

Drawing from an Urn is Isometric [★]

Bart Jacobs

Institute for Computing and Information Sciences,
Radboud University, Nijmegen, NL
bart@cs.ru.nl

Abstract. Drawing (a multiset of) coloured balls from an urn is one of the most basic models in discrete probability theory. Three modes of drawing are commonly distinguished: multinomial (draw-replace), hypergeometric (draw-delete), and Pólya (draw-add). These drawing operations are represented as maps from urns to distributions over multisets of draws. The set of urns is a metric space via the Wasserstein distance. The set of distributions over draws is also a metric space, using Wasserstein-over-Wasserstein. The main result of this paper is that the three draw operations are all isometries, that is, they preserve the Wasserstein distances.

Keywords: probability · urn drawing · Wasserstein distance.

1 Introduction

We start with an illustration of the topic of this paper. We consider a situation with a set $C = \{R, G, B\}$ of three colours: red, green, blue. Assume that we have two urns v_1, v_2 with 10 coloured balls each. We describe these urns as multisets of the form:

$$v_1 = 8|G\rangle + 2|B\rangle \quad \text{and} \quad v_2 = 5|R\rangle + 4|G\rangle + 1|B\rangle.$$

Recall that a multiset is like a set, except that elements may occur multiple times. Here we describe urns as multisets using ‘ket’ notation $| - \rangle$. It separates multiplicities of elements (before the ket) from the elements in the multiset (inside the ket). Thus, urn v_1 contains 8 green balls and 2 blue balls (and no red ones). Similarly, urn v_2 contains 5 red, 4 green, and 1 blue ball(s).

Below, we shall describe the Wasserstein distance between multisets (of the same size). How this works does not matter for now; we simply posit that the Wasserstein distance $d(v_1, v_2)$ between these two urns is $\frac{1}{2}$ — where we assume the discrete distance on the set C of colours.

We turn to draws from these two urns, in this introductory example of size two. These draws are also described as multisets, with elements from the set $C = \{R, G, B\}$ of colours. There are six multisets (draws) of size 2, namely:

$$2|R\rangle \quad 1|R\rangle + 1|G\rangle \quad 2|G\rangle \quad 1|R\rangle + 1|B\rangle \quad 2|B\rangle \quad 1|G\rangle + 1|B\rangle. \quad (1)$$

[★] In: Proceedings of FoSSaCS 2024

As we see, there are three draws with 2 balls of the same colour, and three draws with balls of different colours.

We consider the hypergeometric probabilities associated with these draws, from the two urns. Let's illustrate this for the draw $1|G\rangle + 1|B\rangle$ of one green ball and one blue ball from the urn v_1 . The probability of drawing $1|G\rangle + 1|B\rangle$ is $\frac{16}{45}$; it is obtained as sum of:

- first drawing-and-deleting a green ball from $v_1 = 8|G\rangle + 2|B\rangle$, with probability $\frac{8}{10}$. It leaves an urn $7|G\rangle + 2|B\rangle$, from which we can draw a blue ball with probability $\frac{2}{9}$. Thus drawing “first green then blue” happens with probability $\frac{8}{10} \cdot \frac{2}{9} = \frac{8}{45}$.
- Similarly, the probability of drawing “first blue then green” is $\frac{2}{10} \cdot \frac{8}{9} = \frac{8}{45}$.

We can similarly compute the probabilities for each of the above six draws (1) from urn v_1 . This gives the hypergeometric distribution, which we write using kets-over-kets as:

$$hg[2](v_1) = \frac{28}{45} |2|G\rangle + \frac{16}{45} |1|G\rangle + 1|B\rangle + \frac{1}{45} |2|B\rangle.$$

The fraction written before a big ket is the probability of drawing the multiset (of size 2), written inside that big ket, from the urn v_1 .

Drawing from the second urn v_2 gives a different distribution over these multisets (1). Since urn v_2 contains red balls, they additionally appear in the draws.

$$hg[2](v_2) = \frac{2}{9} |2|R\rangle + \frac{4}{9} |1|R\rangle + 1|G\rangle + \frac{2}{15} |2|G\rangle + \frac{1}{9} |1|R\rangle + 1|B\rangle + \frac{4}{45} |1|G\rangle + 1|B\rangle.$$

We can also compute the distance between these two hypergeometric distributions over multisets. It involves a Wasserstein distance, over the space of multisets (of size 2) with their own Wasserstein distance. Again, details of the calculation are skipped at this stage. The distance between the above two hypergeometric draw-distributions is:

$$d(hg[2](v_1), hg[2](v_2)) = \frac{1}{2} = d(v_1, v_2).$$

This coincidence of distances is non-trivial. It holds, in general, for arbitrary urns (of the same size) over arbitrary metric spaces of colours, for draws of arbitrary sizes. Moreover, the same coincidence of distances holds for the multinomial and Pólya modes of drawing. These coincidences are the main result of this paper, see Theorems 1, 2, and 3 below.

In order to formulate and obtain these results, we describe multinomial, hypergeometric and Pólya distributions in the form of (Kleisli) maps:

$$\mathcal{D}(X) \xrightarrow{mn[K]} \mathcal{D}(\mathcal{M}[K](X)) \begin{array}{c} \xleftarrow{hg[K]} \\ \xleftarrow{pl[K]} \end{array} \mathcal{M}[L](X) \quad (2)$$

They all produce distributions (indicated by \mathcal{D}), in the middle of this diagram, on multisets (draws) of size K , indicated by $\mathcal{M}[K]$, over a set X of colours. Details will be provided below. Using the maps in (2), the coincidence of distances that we saw above

can be described as a preservation property, in terms of distance preserving maps — called isometries. At this stage we wish to emphasise that the representation of these different drawing operations as maps in (2) has a categorical background. It makes it possible to formulate and prove basic properties of drawing from an urn, such as naturality in the set X of colours. Also, as shown in [8] for the multinomial and hypergeometric case, drawing forms a monoidal transformation (with ‘zipping’ for multisets as coherence map). This paper demonstrates that the three draw maps (2) are even more well-behaved: they are all isometries, that is, they preserve Wasserstein distances. This is a new and amazing fact.

This paper concentrates on the mathematics behind these isometry results, and not on interpretations or applications. We do like to refer to interpretations in machine learning [14] where the distance that we consider on colours in an urn is called the *ground distance*. Actual distances between colours are used there, based on experiments in psychophysics, using perceived differences [16].

The Wasserstein — or Wasserstein-Kantorovich, or Monge-Kantorovich — distance is the standard distance on distributions and on multisets, going back to [12]. After some preliminaries on multisets and distributions, and on distances in general, Sections 4 and 5 of this paper recall the Wasserstein distance on distributions and on multisets, together with some basic results. The three subsequent Sections 6 – 8 demonstrate that multinomial, hypergeometric and Pólya drawing are all isometric. Distances occur on multiple levels: on colours, on urns (as multisets or distributions) and on draw-distributions. This may be confusing, but many illustrations are included.

2 Preliminaries on multisets and distributions

A *multiset* over a set X is a finite formal sum of the form $\sum_i n_i |x_i\rangle$, for elements $x_i \in X$ and natural numbers $n_i \in \mathbb{N}$ describing the multiplicities of these elements x_i . We shall write $\mathcal{M}(X)$ for the set of such multisets over X . A multiset $\varphi \in \mathcal{M}(X)$ may equivalently be described in functional form, as a function $\varphi: X \rightarrow \mathbb{N}$ with finite support: $\text{supp}(\varphi) := \{x \in X \mid \varphi(x) \neq 0\}$. Such a function $\varphi: X \rightarrow \mathbb{N}$ can be written in ket form as $\sum_{x \in X} \varphi(x) |x\rangle$. We switch back-and-forth between the ket and functional form and use the formulation that best suits a particular situation.

For a multiset $\varphi \in \mathcal{M}(X)$ we write $\|\varphi\| \in \mathbb{N}$ for the *size* of the multiset. It is the total number of elements, including multiplicities:

$$\|\varphi\| := \sum_{x \in X} \varphi(x).$$

For a number $K \in \mathbb{N}$ we write $\mathcal{M}[K](X) \subseteq \mathcal{M}(X)$ for the subset of multisets of size K . There are ‘accumulation’ maps $\text{acc}: X^K \rightarrow \mathcal{M}[K](X)$ turning lists into multisets via $\text{acc}(x_1, \dots, x_K) := 1|x_1\rangle + \dots + 1|x_K\rangle$. For instance $\text{acc}(c, b, a, c, a, c) = 2|a\rangle + 1|b\rangle + 3|c\rangle$. A standard result (see [10]) is that for a multiset $\varphi \in \mathcal{M}[K](X)$ there are $\binom{K!}{\varphi}$ many sequences $\mathbf{x} \in X^K$ with $\text{acc}(\mathbf{x}) = \varphi$, where $\binom{K!}{\varphi} = \prod_x \varphi(x)!$.

Multisets $\varphi, \psi \in \mathcal{M}(X)$ can be added and compared elementwise, so that $(\varphi + \psi)(x) = \varphi(x) + \psi(x)$ and $\varphi \leq \psi$ means $\varphi(x) \leq \psi(x)$ for all $x \in X$. In the latter case, when $\varphi \leq \psi$, we can also subtract $\psi - \varphi$ elementwise.

