Functional Optimization for Machine Learning: Exercises

Michael Unser*

June 6, 2024

1 Representer Theorems

1.1 Inverse problems with Tikhonov regularization

Let $\mathcal{H}$ be a Hilbert space with inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$. We recall that the Riesz map $\mathrm{R} : \mathcal{H} \to \mathcal{H}'$ is the unitary transform that maps $\mathcal{H}$ into its dual $\mathcal{H}'$ equipped with the inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}'}$. In effect, each point $f \in \mathcal{H}$ is mapped into its (unique) Hilbert conjugate $f^* = \mathrm{R}\{f\} \in \mathcal{H}'$. The unitary nature of R is expressed by the characteristic identity

$$\forall f, g \in \mathcal{H} : \quad \langle f, g \rangle_{\mathcal{H}} = \langle f, g^* \rangle = \langle f^*, g^* \rangle_{\mathcal{H}'} \tag{1}$$

where $f^* = \mathrm{R}\{f\}, g^* = \mathrm{R}\{g\} \in \mathcal{H}'$. Moreover, $\mathrm{R}^* = \mathrm{R}^{-1} : \mathcal{H}' \to \mathcal{H}$ so that the adjoint operator R^* is the Riesz map from $\mathcal{H}' \to \mathcal{H}'' = \mathcal{H}$, due to the reflexivity of Hilbert spaces. This invertibility property is summarized by $f^{**} = f$ for all $f \in \mathcal{H}$.

The goal here is to invoke the general representer theorem in order to solve the generic linear inverse problem

$$f_0 = \arg \min_{f \in \mathcal{H}} \left(\sum_{m=1}^{M} |y_m - \langle \nu_m, f \rangle|^2 + \lambda \|f\|_{\mathcal{H}}^2 \right) \tag{2}$$

where $\lambda \in \mathbb{R}^+$ is an adjustable regularization parameter. The only hypothesis is that the linear measurement functionals are well defined over $\mathcal{H}$—that is, $\nu_1, \dots, \nu_M \in \mathcal{H}' = \mathcal{X}$—and linearly independent.

Specifically, we want you address the following points:

1. Show that the solution is unique and can be written as $f_0 = \sum_{m=1}^{M} a_m \varphi_m$ where the φ_m are suitable basis functions.

2. Recast the problem as a finite dimensional optimization.

*Biomedical Imaging Group, École polytechnique fédérale de Lausanne (EPFL), Station 17, CH-1015, Lausanne, Switzerland (michael.unser@epfl.ch).

3. Solve 2. explicitly by showing that the optimal expansion coefficients in 1. can be evaluated as $\mathbf{a} = (\mathbf{H} + \lambda \mathbf{I}_M)^{-1}\mathbf{y}$ for some suitable "system" matrix $\mathbf{H} \in \mathbb{R}^{M\times M}$.

1.2 Application of Banach conjugates

The signal space of our interest is $\mathcal{X}' = (\mathbb{R}^N, \|\cdot\|_{\ell_p})$ with $p \geq 1$ whose predual is $\mathcal{X} = (\mathbb{R}^N, \|\cdot\|_{\ell_q})$ with $\frac{1}{p} + \frac{1}{q} = 1$. These spaces are strictly convex for $p \in (1, \infty)$.

1.2.1 Exercise on Compressed Sensing

Our goal is to recover an unknown signal $\mathbf{s} = (s_n) \in \mathbb{R}^N$ from a set of corrupted linear measurements $y_m = \mathbf{h}_m^{\mathsf{T}}\mathbf{s} + \epsilon_m$, $m = 1, \dots, M$ where ϵ_m represents some measurement noise. To favour sparse solutions, we formulate the problem as

$$\mathbf{s} = \arg\min_{\mathbf{x}\in\mathbb{R}^N} \|\mathbf{y} - \mathbf{H}\mathbf{x}\|^2 + \lambda\|\mathbf{x}\|^p \tag{3}$$

with p close to one. Here, $\mathbf{y} \in \mathbb{R}^M$ is the data vector, $\mathbf{H} = [\mathbf{h}_1\ \mathbf{h}_2\ \cdots\ \mathbf{h}_M]^{\mathsf{T}} \in \mathbb{R}^{M\times N}$ is the system matrix, and $\lambda \in \mathbb{R}^+$ is an adjustable regularization parameter.

