Lifting Couplings in Wasserstein Spaces

Paolo Perrone

University of Oxford, Computer Science Department, Oxford, UK

This paper makes mathematically precise the idea that conditional probabilities are analogous to *path liftings* in geometry.

The idea of lifting is modeled in terms of the category-theoretic concept of a *lens*, which can be interpreted as a consistent choice of arrow liftings. The category we study is the one of probability measures over a given standard Borel space, with morphisms given by the couplings, or transport plans.

The geometrical picture is even more apparent once we equip the arrows of the category with *weights*, which one can interpret as "lengths" or "costs", forming a so-called *weighted category*, which unifies several concepts of category theory and metric geometry. Indeed, we show that the weighted version of a lens is tightly connected to the notion of *submetry* in geometry.

Every weighted category gives rise to a pseudo-quasimetric space via optimization over the arrows. In particular, Wasserstein spaces can be obtained from the weighted categories of probability measures and their couplings, with the weight of a coupling given by its cost. In this case, conditional distributions allow one to form weighted lenses, which one can interpret as "lifting transport plans, while preserving their cost".

1 Introduction

This paper explores a connection between conditional probabilities and liftings in geometry and topology.

Projections and liftings. Forming a marginal of a probability measure is a destructive operation. Given a probability measure on a product space $X \times Y$, its marginal on X is often called a *projection*, borrowing from the language of geometry. The idea is that the marginal distribution captures only the variability in X, while the original joint measure may exhibit also variability in Y (as well as possible correlations), in a way that's analogous to the projection onto the X-axis of a geometrical figure in the XY-plane. Similar intuitions can be had, more generally, whenever a probability measure is pushed forward along a noninjective measurable map, or when one restricts the measure to a coarser σ-algebra.

There are several situations, in geometry and related fields, where one has a projection $p : E \to B$, and one wants to "translate motions in B to corresponding motions in E". This idea is made mathematically precise in different ways, such as the following.

- In algebraic topology, a *fibration* $p : E \to B$ satisfies in particular a path-lifting property. Given a point $e \in E$ and a curve c in B starting at $p(e)$, we can lift c to a curve $\tilde{c}$ in E starting at e which is projected by p back to c.

- In metric geometry, a *submetry* $p : E \to B$ can be seen as "lifting distances". Consider a point $e \in E$, and a point $b \in B$ which has distance r from $p(e)$. Then we can lift b to a point of E which has distance r from e. (See also Section 3.)

Paolo Perrone: paolo.perrone@cs.ox.ac.uk, www.paoloperrone.org, 0000-0002-9123-9089

arXiv:2110.06591v4 [math.CT] 6 May 2025

- In differential topology, a *bundle* may be equipped with a *connection*, which can be interpreted as a way to lift vectors, or also paths (in an open neighborhood), in a way which preserves concatenation.

We are particularly interested in liftings that *preserve composition* (of curves, displacements, etc.). These can be effectively modeled via the category-theoretic notion of a *lens*, which we recall in Section 3. Intuitively, a lets consists of a functor $F : \mathcal{E} \to \mathcal{B}$ together with a "lifting" assignment which, in analogy to the points above, takes an object E of $\mathcal{E}$ and a morphism of $\mathcal{B}$ with domain $F(E)$, and gives a morphism of $\mathcal{E}$ with domain E, in such a way to preserve identity and composition. (See Section 3 for the details.)

The main focus of this work is the idea that by means of conditional distributions, one can construct a similar lifting to the ones described above. Namely, take a surjective, measurable function between measurable spaces $f : X \to Y$, and consider a probability measure p on X. Take the pushforward measure $f_\sharp p$ on Y, and consider a transport plan (a.k.a. coupling) from $f_\sharp p$ to another measure q on Y. Under what conditions can we lift the transport plan on Y to a transport plan on X? And when does this notion of lifting preserve the composition of transport plans?

We will show that these liftings exist whenever f is equipped with the structure of a *measurable lens* (Theorem 5.5). This in particular includes the case of product projections, which at the level of probability measures give marginals – and in that case, the lifting is given by forming a particular conditional product (Example 5.8).

Optimal transport and couplings. Consider two probability distributions p and q on a measurable space X. A *coupling* or *transport plan* between p and q is a joint probability distribution s on $X \times X$ such that its marginals are p and q. The name "transport plan" is motivated by the intuition that s can be considered a specified way of moving the mass of p to the configuration q. This is most visible in the discrete case. So, for an illustration, suppose that X is a finite set. A probability distribution p on X amounts to a function $p : X \to [0,1]$ such that
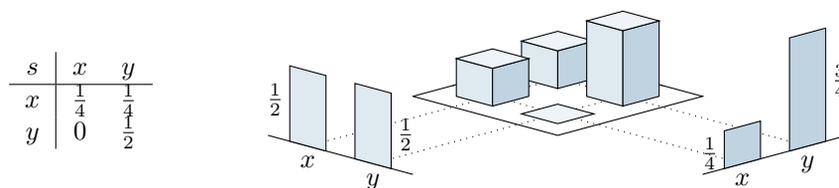
$$\sum_{x \in X} p(x) = 1.$$

So let p and q be distributions on X. A coupling s of p and q is then a function $s : X \times X \to [0,1]$ such that

$$\sum_{x' \in X} s(x, x') = p(x) \qquad \text{and} \qquad \sum_{x' \in X} s(x', x) = q(x)$$

for all $x \in X$. For example, if $X = \{x, y\}$, consider the probability measures p and q given as follows.
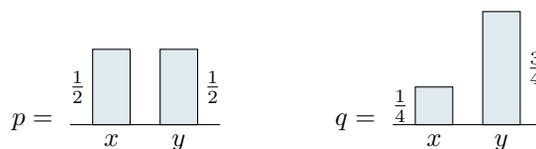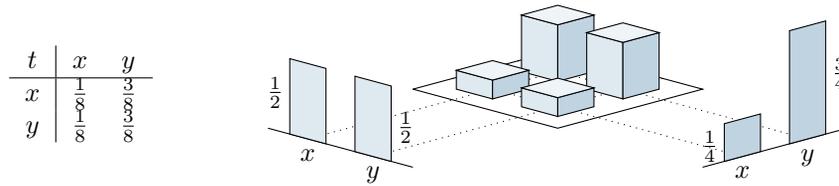

Here is a possible coupling between p and q.

We can interpret this coupling as a transport plan as follows, following the picture from left to right:

- Split the mass at x (which is $\frac{1}{2}$) into $\frac{1}{4}$ and $\frac{1}{4}$, keep $\frac{1}{4}$ at x, and move the rest to y;

- Leave all the mass at y (which is $\frac{1}{2}$) in place.

Here is another possible coupling.

t	x	y
x	$\frac{1}{8}$	$\frac{3}{8}$
y	$\frac{1}{8}$	$\frac{3}{8}$

Again, we can interpret this as follows:

- Split the mass at x (which is $\frac{1}{2}$) into $\frac{1}{8}$ and $\frac{3}{8}$, keep $\frac{1}{8}$ at x, and move $\frac{3}{8}$ to y;

- Split the mass at x (which is $\frac{1}{2}$) into $\frac{1}{8}$ and $\frac{3}{8}$, move $\frac{1}{8}$ to x, and keep $\frac{3}{8}$ at y.

We notice a few things.

(a) A transport plan from p to q can also be read backwards, as a transport plan from q to p.

(b) Transport plans can be composed: provided some technical conditions are satisfied, they form a category (see Section 4). Since a transport plan can be read in both directions, this category has a natural dagger structure.

(c) The first example, intuitively, "moves less mass" than the second example (it only moves $\frac{1}{4}$ from x to y, versus the second example that moves $\frac{3}{8}$ from y to x and $\frac{1}{8}$ from y back to x, for a total of $\frac{1}{2}$). In some sense, the first transport plan is "more optimal" or "less wasteful" than the second one.

To make the last point precise it is convenient to make use of *cost functions*.

Cost functions, distances, and weights. In the context of transport plans, one is often interested in the *cost* of a transport plan whenever the underlying space is equipped with a suitable cost function, for instance, a metric. This is part of the so-called theory of *(optimal) transport* (see the first chapters of [Vil09] for an overview).

A cost function on X is a function $c : X \times X \to [0, \infty]$. For the purposes of the present paper, we can interpret the quantity $c(x, y)$ as the "cost of going from x to y". We are in particular interested in cost functions such that $c(x, x) = 0$, and satisfying the triangle inequality $c(x, z) \le c(x, y) + c(y, z)$. If X is equipped with such a cost function, transport plans can be assigned a cost as well (see Section 4). The cost of a transport plan between the distributions p and q can be interpreted as the cost of "moving the mass from the pile p to the pile q according to the specified plan". If the cost function c satisfies a triangle inequality, the costs of transport plans satisfy one too, whenever two plans are composed. The resulting structure is called a *weighted category*, or *normed category* (see Remark 2.2 for the terminology), a particular kind of enriched category.

Weighted categories have been around at least since [Law73], and have recently sparked some interest in different fields of mathematics, including geometry and topological data analysis [BdSS17, Kub17, LI21]. They can be considered a common generalization of categories and metric spaces. In optimal transport theory, given probability measures p and q on a metric space, one can define their *Wasserstein distance* by taking the infimum of the cost of transport plans between p and q (see Section 4 as well as [Vil09, Section 6]). This is an instance of a more general categorical construction: given objects X and Y of a weighted category, one can consider all the weights of the morphisms $X \to Y$ and take their infimum, the "minimal cost of going from X to Y", and obtain this way a pseudo-quasimetric between the objects. Several metric spaces appearing in mathematics can be obtained in this way: besides Wasserstein spaces, also the Gromov-Hausdorff distance is an example of this construction, as defined for example in [BBI01, Proposition 7.3.16 and Exercise 7.3.26]. The notion of weighted category is therefore very "primitive", it underlies several contemporary geometric constructions.

