EXTENDING CALCULUS: DERIVATIVES

ARICK SHAO

1. Directional Derivatives in $\mathbb{R}^n$

One learns early in first-year calculus the definition of the derivative, and its basic interpretation. To review, consider a function

$$f : (a, b) \to \mathbb{R},$$

and consider the limit

$$\lim_{h \to 0} \frac{f(x+h) - f(x)}{h} = \lim_{h \to 0} \frac{f(x+h) - f(x)}{(x+h) - x}, \qquad a < x < b.$$

Fixing a displacement value h, consider the line through the two points $(x, f(x))$ and $(x+h, f(x+h))$ on the graph of f. From high school algebra, we know that the slope of the line is given precisely by the quotient

$$\frac{f(x+h) - f(x)}{(x+h) - x}.$$

Thus, in the above limit, we are computing the slope of the line through $(x, f(x))$ and $(x+h, f(x+h))$, and we are exploring what happens to this slope when the latter point slides progressively closer to the former along the graph of f.

If the limit of the above quotient exists, then we say that f is *differentiable* at x. In this case, the *derivative* of f at x is defined to be that limit:

$$f'(x) = \lim_{h \to 0} \frac{f(x+h) - f(x)}{h}.$$

If f is sufficiently regular to be differentiable at x, then the family of lines described above tends in the limit $h \to 0$ to a line through $(x, f(x))$, with slope $f'(x)$.

Visually, this aforementioned line can be interpreted as the one tangent to the graph of f at $(x, f(x))$. One can also think of this line as the best linear approximation of the graph of f near $(x, f(x))$. In physics, one often interprets $f'(x)$ to be the instantaneous rate of change of f at x.

A large variety of applications exist for derivatives. Here, we focus on the main application discussed in elementary calculus courses: finding the extrema of a function. Suppose f, as given above, has a *local maximum* at $x \in (a, b)$. More specifically, suppose there is some subinterval

$$I = (x - \epsilon, x + \epsilon)$$

of (a, b), containing x, such that $f(y) \leq f(x)$ for all points $y \in I$. In this case, can we say anything about the derivative $f'(x)$, if it exists?

To answer this question, we consider the following cases:

- If $f'(x) = c > 0$, then near the point $(x, f(x))$, the graph of f behaves like the line with positive slope c. Consequently, f, behaving locally like this line, must be strictly increasing at x. This implies that $f(x+h) > f(x)$ for some small $h > 0$, so f cannot have a local maximum at x.
- Similarly, if $f'(x) = c < 0$, then near $(x, f(x))$, the graph of f behaves like the line with negative slope c. This implies $f(x-h) > f(x)$ for some small $h > 0$, so again, f cannot have a local maximum at x.

As a result, we can conclude that *if f has a local maximum at x, then either $f'(x)$ vanishes, or f fails to be differentiable at x.* By a completely analogous reasoning, the same conclusion also holds if f has a *local minimum* at x.

Therefore, we define the *critical points* of f to be the points $x \in (a, b)$ for which $f'(x)$ either vanishes or fails to exist. This is directly related to this problem of finding the minima and maxima (i.e., the *extrema*) of f. Indeed, *in order to find these extrema, one needs only consider the critical points of f.*

More specifically, in order to find the extrema of f:

(1) Find the critical points of f: solve the equation $f'(x) = 0$, and determine for which points f' fails to exist.
(2) For each such critical point, check (e.g., directly, or using a second derivative test) whether f has a local extremum there.
(3) To find the *absolute extrema* of f, compare the values of f at all its critical points, as well as at the boundaries a and b.

This is a very general and systematic method, which can be extended to higher dimensions. This includes functions on $\mathbb{R}^n$, which one encounters within vector calculus, as well as functions on infinite dimensional spaces.

1.1. **Directional Derivatives.**

Recall that for the single-variable function, its derivative represents the rate of change of that function. However, for functions of multiple variables, the notion of "rate of change" does not quite make sense.