1. Prove the Hölder inequality (7) for $p = 1$.
2. Prove that the duality map $\mathrm{J}_{\ell_q} : (\mathbb{R}^N, \|\cdot\|_{\ell_q}) \to (\mathbb{R}^N, \|\cdot\|_{\ell_p})$ is given by

$$\mathrm{J}_{\ell_q}(\mathbf{y}) = \mathbf{y}^* \quad \text{with} \quad [\mathbf{y}^*]_n = \frac{|y_n|^{q-1}}{\|\mathbf{y}\|_{\ell_q}^{q-2}} \operatorname{sign}(y_n) \tag{4}$$

3. Use the general representer theorem to give the parametric form of the solution for $p > 1$.

1.2.2 Theoretical background

Let $(\mathcal{X}, \mathcal{X}')$ be a dual pair of Banach spaces. Formally, an element of the dual space, $g \in \mathcal{X}'$, is a continuous linear functional on $\mathcal{X}$ whose action $g : \mathcal{X} \to \mathbb{R}$ results in the linear mapping $f \mapsto \langle g, f\rangle_{\mathcal{X}'\times\mathcal{X}} = \langle g, f\rangle$ (for short) with the induced (operator-)norm of this functional being

$$\|g\|_{\mathcal{X}'} = \sup_{f\in\mathcal{X}\setminus\{0\}} \frac{\langle g, f\rangle}{\|f\|_{\mathcal{X}}}. \tag{5}$$

This association also goes the other way around in that, by fixing $f \in \mathcal{X}$, one can specify a linear functional on $\mathcal{X}'$ via the rule $f : g \mapsto \langle g, f\rangle$. This is to say that there is a natural pairing between $\mathcal{X}$ and $\mathcal{X}'$ encoded in the bilinear map $\langle\cdot,\cdot\rangle : \mathcal{X}' \times \mathcal{X} \to \mathbb{R}$ (duality pairing) whose continuity (in both variables) is ruled by the duality bound

$$\forall (f, g) \in \mathcal{X} \times \mathcal{X}' : \quad |\langle g, f\rangle| \leq \|g\|_{\mathcal{X}'}\|f\|_{\mathcal{X}}. \tag{6}$$

This bound is tight because of (5). In fact, we like to view (6) as the Banach counterpart of the Cauchy-Schwartz inequality.

As prototypical example, we now consider the finite-dimensional Banach space $\mathcal{X} = (\mathbb{R}^N, \|\cdot\|_{\ell_p})$ with $p \geq 1$ and

$$\forall \mathbf{x} = (x_n) \in \mathbb{R}^N: \quad \|\mathbf{x}\|_{\ell_p} = \begin{cases} p \in [1, \infty) & \left(\sum_{n=1}^N |x_n|^p\right)^{1/p} \\ p = \infty & \sup_{n \in \{1,\dots,N\}} |x_n|. \end{cases} \tag{7}$$

The dual space is $\mathcal{X}' = (\mathbb{R}^N, \|\cdot\|_{\ell_q})$ with $\frac{1}{p} + \frac{1}{q} = 1$, which is the same set $\mathbb{R}^N$, but equipped with a different norm. The duality pairing between those two topological vector spaces is ruled by the Hölder inequality

$$\forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^N: \quad |\langle \mathbf{x}, \mathbf{y} \rangle| \leq \|\mathbf{x}\|_{\ell_p} \|\mathbf{y}\|_{\ell_q}. \tag{8}$$

Definition 1. *Let $(\mathcal{X}, \mathcal{X}')$ be a dual pair of Banach spaces. Then, the elements $f^* \in \mathcal{X}'$ and $f \in \mathcal{X}$ form a* **conjugate pair** *if*

- $\|f^*\|_{\mathcal{X}'} = \|f\|_{\mathcal{X}}$ *(norm preservation), and*
- $\langle f^*, f \rangle_{\mathcal{X}' \times \mathcal{X}} = \|f^*\|_{\mathcal{X}'} \|f\|_{\mathcal{X}}$ *(sharp duality bound).*

For any given $f \in \mathcal{X}$, the set of admissible conjugates defines the **duality mapping**

$$J(f) = \{f^* \in \mathcal{X}' : \|f^*\|_{\mathcal{X}'} = \|f\|_{\mathcal{X}} \text{ and } \langle f^*, f \rangle_{\mathcal{X}' \times \mathcal{X}} = \|f^*\|_{\mathcal{X}'} \|f\|_{\mathcal{X}}\},$$

which is a non-empty subset of $\mathcal{X}'$. Whenever the duality mapping is single-valued (for instance, when $\mathcal{X}'$ is strictly convex), one also defines the duality operator $\mathrm{J}_{\mathcal{X}} : \mathcal{X} \to \mathcal{X}'$, which is such that $f^* = \mathrm{J}_{\mathcal{X}}(f)$.

