CISC 371 Class 3

Stationarity and Convexity

Texts: [1] pp. 7–8; [2] pp. 7–8, 14–17

Main Concepts:
- *Stationary point, minimizer, and saddle point*
- *First-order condition*
- *Convex function*
- *Gradient inequality*

Sample Problem, Signal Processing: For the step response of a simple dynamical system that oscillates, why does the quadratic approximation change substantially for a small change in the initial estimate of the minimizer?

From prerequisite material in basic calculus, we know that the maximum or minimum value of a function occurs when the derivative is zero. For the remainder of this course, we will assume that an objective function f, that has a scalar argument, is:

Continuous: satisfies the Weierstrass /Jordan definitions of continuity
Differentiable: satisfies the existence condition everywhere
Continuously differentiable: the derivative is continuous

As needed, we can restrict the domain of f to be a subset of the real numbers.

To better understand what can happen when we try to optimize a function that has a scalar argument, we can recall some of the necessary and sufficient conditions for there to be a strict local minimizer of the objective.

3.1 Stationary Point of a Function of a Scalar Argument

A *stationary point* of a function is a point at which the derivative of a function is zero.

Definition: stationary point of $f(t)$

For any $t \in \mathbb{R}$ and any $f : \mathbb{R} \to \mathbb{R}$, a scalar $t^* \in \mathbb{R}$ is a *stationary point* of f is defined as

$$f'(t^*) = 0 \tag{3.1}$$

Observation: By definition, a stationary point is a point at which an infinitesimal change in the argument produces zero change in the value of the function. This implies that stationarity is a necessary condition for optimality.

Theorem: first-order necessary condition for optimality

If $t^* \in \mathbb{R}$ *is a local minimizer of* f*, then* t^* *is a stationary point*

Proof: See Theorem 3.9 in the extra notes for this class

Observation: Stationarity is a necessary condition for a local minimizer, for a strict local maximizer, and for a local maximizer. Theorem 3.9 is often called a first-order condition for the existence of a minimizer.

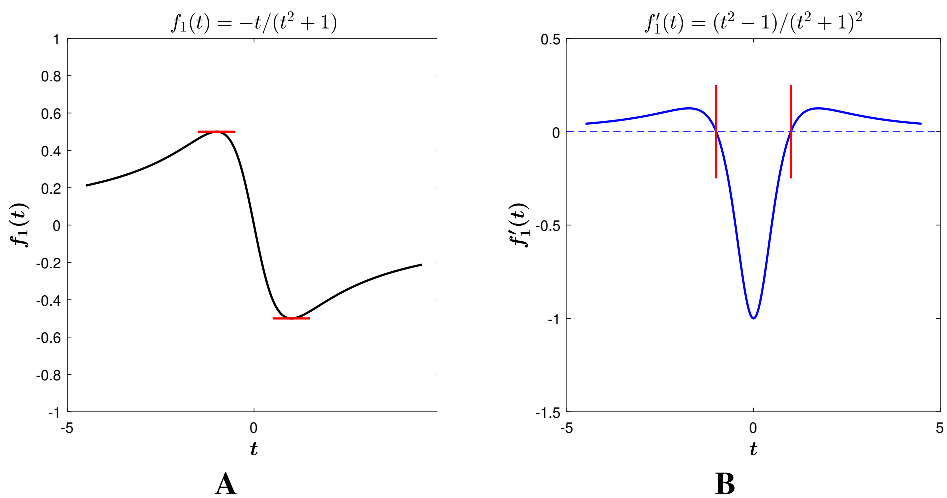

Figure 3.1: The function $f_1(t) = -t/(t^2+1)$ has two stationary points. (A) The function is shown in black and stationary points are shown in red. The stationary point at $t = +1$ is a local minimizer and the stationary point at $t = -1$ is a local maximizer. (B) The derivative of f_1 is shown in blue and its zeros are shown in red. Zeros of the derivative are stationary points of the function.

Stationarity is a necessary condition for optimality but it is not sufficient. An important kind of stationary point is a *saddle point*, which is neither a local minimizer nor a local maximizer.

