

3 Optimization Basics

Goal

Convex sets and functions, gradient and Hessian, Fenchel conjugate, Lagrangian dual and gradient descent.

Alert 3.1: Convention

Gray boxes are not required hence can be omitted for unenthusiastic readers.

[This note is likely to be updated again soon.](#)

Definition 3.2: Vector space

Throughout the course, our universe is typically a (linear) **vector space** V over the real **scalar** field \mathbb{R} . On V we have two operators: addition $+$: $V \times V \rightarrow V$ and (scalar) multiplication \cdot : $\mathbb{R} \times V \rightarrow V$. Together they satisfy the following axioms:

- Addition commutativity: $\forall \mathbf{u}, \mathbf{v} \in V, \mathbf{u} + \mathbf{v} = \mathbf{v} + \mathbf{u}$;
- Addition associativity: $\forall \mathbf{u}, \mathbf{v}, \mathbf{w} \in V, (\mathbf{u} + \mathbf{v}) + \mathbf{w} = \mathbf{u} + (\mathbf{v} + \mathbf{w})$;
- Identity of addition: $\exists \mathbf{0} \in V$ such that $\forall \mathbf{v} \in V, \mathbf{0} + \mathbf{v} = \mathbf{v}$;
- Inverse of addition: $\forall \mathbf{v} \in V, \exists$ unique $\mathbf{u} \in V$ such that $\mathbf{v} + \mathbf{u} = \mathbf{0}$, in which case we **denote** $\mathbf{u} = -\mathbf{v}$;
- Identity of multiplication: $\forall \mathbf{v} \in V, 1 \cdot \mathbf{v} = \mathbf{v}$, where $1 \in \mathbb{R}$ is the usual constant 1;
- Compatibility: $\forall a, b \in \mathbb{R}, \mathbf{v} \in V, a(b\mathbf{v}) = (ab)\mathbf{v}$;
- Distributivity over vector addition: $\forall a \in \mathbb{R}, \mathbf{u}, \mathbf{v} \in V, a(\mathbf{u} + \mathbf{v}) = a\mathbf{u} + a\mathbf{v}$;
- Distributivity over scalar addition: $\forall a, b \in \mathbb{R}, \mathbf{v} \in V, (a + b)\mathbf{v} = a\mathbf{v} + b\mathbf{v}$.

These axioms are so natural that you may question why do we bother to formalize them? Well, take two images/documents/graphs/speeches/DNAs/etc., how do you add them? multiply with scalars? inverse? **Not so obvious...** Representing objects as vectors in a vector space so that we can exploit linear algebra tools is arguably one of the most important lessons in ML/AI.

Example 3.3: Euclidean space

The most common vector space we are going to need is the d -dimensional **Euclidean space** \mathbb{R}^d . Each vector $\mathbf{v} \in \mathbb{R}^d$ can be identified with a d -tuple: $\mathbf{v} = (v_1, v_2, \dots, v_d)$. The addition and multiplication operators are defined element-wise, and we can easily verify the axioms:

- Let $\mathbf{u} = (u_1, u_2, \dots, u_d)$, then $\mathbf{u} + \mathbf{v} = (u_1 + v_1, u_2 + v_2, \dots, u_d + v_d)$;
- Verify associativity yourself;
- $\mathbf{0} = (0, 0, \dots, 0)$;
- $-\mathbf{v} = (-v_1, -v_2, \dots, -v_d)$;
- Obvious;
- $a\mathbf{v} = (av_1, av_2, \dots, av_d)$;
- Verify distributivity yourself;
- Verify distributivity yourself.

In Section 1, we encoded an email as a binary vector $\mathbf{x} \in \{0, 1\}^d \subseteq \mathbb{R}^d$. This crude bag-of-words representation allowed us to use all linear algebra tools.

Definition 3.4: Convex set

A point set $C \subseteq V$ is called **convex** if

$$\forall \mathbf{x}, \mathbf{z} \in C, [\mathbf{x}, \mathbf{z}] := \{\lambda \mathbf{x} + (1 - \lambda) \mathbf{z} : \lambda \in [0, 1]\} \subseteq C.$$

Elements in the “interval” $[\mathbf{x}, \mathbf{z}]$ are called convex combinations of \mathbf{x} and \mathbf{z} .

Theorem 3.5: Intersection and union of convex sets

Arbitrary intersection and **increasing** union of convex sets are convex. \square

Thus, $\liminf_{\alpha} C_{\alpha} := \bigcup_{\alpha} \bigcap_{\beta \geq \alpha} C_{\beta}$ is convex. However, arbitrary union hence $\limsup_{\alpha} C_{\alpha} := \bigcap_{\alpha} \bigcup_{\beta \geq \alpha} C_{\beta}$ may **not** be convex.

Exercise 3.6: Hyperplane and halfspace

Verify the convexity of the **hyperplane** and **halfspace**:

$$\partial H_{\mathbf{w}, b} := \{\mathbf{x} \in \mathbb{R}^d : \langle \mathbf{w}, \mathbf{x} \rangle + b = 0\}$$

$$H_{\mathbf{w}, b} := \{\mathbf{x} \in \mathbb{R}^d : \langle \mathbf{w}, \mathbf{x} \rangle + b \leq 0\}$$

(The partial notation ∂ in front of a set means **boundary**.)

Exercise 3.7: Polyhedron

A **polyhedron** is some **finite** intersection of halfspaces:

$$P := \bigcap_{i=1, \dots, n} H_{\mathbf{w}_i, b_i} = \{\mathbf{x} \in \mathbb{R}^d : W\mathbf{x} + \mathbf{b} \leq \mathbf{0}\}.$$

Prove any polyhedron is convex.

If a polyhedron is bounded, then we call it a polytope. Prove the following:

- the unit ball of ℓ_1 norm $\{\mathbf{x} : \|\mathbf{x}\|_1 \leq 1\}$ is a polytope;
- the unit ball of ℓ_{∞} norm $\{\mathbf{x} : \|\mathbf{x}\|_{\infty} \leq 1\}$ is a polytope;
- the unit ball of ℓ_2 norm $\{\mathbf{x} : \|\mathbf{x}\|_2 \leq 1\}$ is **not** a polytope.

