

Subset Sampling over Joins

Aryan Esmailpour¹ Xiao Hu² Jinchao Huang³ Stavros Sintos¹

¹University of Illinois Chicago ²University of Waterloo ³The Chinese University of Hong Kong

PODS 2026


June 2026

▪ Bengaluru, India

A Familiar Primitive, but in the Relational Setting

Explicit subset sampling


Given items e_1, \dots, e_n , include each item independently with probability $p(e_i)$.



A Familiar Primitive, but in the Relational Setting


Explicit subset sampling

Given items e_1, \dots, e_n , include each item independently with probability $p(e_i)$.



Relational subset sampling


The items are not stored explicitly: they are all tuples in $Join(Q)$.



A Familiar Primitive, but in the Relational Setting


Explicit subset sampling

Given items e_1, \dots, e_n , include each item independently with probability $p(e_i)$.



Relational subset sampling

The items are not stored explicitly: they are all tuples in $Join(Q)$.



Goal: generate exactly the subset sample, without materializing the join.

Why Subset Sampling over Joins?



$|Join(Q)|$ can be much larger than N

Many tasks do not need the full join

- approximate analytics and visualization
- optimizer statistics and selectivity estimation
- relational ML, coresets, and data condensation

Importance-aware sampling

Subset sampling selects join results by considering the “importance” of constituent tuples.

More motivation: A concurrent independent SIGMOD paper, *Poisson Sampling over Acyclic Joins* [BNPV'26], gives systems-oriented examples.

Problem Definition

Input. Acyclic join $Q = \{R_1, \dots, R_k\}$; per-tuple weights $p_j : R_j \rightarrow [0, 1]$.

Join-result weight:

$$p(\mathbf{u}) = \mathcal{F}(p_1(\mathbf{u}[\text{schema}(R_1)]), \dots, p_k(\mathbf{u}[\text{schema}(R_k)])), \quad \mathcal{F} \in \{\text{PRODUCT}, \text{MIN}, \text{MAX}, \text{SUM}\}.$$

Relational Subset Sampling

Return $\mathbf{X} \subseteq \text{Join}(Q)$ where each $\mathbf{u} \in \text{Join}(Q)$ is in \mathbf{X} *independently* w.p. $p(\mathbf{u})$.

Problem Definition

Input. Acyclic join $Q = \{R_1, \dots, R_k\}$; per-tuple weights $p_j : R_j \rightarrow [0, 1]$.

Join-result weight:

$$p(\mathbf{u}) = \mathcal{F}(p_1(\mathbf{u}[\text{schema}(R_1)]), \dots, p_k(\mathbf{u}[\text{schema}(R_k)])), \quad \mathcal{F} \in \{\text{PRODUCT}, \text{MIN}, \text{MAX}, \text{SUM}\}.$$

Relational Subset Sampling

Return $\mathbf{X} \subseteq \text{Join}(Q)$ where each $\mathbf{u} \in \text{Join}(Q)$ is in \mathbf{X} *independently* w.p. $p(\mathbf{u})$.

Three settings:

Static Index

Sample from data structure

One-shot

One sample from scratch

Dynamic (insertions)

Base tuples are streaming in

Key quantities (data complexity).


$N = \sum_j |R_j| =$ input size; $\mu_\Psi = \mathbf{E}[|\mathbf{X}|] = \sum_{\mathbf{u}} p(\mathbf{u}) =$ expected sample size; $|\text{Join}(Q)| =$ join size.

Want runtime in N and μ_Ψ — *not* in $|\text{Join}(Q)|$.

Baseline: Materialize the Join, Then Sample

Reduction to the classic problem


Build the explicit instance $\Psi = \langle \text{Join}(Q), p \rangle$:
compute $\mathbf{u} \in \text{Join}(Q)$ and their probabilities
from the base-tuple weights.



Baseline: Materialize the Join, Then Sample

Reduction to the classic problem

Build the explicit instance $\Psi = \langle \text{Join}(Q), \mathbf{p} \rangle$:
compute $\mathbf{u} \in \text{Join}(Q)$ and their probabilities
from the base-tuple weights.



Setting	Preproc.	Sampling	Space
Static index	$O(N + \text{Join}(Q))$	$O(1 + \mu_\Psi)$	$O(\text{Join}(Q))$
One-shot	—	$O(N + \text{Join}(Q))$	$O(\text{Join}(Q))$

Materialization bottleneck

AGM gives $|\text{Join}(Q)| = \Theta(N^{\rho^*})$ in the worst case. The explicit baseline pays for the whole join even when $\mu_\Psi \ll |\text{Join}(Q)|$.

Our Results at a Glance (PRODUCT, acyclic joins)

Setting	Method	Preproc. / Update	Sampling Time	Space
Static Index	Baseline	$O(N + Join(Q))$	$O(1 + \mu_\Psi)$	$O(Join(Q))$
	Ours	$O(N \log N \log \log N)$	$O(1 + \mu_\Psi \log N)$	$O(N \log N)$
One-shot	Baseline	—	$O(N + Join(Q))$	$O(Join(Q))$
	Ours	—	$O(N \log^2 N + \mu_\Psi)$	$O(N \log^2 N + \mu_\Psi)$
Dynamic (Index)	Basic	$O(\log^k N)$	$O(1 + \mu_\Psi \log N)$	$O(N \log^{k-1} N)$
	Optimized	$O(\log^3 N \log \log N)$	$O(1 + \mu_\Psi \log N)$	$O(N \log N)$
Dynamic (One-shot)	Ours	—	$O(N \log^3 N \log \log N + \mu_\Psi \log N)$	$O(N \log N)$

- **Near-optimal** in N and μ_Ψ up to polylog factors.
- Cyclic joins: replace N by N^w via tree decomposition.
- Extends to MIN, MAX, SUM with an extra $\log N$ factor (Appendix).

