact subset of U is all we are required to assume. In the special case $U = V$, remaining in a compact is the same as boundedness, and so the meaning of this corollary in such cases is that if we have a solution on J and the solution is bounded as we approach an endpoint of J in I , then the solution extends to an open neighborhood of that endpoint in J . Of course, if U is a more general open subset of V then we could run into other problems: perhaps $u(t)$ approaches the boundary of U as $t \rightarrow \tau_0 \in \partial_I J$. This possibility is ruled out by the compactness hypothesis in the corollary.

Proof. If $I' \subseteq I$ is a compact connected neighborhood of τ_0 , we may replace I with I' , so we can assume I is compact. Fix a norm on V . In the proof of local existence, the conditions on a involved upper bounds in terms of parameters r , L , and M . Since K is compact and U is open, there exists $r_0 \in (0, 1)$ such that for *any* $v \in K$, $\overline{B}_{2r_0}(v) \subseteq U$. Let $K' \subseteq V$ be the set of points with distance $\leq 2r_0$ from K , so K' is compact and $K' \subseteq U$ (since compactness of K implies that the “distance” to K for any point of V is attained by some point of K).

By compactness of $I \times K'$, there exist $L_0, M_0 > 0$ that are respectively upper bounds on $\|D\phi(t, v)\|$ and $\|\phi(t, v)\|$ for all $(t, v) \in I \times K'$. In particular, for all $v \in B$ we have $\overline{B}_{2r_0}(v) \subseteq K'$ and hence the parameters r_0, L_0, M_0 are suitable for the proof of local existence with an initial condition $u(\tau) = v$ for *any* $\tau \in I$ and *any* $v \in K$. In particular, letting $a = \min(1/2L_0, r_0/M_0)$, for any $\tau \in I$ and $v \in K$ there is a solution to the initial-value problem

$$\tilde{u}'(x) = \phi(x, \tilde{u}(x)), \quad \tilde{u}(\tau) = v$$

on $I \cap [\tau - a, \tau + a]$, where $a > 0$ is *independent* of $(\tau, v) \in I \times K$.

Returning to the endpoint τ_0 near which we want to extend the solution to the initial-value problem, consider points $\tau \in J$ near τ_0 . By hypothesis we have $u(\tau) \in K$ for all such τ . We may find such τ with $|\tau - \tau_0| < a$, so $\tau_0 \in I \cap (\tau - a, \tau + a)$. Thus, if we let $v = u(\tau)$ then the initial-value problem

$$\tilde{u}'(x) = \phi(x, \tilde{u}(x)), \quad \tilde{u}(\tau) = v$$

has a solution on the open subset $I \cap (\tau - a, \tau + a)$ in I that contains τ_0 (here we are using the “universality” of a for an initial-value condition at any point in I with the initial value equal to any point in K). However, on $J \cap (\tau - a, \tau + a)$ such a solution is given by u ! Hence, we get a solution to our differential equation near τ_0 in I such that on J near τ_0 it coincides with u . This solves our extension problem around τ_0 for our original initial-value problem. ■

For our needs in differential geometry, it remains to address two rather different phenomena for first-order initial-value problems of the form $u'(t) = \phi(t, u(t))$ with $u(t_0) = v_0$ for C^p mappings $\phi : I \times U \rightarrow V$:

- the global existence result in the linear case on all of I (with $U = V$),
- the C^p dependence of u on the initial condition v_0 as well as on auxiliary parameters (when ϕ has C^p dependence on such auxiliary parameters) in the general case.

In the final section of this handout, we take up the first of these problems.

3. LINEAR ODE

Our next goal is to prove a global form of Theorem 2.1 in the *linear* case:

Theorem 3.1. *Let $I \subseteq \mathbf{R}$ be a non-trivial interval and let V be a finite-dimensional vector space over \mathbf{R} . Let $A : I \rightarrow \text{Hom}(V, V)$ and $f : I \rightarrow V$ be C^p mappings with $p \geq 0$. Choose $t_0 \in I$ and $v_0 \in V$. The initial-value problem*

$$u'(t) = (A(t))(u(t)) + f(t), \quad u(t_0) = v_0$$

has a unique solution u on I , and $u : I \rightarrow V$ is C^{p+1} .

Remark 3.2. Observe that in this theorem we allow A and f to be merely continuous, not necessarily differentiable. Thus, the corresponding $\phi(t, v) = (A(t))(v) + f(t)$ is merely continuous and not necessarily differentiable. But this ϕ has the special feature that $\phi(t, \cdot) : V \rightarrow V$ is affine-linear, and so its possible lack of differentiability is not harmful: the *only* reason we needed ϕ to be at least C^1 in the proof of the existence and uniqueness theorem was to locally have the Lipschitz property

$$\|\phi(t, v) - \phi(t, v')\| \leq L\|v - v'\|$$

for some $L > 0$, with t near t_0 and any v, v' near v_0 .

