L^p-Gromov-Hausdorff distances for Shape Comparison

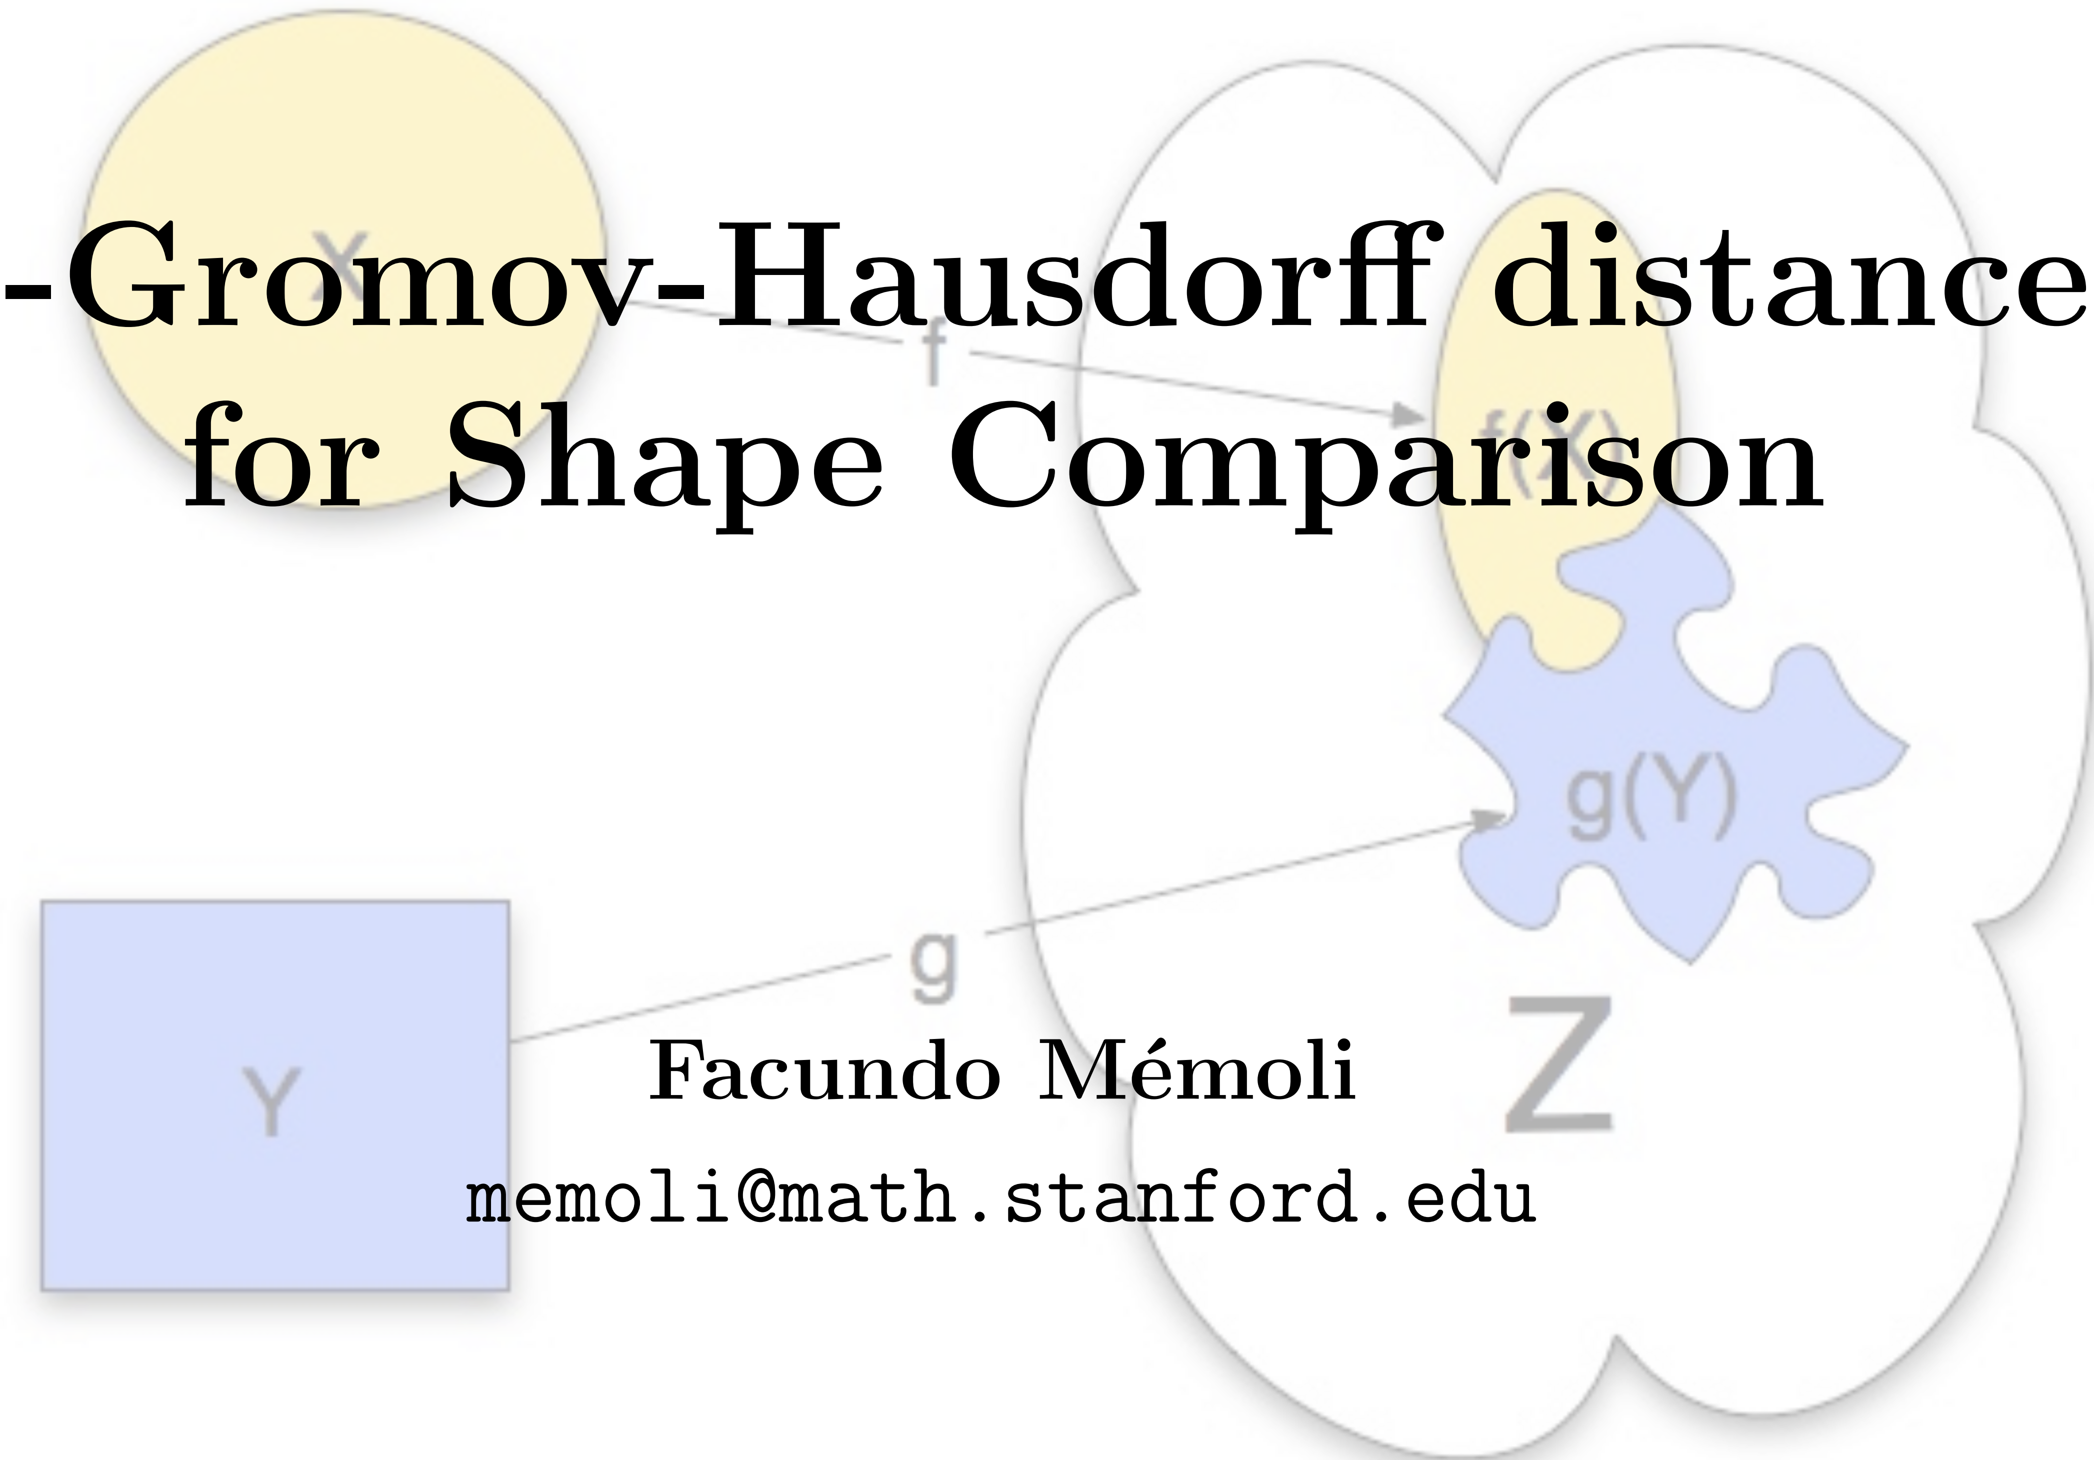

Facundo Mémoli

memoli@math.stanford.edu

The GH distance for Shape Comparison, [MS04,05]

- Regard shapes as (compact) metric spaces. Let $\mathcal{X}$ denote set of all compact metric spaces. Define metric on $\mathcal{X}$, then $(\mathcal{X}, d_{\mathcal{GH}})$ is itself a metric space.
- The metric with which one endows the shapes depends on the desired invariance. For example, if invariance to
  - *rigid isometries* is desired, use Euclidean distance.
  - *bends* is desired, use "intrinsic" distance.
- GH distance provides reasonable framework for Shape Comparison: good theoretical properties.

Main Properties

1. Let (X, d_X), (Y, d_Y) and (Z, d_Z) be metric spaces then

$$d_{\mathcal{GH}}(X,Y) \leq d_{\mathcal{GH}}(X,Z) + d_{\mathcal{GH}}(Y,Z).$$

2. If $d_{\mathcal{GH}}(X,Y) = 0$ and (X, d_X), (Y, d_Y) are compact metric spaces, then (X, d_X) and (Y, d_Y) are isometric.

3. Let $\mathbb{X}_n = \{x_1, \ldots, x_n\} \subset X$ be a finite subset of the compact metric space (X, d_X). Then,

$$d_{\mathcal{GH}}(X, \mathbb{X}_n) \leq d_{\mathcal{H}}(X, \mathbb{X}_n).$$

4. For compact metric spaces (X, d_X) and (Y, d_Y):

$$\begin{aligned} \frac{1}{2}\left|\mathrm{diam}(X) - \mathrm{diam}(Y)\right| &\leq d_{\mathcal{GH}}(X,Y) \\ &\leq \frac{1}{2}\max\left(\mathrm{diam}(X), \mathrm{diam}(Y)\right) \end{aligned}$$

Stability

$$|d_{\mathcal{GH}}(X,Y) - d_{\mathcal{GH}}(X_n, Y_m)| \leq r(X_n) + r(Y_m)$$

for finite samplings $X_n \subset X$ and $Y_m \subset Y$, where $r(X_n)$ and $r(Y_m)$ are the covering radii.

Critique

- Was not able to show connections with (sufficiently many) pre-existing appraches
- Computationally hard: currently only two attempts have been made:
  - [**MS04,MS05**] and [**BBK06**] only for surfaces.
  - [**MS05**] gives probabilistic guarantees for estimator based on sampling parameters.
  - Full generality leads to a hard **combinatorial optimization problem**: QAP.