The mapping $X \mapsto \mathcal{M}(X)$ is functorial: for a function $f: X \rightarrow Y$ we have $\mathcal{M}(f): \mathcal{M}(X) \rightarrow \mathcal{M}(Y)$ given by $\mathcal{M}(f)(\varphi)(y) = \sum_{x \in f^{-1}(y)} \varphi(x)$. This map $\mathcal{M}(f)$ preserves sums and size.

For a multiset $\tau \in \mathcal{M}(X \times Y)$ on a product set we can take its two marginals $\mathcal{M}(\pi_1)(\tau) \in \mathcal{M}(X)$ and $\mathcal{M}(\pi_2)(\tau) \in \mathcal{M}(Y)$ via functoriality, using the two projection functions $\pi_1: X \times Y \rightarrow X$ and $\pi_2: X \times Y \rightarrow Y$. Starting from $\varphi \in \mathcal{M}(X)$ and $\psi \in \mathcal{M}(Y)$, we say that $\tau \in \mathcal{M}(X \times Y)$ is a *coupling* of φ, ψ if φ and ψ are the two marginals of τ . We define the *decoupling* map:

$$\mathcal{M}(X \times Y) \xrightarrow{dcpl := (\mathcal{M}(\pi_1), \mathcal{M}(\pi_2))} \mathcal{M}(X) \times \mathcal{M}(Y) \quad (3)$$

The inverse image $dcpl^{-1}(\varphi, \psi) \subseteq \mathcal{M}(X \times Y)$ is thus the subset of couplings of φ, ψ .

A *distribution* is a finite formal sum of the form $\sum_i r_i |x_i\rangle$ with multiplicities $r_i \in [0, 1]$ satisfying $\sum_i r_i = 1$. Such a distribution can equivalently be described as a function $\omega: X \rightarrow [0, 1]$ with finite support, satisfying $\sum_x \omega(x) = 1$. We write $\mathcal{D}(X)$ for the set of distributions on X . This \mathcal{D} is functorial, in the same way as \mathcal{M} . Both \mathcal{D} and \mathcal{M} are monads on the category **Sets** of sets and functions, but we only use this for \mathcal{D} . The unit and multiplication / flatten maps $unit: X \rightarrow \mathcal{D}(X)$ and $flat: \mathcal{D}^2(X) \rightarrow \mathcal{D}(X)$ are given by:

$$unit(x) := 1|x\rangle \quad flat(\Omega) := \sum_{x \in X} \left(\sum_{\omega \in \mathcal{D}(X)} \Omega(\omega) \cdot \omega(x) \right) |x\rangle. \quad (4)$$

Kleisli maps $c: X \rightarrow \mathcal{D}(Y)$ are also called channels and written as $c: X \dashrightarrow Y$. Kleisli extension $c \gg (-): \mathcal{D}(X) \rightarrow \mathcal{D}(Y)$ for such a channel, is defined on $\omega \in \mathcal{D}(X)$ as:

$$c \gg \omega := flat(\mathcal{D}(c)(\omega)) = \sum_{y \in Y} \left(\sum_{x \in X} \omega(x) \cdot c(x)(y) \right) |y\rangle.$$

Channels $c: X \dashrightarrow Y$ and $d: Y \dashrightarrow Z$ can be composed to $d \circ c: X \dashrightarrow Z$ via $(d \circ c)(x) := d \gg c(x)$. Each function $f: X \rightarrow Y$ gives rise to a deterministic channel $\langle f \rangle := unit \circ f: X \dashrightarrow Y$, that is, via $\langle f \rangle(x) = 1|f(x)\rangle$.

An example of a channel is arrangement $arr: \mathcal{M}[K](X) \rightarrow \mathcal{D}(X^K)$. It maps a multiset $\varphi \in \mathcal{M}[K](X)$ to the uniform distribution of sequences that accumulate to φ .

$$arr(\varphi) := \sum_{\mathbf{x} \in acc^{-1}(\varphi)} \frac{1}{|\varphi|} |\mathbf{x}\rangle = \sum_{\mathbf{x} \in acc^{-1}(\varphi)} \frac{\varphi \parallel}{K!} |\mathbf{x}\rangle. \quad (5)$$

One can show that $\langle acc \rangle \circ arr = \mathcal{D}(acc) \circ arr = unit: \mathcal{M}[K](X) \rightarrow \mathcal{D}(\mathcal{M}[K](X))$. The composite in the other direction produces the uniform distribution of all permutations of a sequence:

$$arr \circ \langle acc \rangle = arr \circ acc = prm \quad \text{where} \quad prm(\mathbf{x}) := \sum_{t: K \cong K} \frac{1}{K!} |t(\mathbf{x})\rangle, \quad (6)$$

in which $\underline{t}(x_1, \dots, x_K) := (x_{t(1)}, \dots, x_{t(K)})$. In writing $t: K \xrightarrow{\cong} K$ we implicitly identify the number K with the set $\{1, \dots, K\}$.

Each multiset $\varphi \in \mathcal{M}(X)$ of non-zero size can be turned into a distribution via normalisation. This operation is called frequentist learning, since it involves learning a distribution from a multiset of data, via counting. Explicitly:

$$\text{Flrn}(\varphi) := \sum_{x \in X} \frac{\varphi(x)}{\|\varphi\|} |x\rangle.$$

For instance, if we learn from an urn with three red, two green and five blue balls, we get the probability distribution for drawing a ball of a particular colour from the urn:

$$\text{Flrn}(3|R + 2|G + 5|B\rangle) = \frac{3}{10}|R\rangle + \frac{1}{5}|G\rangle + \frac{1}{2}|B\rangle.$$

This map Flrn is a natural transformation (but not a map of monads).

Given two distributions $\omega \in \mathcal{D}(X)$ and $\rho \in \mathcal{D}(Y)$, we can form their parallel product $\omega \otimes \rho \in \mathcal{D}(X \times Y)$, given in functional form as:

$$(\omega \otimes \rho)(x, y) := \omega(x) \cdot \rho(y).$$

Like for multisets, we call a joint distribution $\tau \in \mathcal{D}(X \times Y)$ a *coupling* of $\omega \in \mathcal{D}(X)$ and $\rho \in \mathcal{D}(Y)$ if ω, ρ are the two marginals of τ , that is if, $\mathcal{D}(\pi_1)(\tau) = \omega$ and $\mathcal{D}(\pi_2) = \rho$. We can express this also via a decouple map $\text{dcpl} = \langle \mathcal{D}(\pi_1), \mathcal{D}(\pi_2) \rangle$ as in (3).

An *observation* on a set X is a function of the form $p: X \rightarrow \mathbb{R}$. Such a map p , together with a distribution $\omega \in \mathcal{D}(X)$, is called a random variable — but confusingly, the distribution is often left implicit. The map $p: X \rightarrow \mathbb{R}$ will be called a *factor* if it restricts to non-negative reals $X \rightarrow \mathbb{R}_{\geq 0}$. Each element $x \in X$ gives rise to a point observation $\mathbf{1}_x: X \rightarrow \mathbb{R}$, with $\mathbf{1}_x(x') = 1$ if $x = x'$ and $\mathbf{1}_x(x') = 0$ if $x \neq x'$. For a distribution $\omega \in \mathcal{D}(X)$ and an observation $p: X \rightarrow \mathbb{R}$ on the same set X we write $\omega \models p$ for the validity (expected value) of p in ω , defined as (finite) sum: $\sum_{x \in X} \omega(x) \cdot p(x)$. We shall write $\text{Obs}(X) = \mathbb{R}^X$ and $\text{Fact}(X) = (\mathbb{R}_{\geq 0})^X$ for the sets of observations and factors on X .

3 Preliminaries on metric spaces

A metric space will be written as a pair (X, d_X) , where X is a set and $d_X: X \times X \rightarrow \mathbb{R}_{\geq 0}$ is a distance function, also called metric. This metric satisfies:

- $d_X(x, x') = 0$ iff $x = x'$;
- symmetry: $d_X(x, x') = d_X(x', x)$;
- triangular inequality: $d_X(x, x'') \leq d_X(x, x') + d_X(x', x'')$.

Often, we drop the subscript X in d_X if it is clear from the context. We use the standard distance $d(x, y) = |x - y|$ on real and natural numbers.

Definition 1. Let $(X, d_X), (Y, d_Y)$ be two metric spaces.

1. A function $f: X \rightarrow Y$ is called *short* (or also *non-expansive*) if:

$$d_Y(f(x), f(x')) \leq d_X(x, x'), \quad \text{for all } x, x' \in X.$$

Such a map is called an *isometry* or an *isometric embedding* if the above inequality \leq is an actual equality $=$. This implies that the function f is injective, and thus an ‘embedding’.

We write \mathbf{Met}_S for the category of metric spaces with short maps between them.

2. A function $f: X \rightarrow Y$ is *Lipschitz* or *M-Lipschitz*, if there is a number $M \in \mathbb{R}_{>0}$ such that:

$$d_Y(f(x), f(x')) \leq M \cdot d_X(x, x'), \quad \text{for all } x, x' \in X.$$

The number M is sometimes called the *Lipschitz constant*. Thus, a short function is Lipschitz, with constant 1. We write \mathbf{Met}_L for the category of metric spaces with Lipschitz maps between them (with arbitrary Lipschitz constants).