Theorem 3.8: Convex sets as intersection of halfspaces

Any closed convex set is intersection of halfspaces:

$$C = \bigcap_{i \in I} H_{\mathbf{w}_i, b_i}$$

The subtlety is that for non-polyhedra, the index set I is **infinite**. For instance: \square

$$\{\mathbf{x} : \|\mathbf{x}\|_2 \leq 1\} = \bigcap_{\mathbf{w} : \|\mathbf{w}\|_2 = 1} H_{\mathbf{w}, -1}.$$

Definition 3.9: Convex function (Jensen 1905)

The extended real-valued function $f : V \rightarrow (-\infty, \infty]$ is called **convex** if **Jensen's inequality** holds:

$$\forall \mathbf{x}, \mathbf{z} \in V, \forall \lambda \in (0, 1), f(\lambda \mathbf{x} + (1 - \lambda)\mathbf{z}) \leq \lambda f(\mathbf{x}) + (1 - \lambda)f(\mathbf{z}). \quad (3.1)$$

It is necessary that the (effective) domain of f , i.e. $\text{dom } f := \{\mathbf{x} \in V : f(\mathbf{x}) < \infty\}$, is a convex set.

We call f *strictly convex* iff the equality in (3.1) holds only when $\mathbf{x} = \mathbf{z}$.

A function f is (strictly) **concave** iff $-f$ is (strictly) convex.

According to [wikipedia](#), Jensen (Danish) never held any academic position and proved his mathematical results in his spare time.

Jensen, Johan Ludwig William Valdemar (1905). “Om konvekse Funktioner og Uligheder mellem Middelværdier”. *Nyt Tidsskrift for Matematik B*, vol. 16, pp. 49–68.

Exercise 3.10: Affine = convex and concave

Prove that a function is both convex and concave iff it is affine (see Definition 1.12).

Remark 3.11: Convexity is remarkably important!

In the above definition of convexity, we have used the fact that V is a vector space, so that we can add vectors and multiply them with (real) scalars. It is quite remarkable that such a simple definition leads to a **huge** body of interesting results, a tiny part of which we shall be able to present below.

Definition 3.12: Epigraph and sub-level sets

The **epigraph** of a function f is defined as the set of points lying on or above its **graph**:

$$\text{epi } f := \{(\mathbf{x}, t) \in V \times \mathbb{R} : f(\mathbf{x}) \leq t\}.$$

It is clear that the epigraph of a function completely determines it:

$$f(\mathbf{x}) = \min\{t : (\mathbf{x}, t) \in \text{epi } f\}.$$

So two functions are equal iff their epigraphs (sets) are the same. Again, we see that functions and sets are somewhat equivalent.

The **sub-level sets** of f are defined as:

$$\forall t \in \mathbb{R}, L_t := \llbracket f \leq t \rrbracket := \{\mathbf{x} \in V : f(\mathbf{x}) \leq t\}.$$

Sub-level sets also completely determine the function:

$$f(\mathbf{x}) = \min\{t : \mathbf{x} \in L_t\}.$$

Each sub-level set is clearly a section of the epigraph.

Exercise 3.13: Calculus for convexity

Prove the following:

- Any norm is convex;
- If f and g are convex, then for any $\alpha, \beta \geq 0$, $\alpha f + \beta g$ is also convex; (what about $-f$?)
- If $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is convex, then so is $\mathbf{w} \mapsto f(A\mathbf{w} + \mathbf{b})$;

- If f_t is convex for all $t \in T$, then $f := \sup_{t \in T} f_t$ is convex;
- If $f(\mathbf{x}, t)$ is **jointly convex** in \mathbf{x} and t , then $\mathbf{x} \mapsto \inf_{t \in T} f(\mathbf{x}, t)$ is convex;
- f is a convex function iff its epigraph $\text{epi } f$ is a convex set;
- If $f : C \rightarrow \mathbb{R}$ is convex, then the perspective function $g(\mathbf{x}, t) := tf(\mathbf{x}/t)$ is convex on $C \times \mathbb{R}_{++}$;
- If $f(\mathbf{x}, \mathbf{z})$ is convex, then for any \mathbf{x} , $f_{\mathbf{x}} := f(\mathbf{x}, \cdot)$ is convex and similarly for $f^{\mathbf{z}} := f(\cdot, \mathbf{z})$. (what about the converse?)
- All sub-level sets of a convex function are convex, but a function with all sub-level sets being convex need not be convex (these are called quasi-convex functions).

Definition 3.14: Fenchel conjugate function

For any extended real-valued function $f : \mathcal{V} \rightarrow (-\infty, \infty]$ we define its Fenchel conjugate function as:

$$f^*(\mathbf{x}^*) := \sup_{\mathbf{x}} \langle \mathbf{x}, \mathbf{x}^* \rangle - f(\mathbf{x}).$$

According to one of the rules in Exercise 3.13, f^* is always a convex function (of \mathbf{x}^*).

If $\text{dom } f$ is nonempty and **closed**, and f is continuous, then

$$f^{**} := (f^*)^* = f.$$

This remarkable property of convex functions will be used later in the course.

Theorem 3.15: Verifying convexity

Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be twice differentiable. Then f is convex iff its Hessian $\nabla^2 f$ is always **positive semidefinite**. If the Hessian is always positive definite, then f is strictly convex. \square

The function $f(x) = x^4$ is strictly convex at its minimum $x = 0$.

Fix any \mathbf{x} and take any direction \mathbf{d} . Consider the univariate function $h(t) := f(\mathbf{x} + t\mathbf{d})$. We verify that $h''(t) = \langle \mathbf{d}, \nabla^2 f(\mathbf{x} + t\mathbf{d})\mathbf{d} \rangle \geq 0$. In other words, **a convex function has increasing derivative** along any direction \mathbf{d} (starting from any point \mathbf{x}).

Exercise 3.16: Example convex functions

Prove the following functions are convex (Exercise 3.13 and Theorem 3.15 may be handy):

- affine functions: $f(\mathbf{x}) = \mathbf{w}^\top \mathbf{x} + b$;
- exponential function: $f(x) = \exp(x)$;
- entropy: $f(x) = x \log x$ with $x \geq 0$ and $0 \log 0 := 0$;
- log-sum-exp: $f(\mathbf{x}) = \log \sum_{j=1}^d \exp(x_j)$; (its gradient is the so-called **softmax**, to be discussed later)

(You may appreciate the epigraph more after the last exercise.)

Definition 3.17: Optimization problem

Consider a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, we are interested in the minimization problem:

$$\mathbf{p}_* := \min_{\mathbf{x} \in C} f(\mathbf{x}), \quad (3.2)$$

where $C \subseteq \mathbb{R}^d$ represents the constraints that \mathbf{x} *must* satisfy.

Historically, the minimization problem (3.2) was motivated by the need of solving nonlinear equations $h(\mathbf{x}) = 0$ (Cauchy 1847; Curry 1944), which can be reformulated as a least squares minimization problem $\min_{\mathbf{x}} h^2(\mathbf{x})$. As pointed out by Polyak (1963), the minimization problem itself has become ubiquitous, and often does not have to correspond to solving any nonlinear equation.