As we said above, the main theme of this work is the idea of *lifting*, which is common to both category theory and geometry, and therefore well-suited for treatment in terms of weighted categories. Indeed, lenses in category theory and submetries in metric geometry are very much

related. We incorporate the two concepts by defining *weighted lenses* in Section 3, which are lenses where the liftings are weight-preserving.

The main result of this work shows that under some natural conditions, a metric lens between two spaces gives rise to a weighted lens between the respective categories of couplings (Theorem 5.5), making mathematically precise the idea of "lifting transport plans while preserving their cost".

Outline.   In Section 2 we review some of the basic concepts of weighted category theory. The definitions of embedding and of symmetry for weighted categories seem to be new, but can be seen as special cases of the enriched-categorical notion.

In Section 3 we consider *lenses*, the categorical structures formalizing the idea of "lifting" used in this work. We review two related notions of lens, for sets and for categories, and define weighted versions of them. We then show the link with submetries in metric geometry, which seems to be a new idea.

In Section 4 we define, for each standard Borel space X, a category PX whose objects are probability measures on X, whose morphisms are couplings, and where composition is given by means of conditional products. If X is equipped with a cost function, we can turn PX into a weighted category (in different ways, one for each integer, analogously to L^p spaces), which after optimization give the famous Wasserstein spaces.

In Section 5 we state and prove the main statements of this work. Proposition 5.2 and Proposition 5.4 show that pushforwards of probability measures are functorial, and that pushforwards along embeddings are embeddings. The most important result is Theorem 5.5, which shows that a set-based lens between two spaces gives rise to a categorical lens between the respective spaces of probability measures, effectively lifting entire couplings along the chosen lifting of points. This gives in particular a functor between measurable lenses and weighted lenses (Corollary 5.7). In Example 5.8 we show that this includes the special case of product projections, where the lifting is given by taking the conditional product. All these results come with a weighted version in the metric and pseudo-quasimetric case.

2   Weighted categories

Here we recall the main ideas of weighted categories and functors, their relationship with ordinary categories and with metric spaces, and their notion of symmetry, isomorphism, and embedding. Some of the definition given here are technically new (such as the notion of symmetry, Definition 2.8), but very similar to well known structures (such as dagger categories).

2.1   Main definitions

Definition 2.1. *A* weighted category *is a category $\mathcal{C}$ where to each arrow $f : X \to Y$ we assign a value $w(f) \in [0, \infty]$ called the* weight *of f, such that*

- *all identities have weight zero;*

- *for every pair of composable arrows $f : X \to Y$ and $g : Y \to Z$, we have the "triangle inequality"*

$$w(g \circ f) \leq w(g) + w(f). \tag{2.1}$$

In this work we will mostly consider small categories, where the objects form a set. This will always be implicitly assumed from now on, unless otherwise stated.

Remark 2.2. Some other authors use other names, such as "normed categories". Also [Gra07] and [BdSS17] call these "weighted categories". In particular, in [Gra07] it is argued that these structures are more lax than what one may want from a norm, they are more similar to *seminorms*. Probably for similar reasons, in [LI21] these structures are called "seminormed categories". The original works [Law73] and [BG75], call them "normed categories" (the latter in Italian, "categorie normate"), while [Kub17] and [LI21] reserve that term for weighted categories which satisfy (different) special properties. (For the latter, see also our Remark 2.13.) We choose the name "weighted categories" as it seems to be the one giving rise to the least ambiguity.

Readers interested in enriched category theory [Kel82] can view weighted categories as enriched categories, where the base of the enrichment is the category of weighted sets [Gra07, Section 1.4]. We will not explicitly adopt this point of view in this work, except for some conventions. The enrichment approach is worked out in detail in [Gra07].

Every ordinary category can be considered a weighted category where all arrows have weight zero. Because of that, all the statements that we make for weighted categories also hold for ordinary categories. At the other end of the spectrum, weighted categories also generalize metric spaces, and pseudo-quasimetric spaces.

Example 2.3. A *pseudo-quasimetric space* (more briefly, *pq-metric space*) or *Lawvere metric space* is a set X together with a "cost" function $c : X \times X \to [0, \infty]$, such that

- for each $x \in X$, we have $c(x, x) = 0$;

- for each x, y and $z \in X$, we have the triangle inequality

$$c(x, z) \leq c(x, y) + c(y, z).$$

A pq-metric space can be seen as a weighted category with exactly one arrow between any two objects x and y, and the quantity $c(x, y)$ is the weight of the arrow.

In analogy to metric spaces, we will sometimes call c *distance*. Note that in general, for a pq-metric,

- if $c(x, y) = 0$, it is not necessarily the case that $x = y$;

- symmetry is not required: $c(x, y)$ can be different from $c(y, x)$. In particular, one of the two may be zero when the other one is not;

- infinities are allowed.

Here is another example, again from metric geometry.

Example 2.4. Consider the category $\mathcal{L}ip$ where

- objects are metric spaces;

- morphisms are (all) functions;

- the weight of a morphism $f : X \to Y$ is the maximum between 0 and the *logarithm* of its Lipschitz constant,

$$\sup_{x,x' \in X} \log \frac{d\big(f(x), f(x')\big)}{d(x, x')},$$

which is possibly infinite (and with the convention that $0/0 = 1$).

$\mathcal{L}ip$ is a (large) weighted category. (One can also define *multiplicative weights* and avoid taking the logarithm, see [Gra07] for more on this.), as the formula (5.1) will show.

Every weighted category gives a pq-metric space by "optimization over the arrows", as follows.

Definition 2.5. *Let $\mathcal{C}$ be a weighted category. The* optimization *of $\mathcal{C}$ is the pq-metric space* $\mathrm{Opt}(\mathcal{C})$ *where the points are the objects of $\mathcal{C}$, and the distance is given by*

$$c(X, Y) := \inf_{f:X \to Y} w(f), \tag{2.2}$$

where by convention, if there are no arrows from X to Y, the infimum over the empty set gives ∞. As it can be easily checked, this quantity is indeed a pq-metric.

This construction was "left as exercise" in [Law73, Introduction, pg. 139-140], and fully worked out (and generalized) in [BG75]. It is an instance of the "change of enrichment" construction of enriched category theory.

Note that in general the minimum is not realized.

Definition 2.6. *A weighted category $\mathcal{C}$ is called* optimization-complete *if for every two objects X and Y of $\mathcal{C}$, there exists a morphism $X \to Y$ of minimal weight. That is, the infimum in* (2.2) *is attained.*

As functors between weighted categories it is useful to take those functors which are "weight-nonincreasing" (or "1-Lipschitz"):

Definition 2.7. *Let $\mathcal{C}$ and $\mathcal{D}$ be weighted categories. A* weighted functor $F : \mathcal{C} \to \mathcal{D}$ *is a functor such that for each morphism $f : X \to Y$ of $\mathcal{C}$*

$$w(Ff) \leq w(f).$$

We denote by $\mathcal{WCat}$ the category of (small) weighted categories and weighted functors.

If $\mathcal{C}$ and $\mathcal{D}$ are ordinary categories, and we consider them as weighted categories where all the weights are zero, a weighted functor $\mathcal{C} \to \mathcal{D}$ is the same thing as an ordinary functor. Categorically, this gives a fully faithful functor $\mathcal{Cat} \to \mathcal{WCat}$.

If X and Y are pq-metric spaces and we consider them as weighted categories, a function $f : X \to Y$ is a weighted functor if and only if it is 1-Lipschitz. (Since arrows are unique, identities and composition are trivially preserved.) Calling $\mathbf{pqMet}$ the category of pq-metric spaces and 1-Lipschitz maps, we then have a fully faithful functor $\mathbf{pqMet} \to \mathcal{WCat}$. Somewhat conversely, any weighted functor between weighted categories gives a 1-Lipschitz function between the pq-metric spaces obtained via optimization. Therefore, Opt is a functor $\mathcal{WCat} \to \mathbf{pqMet}$.

The following definition can be seen as the many-arrow analogue of the symmetry of a metric, or as a weighted analogue of dagger categories. Similarly to the unweighted case, denote by $\mathcal{C}^{\mathrm{op}}$ the category obtained by reversing all arrows of $\mathcal{C}$, with the same weights.

Definition 2.8. *A* symmetric weighted category *is a weighted category $\mathcal{C}$ equipped with a weighted functor $\dagger : \mathcal{C}^{\mathrm{op}} \to \mathcal{C}$, called the* symmetry *or* dagger*, such that*

- *for each object X, $X^{\dagger} = X$ (i.e. $\dagger$ is the identity on the objects);*

- *for each morphism f, we have that $f^{\dagger\dagger} = f$ (i.e. $\dagger$ is an involution, $\dagger^2 = \mathrm{id}$).*

In other words, in a symmetric weighted category to each morphism $f : X \to Y$ there is a morphism in the opposite direction, $f^{\dagger} : Y \to X$, necessarily of the same weight as f. (And moreover, this correspondence preserves identities and composition).

