Shape Matching, A Metric Geometry Approach
CS468-Fall-2008
Lecture-2

Facundo Mémoli
Department of Mathematics, Stanford University
memoli@math.stanford.edu

October 6th, 2008

Contents

1 Summary of first Lecture and some proofs **1**

2 Introducing invariances **2**

3 Assignment for next class **5**

1 Summary of first Lecture and some proofs

- For a compact metric space (X, d), $d_{\mathcal{H}}^{Z}(,)$ is a metric on $\mathcal{C}(Z)$. I gave as an exercise proving that $d_{\mathcal{H}}^{Z}(,)$ satisfies the triangle inequality. For this, I suggested you use the expression $d_{\mathcal{H}}^{Z}(A, B) = \inf_{R \in \mathcal{R}(A,B)} \|d\|_{L^{\infty}(R)}$ and then the *composition of correspondences.*

- We saw another class of objects: weighted objects that were specified thru a probability measure on Z. These weights can be interpreted as signaling the relative importance of different points in the object. The weights can come from a sampling process. When you don't have a good reason to assign a larger weight to a point, just assign all points the same weight: *uniform distribution.*

- Recall the classification problem that I used as a motivation. The goal is to define a notion **dist** of dissimilarity between objects in a certain class $\mathcal{O}(Z)$ to be specified as well. We will be using the following construction, all the time:

$$(Z, d) \mapsto (\mathcal{O}(Z), \mathbf{dist}).$$

  We've seen two possible choices, compact subsets of Z endowed with Hausdorff distance on the one hand, and weighted-compact subset endowed with Wasserstein metric on the other.

- For a compact metric space (X, d), $d_{\mathcal{W},p}^{Z}(,)$ is a metric on $\mathcal{C}_w(Z)$. In this case, we have that $d_{\mathcal{W},p}^{Z}(A, B) = \min_{\mu \in \mathcal{M}(\mu_A, \mu_B)} \|d\|_{L^p(\mu)}$. We will prove the triangle inequality. This technique will be used in later lectures. Consider $\mu_1 \in \mathcal{M}(\mu_A, \mu_B)$ and $\mu_2 \in \mathcal{M}(\mu_B, \mu_C)$. Consider the measure

$$\mu(a, c) := \sum_{b \in B} \frac{\mu_1(a, b) \cdot \mu_2(b, c)}{\mu_B(b)}.$$

  This works if we assume that supp $[\mu_B] = B$ (so that denominator doesn't gives us trouble.)

1. By definition $\mu(a,c) \geqslant 0$.
2. Notice that $\sum_{a\in A}\mu(a,c) = \mu_C(c)$ and that similarly $\sum_{c\in C}\mu(a,c) = \mu_A(a)$.
3. Then, $\mu \in \mathcal{M}(\mu_A,\mu_C)$.
   So we have something that we can call *composition of measure couplings.* Now, proving the triangle inequality for the Wasserstein metric is a matter of using inequalities, in a manner similar to the proof of the triangle inequality above.

- Notice that $W(\mu) := \|d\|^p_{L^p(\mu)} = \sum_{a,b} d^p(a,b)\mu(a,b)$ is a continuous function of $\mu \in \mathcal{M}(\mu_A,\mu_B)$ and since $\mathcal{M}(\mu_A,\mu_B)$ is *compact*, then we can always find a minimizer and in effect then inf can be replaced by min in the definition of the Wasserstein distance. Arguing compactness in the general case of non-necessarily finitely supported probability measures takes more work.

- So, the proof of the triangle inequality for the Wasserstein distance.

- How is $d^Z_{\mathcal{H}}(,)$ related to $d^Z_{\mathcal{W},p}(,)$? We have that for all $A,B \in \mathcal{C}(Z)$ and μ_A and μ_B s.t. supp $[\mu_A] = A$ and supp $[\mu_B] = B$:
$$d^Z_{\mathcal{H}}(A,B) \leqslant d^Z_{\mathcal{W},\infty}(A,B).$$

- **Why insist on having a metric on objects?** First of all, one must have a notion of equality between objects. This is where the invariances enter the game. When $Z = \mathbb{R}^d$, if you want to admit invariance to Euclidean isometries, then you'd say that $A = B$ whenever $\exists T \in E(d)$ s.t. $A = T(B)$.