We remind that the **minimum value** \mathbf{p}_* is an extended real number in $[-\infty, \infty]$ (where $\mathbf{p}_* = \infty$ iff $C = \emptyset$). When \mathbf{p}_* is finite, any feasible $\mathbf{x} \in C$ such that $f(\mathbf{x}) = \mathbf{p}_*$ is called a **minimizer**. **Minimum value always exists while minimizers may not!**

Cauchy, M. Augustin-Louis (1847). “Méthode générale pour la résolution des systèmes d’équations simultanées”. *Comptes rendus hebdomadaires des séances de l’Académie des sciences*, vol. 25, no. 2, pp. 536–538.

Curry, Haskell B. (1944). “The Method of Steepest Descent for Non-linear Minimization Problems”. *Quarterly of Applied Mathematics*, vol. 2, no. 3, pp. 258–261.

Polyak, Boris Teodorovich (1963). “Gradient methods for the minimization of functionals”. *USSR Computational Mathematics and Mathematical Physics*, vol. 3, no. 4, pp. 643–653.

Definition 3.18: Level-bounded

A function $f : \mathbb{R}^d \rightarrow (-\infty, \infty]$ is said to be level bounded iff for all $t \in \mathbb{R}$, the sublevel set $\{f \leq t\}$ is bounded. Equivalently, f is level bounded iff $\|\mathbf{x}\| \rightarrow \infty \implies f(\mathbf{x}) \rightarrow \infty$.

Theorem 3.19: Existence of minimizer

A continuous and level-bounded function $f : \mathbb{R}^d \rightarrow (-\infty, \infty]$ has a minimizer. □

To appreciate this innocent theorem, let $f(x) = \exp(x)$:

- Is f continuous?
- Is f level-bounded?
- What is the minimum value of f ? Does f has a minimizer on \mathbb{R} ?

Definition 3.20: Extrema of an unconstrained function

Recall that \mathbf{x} is a **local minimizer** of f if there exists an (open) neighborhood \mathcal{N} of \mathbf{x} so that for all $\mathbf{z} \in \mathcal{N}$ we have $f(\mathbf{x}) \leq f(\mathbf{z})$. In case when \mathcal{N} can be chosen to be the entire space \mathbb{R}^d , we say \mathbf{x} is a **global minimizer** of f with the notation $\mathbf{x} \in \text{argmin } f$.

By definition, a global minimizer is always a local minimizer while the converse is true only for a special class of functions (knowns as invex functions). The global minimizer may not be unique (take a constant function) but the **global minimum value** is.

The definition of a local (global) maximizer is analogous.

Exercise 3.21: Properties of extrema

Prove the following:

- \mathbf{x} is a local (global) minimizer of f iff \mathbf{x} is a local (global) maximizer of $-f$.
- \mathbf{x} is a local (global) minimizer of f iff \mathbf{x} is a local (global) minimizer of $\lambda f + c$ for any $\lambda > 0$ and $c \in \mathbb{R}$.
- If \mathbf{x} is a local (global) minimizer (maximizer) of f , then it is a local (global) minimizer (maximizer) of $g(f)$ for any increasing function $g : \mathbb{R} \rightarrow \mathbb{R}$. What if g is *strictly* increasing?
- \mathbf{x} is a local (global) minimizer (maximizer) of a positive function f iff it is a local (global) maximizer (minimizer) of $1/f$.

- \mathbf{x} is a local (global) minimizer (maximizer) of a positive function f iff it is a local (global) minimizer (maximizer) of $\log f$.

Alert 3.22: The epigraph trick

Often, we rewrite the optimization problem

$$\min_{\mathbf{x} \in C} f(\mathbf{x})$$

as the equivalent one:

$$\min_{(\mathbf{x}, t) \in \text{epi } f \cap C \times \mathbb{R}} t, \quad (3.3)$$

where the newly introduced variable t is **jointly optimized** with \mathbf{x} . The advantages of (3.3) include:

- the objective in (3.3) is a simple, canonical linear function $\langle (\mathbf{0}, 1), (\mathbf{x}, t) \rangle$;
- the constraints in (3.3) may reveal more (optimization) insights, as we will see in SVMs.

Theorem 3.23: Local is global under convexity

Any local minimizer of a convex function (over some convex constraint set) is global.

Proof. Let \mathbf{x} be a local minimizer of a convex function f . Suppose there exists \mathbf{z} such that $f(\mathbf{z}) < f(\mathbf{x})$. Take convex combination and appeal to the definition of convexity in Definition 3.9:

$$\forall \lambda \in (0, 1), f(\lambda \mathbf{x} + (1 - \lambda)\mathbf{z}) \leq \lambda f(\mathbf{x}) + (1 - \lambda)f(\mathbf{z}) < f(\mathbf{x}),$$

contradicting to the local minimality of \mathbf{x} . □

In fact, in Theorem 3.24 we will see that any stationary point of a convex function is its global minimizer.

Theorem 3.24: Fermat's necessary condition for extrema

A necessary condition for \mathbf{x} to be a local minimizer of a differentiable function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is

$$\nabla f(\mathbf{x}) = \mathbf{0}.$$

(Such points are called *stationary*, a.k.a. *critical*.) If f is convex, then the necessary condition is also sufficient.

Proof. Suppose t is a local minimizer of a univariate function $g : \mathbb{R} \rightarrow \mathbb{R}$, then apply the definition of derivative we may easily deduce that $g'(t) \leq 0$ and $g'(t) \geq 0$, i.e. $g'(t) = 0$.

Now if \mathbf{x} is a local minimizer of $f : \mathbb{R}^d \rightarrow \mathbb{R}$, then $t = 0$ is a local minimizer of the univariate function $g(t) := f(\mathbf{x} + t\mathbf{w})$ for any \mathbf{w} . Apply the previous result for univariate functions and the chain rule:

$$g'(0) = \mathbf{w}^\top \nabla f(\mathbf{x}) = 0.$$

Since \mathbf{w} was arbitrary, we must have $\nabla f(\mathbf{x}) = \mathbf{0}$.

If f is convex, then we have

$$\langle \nabla f(\mathbf{x}), \mathbf{z} - \mathbf{x} \rangle = \lim_{t \rightarrow 0^+} \frac{f(\mathbf{x} + t(\mathbf{z} - \mathbf{x})) - f(\mathbf{x})}{t} \leq \frac{tf(\mathbf{z}) + (1 - t)f(\mathbf{x}) - f(\mathbf{x})}{t} = f(\mathbf{z}) - f(\mathbf{x}).$$

In other words, we obtain the first order characterization of convexity:

$$\forall \mathbf{x}, \mathbf{z}, \quad f(\mathbf{z}) \geq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{z} - \mathbf{x} \rangle.$$

Plugging in $\nabla f(\mathbf{x}) = \mathbf{0}$ we see that \mathbf{x} is in fact a global minimizer. \square

Take $f(x) = x^3$ and $x = 0$ we see that this necessary condition is not sufficient for nonconvex functions. For local maximizers, we simply negate the function and apply the theorem to $-f$ instead.

