8 Characteristic lines and surfaces

In this section, we revisit the characteristic lines that we studied for first order systems before. We relate them to the existence and uniqueness of solutions to the Cauchy –or initial value– problem and to the propagation of singularities, and use them to classify quasi-linear equations in two independent variables. Then we switch to more that two independent variables, and draw a distinction between characteristic surfaces and their bi-characteristic lines.

8.1 Preliminaries

We begin by drawing new conclusions from a simple example that we studied before. Consider the equation

$$u_t + u_x = 0, \tag{162}$$

which can be restated as follows: along the characteristic lines $x = x(t)$ with slope

$$\frac{dx}{dt} = 1,$$

$u(x, t)$ does not vary:

$$\frac{du}{dt} = 0.$$

In our previous studies, we mainly though of this formulation in terms of characteristics as a simple way of solving a partial differential equation, reducing it to a system of ordinary ones. From initial data at $t = 0$,

$$u(x, 0) = u_0(x),$$

we can find the solution at all times by following the characteristics:

$$u(x, t) = u_0(x - t).$$

Now we change perspective. Consider the more general situation in which we are given initial data not at $t = 0$ but along an arbitrary line $(x(s), t(s))$:

$$u(x(s), t(s)) = u_0(s).$$

Can we still solve equation (162) using characteristics? More fundamentally, will the equation still have solutions consistent with the initial data and, if so, only one?

The answer is simple. If the line $(x(s), t(s))$ is *transversal* –i.e., nowhere tangent– to the characteristic lines $x - t =$ const., then we can still follow characteristics, along which u is constant, to extend the data u_0 to a solution valid throughout the plane (x, t). On the opposite case, when $(x(s), t(s))$ *is* a characteristic (i.e. $x(s) = t(s) + \text{const.}$), two problems arise: first, $u_0(s)$ cannot be given arbitrarily, as it needs to be a constant in order to satisfy the equation.

Second, even if $u_0(s)$ is a constant, the equation tells us nothing about the behavior of $u(x,t)$ nearby, since it only informs about variations of u along the characteristic direction $\frac{dx}{dt} = 1$, which in this case stays on the line $(x(s), t(s))$ where the data are provided.

The intermediate, perhaps more typical case, where $(x(s), t(s))$ is tangent to the characteristic field at isolated points $s = s_j$, brings in two kinds of constraints on the data. On the one hand, we need to have $u_0'(s_j) = 0$, for consistency with the characteristic equation at $(x(s_j), t(s_j))$. On the other, when a characteristic crosses the initial line $(x(s), t(s))$ twice, we would need to have the same value of u_0 on both intersections, since there should be a single value of u on each characteristic. We will not consider this second, *global* problem here though (We have seen characteristics crossing before, bringing in contradictory information, in quasi-linear first order equations, and we found a way to solve this contradiction by extending the notion of a solution –i.e., by allowing shocks.) In this section, we will only concern ourselves with extending the solution locally to a small neighborhood of the initial curve $(x(s), t(s))$.

Clearly, everything that we have said so far in the context of our simple example extends to the more general class of equations

$$u_t + a(x,t,u)u_x = b(x,t,u),$$

with characteristics

$$\frac{dx}{dt} = a(x,t,u),$$

along which u satisfies

$$\frac{du}{dt} = b(x,t,u).$$

The only difference is that now whether or not the initial curve $(x(s), t(s))$ is a characteristic depends not only on $(x(s), t(s))$, but also on $u_0(s)$.

Yet let us stick for a little longer to the simple example in (162) to gain one more perspective on the meaning of characteristics. The differential equation (162) provides information on the derivative of u along characteristics. The initial value, on the other hand, provides u and hence its tangential derivatives along the initial line $(x(s), t(s))$. Putting these two pieces of information together, we know along $(x(s), t(s))$ all first derivatives of u; this makes it possible for us to extend u to a neighborhood of the initial curve. Unless, of course, the two directions coincide: if the initial curve is locally characteristic, then both the equation and the data provide information on the derivative tangent to the curve. Then, for there to be a solution, both pieces of information need to agree. In addition, the derivative of u normal to the curve is still unknown.