Consider a function $f : \mathbb{R}^2 \to \mathbb{R}$ of two variables, and fix a point $(x_0, y_0) \in \mathbb{R}^2$. Starting from (x_0, y_0), the value of f can change in dramatically different ways if one goes in different directions from (x_0, y_0). For example, if $f(x, y) = x^2 - y^2$, then f behaves like an upward-opening parabola in the horizontal direction and like a downward-opening parabola in the vertical direction. More generally, if we consider the function $f(x, y) = g(x)h(y)$, then f behaves like g in horizontal directions and like h in vertical directions. Thus, it is not particularly meaningful to define something to be "the rate of change of f". As a result, in order to properly define some notion derivative of f, we will have to further refine our intuitions.

Consider now a line ℓ through (x_0, y_0). Recall ℓ is described completely by:

- The point (x_0, y_0), which is on ℓ.
- The direction of ℓ, which is given by a vector $(v_x, v_y) \in \mathbb{R}^2$.

Given these two parameters, one can now describe ℓ parametrically by

$$\ell(t) = (x_0, y_0) + t(v_x, v_y).$$

In particular, $\ell(0) = (x_0, y_0)$, and by varying $t \in \mathbb{R}$, one can trace out all of ℓ.

Let us now restrict our map f to ℓ. This reduces f to a single-variable function,

$$f_\ell : \mathbb{R} \to \mathbb{R}, \qquad f_\ell(t) = f(\ell(t)),$$

for which the notions of "rate of change" and derivative make sense. Indeed, the derivative $f_\ell'(0)$, if it exists, gives the rate of change of f at $\ell(0) = (x_0, y_0)$ in the direction of ℓ. More accurately, $f'_\ell(0)$ describes the rate of change of f at (x_0, y_0) as one moves (only) in the direction (v_x, v_y).

By the definition of single-variable derivatives, we see that

$$\begin{aligned} f'_\ell(0) &= \lim_{h\to 0} \frac{f(\ell(h)) - f(\ell(0))}{h} \\ &= \lim_{h\to 0} \frac{f((x_0, y_0) + h(v_x, v_y)) - f(x_0, y_0)}{h} \\ &= \lim_{h\to 0} \frac{f(x_0 + hv_x, y_0 + hv_y) - f(x_0, y_0)}{h}. \end{aligned}$$

This motivates our following definition: we define the *directional*, or *Gâteaux*, [1] *derivative* of f in the direction (v_x, v_y) at the point (x_0, y_0) to be the quantity

$$df|_{(x_0,y_0)}(v_x, v_y) = \lim_{h\to 0} \frac{f((x_0, y_0) + h(v_x, v_y)) - f(x_0, y_0)}{h}.$$

Again, this represents the rate of change of f, *in the direction* (v_x, v_y).

In particular, we define the *partial derivatives* of f to be the directional derivatives in the standard directions $(1, 0)$ and $(0, 1)$:

$$\partial_x f(x_0, y_0) = df|_{(x_0,y_0)}(1, 0), \qquad \partial_y f(x_0, y_0) = df|_{(x_0,y_0)}(0, 1).$$

Exercise 1. *Compute* $\partial_x f(x, y)$, $\partial_y f(x, y)$, *and* $df|_{(x,y)}(v_x, v_y)$ *for*

$$f(x, y) = x + y, \qquad f(x, y) = xy, \qquad f(x, y) = e^{xy} \cos y.$$

Exercise 2. *Let* a *be a real number, and consider the function*

$$f(x, y) = \begin{cases} \frac{2x^2+xy-3y^2}{4x^2+y^2} & (x, y) \neq (0, 0), \\ a & (x, y) = (0, 0). \end{cases}$$

For what value of a *does the partial derivative* $\partial_y f(0, 0)$ *exist?*

Now, there is absolutely nothing stopping us from extending the above directly to higher-dimensional spaces. Suppose $f : \mathbb{R}^n \to \mathbb{R}$ is a function of n variables, and fix a point $\mathbf{r}_0 \in \mathbb{R}^n$. If the vector $\mathbf{v} \in \mathbb{R}^n$ represents a direction in $\mathbb{R}^n$, then the line ℓ through $\mathbf{r}_0$ in the direction $\mathbf{v}$ can be parametrized as

$$\ell(t) = \mathbf{r}_0 + t\mathbf{v}.$$

Therefore, we can define the *directional derivative* of f in the direction $\mathbf{v}$ at $\mathbf{r}_0$ by

$$df|_{\mathbf{r}_0}(\mathbf{v}) = \lim_{h\to 0} \frac{f(\mathbf{r}_0 + h\mathbf{v}) - f(\mathbf{r}_0)}{h}.$$