Example 3.25: The difficulty of satisfying a constraint

Consider the trivial problem:

$$\min_{x \geq 1} x^2,$$

which admits a unique minimizer $x_* = 1$. However, if we ignore the constraint and set the derivative to zero we would obtain $x = 0$, which does not satisfy the constraint!

We can apply Theorem 3.24 only when there is no constraint. We will introduce the Lagrangian to “remove” constraints below.

A common trick is to ignore any constraint and apply Theorem 3.24 anyway to derive a “bogus” minimizer \mathbf{x}_* . Then, we verify if \mathbf{x}_* satisfies the constraint: If it does, then we actually find a minimizer for the *constrained* problem! For instance, had the constraint above been $x \geq -1$ we would be fine by ignoring the constraint. Needless to say, this trick only works occasionally.

Remark 3.26: Iterative algorithm

The prevailing algorithms in machine learning are iterative in nature, i.e., we will construct a sequence $\mathbf{x}_0, \mathbf{x}_1, \dots$, which will hopefully “converge” to something we contend with.

Definition 3.27: Projection to a closed set

Let $C \subseteq \mathbb{R}^d$ be a **closed** set. We define the (Euclidean) projection of a point $\mathbf{x} \in \mathbb{R}^d$ to C as:

$$P_C(\mathbf{x}) := \operatorname{argmin}_{\mathbf{z} \in C} \|\mathbf{z} - \mathbf{x}\|_2,$$

i.e., points in C that are closest to the given point \mathbf{x} . Needless to say, $P_C(\mathbf{x}) = \mathbf{x}$ iff \mathbf{x} lies in C .

The projection is always unique iff C is convex (Bunt 1934; Motzkin 1935). The only if part remains a long open problem when the space is infinite dimensional.

Bunt, L. N. H. (1934). “Bijdrage tot de theorie de convexe puntverzamelingen”. PhD thesis. University of Groningen.
 Motzkin, Theodore Samuel (1935). “Sur quelques propriétés caractéristiques des ensembles convexes”. *Atti della Reale Accademia Nazionale dei Lincei*, vol. 21, no. 6, pp. 562–567.

Exercise 3.28: Projecting to nonnegative orthant

Let $C = \mathbb{R}_+^d$ be the nonnegative orthant. Find the formula for $P_C(\mathbf{x})$. (Exercise 3.21 may be handy.)

Algorithm 3.29: Algorithm of feasible direction

We remind that line 4 below is possible because our universe is a vector space! This testifies again the

(obvious?) importance of making machine learning amenable to linear algebra.

Algorithm: Algorithm of feasible direction

Input: $\mathbf{x}_0 \in \text{dom } f \cap C$

```

1 for  $t = 0, 1, \dots$  do
2   choose direction  $\mathbf{d}_t$  // e.g.  $\limsup_{t \rightarrow \infty} \langle \mathbf{d}_t, \nabla f(\mathbf{x}_t) \rangle > 0$ 
3   choose step size  $\eta_t > 0$ 
4    $\mathbf{x}_{t+1} = \mathbf{x}_t - \eta_t \mathbf{d}_t$  // update
5    $\mathbf{x}_{t+1} = P_C(\mathbf{x}_{t+1})$  // optional projection step
```

Intuitively, Algorithm 3.29 first finds a direction \mathbf{d}_t , and then moves the iterate \mathbf{x}_t along the direction. How far we move away from the current iterate \mathbf{x}_t is determined by the step size (assuming \mathbf{d}_t is normalized). To motivate the selection of the direction, apply **Taylor's expansion**:

$$f(\mathbf{x}_{t+1}) = f(\mathbf{x}_t - \eta_t \mathbf{d}_t) = f(\mathbf{x}_t) - \eta_t \langle \mathbf{d}_t, \nabla f(\mathbf{x}_t) \rangle + o(\eta_t),$$

where $o(\eta_t)$ is the small order term. Clearly, if $\langle \mathbf{d}_t, \nabla f(\mathbf{x}_t) \rangle > 0$ and η_t is small, then $f(\mathbf{x}_{t+1}) < f(\mathbf{x}_t)$, i.e. the algorithm is descending hence making progress. Typical choices for the direction include:

- **Gradient Descent (GD):** $\mathbf{d}_t = \nabla f(\mathbf{x}_t)$;
- **Newton:** $\mathbf{d}_t = [\nabla^2 f(\mathbf{x}_t)]^{-1} \nabla f(\mathbf{x}_t)$;
- **Stochastic Gradient Descent (SGD):** $\mathbf{d}_t = \xi_t$, $E(\xi_t) = \nabla f(\mathbf{x}_t)$.

Remark 3.30: Randomness helps?

Note that if we start from a stationary point \mathbf{x}_0 , i.e. $\nabla f(\mathbf{x}_0) = \mathbf{0}$, then gradient descent will not move, no matter what step size we choose! This is why such points are called stationary in the first place. On the other hand, stochastic gradient descent may still move because of the random noise added to it.

Remark 3.31: More on SGD

In machine learning we typically minimize some average loss over a training set $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}$:

$$\min_{\mathbf{w}} \frac{1}{n} \sum_{i=1}^n \ell(\mathbf{w}; \mathbf{x}_i, y_i) \quad =: \quad \min_{\mathbf{w}} \hat{E} \ell(\mathbf{w}; \mathbf{x}, y), \quad (3.4)$$

where (\mathbf{x}, y) is randomly chosen from the training set \mathcal{D} and the empirical expectation is taken w.r.t. \mathcal{D} . Computing the gradient obviously costs $O(n)$ since we need to go through each sample in the training set \mathcal{D} . On the other hand, if we take a random sample (\mathbf{x}, y) from the training set, and compute

$$\xi = \nabla \ell(\mathbf{w}; \mathbf{x}, y).$$

Obviously, $\hat{E} \xi$ equals the gradient but computing ξ is clearly much cheaper. In practice, one usually sample a mini-batch, and compute the (average) gradient over the mini-batch.

Exercise 3.32: Perceptron as SGD

Recall the perceptron update: if $y(\langle \mathbf{w}, \mathbf{x} \rangle + b) \leq 0$, then

$$\mathbf{w} \leftarrow \mathbf{w} + y\mathbf{x}, \quad b \leftarrow b + y.$$

Construct a loss function $\ell(y(\langle \mathbf{w}, \mathbf{x} \rangle + b))$ in (3.4) so that perceptron reduces to SGD (with step size say $\eta \equiv 1$).

