Approximating the total variation distance between two product distributions

Weiming Feng (UC Berkeley)

Based on joint works with



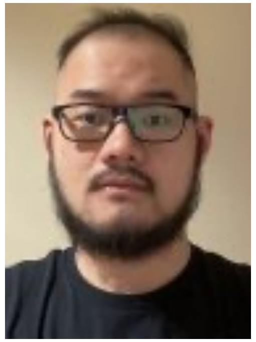

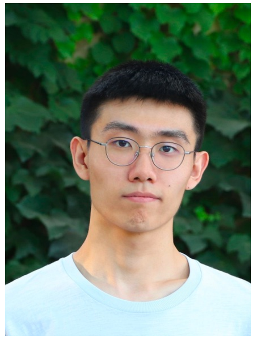

Heng Guo
Edinburgh

Mark Jerrum
QMUL

Jiaheng Wang
Edinburgh

Tianren Liu
Peking

Liqiang Liu
Peking

Randomised Algorithm

Deterministic Algorithm

Probabilistic Circuits and Logic Workshop Berkeley, CA, Oct 20, 2023

Difference between two distributions

Data: two distributions P and Q over state space Ω

Question: how to measure the difference between P and Q

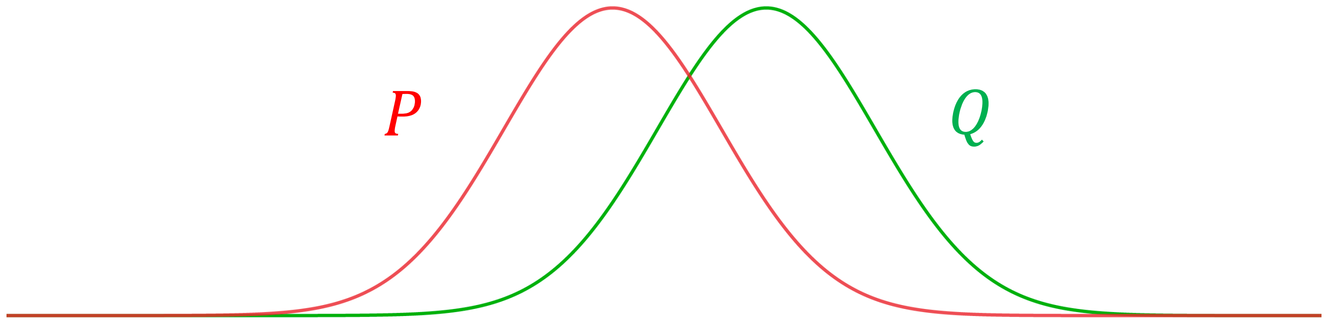

- Total variation distance (TV distance): $d_{TV}(P,Q) = \frac{1}{2}\sum_{x\in\Omega}|P(x) - Q(x)|$
- KL-Divergence (relative entropy): $D_{KL}(P||Q) = \sum_{x\in\Omega} P(x)\log\frac{P(x)}{Q(x)}$
- χ^2-divergence: $D_{\chi^2}(P||Q) = \left(\sum_{x\in\Omega}\frac{P^2(X)}{Q(x)}\right) - 1$

Total Variation distance

Total variation (TV) distance between $\mathbb{P}$ and $\mathbb{Q}$ over state space Ω

$$d_{TV}(\mathbb{P},\mathbb{Q}) = \frac{1}{2}\sum |\mathbb{P}(x) - \mathbb{Q}(x)| = \max_{S\subseteq\Omega} |\mathbb{P}(S) - \mathbb{Q}(S)|$$

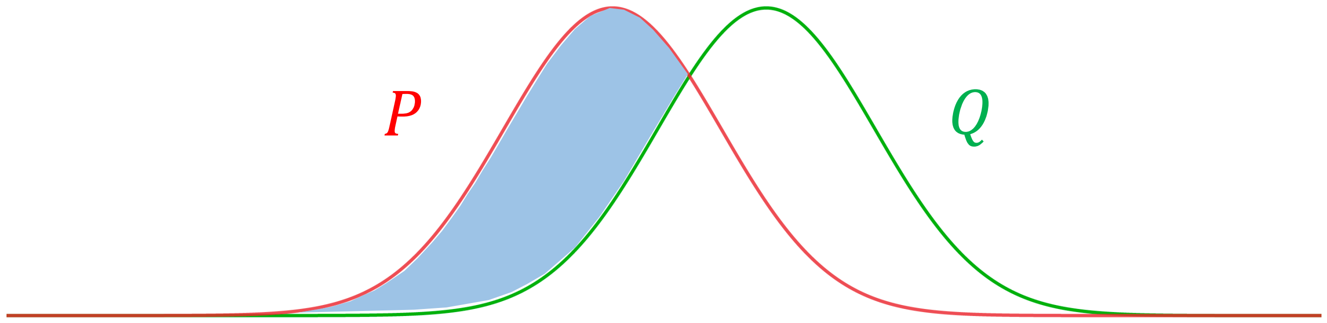

Properties of TV distance

- metric (triangle inequality)
- bounded
- data processing inequality
- various characterisations

Applications of TV distance

- property testing
- Markov chain mixing time
- approximate algorithms
- learning algorithms

Compute TV distance

[Bhattacharyya, Gayen, Meel, Myrisiotis, Pavan, Vinodchandran, 2022]

- **Input:** descriptions of two distributions $\mathbb{P}, \mathbb{Q}$ over Ω
- **Output:** the total variation distance between $\mathbb{P}$ and $\mathbb{Q}$

Trivial algorithm: enumerate all $x \in \Omega$ and add $\frac{1}{2}|\mathbb{P}(x) - \mathbb{Q}(x)|$ together

Challenge: $\mathbb{P}$ and $\mathbb{Q}$ have ***succinct descriptions***

- $|\Omega|$ can be ***exponentially large*** w.r.t. the size of input

Examples: probabilistic graphical models, spin systems.