Like in the 2-dimensional case, we can rewrite this as a single-variable derivative,

$$df|_{\mathbf{r}_0}(\mathbf{v}) = f'_\ell(0) = \frac{d}{dt}[f(\mathbf{r}_0 + t\mathbf{v})]\Big|_{t=0}.$$

Again, the *partial derivatives* are just special directional derivatives:

$$\partial_1 f(\mathbf{r}_0) = df|_{(\mathbf{r}_0)}(1, 0, \ldots, 0), \qquad \partial_2 f(\mathbf{r}_0) = df|_{(\mathbf{r}_0)}(0, 1, 0, \ldots, 0), \qquad \text{etc.}$$

Exercise 3. *Compute* $\partial_k f(\mathbf{r})$ *and* $df|_{\mathbf{r}}(\mathbf{v})$ *for the function*

$$f : \mathbb{R}^n \setminus \{(0, \ldots, 0)\} \to \mathbb{R}, \qquad f(\mathbf{r}) = |\mathbf{r}|.$$

[1] The name "Gâteaux derivative" is generally reserved for analogous constructions in infinite-dimensional spaces. In vector calculus, the term "directional derivative" is almost always used.

1.2. **Finding Extrema.** For single-variable functions, one could greatly simplify the task of finding local or global extrema of that function by considering only its critical points, i.e., points for which the derivative of the function either vanishes or fails to exist. Now, we see how one can extend this to functions of n variables.

Let D be some sufficiently nice region in $\mathbb{R}^n$, and consider a function $f : D \to \mathbb{R}$ of n variables. Suppose that f has a local maximum, say, at some point $\mathbf{r}_0$ in the interior of D. In other words, near $\mathbf{r}_0$, the values of f are at most that of $f(\mathbf{r}_0)$. One can rephrase this more rigorously using distances and epsilons:

There is some $\epsilon > 0$ such that if $\mathbf{r} \in D$ and $|\mathbf{r} - \mathbf{r}_0| < \epsilon$, then $f(\mathbf{r}) \leq f(\mathbf{r}_0)$.

Again, this is the direct analogue of the definition for single variable functions. Local minima of f are defined in a completely analogous manner.

Let us now relate these notions to directional derivatives. Fix $\mathbf{v} \in \mathbb{R}^n$, and consider the line $\ell_{\mathbf{r}_0,\mathbf{v}}$ through $\mathbf{r}_0$ in the direction $\mathbf{v}$, given parametrically by

$$\ell_{\mathbf{r}_0,\mathbf{v}}(t) = \mathbf{r}_0 + t\mathbf{v}.$$

As $\mathbf{r}_0$ is a local maximum of f, then $\mathbf{r}_0$ is also a local maximum of the restriction of f to $\ell_{\mathbf{r}_0,\mathbf{v}}$. Consequently, the single variable function

$$t \mapsto g(t) = f(\ell_{\mathbf{r}_0,\mathbf{v}}(t))$$

must have a local maximum at $t = 0$. From the single variable theory, we have

$$g'(t) = df|_{\mathbf{r}_0}(\mathbf{v}) = 0.$$

Here, the first equality follows from the definition of directional derivatives.

Note that we did not assume anything special about the direction $\mathbf{v}$. In fact, the above holds for *any direction* $\mathbf{v}$ from $\mathbf{r}_0$. Therefore, we have proved the following:

Proposition 1.1. *If f has a local extremum (maximum or minimum) at $\mathbf{r}_0 \in \mathbb{R}^n$, then for any direction $\mathbf{v} \in \mathbb{R}^n$ for which $df|_{\mathbf{r}_0}(\mathbf{v})$ exists, we have*

$$df|_{\mathbf{r}_0}(\mathbf{v}) = 0.$$

As a consequence of Proposition 1.1, in order to find the local extrema of f, we need only look at the *critical points* of f, that is, the points $\mathbf{r}$ where either

- All directional derivatives of f at $\mathbf{r}$ vanish, or
- Some directional derivative of f at $\mathbf{r}$ fails to exist.