Example 3.33: Descending alone does NOT guarantee convergence to stationary point

Consider the function

$$f = \begin{cases} \frac{3}{4}(1-x)^2 - 2(1-x), & x > 1 \\ \frac{3}{4}(1+x)^2 - 2(1+x), & x < -1 \\ x^2 - 1, & -1 \leq x \leq 1 \end{cases}, \text{ with gradient } f' = \begin{cases} \frac{3}{2}x + \frac{1}{2}, & x > 1 \\ \frac{3}{2}x - \frac{1}{2}, & x < -1 \\ 2x, & -1 \leq x \leq 1 \end{cases}.$$

Clearly, f is convex and has a unique minimizer $x^* = 0$ (with $f^* = -1$). It is easy to verify that

$$f(x) < f(y) \iff |x| < |y|.$$

Start with $x_0 > 1$, set the step size $\eta \equiv 1$, and choose $d = f'(x)$, then $x_1 = x_0 - f'(x_0) = -\frac{x_0+1}{2}$. By induction it is easy to show that $x_{t+1} = -\frac{1}{2}(x_t - (-1)^t)$. Thus, $x_t > 1$ if t is odd and $x_t < -1$ if t is even. Moreover, $|x_{t+1}| < |x_t|$, implying $f(x_{t+1}) < f(x_t)$. It is easy to show that $|x_t| \rightarrow 1$ and $f(x_t) \rightarrow 0$, hence the algorithm is not converging to a stationary point.

Remark 3.34: Step sizes

Let us mention a few ways to choose the step size η_t :

- Cauchy's rule (Cauchy 1847), where the existence of the minimizer is assumed:

$$\eta_t \in \operatorname{argmin}_{\eta \geq 0} f(\mathbf{x}_t - \eta \mathbf{d}_t).$$

- Curry's rule (Curry 1944), where the finiteness of η_t is assumed:

$$\eta_t = \inf\{\eta \geq 0 : f'(\mathbf{x}_t - \eta \mathbf{d}_t) = 0\}.$$

- Constant rule: $\eta_t \equiv \eta > 0$.
- Summable rule: $\sum_t \eta_t = \infty, \sum_t \eta_t^2 < \infty$, e.g. $\eta_t = O(1/t)$;
- Diminishing rule: $\sum_t \eta_t = \infty, \lim_t \eta_t = 0$, e.g. $\eta_t = O(1/\sqrt{t})$.

The latter three are most common, especially for SGD.

Note that under the condition $\langle \mathbf{d}_t, \nabla f(\mathbf{x}_t) \rangle > 0$, the function $h(\eta) := f(\mathbf{x} - \eta \mathbf{d}_t)$ is decreasing for small positive η . Therefore, Curry's rule essentially selects the **smallest** local minimizer of h while Cauchy's rule selects the **global** minimizer of h . Needless to say, Cauchy's rule leads to larger per-step decrease of the function value but is also computationally more demanding. Under both Cauchy's and Curry's rule, we have the orthogonality property:

$$\langle \mathbf{d}_t, \nabla f(\mathbf{x}_{t+1}) \rangle = 0,$$

i.e., the gradient at the next iterate \mathbf{x}_{t+1} is orthogonal to the current direction vector \mathbf{d}_t . This explains the zigzag behaviour in gradient algorithms.

Cauchy, M. Augustin-Louis (1847). "Méthode générale pour la résolution des systèmes d'équations simultanées".

Comptes rendus hebdomadaires des séances de l'Académie des sciences, vol. 25, no. 2, pp. 536–538.

Curry, Haskell B. (1944). "The Method of Steepest Descent for Non-linear Minimization Problems". *Quarterly of Applied Mathematics*, vol. 2, no. 3, pp. 258–261.

Definition 3.35: L -smoothness

For twice differentiable functions f , define its smoothness parameter (a.k.a. the Lipschitz constant of the

gradient):

$$L = \max_{\mathbf{x}} \|\nabla^2 f(\mathbf{x})\|_{\text{sp}}$$

where the norm $\|\cdot\|_{\text{sp}}$ is the spectral norm (i.e., the largest singular value, see Definition 2.13). For such functions, choosing $\eta \in (0, \frac{1}{L})$ will guarantee convergence of gradient descent. For convex functions, the step size can be enlarged to $\eta \in (0, \frac{2}{L})$.

With $\eta = 2/L$, gradient descent may not converge (although the function value still converges to some value): simply revisit Example 3.33.

Example 3.36: Iterates of gradient descent may NOT converge

While the function values $f(\mathbf{x}_t)$ of an iterative algorithm (such as gradient descent) usually converge, the iterate \mathbf{x}_t itself may not:

- when there is no minimizer at all: $f(x) = \exp(x)$;
- when there is a unique minimizer but the function is non-differentiable: consider the convex function

$$f(x) = \begin{cases} e^{-x}, & x \leq 0; \\ x + 1, & x > 0; \end{cases}$$

- even when the function is smooth but there are many minimizers: see the intriguing “smoothed Mexican hat” function in (Absil et al. 2005, Fig 2.1).

Absil, P., R. Mahony, and B. Andrews (2005). “Convergence of the Iterates of Descent Methods for Analytic Cost Functions”. *SIAM Journal on Optimization*, vol. 16, no. 2, pp. 531–547.

Example 3.37: Gradient descent may converge to saddle point

Consider the function $f(x, y) = \frac{1}{2}x^2 + \frac{1}{4}y^4 - \frac{1}{2}y^2$, whose gradient is $\nabla f(x, y) = \begin{pmatrix} x \\ y^3 - y \end{pmatrix}$ and Hessian is $\begin{bmatrix} 1 & 0 \\ 0 & 3y^2 - 1 \end{bmatrix}$. Clearly, there are three stationary points $(0, 0), (0, 1), (0, -1)$, where the last two are global minimizers whereas the first is a saddle point. Take any $\mathbf{x}_0 = (t, 0)$, then $\mathbf{x}_t = (x_t, 0)$ hence it can only converge to $(0, 0)$, which is a saddle point.

Definition 3.38: Lagrangian Dual

Consider the canonical optimization problem:

$$\min_{\mathbf{x} \in C \subseteq \mathbb{R}^d} f(\mathbf{x}) \quad (3.5)$$

$$\text{s.t. } \mathbf{g}(\mathbf{x}) \leq \mathbf{0}, \quad (3.6)$$

$$\mathbf{h}(\mathbf{x}) = \mathbf{0}, \quad (3.7)$$

where $f : \mathbb{R}^d \rightarrow \mathbb{R}$, $\mathbf{g} : \mathbb{R}^d \rightarrow \mathbb{R}^n$, and $\mathbf{h} : \mathbb{R}^d \rightarrow \mathbb{R}^m$. The set C is retained here to represent “simple” constraints that is more convenient to deal with directly than put into either (3.6) or (3.7). (Alternatively, one may always put $C = \mathbb{R}^d$.)