For example, if we take a pq-metric space and consider it as a weighted category as in Example 2.3, it is symmetric as a weighted category if and only if the cost function is symmetric (for example, if it is a metric). Similarly, if $\mathcal{C}$ is a symmetric weighted category, then the cost function of the optimization $\mathrm{Opt}(\mathcal{C})$ will be symmetric.

2.2   Isomorphisms

Of particular interest is the notion of isomorphism in a weighted category, which is again bringing together category-theoretic and geometric concepts.

Definition 2.9. *Let X and Y be objects in a weighted category $\mathcal{C}$. A* quasi-isomorphism *between X and Y is a pair of morphisms $f : X \to Y$ and $g : Y \to X$ of finite weight, such that $g \circ f = \mathrm{id}_X$ and $f \circ g = \mathrm{id}_Y$.*

An isomorphism *between X and Y is a quasi-isomorphism pair (f, g) such that both f and g have weight zero.*

We also write $g = f^{-1}$, $f = g^{-1}$, and sometimes we denote the pair (f, g) simply by f. If an isomorphism (resp. quasi-isomorphism) between X and Y exists, we say that X and Y are isomorphic *(resp.* quasi-isomorphic*), and we write $X \cong Y$ (resp. $X \simeq Y$).*

Our term "quasi-isomorphism" is unrelated to the notion in homological algebra with the same name.

Remark 2.10. Just as for the definition of weighted categories, different authors refer to weighted isomorphisms in different ways. In general, only invertible morphisms *of weight zero* tend to behave like "full" isomorphisms. (See the next two examples.) This intuition can be made mathematically precise: considering weighted categories as enriched categories, the invertible morphisms of weight zero are exactly the isomorphisms in the sense of enriched category theory (as defined for example in [Kel82]). In [LI21] these morphisms are called *norm isomorphisms*, in [Kub17] they are called *0-isomorphisms*.

Example 2.11. If X is a pq-metric space, considered as a weighted category, the points $x, y \in X$ are isomorphic if and only if $c(x, y) = 0$ and $c(y, x) = 0$. In particular, in a metric space, $x \cong y$ of and only if $x = y$. The points x and y are quasi-isomorphic if and only if they have finite distance in both directions.

Example 2.12. The isomorphisms in the category $\mathcal{L}ip$ of Example 2.4 are precisely the isometries. The quasi-isomorphisms are precisely the bi-Lipschitz maps.

Remark 2.13. Given isomorphic objects X and Y in a weighted category $\mathcal{C}$, they will also be isomorphic in $\mathrm{Opt}(\mathcal{C})$, i.e. have distance zero in both directions. Instead, if X and Y are isomorphic in $\mathrm{Opt}(\mathcal{C})$, they are not necessarily isomorphic in $\mathcal{C}$: all we can say is that for every $\varepsilon > 0$, there are morphisms $f_\varepsilon : X \to Y$ and $g_\varepsilon : Y \to X$ of weight less than ε. If $\mathcal{C}$ is optimization-complete, then there are also morphisms $f : X \to Y$ and $g : Y \to X$ of weight zero realizing the two minima. However, this still does not guarantee that f and g are inverses. Similar consideration can be given for quasi-isomorphisms.

The authors of [LI21] define a *normed category* as a weighted category $\mathcal{C}$ where, in our terminology, for every two objects X and Y we have that $X \cong Y$ in $\mathcal{C}$ if and only if $X \cong Y$ in $\mathrm{Opt}(\mathcal{C})$. (Recall however that the term "normed category" is used differently by different authors, see Remark 2.2.)

2.3 Embeddings

Embeddings of weighted categories are the common generalization of isometric embeddings of metric spaces and of fully faithful functors.

Definition 2.14. *Let $\mathcal{C}$ and $\mathcal{D}$ be weighted categories. An* embedding *of $\mathcal{C}$ into $\mathcal{D}$ is a weighted functor $F : \mathcal{C} \to \mathcal{D}$ such that for every pair of objects X and Y of $\mathcal{C}$, the function*

$$\mathrm{Hom}_{\mathcal{C}}(X, Y) \longrightarrow \mathrm{Hom}_{\mathcal{D}}(FX, FY)$$
$$f \longmapsto Ff$$

is a weight-preserving bijection.

Example 2.15. If X and Y are metric spaces (or even pq-metric spaces), considered as weighted categories, a map $f : X \to Y$ is an embedding in the sense above exactly if it is an isometric embedding. Somewhat conversely, an embedding of weighted categories $F : \mathcal{C} \to \mathcal{D}$ induces an isometric embedding of the pq-metric spaces $\mathrm{Opt}(\mathcal{C}) \to \mathrm{Opt}(\mathcal{D})$.

In ordinary category theory, fully faithful functors are injective up to isomorphism, and in metric geometry, isometric embeddings are injective up to zero distances. Both these statements can be considered a special case of the following.

Proposition 2.16. *Let $F : \mathcal{C} \to \mathcal{D}$ be an embedding of weighted categories, and let X and Y be objects of $\mathcal{C}$. If FX and FY are isomorphic (resp. quasi-isomorphic) in $\mathcal{D}$, then X and Y are already isomorphic (resp. quasi-isomorphic) in $\mathcal{C}$.*

Proof. Let $f : FX \to FY$ be an isomorphism of $\mathcal{D}$, with inverse $g : FY \to FX$, both of weight zero. Since F is an embedding, it is surjective on morphisms, and so necessarily $f = F\tilde{f}$ and $g = F\tilde{g}$ for some $\tilde{f} : X \to Y$ and $\tilde{g} : Y \to Z$, also of weight zero. Now,

$$F(\tilde{g} \circ \tilde{f}) = F\tilde{g} \circ F\tilde{f} = g \circ f = \mathrm{id}_{FX} = F(\mathrm{id}_X),$$

but since F is injective on morphisms, we have $\tilde{g} \circ \tilde{f} = \mathrm{id}_X$. Similarly, $\tilde{f} \circ \tilde{g} = \mathrm{id}_Y$. Therefore $\tilde{f}$ and $\tilde{g}$ form an isomorphism pair between X and Y.

The quasi-isomorphism case works analogously. $\qquad\square$

3 Lenses and submetries

In the setting of weighted categories, we can bring together the concept of *lens* (from category theory) and the concept of *submetry* (from metric geometry).

When it comes to lenses, different authors mean different specific concepts under the same name. Here we will use two (related) flavors of lenses: *set-based* lenses and *delta* lenses, and we will provide weighted versions of both.

In our context, we can interpret lenses, as well as submetries, in terms of "lifting". Indeed, suppose that we have two sets X and Y, and a "projection" of X onto Y. We sketch the situation in the following picture.
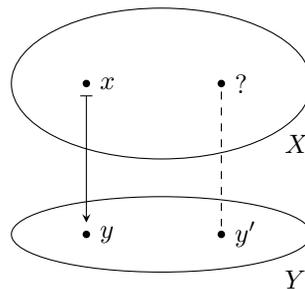
Given a point x of X, as in the picture above, we can project it down to a point y of Y. If we now move from y to a different point y' in Y, can we lift y' back to a point of X which "corresponds to x in a canonical way"? Here, "canonical" may mean different things in different contexts. For metric spaces we may want, for example, that $d(x, x') = d(y, y')$. In other situations, we may want a form of parallel transport instead. Lenses and submetries make this intuition precise in these different settings. In particular, *set-based lenses* can be seen as lifting *points*, while *delta lenses* can be seen as lifting *paths*, in a way that preserves identities and composition.

3.1 Lifting points

Here we use the notion of (set-based) lens first defined in [FGM$^+$07], and analyzed in detail for example in [Cla20a] (where they are called "state-based lenses").

Definition 3.1. *A* set-based lens *consists of*

- *two sets X and Y;*

- *a function $f : X \to Y$;*

- *a function $\varphi : X \times Y \to X$,*

such that the following equations are satisfied for every $x \in X$ and $y \in Y$,

$$f\big(\varphi(x, y)\big) = y \qquad\qquad \varphi\big(x, f(x)\big) = x \qquad\qquad \varphi\big(\varphi(x, y), y'\big) = \varphi(x, y') \qquad (3.1)$$

or equivalently, such that the following diagrams commute,

$$
\begin{array}{ccc}
X \times Y \xrightarrow{\varphi} X & \quad X \xrightarrow{(\mathrm{id},f)} X \times Y & \quad X \times Y \times Y \xrightarrow{\varphi \times \mathrm{id}} X \times Y \\
\end{array}
\tag{3.2}
$$

where the maps π are the product projections onto the respective factors.

We will refer to the conditions (3.1) (or equivalently, (3.2)) as the *lens laws.* In particular we call them, in order, the *lifting law*, the *identity law*, and the *composition law*. We can interpret the maps and the laws as follows.

- The map f is a projection from X to Y.

- The map φ takes a point $x \in X$ and a point $y \in Y$, and gives a lifting of y back to a point $\varphi(x,y) \in X$ which in some sense "corresponds to x, but above y".

- The lifting law says that indeed, the chosen lifting $\varphi(x,y)$ is in the preimage $f^{-1}(y)$, i.e. "lies above y".

- The identity law says that the point over $f(x)$ corresponding to x is x itself.

- The composition law says that the procedure can be iterated, and the resulting lift does not change.

Lenses can be composed. Indeed, if (f,φ) is a lens between X and Y and (g,ψ) is a lens between Y and Z, the composite lens between X and Z is given by

$$
X \xrightarrow{f} Y \xrightarrow{g} Z
$$

and

$$
X \times Z \xrightarrow{(\mathrm{id},f) \times \mathrm{id}} X \times Y \times Z \xrightarrow{\mathrm{id} \times \psi} X \times Y \xrightarrow{\varphi} X.
$$

This can be interpreted in terms of sequential liftings: given $x \in X$ and $z \in Z$, we first lift z to Y using ψ, and then lift the result to X using φ.