This latter conclusion allows us to rethink the characteristics in terms that do not involve the Cauchy problem. Consider a solution $u(x,t)$ of equation (162). For it to make sense as a classical solution, its first derivatives need to exist, but they do not necessarily need to be continuous. Could there be a line $(x(s), t(s))$ across which the first derivatives jump? The answer is that, for this to be possible, the line in question needs to be a characteristic. For otherwise, since u

adopts some value $u(x(s),t(s))$ on the line, this and the equation determine all first derivatives of u, which are therefore the same on both sides of the line. It is only when $(x(s),t(s))$ is a characteristic that the normal derivative of u can jump across it. Hence, from this perspective,

Characteristic lines are those along which weak singularities propagate; i.e., across which the normal derivative of u may jump.

In order to extend the notion of characteristics to more general situations, it is convenient to re-compute them using this latter characterization. A prescribed value $u_0(s)$ along the initial curve $(x(s),t(s))$ implies the following compatibility constraint on the first derivatives of u on the curve:

$$t'(s)\ u_t + x'(s)\ u_x = u_0'(s).$$

Then this constraint and the differential equation provide a system of equations for u_t, u_x:

$$\begin{pmatrix} t'(s) & x'(s) \\ 1 & 1 \end{pmatrix} \begin{pmatrix} u_t \\ u_x \end{pmatrix} = \begin{pmatrix} u_0'(s) \\ 0 \end{pmatrix},$$

a system that always has a solution unless the determinant

$$t'(s) - x'(s)$$

vanishes, yielding the characteristic lines

$$x - t = \text{const.}$$

If the line $(x(s),t(s))$ is characteristic, then the system above only has solutions when the right-hand side is orthogonal to the left null-space of the matrix (the Fredholm alternative), yielding the condition

$$u_0'(s) = 0,$$

consistent with the evolution of u along characteristics. When this constraint holds, the system has the infinitely many solutions

$$\begin{pmatrix} u_t \\ u_x \end{pmatrix} \propto \begin{pmatrix} -1 \\ 1 \end{pmatrix},$$

i.e. the derivative normal to the curve is arbitrary. Then we can make this normal derivative jump, adopting one value on each side of $(x(s),t(s))$. This jump, however, can only be prescribed at one point s. For consider a general solution with a prescribed jump $f(s)$ of the normal derivative across a characteristic:

$$[u_t(x(s),t(s)) - u_x(x(s),t(s))] = f(s),$$

where the brackets denote the intensity of the jump. Since, along the line $(x(s),t(s))$, we have that $u_0'(s) = 0$, differentiating this equations with respect to s yields

$$0 = \frac{d}{ds}\left[u_t(x(s),t(s)) - u_x(x(s),t(s))\right] = f'(s).$$

So, even though the normal derivative of u may jump across a characteristic, the evolution along the characteristic of the strength of the jump is constrained; in our simple example, the intensity of the jump is constant throughout.

It should be clear how to extend these results to general first order equations for a function of two independent variables. Our next job, therefore, is to extend them to equations of higher order and to more dimensions.

8.2 Classification of second-order, quasi-linear equations in two independent variables

Consider the equation

$$au_{xx} + 2bu_{xy} + cu_{yy} = d, \tag{163}$$

where a, b, c and d can be functions of x, y, u, u_x and u_y. The Cauchy problem for this equation consists of $u(s)$ and its normal derivative $\frac{\partial u}{\partial n}(s)$ along a line $(x(s), y(s))$. We can mimic the argument above, and write two compatibility conditions (one for u_x and one for u_y) and the equation itself as a system for the second derivatives of u. To this end, it is best to think that we are given $u_x = f(s)$ and $u_y = g(s)$ along the initial line, which follow from knowing $u(s)$ and $\frac{\partial u}{\partial n}(s)$. Then the system becomes

$$\begin{pmatrix} x'(s) & y'(s) & 0 \\ 0 & x'(s) & y'(s) \\ a & 2b & c \end{pmatrix} \begin{pmatrix} u_{xx} \\ u_{xy} \\ u_{yy} \end{pmatrix} = \begin{pmatrix} f'(s) \\ g'(s) \\ d \end{pmatrix},$$

with a unique solution unless the determinant is zero, which yields the characteristic condition

$$ay'(s)^2 - 2bx'(s)y'(s) + cx'(s)^2 = 0,$$

which can be re-written as a differential equation for $y(x)$:

$$\frac{dy}{dx} = \frac{y'(s)}{x'(s)} = \frac{b \pm \sqrt{b^2 - ac}}{a}.$$

We notice that, through each point, there are two, one or no real characteristic slopes, depending on the sign of the discriminant $\Delta = b^2 - ac$. We classify the equation (163) accordingly as *hyperbolic*, *parabolic* or *elliptic*, depending on whether Δ is positive, zero or negative. Then hyperbolic equations have two real characteristic families, parabolic equations only one, and elliptic equations none. It follows that Cauchy data should always determine a local solution for elliptic equations, but only on non-characteristic lines for equations that are parabolic or hyperbolic. We will qualify this statement later through examples and through the theorem of Cauchy and Kovalewski.