2 Deep splines with stability control

2.1 Simple convolution layer

We shall investigate the stability properties of the convolution operator $\mathrm{T}_h : \ell_2(\mathbb{Z}) \to \ell_2(\mathbb{Z})$, which is such that $\mathrm{T}_h\{x\}[n] = (h * x)[n] = \sum_{k\in\mathbb{Z}} h[k]x[n-k]$. The Lipschitz constant of T_h is the operator norm

$$\mathrm{Lip}(\mathrm{T}_h) = \|\mathrm{T}_h\| = \sup_{x\in\ell_2(\mathbb{Z})\setminus\{0\}} \frac{\|h * x\|_{\ell_2}}{\|x\|_{\ell_2}}. \tag{9}$$

This convolution operator can also be characterized by its frequency response $\hat{h}(\omega) = \mathcal{F}_\mathrm{d}\{h\} = \sum_{k\in\mathbb{Z}} h[k]\mathrm{e}^{-\mathrm{j}\omega k}$ where the operator $\mathcal{F}_\mathrm{d} : \ell_2(\mathbb{Z}) \to L_2([-\pi,\pi])$ is the discrete Fourier transform. The latter is an isometry, as expressed Parseval's formula:

$$\forall x \in \ell_2(\mathbb{Z}) : \quad \|x\|_{\ell_2}^2 = \sum_{k\in\mathbb{Z}} |x[k]|^2 = \int_{-\pi}^{+\pi} |\hat{x}(\omega)|^2 \frac{\mathrm{d}\omega}{2\pi} = \|\hat{x}\|^2_{L_2([-\pi,\pi])}. \tag{10}$$

1. Use the discrete Fourier transform and Parseval's formula to get an estimate of $\mathrm{Lip}(\mathrm{T}_h)$.
2. Show that $\mathrm{Lip}(\mathrm{T}_h) = \|\mathrm{T}_h\| \le \sum_{k\in\mathbb{Z}} |h[k]| = \|h\|_{\ell_1(\mathbb{Z})}$, the advantage of the latter being that it is very easy to calculate.
3. Can you ensure that the bound in Item 1 is tight? Are there scenarios where the simpler bound in Item 2 is tight as well?

2.2 Pointwise nonlinearities

Let us consider the following functions, which are shown in Fig. 1:

$$\mathrm{tri}(x) = \begin{cases} 0, & x \le -1 \\ 1+x, & x \in [-1,0) \\ 1-x, & x\in[0,1) \\ 0, & x \ge 1 \end{cases} \tag{11}$$

$$f(x) = \tfrac{1}{2} + x - \mathrm{ReLU}(x-1) \tag{12}$$

$$g(x) = \begin{cases} 0, & x \le 0 \\ x, & x \in [0,1) \\ 1-\frac{1}{2}(x-1), & x\in[1,3) \\ 0, & x \ge 3 \end{cases} \tag{13}$$

where $\mathrm{ReLU}(x) = (x)_+ = \max(x, 0)$.

1. Compute f', $\mathrm{Lip}(f)$, f'', $\mathrm{TV}^{(2)}(f) = \|\mathrm{D}^2 f\|_{\mathcal{M}}$.
   Hint: Use the property that $\|\delta(\cdot - x_0)\|_{\mathcal{M}} = 1$ *for any* $x_0 \in \mathbb{R}$.

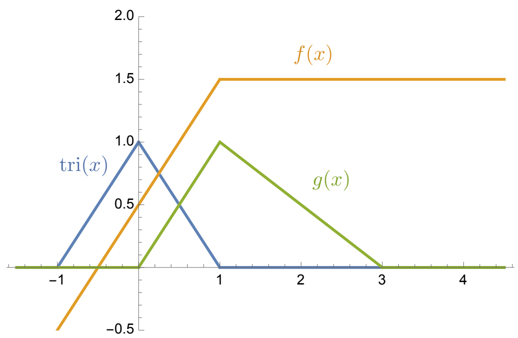

Figure 1: Continuous piece-wise linear (CPWL) functions/activations.

2. Compute g', $\mathrm{Lip}(g)$, $\mathrm{TV}^{(2)}(g) = \|\mathrm{D}^2 g\|_{\mathcal{M}}$.

3. Express f as $f(x) = b_0 + b_1 x + a|x-1|$ with suitable values of $b_0, b_1, a \in \mathbb{R}$,.

4. Can you express f as $f(x) = \sum_{k\in\mathbb{Z}} c[k]\mathrm{tri}(x-k)$? If yes, give the explicit value of the B-spline coefficients $c[k]$.

5. Express g as $g(x) = \sum_k a_k \mathrm{ReLU}(x - x_k)$ with suitable values of (a_k, x_k) and show that $\mathrm{TV}^{(2)}(g) = \sum_k |a_k|$.

6. Let $\sigma(x) = \sum_{k\in\mathbb{Z}} c[k]\mathrm{tri}(x-k)$. Give the explicit expression of $\mathrm{Lip}(\sigma)$ and $\mathrm{TV}^{(2)}(\sigma)$ in terms of the B-spline coefficients $c[k]$.
Hint: Express $\sigma'(x)$ in terms of rectangular basis functions.