Remark 3.2. Set-based lenses can always be turned into product projections, making the set X isomorphic to a product $Y \times Z$ [SN22, Example 6.74]. To see it, take for example Z to be the set of sections $s : Y \to X$ of f (such that $f \circ s = \mathrm{id}_Y$) with the "rigidity" property that $\varphi(s(y), y') = s(y')$ for all $y, y' \in Y$. This way, the map $f : X \to Y$ corresponds to the product projection

$$
\begin{aligned}
Y \times Z &\xrightarrow{f} Y \\
(y,z) &\longmapsto y
\end{aligned}
\tag{3.3}
$$

and the map ϕ is then the canonical lifting

$$
\begin{aligned}
(Y \times Z) \times Y &\xrightarrow{\varphi} Y \times Z \\
\big((y,z),y'\big) &\longmapsto (y',z).
\end{aligned}
\tag{3.4}
$$

One can interpret this as "choosing the point in $Y \times Z$ which lies over y', but which is at the same height as (y,z)". Geometrically, φ can be seen as a choice of the "horizontal" direction in the product $Y \times Z$, analogous to a choice of connection on a (trivial) bundle.

The definition of set-based lens can be given in any category with products [JRW10]. For example, we are interested in *measurable* lenses.

Definition 3.3. *A* measurable lens *consists of a lens (f,φ) between measurable spaces X and Y, such that the functions f and φ are measurable functions.*

Given a measurable lens (f, φ) between measurable spaces X and Y, by Remark 3.2 it is still true that on the underlying sets it can be seen as a product projection $Y \times Z \to Y$. However, in principle, the σ-algebra on X might be different than the product σ-algebra of $Y \times Z$. The question of whether measurable lenses can be always turned into product projections in the category of measurable spaces is still open, and we leave it to future work. (See however Remark 3.8.)

Let's now move to the metric setting, and let's recall the notion of submetry. Traditionally, it is given as follows [Ber87, BP10].[1]

Definition 3.4. *Let X and Y be metric spaces. Denote by $B_X(x, r)$ the closed ball in X of center x and radius r, and define B_Y analogously. A map $f : X \to Y$ is a* submetry *if for every $x \in X$ and for every $r > 0$ we have that*

$$f\big(B_X(x, r)\big) = B_Y\big(f(x), r\big).$$

As it can be checked, this definition can be rephrased in the following equivalent way, which is closer to our idea of "lifting".

Definition 3.5 (alternative)**.** *Let X and Y be metric spaces. A function $f : X \to Y$ is a* submetry *if it is 1-Lipschitz, and for every $x \in X$ and $y' \in Y$ there exists $x' \in X$ such that*

$$f(x') = y' \qquad and \qquad d(x, x') = d\big(f(x), y'\big).$$

This definition allows us to draw a link between submetries and lenses.[2]

Definition 3.6. *A* metric lens *is a lens (X, Y, f, φ) where X and Y are metric spaces, f is 1-Lipschitz, and f and φ satisfy the requirement that*

$$d\big(x, \varphi(x, y)\big) = d\big(f(x), y\big). \tag{3.5}$$

Note that this is not the internal notion of lens in a category of metric spaces: φ is not required to be 1-Lipschitz, or even continuous. (It is rather a special case of weighted lens, more on that later.)

Every metric lens is canonically a submetry. Indeed, let (X, Y, f, φ) be a metric lens, and let $x \in X$ and $y' \in Y$. A point $x' \in X$ such that

$$f(x') = y' \qquad and \qquad d(x, x') = d\big(f(x), y'\big).$$

is simply given by $x' \coloneqq \varphi(x, y')$. Here is a partial converse to this fact.

Proposition 3.7. *Let $f : X \to Y$ be a submetry. For every $x \in X$ and $y' \in Y$, choose an $x' \in X$ such that*

$$f(x') = y' \qquad and \qquad d(x, x') = d\big(f(x), y'\big).$$

Denote this chosen x' by $\varphi(x, y')$, so that we have a function $\varphi : X \times Y \to Y$.

Then the pair (f, φ) satisfies the lifting and identity axioms of lenses (3.1).

Proof. The lifting law says that

$$f\big(\varphi(x, y)\big) = y,$$

which is just the condition that $f(x') = y'$.

The identity law says that

$$\varphi\big(x, f(x)\big) = x.$$

To prove it, note that by the condition on distances,

$$d\big(x, \varphi\big(x, f(x)\big)\big) = d\big(f(x), f(x)\big) = 0,$$

therefore $x = \varphi\big(x, f(x)\big)$. □

---

[1]In the Riemannian geometry literature, sometimes people use a definition similar to this one, but where the condition holds only for $r \leq R$ for some fixed bound R. We will not use such notion here.

[2]Callum Reader, in personal communication, pointed out to me that there is an additional, very deep connection between submetries and lenses. Indeed, one can define *enriched* lenses between enriched categories [CDM22], and if one instantiates the definition in the case of Lawvere metric spaces, one obtains exactly *weak submetries*, which are defined by replacing closed balls with open balls in Definition 3.4, see [BP10, Section 1].

Therefore, one can view a metric lens as "a submetry with a choice of liftings, compatible with composition".

One can generalize the notion of metric lens to arbitrary pq-metric spaces. (Note that for pq-metric spaces a "submetry" only satisfies the identity law up to distance zero, i.e. $\varphi\big(x, f(x)\big)$ and x are not necessarily the same point, but have distance zero.)

Remark 3.8. Differently from the set case, metric lenses are not necessarily just product projections. Once again, they are product projections on the underlying sets, but the set $X \cong Y \times Z$ may be equipped with a different metric than the product one. In particular, we get a metric lens as long as we equip $Y \times Z$ with a metric such that for every $z \in Z$,

$$d\big((y, z), (y', z)\big) = d(y, y').$$

Moreover, the map φ is 1-Lipschitz whenever the metric is larger or equal than the product metric, i.e.

$$d((y, z), (y', z')) \geq \max\{d(y, y'), d(z, z')\}.$$

These two conditions are very natural conditions for a metric on $Y \times Z$, and most metrics on a product used in the literature do satisfy these requirements.

3.2 Lifting paths

We now turn to categories and weighted categories. This time we don't just want to lift points, but also paths, or arrows. We will consider a weighted analogue of the structure called *delta lens*, first defined in [DXC11]. (See also [Cla20b] for the internal version of delta lenses.)

We recall the original definition, and then present the weighted version.

Definition 3.9. *Let $\mathcal{E}$ and $\mathcal{B}$ be categories. Let $F : \mathcal{E} \to \mathcal{B}$ be a functor, let $b : B \to B'$ be a morphism of $\mathcal{B}$, and let E be an object of $\mathcal{E}$ with $F(E) = B$. A lifting of b at E is an object E' of $\mathcal{E}$ such that $F(E') = B'$, together with an arrow $\Phi(E, b) : E \to E'$ of $\mathcal{E}$ such that $F\big(\Phi(E, b)\big) = b$.*
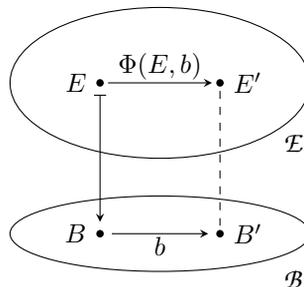
Definition 3.10. *Let $\mathcal{E}$ and $\mathcal{B}$ be categories. A* delta lens *or* categorical lens *from $\mathcal{E}$ to $\mathcal{B}$ consists of the following data.*

- *A functor $F : \mathcal{E} \to \mathcal{B}$.*

- *For each morphism $b : B \to B'$ of $\mathcal{B}$ and each object E of $\mathcal{E}$ with $F(E) = B$, a chosen lifting $\Phi(E, b) : E \to E'$ of b satisfying the following conditions.*

  - *The identities are lifted to identities, i.e. $\Phi(E, \mathrm{id}_B) = \mathrm{id}_E$.*

  - *The choice of liftings preserves composition. That is, given $b : B \to B'$ and $b' : B' \to B''$, consider the chosen lifting $\Phi(E, b) : E \to E'$ of b, and then, at its codomain E' consider the lifting $\Phi(E', b')$ of b'. We require that $\Phi(E', b') \circ \Phi(E, b) = \Phi(E, b' \circ b)$.*

A famous particular case of lens is given by *split Grothendieck opfibrations*, where liftings are moreover *initial* in some sense. A somewhat trivial case is given by set-based lenses, which are equivalently delta lenses between *codiscrete* categories (a unique arrow between any two objects). See for example [Cla20b] for more on this.

Here is the weighted analogue of liftings and lenses.[3]

---

[3] Once again, this is an instance of *enriched lens* as given in [CDM22], for the case of weighted categories.

Definition 3.11. *Let $\mathcal{E}$ and $\mathcal{B}$ be weighted categories. Let $F : \mathcal{E} \to \mathcal{B}$ be a weighted functor, let $b : B \to B'$ be a morphism of $\mathcal{B}$, and let E be an object of $\mathcal{E}$ with $F(E) = B$. A* weight-preserving *lifting of b at E is a lifting $\Phi(E, b)$ of b such that $w\big(\Phi(E, b)\big) = w(b)$.*

A weighted lens *is defined similarly to a delta lens, but where all the liftings are required to be weight-preserving.*

Consider now pq-metric spaces X and Y as weighted categories as in Example 2.3. A weighted lens between them consists of

- a 1-Lipschitz map $f : X \to Y$;

- for every $x \in X$ and $y' \in Y$, a chosen lifting of y' to a point $x' \in f^{-1}(y')$ such that $d(f(x), y') = d(x, x')$.