Lemma 1. For two metric spaces (X_1, d_1) and (X_2, d_2) we equip the cartesian product $X_1 \times X_2$ of sets with the sum of the two metrics:

$$d\left((x_1, x_2), (x'_1, x'_2)\right) := d_{X_1}(x_1, x'_1) + d_{X_2}(x_2, x'_2). \quad (7)$$

With the usual projections and tuples this forms a product in the category \mathbf{Met}_L . \square

The product \times also exists in the category \mathbf{Met}_S of metric spaces with short maps. There, it forms a *monoidal* product (a tensor \otimes) since there are no diagonals. In the setting of $[0, 1]$ -bounded metrics (with short maps) one uses the maximum instead of the sum (7) in order to form products (possibly infinite). In the category \mathbf{Met}_L the products $X_1 \times X_2$ with maximum and with sum of distances are isomorphic, via the identity maps. This works since for $r, s \in \mathbb{R}_{\geq 0}$ one as $\max(r, s) \leq r + s$ and $r + s \leq 2 \cdot \max(r, s)$.

4 The Wasserstein distance between distributions

This section introduces the Wasserstein distance between probability distributions and recalls some basic results. There are several equivalent formulations for this distance. We express it in terms of validity and couplings, see also *e.g.* [1,3,6,4].

Definition 2. Let (X, d_X) be a metric space. The Wasserstein metric $d: \mathcal{D}(X) \times \mathcal{D}(X) \rightarrow \mathbb{R}_{\geq 0}$ is defined by any of the three equivalent formulas:

$$\begin{aligned} d(\omega, \omega') &:= \bigwedge_{\tau \in \text{depl}^{-1}(\omega, \omega')} \tau \models d_X \\ &= \bigvee_{p, p' \in \text{Obs}(X), p \oplus p' \leq d_X} \omega \models p + \omega' \models p' \\ &= \bigvee_{q \in \text{Fact}_S(X)} |\omega \models q - \omega' \models q|. \end{aligned} \quad (8)$$

This turns $\mathcal{D}(X)$ into a metric space. The operation \oplus in the second formulation is defined as $(p \oplus p')(x, x') = p(x) + p'(x')$. The set $\text{Fact}_S(X) \subseteq \text{Fact}(X)$ in the third formulation is the subset of short factors $X \rightarrow \mathbb{R}_{\geq 0}$. To be precise, we should write $\text{Fact}_S(X, d_X)$ since the distance d_X on X is a parameter, but we leave it implicit for convenience. The meet \wedge and joins \vee in (8) are actually reached, by what are called the optimal coupling and the optimal observations / factor.

In this definition it is assumed that X is a metric space. This includes the case where X is simply a set, with the discrete metric (where different elements have distance 1). The above Wasserstein distance can then be formulated as what is often called the *total variation distance*. For distributions $\omega, \omega' \in \mathcal{D}(X)$ it is:

$$d(\omega, \omega') = \frac{1}{2} \sum_{x \in X} |\omega(x) - \omega'(x)|.$$

This discrete case is quite common, see *e.g.* [11] and the references given there.

The equivalence of the first and second formulation in (8) is an instance of strong duality in linear programming, which can be obtained via Farkas' Lemma, see *e.g.* [13]. The second formulation is commonly associated with Monge. The single factor q in the third formulation can be obtained from the two observations p, p' in the second formulation, and vice-versa. What we call the Wasserstein distance is also called the Monge-Kantorovich distance.

We do not prove the equivalence of the three formulations for the Wasserstein distance $d(\omega, \omega')$ between two distributions ω, ω' in (8), one with a meet \wedge and two with a join \vee . This is standard and can be found in the literature, see *e.g.* [15]. These three formulations do not immediately suggest how to calculate distances. What helps is that the minimum and maxima are actually reached and can be computed. This is done via linear programming, originally introduced by Kantorovich, see [13,15,3]. In the sequel, we shall see several examples of distances between distributions. They are obtained via our own Python implementation of the linear optimisation, which also produces the optimal coupling, observations or factor. This implementation is used only for illustrations.

Example 1. Consider the set X containing the first eight natural numbers, so $X = \{0, 1, \dots, 7\} \subseteq \mathbb{N}$, with the usual distance, written as d_X , between natural numbers: $d_X(n, m) = |n - m|$. We look at the following two distributions on X .

$$\omega = \frac{1}{2}|0\rangle + \frac{1}{2}|4\rangle \quad \omega' = \frac{1}{8}|2\rangle + \frac{1}{8}|3\rangle + \frac{1}{8}|6\rangle + \frac{5}{8}|7\rangle.$$

We claim that the Wasserstein distance $d(\omega, \omega')$ is $\frac{15}{4}$. This will be illustrated for each of the three formulations in Definition 2.

– The optimal coupling $\tau \in \mathcal{D}(X \times X)$ of ω, ω' is:

$$\tau = \frac{1}{8}|0, 2\rangle + \frac{1}{8}|0, 3\rangle + \frac{1}{8}|0, 6\rangle + \frac{1}{8}|0, 7\rangle + \frac{1}{2}|4, 7\rangle.$$

It is not hard to see that τ 's first marginal is ω , and its second marginal is ω' . We compute the distances as:

$$\begin{aligned} d(\omega, \omega') &= \tau \models d_X \\ &= \frac{1}{8} \cdot d_X(0, 2) + \frac{1}{8} \cdot d_X(0, 3) + \frac{1}{8} \cdot d_X(0, 6) + \frac{1}{8} \cdot d_X(0, 7) + \frac{1}{2} \cdot d_X(4, 7) \\ &= \frac{2}{8} + \frac{3}{8} + \frac{6}{8} + \frac{7}{8} + \frac{3}{2} = \frac{18}{8} + \frac{3}{2} = \frac{9}{4} + \frac{6}{4} = \frac{15}{4}. \end{aligned}$$

- There are the following two optimal observations $p, p': X \rightarrow \mathbb{R}$, described as sums of weighted point predicates:

$$\begin{aligned} p &= -1 \cdot \mathbf{1}_1 - 2 \cdot \mathbf{1}_2 - 3 \cdot \mathbf{1}_3 - 4 \cdot \mathbf{1}_4 - 5 \cdot \mathbf{1}_5 - 6 \cdot \mathbf{1}_6 - 7 \cdot \mathbf{1}_7 \\ p' &= 1 \cdot \mathbf{1}_1 + 2 \cdot \mathbf{1}_2 + 3 \cdot \mathbf{1}_3 + 4 \cdot \mathbf{1}_4 + 5 \cdot \mathbf{1}_5 + 6 \cdot \mathbf{1}_6 + 7 \cdot \mathbf{1}_7. \end{aligned}$$

It is not hard to see that $(p \oplus p')(i, j) := p(i) + p'(j) \leq d_X(i, j)$ holds for all $i, j \in X$. Using the second formulation in (8) we get:

$$\begin{aligned} (\omega \models p) + (\omega' \models p') &= \frac{1}{2} \cdot p(0) + \frac{1}{2} \cdot p(4) + \frac{1}{8} \cdot p'(2) + \frac{1}{8} \cdot p'(3) + \frac{1}{8} \cdot p'(6) + \frac{5}{8} \cdot p'(7) \\ &= \frac{-4}{2} + \frac{2}{8} + \frac{3}{8} + \frac{6}{8} + \frac{35}{8} = -2 + \frac{46}{8} = \frac{30}{8} = \frac{15}{4}. \end{aligned}$$

- Finally, there is a (single) short factor $q: X \rightarrow \mathbb{R}_{\geq 0}$ given by:

$$q = 7 \cdot \mathbf{1}_0 + 6 \cdot \mathbf{1}_1 + 5 \cdot \mathbf{1}_2 + 4 \cdot \mathbf{1}_3 + 3 \cdot \mathbf{1}_4 + 2 \cdot \mathbf{1}_5 + 1 \cdot \mathbf{1}_6.$$

Then:

$$\begin{aligned} (\omega \models q) - (\omega' \models q) &= \frac{1}{2} \cdot q(0) + \frac{1}{2} \cdot q(4) - \left(\frac{1}{8} \cdot q(2) + \frac{1}{8} \cdot q(3) + \frac{1}{8} \cdot q(6) + \frac{5}{8} \cdot q(7) \right) \\ &= \frac{7}{2} + \frac{3}{2} - \left(\frac{5}{8} + \frac{4}{8} + \frac{1}{8} \right) = \frac{10}{2} - \frac{10}{8} = \frac{20}{4} - \frac{5}{4} = \frac{15}{4}. \end{aligned}$$

From the fact that the coupling τ , the two observations p, p' , and the single factor q produce the same distance one can deduce that they are optimal, using the formula (8).

We proceed with several standard properties of the Wasserstein distance on distributions.