Definition: saddle point of $f(t)$

For any $t \in \mathbb{R}$, and any $f : \mathbb{R} \to \mathbb{R}$, a scalar $t^* \in \mathbb{R}$ is a *saddle point* of f is defined as

$$(t^* \text{ is a stationary point}) \wedge \neg(t^* \text{ is a local maximizer}) \wedge \neg(t^* \text{ is a local minimizer}) \tag{3.2}$$

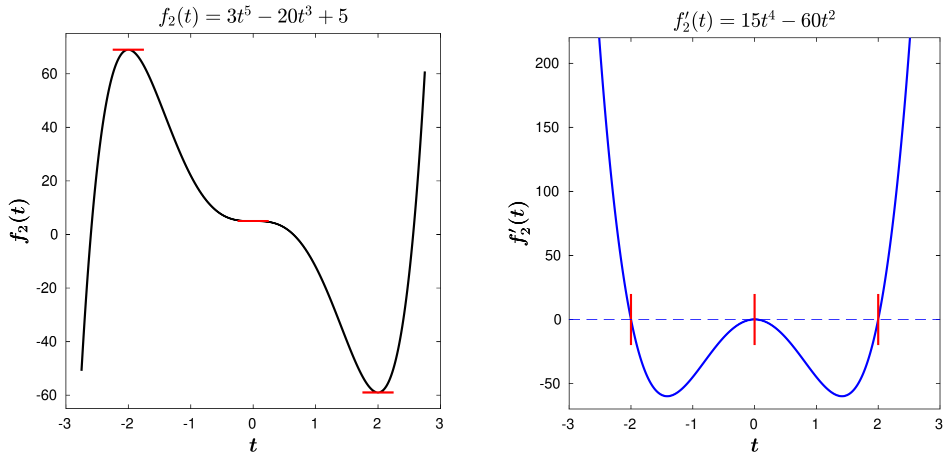

A **B**

Figure 3.2: The function $f_2(t) = 3t^5 - 20t^3 + 5$ has three stationary points. (A) The function is shown in black and stationary points are shown in red. The stationary point at $t = +2$ is a local minimizer; the stationary point at $t = -2$ is a local maximizer; the stationary point at $t = 0$ is a saddle point. (B) The derivative of f_2 is shown in blue and its zeros are shown in red.

For a function that has a scalar argument, using a higher derivative is suggestive and may not be conclusive. For example, the function $f_2 = 3t^5 - 20t^3 + 5$ has three stationary points, shown in Figure 3.2. The second derivative correctly identifies the local maximizer and the local minimizer but is inconclusive at the saddle point. Another example is the function $f_3(t)$ that is illustrated in Figure 3.3, for which all of the derivatives are zero at the saddle point.

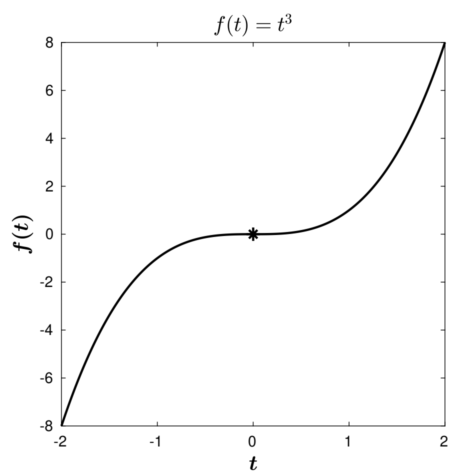

Figure 3.3: A graph of the function $f_3(t) = t^3$ shows that the only stationary point, which is $t^* = 0$, is a saddle point. At $t^* = 0$, the second derivative is also zero and a second-order test is inconclusive.

From prerequisite elementary calculus, we recall that a function with a scalar argument has three fundamental kinds of stationary points t^* that can also be classified according to the sign of the second derivative of the function:

- A *strict local minimizer* has $f''(t^*) > 0$
- A *strict local maximizer* has $f''(t^*) < 0$
- A *saddle point* has $f''(t^*) = 0$, with t^* not a local minimizer and not a local maximizer

3.2 Convexity of a Function of a Scalar Argument

In optimization, the concept of being "concave-up" is called *convexity*. Functions can be convex and sets can be convex; for now, we only need to understand the idea of a convex function that has a scalar argument. In simple words, a convex function is where, if we pick any two points, function is "below" the line that "connects" the points.