Remark. *If f is sufficiently regular, then all directional derivatives of f are determined by merely the partial derivatives of f. More specifically, if*

$$\mathbf{v} = (v_1, v_2, \dots, v_n),$$

and if f is "differentiable", [2] *then*

$$df|_{\mathbf{r}_0}(\mathbf{v}) = v_1\partial_1 f(\mathbf{r}_0) + v_2\partial_2 f(\mathbf{r}_0) + \cdots + v_n\partial_n f(\mathbf{r}_0).$$

In particular, in order to find a critical point $\mathbf{r}_0$ of f, then one needs only show that all the partial derivatives of f at $\mathbf{r}_0$ vanishes.

Remark. *Like for the single variable case, this* first derivative test *only applies to points in the* interior *of the domain D of f. For points on the boundary of D, one must resort to other methods, e.g., direct inspection or Lagrange multipliers.*

[2]We do not provide the definition of "differentiability" of multivariable functions here, as it would require a somewhat lengthy explanation. However, this "differentiability" is a stronger condition than simply the existence of certain directional derivatives.

Exercise 4. *Let D be the interior of the unit disk in $\mathbb{R}^2$ about the origin, i.e.,*

$$D = \{(x, y) \mid x^2 + y^2 < 1\}.$$

(1) *Find the critical points of the functions*

$$f(x, y) = x^2 + y^2, \qquad g(x, y) = (x - \frac{1}{2})^2 - (y + \frac{1}{2})^2$$

in the interior of D.

(2) *For each critical point of f, determine directly whether it is a local maximum of f, a local maximum of f, or neither.*

(3) *For each critical point of g, determine directly whether it is a local maximum of g, a local maximum of g, or neither.*

2. An Infinite-Dimensional Model Problem

The next task is to extend this notion of directional derivatives we have developed on the finite-dimensional spaces $\mathbb{R}^n$ to infinite-dimensional spaces. In order to keep the amount of technical details down, we restrict our attention to a single "model problem", which we will proceed to solve using these directional derivatives. Here, we will demonstrate that *the shortest curve between two points $A, B \in \mathbb{R}^n$ must be the line segment between A and B.* Although this is a rather particular problem, with an intuitively obvious answer, it does serve to demonstrate the role played by these generalized derivatives in infinite-dimensional settings.

2.1. **The Arc Length Functional.**

Fix distinct points $A, B \in \mathbb{R}^n$. Let

$$\alpha : [a, b] \to \mathbb{R}^n$$

denote a sufficiently smooth parametrized curve from A to B, i.e.,

$$\alpha(a) = A, \qquad \alpha(b) = B.$$

Then, the total *arc length* of the curve α is given by

$$\mathcal{L}(\alpha) = \int_a^b |\alpha'(t)| dt.$$

This formula can be justified as follows. Suppose $\alpha(t)$ represents the position of a particle at time t, so that the curve defined by α represents the trajectory of this particle. Then, the derivative $\alpha'(t)$ represents the velocity of the particle at time t, and its magnitude $|\alpha'(t)|$ is the speed of this particle. Therefore, the above formula for arc length simply states that if you integrate (i.e., "sum") the speed of this particle over time, then you obtain the total distance traveled by the particle.

Now, there is an extra degree of freedom in our setting that we have not exploited yet: the fact that a single curve can be parametrized in (infinitely) many different ways. For example, the curve represented by α can also be parametrized as

$$\alpha_1 : [0, b - a] \to \mathbb{R}^n, \qquad \alpha_1(t) = \alpha(t + a),$$
$$\alpha_2 : [0, 1] \to \mathbb{R}^n, \qquad \alpha_2(t) = \alpha_1[(b - a)t].$$

Consequently, given a fixed curve C in $\mathbb{R}^n$, we have the freedom to choose a convenient parametrization α of C which is best adapted to the situation at hand.