Lemma 2. *In the context of Definition 2, the following properties hold.*

1. For an M -Lipschitz function $f: X \rightarrow Y$, the pushforward map $\mathcal{D}(f): \mathcal{D}(X) \rightarrow \mathcal{D}(Y)$ is also M -Lipschitz; as a result, \mathcal{D} lifts to a functor $\mathcal{D}: \mathbf{Met}_L \rightarrow \mathbf{Met}_L$, and also to $\mathcal{D}: \mathbf{Met}_S \rightarrow \mathbf{Met}_S$.
2. If $f: X \rightarrow Y$ is an isometry, then so is $\mathcal{D}(f): \mathcal{D}(X) \rightarrow \mathcal{D}(Y)$.
3. For an M -Lipschitz factor $q: X \rightarrow \mathbb{R}_{\geq 0}$, the validity-of- q factor $(-) \models q: \mathcal{D}(X) \rightarrow \mathbb{R}_{\geq 0}$ is also M -Lipschitz.
4. For each element $x \in X$ and distribution $\omega \in \mathcal{D}(X)$ one has: $d(1|x, \omega) = \omega \models d_X(x, -)$; especially, $d(1|x, 1|x') = d_X(x, x')$, making the map $\text{unit}: X \rightarrow \mathcal{D}(X)$ an isometry.

5. The monad multiplication $\text{flat} : \mathcal{D}^2(X) \rightarrow \mathcal{D}(X)$ is short, so that \mathcal{D} lifts from a monad on **Sets** to a monad on \mathbf{Met}_S and on \mathbf{Met}_L .
6. If a channel $c : X \rightarrow \mathcal{D}(Y)$ is M -Lipschitz, then so is its Kleisli extension $c \gg (-) := \text{flat} \circ \mathcal{D}(c) : \mathcal{D}(X) \rightarrow \mathcal{D}(Y)$.
7. If channel $c : X \rightarrow Y$ is M -Lipschitz and channel $d : Y \rightarrow Z$ is K -Lipschitz, then their (channel) composite $d \circ c : X \rightarrow Z$ is $(M \cdot K)$ -Lipschitz.
8. For distributions $\omega_i, \omega'_i \in \mathcal{D}(X)$ and numbers $r_i \in [0, 1]$ with $\sum_i r_i = 1$ one has:

$$d\left(\sum_i r_i \cdot \omega_i, \sum_i r_i \cdot \omega'_i\right) \leq \sum_i r_i \cdot d(\omega_i, \omega'_i).$$

9. The permutation channel $\text{prm} : X^K \rightarrow \mathcal{D}(X^K)$ from (6) is short.

Proof. We skip the first two points since they are standard.

3. Let $q : X \rightarrow \mathbb{R}_{\geq 0}$ be M -Lipschitz, then $\frac{1}{M} \cdot q : X \rightarrow \mathbb{R}_{\geq 0}$ is short. The function $(-) \models q : \mathcal{D}(X) \rightarrow \mathbb{R}_{\geq 0}$ is then also M -Lipschitz, since for $\omega, \omega' \in \mathcal{D}(X)$,

$$\begin{aligned} |\omega \models q - \omega' \models q| &= M \cdot \left| \omega \models \frac{1}{M} \cdot q - \omega' \models \frac{1}{M} \cdot q \right| \\ &\leq M \cdot \bigvee_{p \in \text{Fact}_S(X)} \left| \omega \models p - \omega' \models p \right| \\ &= M \cdot d(\omega, \omega'). \end{aligned}$$

4. The only coupling of $1|x\rangle, \omega \in \mathcal{D}(X)$ is $1|x\rangle \otimes \omega \in \mathcal{D}(X \times X)$. Hence:

$$d(1|x\rangle, \omega) = 1|x\rangle \otimes \omega \models d_X = \sum_{x' \in X} \omega(x') \cdot d_X(x, x') = \omega \models d_X(x, -).$$

5. We first note that for a distribution of distributions $\Omega \in \mathcal{D}^2(X)$ and a short factor $p : X \rightarrow \mathbb{R}_{\geq 0}$ the validity in Ω of the short validity factor $(-) \models p : \mathcal{D}(X) \rightarrow \mathbb{R}_{\geq 0}$ from item 3 satisfies:

$$\begin{aligned} \Omega \models ((-) \models p) &= \sum_{\omega \in \mathcal{D}(X)} \Omega(\omega) \cdot (\omega \models p) \\ &= \sum_{\omega \in \mathcal{D}(X)} \sum_{x \in X} \Omega(\omega) \cdot \omega(x) \cdot p(x) \\ &\stackrel{(4)}{=} \sum_{x \in X} \text{flat}(\Omega)(x) \cdot p(x) \\ &= \text{flat}(\Omega) \models p. \end{aligned}$$

Thus for $\Omega, \Omega' \in \mathcal{D}^2(X)$,

$$\begin{aligned} &d_X\left(\text{flat}(\Omega), \text{flat}(\Omega')\right) \\ &= \bigvee_{p \in \text{Fact}_S(X)} \left| \text{flat}(\Omega) \models p - \text{flat}(\Omega') \models p \right| \\ &= \bigvee_{p \in \text{Fact}_S(X)} \left| \Omega \models ((-) \models p) - \Omega' \models ((-) \models p) \right| \quad \text{as just shown} \\ &\leq \bigvee_{Q \in \text{Fact}_S(\mathcal{D}(X))} \left| \Omega \models Q - \Omega' \models Q \right| \quad \text{by item 3} \\ &= d_{\mathcal{D}(X)}(\Omega, \Omega'). \end{aligned}$$

6. Directly by points (1) and (5).
7. The channel composite $d \circ c = flat \circ \mathcal{D}(d) \circ c$ consists of a functional composite of M -Lipschitz, K -Lipschitz, and 1-Lipschitz maps, and is thus $(M \cdot K \cdot 1)$ -Lipschitz. This uses items 1 and (5).
8. If we have couplings τ_i for ω_i, ω'_i , then $\sum_i r_i \cdot \tau_i$ is a coupling of $\sum_i r_i \cdot \omega_i$ and $\sum_i r_i \cdot \omega'_i$. Moreover:

$$d\left(\sum_i r_i \cdot \omega_i, \sum_i r_i \cdot \omega'_i\right) \leq \left(\sum_i r_i \cdot \tau_i\right) \models d_X = \sum_i r_i \cdot \left(\tau_i \models d_X\right).$$

Since this holds for all τ_i , we get: $d\left(\sum_i r_i \cdot \omega_i, \sum_i r_i \cdot \omega'_i\right) \leq \sum_i r_i \cdot d(\omega_i, \omega'_i)$.

9. We unfold the definition of the *prm* map from (6) and use the previous item in the first step below. We also use that the distance between two sequences is invariant under permutation (of both).

$$\begin{aligned} d_{\mathcal{D}(X^K)}(\text{prm}(\mathbf{x}), \text{prm}(\mathbf{y})) &\leq \sum_{t: K \overset{\cong}{\rightarrow} K} \frac{1}{K!} \cdot d_{\mathcal{D}(X^K)}(1|\underline{t}(\mathbf{x})\rangle, 1|\underline{t}(\mathbf{y})\rangle) \\ &= \sum_{t: K \overset{\cong}{\rightarrow} K} \frac{1}{K!} \cdot d_{X^K}(\underline{t}(\mathbf{x}), \underline{t}(\mathbf{y})) \quad \text{by item 4} \\ &= \sum_{t: K \overset{\cong}{\rightarrow} K} \frac{1}{K!} \cdot d_{X^K}(\mathbf{x}, \mathbf{y}) = d_{X^K}(\mathbf{x}, \mathbf{y}). \quad \square \end{aligned}$$

Later on we need the following facts about tensors of distributions.

Proposition 1. *Let X, Y be metric spaces, and K be a positive natural number.*

1. *The tensor map $\otimes: \mathcal{D}(X) \times \mathcal{D}(Y) \rightarrow \mathcal{D}(X \times Y)$ is an isometry.*
2. *The K -fold tensor map $\text{iid}[K]: \mathcal{D}(X) \rightarrow \mathcal{D}(X^K)$, given by $\text{iid}[K](\omega) := \omega^K = \omega \otimes \dots \otimes \omega$, is K -Lipschitz. Actually, there is an equality: $d(\omega^K, \rho^K) = K \cdot d(\omega, \rho)$.*

Proof. 1. Let distributions $\omega, \omega' \in \mathcal{D}(X)$ and $\rho, \rho' \in \mathcal{D}(Y)$ be given. For the inequality $d_{\mathcal{D}(X) \times \mathcal{D}(Y)}((\omega, \rho), (\omega', \rho')) \leq d_{\mathcal{D}(X \times Y)}(\omega \otimes \rho, \omega' \otimes \rho')$ one uses that a coupling $\tau \in \mathcal{D}((X \times Y) \times (X \times Y))$ of $\omega \otimes \rho, \omega' \otimes \rho' \in \mathcal{D}(X \times Y)$ can be turned into two couplings τ_1, τ_2 of ω, ω' and of ρ, ρ' , namely as $\tau_i := \mathcal{D}(\pi_i \times \pi_i)(\tau)$. For the reverse inequality one turns two couplings τ_1, τ_2 of ω, ω' and ρ, ρ' into a coupling τ of $\omega \otimes \rho, \omega' \otimes \rho'$ via $\tau := \mathcal{D}(\langle \pi_1 \times \pi_1, \pi_2 \times \pi_2 \rangle)(\tau_1 \otimes \tau_2)$.