The (nonlinear) constraints (3.6) are difficult to deal with. Fortunately, we can introduce the **Lagrangian multipliers** (a.k.a. **dual variables**) $\boldsymbol{\mu} \in \mathbb{R}_+^n$, $\boldsymbol{\nu} \in \mathbb{R}^m$ to move constraints into the Lagrangian:

$$L(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\nu}) := f(\mathbf{x}) + \boldsymbol{\mu}^\top \mathbf{g}(\mathbf{x}) + \boldsymbol{\nu}^\top \mathbf{h}(\mathbf{x}).$$

We can now rewrite the original problem (3.5) as the following fancy **min-max** problem:

$$p_* := \min_{\mathbf{x} \in C} \max_{\boldsymbol{\mu} \geq \mathbf{0}, \boldsymbol{\nu}} L(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\nu}). \quad (3.8)$$

(Here p stands for **primal**, as opposed to the dual below.) Indeed, if we choose some $\mathbf{x} \in C$ that does not satisfy either of the two constraints in (3.6)-(3.7), then there exist $\boldsymbol{\mu} \in \mathbb{R}_+^n$ and $\boldsymbol{\nu} \in \mathbb{R}^m$ such that $L(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\nu}) \rightarrow \infty$. Thus, in order to minimize w.r.t. \mathbf{x} , we are forced to choose $\mathbf{x} \in C$ to satisfy both constraints in (3.6)-(3.7), in which case $\max_{\boldsymbol{\mu} \geq \mathbf{0}, \boldsymbol{\nu}} L(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\nu})$ simply reduces to $f(\mathbf{x})$ (by setting for instance $\boldsymbol{\mu} = \mathbf{0}, \boldsymbol{\nu} = \mathbf{0}$). **In other words, the explicit constraints in (3.5) have become implicit in (3.8)!**

The **Lagrangian dual** simply swaps the order of min and max:

$$d^* := \max_{\boldsymbol{\mu} \geq \mathbf{0}, \boldsymbol{\nu}} \min_{\mathbf{x} \in C} L(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\nu}). \quad (3.9)$$

(Here d stands for **dual**.) We emphasize **the dimension of the dual variable $\boldsymbol{\mu}$ is the number of inequality constraints while that of $\boldsymbol{\nu}$ is the number of equality constraints; they are different from the dimension of the (primal) variable \mathbf{x}** . Note also that **there is no constraint on \mathbf{x} in the dual problem (3.9) (except the simple one $\mathbf{x} \in C$ which we do not count ☺), implicit or explicit!**

In some sense, the Lagrangian dual is “the trick” that allows us to remove complicated constraints and apply Theorem 3.24.

Theorem 3.39: Weak duality

For any function $f : X \times Y \rightarrow \mathbb{R}$, we have

$$\min_{\mathbf{x}} \max_{\mathbf{y}} f(\mathbf{x}, \mathbf{y}) \geq \max_{\mathbf{y}} \min_{\mathbf{x}} f(\mathbf{x}, \mathbf{y})$$

Proof. Left as exercise. □

We immediately deduce that the Lagrangian dual (3.9) is a **lower bound** of the original problem (3.5), i.e. $p_* \geq d^*$. When equality is achieved, we say **strong duality** holds. For instance, if f, g_i for all $i = 1, \dots, n$, and C are convex, **h is affine**, and some mild regularity condition holds (e.g. **Slater’s condition**), then we have strong duality for the Lagrangian in Definition 3.38.

Exercise 3.40: Does the direction matter?

Derive the Lagrangian dual of the following problem:

$$\begin{aligned} \max_{\mathbf{x} \in C} f(\mathbf{x}) \\ \text{s.t. } \mathbf{g}(\mathbf{x}) \geq \mathbf{0} \\ \mathbf{h}(\mathbf{x}) = \mathbf{0}. \end{aligned}$$

(Start with reducing to (3.5) and then see if you can derive directly.)

Definition 3.41: KKT conditions (Karush 1939; Kuhn and Tucker 1951)

The inner minimization in the Lagrangian dual (3.9) can usually be solved in closed-form:

$$\tilde{\mathbf{x}}(\boldsymbol{\mu}, \boldsymbol{\nu}) := \operatorname{argmin}_{\mathbf{x} \in C} L(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\nu}). \quad (3.10)$$

When $C = \mathbb{R}^d$, applying Theorem 3.24 we obtain the stationary condition (3.12) below. Plugging any

minimizer $\mathfrak{X}(\boldsymbol{\mu}, \boldsymbol{\nu})$ back into (3.9) we obtain the dual problem:

$$\max_{\boldsymbol{\mu} \in \mathbb{R}_+^n, \boldsymbol{\nu} \in \mathbb{R}^m} L(\mathfrak{X}(\boldsymbol{\mu}, \boldsymbol{\nu}); \boldsymbol{\mu}, \boldsymbol{\nu}) \quad \equiv \quad - \min_{\boldsymbol{\mu} \in \mathbb{R}_+^n, \boldsymbol{\nu} \in \mathbb{R}^m} -L(\mathfrak{X}(\boldsymbol{\mu}, \boldsymbol{\nu}); \boldsymbol{\mu}, \boldsymbol{\nu}). \quad (3.11)$$

Compared to the original problem (3.5) that comes with complicated constraints (3.6)-(3.7), the dual problem (3.11) has only very simple nonnegative constraint on $\boldsymbol{\mu}$. Thus, we may try to solve the Lagrangian dual (3.11) instead! In fact, we have the following necessary conditions for \mathbf{x} to minimize the original problem (3.5) and for $\boldsymbol{\mu}, \boldsymbol{\nu}$ to maximize the dual problem (3.11):

- primal feasibility:

$$\mathbf{g}(\mathbf{x}) \leq \mathbf{0}, \quad h(\mathbf{x}) = \mathbf{0}, \quad \mathbf{x} \in C;$$

- dual feasibility:

$$\boldsymbol{\mu} \geq \mathbf{0};$$

- stationarity: $\mathbf{x} = \mathfrak{X}(\boldsymbol{\mu}, \boldsymbol{\nu})$; for $C = \mathbb{R}^d$, we simply have

$$\nabla f(\mathbf{x}) + \sum_{i=1}^n \mu_i \nabla g_i(\mathbf{x}) + \sum_{j=1}^m \nu_j \nabla h_j(\mathbf{x}) = \mathbf{0}; \quad (3.12)$$

- complementary slackness:

$$\langle \boldsymbol{\mu}, \mathbf{g}(\mathbf{x}) \rangle = 0.$$

Note that from primal and dual feasibility we always have

$$\forall i = 1, \dots, n, \quad \mu_i g_i(\mathbf{x}) \leq 0.$$

Thus, complementary slackness actually implies equality in all n inequalities above.