Product distribution

Product distribution $\mathbb{P}$ over $[s]^n$

$$\mathbb{P} = \mathbb{P}_1 \times \mathbb{P}_2 \times \cdots \times \mathbb{P}_n$$

$$\forall X, \mathbb{P}(X) = \prod_{i=1}^{n} \mathbb{P}_i(X_i)$$

$\mathbb{P}_i$ is a distribution over $[s] = \{1,2,\dots,s\}$

Random sample $X = (X_1, X_2, \dots, X_n) \sim \mathbb{P}$

$X \in [s]^n$: n-dimensional random vector

$X_i \in [s]$: independent sample from $\mathbb{P}_i$

$\mathbb{P}$ can be described by $\{\, \mathbb{P}_i: [s] \to [0,1] \mid 1 \le i \le n \,\}$

description size
sn

state space size
s^n

Compute TV distance between product distributions

[Bhattacharyya, Gayen, Meel, Myrisiotis, Pavan, Vinodchandran, 2022]

- **Input:** distributions $\{\mathbb{P}_i, \mathbb{Q}_i | 1 \leq i \leq n\}$ specifying $\mathbb{P}$ and $\mathbb{Q}$ over $[s]^n$
- **Output:** the total variation distance between $\mathbb{P}$ and $\mathbb{Q}$

Theorem [BGMMPV22]: the problem is **#P-complete** even for Boolean case $(s = 2)$

FPTAS (Full Poly-time Approximation Scheme)

A *deterministic* algorithm outputs a $\hat{d}$ in time $\mathrm{poly}(n, s, 1/\epsilon)$

$$(1-\epsilon)d_{TV}(\mathbb{P},\mathbb{Q}) \leq \hat{d} \leq (1+\epsilon)d_{TV}(\mathbb{P},\mathbb{Q})$$

FPRAS (Full Poly-time Randomised Approximation Scheme)

A *randomised* algorithm outputs a random $\hat{d}$ in time $\mathrm{poly}(n, s, 1/\epsilon)$

$$\Pr\left[(1-\epsilon)d_{TV}(\mathbb{P},\mathbb{Q}) \leq \hat{d} \leq (1+\epsilon)d_{TV}(\mathbb{P},\mathbb{Q})\right] \geq 2/3$$

Previous results

Theorem [BGMMPV22] ***FPTAS/FPRAS*** exists for product distributions $\mathbb{P}, \mathbb{Q}$ such that

- $\mathbb{P}$ and $\mathbb{Q}$ are *Boolean* distributions ($s = 2$)
- $\mathbb{Q}$ has *constant number* of distinct marginals (e.g. uniform distribution over $\{0,1\}^n$)

Theorem [BGMMPV22] ***FPRAS*** exists for product distributions $\mathbb{P}, \mathbb{Q}$ such that

- $\mathbb{P}$ and $\mathbb{Q}$ are *Boolean* distributions ($s = 2$)
- $\forall i \in [n], \underbrace{\mathbb{P}_i(1) \geq \mathbb{Q}_i(1)}_{\textit{break symmetry}}$ and $\underbrace{\mathbb{P}_i(1) \geq 1/2}_{\textit{marginal lower bound}}$

Open problem [BGMMPV22]**:**

Do FPTAS/FPRAS exist for ***general*** product distributions?

Our results [F., Guo, Jerrum, Wang 2023] [F., Liu, Liu 2023] **:**

FPTAS/FPRAS exist for ***general*** product distributions

Product distributions $\mathbb{P}, \mathbb{Q}$ over $[s]^n$ and error bound $0 < \epsilon < 1$

- FPTAS running time: $\tilde{O}\left(\frac{sn^2}{\epsilon} \log \frac{1}{d_{TV}(\mathbb{P},\mathbb{Q})}\right)$
- FPRAS running time : $\tilde{O}\left(\frac{sn^2}{\epsilon^2}\right)$

Extension: Markov chains [F., Liu, Liu 2023]

- distributions π_1, π_2 and transition Matrices M_1, M_2 over state space $[s]$
- approximate $d_{TV}((X_k)_{k=1}^n, (Y_k)_{k=1}^n)$ such that
  - $X_1 \sim \pi_1$ and $X_k \sim M_1(X_{k-1},\cdot)$ / $Y_1 \sim \pi_2$ and $Y_k \sim M_2(Y_{k-1},\cdot)$

FPTAS exists for TV-distance between Markov chains

A natural estimator

Total variation (TV) distance between $\mathbb{P}$ and $\mathbb{Q}$ over state space Ω

$$d_{TV}(\mathbb{P},\mathbb{Q}) = \frac{1}{2}\sum_{x\in\Omega} |\mathbb{P}(x) - \mathbb{Q}(x)| = \sum_{x\in\Omega:\mathbb{Q}(X)>\mathbb{P}(X)} |\mathbb{Q}(x) - \mathbb{P}(x)| = \sum_{x\in\Omega:\mathbb{Q}(X)>\mathbb{P}(X)} \mathbb{Q}(X)\left(1 - \frac{\mathbb{P}(X)}{\mathbb{Q}(X)}\right)$$

Ratio $R \sim \mathbb{R} = (\mathbb{P}||\mathbb{Q})$

$$R = \frac{\mathbb{P}(X)}{\mathbb{Q}(X)}, \qquad \text{where } X \sim \mathbb{Q}$$

$$d_{TV}(\mathbb{P},\mathbb{Q}) = \mathbb{E}[\max(0,1-R)]$$

- sample R independent
- take average of $\max(0,1-R)$

unbiased estimator of $d_{TV}(\mathbb{P},\mathbb{Q})$

- Approximate the TV distance with ***additive*** error $\hat{d} \in d_{TV}(\mathbb{P},\mathbb{Q}) \pm \epsilon$
- ***Relative*-error** approximation requires many samples because $d_{TV}(\mathbb{P},\mathbb{Q})$ can be exponentially small

$\hat{d} \in (1 \pm \epsilon) d_{TV}(\mathbb{P},\mathbb{Q})$

TV distance and coupling

- **Distributions**: P and Q over the domain Ω
- **Coupling**: a joint distribution $(X, Y) \in \Omega \times \Omega$ such that $X \sim P$ and $Y \sim Q$

x	0	1
$P(x)$	$1/2$	$1/2$

y	0	1
$Q(y)$	$1/3$	$2/3$

Example: Independent Coupling

Sample $X \sim P$ and $Y \sim Q$ independently

$$\Pr[X \neq Y] = \frac{1}{2}$$

Can we make this prob. smaller?