2. For $\omega, \rho \in \mathcal{D}(X)$ and $K \in \mathbb{N}$, using the previous item, we get:

$$d_{\mathcal{D}(X^K)}(\omega^K, \rho^K) \stackrel{1}{=} d_{\mathcal{D}(X) \times \dots \times \mathcal{D}(X)}((\omega, \dots, \omega), (\rho, \dots, \rho)) \stackrel{(7)}{=} K \cdot d_{\mathcal{D}(X)}(\omega, \rho). \quad \square$$

5 The Wasserstein distance between multisets

There is also a Wasserstein distance between multisets of the same size. This section recalls the definition and the main results.

Definition 3. Let (X, d_X) be a metric space and $K \in \mathbb{N}$ a natural number. We can turn the metric $d_X: X \times X \rightarrow \mathbb{R}_{\geq 0}$ into the Wasserstein metric $d: \mathcal{M}[K](X) \times \mathcal{M}[K](X) \rightarrow \mathbb{R}_{\geq 0}$ on multisets (of the same size), via:

$$\begin{aligned} d(\varphi, \varphi') &:= \bigwedge_{\tau \in \text{dcp}l^{-1}(\varphi, \varphi')} \text{Flrn}(\tau) \models d_X \\ &= \frac{1}{K} \cdot \bigwedge_{\mathbf{x} \in \text{acc}^{-1}(\varphi), \mathbf{x}' \in \text{acc}^{-1}(\varphi')} d_{X^K}(\mathbf{x}, \mathbf{x}') \\ &\stackrel{(7)}{=} \frac{1}{K} \cdot \bigwedge_{\mathbf{x} \in \text{acc}^{-1}(\varphi), \mathbf{x}' \in \text{acc}^{-1}(\varphi')} \sum_{0 \leq i < K} d_X(x_i, x'_i). \end{aligned} \quad (9)$$

All meets in (9) are finite and can be computed via enumeration. Alternatively, one can use linear optimisation. We give an illustration below. The equality of the first two formulations is standard, like in Definition 2, and is used here without proof. There is an alternative formulation of the above distance between multisets that uses bistochastic matrices, see *e.g.* [2,6], but we do not need it here.

Example 2. Consider the following two multisets of size 4 on the set $X = \{1, 2, 3\} \subseteq \mathbb{N}$, with standard distance between natural numbers.

$$\varphi = 3|1\rangle + 1|2\rangle \qquad \varphi' = 2|1\rangle + 1|2\rangle + 1|3\rangle.$$

The optimal coupling $\tau \in \mathcal{M}[4](X \times X)$ is:

$$\tau = 2|1, 1\rangle + 1|1, 2\rangle + 1|2, 3\rangle.$$

The resulting Wasserstein distance $d(\varphi, \varphi')$ is:

$$\text{Flrn}(\tau) \models d_X = \frac{1}{2} \cdot d_X(1, 1) + \frac{1}{4} \cdot d_X(1, 2) + \frac{1}{4} \cdot d_X(2, 3) = \frac{1}{4} \cdot 1 + \frac{1}{4} \cdot 1 = \frac{1}{2}.$$

Alternatively, we may proceed as follows. There are $\binom{4!}{3! \cdot 1!} = 4$ lists that accumulate to φ , and $\binom{4!}{2! \cdot 1! \cdot 1!} = 12$ lists that accumulate to φ' . We can align them all and compute the minimal distance. It is achieved for instance at:

$$\frac{1}{4} \cdot d_{X^4} \left((1, 1, 1, 2), (1, 1, 2, 3) \right) \stackrel{(7)}{=} \frac{1}{4} \cdot (0 + 0 + 1 + 1) = \frac{2}{4} = \frac{1}{2}.$$

Lemma 3. We consider the situation in Definition 3.

1. Frequentist learning $\text{Flrn}: \mathcal{M}[K](X) \rightarrow \mathcal{D}(X)$ is an isometry, for $K > 0$.
2. For numbers $K, n \geq 1$ the scalar multiplication function $n \cdot (-): \mathcal{M}[K](X) \rightarrow \mathcal{M}[n \cdot K](X)$ is an isometry.
3. The sum of distributions $+: \mathcal{M}[K](X) \times \mathcal{M}[L](X) \rightarrow \mathcal{M}[K + L](X)$ is short.
4. If $f: X \rightarrow Y$ is M -Lipschitz, then $\mathcal{M}[K](f): \mathcal{M}[K](X) \rightarrow \mathcal{M}[K](Y)$ is M -Lipschitz too. Thus, the fixed size multiset functor $\mathcal{M}[K]$ lifts to categories of metric spaces \mathbf{Met}_S and \mathbf{Met}_L .
5. For $K > 0$ the accumulation map $\text{acc}: X^K \rightarrow \mathcal{M}[K](X)$ is $\frac{1}{K}$ -Lipschitz, and thus short.

6. The arrangement channel $\text{arr}: \mathcal{M}[K](X) \rightarrow X^K$ is K -Lipschitz; in fact there is an equality $d(\text{arr}(\varphi), \text{arr}(\varphi')) = K \cdot d(\varphi, \varphi')$.

Proof. 1. Via naturality of frequentist learning: if $\tau \in \mathcal{M}[K](X \times X)$ is a coupling of $\varphi, \varphi' \in \mathcal{M}[K](X)$, then $\text{Flrn}(\tau) \in \mathcal{D}(X \times X)$ is a coupling of $\text{Flrn}(\varphi), \text{Flrn}(\varphi') \in \mathcal{D}(X)$. This gives $d(\varphi, \varphi') \leq d(\text{Flrn}(\varphi), \text{Flrn}(\varphi'))$. The reverse inequality is a bit more subtle. Let $\sigma \in \mathcal{D}(X \times X)$ be an optimal coupling of $\text{Flrn}(\varphi), \text{Flrn}(\varphi')$. Then, since any coupling $\tau \in \mathcal{M}[K](X \times X)$ of φ, φ' gives, as we have just seen, a coupling $\text{Flrn}(\tau) \in \mathcal{D}(X \times X)$ of $\text{Flrn}(\varphi), \text{Flrn}(\varphi')$, we obtain, by optimality:

$$d(\text{Flrn}(\varphi), \text{Flrn}(\varphi')) = \sigma \models d_X \leq \text{Flrn}(\tau) \models d_X.$$

Since this holds for any coupling τ , we get $d(\text{Flrn}(\varphi), \text{Flrn}(\varphi')) \leq d(\varphi, \varphi')$.

2. For multisets $\varphi, \varphi' \in \mathcal{M}[K](X)$, by the previous item:

$$\begin{aligned} d_{\mathcal{M}[K](X)}(\varphi, \varphi') &= d_{\mathcal{D}(X)}(\text{Flrn}(\varphi), \text{Flrn}(\varphi')) \\ &= d_{\mathcal{D}(X)}(\text{Flrn}(n \cdot \varphi), \text{Flrn}(n \cdot \varphi')) \\ &= d_{\mathcal{M}[n \cdot K](X)}(n \cdot \varphi, n \cdot \varphi'). \end{aligned}$$

3. For multisets $\varphi, \varphi' \in \mathcal{M}[K](X)$ and $\psi, \psi' \in \mathcal{M}[L](X)$, using Lemma 2 (8),

$$\begin{aligned} &d(\varphi + \psi, \varphi' + \psi') \\ &\stackrel{1}{=} d(\text{Flrn}(\varphi + \psi), \text{Flrn}(\varphi' + \psi')) \\ &= d\left(\frac{K}{K+L} \cdot \text{Flrn}(\varphi) + \frac{L}{K+L} \cdot \text{Flrn}(\psi), \frac{K}{K+L} \cdot \text{Flrn}(\varphi') + \frac{L}{K+L} \cdot \text{Flrn}(\psi')\right) \\ &\leq \frac{K}{K+L} \cdot d(\text{Flrn}(\varphi), \text{Flrn}(\varphi')) + \frac{L}{K+L} \cdot d(\text{Flrn}(\psi), \text{Flrn}(\psi')) \\ &\stackrel{1}{=} \frac{K}{K+L} \cdot d(\varphi, \varphi') + \frac{L}{K+L} \cdot d(\psi, \psi') \\ &\leq d(\varphi, \varphi') + d(\psi, \psi') \\ &\stackrel{(7)}{=} d((\varphi, \psi), (\varphi', \psi')). \end{aligned}$$

4. Let $f: X \rightarrow Y$ be M -Lipschitz. We use that frequentist learning Flrn is an isometry and a natural transformation $\mathcal{M}[K] \Rightarrow \mathcal{D}$. For multisets $\varphi, \varphi' \in \mathcal{M}[K](X)$,

$$\begin{aligned} &d_{\mathcal{M}[K](Y)}(\mathcal{M}(f)(\varphi), \mathcal{M}(f)(\varphi')) \\ &\stackrel{1}{=} d_{\mathcal{D}(Y)}(\text{Flrn}(\mathcal{M}(f)(\varphi)), \text{Flrn}(\mathcal{M}(f)(\varphi'))) \\ &= d_{\mathcal{D}(Y)}(\mathcal{D}(f)(\text{Flrn}(\varphi)), \mathcal{D}(f)(\text{Flrn}(\varphi'))) \quad \text{by naturality of Flrn} \\ &\leq M \cdot d_{\mathcal{D}(X)}(\text{Flrn}(\varphi), \text{Flrn}(\varphi')) \quad \text{by Lemma 2 (1)} \\ &\stackrel{1}{=} d_{\mathcal{M}[K](X)}(\varphi, \varphi'). \end{aligned}$$