It is convenient to compute the characteristics in a different way. First we rewrite equation (163) in a new set of coordinates, $\xi(x, y)$ and $\eta(x, y)$. The resulting equation adopts an entirely similar form:

$$Au_{\xi\xi} + 2Bu_{\xi\eta} + Cu_{\eta\eta} = D, \tag{164}$$

where

$$A = a{\xi_x}^2 + 2b\xi_x\xi_y + c{\xi_y}^2,$$

$$B = a\xi_x\eta_x + b\left(\xi_x\eta_y + \xi_y\eta_x\right) + c\xi_y\eta_y,$$

$$C = a{\eta_x}^2 + 2b\eta_x\eta_y + c{\eta_y}^2$$

and, for our purposes here, we do not need to compute D. So far the coordinates ξ and η have been arbitrary. But say we want ξ to represent the characteristic field, so that $\xi(x, y) =$ const. are characteristic lines. Then, by definition, we shouldn't be able to compute the second derivative of u in the direction normal to $\xi =$ const. from the equation, i.e. $u_{\xi\xi}$. It follows that A should be zero, so ξ needs to satisfy the PDE

$$a{\xi_x}^2 + 2b\xi_x\xi_y + c{\xi_y}^2 = 0. \tag{165}$$

Then, along a characteristic, we have

$$\frac{dy}{dx} = -\frac{\xi_x}{\xi_y} = \frac{b \pm \sqrt{b^2 - ac}}{a}$$

as before.

Notice that (165) is a nonlinear, first order PDE. Then we can compute its own characteristics, as in section (3.3). These are given by

$$\dot{x} = 2a\xi_x + 2b\xi_y,$$

$$\dot{y} = 2b\xi_x + 2c\xi_y,$$

and

$$\dot{\xi} = \xi_x\dot{x} + \xi_y\dot{y} = 2\left[a{\xi_x}^2 + 2b\xi_x\xi_y + c{\xi_y}^2\right] = 0,$$

so ξ is constant along the characteristics of (165), and then these agree with the characteristics of (163). This sounds like a tongue twister; we will come back to it, understand it more deeply and prove it in more generality, when we consider situations with more that two independent variables, where we will need to discriminate between characteristic surfaces and bi-characteristic lines.

When the equation is hyperbolic, there are two characteristic families. We can use one for ξ and one for η, thus eliminating both A and C from (164), which acquires the canonical form

$$u_{\xi\eta} = E,$$

where $E = \frac{D}{2B}$. The simplest hyperbolic equation has $E = 0$:

$$u_{\xi\eta} = 0,$$

with general solution

$$u = F(\xi) + G(\eta).$$

The change of variables $\xi = x - t$, $\eta = x + t$ reveals that we have just solved again the one-dimensional wave equation,

$$u_{tt} - u_{xx} = 0,$$

a prototype for hyperbolic equations. The prototypical elliptic and parabolic equations are Laplace's

$$u_{xx} + u_{yy} = 0,$$

and the heat equation

$$u_t = u_{xx}$$

respectively.

8.3 More than two independent variables and equations of higher order

Much of what we said above extends to equations in more that two independent variables. The main difference is that now the Cauchy problem involves initial data given on surfaces, not lines, and so those surfaces where Cauchy data are insufficient –and possibly inconsistent– are now *characteristic surfaces*. Instead of attempting a general discussion, let us work out the details in a representative –and significant– example: the 2-dimensional wave equation in an inhomogeneous medium:

$$u_{tt} - c(x, y)^2 (u_{xx} + u_{yy}) = 0.$$

A surface $\phi(x, y, t) = \text{const.}$ will be characteristic if we cannot figure out $u_{\phi\phi}$ from the equation. As before, this implies that the following nonlinear first order PDE is satisfied:

$$F(\phi_x, \phi_y, \phi_t, x, y) = {\phi_t}^2 - c(x, y)^2 \left({\phi_x}^2 + {\phi_y}^2\right) = 0,$$