TV distance and coupling

- **Distributions**: P and Q over the domain Ω
- **Coupling**: a joint distribution $(X, Y) \in \Omega \times \Omega$ such that $X \sim P$ and $Y \sim Q$

x	0	1
$P(x)$	$1/2$	$1/2$

y	0	1
$Q(y)$	$1/3$	$2/3$

Example: Optimal Coupling

- Sample $r \in (0,1)$ uniformly at random
- Let $X = 0$ iff $r < P(0)$
- Let $Y = 0$ iff $r < Q(0)$

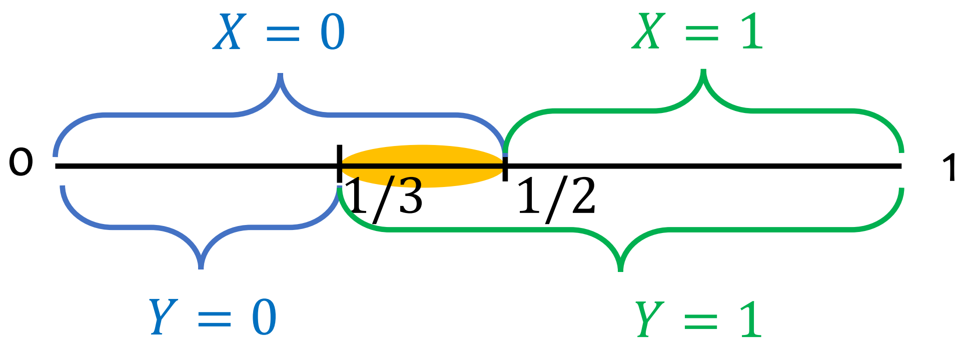

TV distance and coupling

- **Distributions**: P and Q over the domain Ω
- **Coupling**: a joint distribution $(X, Y) \in \Omega\times\Omega$ such that $X \sim P$ and $Y \sim Q$

x	0	1
$P(x)$	$1/2$	$1/2$

y	0	1
$Q(y)$	$1/3$	$2/3$

Example: Optimal Coupling

- Sample $r \in (0,1)$ uniformly at random
- Let $X = 0$ iff $r < P(0)$
- Let $Y = 0$ iff $r < Q(0)$

$$\Pr[X \neq Y] = \frac{1}{6} = d_{TV}(P, Q)$$

TV distance and coupling

- **Distributions**: $\mathbb{P}$ and $\mathbb{Q}$ over the domain Ω
- **Coupling**: a joint $(X, Y) \in \Omega \times \Omega$ such that $X \sim \mathbb{P}$ and $Y \sim \mathbb{Q}$

Coupling Inequality (Coupling Lemma)

$$\forall \text{coupling } (X, Y), \qquad d_{TV}(\mathbb{P}, \mathbb{Q}) \leq \Pr[X \neq Y]$$

$$\exists \text{optimal coupling } (X, Y), \qquad d_{TV}(\mathbb{P}, \mathbb{Q}) = \Pr[X \neq Y]$$

The optimal coupling may ***not*** be unique

Given two **product distributions** $\mathbb{P}, \mathbb{Q}$ over $[s]^n$,

what is their optimal coupling?

A ***greedy*** coupling $(X, Y) = ((X_1, X_2, \dots, X_n), (Y_1, Y_2, \dots, Y_n))$ of $\mathbb{P}, \mathbb{Q}$

each (X_i, Y_i) is coupled ***optimally*** and ***independently***

- Greedy coupling is **not** an optimal coupling

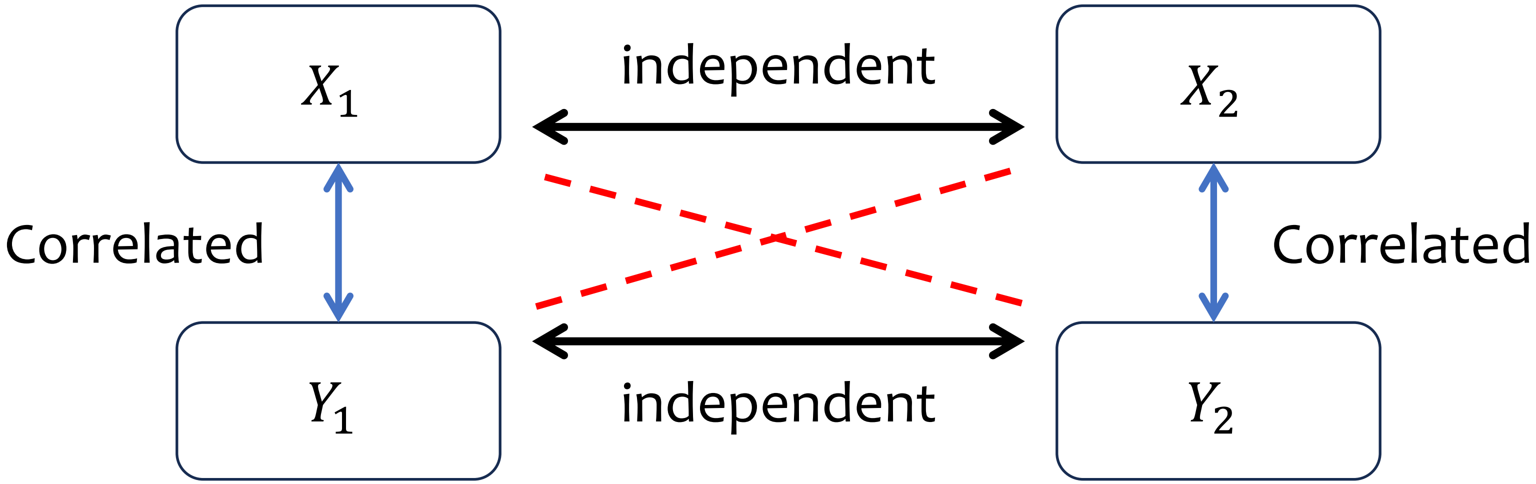

Optimal coupling can utilise the correlations of (X_1, Y_2) and (Y_1, X_2)