5. The map $\text{acc}: X^K \rightarrow \mathcal{M}[K](X)$ is $\frac{1}{K}$ -Lipschitz since for $\mathbf{y}, \mathbf{y}' \in X^K$,

$$\begin{aligned} d(\text{acc}(\mathbf{y}), \text{acc}(\mathbf{y}')) &= \frac{1}{K} \cdot \bigwedge_{\mathbf{x} \in \text{acc}^{-1}(\text{acc}(\mathbf{y})), \mathbf{x}' \in \text{acc}^{-1}(\text{acc}(\mathbf{y}'))} d_{X^K}(\mathbf{x}, \mathbf{x}') \\ &\leq \frac{1}{K} \cdot d_{X^K}(\mathbf{y}, \mathbf{y}'). \end{aligned}$$

6. For fixed $\varphi, \varphi' \in \mathcal{M}[K](X)$, take arbitrary $\mathbf{x} \in \text{acc}^{-1}(\varphi)$ and $\mathbf{x}' \in \text{acc}^{-1}(\varphi')$. Then:

$$\begin{aligned} d_{\mathcal{D}(X^K)}(\text{arr}(\varphi), \text{arr}(\varphi')) &= d_{\mathcal{D}(X^K)}(\text{arr}(\text{acc}(\mathbf{x})), \text{arr}(\text{acc}(\mathbf{x}'))) \\ &\stackrel{(6)}{=} d_{\mathcal{D}(X^K)}(\text{prm}(\mathbf{x}), \text{prm}(\mathbf{x}')) \\ &\leq d_{X^K}(\mathbf{x}, \mathbf{x}') \quad \text{by Lemma 2 (9)}. \end{aligned}$$

Since this holds for all $\mathbf{x} \in \text{acc}^{-1}(\varphi)$, $\mathbf{x}' \in \text{acc}^{-1}(\varphi')$ we get an inequality $d_{\mathcal{D}(X^K)}(\text{arr}(\varphi), \text{arr}(\varphi')) \leq K \cdot d_{\mathcal{M}[K](X)}(\varphi, \varphi')$, see Definition 3. This inequality is an actual equality since acc , and thus $\mathcal{D}(\text{acc})$, is $\frac{1}{K}$ -Lipschitz:

$$\begin{aligned} d_{\mathcal{M}[K](X)}(\varphi, \varphi') &= d_{\mathcal{D}(\mathcal{M}[K](X))}(1|\varphi, 1|\varphi') \\ &= d_{\mathcal{D}(\mathcal{M}[K](X))}(\mathcal{D}(\text{acc})(\text{arr}(\varphi)), \mathcal{D}(\text{acc})(\text{arr}(\varphi'))) \\ &\leq \frac{1}{K} \cdot d_{\mathcal{D}(X^K)}(\text{arr}(\varphi), \text{arr}(\varphi')) \quad \square \end{aligned}$$

6 Multinomial drawing is isometric

Multinomial draws are of the draw-and-replace kind. This means that a drawn ball is returned to the urn, so that the urn remains unchanged. Thus we may use a distribution $\omega \in \mathcal{D}(X)$ as urn. For a draw size number $K \in \mathbb{N}$, the multinomial distribution $mn[K](\omega) \in \mathcal{D}(\mathcal{M}[K](X))$ on multisets / draws of size K can be defined via accumulated sequences of draws:

$$\begin{aligned} mn[K](\omega) &:= \mathcal{D}(\text{acc})(\omega^K) \\ &= \mathcal{D}(\text{acc})(iid[K](\omega)) \\ &= \sum_{\varphi \in \mathcal{M}[K](X)} (\varphi) \cdot \prod_{x \in X} \omega(x)^{\varphi(x)} | \varphi \rangle. \end{aligned} \tag{10}$$

We recall that $(\varphi) = \frac{K!}{\prod_x \varphi(x)!}$ is the number of sequences that accumulate to a multiset / draw $\varphi \in \mathcal{M}[K](X)$. A basic result from [8, Prop. 3] is that applying frequentist learning to the draws yields the original urn:

$$Flm \gg mn[K](\omega) = \omega. \tag{11}$$

We can now formulate and prove our first isometry result.

Theorem 1. *Let X be an arbitrary metric space (of colours), and $K > 0$ be a positive natural (draw size) number. The multinomial channel*

$$\mathcal{D}(X) \xrightarrow{mn[K]} \mathcal{D}(\mathcal{M}[K](X))$$

is an isometry. This involves the Wasserstein metric (8) for distributions over X on the domain $\mathcal{D}(X)$, and the Wasserstein metric for distributions over multisets of size K , with their Wasserstein metric (9), on the codomain $\mathcal{D}(\mathcal{M}[K](X))$.

Proof. Let distributions $\omega, \omega' \in \mathcal{D}(X)$ be given. The map $mn[K]$ is short since:

$$\begin{aligned}
& d_{\mathcal{D}(\mathcal{M}[K](X))} \left(mn[K](\omega), mn[K](\omega') \right) \\
& \stackrel{(10)}{=} d_{\mathcal{D}(\mathcal{M}[K](X))} \left(\mathcal{D}(\text{acc})(iid[K](\omega)), \mathcal{D}(\text{acc})(iid[K](\omega')) \right) \\
& \leq \frac{1}{K} \cdot d_{\mathcal{D}(X^K)} \left(iid[K](\omega), iid[K](\omega') \right) && \text{by Lemma 3 (5)} \\
& = \frac{1}{K} \cdot K \cdot d_{\mathcal{D}(X)}(\omega, \omega') && \text{by Proposition 1 (2)} \\
& = d_{\mathcal{D}(X)}(\omega, \omega').
\end{aligned}$$

There is also an inequality in the other direction, via:

$$\begin{aligned}
d_{\mathcal{D}(X)}(\omega, \omega') & \stackrel{(11)}{=} d_{\mathcal{D}(X)} \left(Flrn \gg mn[K](\omega), Flrn \gg mn[K](\omega') \right) \\
& \leq d_{\mathcal{D}(\mathcal{M}[K](X))} \left(mn[K](\omega), mn[K](\omega') \right).
\end{aligned}$$

The latter inequality follows from the fact that frequentist learning $Flrn$ is short, see Lemma 3 (1), and that Kleisli extension $Flrn \gg (-)$ is thus short too, see Lemma 2 (6).

□

Example 3. Consider the following two distributions $\omega, \omega' \in \mathcal{D}(\mathbb{N})$.

$$\omega = \frac{1}{3}|0\rangle + \frac{2}{3}|2\rangle \quad \text{and} \quad \omega' = \frac{1}{2}|1\rangle + \frac{1}{2}|2\rangle \quad \text{with} \quad d(\omega, \omega') = \frac{1}{2}.$$

This distance $d(\omega, \omega')$ involves the standard distance on \mathbb{N} , using the optimal coupling $\frac{1}{3}|0, 1\rangle + \frac{1}{6}|2, 1\rangle + \frac{1}{2}|2, 2\rangle \in \mathcal{D}(\mathbb{N} \times \mathbb{N})$.

We take draws of size $K = 3$. There are 10 multisets of size 3 over $\{0, 1, 2\}$:

$$\begin{aligned}
\varphi_1 &= 3|0\rangle & \varphi_2 &= 2|0\rangle + 1|1\rangle & \varphi_3 &= 1|0\rangle + 2|1\rangle & \varphi_4 &= 3|1\rangle \\
\varphi_5 &= 2|0\rangle + 1|2\rangle & \varphi_6 &= 1|0\rangle + 1|1\rangle + 1|2\rangle & \varphi_7 &= 2|1\rangle + 1|2\rangle \\
\varphi_8 &= 1|0\rangle + 2|2\rangle & \varphi_9 &= 1|1\rangle + 2|2\rangle & \varphi_{10} &= 3|2\rangle.
\end{aligned}$$

These multisets occur in the following multinomial distributions of draws of size 3.

$$\begin{aligned}
mn[3](\omega) &= \frac{1}{27}| \varphi_1 \rangle + \frac{2}{9}| \varphi_5 \rangle + \frac{4}{9}| \varphi_8 \rangle + \frac{8}{27}| \varphi_{10} \rangle \\
mn[3](\omega') &= \frac{1}{8}| \varphi_4 \rangle + \frac{3}{8}| \varphi_7 \rangle + \frac{3}{8}| \varphi_9 \rangle + \frac{1}{8}| \varphi_{10} \rangle.
\end{aligned}$$