The specific parametrization we wish to consider is the *arc length parametrization.* If C is a curve from A to B, then this is the parametrization

$$\beta : [0, T] \to \mathbb{R}^n$$

which satisfies the conditions

$$\beta(0) = A, \qquad \beta(T) = B, \qquad |\beta'(t)| \equiv 1.$$

As the name suggests, this parametrization is related to the arc length. For any $0 \leq t \leq T$, the length of the segment of the curve between $\beta(0)$ and $\beta(t)$ is

$$\int_0^t |\beta'(s)| ds = \int_0^t ds = t.$$

In particular, the length of C is just

$$\int_0^T |\beta'(s)| ds = T.$$

The first question to ask is whether such an arc length parametrization actually exists. Suppose $\alpha : [0, b] \to \mathbb{R}^n$ is an arbitrary parametrization of C. We wish to reparametrize α into our desired arc length parametrization, that is, we wish to find some monotone increasing function $\sigma : [0, T] \to [0, b]$ such that $\beta(t) = \alpha(\sigma(t))$ is the arc length parametrization. Since we require that

$$1 = |\beta'(t)| = \left| \frac{d}{dt} \alpha(\sigma(t)) \right| = |\alpha'(\sigma(t))| \sigma'(t)$$

for all t, where in the last step, we applied the chain rule, this implies that in order to obtain the arc length parametrization β, we must solve the differential equation

$$\sigma'(t) = \frac{1}{|\alpha'(\sigma(t))|}, \qquad \sigma(0) = 0.$$

This can always be done using iteration-type methods that we have discussed, as long as α itself is a "reasonable" parametrization, in that α' does not vanish. [3]

2.2. **Critical Points.** Now that we have defined the arc length functional $\mathcal{L}$ and constructed general arc length parametrizations, we can now turn our attention to the main problem. [4] Again, we wish to find the shortest curve from A to B. In other words, *we want to minimize the functional $\mathcal{L}$, given the constraint that the curves we consider all must travel from A to B.*

The first step is to connect this to the finite-dimensional problem discussed in the preceding section. Suppose a curve $\alpha : [0, b] \to \mathbb{R}^n$ is such a minimizer of $\mathcal{L}$. If we were to deform α just a little bit, but still as curve from A to B, then this deformed curve $\bar{\alpha}$ from A to B cannot be shorter than α, i.e.,

$$\mathcal{L}(\bar{\alpha}) \geq \mathcal{L}(\alpha).$$

In this sense, we can think of α as a minimum of $\mathcal{L}$ (with the constraint that α goes from A to B). We would like to show that α can only be a minimum of $\mathcal{L}$ if α is a "critical point" of $\mathcal{L}$, i.e., that all "directional derivatives" of $\mathcal{L}$ vanish at α.

Let us be now a bit more specific, and consider a *linear* deformation of α. If $\gamma : [0, b] \to \mathbb{R}^n$ is another curve, then $\alpha + \varepsilon\gamma$ represents a deformation of α for any $\varepsilon \in \mathbb{R}$. In order for $\alpha + \varepsilon\gamma$ to also be a curve from A to B, though, we require that this deformation satisfies $\gamma(0) = \gamma(b) = \mathbf{0}$, i.e., that γ has vanishing endpoints.

Since α is assumed to be a minimizer of $\mathcal{L}$, then

$$\mathcal{L}(\alpha + \varepsilon\gamma) \geq \mathcal{L}(\alpha).$$

[3] Note that T here must be the length of α.

[4] The term "functional" refers to functions whose arguments are functions rather than numbers.

In other words, if we define the function

$$g : \mathbb{R} \to \mathbb{R}, \qquad g(\varepsilon) = \mathcal{L}(\alpha + \varepsilon\gamma),$$

then g must have a minimum at 0. As a result,

$$g'(0) = \frac{d}{d\varepsilon}\mathcal{L}(\alpha + \varepsilon\gamma)\bigg|_{\varepsilon=0} = 0.$$

Moreover, if we think of the function α as a "point" in our "infinite-dimensional space of curves", then one can think of the map

$$\ell(\varepsilon) = \alpha + \varepsilon\gamma$$

as parametrizing a line in this infinite-dimensional space. Then, the function g given above is precisely the restriction of $\mathcal{L}$ to this line ℓ! Therefore, this quantity

$$g'(0) = \frac{d}{d\varepsilon}\mathcal{L}(\alpha + \varepsilon\gamma)\bigg|_{\varepsilon=0} = \frac{d}{d\varepsilon}\mathcal{L}(\ell(\varepsilon))\bigg|_{\varepsilon=0}$$

represents the *directional derivative of $\mathcal{L}$ in the direction γ*.