A ***greedy*** coupling $(X, Y) = ((X_1, X_2, \dots, X_n), (Y_1, Y_2, \dots, Y_n))$ of $\mathbb{P}, \mathbb{Q}$

each (X_i, Y_i) is coupled ***optimally*** and ***independently***

- Greedy coupling is **not** an optimal coupling
- Greedy coupling can **approximate** the optimal coupling

$$d_{TV}(\mathbb{P}, \mathbb{Q}) \le \Pr_{\text{Greedy}}[X \ne Y] \le n d_{TV}(\mathbb{P}, \mathbb{Q})$$

Proof.

$$\Pr_{\text{Greedy}}[X \ne Y] \le \sum_{i=1}^{n} \Pr[X_i \ne Y_i] = \sum_{i=1}^{n} d_{TV}(\mathbb{P}_i, \mathbb{Q}_i) \le n d_{TV}(\mathbb{P}, \mathbb{Q})$$

local optimal coupling

A ***greedy*** coupling $(X, Y) = ((X_1, X_2, \dots, X_n), (Y_1, Y_2, \dots, Y_n))$ of $\mathbb{P}, \mathbb{Q}$

each (X_i, Y_i) is coupled ***optimally*** and ***independently***

- Greedy coupling is **not** an optimal coupling
- Greedy coupling can **approximate** the optimal coupling

$$d_{TV}(\mathbb{P}, \mathbb{Q}) \le \Pr_{\text{greedy}}[X \ne Y] \le n d_{TV}(\mathbb{P}, \mathbb{Q})$$

- Discrepancy of greedy coupling can be computed ***efficiently***

$$\Pr_{\text{greedy}}[X \ne Y] = 1 - \Pr_{\text{greedy}}[X = Y] = 1 - \prod_{i=1}^{n} \left(1 - d_{TV}(\mathbb{P}_i, \mathbb{Q}_i)\right)$$

- Our ideal: estimate $\frac{\Pr_{\text{opt}}[X \ne Y]}{\Pr_{\text{greedy}}[X \ne Y]} = \frac{d_{TV}(\mathbb{P}, \mathbb{Q})}{\Pr_{\text{greedy}}[X \ne Y]} \ge \frac{1}{n}$

Our Estimator [F., Guo, Jerrum, Wang 2023]

- π: the distribution of X in the greedy coupling conditional on $X \neq Y$

$$\forall \sigma \in [s]^n, \qquad \pi(\sigma) = \Pr_{\text{greedy}}[X = \sigma \mid X \neq Y]$$

- f: a function $[s]^n \to \mathbb{R}_{>0}$ such that

$$\forall \sigma \in [s]^n, \qquad f(\sigma) = \frac{\Pr_{\text{opt}}[X = \sigma \wedge X \neq Y]}{\Pr_{\text{greedy}}[X = \sigma \wedge X \neq Y]} = \frac{\max\{0, \mathbb{P}(\sigma) - \mathbb{Q}(\sigma)\}}{\Pr_{\text{greedy}}[X = \sigma \wedge X \neq Y]}$$

- Estimator: $f(\sigma)$ where $\sigma \sim \pi$

Properties of the estimator

- **Correct expectation**

$$\mathbb{E}_{\sigma \sim \pi}[f(\sigma)] = \frac{\Pr_{\text{opt}}[X \neq Y]}{\Pr_{\text{greedy}}[X \neq Y]} = \frac{d_{TV}(\mathbb{P}, \mathbb{Q})}{\Pr_{\text{greedy}}[X \neq Y]} \geq \frac{1}{n}$$

- **Low variance**

$$\mathrm{Var}_{\sigma \sim \pi}[f(\sigma)] \leq 1$$

$$\forall \sigma \in [s]^V, \qquad \Pr_{\text{opt}}[X = \sigma \wedge X \neq Y] \leq \Pr_{\text{greedy}}[X = \sigma \wedge X \neq Y]$$

$\Rightarrow$

$$\forall \sigma \in [s]^V, \qquad 0 \leq f(\sigma) \leq 1$$

Our Estimator [F., Guo, Jerrum, Wang 2023]

- π: the distribution of X in the greedy coupling conditional on $X \neq Y$

$$\forall \sigma \in [s]^n, \qquad \pi(\sigma) = \Pr_{\text{greedy}}[X = \sigma \mid X \neq Y]$$

- f: a function $[s]^V \to \mathbb{R}_{>0}$ such that

$$\forall \sigma \in [s]^n, \qquad f(\sigma) = \frac{\Pr_{\text{opt}}[X = \sigma \wedge X \neq Y]}{\Pr_{\text{greedy}}[X = \sigma \wedge X \neq Y]} = \frac{\max\{0, \mathbb{P}(\sigma) - \mathbb{Q}(\sigma)\}}{\Pr_{\text{greedy}}[X = \sigma \wedge X \neq Y]}$$

- Estimator: $f(\sigma)$ where $\sigma \sim \pi$

Properties of the estimator

- **Correct expectation**

$$\mathbb{E}_{\sigma\sim\pi}[f(\sigma)] = \frac{\Pr_{\text{opt}}[X \neq Y]}{\Pr_{\text{greedy}}[X \neq Y]} = \frac{d_{TV}(\mathbb{P}, \mathbb{Q})}{\Pr_{\text{greedy}}[X \neq Y]} = R \geq \frac{1}{n}$$

- **Low variance**

$$\mathrm{Var}_{\sigma\sim\pi}[f(\sigma)] \leq 1$$

$T = O\left(\frac{n}{\epsilon^2}\right)$ samples

- Sample $\sigma_1, \sigma_2, \ldots \sigma_T$ independently from π
- Return $\frac{1}{T}\left(f(\sigma_1) + f(\sigma_2) + \cdots + f(\sigma_n)\right)$

Our Estimator [F., Guo, Jerrum, Wang 2023]

- π: the distribution of X in the greedy coupling conditional on $X \neq Y$

$$\forall \sigma \in [s]^n, \qquad \pi(\sigma) = \Pr_{\text{greedy}}[X = \sigma \mid X \neq Y]$$

- f: a function $[s]^V \to \mathbb{R}_{>0}$ such that

$$\forall \sigma \in [s]^n, \qquad f(\sigma) = \frac{\Pr_{\text{opt}}[X = \sigma \wedge X \neq Y]}{\Pr_{\text{greedy}}[X = \sigma \wedge X \neq Y]} = \frac{\max\{0, \mathbb{P}(\sigma) - \mathbb{Q}(\sigma)\}}{\Pr_{\text{greedy}}[X = \sigma \wedge X \neq Y]}$$