The optimal coupling $\tau \in \mathcal{D}(\mathcal{M}[3](\mathbb{N}) \times \mathcal{M}[3](\mathbb{N}))$ between these two multinomial distributions is:

$$\begin{aligned}
\tau &= \frac{1}{27} | \varphi_1, \varphi_4 \rangle + \frac{19}{216} | \varphi_5, \varphi_4 \rangle + \frac{1}{8} | \varphi_{10}, \varphi_{10} \rangle + \frac{29}{216} | \varphi_5, \varphi_7 \rangle \\
&\quad + \frac{5}{72} | \varphi_8, \varphi_7 \rangle + \frac{3}{8} | \varphi_8, \varphi_9 \rangle + \frac{37}{216} | \varphi_{10}, \varphi_7 \rangle.
\end{aligned}$$

We compute the distance between the multinomial distributions, using $d_{\mathcal{M}} = d_{\mathcal{M}[3](\mathbb{N})}$.

$$\begin{aligned}
d(mn[3](\omega), mn[3](\omega')) &= \tau \models d_{\mathcal{M}} \\
&= \frac{1}{27} \cdot d_{\mathcal{M}}(\varphi_1, \varphi_4) + \frac{19}{216} \cdot d_{\mathcal{M}}(\varphi_5, \varphi_4) + \frac{1}{8} \cdot d_{\mathcal{M}}(\varphi_{10}, \varphi_{10}) + \frac{29}{216} \cdot d_{\mathcal{M}}(\varphi_5, \varphi_7) \\
&\quad + \frac{5}{72} \cdot d_{\mathcal{M}}(\varphi_8, \varphi_7) + \frac{3}{8} \cdot d_{\mathcal{M}}(\varphi_8, \varphi_9) + \frac{37}{216} \cdot d_{\mathcal{M}}(\varphi_{10}, \varphi_7) \\
&= \frac{1}{27} \cdot 1 + \frac{19}{216} \cdot 1 + \frac{1}{8} \cdot 0 + \frac{29}{216} \cdot \frac{2}{3} + \frac{5}{72} \cdot \frac{2}{3} + \frac{3}{8} \cdot \frac{1}{3} + \frac{37}{216} \cdot \frac{2}{3} = \frac{1}{2}.
\end{aligned}$$

As predicted in Theorem 1, this distance coincides with the distance $d(\omega, \omega') = \frac{1}{2}$ between the original urn distributions. One sees that the computation of the distance between the draw distributions is more complex, involving ‘Wasserstein over Wasserstein’.

7 Hypergeometric drawing is isometric

We start with some preparatory observations on probabilistic projection and drawing.

Lemma 4. *For a metric space X and a number K , consider the probabilistic projection-delete PD and probabilistic draw-delete DD channels.*

$$X^{K+1} \xrightarrow{PD} \mathcal{D}(X^K) \qquad \mathcal{M}[K+1](X) \xrightarrow{DD} \mathcal{D}(\mathcal{M}[K](X))$$

They are defined via deletion of elements from sequences and from multisets:

$$\begin{aligned} PD(x_1, \dots, x_{K+1}) &:= \sum_{1 \leq i \leq K+1} \frac{1}{K+1} |x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_{K+1}\rangle \\ DD(\psi) &:= \sum_{x \in \text{supp}(\psi)} \frac{\psi(x)}{K+1} |\psi - 1|x\rangle \\ &= \sum_{x \in \text{supp}(\psi)} \text{Flrn}(\psi)(x) |\psi - 1|x\rangle. \end{aligned}$$

Then:

1. $\langle \text{acc} \rangle \circ PD = DD \circ \langle \text{acc} \rangle$;
2. $\text{Flrn} \gg DD(\psi) = \text{Flrn}(\psi)$;
3. PD is $\frac{K}{K+1}$ -Lipschitz, and thus short;
4. DD is an isometry.

Proof. The first point is easy and the second one is [8, Lem. 5 (ii)].

3. For $\mathbf{x}, \mathbf{y} \in X^{K+1}$, via Lemma 2 (8) and (4),

$$\begin{aligned} d(PD(\mathbf{x}), PD(\mathbf{y})) &= d \left(\sum_{1 \leq i \leq K+1} \frac{1}{K+1} |x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_{K+1}\rangle, \right. \\ &\quad \left. \sum_{1 \leq i \leq K+1} \frac{1}{K+1} |y_1, \dots, y_{i-1}, y_{i+1}, \dots, y_{K+1}\rangle \right) \\ &\leq \sum_{1 \leq i \leq K+1} \frac{1}{K+1} \cdot d \left(1|x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_{K+1}\rangle, \right. \\ &\quad \left. 1|y_1, \dots, y_{i-1}, y_{i+1}, \dots, y_{K+1}\rangle \right) \\ &= \sum_{1 \leq i \leq K+1} \frac{1}{K+1} \cdot d_{X^K} \left((x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_{K+1}), \right. \\ &\quad \left. (y_1, \dots, y_{i-1}, y_{i+1}, \dots, y_{K+1}) \right) \\ &= \sum_{1 \leq i \leq K+1} \frac{1}{K+1} \cdot K \cdot d_X(x_i, y_i) \\ &\stackrel{(7)}{=} \frac{K}{K+1} \cdot d_{X^{K+1}}(\mathbf{x}, \mathbf{y}). \end{aligned}$$

4. Via item 1 we get:

$$\langle \text{acc} \rangle \circ PD \circ \text{arr} = DD \circ \langle \text{acc} \rangle \circ \text{arr} = DD \circ \text{unit} = DD. \quad (*)$$

Now we can show that DD is short: for $\psi, \psi' \in \mathcal{M}[K+1](X)$

$$\begin{aligned} & d_{\mathcal{D}(\mathcal{M}[K](X))}(DD(\psi), DD(\psi')) \\ & \stackrel{(*)}{=} d_{\mathcal{D}(\mathcal{M}[K](X))}\left(\mathcal{D}(\text{acc})(PD \gg \text{arr}(\psi)), \mathcal{D}(\text{acc})(PD \gg \text{arr}(\psi'))\right) \\ & \leq \frac{1}{K} \cdot d_{\mathcal{D}(X^K)}(PD \gg \text{arr}(\psi), PD \gg \text{arr}(\psi')) \\ & \leq \frac{1}{K} \cdot \frac{K}{K+1} \cdot d_{\mathcal{D}(X^{K+1})}(\text{arr}(\psi), \text{arr}(\psi')) \\ & = \frac{1}{K+1} \cdot (K+1) \cdot d_{\mathcal{M}[K+1](X)}(\psi, \psi') \\ & = d_{\mathcal{M}[K+1](X)}(\psi, \psi'). \end{aligned}$$

For the reverse inequality we use item 2 and the fact that $Flrn$ is a short:

$$\begin{aligned} & d_{\mathcal{D}(\mathcal{M}[K](X))}(DD(\psi), DD(\psi')) \\ & \geq d_{\mathcal{D}(\mathcal{M}[K](X))}(Flrn \gg DD(\psi), Flrn \gg DD(\psi')) \\ & = d_{\mathcal{D}(X)}(Flrn(\psi), Flrn(\psi')) \\ & = d_{\mathcal{M}[K+1](X)}(\psi, \psi'). \quad \square \end{aligned}$$

The hypergeometric channel $hg[K]: \mathcal{M}[L](X) \rightarrow \mathcal{D}(\mathcal{M}[K](X))$, for urn size $L \geq K$, where K is the draw size, is an iteration of draw-delete's, see [8, Thm. 6]:

$$hg[K](v) := \underbrace{DD \circ \dots \circ DD}_{L-K \text{ times}} = \sum_{\varphi \in \mathcal{M}[K](X), \varphi \leq v} \frac{\binom{v}{\varphi}}{\binom{L}{K}} |\varphi\rangle, \quad (12)$$

where $\binom{v}{\varphi} := \prod_{x \in X} \binom{v(x)}{\varphi(x)}$.

Theorem 2. *The hypergeometric channel $hg[K]: \mathcal{M}[L](X) \rightarrow \mathcal{D}(\mathcal{M}[K](X))$ defined in (12), for $L \geq K$, is an isometry.*

Proof. We see in (12) that $hg[K]$ is a (channel) iteration of isometries DD , and thus of short maps; hence it is short itself. Via iterated use of Lemma 4 (2) we get $Flrn \gg hg[K](\psi) = Flrn(\psi)$. This gives the inequality in the other direction, like in the proof of Lemma 4 (2):

$$\begin{aligned} d_{\mathcal{M}[K+1](X)}(\psi, \psi') & = d_{\mathcal{D}(X)}(Flrn(\psi), Flrn(\psi')) \\ & = d_{\mathcal{D}(\mathcal{M}[K](X))}(Flrn \gg hg[K](\psi), Flrn \gg hg[K](\psi')) \\ & \leq d_{\mathcal{D}(\mathcal{M}[K](X))}(hg[K](\psi), hg[K](\psi')). \quad \square \end{aligned}$$

The very beginning of this paper contains an illustration of this result, for urns over the set of colours $C = \{R, G, B\}$, considered as a discrete metric space.