Critique

- Was not able to show connections with (sufficiently many) pre-existing appraches
- Computationally hard: currently only two attempts have been made:
  - [**MS04,MS05**] and [**BBK06**] only for surfaces.
  - [**MS05**] gives probabilistic guarantees for estimator based on sampling parameters.
  - Full generality leads to a hard **combinatorial optimization problem**: QAP.

Desiderata

- Obtain an L^p version of the GH distance that:
  - retains theoretical underpinnings
  - its implementation leads to easier (continuous, quadratic, with linear constraints) optimization problems
  - can be related to pre-existing approaches (shape contexts, shape distributions, Hamza-Krim,..) via lower/upper bounds.

Gromov-Hausdorff

$$\Downarrow$$

Gromov-<u>Wasserstein</u>

(Kantorovich, Rubinstein, Earth Mover's Distance, Mass Transportation)

correspondences and the Hausdorff distance

Definition [Correspondences]

For sets A and B, a subset $R \subset A \times B$ is a *correspondence* (between A and B) if and and only if

- $\forall\ a \in A$, there exists $b \in B$ s.t. $(a, b) \in R$
- $\forall\ b \in B$, there exists $a \in A$ s.t. $(a, b) \in R$

Let $\mathcal{R}(A, B)$ denote the set of all possible correspondences between sets A and B. Note that in the case $n_A = n_B$, correspondences are larger than bijections.

correspondences

Note that when A and B are finite, $R \in \mathcal{R}(A, B)$ can be represented by a matrix $((r_{a,b})) \in \{0, 1\}^{n_A \times n_B}$ s.t.

$$\sum_{a \in A} r_{ab} \geq 1 \;\; \forall b \in B$$

$$\sum_{b \in B} r_{ab} \geq 1 \;\; \forall a \in A$$

correspondences

Note that when A and B are finite, $R \in \mathcal{R}(A,B)$ can be represented by a matrix $((r_{a,b})) \in \{0,1\}^{n_A \times n_B}$ s.t.

$$\sum_{a \in A} r_{ab} \geq 1 \;\; \forall b \in B$$

$$\sum_{b \in B} r_{ab} \geq 1 \;\; \forall a \in A$$

Proposition
Let (X,d) be a compact metric space and $A, B \subset X$ be compact. Then

$$d_{\mathcal{H}}(A,B) = \inf_{R \in \mathcal{R}(A,B)} \|d\|_{L^\infty(R)}$$

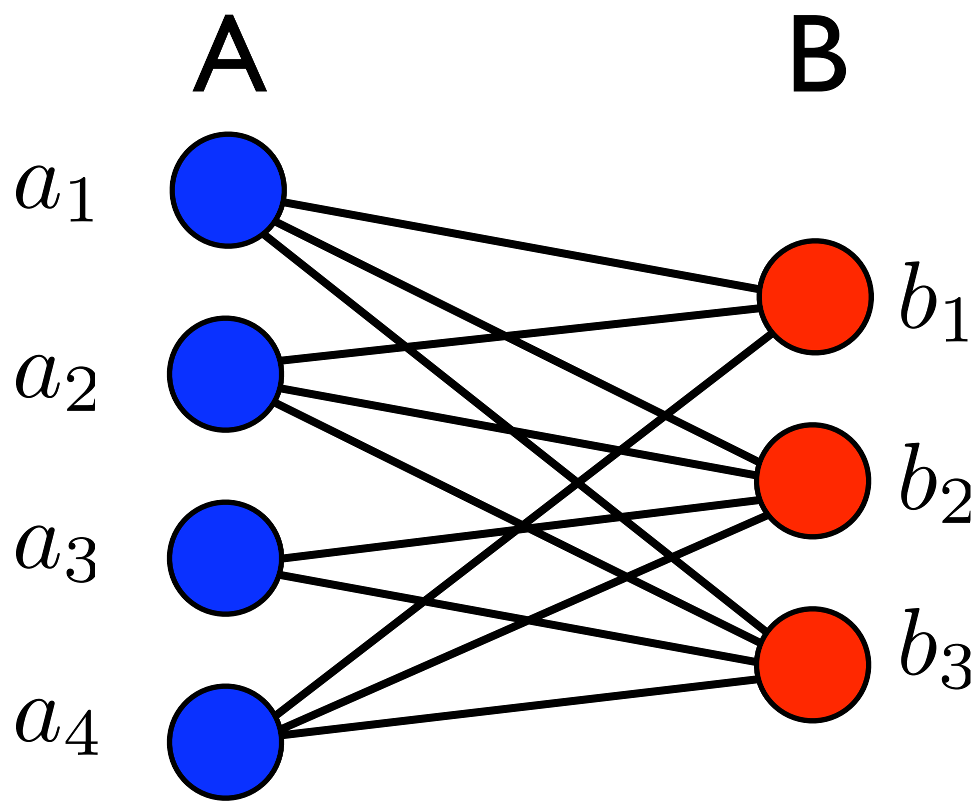

- Edges have weights: if $e = (i, j)$, $w_e = d(a_i, b_j)$.
- Interpret A and B as two groups of people that know each other.
- Interpret the value of w_e as the degree of *animosity* between a_i and b_j.
- What is the subset L of edges that leaves no point in $A \cup B$ isolated that minimizes the maximal weight:

$$\max_{e \in L} w_e$$

  that is

- We want minimize the maximal animosity.

correspondences and measure couplings

Let (A, μ_A) and (B, μ_B) be compact subsets of the compact metric space (X, d) and μ_A and μ_B be **probability measures** supported in A and B respectively.

Definition [Measure coupling] Is a probability measure μ on $A \times B$ s.t. (in the finite case this means $((\mu_{a,b})) \in [0,1]^{n_A \times n_B}$)

- $\sum_{a \in A} \mu_{ab} = \mu_B(b) \;\; \forall b \in B$
- $\sum_{b \in B} \mu_{ab} = \mu_A(a) \;\; \forall a \in A$

Let $\mathcal{M}(\mu_A, \mu_B)$ be the set of all couplings of μ_A and μ_B.
Notice that in the finite case, $((\mu_{a,b}))$ must satisfy $n_A + n_B$ *linear* constraints.

correspondences and measure couplings

Proposition $[(\mu \leftrightarrow R)]$

- Given (A, μ_A) and (B, μ_B), and $\mu \in \mathcal{M}(\mu_A, \mu_B)$, then

$$R(\mu) := \text{supp}(\mu) \in \mathcal{R}(A, B).$$

- König's Lemma. [gives conditions for $R \to \mu$]

Wasserstein distance

$$d_{\mathcal{H}}(A, B) = \inf_{R \in \mathcal{R}(A,B)} \|d\|_{L^\infty(R)}$$

$$\Downarrow \ (R \leftrightarrow \mu)$$

$$d_{\mathcal{W},\infty}(A, B) = \inf_{\mu \in \mathcal{M}(\mu_A,\mu_B)} \|d\|_{L^\infty(R(\mu))}$$

$$\Downarrow \ (L^\infty \leftrightarrow L^p)$$

$$d_{\mathcal{W},p}(A, B) = \inf_{\mu \in \mathcal{M}(\mu_A,\mu_B)} \|d\|_{L^p(A \times B,\mu)}$$

Wasserstein distance

$$d_{\mathcal{H}}(A,B) = \inf_{R \in \mathcal{R}(A,B)} \|d\|_{L^\infty(R)}$$

$$\Downarrow \; (R \leftrightarrow \mu)$$

$$d_{\mathcal{W},\infty}(A,B) = \inf_{\mu \in \mathcal{M}(\mu_A,\mu_B)} \|d\|_{L^\infty(R(\mu))}$$

$$\Downarrow \; (L^\infty \leftrightarrow L^p)$$

$$d_{\mathcal{W},p}(A,B) = \inf_{\mu \in \mathcal{M}(\mu_A,\mu_B)} \|d\|_{L^p(A\times B,\mu)}$$

Wasserstein distance

$$d_{\mathcal{H}}(A, B) = \inf_{R \in \mathcal{R}(A,B)} \|d\|_{L^\infty(R)}$$

$$\Downarrow \ (R \leftrightarrow \mu)$$

$$d_{\mathcal{W},\infty}(A, B) = \inf_{\mu \in \mathcal{M}(\mu_A, \mu_B)} \|d\|_{L^\infty(R(\mu))}$$

$$\Downarrow \ (L^\infty \leftrightarrow L^p)$$

$$d_{\mathcal{W},p}(A, B) = \inf_{\mu \in \mathcal{M}(\mu_A, \mu_B)} \|d\|_{L^p(A \times B, \mu)}$$

Wasserstein distance

$$d_{\mathcal{H}}(A,B) = \inf_{R \in \mathcal{R}(A,B)} \|d\|_{L^\infty(R)}$$

$$\Downarrow \; (R \leftrightarrow \mu)$$

$$d_{\mathcal{W},\infty}(A,B) = \inf_{\mu \in \mathcal{M}(\mu_A,\mu_B)} \|d\|_{L^\infty(R(\mu))}$$

$$\Downarrow \; (L^\infty \leftrightarrow L^p)$$

$$d_{\mathcal{W},p}(A,B) = \inf_{\mu \in \mathcal{M}(\mu_A,\mu_B)} \|d\|_{L^p(A\times B,\mu)}$$