From this perspective, we see that if α is a minimizer of $\mathcal{L}$, then the *directional derivative* of $\mathcal{L}$ in the direction γ, i.e., the quantity

$$d\mathcal{L}|_\alpha(\gamma) = \frac{d}{d\varepsilon}\mathcal{L}(\alpha + \varepsilon\gamma)\bigg|_{\varepsilon=0},$$

must vanish. Furthermore, the above is true for *any* "direction curve" γ with vanishing endpoints. In this sense, we can think of α as a *critical point* of $\mathcal{L}$.

Finally, from this argument, we have shown that in order to find potential minimizers of $\mathcal{L}$, for curves from A to B, then we need only look at corresponding critical points of $\mathcal{L}$, i.e., curves α such that the directional derivative $d\mathcal{L}|_\alpha(\gamma)$ vanishes for all directions γ with vanishing endpoints. This greatly reduces the number of curves which could possibly be a minimizer. Note this is the direct analogue of the finite-dimensional case - that local minima can only occur at critical points.

2.3. **Finding Critical Points.** It remains to solve for the critical points of $\mathcal{L}$. In other words, we wish to find a curve from A to B, parametrized by α, such that

$$d\mathcal{L}|_\alpha(\gamma) = \frac{d}{d\varepsilon}\mathcal{L}(\alpha + \varepsilon\gamma)\bigg|_{\varepsilon=0} = 0$$

for all curves $\gamma : [0, b] \to \mathbb{R}^n$ with vanishing endpoints. Since we have the freedom to parametrize α as we see fit, we will take the arc length parametrization described earlier. More specifically, we assume that

$$b = T, \qquad |\alpha'(t)| = 1.$$

Given a direction curve γ as described above, then

$$\begin{aligned}
\frac{d}{d\varepsilon}\mathcal{L}(\alpha + \varepsilon\gamma)\bigg|_{\varepsilon=0} &= \int_0^T \frac{d}{d\varepsilon}|\alpha'(t) + \varepsilon\gamma'(t)|dt\bigg|_{\varepsilon=0} \\
&= \int_0^T \frac{d}{d\varepsilon}\sqrt{[\alpha'(t) + \varepsilon\gamma'(t)] \cdot [\alpha'(t) + \varepsilon\gamma'(t)]}dt\bigg|_{\varepsilon=0} \\
&= \frac{1}{2}\int_0^T \frac{\frac{d}{d\varepsilon}\{[\alpha'(t) + \varepsilon\gamma'(t)] \cdot [\alpha'(t) + \varepsilon\gamma'(t)]\}}{|\alpha'(t) + \varepsilon\gamma'(t)|}\bigg|_{\varepsilon=0}
\end{aligned}$$

$$= \int_0^T \frac{[\alpha'(t) + \varepsilon\gamma'(t)] \cdot \gamma'(t)}{|\alpha'(t) + \varepsilon\gamma'(t)|}\Bigg|_{\varepsilon=0}$$
$$= \int_0^T \frac{\alpha'(t)}{|\alpha'(t)|} \cdot \gamma'(t) dt.$$

In the above, we applied both the power rule and the chain rule. In particular, "$\cdot$" represents the dot product of n-dimensional vectors. Recalling the arc length parametrization condition $|\alpha'(t)| \equiv 1$, then we wish to solve

$$0 = \frac{d}{d\varepsilon} \mathcal{L}(\alpha + \varepsilon\gamma)\Bigg|_{\varepsilon=0} = \int_0^T \alpha'(t) \cdot \gamma'(t) dt.$$

Since γ has vanishing endpoints, then we can integrate by parts:

$$\begin{aligned} 0 &= \int_0^T \alpha'(t) \cdot \gamma'(t) dt \\ &= \alpha'(T) \cdot \gamma(T) - \alpha'(0) \cdot \gamma(T) - \int_0^T \alpha''(t) \cdot \gamma(t) dt \\ &= - \int_0^T \alpha''(t) \cdot \gamma(t) dt. \end{aligned}$$

Note that the above must hold for *any* γ with vanishing endpoints. Since we have such a vast variety of choices for γ, then the above equation can hold only when $\alpha''(t) \equiv 0$. Indeed, if $\alpha''(t)$ fails to vanish everywhere, then one can easily construct such a direction curve γ such that the integral of $\alpha'' \cdot \gamma$ fails to vanish.