When strong duality mentioned in Theorem 3.39 holds, the above KKT conditions are also sufficient!

Karush, W. (1939). “Minima of Functions of Several Variables with Inequalities as Side Constraints”. MA thesis. University of Chicago.

Kuhn, H. W. and A. W. Tucker (1951). “Nonlinear programming”. In: *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pp. 481–492.

Algorithm 3.42: Dual gradient descent (Uzawa58)

We can apply Algorithm 3.29 to the dual problem (3.11), equipped with the projection onto nonnegative orthants in Exercise 3.28. After solving the dual, we may **attempt** to recover the primal through (3.10). In both steps we sidestep any complicated constraints! However, one needs to **always keep the following Alert 3.43 in mind**.

Alert 3.43: Instability

A popular way to solve the primal problem (3.5) is to solve the dual (3.11) first. Then, with the optimal dual variable $(\boldsymbol{\mu}^*, \boldsymbol{\nu}^*)$ at hand, we “recover” the primal solution by

$$\mathfrak{X}(\boldsymbol{\mu}^*, \boldsymbol{\nu}^*) = \operatorname{argmin}_{\mathbf{x} \in C} L(\mathbf{x}; \boldsymbol{\mu}^*, \boldsymbol{\nu}^*).$$

The minimizing set \mathfrak{X} always contains all minimizers of the primal problem (3.5). Thus, **if \mathfrak{X} is a singleton, everything is fine.** Otherwise **\mathfrak{X} may actually contain points that are not minimizers of the primal problem!** Fortunately, we need only verify primal feasibility in order to identify the true minimizers of the primal (3.5) (when strong duality holds). The problem is, in practice, our algorithm only returns one “solution” from \mathfrak{X} ,

and if it fails primal feasibility, we may not be able to fetch a different “solution” from \mathfrak{X} .

Example 3.44: Instability

Let us consider the trivial problem:

$$\min_x 0, \text{ s.t. } x \geq 0,$$

whose minimizers are \mathbb{R}_+ . We derive the Lagrangian dual:

$$\max_{\mu \geq 0} \min_x -\mu x = \max_{\mu \geq 0} \begin{cases} 0, & \text{if } \mu = 0 \\ -\infty, & \text{o.w.} \end{cases}.$$

Clearly, we have $\mu^* = 0$. Fixing $\mu^* = 0$ and solving

$$\min_x -\mu^* x$$

gives us $\mathfrak{X} = \mathbb{R}$ which strictly contains the primal solutions \mathbb{R}_+ ! Verifying primal feasibility $x \geq 0$ then identifies true minimizers.

Example 3.45: (Kernel) ridge regression

Recall ridge regression:

$$\begin{aligned} \min_{\mathbf{w}, \mathbf{z}} \quad & \frac{1}{2} \|\mathbf{z}\|_2^2 + \frac{\lambda}{2} \|\mathbf{w}\|_2^2 \\ \text{s.t.} \quad & X\mathbf{w} - \mathbf{y} = \mathbf{z}, \end{aligned}$$

where we *introduced* an “artificial” constraint (and variable). Derive the Lagrangian dual:

$$\max_{\boldsymbol{\alpha}} \min_{\mathbf{w}, \mathbf{z}} \frac{1}{2} \|\mathbf{z}\|_2^2 + \frac{\lambda}{2} \|\mathbf{w}\|_2^2 + \boldsymbol{\alpha}^\top (X\mathbf{w} - \mathbf{y} - \mathbf{z}),$$

Applying Theorem 3.24 to the inner minimization problem:

$$\mathbf{w}_* = -X^\top \boldsymbol{\alpha} / \lambda, \quad \mathbf{z}_* = \boldsymbol{\alpha} \tag{3.13}$$

Plugging it back in (and simplify) we obtain the dual:

$$\max_{\boldsymbol{\alpha}} -\frac{1}{2\lambda} \|X^\top \boldsymbol{\alpha}\|_2^2 - \boldsymbol{\alpha}^\top \mathbf{y} - \frac{1}{2} \|\boldsymbol{\alpha}\|_2^2$$

Applying Theorem 3.24 again we obtain:

$$\boldsymbol{\alpha}^* = -(XX^\top / \lambda + I)^{-1} \mathbf{y}$$

and hence from (3.13) we have

$$\mathbf{w}_* = X^\top (XX^\top + \lambda I)^{-1} \mathbf{y}.$$

From Section 2 we know the solution of ridge regression is

$$\mathbf{w}_* = (X^\top X + \lambda I)^{-1} X^\top \mathbf{y}$$

Verify the two solutions of \mathbf{w}_* are the same. (In fact, we have accidentally proved the [Sherman-Morrison formula](#).) The purpose of this “linear algebra exercise” will become evident when we discuss reproducing kernels.

Algorithm 3.46: Gradient descent ascent (GDA)

When we cannot solve the inner minimization problem in the Lagrangian dual (3.9), an alternative is to iteratively perform one gradient descent step on the inner minimization and then perform another gradient ascent step on the outer maximization. This idea can be traced back to (at least) (**BrownvonNeumann50; ArrowHurwicz58**).

More generally, let us consider the min-max optimization problem:

$$\min_{\mathbf{x} \in X} \max_{\mathbf{y} \in Y} f(\mathbf{x}, \mathbf{y}).$$

Algorithm: Gradient descent ascent for min-max

Input: $(\mathbf{x}_0, \mathbf{y}_0) \in \text{dom } f \cap X \times Y$

```

1  $s_{-1} = 0, (\bar{\mathbf{x}}_{-1}, \bar{\mathbf{y}}_{-1}) = (\mathbf{0}, \mathbf{0})$  // optional
2 for  $t = 0, 1, \dots$  do
3   choose step size  $\eta_t > 0$ 
4    $\mathbf{x}_{t+1} = P_X[\mathbf{x}_t - \eta_t \nabla_{\mathbf{x}} f(\mathbf{x}_t, \mathbf{y}_t)]$  // GD on minimization
5    $\mathbf{y}_{t+1} = P_Y[\mathbf{y}_t + \eta_t \nabla_{\mathbf{y}} f(\mathbf{x}_t, \mathbf{y}_t)]$  // GA on maximization
6    $s_t = s_{t-1} + \eta_t$ 
7    $(\bar{\mathbf{x}}_t, \bar{\mathbf{y}}_t) = \frac{s_{t-1}(\bar{\mathbf{x}}_{t-1}, \bar{\mathbf{y}}_{t-1}) + \eta_t(\mathbf{x}_t, \mathbf{y}_t)}{s_t}$  // averaging:  $(\bar{\mathbf{x}}_t, \bar{\mathbf{y}}_t) = \frac{\sum_{k=1}^t \eta_k(\mathbf{x}_k, \mathbf{y}_k)}{\sum_k \eta_k}$ 

```

Variations of Line 7 include (but are not limited to):

- use different step sizes on \mathbf{x} and \mathbf{y} ;
- use \mathbf{x}_{t+1} in the update on \mathbf{y} (or vice versa);
- use stochastic gradients in both steps;
- after every update in \mathbf{x} , perform k updates in \mathbf{y} (or vice versa);

Example 3.47: Vanilla GDA may never converge for any step size

Let us consider the following simple problem:

$$\min_{x \in [-1, 1]} \max_{y \in [-1, 1]} xy \quad \equiv \quad \max_{y \in [-1, 1]} \min_{x \in [-1, 1]} xy.$$

The left-hand side is equivalent as $\min_{x \in [-1, 1]} |x|$ hence with unique minimizer $x^* = 0$. Similarly, the right-hand side is equivalent as $\max_{y \in [-1, 1]} -|y|$ hence with unique maximizer $y^* = 0$. It follows that the optimal saddle-point (equilibrium) is $x^* = y^* = 0$.