an equation that we have called Eikonal in one of our old homeworks. The characteristic curves of this equation are called *bi-characteristics* of the original PDE; they correspond to the rays of light in the geometrical optics approximation to the wave equation. A seemingly magical fact is that these bi-characteristic lines lie on the characteristic surfaces –this was our tongue-twister before, when all characteristics were lines, and the statement was that characteristics and bi-characteristics agreed. To prove this fact in general –not just for the present example–, it is enough to notice that, along a bi-characteristic,

$$\dot{\phi} = \sum_{j=1}^{d} \phi_{x_j} \frac{\partial F}{\partial \phi_{x_j}} = nF$$

for functions F that are homogeneous of degree n (2 in our case) in the ϕ_{x_j}'s. Since $F = 0$, ϕ is constant along bi-characteristics, and hence these lie on the characteristic surfaces, that are precisely defined by the constancy of ϕ. Note that this argument applies to equations of any order –hence the n– and in any number d of dimensions; in our case, $d = 3$, and $x_j,\ j = 1, 2, 3$ stand for x, y and t.

8.4 Can we actually give Cauchy data? Some typical examples

From the discussion above, it would appear that we can always give Cauchy data on non-characteristic surfaces and solve the differential equation nearby, while we cannot do the same on characteristic surfaces. Yet this seems to contradict things we have learned before while studying some classical equations. In this subsection, we attempt to clarify these apparent contradictions through examples.

8.4.1 One-directional wave equation

For the first-order (i.e. one-directional) wave equation

$$u_t + u_x = 0,$$

the intuition that we have built is clearly true; in fact, this is the equation that we used for building our intuition. The characteristics are

$$\xi(x, t) = x - t = \text{const.},$$

along which u is constant. Hence, when Cauchy data $u_0(s) = u(x(s), t(s))$ are given on a non-characteristic line, we can extend them to a global solution by following these characteristics:

$$u(x, t) = f(x - t),$$

with the function f defined by the condition that

$$f(x(s) - t(s)) = u_0(s).$$

If, on the other hand, Cauchy data $u_0(s)$ are given on a characteristic line with $x(s) - t(s) = \text{const.}$, then continuous solutions –in fact, infinitely many– will only exist when $u_0(s)$ is constant. In the exact solution above, only the value of f at one point is provided by the data.

8.4.2 The one-dimensional wave equation

For the full wave equation in one spatial dimension,

$$u_{tt} - u_{xx} = 0,$$

the intuition also holds. It is easiest to see this from the exact general solution,

$$u(x,t) = F(x-t) + G(x+t),$$

where F is constant along the characteristic family $x - t =$ const. and G along the other family, $x+t =$ const. If Cauchy data $[u(x(s),t(s)), \frac{\partial u}{\partial n} u(x(s),t(s))]$ are given along a non-characteristic line, they can be used to find F and G uniquely. Otherwise, the data need to satisfy a compatibility condition for solutions to exist, and then there will be infinitely many, since one of the functions F and G will be specified at only one point.

This is a general situation for hyperbolic systems: Cauchy data provided along non-characteristic surfaces can be extended to a solution nearby, while data along characteristics need to satisfy constraints and do not fully determine a local solution.

8.4.3 The two-dimensional Laplace's equation

Laplace's equation

$$u_{xx} + u_{yy} = 0,$$

as all elliptic equations, has no real characteristic lines, and so it would appear that Cauchy data could be safely provided along any line (or surface for the corresponding equation in higher dimensions.) How would this work, for instance, for data provided on the $y = 0$ line? We would be given

$$u(x,0) = f(x)$$

and

$$u_y(x,0) = g(x),$$

and asked to compute $u(x,y)$ nearby. To this end, we can differentiate the equation with respect to y, and write

$$\frac{\partial^{n+2} u(x,0)}{\partial y^{n+2}} = -\frac{\partial^2}{\partial x^2} \frac{\partial^n u(x,0)}{\partial y^n},$$

which allows us to compute all normal derivatives of u recursively. Cauchy-Kovalewky's theorem tells us that the correspondingly built Taylor expansion for $u(x,y)$ will converge, for $|y|$ small enough, to a solution $u(x,y)$, provided that $f(x)$ and $g(x)$ are analytic functions.