8 Pólya drawing is isometric

Hypergeometric distributions use the draw-delete mode: a drawn ball is removed from the urn. The less well-known Pólya draws [7] use the draw-add mode. This means that a drawn ball is returned to the urn, together with another ball of the same colour (as the drawn ball). Thus, with hypergeometric draws the urn decreases in size, so that only finitely many draws are possible, whereas with Pólya draws the urn grows in size, and the drawing may be repeated arbitrarily many times. As a result, for Pólya distributions we do not need to impose restrictions on the size K of draws. We do have to restrict draws from urn v to multisets $\varphi \in \mathcal{M}[K](X)$ with $\text{supp}(\varphi) \subseteq \text{supp}(v)$ since we can only draw balls of colours that are in the urn. Pólya distributions are formulated in terms of multi-choose binomials $\binom{n}{m} := \binom{n+m-1}{m} = \frac{(n+m-1)!}{m!(n-1)!}$, for $n > 0$. This multi-choose number $\binom{n}{m}$ is the number of multisets of size m over a set with n elements, see [9,10] for details.

$$pl[K](v) := \sum_{\varphi \in \mathcal{M}[K](X), \text{supp}(\varphi) \subseteq \text{supp}(v)} \frac{\binom{v}{\varphi}}{\binom{L}{K}} |\varphi\rangle, \quad (13)$$

where $\binom{v}{\varphi} := \prod_{x \in \text{supp}(v)} \binom{v(x)}{\varphi(x)}$.

Theorem 3. *Each Pólya channel $pl[K]: \mathcal{M}[L](X) \rightarrow \mathcal{D}(\mathcal{M}[K](X))$, for urn and draw sizes $L > 0, K > 0$, is an isometry.*

Proof. One inequality follows by exploiting the equation $Flrn \gg pl[K](\psi) = Flrn(\psi)$ like in previous sections. The reverse inequality, for shortness of $pl[K]$, involves an auxiliary draw-store-add channel of the form:

$$\mathcal{M}[L](X) \times \mathcal{M}[N](X) \xrightarrow{DSA} \mathcal{D}(\mathcal{M}[L](X) \times \mathcal{M}[N+1](X))$$

It is defined as:

$$\begin{aligned} DSA(v, \varphi) &:= \sum_{x \in \text{supp}(v+\varphi)} Flrn(v+\varphi)(x) |v, \varphi+1|x\rangle \\ &= 1|v\rangle \otimes \left(\sum_{x \in \text{supp}(v+\varphi)} Flrn(v+\varphi)(x) |\varphi+1|x\rangle \right). \end{aligned}$$

With some effort one shows that this channel DSA is short and that the Pólya channel can be expressed via iterated draw-store-add's, namely as:

$$pl[K](v) = \mathcal{D}(\pi_2) \left(\underbrace{(DSA \circ \dots \circ DSA)}_{K \text{ times}}(v, \mathbf{0}) \right),$$

where $\mathbf{0} \in \mathcal{M}[0](X)$ is the empty multiset. This makes the Pólya channel $pl[K]$ short, and thus an isometry. \square

We illustrate that the Pólya channel is an isometry.

Example 4. We take as space of colours $X = \{0, 10, 50\} \subseteq \mathbb{N}$ with two urns:

$$v_1 = 3|0\rangle + 1|10\rangle \quad v_2 = 1|0\rangle + 2|10\rangle + 1|50\rangle.$$

The distance between these urns is 15, via the optimal coupling $1|0, 0\rangle + 2|0, 10\rangle + 1|10, 50\rangle$, yielding $\frac{1}{4} \cdot (0 - 0) + \frac{1}{2} \cdot (10 - 0) + \frac{1}{4} \cdot (50 - 10) = 5 + 10 = 15$.

We look at Pólya draws of size $K = 2$. This gives distributions:

$$\begin{aligned} pl[2](v_1) &= \frac{3}{5} |2|0\rangle\rangle + \frac{3}{10} |1|0\rangle + 1|10\rangle\rangle + \frac{1}{10} |2|10\rangle\rangle \\ pl[2](v_2) &= \frac{1}{10} |2|0\rangle\rangle + \frac{1}{5} |1|0\rangle + 1|10\rangle\rangle + \frac{3}{10} |2|10\rangle\rangle + \frac{1}{10} |1|0\rangle + 1|50\rangle\rangle \\ &\quad + \frac{1}{5} |1|10\rangle + 1|50\rangle\rangle + \frac{1}{10} |2|50\rangle\rangle \end{aligned}$$

We compute the distance between these two distributions via the last formulation in (8), using the optimal short factor $p: \mathcal{M}[2](X) \rightarrow \mathbb{R}_{\geq 0}$ given by:

$$\begin{array}{lll} p(2|0\rangle) = 0 & p(1|0\rangle + 1|10\rangle) = 5 & p(2|10\rangle) = 10 \\ p(1|0\rangle + 1|50\rangle) = 25 & p(1|10\rangle + 1|50\rangle) = 30 & p(2|50\rangle) = 50. \end{array}$$

Then:

$$\begin{aligned} pl[2](v_1) \models p &= \frac{3}{5} \cdot 0 + \frac{3}{10} \cdot 5 + \frac{1}{10} \cdot 10 = \frac{5}{2} \\ pl[2](v_2) \models p &= \frac{1}{10} \cdot 0 + \frac{1}{5} \cdot 5 + \frac{3}{10} \cdot 10 + \frac{1}{10} \cdot 25 + \frac{1}{5} \cdot 30 + \frac{1}{10} \cdot 50 = \frac{35}{2}. \end{aligned}$$

As predicted by Theorem 3, the distance between the Pólya distributions then coincides with the distance between the urns:

$$\begin{aligned} d(pl[2](v_1), pl[2](v_2)) &= |pl[2](v_1) \models p - pl[2](v_2) \models p| \\ &= \frac{35}{2} - \frac{5}{2} = 15 = d(v_1, v_2). \end{aligned}$$

9 Conclusions

Category theory provides a fresh look at the area of probability theory, see *e.g.* [5] or [10] for an overview. Its perspective allows one to formulate and prove new results. This paper demonstrates that draw operations, viewed as (Kleisli) maps, are incredibly well-behaved: they preserve Wasserstein distances. Such distances on urns filled with coloured balls are relatively simple, starting from a ‘ground’ metric on the set of colours. But on draw distributions, the distances involve Wasserstein-over-Wasserstein. This paper concentrates on drawing from an urn. A natural question is whether other probabilistic operations, as Kleisli maps, preserve distance. This is a topic for further investigation.

Acknowledgments

Thanks are due to the anonymous reviewers for their detailed comments that improved an earlier version of this work.

References

1. F. van Breugel. An introduction to metric semantics: operational and denotational models for programming and specification languages. *Theor. Comp. Sci.*, 258(1-2):1–98, 2001. doi:10.1016/S0304-3975(00)00403-5.
2. H. Brezis. Remarks on the Monge-Kantorovich problem in the discrete setting. *Comptes Rendus Mathématique*, 356(2):207–213, 2018. doi:10.1016/j.crma.2017.12.008.
3. Y. Deng and W. Du. The Kantorovich metric in computer science: A brief survey. In C. Baier and A. di Pierro, editors, *Quantitative Aspects of Programming Languages*, number 253(3) in *Elect. Notes in Theor. Comp. Sci.*, pages 73–82. Elsevier, Amsterdam, 2009. doi:10.1016/j.entcs.2009.10.006.
4. J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. Metrics for labelled Markov processes. *Theor. Comp. Sci.*, 318:232–354, 2004.
5. T. Fritz. A synthetic approach to Markov kernels, conditional independence, and theorems on sufficient statistics. *Advances in Math.*, 370:107239, 2020. doi:10.1016/J.AIM.2020.107239.
6. T. Fritz and P. Perrone. A probability monad as the colimit of spaces of finite samples. *Theory and Appl. of Categories*, 34(7):170–220, 2019. doi:10.48550/arXiv.1712.05363.
7. F. Hoppe. Pólya-like urns and the Ewens’ sampling formula. *Journ. Math. Biology*, 20:91–94, 1984. doi:10.1007/BF00275863.
8. B. Jacobs. From multisets over distributions to distributions over multisets. In *Logic in Computer Science*. IEEE, Computer Science Press, 2021. doi:10.1109/lics52264.2021.9470678.
9. B. Jacobs. Urns & tubes. *Compositionality*, 4(4), 2022. doi:10.32408/compositionality-4-4.
10. B. Jacobs. Structured probabilistic reasoning. Book, in preparation, see <http://www.cs.ru.nl/B.Jacobs/PAPERS/ProbabilisticReasoning.pdf>, 2023.
11. B. Jacobs and A. Westerbaan. Distances between states and between predicates. *Logical Methods in Comp. Sci.*, 16(1), 2020. See <https://lmcs.episciences.org/6154>.
12. L. Kantorovich and G. Rubinshtein. On a space of totally additive functions. *Vestnik Leningrad Univ.*, 13:52–59, 1958.
13. J. Matoušek and B. Gärtner. *Understanding and Using Linear Programming*. Springer Verlag, Berlin, 2006. doi:10.1007/978-3-540-30717-4.
14. Y. Rubner, C. Tomasi, and L. Guibas. The Earth Mover’s Distance as a metric for image retrieval. *Int. Journ. of Computer Vision*, 40:99–121, 2000. doi:10.1023/A:1026543900054.
15. C. Villani. *Optimal Transport — Old and New*. Springer, Berlin Heidelberg, 2009. doi:10.1007/978-3-540-71050-9.
16. G. Wyszecki and W. Stiles. *Color Science: Concepts and Methods, Quantitative Data and Formulae*. Wiley, 1982.