From prerequisite material in linear algebra, a line segment is the linear interpolation between two vectors. If we have two vectors in a vector space, such as $\vec{u}$ and $\vec{v}$, then the line segment that "connects" the vectors can be described by a free scalar variable θ, where $0 \leq \theta \leq 1$. A vector on this line segment can be written as

$$(1-\theta)\vec{u} + \theta\vec{v} \quad \text{for} \quad 0 \leq \theta \leq 1 \tag{3.3}$$

For a function f that has a scalar argument, which is $f : \mathbb{R} \rightarrow \mathbb{R}$, we can select any point $t_1 \in \mathbb{R}$ and any point $t_2 \in \mathbb{R}$. We can define 2D vectors that are

$$\vec{u} \stackrel{\text{def}}{=} \begin{bmatrix} t_1 \\ f(f_1) \end{bmatrix} \qquad \vec{v} \stackrel{\text{def}}{=} \begin{bmatrix} t_2 \\ f(f_2) \end{bmatrix} \tag{3.4}$$

Substituting the terms of Equation 3.4 into Equation 3.3, a point $(1-\theta)t_1 + \theta t_2$ corresponds to a point on the line segment that has a second entry equal to $(1-\theta)f(t_1) + \theta f(t_2)$. This is illustrated in Figure 3.4, showing the case of $\theta = 0.6$ as an example.

We can define a strictly convex function using these ideas.

Definition: strictly convex function $f(t)$

> For any $t_1 \in \mathbb{R}$, any $t_2 \in \mathbb{R}$, any $\theta \in \mathbb{R}_{++}$, and any $f : \mathbb{R} \rightarrow \mathbb{R}$, that function f is a *strictly convex function* is defined as
>
> $$((t_1 \neq t_2) \wedge (0 < \theta < 1)) \rightarrow (f((1-\theta)t_1 + \theta t_2)) < ((1-\theta)f(t_1) + \theta f(t_2))) \tag{3.5}$$

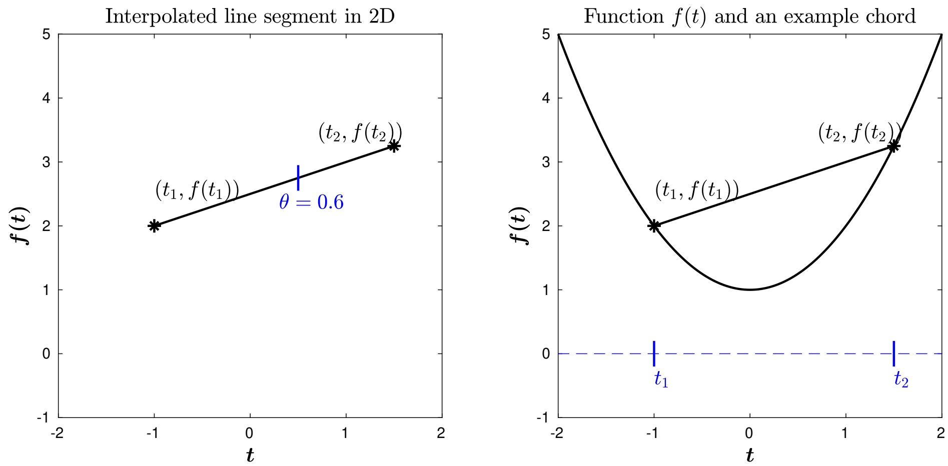

A **B**

Figure 3.4: A function $f(t)$ is evaluated at two points. (A) The points, and the values of the function, can be interpreted as 2D vectors that a line segment "connects". (B) Any value of $0 \leq \theta \leq 1$ interpolates the vectors; the interpolation for $\theta = 0.6$ is shown in blue.

Observation: The line segment that "connects" the points and valuations is a *chord* of the graph of the function. A strictly convex function that is evaluated at any point t_3, such that $t_1 < t_3 < t_2$, will have $f(t_3)$ that is strictly less than the value on the chord. This is illustrated in Figure 3.5.

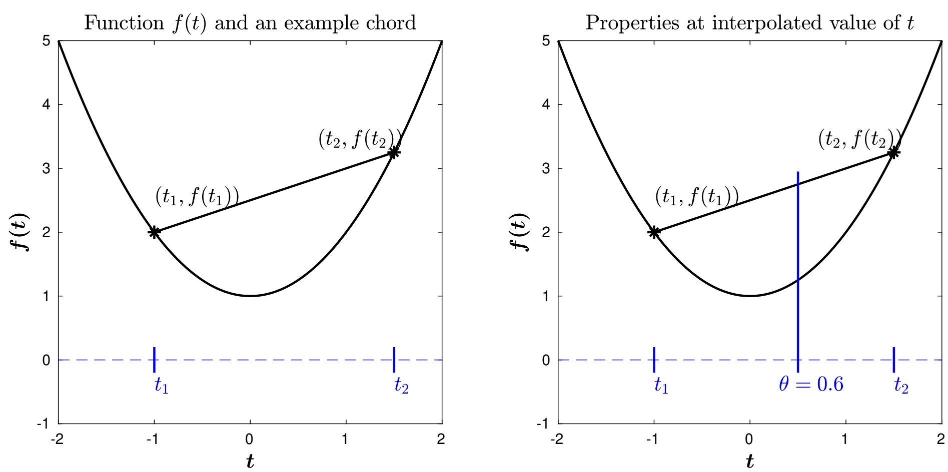

A **B**

Figure 3.5: The graph of a function $f(t)$ that is evaluated at two points. (A) The line segment that "connects" the evaluations is "above" the graph of the function. (B) For $\theta = 0.6$, the value $f((1-\theta)t_1 + \theta t_2)$ is strictly less than $(1-\theta)f(t_1) + \theta f(t_2)$.