Yet analytic functions are very special; what will happen with more general Cauchy data? This is most easily seen by writing a general solution $u(x,t)$ as a superposition of Fourier modes:

$$u(x,t) = \int \left[A(k)e^{|k|y} + B(k)e^{-|k|y} \right] e^{ikx} \, dk,$$

where the functions $A(k)$ and $B(k)$ are a linear combination of the Fourier coefficients of $f(x)$ and $g(x)$. Clearly, for this integral to converge for $|y| \neq 0$, no matter how small, we need the coefficients A and B to decay exponentially,

$$|A(k)|, |B(k)| < e^{-\alpha k},$$

with $\alpha > |y|$. This will only be the case if $f(x)$ and $g(x)$ are analytic. So, for non-analytic Cauchy data, a solution does not exist even locally, even though the line $y = 0$ is non-characteristic.

For analytic data, a solution exists, but the problem is ill-posed. Because, for any $y \neq 0$, there will be an analytic small perturbation of the data that will make the solution not only change by an arbitrary amount, but even cease to exist: it is enough to make an analytic perturbation with $\alpha < |y|$.

Thus, for elliptic problems, it is not right to conclude from the absence of real characteristics that the Cauchy problem is fine. Rather, a more robust conclusion is that all solutions to the equation are very smooth in the interior of their domain of existence, since weak singularities can only occur across characteristic surfaces, and elliptic equations have none.

8.4.4 The heat equation

For the one-dimensional heat equation,

$$u_t = u_{xx},$$

the only characteristic family is given by $t =$ constant. Consider, in particular, the line $t = 0$. Clearly, Cauchy data

$$u(x, 0) = f(x), \quad u_t(x, 0) = g(x)$$

need to satisfy the constraint given by the equation itself:

$$g(x) = f''(x).$$

Hence, it is enough to provide $u(x, 0) = f(x)$. But will this determine the solution for $t > 0$? (We have seen before that, for $t < 0$, the problem is ill-posed. Moreover, an argument in Fourier space like the one carried out before for Laplace's equation shows that solutions typically do not even exist for $t < 0$.)

It would appear that the solution is uniquely determined by $f(x) = u(x, 0)$. On the one hand, we have computed the solution to this problem in closed form before, both through a convolution with the fundamental solution $G(x, t)$ and through Fourier synthesis. On the other, for smooth enough data $f(x)$, we could compute all normal derivatives of u at $y = 0$:

$$\frac{\partial^{n+1} u(x, 0)}{\partial y^{n+1}} = \frac{\partial^2}{\partial x^2} \frac{\partial^n u(x, 0)}{\partial y^n},$$

and do a Taylor expansion à la Cauchy-Kovalewsky. Yet the problem here is that, for the heat equation, information propagates infinitely fast. Hence, even if $f(x)$ is zero in an interval, making the expansion above equal to zero to all orders, the solution for any positive time t may be made arbitrarily far from zero by picking initial data sufficiently big outside of the interval where $f(x)$ is zero. This can be done explicitly, for instance using the exact solution that we computed long ago for initial data given by the step function (page 49): the function $\phi(x, y, t)$ satisfying the heat equation

$$\phi_t = \phi_{xx}$$

with piecewise constant initial data

$$\phi(x, y, 0) = \begin{cases} f(y) & \text{for } x \le y \\ 0 & \text{for } x > y \end{cases}$$

is given by the complementary error function

$$\phi(x, y, t) = \frac{f(y)}{\sqrt{4\pi}} \int_{\frac{x-y}{\sqrt{t}}}^{\infty} e^{-\frac{s^2}{4}}\, ds.$$

Fixing $x = 0$ and $t = \tau$ and letting y approach $-\infty$, we obtain to leading order

$$\phi(0, y, \tau) \approx \sqrt{\frac{\tau}{\pi}} \frac{f(y)}{|y|}\; e^{-\frac{y^2}{4\tau}}.$$

Then we can define

$$f(y) = \sqrt{\frac{\pi}{\tau}} |y|\; e^{\frac{y^2}{4\tau}},$$

and obtain a family of initial value problems for the heat equation that, as $y \to -\infty$, has initial value equal to zero for all finite values of x, yet solution equal to one at $x = 0$, $t = \tau$, contradicting the result based on the Taylor expansion above.

The solution that we have built grows unboundedly for large negative values of x. It can shown that the exact answer that we found through convolution with $G(x, t)$ is the only one consistent with a decay condition at infinity.