The plan

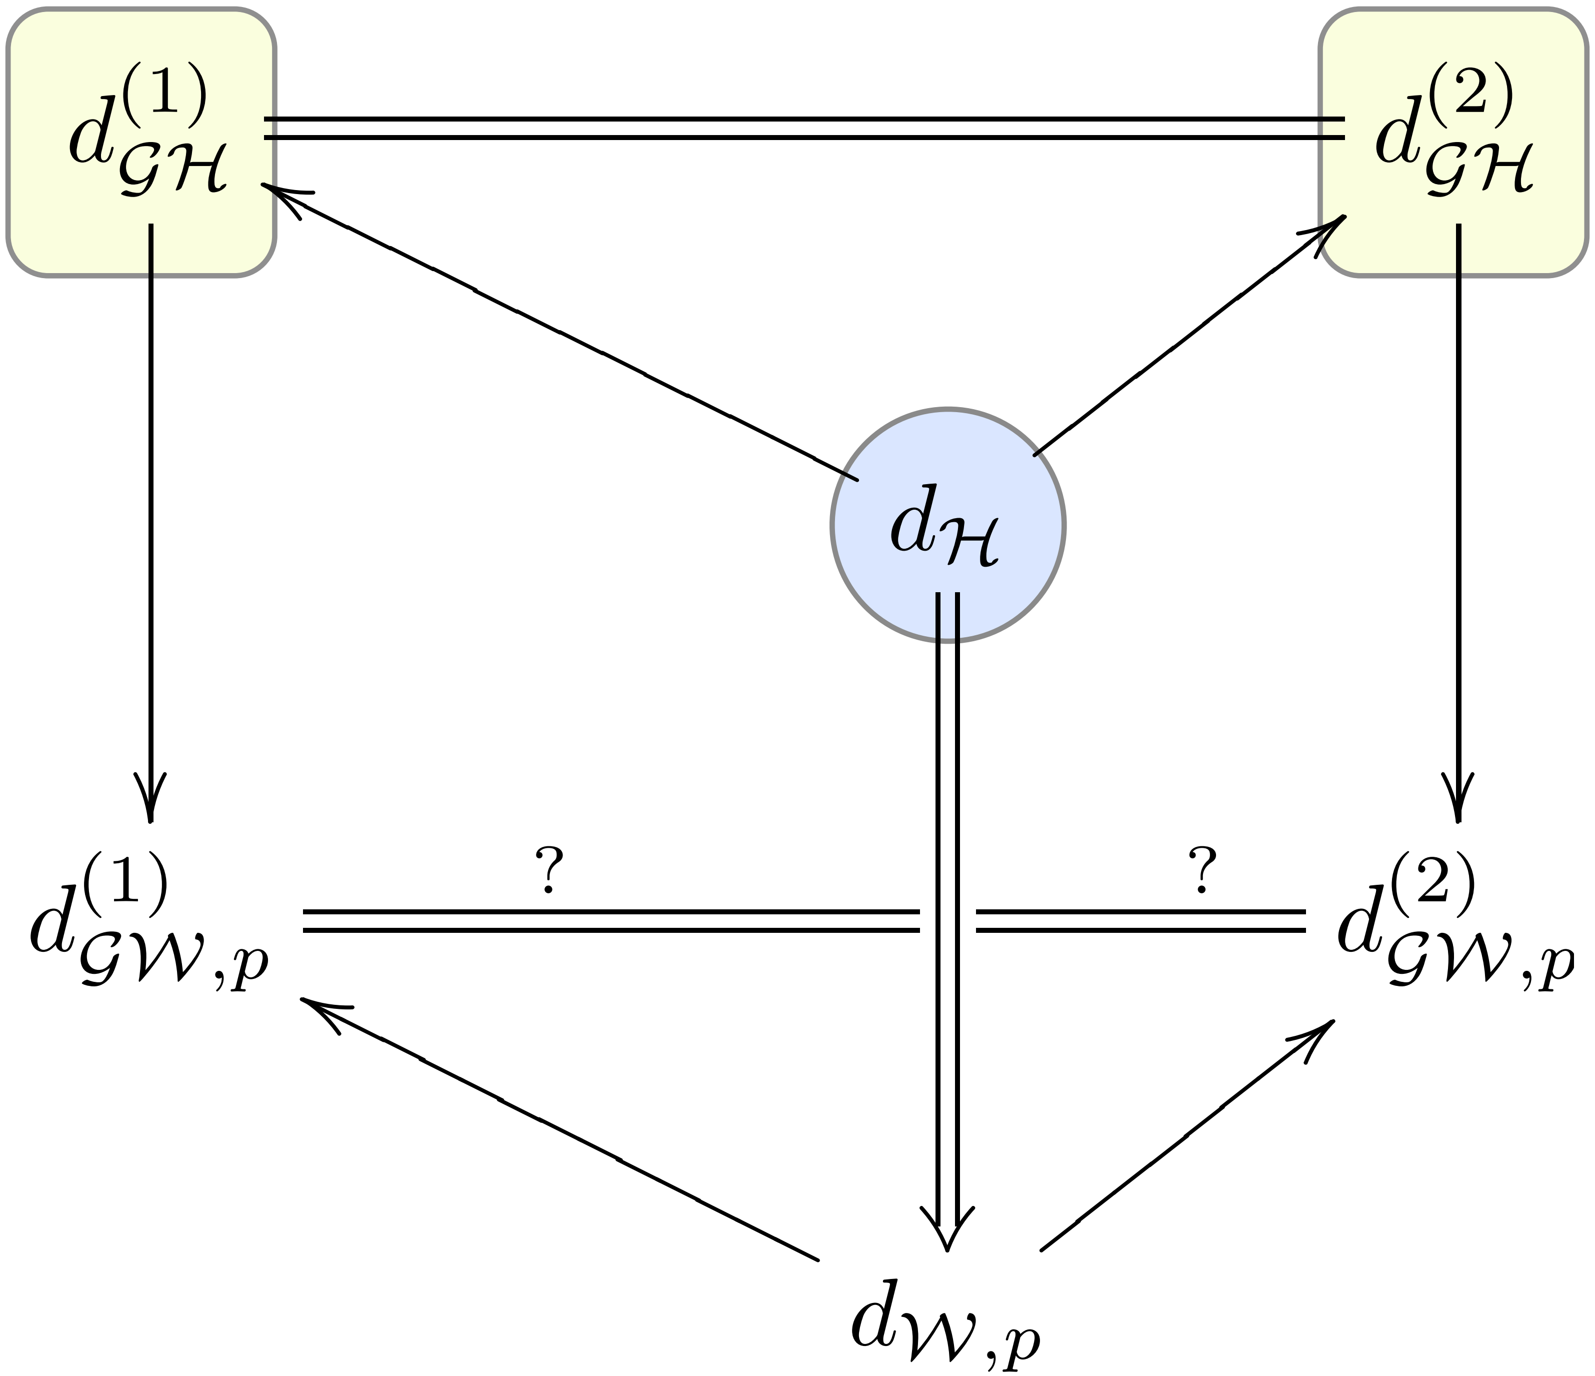

$d_{\mathcal{GH}}^{(1)}$
$d_{\mathcal{GH}}^{(2)}$
$d_{\mathcal{H}}$
$d_{\mathcal{GW},p}^{(1)}$
?
?
$d_{\mathcal{GW},p}^{(2)}$
$d_{\mathcal{W},p}$

The plan

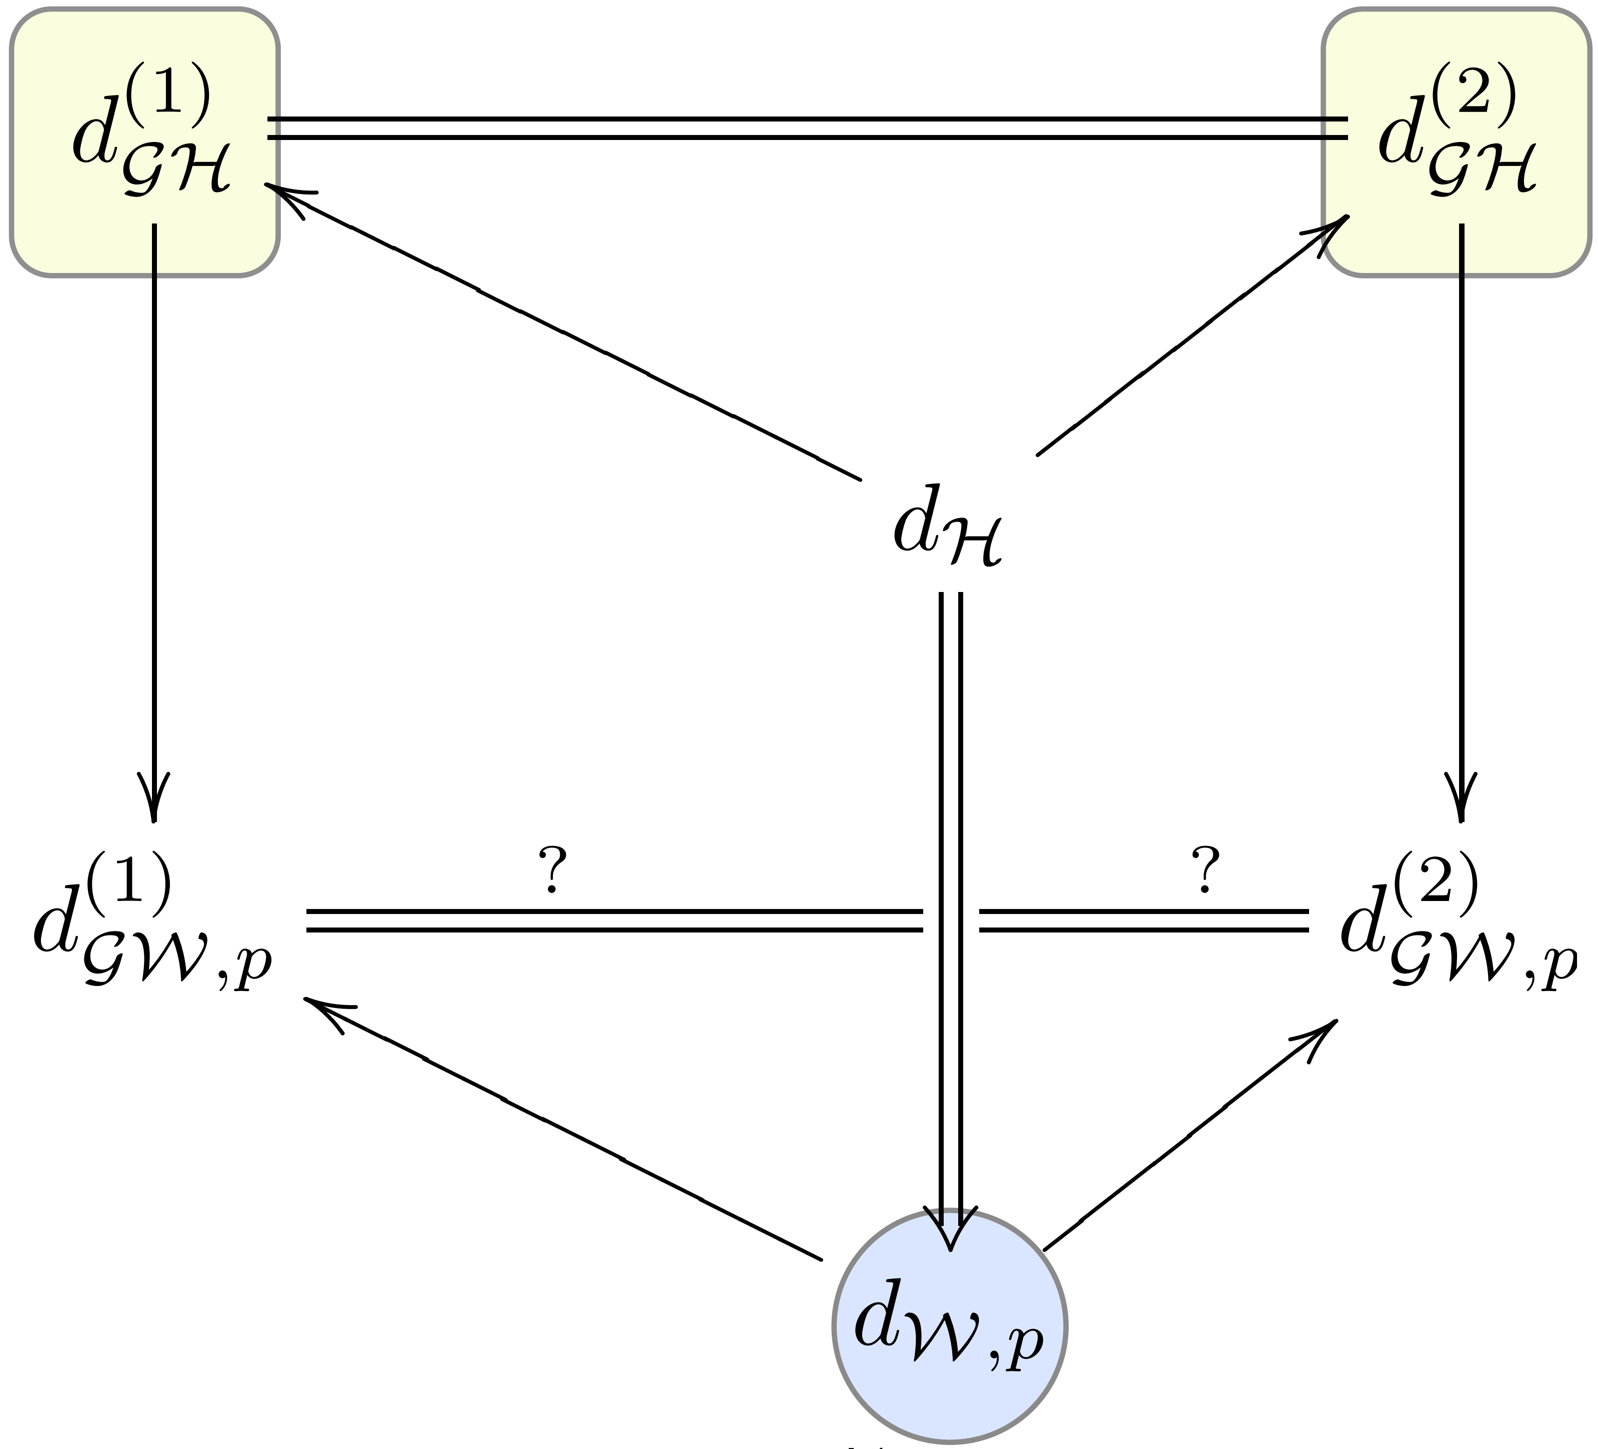

$d^{(1)}_{\mathcal{GH}}$
$d^{(2)}_{\mathcal{GH}}$
$d_{\mathcal{H}}$
?
?
$d^{(1)}_{\mathcal{GW},p}$
$d^{(2)}_{\mathcal{GW},p}$
$d_{\mathcal{W},p}$