  1. $\mathbf{dist}(A,B) = 0$ if and only if $A = B$. You require this so that your metric will declare that two points are at zero distance iff they are the same according to your notion of invariance.
  2. The triangle inequality. This property is important for more than one reason. A very important reason is that it guarantees *stability.* To fix ideas assume that $\mathcal{O}(Z) = \mathcal{C}(Z)$ and $\mathbf{dist} = d^Z_{\mathcal{H}}(,)$. Assume $\mathbb{A} \subset A$ and $\mathbb{B} \subset B$ are ε-nets for A and B respectively. Then,
  $$|d^Z_{\mathcal{H}}(A,B) - d^Z_{\mathcal{H}}(\mathbb{A},\mathbb{B})| \leqslant d^Z_{\mathcal{H}}(A,\mathbb{A}) + d^Z_{\mathcal{H}}(B,\mathbb{B}) \leqslant 2\varepsilon.$$

     What you'd like to compute is $d^Z_{\mathcal{H}}(A,B)$, but if you have access only to $\mathbb{A}$ and $\mathbb{B}$, then your error is bounded above by the error of approximation of the objects by their samples, ε. In other words, if you sample more finely, your error will become smaller. This is stability.

  The triangle inequality can be replaced by other structurally similar conditions, such as $\mathbf{dist}(A,B) \leqslant K\cdot(\mathbf{dist}(A,C)+\mathbf{dist}(C,B))$, $\mathbf{dist}(A,B) \leqslant K\cdot\max(\mathbf{dist}(A,C),\mathbf{dist}(C,B))$, etc.

2 Introducing invariances

- Assume $Z = \mathbb{R}^d$ and that objects are those in $\mathcal{C}(Z)$. Further, assume that your notion of equality is that given by the existence of a Euclidean isometry between the objects. A choice of **dist** compatible with these assumtptions is
$$d^{\mathbb{R}^d,rigid}_{\mathcal{H}}(A,B) := \inf_{T\in E(d)} d^Z_{\mathcal{H}}(A,T(B)).$$

- Assume $Z = \mathbb{R}^d$ and that objects are those in $\mathcal{C}_w(Z)$. Assume that your notion of equality is that given by the existence of a Euclidean isometry between the objects. A choice of **dist** compatible with these assumtptions is
$$d^{\mathbb{R}^d,rigid}_{\mathcal{W},p}(A,B) := \inf_{T\in E(d)} d^Z_{\mathcal{W},p}(A,T(B)).$$

- These family of ideas work well when one has an *embedding space*, Euclidean space in this case, but roughly the same applies to any compact metric space.

- But what can we do in order to define a reasonable **dist**, which is a true metric that accounts for invariances to *bends*? The papers by Elad and Kimmel, Hamza and Krim, Hilaga et al. all provide partial answers to this question, central to this course.

- This answers are partial in that the notions of dissimilarity they propose are not metrics. They fail to be metrics because they do not satisfy neither triangle inequality nor the property that two shapes will be at zero "distance" if they are bend-isometric.

- The idea of using the Gromov-Hausdorff distance between metric spaces takes care of this difficulty, [**MS05**].

- We now argue about this, But let's go back to Euclidean subsets for a while. Remember the *Folklore Lemma*. If $A, B \subset \mathbb{R}^d$ are finite sets of points with the same cardinality n are s.t. there exist a permutation $\pi \in \Pi(n)$ s.t. $\|a_i - a_j\| = \|b_{\pi(i)} - b_{\pi(j)}\|$ for all i, j, then there exists $T \in E(d)$ s.t. $A = T(B)$.

- We use the following intuition. Let A and B be two finite sets of points in $\mathbb{R}^d$. In order to ascertain how close these points are to being isometric in the Euclidean sense (existence of $T \in E(d)$ s.t. $A = T(B)$), we try to measure how different are the distance matrices

$$\mathbf{D}(A) \text{ and } \mathbf{D}(B).$$

- Then one may think of computing an L^2 measure

$$\min_{\pi} \sum_{ij} (\|a_i - a_j\| - \|b_{\pi(i)} - b_{\pi(j)}\|)^2,$$

  or a worst case kind of measure

$$\min_{\pi} \max_{ij} \left| \|a_i - a_j\| - \|b_{\pi(i)} - b_{\pi(j)}\| \right|$$

  or something of the sort.