For ϕ of the special “affine-linear” form, this Lipschitz condition can be verified without conditions on A and f beyond continuity:

$$\|\phi(t, v) - \phi(t, v')\| = \|(A(t))(v - v')\| \leq \|A(t)\| \cdot \|v - v'\|,$$

so we just need a uniform positive upper bound on the operator norm of $A(t)$ for t near t_0 . The existence of such a bound follows from the continuity of A and the continuity of the operator norm on the finite-dimensional vector space $\text{Hom}(V, V)$. In particular, for such special ϕ we may apply the results of the preceding section without differentiability restrictions on the continuous A and f .

One point we should emphasize is that the proof of Theorem 3.1 will *not* involve reproving the local existence theorem in a manner that exploits linearity to get a bigger domain of existence. Rather, the proof will use the local existence/uniqueness theorem as *input* (via Remark 3.2) and exploit the linear structure of the differential equation to push out the domain on which the solution exists. In particular, the contraction mapping technique (as used for the local version in the general case without linearity assumptions) does not construct the global solution in the linear case: the integral operator used in the local proof fails to be a contraction mapping when the domain is “too big”, and hence we have to use an entirely new idea to prove the global existence theorem in the linear case. (It is conceivable that a better choice of norms could rescue the contraction method, but I do not see such a choice, and anyway it is instructive to see how alternative methods can bypass this difficulty.)

Let us now begin the proof of Theorem 3.1. The uniqueness and C^{p+1} properties are special cases of Theorem 2.1. The problem is therefore one of existence on I , and so in view of the uniqueness it suffices to solve the problem on arbitrary *bounded* subintervals of I around t_0 . Moreover, in case t_0 is an endpoint of I we may extend A and f to C^p mappings slightly past this endpoint (for $p > 0$ either use the usual cheap definition of the notion of C^p mapping on a half-closed interval or else use the Whitney extension theorem if you are a masochist, and for $p = 0$ extend by a “constant” past the endpoint), so in such cases it suffices to consider the existence problem on a larger interval with t_0 as an interior point. Hence, for the existence problem we may and do assume I is a bounded open interval in \mathbf{R} . A linear change of variable on t is harmless (why?), so we may and do now suppose $I = (t_0 - r, t_0 + r)$ for some $r > 0$. By the *local* existence theorem, we can solve the problem on $(t_0 - \varepsilon, t_0 + \varepsilon)$ for some $0 < \varepsilon \leq r$.

Let $\rho \in (0, r]$ be the supremum of the set S of $c \in (0, r]$ such that our initial-value problem has a solution on $I_c = (t_0 - c, t_0 + c)$. Of course, the solution on such an I_c is unique, and by uniqueness in general its restriction to any open subinterval centered at t_0 is the unique solution on that subinterval. Hence, $S = (0, \rho)$ or $S = (0, \rho]$. Before we compute ρ , let us show that the latter option must occur. For $0 < c < \rho$ we have a unique solution u_c on I_c for our initial-value problem, and uniqueness ensures that if $c < c' < \rho$ then $u_{c'}|_{I_c} = u_c$. Thus, the u_c 's “glue” to a solution u_ρ on the open union I_ρ of the I_c 's for $c \in (0, \rho)$. This forces $\rho \in S$, so $S = (0, \rho]$. Our problem is therefore to prove $\rho = r$. Put another way, since we have a solution on I_ρ , to get a contradiction if $\rho < r$ it suffices to prove generally that if there is a solution u on I_c for some $c \in (0, r)$ then there is a solution on $I_{c'}$ for some $c < c' < r$. This is what we shall now prove.

The key to the proof is that the solution does not “blow up” in finite time. More specifically, by Corollary 2.5 with $U = V$, it suffices to fix a norm on V and prove:

Lemma 3.3. *The mapping $\|u\| : I_c \rightarrow \mathbf{R}$ defined by $t \mapsto \|u(t)\|$ is bounded.*