- Estimator: $f(\sigma)$ where $\sigma \sim \pi$

Properties of the estimator

- **Correct expectation**

$$\mathbb{E}_{\sigma \sim \pi}[f(\sigma)] = \frac{\Pr_{\text{opt}}[X \neq Y]}{\Pr_{\text{greedy}}[X \neq Y]} = \frac{d_{TV}(\mathbb{P}, \mathbb{Q})}{\Pr_{\text{greedy}}[X \neq Y]} = R \geq \frac{1}{n}$$

- **Low variance**

$$\mathrm{Var}_{\sigma \sim \pi}[f(\sigma)] \leq 1$$

- **Efficient computation**
  - a random sample of $\sigma \sim \pi$ can be generated in time $O(n)$
  - given any $\sigma \in \{0,1\}^n$, $f(\sigma)$ can be computed in time $O(n)$

$T = O\left(\frac{n}{\epsilon^2}\right)$ samples

Ratio and deterministic algorithm

Ratio $R \sim \mathbb{R} = (\mathbb{P}||\mathbb{Q})$

$$R = \frac{\mathbb{P}(X)}{\mathbb{Q}(X)}, \qquad \text{where } X \sim Q$$

$$d_{TV}(\mathbb{P},\mathbb{Q}) = d_{TV}(\mathbb{R}) = \mathbb{E}[\max(0,1-R)]$$

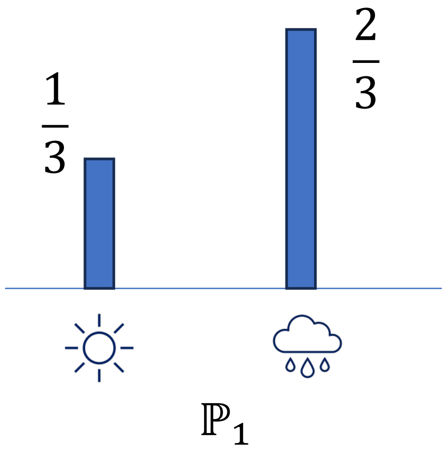

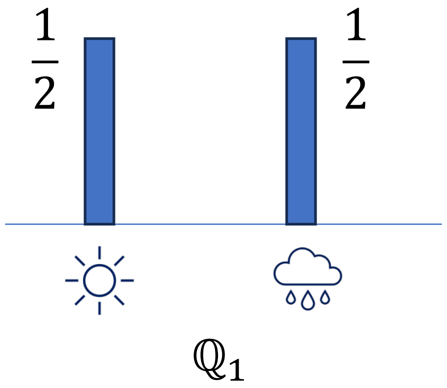

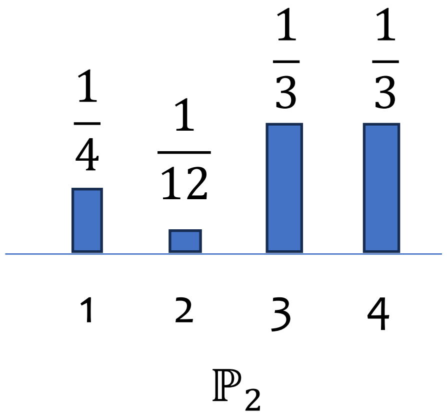

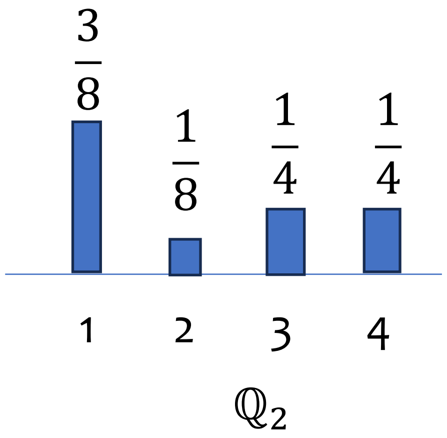

$$\mathbb{R} = (\mathbb{P}_1||\mathbb{Q}_1) = (\mathbb{P}_2||\mathbb{Q}_2)$$

$$d_{TV}(\mathbb{R}) = d_{TV}(\mathbb{P}_1,\mathbb{Q}_1) = d_{TV}(\mathbb{P}_2,\mathbb{Q}_2) = \frac{1}{6}$$

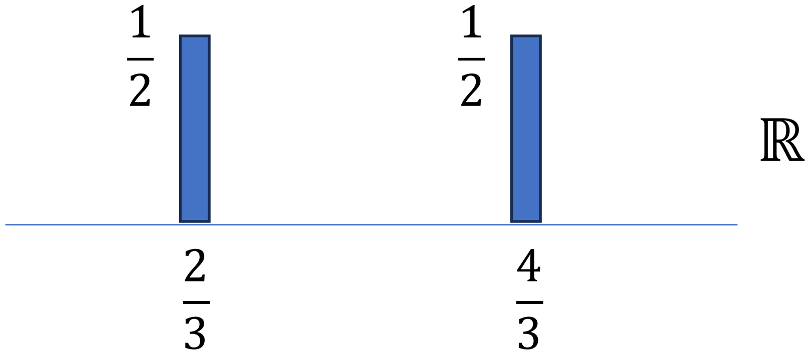

Ratio and deterministic algorithm

Ratio $R \sim \mathbb{R} = (\mathbb{P}||\mathbb{Q})$

$$R = \frac{\mathbb{P}(X)}{\mathbb{Q}(X)}, \qquad \text{where } X \sim Q$$

$$d_{TV}(\mathbb{P},\mathbb{Q}) = d_{TV}(\mathbb{R}) = \mathbb{E}[\max(0,1-R)]$$

$\mathbb{R}$ ***preserves*** $d_{TV}(\mathbb{P},\mathbb{Q})$, may compress some redundant information

If $\mathbb{R}_1 = (\mathbb{P}_1||\mathbb{Q}_1)$ and $\mathbb{R}_2 = (\mathbb{P}_2||\mathbb{Q}_2)$, then

$$\mathbb{R}_1 \cdot_{\text{ind}} \mathbb{R}_2 = (\mathbb{P}_1 \times \mathbb{P}_2 || \mathbb{Q}_1 \times \mathbb{Q}_2)$$