This is, in other words, a metric lens in the sense of Definition 3.6.

Also delta lenses can be composed, analogously to set-based lenses. For the details, see [Cla20b]. For convenience, let's define the following categories.

- *SetLens* is the category of sets and set-based lenses.

- *pqMetLens* is the category of pq-metric spaces and metric lenses.

- *CatLens* is the category of (small) categories and delta lenses.

- *WCatLens* is the category of (small) weighted categories and weighted lenses.

Between these categories we have the following fully faithful functors.

- By considering two sets as a codiscrete categories, set-based lenses between them are the same as delta lenses. This gives a fully faithful functor *SetLens → CatLens*,

- Analogously, by considering pq-metric spaces as weighted categories with a single arrow between any two objects, we get a fully faithful functor *pqMetLens → WCatLens*.

- Considering ordinary categories as having weight zero for all morphisms gives a fully faithful functor *CatLens → WCatLens*.

- One can construct a similar functor *SetLens → pqMetLens*, making the following diagram commute.

$$
\begin{array}{ccc}
\mathit{SetLens} & \longrightarrow & \mathit{pqMetLens} \\
\downarrow & & \downarrow \\
\mathit{CatLens} & \longrightarrow & \mathit{WCatLens}
\end{array}
$$

4   The category of couplings

Here we briefly review some basic notions of optimal transport theory, mostly to fix the notation. For most of the constructions in this section we require the measurable spaces to be standard Borel, also known as "Polish": those measurable spaces whose σ-algebra can be obtained as the Borel σ-algebra of a complete, separable metric space. This is needed in particular to have conditional distributions, more on that later in this section.

For X a (generic) measurable space, denote by $\Sigma(X)$ its σ-algebra, and by PX the set of probability measures over it. If X and Y are measurable spaces, denote by $X \times Y$ their Cartesian product, equipped with the product σ-algebra.

Definition 4.1. *Let X and Y be measurable spaces. Let $p \in PX$ and $q \in PY$ be probability measures. A* coupling *of p and q is a probability measure $s \in P(X \times Y)$ such that its marginals on X and Y are p and q, respectively.*

We denote by $\Gamma(p, q)$ the space of couplings of p and q.

Given a coupling s of p and q, and if $A \subseteq X$ and $B \subseteq Y$ are measurable subsets, we will write $s(A \times B)$ as shorthand for the measure of the "cylinder"

$$s\Big(\big(\pi_1^{-1}(A) \times Y\big) \cap \big(X \times \pi_2^{-1}(B)\big)\Big),$$

where $\pi_1 : X \times Y \to X$ and $\pi_2 : X \times Y \to Y$ are the product projections.

Definition 4.2. *Let X be a measurable space, and let $c : X \times X \to [0, \infty]$ be a pq-metric. Let p and q be Borel probability measures on X, and let s be a coupling of p and q. For an integer $k \geq 1$, the k-cost of the coupling s is the quantity*

$$\mathrm{cost}_k(s) := \sqrt[k]{\iint_{X \times X} c(x,y)^k \, s(dx\,dy)}.$$

Given a standard Borel space X, and fixing k, we now construct a category whose objects are probability measures on X, and whose morphisms are couplings with the weight given by the k-cost.

Given $p \in PX$, the identity coupling I_p is given, intuitively, by "a copy of p supported on the diagonal". We can express this rigorously as

$$I_p(A \times A') := p(A \cap A').$$

Therefore, the cost of the identity is zero for every k:

$$\mathrm{cost}_k(I_p)^k = \int_{X \times X} c(x,y)^k \, I_p(dx\,dy) = \int_X c(x,x)^k \, p(dx) = 0.$$

In order to define composites, we have to use conditional probabilities.

4.1   Composition

Let X and Y be measurable spaces. A *Markov kernel* from X to Y is a mapping

$$X \times \Sigma(Y) \xrightarrow{\ \ h\ \ } [0,1]$$
$$(x, B) \longmapsto h(B|x)$$

such that, for each $x \in X$, $h(-|x) : \Sigma(Y) \to [0,1]$ is a probability measure, and for each $B \in \Sigma(Y)$, $h(B|-) : X \to [0,1]$ is a measurable function.

Markov kernels are often used to denote conditional probabilities, and this works particularly well in standard Borel spaces. We use the following standard "conditioning" result, which will underlie most of the constructions in the rest of this work. (For a reference, see for example [Fre06, Section 452] and [Bog00, Section 10.4].)

Theorem 4.3. *Let X and Y be standard Borel spaces, let p and q be probability measures on X and Y respectively, and let s be a coupling of p and q. There exists a Markov kernel $\vec{s}$ from X to Y such that for every pair measurable subsets $A \subseteq X$ and $B \subseteq Y$,*

$$\int_A \vec{s}(B|x) \, p(dx) = s(A \times B);$$

Moreover, such a kernel $\vec{s}$ is unique p-almost surely.

We call $\vec{s}$ the *conditional distribution* of Y given X (associated to the coupling s).

One can also condition s on Y, obtaining a kernel from Y to X. We denote this by $\overleftarrow{s}(A|y)$. We then have a form of Bayes' theorem, saying that for $A \subseteq X$ and $B \subseteq Y$ measurable,

$$\int_A \vec{s}(B|x) \, p(dx) = s(A \times B) = \int_B \overleftarrow{s}(A|y) \, q(dy). \tag{4.1}$$

A conditional distribution of the identity coupling is given by the indicator function:

$$\vec{I}_p(A|x) = \reflectbox{$\vec{\reflectbox{$I$}}$}_p(A|x) = 1_A(x) = \begin{cases} 1 & x \in A \\ 0 & x \notin A. \end{cases}$$

or equivalently by the map assigning to $x \in X$ the Dirac delta measure δ_x.

Let's now turn to the composition of couplings. This is a standard construction, it is used in optimal transport, under the name "gluing" [Vil09, Chapter 1], as well as in statistics, where it's related to the "conditional product" construction [DS99], and in categorical probability (see [Sim18, Example 6.2] and the references therein, as well as the more general construction of [Fri20, Definition 12.8]).

Definition 4.4. *Let p, q and r be probability measures on the standard Borel spaces X, Y and Z, respectively. Let s be a coupling of p and q, and let t be a coupling of q and r. The distribution $t \circ s$ on $X \times Z$ is defined by*

$$(t \circ s)(A \times C) := \int_Y \overleftarrow{s}(A|y)\,\vec{t}(C|y)\,q(dy)$$

for all measurable subsets $A \subseteq X$ and $C \subseteq Z$.

We have that the marginal of $t \circ s$ on X, evaluated on the measurable subset $A \subseteq X$, is

$$(t \circ s)(A \times Z) = \int_Y \overleftarrow{s}(A|y) \cdot 1\, q(dy) = s(A \times Y) = p(A),$$

and the same can be done for Z, so that $t \circ s$ is a coupling of p and r. Moreover, $t \circ s$ does not depend on the choice of the conditional distributions $\overleftarrow{s}$ and $\vec{t}$, since they are unique q-almost everywhere.

The composite coupling gives, up to measure zero, the (Chapman-Kolmogorov) composition of conditional distributions. That is, for p-almost all x, and for all measurable $C \subseteq Z$,

$$\overrightarrow{t \circ s}(C|x) = \int_Y \vec{t}(C|y)\,\vec{s}(dy|x),$$

where the expression on the right-hand side denotes the Lebesgue integral of the measurable function $\vec{t}(C|-) : Y \to \mathbb{R}$ with measure $\vec{s}(-|x) \in PY$.

Proposition 4.5. *Let X be a standard Borel space, and let $c : X \times X \to [0, \infty]$ be a measurable pq-metric (i.e. measurable as a function). Let $p, q, r \in PX$, and let s, t be couplings of p and q and of q and r, respectively. Then*

$$\mathrm{cost}_k(t \circ s) \le \mathrm{cost}_k(s) + \mathrm{cost}_k(t).$$

Proof. Using Minkowski's inequality,

$$
\begin{aligned}
\mathrm{cost}(t \circ s) &= \sqrt[k]{\iint_{X \times X} c(x,z)^k\,(t \circ s)(dx\,dz)} \\
&= \sqrt[k]{\iint_{X \times X \times X} c(x,z)^k\,\overleftarrow{s}(dx|y)\,\vec{t}(dz|y)\,q(dy)} \\
&\le \sqrt[k]{\iint_{X \times X \times X} \big(c(x,y) + c(y,z)\big)^k\,\overleftarrow{s}(dx|y)\,\vec{t}(dz|y)\,q(dy)} \\
&\le \sqrt[k]{\iint_{X \times X} c(x,y)^k\,\overleftarrow{s}(dx|y)\,q(dy)} + \sqrt[k]{\iint_{X \times X} c(y,z)^k\,\vec{t}(dz|y)\,q(dy)} \\
&= \mathrm{cost}(s) + \mathrm{cost}(t). \qquad \square
\end{aligned}
$$

Using this fact we can now show that for each $k \ge 1$, couplings and their costs form a weighted category.