The plan

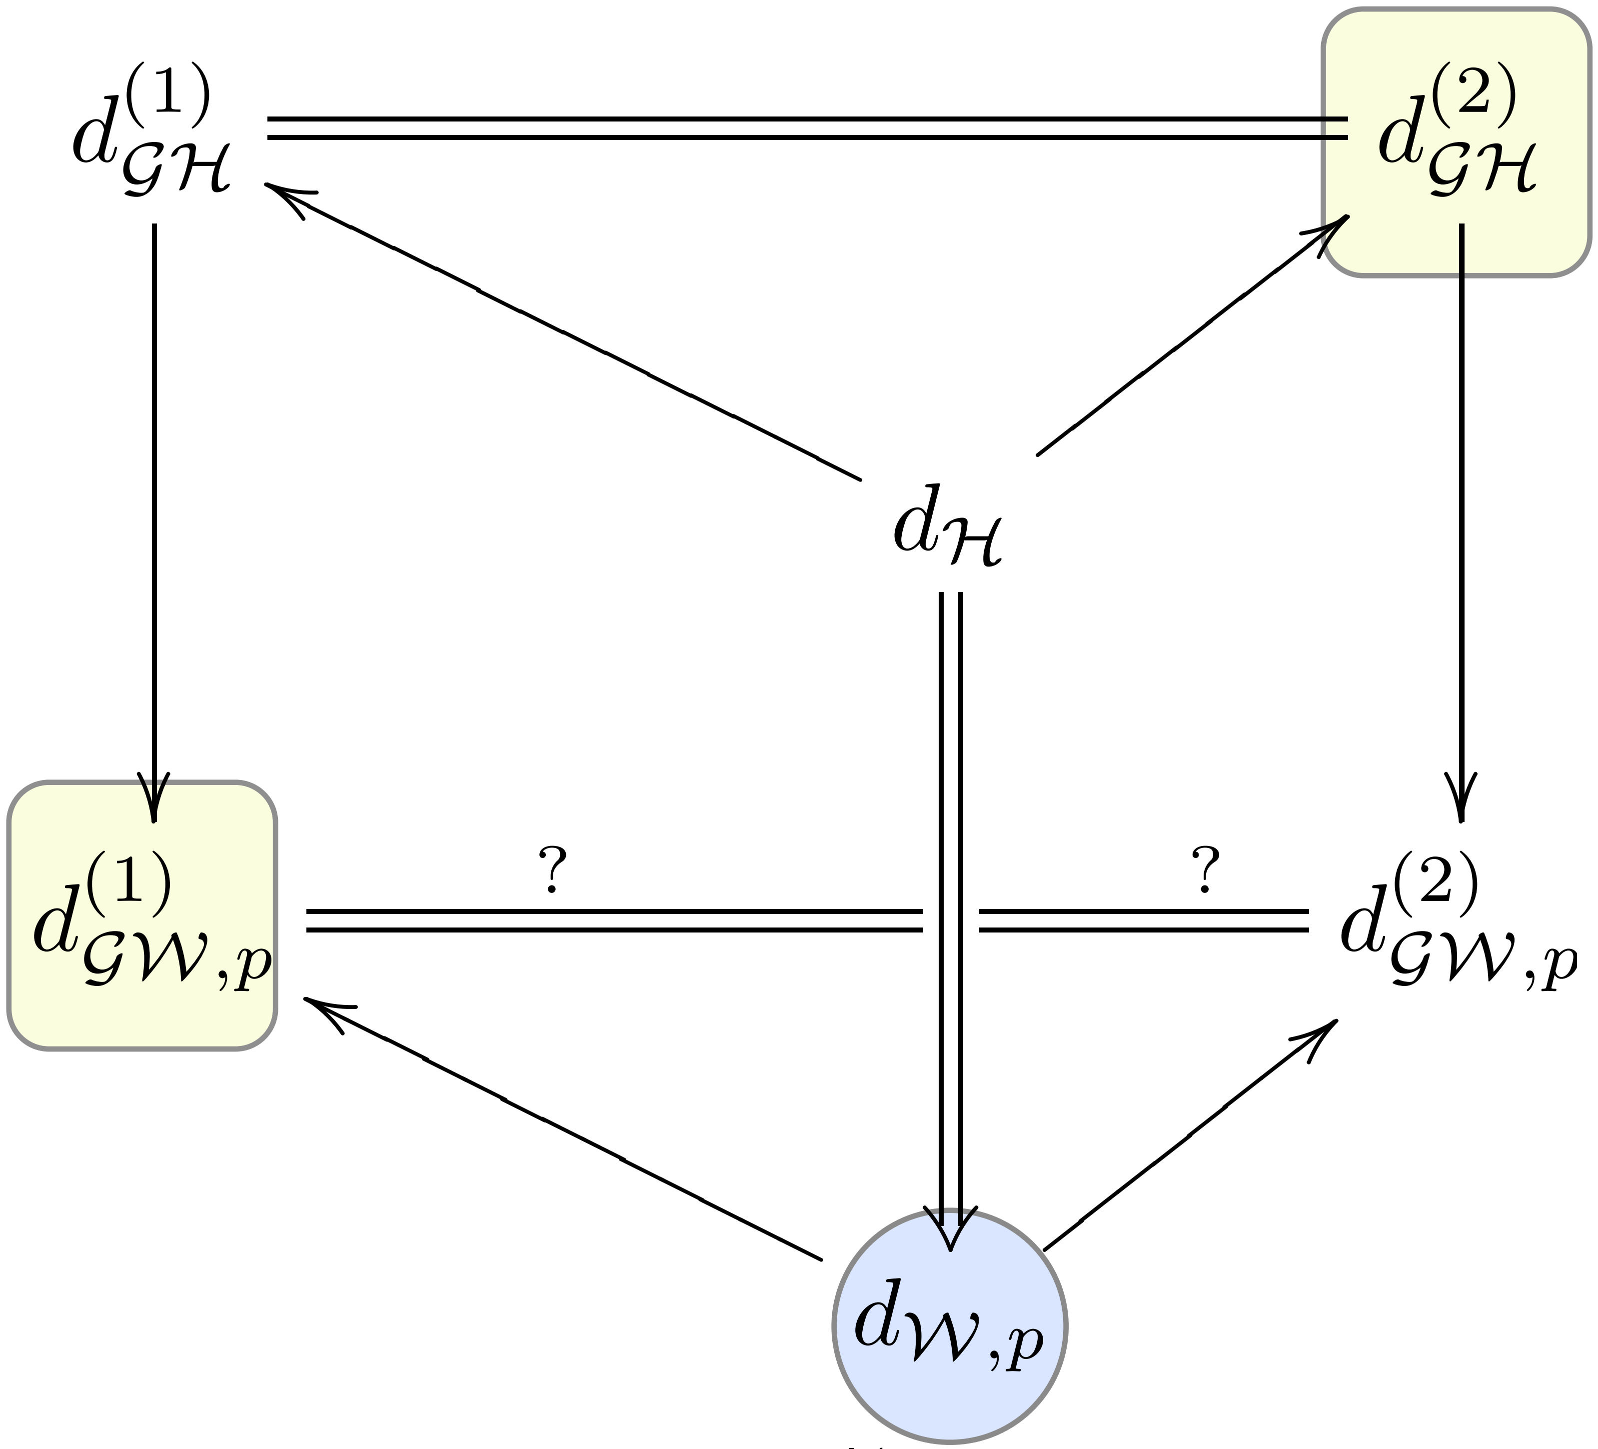
$d^{(1)}_{\mathcal{GH}}$
$d^{(2)}_{\mathcal{GH}}$
$d_{\mathcal{H}}$
$d^{(1)}_{\mathcal{GW},p}$
?
?
$d^{(2)}_{\mathcal{GW},p}$
$d_{\mathcal{W},p}$