- $\mathbb{P}_1 \times \mathbb{P}_2$ is the **product distribution** $\mathbb{P}_1$ and $\mathbb{P}_2$
- $\mathbb{R}_1 \cdot_{\text{ind}} \mathbb{R}_2$ is the distribution of the **product of two independent random real number**

$$R_1 R_2 \sim \mathbb{R}_1 \cdot_{\text{ind}} \mathbb{R}_2, \text{ where } R_1 \sim \mathbb{R}_1 \text{ and } R_2 \sim \mathbb{R}_2 \text{ are ind. samples}$$

A naïve deterministic algorithm

- **Input**: distributions of $\mathbb{P}_1, \mathbb{P}_2, \ldots, \mathbb{P}_n, \mathbb{Q}_1, \mathbb{Q}_2, \ldots, \mathbb{Q}_n$ and an error bound ϵ
- **Output**: an approximation of $d_{TV}(\mathbb{P}, \mathbb{Q})$

- Compute $\mathbb{R}_i \leftarrow (\mathbb{P}_i || \mathbb{Q}_i)$ for all $i \in [n]$

- Compute $\mathbb{R}_{1:1} \leftarrow \mathbb{R}_1$
- Compute $\mathbb{R}_{1:2} \leftarrow \mathbb{R}_{1:1} \cdot_{ind} \mathbb{R}_2$

…

- Compute $\mathbb{R}_{1:i} \leftarrow \mathbb{R}_{1:i-1} \cdot_{ind} \mathbb{R}_i$

…

- Compute distribution $\mathbb{R}_{1:n} \leftarrow \mathbb{R}_{1:n-1} \cdot_{ind} \mathbb{R}_n$
- Return $d_{TV}(\mathbb{R}_{1:n}) = \mathbb{E}_{R \sim \mathbb{R}_{1:n}}[\max(0,1-R)]$

Exact computing of $d_{TV}(\mathbb{P}, \mathbb{Q})$

The support size is ***large***

$$|\mathrm{supp}(\mathbb{R}_{1:i})| = \exp(\Omega(i))$$

Efficiently compute $\widehat{\mathbb{R}}_{\mathbf{1:n}}$ s.t.

$$d_{TV}\left(\widehat{\mathbb{R}}_{1:\boldsymbol{n}}\right) \approx d_{TV}(\mathbb{R}_{1:n})$$

- Compute $\mathbb{R}_i \leftarrow (\mathbb{P}_i || \mathbb{Q}_i)$ for all $i \in [n]$

- $\widehat{\mathbb{R}}_{1:1} \leftarrow \mathbb{R}_1$

- Compute $\mathbb{R}'_{1:2} \leftarrow \widehat{\mathbb{R}}_{1:1} \cdot_{ind} \mathbb{R}_2$
- $\widehat{\mathbb{R}}_{1:2} \leftarrow \textbf{Sparsify}(\mathbb{R}'_{1:2})$

…

- Compute $\mathbb{R}'_{1:i} \leftarrow \widehat{\mathbb{R}}_{1:i-1} \cdot_{ind} \mathbb{R}_i$
- $\widehat{\mathbb{R}}_{1:i} \leftarrow \textbf{Sparsify}(\mathbb{R}'_{1:i})$

…

- Compute $\mathbb{R}'_{1:n} = \widehat{\mathbb{R}}_{1:n-1} \cdot_{ind} \mathbb{R}_n$
- $\widehat{\mathbb{R}}_{1:n} \leftarrow \textbf{Sparsify}(\mathbb{R}'_{1:n})$

- Return $d_{TV}\left(\widehat{\mathbb{R}}_{1:n}\right) = \mathbb{E}_{R \sim \widehat{\mathbb{R}}_{1:n}}[\max(0, 1 - R)]$