Proposition 4.6. *Let X be a standard Borel space, and let $c : X \times X \to [0, \infty]$ be a measurable pq-metric. For each k, the set of probability measures PX forms a weighted category, where the arrows are couplings with k-cost as their weight.*

We will denote the weighted category by $P_k(X, c)$, or more briefly $P_k X$ then this causes no ambiguity. For the unweighted case, we will simply denote the category by PX.

Proof. The only properties that need to be checked are associativity and unitality of the composition. Note that those need to hold *strictly*, not just almost surely (since the morphisms of PX are joint distributions, not conditional).

For unitality, let $p, q \in PX$, and let $s \in \Gamma(p, q)$. We have that for all measurable sets $A, B \subseteq X$, using formula (4.1),

$$(s \circ I_p)(A \times B) = \int_X 1_A(x) \, \vec{s}(B|x) \, p(dx) = \int_A \vec{s}(B|x) \, p(dx) = s(A \times B),$$

and similarly for $I_q \circ s$, so that $s \circ I_p = s = I_q \circ s$.

For associativity, let $p, q, m, n \in PX$, and consider composable couplings $s \in \Gamma(p, q)$, $t \in \Gamma(q, m)$, and $u \in \Gamma(m, n)$. Then for all measurable sets $A, B \subseteq X$, again using (4.1),

$$
\begin{aligned}
((u \circ t) \circ s)(A \times B) &= \int_X \overleftarrow{s}(A|y) \, \overrightarrow{u \circ t}(B|y) \, q(dy) \\
&= \int_X \int_X \overleftarrow{s}(A|y) \, \vec{u}(B|z) \, \vec{t}(dz|y) \, q(dy) \\
&= \int_X \int_X \overleftarrow{s}(A|y) \, \vec{u}(B|z) \, \overleftarrow{t}(dy|z) \, m(dz) \\
&= \int_X \overleftarrow{t \circ s}(A|z) \, \vec{u}(B|z) \, m(dz) \\
&= (u \circ (t \circ s))(A \times B). \qquad \square
\end{aligned}
$$

Optimization on the arrows gives the following pq-metrics,

$$c_k(p, q) = \inf_{s \in \Gamma(p, q)} \mathrm{cost}_k(s) = \inf_{s \in \Gamma(p, q)} \sqrt[k]{\int_{X \times X} c(x, y) \, s(dx \, dy)}, \qquad (4.2)$$

which are the celebrated Wasserstein pq-metrics. Moreover, the following are true.

- If the cost function c on X is symmetric, then $P_k X$ is a symmetric weighted category, and hence (4.2) is a symmetric cost function as well (i.e. a pseudometric).

- It is well known that the infimum in (4.2) is actually attained in many cases [Vil09, Theorem 4.1], and in that case $P_k X$ is an optimization-complete weighted category.

- If we see X as a weighted category (with unique morphisms between any two points), the map $\delta : X \to PX$ assigning to each point x the Dirac measure δ_x is an embedding of weighted categories. The reason why we have an embedding is that any two Dirac measures δ_x and δ_y admit a unique coupling $\delta_{(x,y)}$.

- The map δ is natural in X, on objects it can be seen as the unit of the Kantorovich monad [vB05, FP19]. Moreover, the multiplication of the Kantorovich monad is a nontrivial example of a submetry, as proven in [FP20, Appendix A].

5 Main statements

Here are some results involving the constructions defined in the previous sections. In Section 5.1 we show that pushforwards of probability measures are functorial (Corollary 5.3), and that the pushforward of an embedding is an embedding of categories (both for the weighted and unweighted case). Most importantly, in Section 5.2 we show that lenses between the underlying spaces induce (weighted) lenses between the respective categories of couplings, where the liftings are formed by suitable conditional distributions, and once again that such an assignment is functorial (Corollary 5.7).

5.1   Pushforward measures

Recall that given a measurable function $f : X \to Y$ and a probability measure $p \in PX$, the *pushforward* of p along f is the probability measure $f_\sharp p \in PY$ given by

$$f_\sharp p(B) := p(f^{-1}(B))$$

for all measurable subsets $B \subseteq Y$. The assignment $p \mapsto f_\sharp p$ gives then a function $f_\sharp : PX \to PY$. Examples of pushforward measures are given by the marginal projections: the marginalization $P(X \times Y) \to PX$ is the pushforward along the product projection $\pi_1 : X \times Y \to X$.

This assignment is functorial, in the sense that if we denote by *Meas* the category of measurable spaces and functions, the pushforward gives a functor $P : \textit{Meas} \to \textit{Set}$ which maps a measurable set X to the set PX and a measurable function $f : X \to Y$ to its pushforward. One can also equip the sets PX themselves with a measurable structure, and obtain an endofunctor on *Meas*, the underlying functor of the Giry monad [Gir82]. Since in the previous section we made PX a (weighted) category, we would now like to show that the pushforward $f_\sharp : PX \to PY$ is a (weighted) functor.

Before that, let's look at a useful technical result.

Lemma 5.1. *Let $f : X \to Y$ be a measurable function between standard Borel spaces. Let $s \in P(X \times X)$ be a coupling of p and $q \in PX$. A conditional distribution for the pushforward coupling $f_\sharp^2 s \in P(Y \times Y)$ between $f_\sharp p$ and $f_\sharp q$ satisfies*

$$\overrightarrow{f_\sharp^2 s}\big(B'|f(x)\big) = \vec{s}\big(f^{-1}(B')|x\big)$$

for p-almost all $x \in X$ and all measurable $B' \subseteq Y$.

Proof. For all measurable subsets $B, B' \subseteq Y$, we have that

$$
\begin{aligned}
\int_{f^{-1}(B)} \overrightarrow{f_\sharp^2 s}\big(B'|f(x)\big)\, p(dx) &= \int_B \overrightarrow{f_\sharp^2 s}(B'|y)\, f_\sharp p(dy) \\
&= f_\sharp^2 s(B \times B') \\
&= s\big(f^{-1}(B) \times f^{-1}(B')\big) \\
&= \int_{f^{-1}(B)} \vec{s}\big(f^{-1}(B')|x\big)\, p(dx). \qquad \square
\end{aligned}
$$

Now let's apply this result to the following proposition.

Proposition 5.2. *Let X and Y be standard Borel spaces. Let $f : X \to Y$ be a measurable function. The pushforward mapping $f_\sharp : PX \to PY$ induces a functor between the respective categories of couplings, where the pushforward of a coupling is given by the map*

$$P(X \times X) \xrightarrow{(f \times f)_\sharp} P(Y \times Y),$$

or more briefly, $f_\sharp^2 : P(X \times X) \to P(Y \times Y)$.

Moreover, if X and Y are equipped with measurable cost functions, and f is 1-Lipschitz, then $f_\sharp$ is a weighted functor $P_k X \to P_k Y$ for all k.

Explicitly, given $s \in P(X \times X)$ and measurable subsets $B, B' \subseteq Y$ the function $f_\sharp^2$ gives,

$$f_\sharp^2 s(B \times B') = s\big(f^{-1}(B) \times f^{-1}(B')\big).$$

Proof. For the identity coupling I_p, we have that

$$f_\sharp^2 I_p(B \times B') = p\big(f^{-1}(B) \cap f^{-1}(B')\big) = p\big(f^{-1}(B \cap B')\big) = I_{f_\sharp p}(B \times B')$$

for all measurable subsets $B, B' \subseteq Y$. For composition, let $p, q, r \in PX$, let s be a coupling from p to q, and let t be a coupling between q and r. By Lemma 5.1,

$$f_\sharp^2(t \circ s)(B \times B') = (t \circ s)\big(f^{-1}(B) \times f^{-1}(B')\big)$$

$$= \int_X \overleftarrow{s}\big(f^{-1}(B)|x\big)\, \overrightarrow{t}\big(f^{-1}(B')|x\big)\, q(dx)$$

$$= \int_X \overleftarrow{f_\sharp^2 s}\big(B|f(x)\big)\, \overrightarrow{f_\sharp^2 t}\big(B'|f(x)\big)\, q(dx)$$

$$= \int_Y \overleftarrow{f_\sharp^2 s}(B|y)\, \overrightarrow{f_\sharp^2 t}(B'|y)\, f_\sharp q(dy)$$

$$= (f_\sharp^2 t) \circ (f_\sharp^2 s)(B \times B'),$$

again for all measurable subsets $B, B' \subseteq Y$. Therefore $f_\sharp : PX \to PY$ is functorial.

Suppose now that X and Y are equipped with cost functions c_X and c_Y. The (k-th power of the) cost of $f_\sharp^2 s \in P(Y \times Y)$ is

$$\text{cost}(f_\sharp^2 s)^k = \int_{Y \times Y} c_Y(y, y')^k\, f_\sharp^2 s(dy\, dy')$$

$$= \int_{X \times X} c_Y\big(f(x), f(x')\big)^k\, s(dx\, dx')$$

$$\leq \int_{X \times X} c_X(x, x')^k\, s(dx\, dx')$$

$$= \text{cost}(s)^k. \qquad \square$$

It is easy to see that the assignment $f \mapsto f_\sharp$ also preserves identities and composition.

Corollary 5.3. *Denote by $\mathcal{SBS}$ the category of standard Borel spaces and measurable maps. The assignment $X \mapsto PX$ and $f \mapsto f_\sharp$ defines a functor $\mathcal{SBS} \to \mathcal{Cat}$.*

Just as well, denote by $\mathcal{SBpqMet}$ the category of standard Borel spaces equipped with measurable cost functions and 1-Lipschitz measurable maps. For each k, the assignment $X \mapsto P_k X$ and $f \mapsto f_\sharp$ defines a functor $\mathcal{SBpqMet} \to \mathcal{WCat}$.