To prove the lemma, first note that since $c < r$ we have $\bar{I}_c = [t_0 - c, t_0 + c] \subseteq I$, so by compactness of \bar{I}_c there exist constants $M, m > 0$ such that $\|A(t)\| \leq M$ and $\|f(t)\| \leq m$ for all $t \in \bar{I}_c$; we are using the sup-norm on $\text{Hom}(V, V)$ arising from the choice of norm on V . By the differential equation,

$$\|u'(t)\| \leq M\|u(t)\| + m$$

for all $t \in I_c$. By the Fundamental Theorem of Calculus (for maps $I \rightarrow V$), for $t \in I_c$ we have

$$u(t) = u(t_0) + \int_{t_0}^t u'(y)dy = v_0 + \int_{t_0}^t u'(y)dy,$$

so for $t \geq t_0$

$$\|u(t)\| \leq \|v_0\| + \int_{t_0}^t (M\|u(y)\| + m)dy = (\|v_0\| + m|t - t_0|) + M \int_{t_0}^t \|u(y)\|dy$$

and likewise for $t \leq t_0$ with the final integral given by $\int_t^{t_0}$. Let us briefly grant the following lemma:

Lemma 3.4. *Let $\alpha, \beta, h : [0, a] \rightarrow \mathbf{R}_{\geq 0}$ be continuous functions (with $a > 0$) such that*

$$(3.1) \quad h(\tau) \leq \alpha(\tau) + \int_0^\tau h(y)\beta(y)dy$$

for all $\tau \in [0, a]$. Then $h(\tau) \leq \alpha(\tau) + \int_0^\tau \alpha(y)\beta(y)e^{\int_y^\tau \beta} dy$ for all $\tau \in [0, a]$.

Using this Lemma with $h = \|u\|$, $\alpha(\tau) = \|v_0\| + m|\tau - t_0|$, and $\beta(\tau) = M$ for $\tau \in [t_0, t]$ (with fixed $t \in (t_0, t_0 + c)$), by letting $\mu = \|v_0\| + mc$ we get

$$\|u(t)\| \leq \mu + \int_{t_0}^t M\mu e^{M(t-y)} dy \leq \mu + \mu(e^{M(t-t_0)} - 1) \leq \mu e^{Mc}$$

for all $t \in [t_0, t_0 + c]$. The same method gives the same constant bound for $t \in (t_0 - c, t_0]$. Thus, this proves the boundedness of $\|u(t)\|$ as t varies in $(t_0 - c, t_0 + c)$, conditional on Lemma 3.4 that we must now prove:

Proof. Let $I(\tau) = \int_0^\tau h(y)\beta(y)dy$; this is a C^1 function of $\tau \in [0, a]$ since βh is continuous on $[0, a]$. By direct calculation,

$$I' - \beta I = \beta \cdot (h - I) \leq \beta \alpha$$

by (3.1) (and the non-negativity of β). Hence, letting $q(\tau) = I(\tau)e^{-\int_0^\tau \beta}$, clearly q is C^1 on $[0, a]$ with $q(0) = 0$ and

$$q'(\tau) = e^{-\int_0^\tau \beta} (I'(\tau) - \beta(\tau)I(\tau)) \leq e^{-\int_0^\tau \beta} \beta(\tau)\alpha(\tau).$$

Since $q(0) = 0$, so $q(\tau) = \int_0^\tau q'$, we have $q(\tau) \leq \int_0^\tau \alpha(y)\beta(y)e^{-\int_0^y \beta} dy$, whence multiplying by the number $e^{-\int_0^\tau \beta}$ gives

$$I(\tau) \leq \int_0^\tau \alpha(y)\beta(y)e^{\int_y^\tau \beta} dy.$$

By (3.1) and the definition of I , we are done. ■

This completes the proof of the global existence/uniqueness theorem in the linear case. Let us record a famous consequence and give an example to illustrate it.

Corollary 3.5. *Let $I \subseteq \mathbf{R}$ be a nontrivial interval and let $a_0, \dots, a_{n-1} : I \rightarrow \mathbf{R}$ be smooth functions. Let $D : C^\infty(I) \rightarrow C^\infty(I)$ be the \mathbf{R} -linear map $u \mapsto u^{(n)} + a_{n-1}u^{(n-1)} + \dots + a_1u' + a_0u$.*

The equation $Du = h$ has a solution for all $h \in C^\infty(I)$, and $\ker D$ is n -dimensional. More specifically, for any $t_0 \in I$ the mapping

$$(3.2) \quad u \mapsto (u(t_0), u'(t_0), \dots, u^{(n-1)}(t_0)) \in \mathbf{R}^n$$

is a bijection from the set of solutions to $Du = h$ onto the “space” \mathbf{R}^n of initial conditions.