The Sparsify subroutine

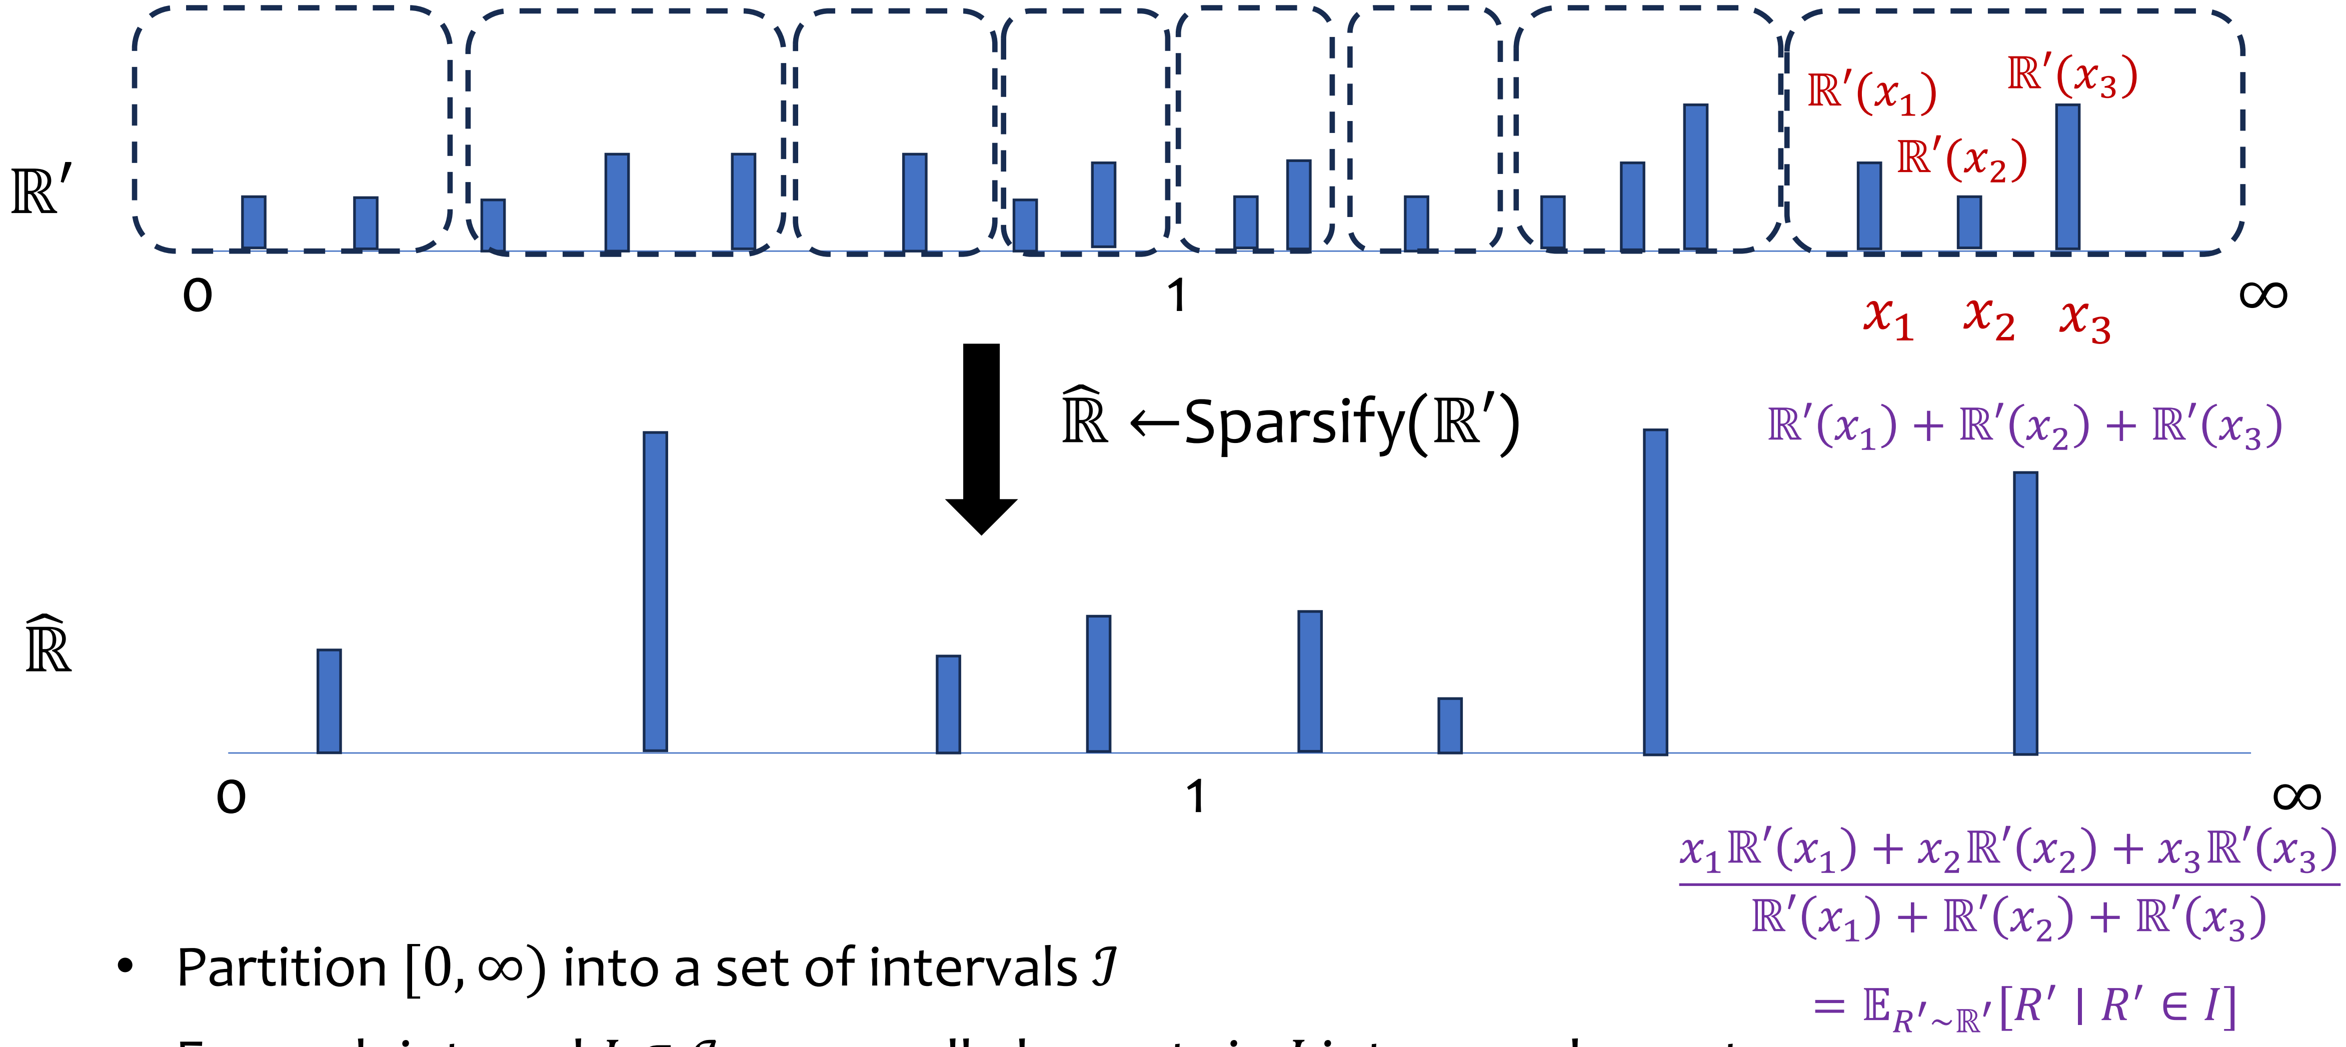

- Partition $[0,\infty)$ into a set of intervals $\mathcal{I}$
- For each interval $I \in \mathcal{I}$, merge all elements in I into one element