This assignment $f \mapsto f_\sharp$ also preserves embeddings, as the next proposition shows.

Proposition 5.4. *Let X and Y be standard Borel spaces, and let $i : X \to Y$ be a measurable embedding (i.e. an injective function, such that $\Sigma(X)$ can be written as the pullback σ-algebra of $\Sigma(Y)$ along the inclusion i). Then the map $i_\sharp : PX \to PY$ is a full embedding of categories.*

If moreover X and Y are equipped with measurable cost functions and i is an isometric embedding, then $i_\sharp : P_k X \to P_k Y$ is an embedding of weighted categories for every k.

Proof. Let $p, q \in PX$. Consider the restriction of $i_\sharp^2 : P(X \times X) \to P(Y \times Y)$ to the couplings of p and q, call it again $i_\sharp^2 : \Gamma(p, q) \to \Gamma(i_\sharp p, i_\sharp q)$. We have to prove that this new mapping is a (weight-preserving) bijection. To prove that it is injective, consider $r, r' \in \Gamma(p, q)$ and suppose that $i_\sharp^2 r = i_\sharp^2 r'$. Since i is a measurable embedding, for all measurable subsets $A, A' \subseteq X$ we can find subsets $B, B' \subseteq Y$ such that $A = i^{-1}(B)$ and $A' = i^{-1}(B')$. Therefore, for all measurable subsets $A, A' \subseteq X$,

$$r(A \times A') = r(i^{-1}(B) \times i^{-1}(B')) = i_\sharp^2 r(B \times B') = i_\sharp^2 r'(B \times B') = r'(A \times A'),$$

so that $r = r'$.

To prove surjectivity, let s be a coupling of $i_\sharp p$ and $i_\sharp q$. Since i is a measurable embedding, for each measurable subset $A \subset X$, the (non-inverse) image $i(A) \subseteq Y$ is again measurable. Construct now the coupling $r \in \Gamma(p, q)$ by

$$r(A \times A') := s\big(i(A) \times i(A')\big),$$

so that $i_\sharp^2 r = s$. Therefore $i_\sharp^2 : \Gamma(p, q) \to \Gamma(i_\sharp p, i_\sharp q)$ is a bijection.

Suppose now that X and Y have cost functions c_X and c_Y, and that i is an isometric embedding. Then given $r \in \Gamma(p, q)$,

$$\text{cost}_k(i_\sharp^2 s)^k = \int_{Y \times Y} c_Y(y, y')^k\, i_\sharp^2 r(dy\, dy')$$

$$= \int_{X \times X} c_Y\big(i(x), i(x')\big)^k r(dx\, dx')$$

$$= \int_{X \times X} c_X(x, x')^k r(dx\, dx')$$

$$= \mathrm{cost}_k(s)^k,$$

which means that the bijection is weight-preserving. $\qquad\qquad\qquad\qquad\qquad\qquad\square$

5.2   Lenses to lenses

We now come to the main result of this work. Namely, we show that measurable lenses give rise to lenses between the categories of couplings, making $X \mapsto PX$ functorial on lenses as well. Again, conditional distributions are the key construction.

Theorem 5.5. *Let $(f, \varphi) : X \to Y$ be a measurable lens between standard Borel spaces. There is a categorical lens $(f_\sharp, \tilde{\varphi}_\sharp)$ between PX and PY where*

- *the projection $f_\sharp : PX \to PY$ is the usual pushforward of measures;*

- *the lifting $\tilde{\varphi}_\sharp : PX \times_{PY} P(Y \times Y) \to P(X \times X)$ takes $p \in PX$ and a coupling $s \in P(Y \times Y)$ whose first marginal has to equal $f_\sharp p$, and returns the coupling $\tilde{\varphi}_\sharp(p, s) \in P(X \times X)$ given by*

$$\tilde{\varphi}_\sharp(p, s)(A \times A') := \int_A \int_Y 1_{A'}\big(\varphi(x, y)\big)\, \vec{s}\big(dy | f(x)\big)\, p(dx) \tag{5.1}$$

for all measurable subsets $A, A' \subseteq X$.

Moreover, if X and Y are equipped with measurable cost functions and (f, φ) is a measurable metric lens, $(f_\sharp, \tilde{\varphi}_\sharp)$ is a weighted lens between $P_k X$ and $P_k Y$ for each k.

In order to better understand the formula (5.1), we can look at its conditional form,

$$\overrightarrow{\tilde{\varphi}_\sharp(s)}(A | x) = \int_Y 1_A\big(\varphi(x, y)\big)\, \vec{s}\big(dy | f(x)\big). \tag{5.2}$$

First of all, note that this conditional distribution does not depend on p (the joint does). Now (5.2) says that the probability of transitioning from x to a point of A is intuitively obtained as follows.

- For each $y \in Y$, we look at the probability (given by $\vec{s}$) of going from $f(x)$ to y.

- We lift each point $y \in Y$ to the point $\varphi(x, y)$ as prescribed by the lens (f, φ). We can do so measurably.

- We construct a kernel that maps x to $\varphi(x, y)$ with the same probability as s maps $f(x)$ to y.

- We look at the probability, as prescribed by this new kernel, that $\varphi(x, y)$ is in A.

Let's now look at the rigorous proof.

Proof. We already know by Proposition 5.2 that $f_\sharp : PX \to PY$ is a functor, weighted if X and Y have cost functions and f is 1-Lipschitz.

Let us now turn to $\tilde{\varphi}_\sharp$. We first have to show that for every $p \in PX$ and for every coupling $s \in P(Y \times Y)$ with $f_\sharp p$ as first marginal, the coupling $\tilde{\varphi}_\sharp(p, s)$ is a lifting of s at p. For that, its first marginal needs to be equal to p, and its pushforward along f down to $Y \times Y$ needs to be equal to s. Now for every measurable $A \subseteq X$,

$$\tilde{\varphi}_\sharp(p, s)(A \times X) = \int_A \int_Y 1\, \vec{s}\big(dy | f(x)\big)\, p(dx) = \int_A p(dx) = p(A),$$

and for every measurable $B, B' \subseteq Y$,

$$(f \times f)_\sharp \tilde{\varphi}_\sharp(p, s)(B \times B') = \tilde{\varphi}_\sharp(p, s)\big(f^{-1}(B) \times f^{-1}(B')\big)$$

$$= \int_{f^{-1}(B)} \int_Y 1_{f^{-1}(B')}\big(\varphi(x,y)\big)\, \vec{s}\big(dy|f(x)\big)\, p(dx)$$

$$= \int_{f^{-1}(B)} \int_Y 1_{B'}\big(f(\varphi(x,y))\big)\, \vec{s}\big(dy|f(x)\big)\, p(dx)$$

$$= \int_{f^{-1}(B)} \int_Y 1_{B'}(y)\, \vec{s}\big(dy|f(x)\big)\, p(dx)$$

$$= \int_{f^{-1}(B)} \vec{s}\big(B'|f(x)\big)\, p(dx)$$

$$= \int_B \vec{s}\big(B'|y\big)\, f_\sharp p(dx)$$

$$= s(B \times B'),$$

where between the third and fourth line we used that $f(\varphi(x,y)) = y$, i.e. the lifting law of the lens (f,φ). Therefore $\tilde{\varphi}_\sharp(p,s)$ is a lifting of s at p.

In order to show that that $\tilde{\varphi}_\sharp$ preserves identity and composition, it is better to use the conditional form of the coupling (5.2). Now, to show that the identity is lifted to the identity,

$$\overrightarrow{\tilde{\varphi}_\sharp(p,I_q)}(A|x) = \int_Y 1_A\big(\varphi(x,y)\big)\, \delta_{f(x)}(dy) = 1_A\big(\varphi(x,f(x))\big) = 1_A(x),$$

where we used that $\varphi(x,f(x)) = x$, i.e. the identity law of the lens (f,φ). To show the composition property, consider the composable couplings $s,s' \in P(Y \times Y)$. Then

$$\int_X \overrightarrow{\tilde{\varphi}_\sharp(s')}(A'|x')\, \overrightarrow{\tilde{\varphi}_\sharp(s)}(dx'|x) = \int_X \int_Y 1_{A'}\big(\varphi(x',y')\big)\, \vec{s}'\big(dy'|f(x')\big) \int_Y \delta_{\varphi(x,y)}(dx')\, \vec{s}\big(dy|f(x)\big)$$

$$= \int_Y 1_{A'}\big(\varphi(\varphi(x,y),y')\big) \int_Y \vec{s}'\big(dy'|f(\varphi(x,y))\big)\, \vec{s}\big(dy|f(x)\big)$$

$$= \int_Y 1_{A'}\big(\varphi(x,y')\big) \int_Y \vec{s}'(dy'|y)\, \vec{s}\big(dy|f(x)\big)$$

$$= \int_Y 1_{A'}\big(\varphi(x,y')\big)\, \overrightarrow{s' \circ s}\big(dy'|f(x)\big),$$

where between the second and third line we used that $\varphi(\varphi(x,y),y') = \varphi(x,y')$, i.e. the composition law of the lens (f,φ), and again its lifting law.