A convex function has conditions that differ from those of a strictly convex function. The points t_1 and t_2 can be the same, and the function must be less than or equal to the interpolated value.

Definition: convex function $f(t)$

> For any $t_1 \in \mathbb{R}$, any $t_2 \in \mathbb{R}$, any $\theta \in \mathbb{R}_+$, and any $f : \mathbb{R} \to \mathbb{R}$, that function f is a *convex function* is defined as
>
> $$(0 \leq \theta \leq 1) \to (f((1-\theta)t_1 + \theta t_2)) \leq ((1-\theta)f(t_1) + \theta f(t_2))) \tag{3.6}$$

In this course, the practical difference between a convex function and a strictly convex function occurs at a local minimizer. A convex function can be "flat" around a local minimizer, which means that the value of the function is unchanged. A strictly convex function cannot be "flat" at a minimizer because the value of the function must be strictly less than the interpolated value.

A function can be mathematically "flat", which means that equality to an interpolation can be proved. A function can also be computationally "flat", which means that computations near the local minimizer result in values that are numerically indistinguishable from the value of the function at the minimizer. A family of functions that are mathematically strictly convex, but which computationally can be treated as "flat" near the origin, can be parameterized with a positive integer k that specifies $f_k = t^{2k+2} + 1$. Three functions in this family are illustrated in Figure 3.6, which have increasingly higher second derivatives that are strictly convex.

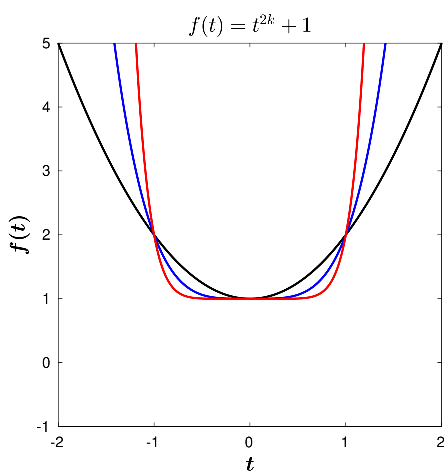

Figure 3.6: Graphs of a family of functions $f_k = t^{2k+2} + 1$. As k increases, the function specified by k is increasingly "flat" at the origin.

With these definitions and understanding, we can now simply state the difficulty we encountered in estimating a local minimizer for the dynamical response function: the function is not locally convex.

3.3 Basic Gradient Inequality

We can easily derive a result that will explain how a function behaves near a local minimizer. We will use this result, in a future class, to understand how a simple and widely used optimization algorithms works.

First we will need to recall, from prerequisite material in differential calculus, the Lagrange form of the remainder term in the Taylor series expansion for a continuously differentiable and analytic function that has a scalar argument. That is, we need to recall how to find a local linear approximation to a "smooth well behaved" function.

Consider a function with a scalar argument, $f : \mathbb{R} \to \mathbb{R}$, that has an infinite number of derivatives that are all continuous – this is a *smooth* function. Suppose that the Taylor series for this function converges everywhere – this is an *analytic* function. This function has a Taylor series that may have an infinite number of terms, which is another way of saying that it is an infinite series. We can *truncate* the Taylor series so that the function f is approximated by a finite sum of terms.

We want to understand how the truncated Taylor series of the function behaves on some interval of interest, such as $[t_L\,,\,t_R]$. There are many ways to represent the truncation, including simply neglecting the higher-order terms or using a "little-o" representation. Here, we will use the Lagrange form of the remainder. This remainder requires that we can place an upper bound on the absolute value of some derivative.