The plan

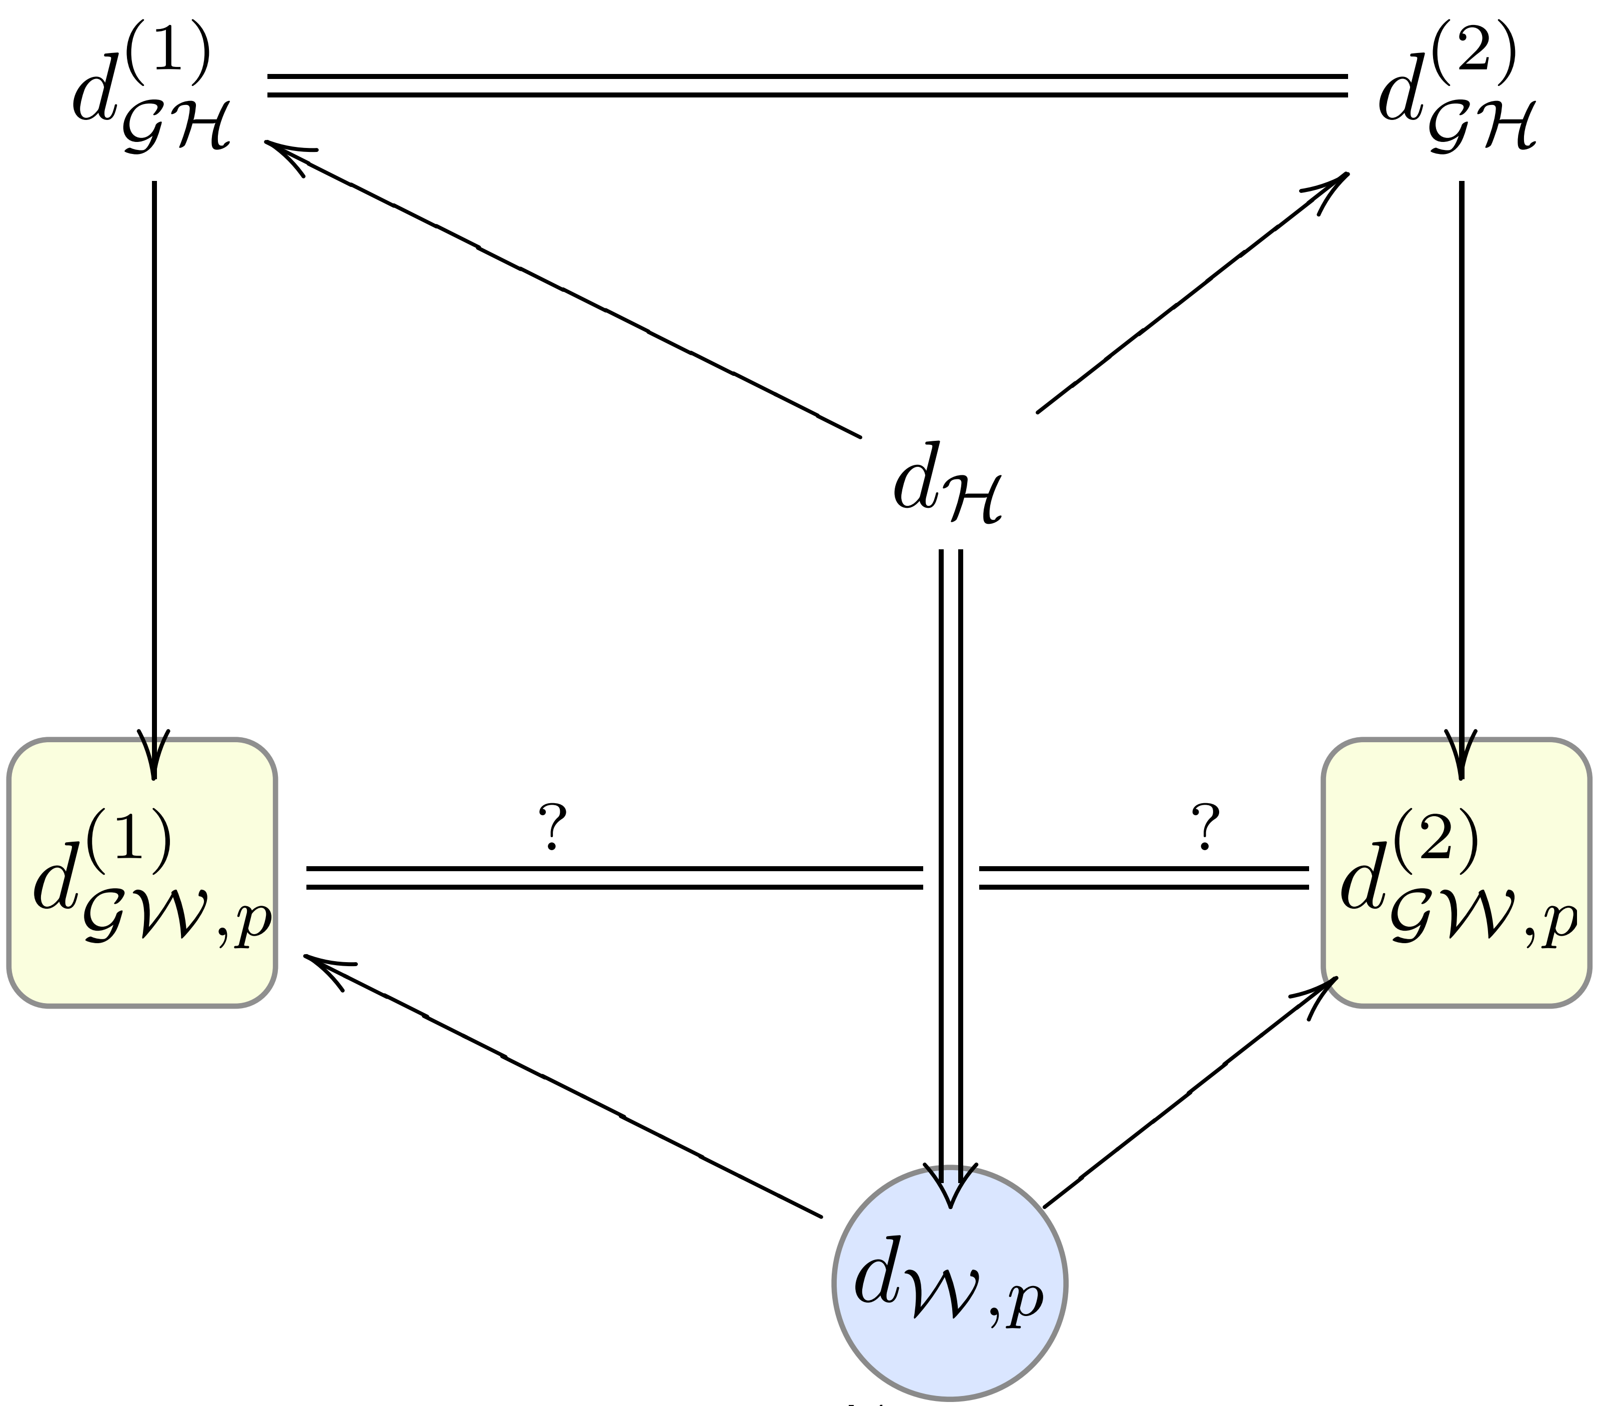

$d^{(1)}_{\mathcal{GH}}$
$d^{(2)}_{\mathcal{GH}}$
$d_{\mathcal{H}}$
$d^{(1)}_{\mathcal{GW},p}$
?
?
$d^{(2)}_{\mathcal{GW},p}$
$d_{\mathcal{W},p}$

The plan

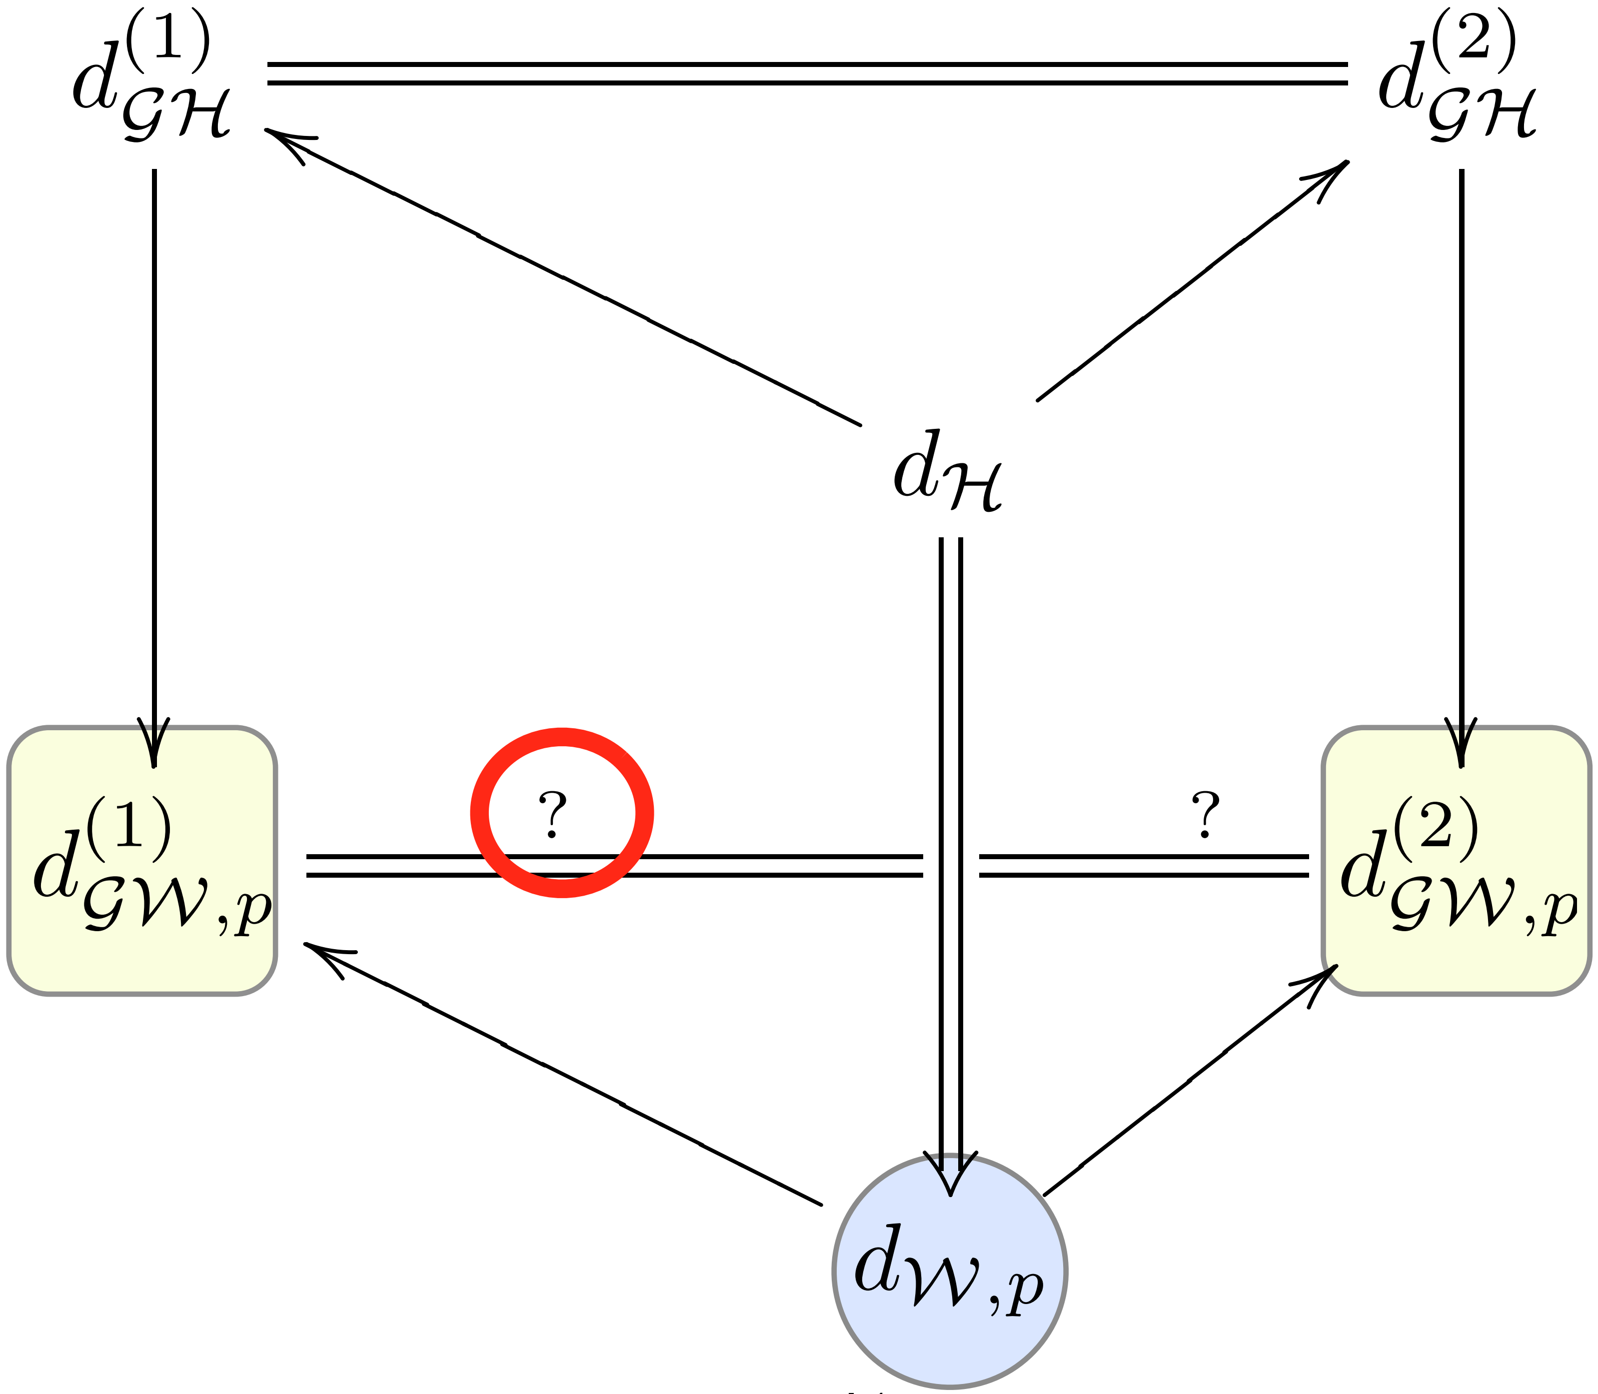
$d^{(1)}_{\mathcal{GH}}$
$d^{(2)}_{\mathcal{GH}}$
$d_{\mathcal{H}}$
$d^{(1)}_{\mathcal{GW},p}$
?
?
$d^{(2)}_{\mathcal{GW},p}$
$d_{\mathcal{W},p}$

GH distance

GH: definition

$$d_{\mathcal{GH}}(X,Y) = \inf_{Z,f,g} d_{\mathcal{H}}^{Z}(f(X), g(Y))$$

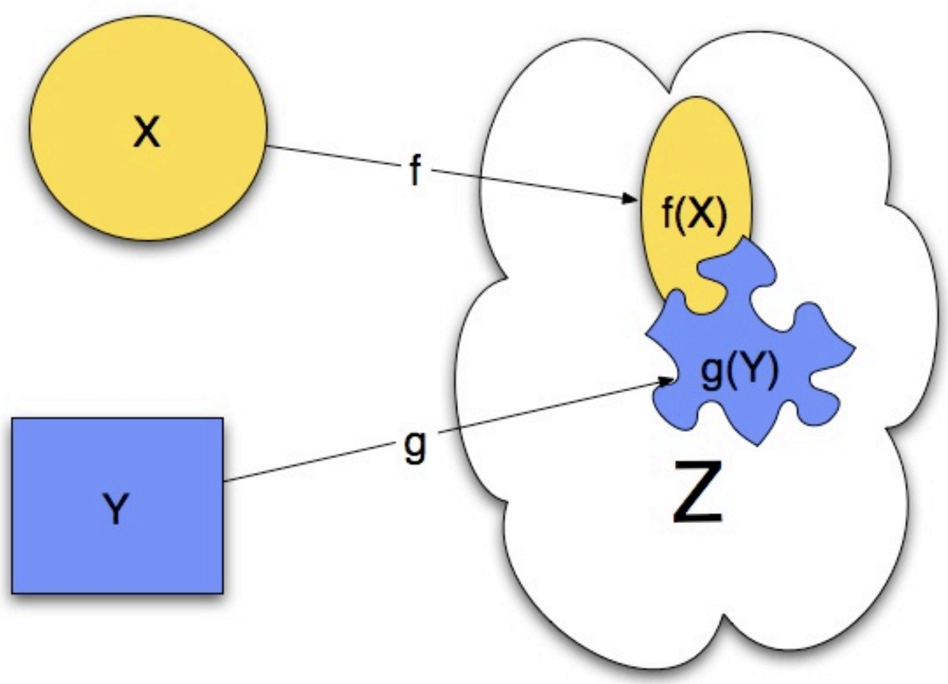

It is enough to consider $Z = X \sqcup Y$ and then we obtain

$$d_{\mathcal{GH}}(X,Y) = \inf_{d} d_{\mathcal{H}}^{(Z,d)}(X,Y)$$

Recall:

Proposition

Let (X,d) be a compact metric space and $A, B \subset X$ be compact. Then

$$d_{\mathcal{H}}(A,B) = \inf_{R \in \mathcal{R}(A,B)} \|d\|_{L^{\infty}(R)}$$

correspondences and GH distance

The GH distance between (X, d_X) and (Y, d_Y) admits the following expression:

$$d_{\mathcal{GH}}^{(1)}(X,Y) = \inf_{d \in \mathcal{D}(d_X, d_Y)} \inf_{R \in \mathcal{R}(X,Y)} \|d\|_{L^\infty(R)}$$

where $\mathcal{D}(d_X, d_Y)$ is a metric on $X \sqcup Y$ that reduces to d_X and d_Y on $X \times X$ and $Y \times Y$, respectively.

$$\begin{array}{cc} & \begin{array}{cc} X & Y \end{array} \\ \begin{array}{c} X \\ Y \end{array} & \begin{pmatrix} d_X & \mathbf{D} \\ \mathbf{D}^T & d_Y \end{pmatrix} = d \end{array}$$

In other words: you need to **glue** X and Y in an optimal way. Note that $\mathbf{D}$ consists of $n_X \times n_Y$ positive reals that must satisfy $\sim n_X \cdot C_2^{n_Y} + n_Y \cdot C_2^{n_X}$ linear constraints.