- In general, A and B may have different cardinality. So, we can extend the above to that case by considering *correspondences* instead of permutations. We will keep track only of the worst case measure above. The corresponding generalization would be

$$D_E(A, B) := \min_{R \in \mathcal{R}(A,B)} f(R)$$

  where

$$f(R) := \max_{(a,b),(a',b') \in R} \left| \|a - a'\| - \|b - b'\| \right|.$$

- A couple of questions are in order, (1) does $D_E(A, B)$ above define a metric, (2) how does it compare to $d_{\mathcal{H}}^{\mathbb{R}^d, rigid}(A, B)$.

- We will see that the same idea can be generalized, and we will prove the triangle inequality for it. Regarding (2), we can make the following argument. Fix $T \in E(d)$. Then, $\|T(b - b')\| = \|b - b'\|$ and hence, for all $(a, b), (a', b') \in R$

$$\left| \|a - a'\| - \|b - b'\| \right| = \left| \|a - a'\| - \|T(b) - T(b')\| \right| \leqslant \|a - T(b)\| + \|a' - T(b')\|.$$

Then,

$$\begin{aligned} f(R) &\leqslant \max\,(a,b),(a',b') \in R\left(\|a-T(b)\| + \|a'-T(b')\|\right) \\ &= \max_{(a,b)\in R}\|a-T(b)\| + \max_{(a',b')\in R}\|a'-T(b')\| \\ &= 2\max_{(a,b)\in R}\|a-T(b)\|. \end{aligned}$$

Then, $D_E(A,B) \leqslant 2d_{\mathcal{H}}^{\mathbb{R}^d}(A,T(B))$. Since, T was arbitrary in $E(d)$, we conclude that

$$(*) \qquad D_E(A,B) \leqslant 2 \cdot d_{\mathcal{H}}^{\mathbb{R}^d,rigid}(A,B).$$

- Note that the inequality will not hold in general. In fact, consider A equal to an equilateral triangle with side length 1 and $B = \{p\}$. Then, $d_{\mathcal{H}}^{\mathbb{R}^d,rigid}(A,B) = 1/\sqrt{3}$ but $D_E(A,B) = 1$ and $1 < 2/\sqrt{3}$.
- We will se later how one can obtain an upper bound for $d_{\mathcal{H}}^{\mathbb{R}^d,rigid}(A,B)$ in terms of $D_E(A,B)$. This and related material can be read in [**M08-euclidean**].
- Let's look into $(*)$ above. $D_E(A,B)$ is to be regarded as an *intrinsic quantity* as it depends on the ambient space $Z = \mathbb{R}^d$ only via the metric. In fact, one can compute the same expression for any pair of metric spaces!
- $d_{\mathcal{H}}^{\mathbb{R}^d,rigid}(A,B)$ on the other hand, is an *extrinsic quantity* since it's computation depends on being able to perform *ambient space isometries*. This is a drawback when optimizing over $E(d)$ is expensive, think of a very large d. So, developing methods for computing *intrinsic* (as opposed to extrinsic) dissimilarity measures is interesting.
- For (compact) metric spaces (X,d_X) and (Y,d_Y) define

  $$\mathbf{dist}_1(X,Y) := \frac{1}{2}\min_{R\in\mathcal{R}(X,Y)}\max_{(x,y),(x',y')\in R}|d_X(x,x') - d_Y(y,y')|.$$

  Risking being tremendously obvious, notice that when $X,Y \subset \mathbb{R}^d$ are endowed with the Euclidean metric, then $D_E(X,Y) = 2 \cdot \mathbf{dist}_1(X,Y)$.
- Another idea for defining a reasonable **dist** is the following. Imagine that out of $\mathbf{D}(A)$ and $\mathbf{D}(B)$ you try to construct a bigger matrix $\mathbf{D}$ that has $\mathbf{D}(A)$ and $\mathbf{D}(B)$ as submatrices as follows:

  $$\mathbf{D} = \begin{pmatrix} \mathbf{D}(A) & (*) \\ (*)' & \mathbf{D}(B) \end{pmatrix}.$$