Lastly, suppose that X and Y are equipped with cost functions c_X and c_Y. Then the (k-th power of the) cost of the lifting is

$$\text{cost}_k(\tilde{\varphi}_\sharp(p,s))^k = \int_{X \times X} c_X(x,x')^k\, \tilde{\varphi}_\sharp(p,s)(dx\,dx')$$

$$= \int_{X \times X} \int_Y c_X(x,x')^k\, \delta_{\varphi(x,y)}(dx')\, \vec{s}\big(dy|f(x)\big)\, p(dx)$$

$$= \int_X \int_Y c_X\big(x,\varphi(x,y)\big)^k\, \vec{s}\big(dy|f(x)\big)\, p(dx)$$

$$= \int_X \int_Y c_Y\big(f(x),y\big)^k\, \vec{s}\big(dy|f(x)\big)\, p(dx)$$

$$= \int_{Y \times Y} c_Y\big(y',y\big)^k\, \vec{s}\big(dy|y'\big)\, f_\sharp p(dy)$$

$$= \int_{Y \times Y} c_Y\big(y',y\big)^k\, s(dy\,dy') = \text{cost}_k(s)^k,$$

where between the third and the fourth line we used that $c_X\big(x,\varphi(x,y)\big) = c_Y\big(f(x),y\big)$, i.e. the metric lens property of (f,φ). Therefore the $\tilde{\varphi}_\sharp(p,s)$ is a weight-preserving lifting, and so $(f_\sharp, \tilde{\varphi}_\sharp)$ is a weighted lens. $\qquad\square$

Corollary 5.6. *If (f,φ) is a metric, measurable lens between X and Y as in Theorem 5.5, then $f : PX \to PY$ is in particular a submetry. This remains true if we drop from (f,φ) the composition requirement, as in Proposition 3.7.*

Just as for the case of pushforwards, the assignment $(f, \varphi) \mapsto (f_\sharp, \tilde{\varphi}_\sharp)$ preserves identities and composition of lenses.

Corollary 5.7. *Denote by $\mathscr{SBSLens}$ the category of standard Borel spaces and measurable lenses. The assignment $X \mapsto PX$ and $(f, \varphi) \mapsto (f_\sharp, \tilde{\varphi}_\sharp)$ defines a functor $\mathscr{SBSLens} \to \mathscr{CatLens}$.*

Just as well, denote by $\mathscr{SBpqMetLens}$ the category of standard Borel spaces equipped with measurable cost functions and metric, measurable lenses. For each k, the assignment $X \mapsto P_k X$ and $(f, \varphi) \mapsto (f_\sharp, \tilde{\varphi}_\sharp)$ defines a functor $\mathscr{SBpqMetLens} \to \mathscr{WCatLens}$.

Finally, let's see what happens for the case of marginal distributions, the pushforwards along product projections.

Example 5.8. We have seen in Remark 3.2 that product projections admit a canonical lens structure, which is metric in the case of pq-metric spaces. Let's see the explicit construction of Theorem 5.5 for this case. So again set $X = Y \times Z$, and consider the projection $\pi_1 : Y \times Z \to Y$.

- The map $(\pi_1)_\sharp : P(Y \times Z) \to PY$ is just marginalization, it takes a measure on $Y \times Z$ and gives the marginal on Y.

- The map $\tilde{\varphi}_\sharp : P(Y \times Z) \times_{PY} P(Y \times Y) \to P(Y \times Z \times Y \times Z)$ takes a joint $r \in Y \times Z$ and a coupling $s \in Y \times Y$ with common first marginal $p \in PY$, and gives the coupling $\tilde{\varphi}_\sharp(r, s)$ on $Y \times Z \times Y \times Z$ given as follows. For every measurable $A, A' \subseteq Y$ and $B, B' \subseteq Z$,

$$
\begin{aligned}
\tilde{\varphi}_\sharp(r,s)(A \times B \times A' \times B') &= \int_{A \times B} \int_Y 1_{A' \times B'}\big(\varphi((y,z), y')\big)\, \vec{s}(dy'|\pi_1(y,z))\, r(dy\, dz) \\
&= \int_{A \times B} \int_Y 1_{A' \times B'}(y', z)\, \vec{s}(dy'|y)\, r(dy\, dz) \\
&= \int_{A \times B} 1_{B'}(z) \int_Y 1_{A'}(y')\, \vec{s}(dy'|y)\, r(dy\, dz) \\
&= \int_{A \times B} 1_{B'}(z)\, \vec{s}(A'|y)\, r(dy\, dz) \\
&= \int_A \int_B 1_{B'}(z)\, \vec{r}(dz|y)\, \vec{s}(A'|y)\, p(dy) \\
&= \int_A \vec{r}(B \cap B'|y)\, \vec{s}(A'|y)\, p(dy).
\end{aligned}
$$

In other words, the lifting given by $\tilde{\varphi}_\sharp$ is a coupling that is the identity between the Z-components, and is a conditional product of r and s along p in the Y-components.

References

[BBI01]   D Burago, Y. Burago, and S Ivanov. *A Course in Metric Geometry*, volume 33 of *Graduate Studies in Mathematics*. American Mathematical Society, 2001.

[BdSS17] Peter Bubenik, Vin de Silva, and Jonathan Scott. Interleaving and gromov-hausdorff distance, 2017. arXiv:1707.06288.

[Ber87]   V. N. Berestovskiĭ. "Submetries" of three-dimensional forms of nonnegative curvature. *Sibirsk. Mat. Zh.*, 28(4):44–56, 1987. In Russian.

[BG75]   Renato Betti and Massimo Galuzzi. Categorie normate. *Bollettino dell'Unione Matematica Italiana*, 11(1):66–75, 1975.

[Bog00]   V. I. Bogachev. *Measure Theory. Vol. I, II.* Springer, 2000.

[BP10]   V. N. Berestovskiĭ and C. P. Plaut. Covering R-trees, R-free groups, and dendrites. *Advances in Mathematics*, 224(5):1765–1783, 2010.

[CDM22]   Bryce Clarke and Matthew Di Meglio. An introduction to enriched cofunctors, 2022. arXiv: 2209.01144.

[Cla20a]   Bryce Clarke. Internal lenses as functors and cofunctors. In *Proceedings of Applied Category Theory 2019*, volume 323 of *Electronic Proceedings in Theoretical Computer Science*, pages 183–185. 2020.

[Cla20b]   Bryce Clarke. Internal split opfibrations and cofunctors. *Theory and Applications of Categories*, 35(44):1608–1633, 2020.

[DS99]   A. Philip Dawid and Milan Studený. Conditional products: an alternative approach to conditional independence. In *Artificial Intelligence and Statistics '99*, pages 32–40. 1999.

[DXC11]   Zinovy Diskin, Yingfei Xiong, and Krzysztof Czarnecki. From state- to delta-based bidirectional model transformations: the asymmetric case. *Journal of Object Technology*, 10, 2011.

[FGM+07]   J. Nathan Foster, Michael B. Greenwald, Jonathan T. Moore, Benjamin C. Pierce, and Alan Schmitt. Combinators for bidirectional tree transformations: A linguistic approach to the view-update problem. *ACM Transactions on Programming Languages and Systems*, 29(3), 2007.

[FP19]   Tobias Fritz and Paolo Perrone. A probability monad as the colimit of spaces of finite samples. *Theory and Applications of Categories*, 34(7):170–220, 2019. arXiv:1712.05363.

[FP20]   Tobias Fritz and Paolo Perrone. Stochastic order on metric spaces and the ordered kantorovich monad. *Advances in Mathematics*, 370, 2020. arXiv:1808.09898.

[Fre06]   D. H. Fremlin. *Measure theory. Vol. 4*. Torres Fremlin, Colchester, 2006. Topological measure spaces. Part I, II, Corrected second printing of the 2003 original.

[Fri20]   Tobias Fritz. A synthetic approach to Markov kernels, conditional independence and theorems on sufficient statistics. *Adv. Math.*, 370:107239, 2020. arXiv:1908.07021.

[Gir82]   Michèle Giry. A Categorical Approach to Probability Theory. In *Categorical aspects of topology and analysis*, volume 915 of *Lecture Notes in Mathematics*. 1982.

[Gra07]   Marco Grandis. Categories, norms and weights. *Journal of homotopy and related structures*, 2(2):171–186, 2007.

[JRW10]   Michael Johnson, Robert Rosebrugh, and Richard Wood. Algebras and update strategies. *Journal of Universal Computer Science*, 16(5):729–748, 2010.

[Kel82]   G. M. Kelly. *Basic concepts of enriched category theory*. Cambridge University Press, 1982.

[Kub17]   Wiesław Kubiś. Categories with norms, 2017. arXiv:1705.10189.

[Law73]   William Lawvere. Metric spaces, generalized logic and closed categories. *Rendiconti del seminario matematico e fisico di Milano*, 43, 1973.

[LI21]   Daniel Luckhardt and Matt Insall. Norms on categories and analogs of the schröder-bernstein theorem, 2021. arXiv:2105.06832.

[Sim18]   Alex Simpson. Category-theoretic structure for independence and conditional independence. In *Proceedings of MFPS 2017*, volume 336 of *Electronic Notes in Theoretical Computer Science*, pages 281–297. 2018.

[SN22]   David I. Spivak and Nelson Niu. Polynomial functors: a general theory of interaction. Available online at topos.site/poly-book.pdf, 2022.

[vB05]   Franck van Breugel. The Metric Monad for Probabilistic Nondeterminism. Available at http://www.cse.yorku.ca, 2005.

[Vil09]   Cédric Villani. *Optimal transport: old and new*, volume 338 of *Grundlehren der mathematischen Wissenschaften*. Springer, 2009.