Another expression for the GH distance

For compact spaces (X, d_X) and (Y, d_Y) let

$$d_{\mathcal{GH}}^{(2)}(X,Y) = \frac{1}{2}\inf_R \max_{(x,y),(x',y')\in R} |d_X(x,x') - d_Y(y,y')|$$

We write, compactly,

$$d_{\mathcal{GH}}^{(2)}(X,Y) = \frac{1}{2}\inf_R \|d_X - d_Y\|_{L^\infty(R\times R)}$$

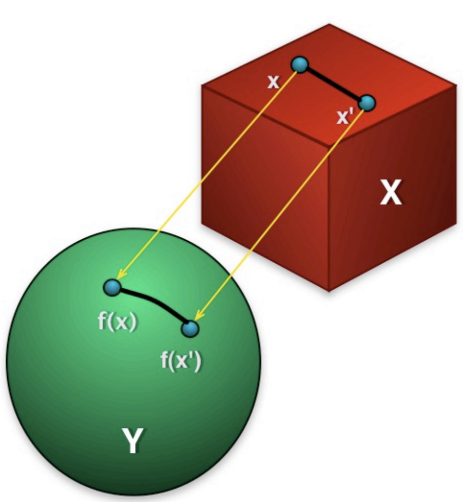

Equivalence thm:

Theorem [Kalton-Ostrovskii]
For all X, Y compact,

$$\begin{array}{ccc} d_{\mathcal{GH}}^{(1)} & = & d_{\mathcal{GH}}^{(2)} \\ \| & & \| \\ \inf_{d,R} \|d\|_{L^\infty(R)} & & \frac{1}{2} \inf_R \|d_X - d_Y\|_{L^\infty(R\times R)} \end{array}$$

Relaxing the notion of correspondence

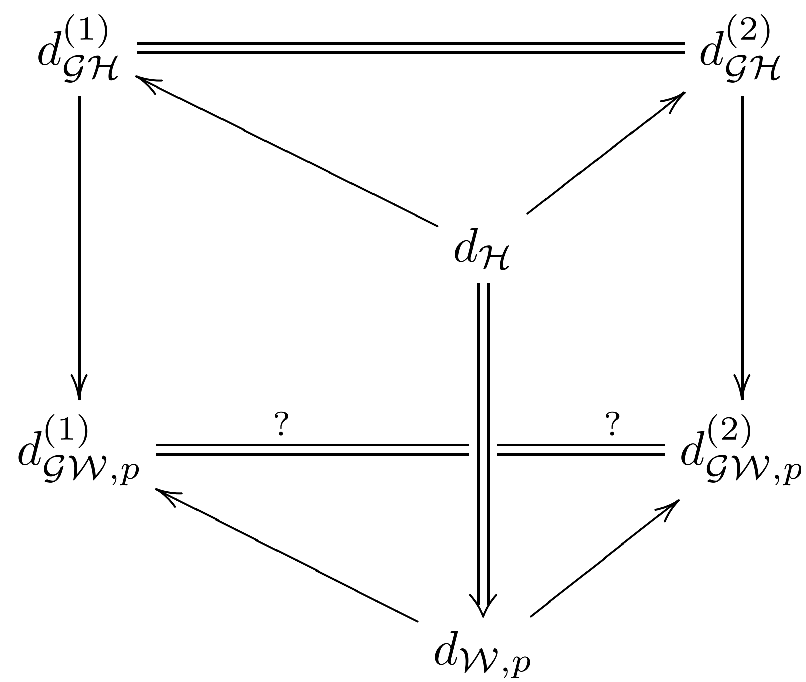

$$\begin{array}{ccc} d_{\mathcal{GH}}^{(1)} & = & d_{\mathcal{GH}}^{(2)} \\ \| & & \| \\ \inf_{d,R} \|d\|_{L^\infty(R)} & & \frac{1}{2} \inf_{R} \|d_X - d_Y\|_{L^\infty(R \times R)} \end{array}$$

$$
\begin{array}{ccc}
d^{(1)}_{\mathcal{GH}} & = & d^{(2)}_{\mathcal{GH}} \\
\Vert & & \Vert \\
\inf_{d,R} \|d\|_{L^{\infty}(R)} & & \frac{1}{2}\inf_{R} \|d_X - d_Y\|_{L^{\infty}(R\times R)} \\
\downarrow & & \downarrow \\
\inf_{d,\mu} \|d\|_{L^{p}(\mu)} & & \frac{1}{2}\inf_{\mu} \|d_X - d_Y\|_{L^{p}(\mu\otimes\mu)} \\
\Vert & & \Vert \\
d^{(1)}_{\mathcal{GW},p} & & d^{(2)}_{\mathcal{GW},p}
\end{array}
$$

Now, one works with **mm-spaces**: triples (X, d, ν) where (X, d) is a compact metric space and ν is a Borel probability measure. Two mm-spaces are *isomorphic* iff there exists isometry $\Phi : X \to Y$ s.t. $\mu_X(\Phi^{-1}(B)) = \mu_Y(B)$ for all measurable $B \subset Y$.

Now, one works with *mm-spaces*: triples (X, d, ν) where (X, d) is a compact metric space and ν is a Borel probability measure. Two mm-spaces are *isomorphic* iff there exists isometry $\Phi : X \to Y$ s.t. $\mu_X(\Phi^{-1}(B)) = \mu_Y(B)$ for all measurable $B \subset Y$.

The first option, proposed and analyzed by K.L Sturm [**St06**], reads

$$d^{(1)}_{\mathcal{GW},p}(X,Y) = \inf_{d \in \mathcal{D}(d_X, d_Y)} \inf_{\mu \in \mathcal{M}(\mu_X, \mu_Y)} \left(\sum_{x,y} d^p(x,y) \mu_{x,y} \right)^{1/p}$$

Now, one works with *mm-spaces*: triples (X, d, ν) where (X, d) is a compact metric space and ν is a Borel probability measure. Two mm-spaces are *isomorphic* iff there exists isometry $\Phi : X \to Y$ s.t. $\mu_X(\Phi^{-1}(B)) = \mu_Y(B)$ for all measurable $B \subset Y$.

The first option, proposed and analyzed by K.L Sturm [St06], reads

$$d^{(1)}_{\mathcal{GW},p}(X,Y) = \inf_{d \in \mathcal{D}(d_X,d_Y)} \inf_{\mu \in \mathcal{M}(\mu_X,\mu_Y)} \left(\sum_{x,y} d^p(x,y)\mu_{x,y} \right)^{1/p}$$

The second option reads **[M07]**

$$d^{(2)}_{\mathcal{GW},p}(X,Y) = \inf_{\mu \in \mathcal{M}(\mu_X,\mu_Y)} \left(\sum_{x,y} \sum_{x',y'} |d_X(x,x') - d_Y(y,y')|^p \mu_{x,y}\mu_{x',y'} \right)^{1/p}$$

The **first** option,

$$d^{(1)}_{\mathcal{GW},p}(X,Y) = \inf_{d\in\mathcal{D}(d_X,d_Y)} \inf_{\mu\in\mathcal{M}(\mu_X,\mu_Y)} \left(\sum_{x,y} d^p(x,y)\mu_{x,y}\right)^{1/p}$$

requires $\mathbf{2}(\mathbf{n_X}\times\mathbf{n_Y})$ variables and $\mathbf{n_X}+\mathbf{n_Y}$ plus $\sim \mathbf{n_Y}\cdot\mathbf{C_2^{n_X}}+\mathbf{n_X}\cdot\mathbf{C_2^{n_Y}}$ linear constraints. When $p=1$ it yields a *bilinear* optimization problem.

Our **second** option,

$$d^{(2)}_{\mathcal{GW},p}(X,Y) = \inf_{\mu\in\mathcal{M}(\mu_X,\mu_Y)} \left(\sum_{x,y}\sum_{x',y'} |d_X(x,x')-d_Y(y,y')|^p \mu_{x,y}\mu_{x',y'}\right)^{1/p}$$

requires $\mathbf{n_X}\times\mathbf{n_Y}$ variables and $\mathbf{n_X}+\mathbf{n_Y}$ linear constraints. It is a *quadratic* (generally non-convex :-( ) optimization problem (with linear and bound constraints) for all p.

Then one would argue for using $d^{(2)}_{\mathcal{GW},p}$.

Numerical Implementation

The numerical implementation of the second option leads to solving a **QOP** with linear constraints:

$$\min_U \tfrac{1}{2} U^T \mathbf{\Gamma} U$$
$$\text{s.t. } U_{ij} \in [0,1],\ U\mathbf{A} = \mathbf{b}$$

where $U \in \mathbb{R}^{n_X \times n_Y}$ is the *unrolled* version of μ, $\mathbf{\Gamma} \in \mathbb{R}^{n_X \times n_Y \times n_X \times n_Y}$ is the unrolled version of $\Gamma_{X,Y}$ and $\mathbf{A}$ and $\mathbf{b}$ encode the linear constrains $\mu \in \mathcal{M}(\mu_X, \mu_Y)$.

This can be approached for example via gradient descent. The QOP is non-convex in general!

Initialization is done via solving one of the several *lower bounds* (discussed ahead). All these lower bounds lead to solving **LOP**s.

For details see [**M07**].

Can GW (1) be equal to GW (2)?

- Using the same proof as in the Kalton-Ostrovskii Thm., one can prove that
$$d^{(1)}_{\mathcal{GW},\infty} = d^{(2)}_{\mathcal{GW},\infty}.$$
- Also, it is obvious that for all $p \geq 1$
$$d^{(1)}_{\mathcal{GW},p} \geq d^{(2)}_{\mathcal{GW},p}.$$
- But the equality does not hold in general. One counterexample is as follows: take $X = (\Delta_{n-1}, ((d_{ij} = 1)), (\nu_i = 1/n))$ and $Y = (\{q\}, ((0)), (1))$. Then, for $p \in [1, \infty)$
$$d^{(1)}_{\mathcal{GW},p}(X,Y) = \frac{1}{2} > \frac{1}{2}\left(\frac{n-1}{n}\right)^{1/p} = d^{(2)}_{\mathcal{GW},p}(X,Y)$$
- Furthermore, these two (tentative) distances are **not equivalent**!! This forces us to analyze them separately. The delicate step is proving that $\text{dist}(X,Y) = 0$ implies $X \simeq Y$.
- K. T. Sturm has analyzed GW (1).