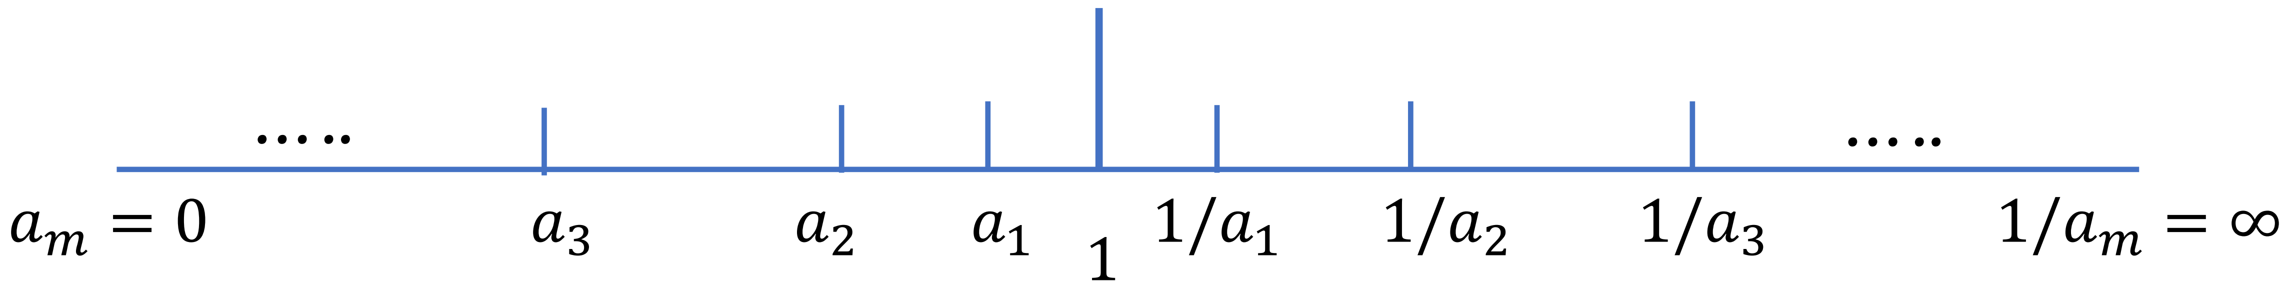

Partition of $[0,1]$

$$\begin{cases} [a_1, a_0 = 1) \\ [a_2, a_1) \\ [a_3, a_2) \\ \dots \\ [a_m = 0, a_{m-1}) \end{cases}$$

Partition of $(1, \infty)$

$$\begin{cases} [a_0 = 1, 1/a_1) \\ [1/a_1, 1/a_2) \\ [1/a_2, 1/a_3) \\ \dots \\ [1/a_{m-1}, 1/a_m = \infty) \end{cases}$$

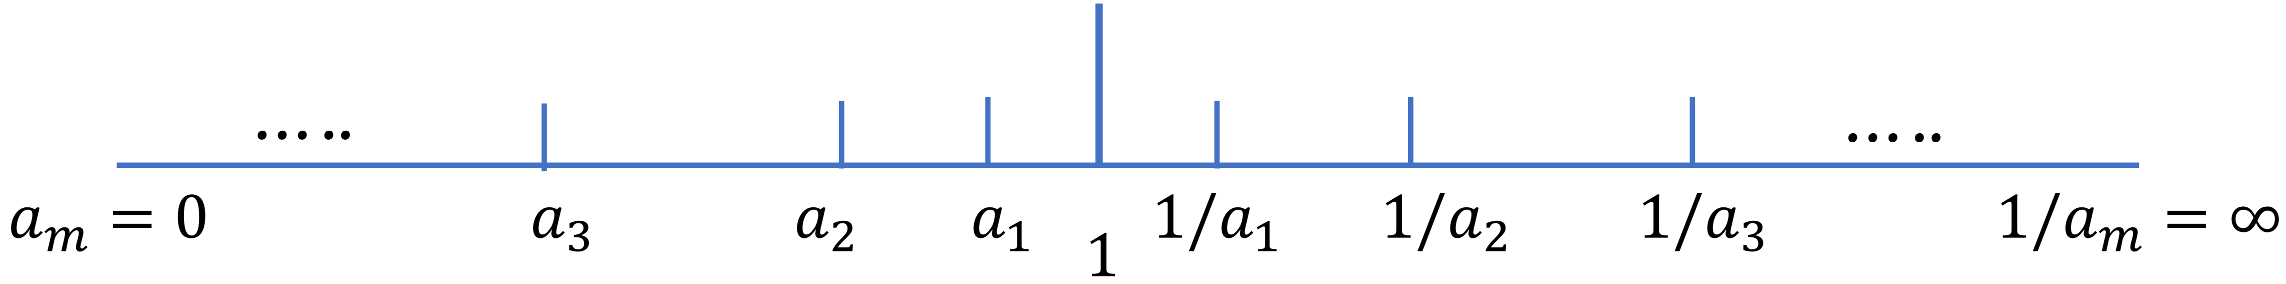

- The first interval is small

$$1 - a_1 \leq \delta_s$$

- The length of $[a_i, a_{i-1}]$ is small w.r.t. $1 - a_{i-1}$

$$\forall i > 1, \qquad |a_i - a_{i-1}| \leq \epsilon_s \cdot |1 - a_{i-1}|$$

$$\delta_s = \Theta\left(\frac{\epsilon d_{TV}(\mathbb{P}, \mathbb{Q})}{2n}\right)$$

$$\epsilon_s = \frac{\epsilon}{2n}$$

Domain size is ***small*** after the sparsification

$$m = O\left(\frac{1}{\epsilon_s}\log\frac{1}{\delta_s}\right) = O\left(\frac{2n}{\epsilon}\log\frac{n}{\epsilon d_{TV}(\mathbb{P}, \mathbb{Q})}\right)$$

Error introduced by sparsification is ***small***

- introduce a ***new metric Δ between ratios***
- If $\mathbb{R}' \leftarrow \text{Sparsify}(\mathbb{R})$, then
  $\Delta(\mathbb{R}', \mathbb{R})$ is small

Summary

- **Problem:** Compute the TV distance between two ***product distributions***
- **Algorithms:** FPTAS and FPRAS
- **Extension:** TV distance between two ***Markov chains***

Open problems

- ***Better running time*** of FPTAS: remove $\log \frac{1}{d_{TV}(\mathbb{P},\mathbb{Q})}$ in $\tilde{O}(n^2 \log \frac{1}{d_{TV}(\mathbb{P},\mathbb{Q})})$?
- Algorithm/complexity for approximating TV distance of ***general models***
  - ➢Bayes networks [BGMMPV arXiv:2309.09134]
- Relation between ***approximating TV distance*** and ***sampling/counting***