We have now shown that the critical points of $\mathcal{L}$ are the curves α (from A to B) such that α'' vanishes identically. [5] There are only a limited number of candidates for such curves, since by integrating, we see that

$$\alpha'(t) \equiv \mathbf{v}_0, \qquad \mathbf{v}_0 \in \mathbb{R}^n,$$
$$\alpha(t) \equiv \mathbf{r}_0 + t\mathbf{v}_0, \qquad \mathbf{r}_0 \in \mathbb{R}^n.$$

Since $\alpha(0) = A$ and $\alpha(T) = B$, then $\mathbf{r}_0$ and $\mathbf{v}_0$ are uniquely determined:

$$\mathbf{r}_0 = A, \qquad \mathbf{v}_0 = \frac{B - A}{T}.$$

To summarize, the above reasoning implies that the only critical point of $\mathcal{L}$ is the curve C with the (arc length) parametrization

$$\alpha : [0, T] \to \mathbb{R}^n, \qquad \alpha(t) = A + t \cdot \frac{B - A}{T}.$$

Note, however, that this curve C is precisely the line segment from A to B!

It is rather intuitively clear that for any two points A and B, there must be a shortest curve from A to B. [6] Since this shortest curve must be a critical point of $\mathcal{L}$, as described above, then it must in fact be the line segment from A to B. This completes the solution of our model problem.

[5] Again, we are also adding a "free" constraint of requiring arc length parametrization.

[6] We will not give a technical proof of this here. This is in fact a part of the *Hopf-Rinow theorem*, a well-known result in differential geometry.

2.4. **Additional Remarks.** In our model problem, we minimized the arc length functional $\mathcal{L}$ on the flat space $\mathbb{R}^n$. This same process can be extended to curved spaces as well, such as spheres, hyperboloids, and so on. We can similarly define the arc length functional on these curved spaces. Thus, we can pose the same question: what are the minimizing curves of this arc length functional?

The answer is once again tied to finding the corresponding critical points of $\mathcal{L}$. By performing a similar calculation as before, one can see that critical points are the curves (or rather, the arc length parametrization α of the curve) which satisfy a differential equation. [7] The solution curves of this differential equation, i.e., the critical points of $\mathcal{L}$, are called *geodesics*. Thus, any minimizing curve from two points A and B along a curved space must be such a geodesic. In particular, in $\mathbb{R}^n$, the geodesic curves are precisely the straight lines.

In contrast to the flat case $\mathbb{R}^n$, in curved spaces, a geodesic need not necessarily be a minimizer of $\mathcal{L}$. For example, if our curved space is the unit sphere

$$\mathbb{S}^2 = \{(x, y, z) \mid x^2 + y^2 + z^2 = 1\},$$

then the geodesics of $\mathbb{S}^2$ are precisely the great circles, i.e., the circles in $\mathbb{S}^2$ of unit radius. Thus, for any two points $A, B \in \mathbb{S}^2$, any segment along a great circle connecting A and B is a geodesic segment. If A and B are nonantipodal points, then A and B can be connected by two arcs of differing lengths which together form a great circle. Then, at least one of these arcs will not minimize length.

In differential geometry, a major question (which has been answered) is to determine when a geodesic between two points in a curved space actually minimizes the length. The answers, in fact, depend fundamentally on how the space is curved.

References

1. B. O'Neill, *Semi-Riemannian geometry with applications to relativity*, Academic Press, 1983.
2. J. Stewart, *Multivariable calculus*, 7 ed., Brooks/Cole, 2012.

[7] In the case of $\mathbb{R}^n$, the differential equation is $\alpha''(t) \equiv 0$.