$\{\Delta_n\}_{n=1}^{\infty}$

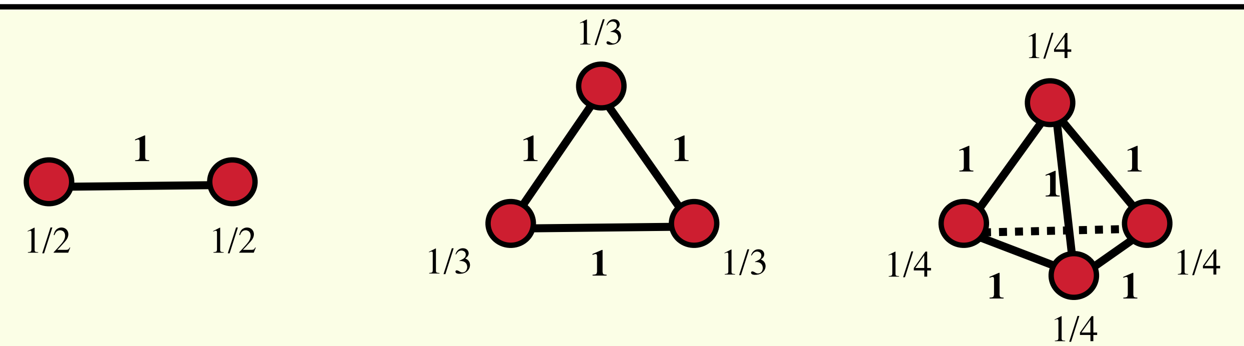

Properties of $d_{\mathcal{GW},p}^{(2)}$

1. Let X, Y and Z mm-spaces then

$$d_{\mathcal{GW},p}(X,Y) \leq d_{\mathcal{GW},p}(X,Z) + d_{\mathcal{GW},p}(Y,Z).$$

2. If $d_{\mathcal{GW},p}(X,Y) = 0$ then X and Y are isomorphic.

3. Let $\mathbb{X}_n = \{x_1, \ldots, x_n\} \subset X$ be a subset of the mm-space (X, d, ν). Endow $\mathbb{X}_n$ with the metric d and a prob. measure ν_n, then

$$d_{\mathcal{GW},p}(X, \mathbb{X}_n) \leq d_{\mathcal{W},p}(\nu, \nu_n).$$

The parameter p is not superfluous

The simplest lower bound one has is based on the triangle inequality plus

$$2 \cdot d^{(2)}_{\mathcal{GW},p}(X,\{q\}) = \left(\int_{X\times X} d_X(x,x')\,\nu(dx)\nu(dx')\right)^{1/p} := \mathbf{diam}_p(X)$$

That is

$$d^{(2)}_{\mathcal{GW},p}(X,Y) \geq \frac{1}{2}|\mathbf{diam}_p(X) - \mathbf{diam}_p(Y)|$$

For example, when $X = S^n$ (spheres with uniform measure and usual intrinsic metric):

- $p = \infty$ gives $\mathbf{diam}_\infty(S^n) = \pi$ for all $n \in \mathbb{N}$
- $p = 1$ gives $\mathbf{diam}_1(S^n) = \pi/2$ for all $n \in \mathbb{N}$
- $p = 2$ gives $\mathbf{diam}_2(S^1) = \pi/\sqrt{3}$ and $\mathbf{diam}_2(S^2) = \sqrt{\pi^2/2 - 2}$

Upper and Lower bounds Let (X, d, ν) be an mm-space.

- **Shape Distributions [Osada-et-al]**: construct histogram of interpoint distances, $F_X : \mathbb{R} \to [0,1]$ given by

$$t \mapsto \nu \otimes \nu \left(\{(x, x') |\, d(x, x') \leq t\} \right)$$

- **Shape Contexts [SC]**: at each $x \in X$, construct histogram of $d(x, \cdot)$, $C_X : X \times \mathbb{R} \to [0,1]$ given by

$$(x, t) \mapsto \nu \left(\{x' |\, d(x, x') \leq t\} \right)$$

- **Hamza-Krim [HK]**: at each $x \in X$ compute mean distance to rest of points, $H_X : X \to \mathbb{R}$

$$x \mapsto \left(\int_X d^p(x, x') \nu(dx') \right)^{1/p}$$

- **Wasserstein under Euclidean isometries**: consider $X, Y \subset \mathbb{R}^d$ and compute

$$d^{iso}_{\mathcal{W},p}(X, Y) = \inf_T d_{\mathcal{W},p}(X, T(Y))$$

- **Gromov-Hausdorff distance [MS04],[MS05]**

Upper and Lower bounds Let (X, d, ν) be an mm-space.

- **Shape Distributions [Osada-et-al]**: construct histogram of interpoint distances, $F_X : \mathbb{R} \to [0,1]$ given by

$$t \mapsto \nu \otimes \nu \left(\{(x, x') |\, d(x, x') \le t\} \right)$$

- **Shape Contexts [SC]**: at each $x \in X$, construct histogram of $d(x, \cdot)$, $C_X : X \times \mathbb{R} \to [0,1]$ given by

$$(x, t) \mapsto \nu \left(\{x' |\, d(x, x') \le t\} \right)$$

- **Hamza-Krim [HK]**: at each $x \in X$ compute mean distance to rest of points, $H_X : X \to \mathbb{R}$

$$x \mapsto \left(\int_X d^p(x, x') \nu(dx') \right)^{1/p}$$

- **Wasserstein under Euclidean isometries**: consider $X, Y \subset \mathbb{R}^d$ and compute

$$d_{\mathcal{W},p}^{iso}(X, Y) = \inf_T d_{\mathcal{W},p}(X, T(Y))$$

- **Gromov-Hausdorff distance [MS04],[MS05]**

Upper and Lower bounds Let (X, d, ν) be an mm-space.

- **Shape Distributions [Osada-et-al]**: construct histogram of interpoint distances, $F_X : \mathbb{R} \to [0,1]$ given by

$$t \mapsto \nu \otimes \nu \left(\{(x, x') | \, d(x, x') \leq t\}\right)$$

- **Shape Contexts [SC]**: at each $x \in X$, construct histogram of $d(x, \cdot)$, $C_X : X \times \mathbb{R} \to [0,1]$ given by

$$(x, t) \mapsto \nu \left(\{x' | \, d(x, x') \leq t\}\right)$$

- **Hamza-Krim [HK]**: at each $x \in X$ compute mean distance to rest of points, $H_X : X \to \mathbb{R}$

$$x \mapsto \left(\int_X d^p(x, x') \nu(dx')\right)^{1/p}$$

- **Wasserstein under Euclidean isometries**: consider $X, Y \subset \mathbb{R}^d$ and compute

$$d^{iso}_{\mathcal{W},p}(X, Y) = \inf_T d_{\mathcal{W},p}(X, T(Y))$$

- **Gromov-Hausdorff distance [MS04],[MS05]**

Upper and Lower bounds Let (X, d, ν) be an mm-space.

- **Shape Distributions [Osada-et-al]**: construct histogram of interpoint distances, $F_X : \mathbb{R} \to [0,1]$ given by

$$t \mapsto \nu \otimes \nu \left(\{(x, x') |\, d(x, x') \leq t\}\right)$$

- **Shape Contexts [SC]**: at each $x \in X$, construct histogram of $d(x, \cdot)$, $C_X : X \times \mathbb{R} \to [0,1]$ given by

$$(x, t) \mapsto \nu \left(\{x' |\, d(x, x') \leq t\}\right)$$

- **Hamza-Krim [HK]**: at each $x \in X$ compute mean distance to rest of points, $H_X : X \to \mathbb{R}$

$$x \mapsto \left(\int_X d^p(x, x') \nu(dx') \right)^{1/p}$$

- **Wasserstein under Euclidean isometries**: consider $X, Y \subset \mathbb{R}^d$ and compute

$$d^{iso}_{\mathcal{W},p}(X, Y) = \inf_T d_{\mathcal{W},p}(X, T(Y))$$

- **Gromov-Hausdorff distance [MS04],[MS05]**

Some Experiments

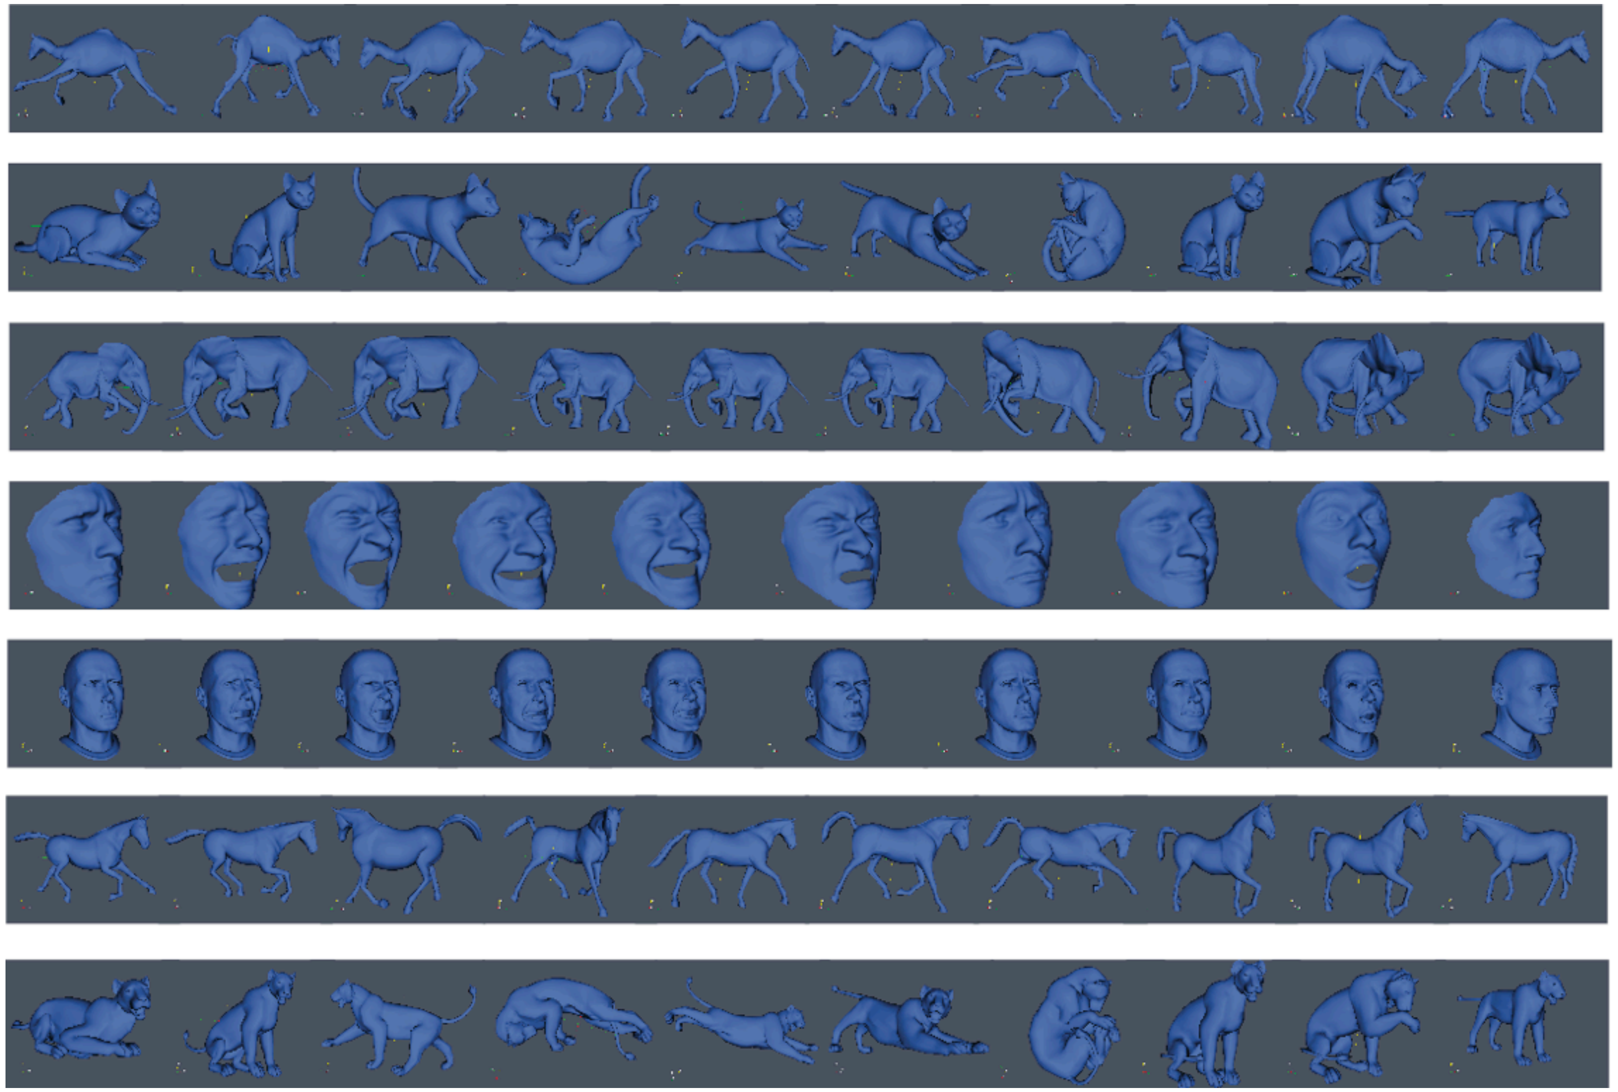

Some experimentation: ~ 70 models in 7 classes. Classification using 1-nn: $P_e \sim 2\%$. Hamza-Krim gave $\sim 15\%$ on same db with all same parameters etc.