Proof. The existence of a solution to $Du = h$ on all of I follows from Theorem 3.1 applied to the first-order reformulation of our problem with $V = \mathbf{R}^n$. Since $D(u_1) = D(u_2)$ if and only if $u_1 - u_2 \in \ker D$, the proposed description of the set of all solutions is exactly the statement that the vector space $\ker D$ maps isomorphically onto \mathbf{R}^n via the mapping (3.2). Certainly this is an \mathbf{R} -linear mapping, so the problem is one of bijectivity. But this is precisely the statement that the equation $Du = 0$ admits a unique solution for each specification of the $u^{(j)}(t_0)$'s for $0 \leq j \leq n - 1$,

and this follows from applying Theorem 3.1 to the first-order reformulation of our problem (using $V = \mathbf{R}^n$). ■

Example 3.6. Consider the general second-order equation with *constant* coefficients

$$y'' + Ay' + By = 0$$

on \mathbf{R} . Let $\delta = A^2 - 4B$. The nature of the solution space depends on the trichotomy of possibilities $\delta > 0$, $\delta < 0$, and $\delta = 0$. Such simple equations are easily solved by the method of the characteristic polynomial if we admit \mathbf{C} -valued functions (in case $\delta < 0$), but rather than discuss that technique (which applies to any constant-coefficient linear ODE) we shall simply exhibit some solutions by inspection and use a dimension-count to ensure we've found all solutions. Corollary 3.5 ensures that the equation has a 2-dimensional solution space in $C^\infty(\mathbf{R})$, with each solution uniquely determined by $y(0)$ and $y'(0)$ (or even $y(x_0)$ and $y'(x_0)$ for any particular $x_0 \in \mathbf{R}$), so to find all solutions we just have to exhibit 2 linearly independent solutions.

If $\delta = 0$ then the left side of the equation is $y'' + Ay' + (A/2)^2y = (\partial_x + A/2)^2y$. In this case inspection or iterating solutions to the trivial equation $(\partial_x + C)y = h$ yields solutions $e^{-Ax/2}$ and $x e^{-Ax/2}$. These are independent because if $c_0 e^{-Ax/2} + c_1 x e^{-Ax/2} = 0$ in $C^\infty(\mathbf{R})$ with $c_0, c_1 \in \mathbf{R}$ then multiplying by $e^{Ax/2}$ gives $c_0 + c_1 x = 0$ in $C^\infty(\mathbf{R})$, an impossibility except if $c_0 = c_1 = 0$. If $\delta \neq 0$, let $k = \sqrt{|\delta|}/2 > 0$. Two solutions are $e^{-Ax/2}y_1$ and $e^{-Ax/2}y_2$ where $y_1(x) = e^{kx}$ and $y_2(x) = e^{-kx}$ if $\delta > 0$, and $y_1(x) = \cos(kx)$ and $y_2(x) = \sin(kx)$ if $\delta < 0$. The linear independence is trivial to verify in each case. Thus, we have found all of the solutions.

We push this example a bit further by imposing initial conditions $y(0) = c_0$ and $y'(0) = c_1$ and studying how the solution depends on A and B viewed as “parameters”. What is the unique associated solution $y_{A,B}$? Some simple algebra gives

$$y_{A,B} = \begin{cases} e^{-Ax/2}(c_0 + (c_1 + c_0 A/2)x), & \delta = 0 \\ e^{-Ax/2} \left(c_0 \cos(kx) + (c_1 + c_0 A/2)x \cdot \frac{\sin(kx)}{kx} \right), & \delta < 0 \\ e^{-Ax/2} \left(c_0 \cdot \frac{e^{kx} + e^{-kx}}{2} + (c_1 + c_0 A/2)x \cdot \frac{e^{kx} - e^{-kx}}{2kx} \right), & \delta > 0 \end{cases}$$

with $k = \sqrt{|\delta|}/2 = \sqrt{|A^2 - 4B|}/2$. Observe that $(A, B, x) \mapsto y_{A,B}(x)$ is *continuous* in the triple (A, B, x) , the real issue being at triples for which $A^2 - 4B = 0$. (Recall that $\sin(v)/v$ is continuous at $v = 0$.) What is perhaps less evident by inspection of these formulas (due to the trichotomous nature of $y_{A,B}(x)$ as a function of (A, B, x) , and the intervention of $\sqrt{|A^2 - 4B|}$) is that $y_{A,B}(x)$ is C^∞ in (A, B, x) ! In fact, even as a function of the 5-tuple (A, B, c_0, c_1, x) it is C^∞ .

Why is there “good” dependence of solutions on initial conditions and auxiliary parameters as they vary? Applications in differential geometry will require affirmative answers to such questions in the non-linear case. Hence, we need to leave the linear setting and turn to the study of properties of solutions to general first-order initial-value problems as we vary the equation (through parameters or initial conditions). This is taken up in the handout on integral curves.