If we run vanilla (projected) GDA with step size $\eta_t \geq 0$, then

$$\begin{aligned} x_{t+1} &= [x_t - \eta_t y_t]_{-1}^1 \\ y_{t+1} &= [y_t + \eta_t x_t]_{-1}^1, \end{aligned}$$

where $[z]_{-1}^1 := (z \wedge 1) \vee (-1)$ is the projection of z onto the interval $[-1, 1]$. Thus, we have

$$x_{t+1}^2 + y_{t+1}^2 \geq 1 \wedge [(x_t - \eta_t y_t)^2 + (y_t + \eta_t x_t)^2] = 1 \wedge [(1 + \eta_t^2)(x_t^2 + y_t^2)] \geq 1 \wedge (x_t^2 + y_t^2).$$

Therefore, if we do *not* initialize at the equilibrium $x^* = y^* = 0$, then the norm of (x_t, y_t) will always be lower bounded by $1 \wedge \|(x_0, y_0)\| > 0 = \|(x^*, y^*)\|$. In other words, (x_t, y_t) will not converge to (x^*, y^*) .

In fact, we can generalize the above failure to any bilinear matrix game:

$$\min_{\mathbf{x} \in C} \max_{\mathbf{y} \in D} \mathbf{x}^\top A \mathbf{y},$$

where C and D are closed sets and $A \in \mathbb{R}^{d \times d}$ is nonsingular. Suppose there exists $\epsilon > 0$ so that any equilibrium $(\mathbf{x}^*, \mathbf{y}^*)$ is ϵ away from the boundary: $\text{dist}(\mathbf{x}^*, \partial C) \wedge \text{dist}(\mathbf{y}^*, \partial D) > \epsilon$, then we know $A\mathbf{y}^* = \mathbf{0}$ hence $\mathbf{y}^* = \mathbf{0}$ and similarly $\mathbf{x}^* = \mathbf{0}$. It follows that vanilla projected gradient

$$\begin{aligned}\mathbf{x}_{t+1} &= P_C(\mathbf{x}_t - \eta_t A \mathbf{y}_t) \\ \mathbf{y}_{t+1} &= P_D(\mathbf{y}_t + \eta_t A^\top \mathbf{x}_t)\end{aligned}$$

will not converge to any equilibrium. Indeed, for vanilla gradient to converge, the projection will eventually be vacuous (otherwise we are ϵ away), but then for the equilibrium $(\mathbf{x}^* = \mathbf{0}, \mathbf{y}^* = \mathbf{0})$:

$$\|\mathbf{x}_{t+1}\|_2^2 + \|\mathbf{y}_{t+1}\|_2^2 = \|\mathbf{x}_t\|_2^2 + \|\mathbf{y}_t\|_2^2 + \eta_t^2 \left\| \begin{bmatrix} \mathbf{0} & A^\top \\ A & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{y}_t \\ \mathbf{x}_t \end{bmatrix} \right\|_2^2 = (1 + \eta_t^2 \sigma_{\min}^2(A))(\|\mathbf{x}_t\|_2^2 + \|\mathbf{y}_t\|_2^2),$$

which is strictly lower bounded by $\|\mathbf{x}_0\|_2^2 + \|\mathbf{y}_0\|_2^2$ if the starting point $(\mathbf{x}_0, \mathbf{y}_0)$ does not coincide with the equilibrium.

Algorithm 3.48: Alternating

The following alternating algorithm is often applied in practice for solving the **joint minimization** problem:

$$\min_{\mathbf{x} \in X} \min_{\mathbf{y} \in Y} f(\mathbf{x}, \mathbf{y}).$$

Algorithm: Alternating minimization for min-min

Input: $(\mathbf{x}_0, \mathbf{y}_0) \in \text{dom } f \cap X \times Y$

```

1 for  $t = 0, 1, \dots$  do
2    $\mathbf{x}_{t+1} = \text{argmin}_{\mathbf{x} \in X} f(\mathbf{x}, \mathbf{y}_t)$  // exact minimization
3    $\mathbf{y}_{t+1} = \text{argmin}_{\mathbf{y} \in Y} f(\mathbf{x}_{t+1}, \mathbf{y})$  // exact maximization
```

It is tempting to adapt alternating to solve min-max problems. The resulting algorithm, when compared to dual gradient (see Algorithm 3.42), is more aggressive in optimizing \mathbf{y} : dual gradient only takes a gradient ascent step while the latter finds the exact maximizer. Surprisingly, being more aggressive (and spending more effort) here actually hurts (sometimes).

Example 3.49: When to expect alternating to work?

Let us consider the simple problem

$$\min_{x \in [-1, 1]} \min_{y \in [-1, 1]} xy,$$

where two minimizers $\{(1, -1), (-1, 1)\}$ exist. Let us apply the alternating Line 3. Choose any $x > 0$, we obtain

$$y \leftarrow -1, x \leftarrow 1, y \leftarrow -1,$$

which converges to an optimal solution after 1 iteration!

To compare, let us consider the “twin” problem:

$$\min_{x \in [-1, 1]} \max_{y \in [-1, 1]} xy,$$

where a unique equilibrium $(x^*, y^*) = (0, 0)$ exists. Choose any $x > 0$ (the analysis for $x < 0$ is similar). Let us perform the inner maximization exactly to obtain $y = 1$. With $y = 1$ in mind, we perform the outer inner minimization exactly to obtain $x = -1$. Keep iterating, we obtain the cycle

$$y \leftarrow -1, x \leftarrow 1, y \leftarrow 1, x \leftarrow -1, y \leftarrow -1,$$

which is bounded away from the equilibrium! Of course, if we perform averaging along the trajectory we again obtain convergence. Dual gradient essentially avoids the oscillating behaviour of alternating by (implicitly) averaging.