Discussion

- Implementation is easy: Gradient descent or alternate opt.
- Solving lower bounds yields a seed for the gradient descent. These lower bounds are compatible with the metric in the sense that a layered recognition system is possible: given two shapes, (1) solve for a LB (this gives you a μ), if value small enough, then (2) solve for GW using the μ as seed for your favorite iterative algorithm.
- Easy extension to partial matching.
- Interest in relating GH/GW ideas to other methods in the literature. Interrelating methods is important also for applications: when confronted with N methods, how do they compare to each other? which one is better for the situation at hand?
- Latest developments:
  - Partial matching [**M08-partial**].
  - Euclidean case [**M08-euclidean**].
  - Persistent Topology based methods (Frosini et al., Carlsson et al.)
- No difference between continuous and discrete. Probability measures take care of the 'transition'.

Bibliography

[**BBI**] Burago, Burago and Ivanov. A course on Metric Geometry. AMS, 2001

[**EK**] A. Elad (Elbaz) and R. Kimmel. On bending invariant signatures for surfaces, IEEE Trans. on PAMI, 25(10):1285-1295, 2003.

[**BBK**] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. Proc. Natl. Acad. Sci. USA, 103(5):1168–1172 (electronic), 2006.

[**HK**] A. Ben Hamza, Hamid Krim: Probabilistic shape descriptor for triangulated surfaces. ICIP (1) 2005: 1041-1044

[**Osada-et-al**] Robert Osada, Thomas A. Funkhouser, Bernard Chazelle, David P. Dobkin: Matching 3D Models with Shape Distributions. Shape Modeling International 2001: 154-166

[**SC**] S. Belongie and J. Malik (2000). "Matching with Shape Contexts". IEEE Workshop on Contentbased Access of Image and Video Libraries (CBAIVL-2000).

[**Goodrich**] M. Goodrich, J. Mitchell, and M. Orletsky. Approximate geometric pattern matching under rigid motions. IEEE Transactions on Pattern Analysis and Machine Intelligence, 21(4):371–376, 1999.

[**M08-partial**] F.Memoli. Lp Gromov-Hausdorff distances for **partial** shape matching, preprint.

[**M08-euclidean**]F. Mémoli. Gromov-Hausdorff distances in Euclidean spaces. NORDIA-CVPR-2008.

[**M07**] F.Memoli. On the use of gromov-hausdorff distances for shape comparison. In Proceedings of PBG 2007, Prague, Czech Republic, 2007.

[**MS04**] F. Memoli and G. Sapiro. Comparing point clouds. In SGP '04: Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing, pages 32–40, New York, NY, USA, 2004. ACM.

[**MS05**] F. Memoli and G. Sapiro. A theoretical and computational framework for isometry invariant recognition of point cloud data. Found. Comput. Math., 5(3):313–347, 2005.

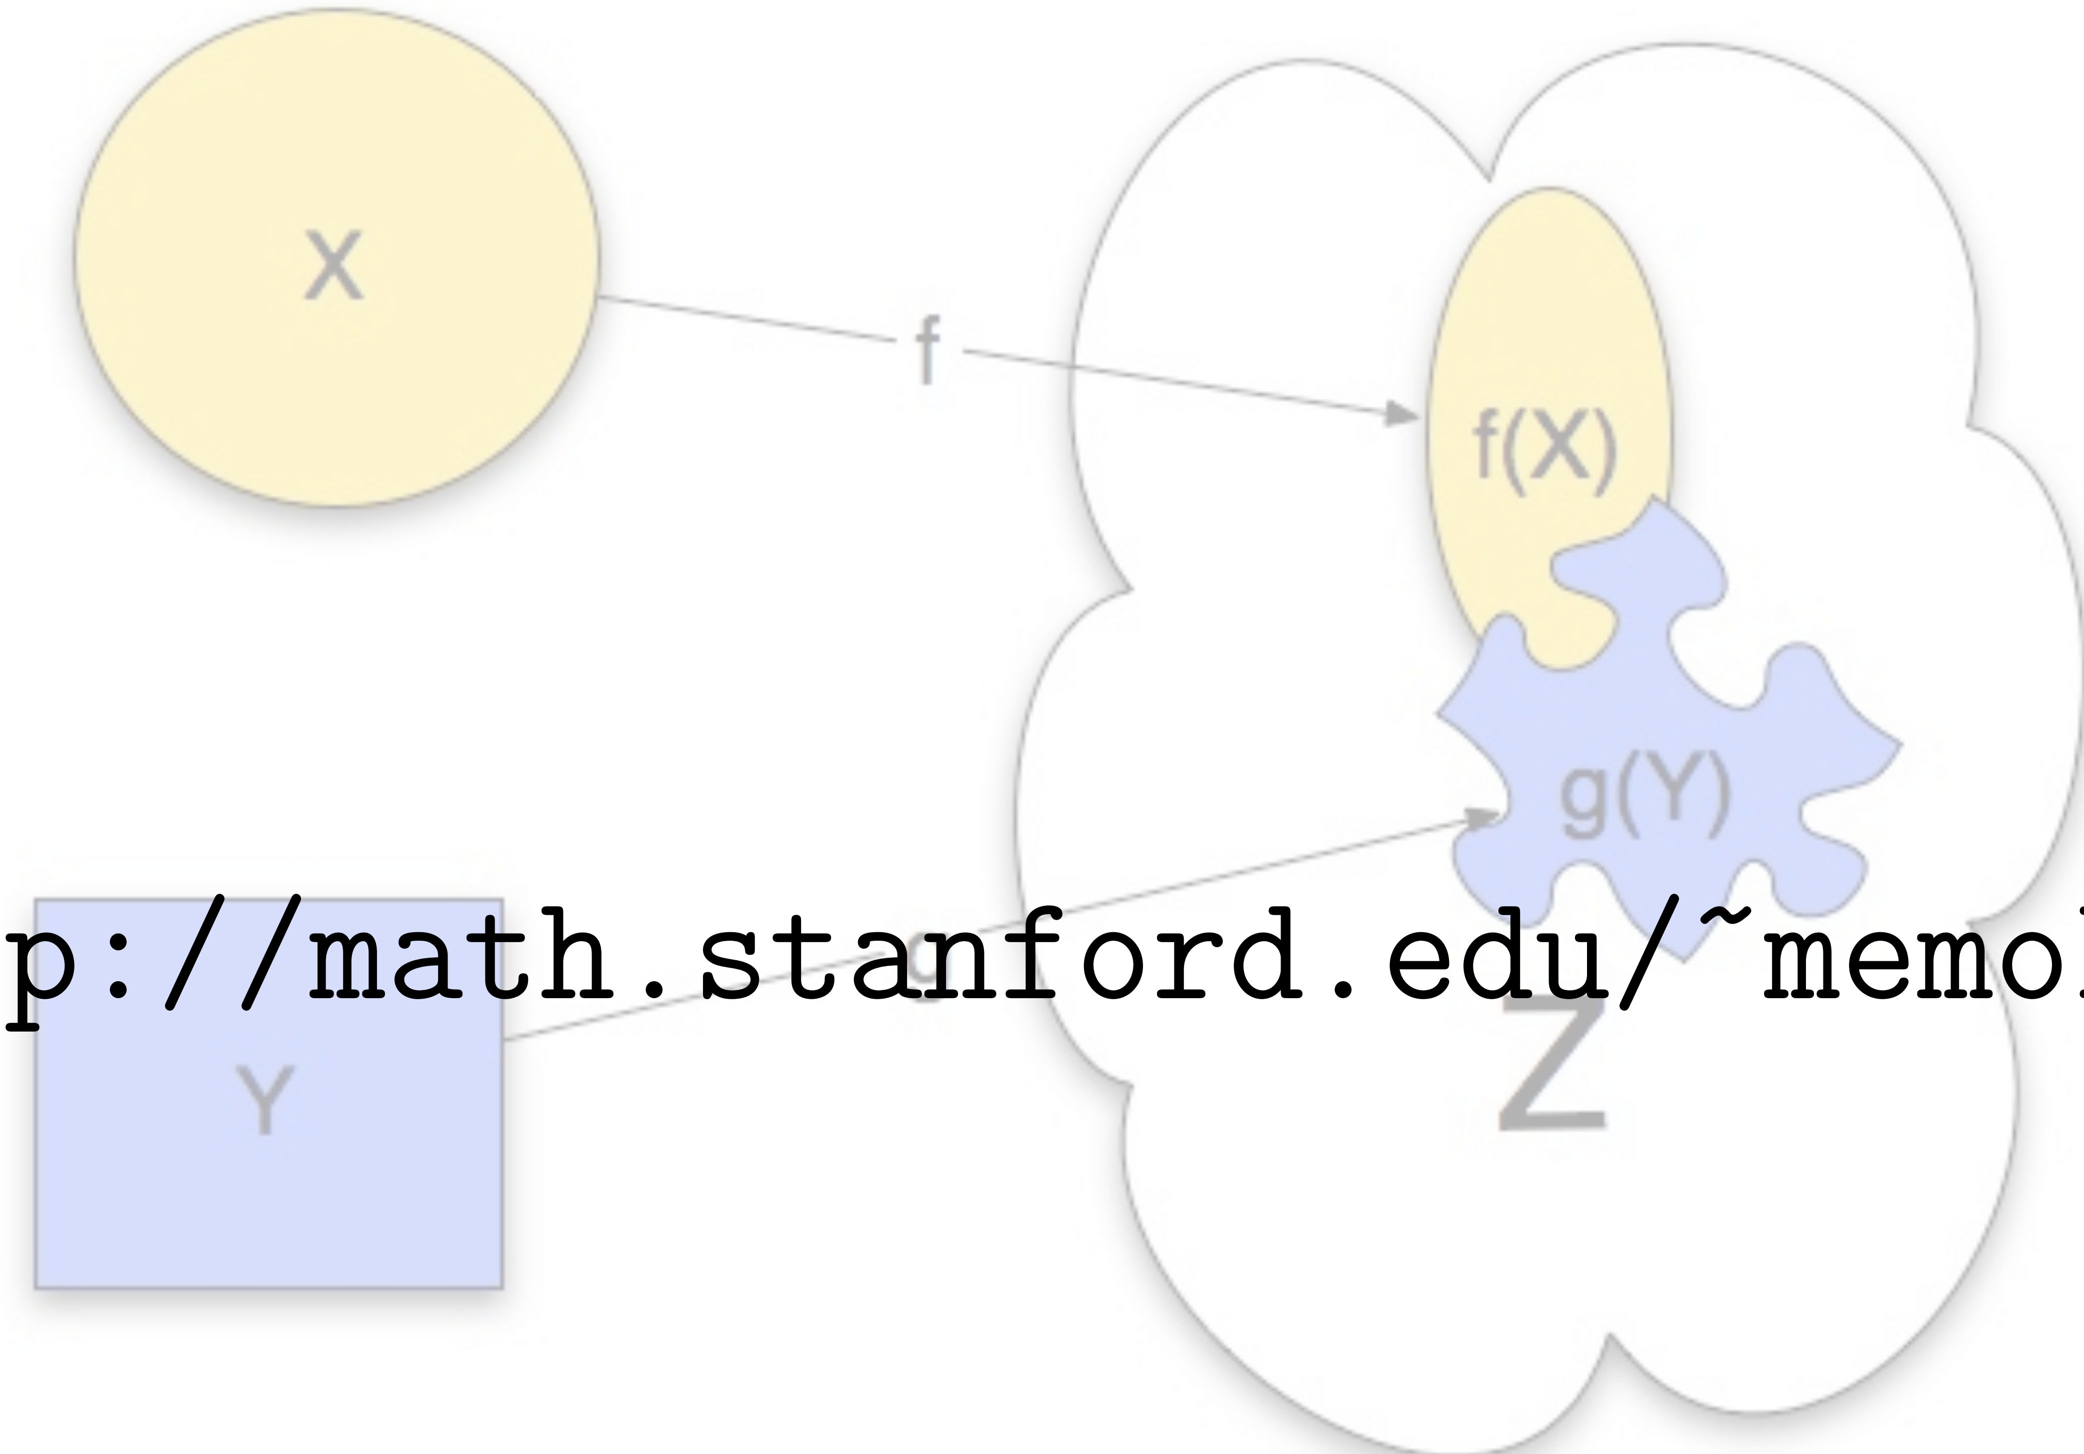

http://math.stanford.edu/~memoli