  Assume that in addition you impose that $\mathbf{D}$ be a distance matrix, namely, that is has zeros along the diagonal and that is satisfies all triangle inequalities. If you managed to do this, you would have constructed $\mathbf{D}$ that is a metric on $A \cup B$! We call all such choices of $\mathbf{D}$ *admissible* for A and B.
- One can then view $(A\cup B,\mathbf{D})$ as a metric space inside which A and B sit isometrically, since $\mathbf{D}(a,a') = \mathbf{D}_A(a,a')$ and $\mathbf{D}(b,b') = \mathbf{D}_B(b,b')$, for all $a,a' \in A$ and $b,b' \in B$, by construction. One says that $\mathbf{D}$ *glues* (A,d_A) and (B,d_B) together.
- Then, one can compute

  $$g(\mathbf{D}) := d_{\mathcal{H}}^{(A\cup B,\mathbf{D})}(A,B)$$

  and eventually minimize over all choices of admissible $\mathbf{D}$.(!) So, let

  $$\mathbf{dist}_2(A,B) := \min_{\mathbf{D}} g(\mathbf{D}).$$

- A few comments are in order.

  1. Note that dist_2 is *intrinsic*, so it can be defined for every pair of compact metric spaces.
  2. Let $\mathcal{D}(d_A, d_B)$ denote the set of all metrics on $A \cup B$ admissible for A and B. We must check that
  $$\mathcal{D}(d_A, d_B) \neq \varnothing$$
  in order for all this to make sense. But, in fact, for $a \in A$ and $b \in B$, let $d \in \mathcal{D}(d_A, d_B)$ be given by
  $$d(a,b) = \frac{\max(\mathbf{diam}\,(X)\,, \mathbf{diam}\,(Y))}{2}.$$
  3. is $\mathbf{dist}_2$ really a metric?
  4. how does $\mathbf{dist}_2$ relate to $\mathbf{dist}_1$?

- Finally, a more abstract approach is the following. Assume (Z, d_Z) is a *sufficiently rich* metric space inside which one can find $A', B' \subset Z$ s.t. A' is isometric to A and B' is isometric to B. Then, compute
$$h(Z, A', B') = d^Z_{\mathcal{H}}(A', B')$$
and finally choose the Z, A', B' that minimize $h(Z, A', B')$. This is the definition of the *Gromov-Hausdorff* distance between metric spaces, [**Gromov**]. We will denote the infimal value of h by $\mathbf{dist}_3(A, B)$ and also, by $d_{\mathcal{GH}}(A, B)$.

- The same comments apply to $\mathbf{dist}_3$.

  1. Is this construction feasible, that is, can I always find Z, A' and B' in Z with the property that $A \sim_{iso} A'$ and $B \sim_{iso} B'$? The answer is yes! Clearly, the construction of $\mathbf{dist}_2$ yields $Z = A \cup B$ and a metric that reduces to d_A and d_B on $A \times A$ and $B \times B$.
  2. It follows from the previous item that therefore,
  $$\mathbf{dist}_3 \leqslant \mathbf{dist}_2.$$
  3. Does $\mathbf{dist}_3$ yield a metric for compact metric spaces?
  4. How are all the $\mathbf{dist}_i$ related?

- From now on, let $\mathfrak{M}$ denote the collection of all compact metric spaces. We have the following

  Proposition 1. *For all $A, B \in \mathfrak{M}$,*
  $$\mathbf{dist}_1(A, B) = \mathbf{dist}_2(A, B) = \mathbf{dist}_3(A, B) \;\; (= d_{\mathcal{GH}}(A, B))\,.$$

3 Assignment for next class

1. Finish the proof of the triangle inequality for $d^Z_{\mathcal{W},p}(,)$.
2. Finish [**Elad-Kimmel**] in order to get a good idea of the algorithm they proposed. You must understand this!
3. Read all the theory sections of [**MS05**].
4. Next class: we delve into the GH distance and its properties.