The first-order Lagrange error term for a Taylor series can be written by knowing that there exists, in the interval of interest, some value $\xi \in [t_L\,,\,t_R]$ that can be used to express the truncated terms of the series. This a basic calculus result allows us to express the Taylor series around an argument $t_0 \in [t_L\,,\,t_R]$, for any argument $t \in [t_L\,,\,t_R]$, as

$$f(t) = f(t_0) + f'(t_0)(t - t_0) + \frac{1}{2} f''(\xi)(t - t_0)^2 \tag{3.7}$$

If f is convex on the interval $[t_L\,,\,t_R]$, and $\xi \in [t_L\,,\,t_R]$, then $0 \leq f''(\xi)$. We can write Equation 3.7 as an inequality

$$f(t) \;\; \geq \;\; f(t_0) + f'(t_0)(t - t_0) \tag{3.8}$$

Equation 3.8 is remarkable and useful. It is remarkable because it implies that, if a function f is convex on an interval then, for every point t in the interval, the function has the tangent line at t as a lower bound. Another way of saying this result is that a locally convex function f is "above" every possible tangent line. This is illustrated in Figure 3.7 for two points in an interval.

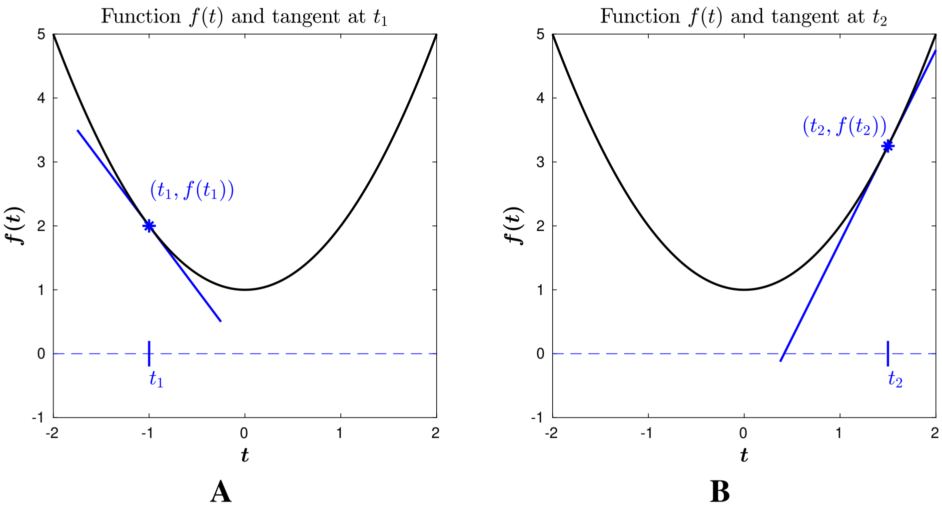

Figure 3.7: Graphs of a locally convex function and the tangent lines at two points. (A) the derivative at t_1 is negative and the function is "above" the tangent line at t_1. (A) the derivative at t_2 is positive and the function is "above" the tangent line at t_2.

Extra Notes

3.4 Extra Notes on Stationarity

The first-order optimality condition for a function with a scalar argument can be proved using the definition of the derivative.

Theorem: first-order necessary condition for optimality

For any $t \in \mathbb{R}$, any $t^* \in \mathbb{R}$, and any continuously differentiable $f : \mathbb{R} \to \mathbb{R}$,

$$(t^* \text{ is a strict local minimizer of } f) \to (t^* \text{ is a stationary point}) \tag{3.9}$$

Proof: Assume that t^* is a local minimizer of f. By definition, there exists $r \in \mathbb{R}_+$ such that, for all $u \in \mathbb{R} : (|u - t^*| < r)$, $f(u) \geq f(t^*)$.

Let $h = u - t^*$. Then

$$\begin{array}{rrcl} & f(u) & \geq & f(t^*) \\ \equiv & f(u) - f(t^*) & \geq & 0 \\ \equiv & f(t^* + h) - f(t^*) & \geq & 0 \end{array}$$

We can bound $f'(t^*)$ by taking limits from the left and from the right, using the above bounds:

$$\begin{aligned}
& \lim_{h\to 0^-} \frac{f(t^*+h)-f(t^*)}{h} && \geq && 0 \\
\wedge \quad & \lim_{h\to 0^+} \frac{f(t^*+h)-f(t^*)}{h} && \geq && 0 \\
\to \quad & \lim_{h\to 0} \frac{f(t^*+h)-f(t^*)}{h} && = && 0 \\
\text{so} \quad & f'(t^*) && = && 0
\end{aligned}$$

End of Extra Notes

References

[1] Beck A: Introduction to Nonlinear Optimization: Theory, Algorithms, and Applications with MATLAB. Siam Press, 2014

[2] Nocedal J, Wright S: Numerical Optimization. Springer Science & Business Media, 2006