Building Blocks: Sampling by DirectAccess

Uniform Probability

For a perfect-uniform subset sampling instance $\Psi = \langle S, p \rangle$ with DirectAccess oracle, each subset sampling query can be answered in $O(1 + \mu_\Psi)$ expected time.



$i \leftarrow i + 1 + \text{Geo}(p)$, then output $\text{DirectAccess}(i)$

Building Blocks: Sampling by DirectAccess

Uniform Probability

For a perfect-uniform subset sampling instance $\Psi = \langle S, p \rangle$ with DirectAccess oracle, each subset sampling query can be answered in $O(1 + \mu_\Psi)$ expected time.



$i \leftarrow i + 1 + \text{Geo}(p)$, then output $\text{DirectAccess}(i)$

Non-uniform Probability

For a subset sampling instance $\Psi = \langle S, p \rangle$ with DirectAccess oracle and $\max p \leq p^+$, each subset sampling query can be answered in $O(1 + |S| \cdot p^+)$ expected time.

near-uniform

$\max p \leq \beta \min p$
cost $O(1 + \beta \mu_\Psi)$

light

$\max p \leq 1/|S|$
cost $O(1)$

Divide and Conquer

Partition a general problem instance into multiple subproblems with nice property.

If a subproblem is near-uniform or light *and* supports $|S|$ plus DirectAccess, we can sample efficiently from it.

First Attempt: Partition the Input Relations

Let $L = \lceil 2\rho^* \log N \rceil$, so $2^L \geq |\text{Join}(Q)|$. For every relation,

$$R_i^{(j)} = \{\mathbf{u} \in R_i : 2^{-j-1} < p_i(\mathbf{u}) \leq 2^{-j}\} \quad (j < L), \quad R_i^{(L)} = \{\mathbf{u} : p_i(\mathbf{u}) \leq 2^{-L}\}.$$

A bucket vector $\mathbf{j} = (j_1, \dots, j_k)$ defines $Q_{\mathbf{j}} = \{R_1^{(j_1)}, \dots, R_k^{(j_k)}\}$, and $\text{Join}(Q) = \bigsqcup_{\mathbf{j} \in \{0, \dots, L\}^k} \text{Join}(Q_{\mathbf{j}})$.

R_1	0	1	2	...	L
R_2	0	1	2	...	L
R_3	0	1	2	...	L

Two Examples

→ $\mathbf{j} = (1, 0, 2)$
 $Q_{\mathbf{j}} = R_1^{(1)} \bowtie R_2^{(0)} \bowtie R_3^{(2)}$

R_1	0	1	2	...	L
R_2	0	1	2	...	L
R_3	0	1	2	...	L

→ $\mathbf{j}' = (1, 2, 0)$
 $Q_{\mathbf{j}'} = R_1^{(1)} \bowtie R_2^{(2)} \bowtie R_3^{(0)}$

First Attempt: Partition the Input Relations

Let $L = \lceil 2\rho^* \log N \rceil$, so $2^L \geq |\text{Join}(Q)|$. For every relation,

$$R_i^{(j)} = \{\mathbf{u} \in R_i : 2^{-j-1} < p_i(\mathbf{u}) \leq 2^{-j}\} \quad (j < L), \quad R_i^{(L)} = \{\mathbf{u} : p_i(\mathbf{u}) \leq 2^{-L}\}.$$

A bucket vector $\mathbf{j} = (j_1, \dots, j_k)$ defines $Q_{\mathbf{j}} = \{R_1^{(j_1)}, \dots, R_k^{(j_k)}\}$, and $\text{Join}(Q) = \bigsqcup_{\mathbf{j} \in \{0, \dots, L\}^k} \text{Join}(Q_{\mathbf{j}})$.

R_1	0	1	2	...	L
R_2	0	1	2	...	L
R_3	0	1	2	...	L

Two Examples

R_1	0	1	2	...	L
R_2	0	1	2	...	L
R_3	0	1	2	...	L

→ $\mathbf{j} = (1, 0, 2)$
 $Q_{\mathbf{j}} = R_1^{(1)} \bowtie R_2^{(0)} \bowtie R_3^{(2)}$

→ $\mathbf{j}' = (1, 2, 0)$
 $Q_{\mathbf{j}'} = R_1^{(1)} \bowtie R_2^{(2)} \bowtie R_3^{(0)}$

Each sub-join has nice property

For $\mathbf{u} \in \text{Join}(Q_{\mathbf{j}})$:

$$2^{-\sum_i j_i - k} < p(\mathbf{u}) \leq 2^{-\sum_i j_i}.$$

- if $\sum_i j_i < L$: $Q_{\mathbf{j}}$ is 2^k -uniform;
- if $\sum_i j_i \geq L$: $Q_{\mathbf{j}}$ is **light**.


For each sub-join, we build index to support DirectAccess [Zhao et al.'18, Carmeli et al.'20].

First Attempt: Partition the Input Relations


Let $L = \lceil 2\rho^* \log N \rceil$, so $2^L \geq |\text{Join}(Q)|$. For every relation,

$$R_i^{(j)} = \{\mathbf{u} \in R_i : 2^{-j-1} < p_i(\mathbf{u}) \leq 2^{-j}\} \quad (j < L), \quad R_i^{(L)} = \{\mathbf{u} : p_i(\mathbf{u}) \leq 2^{-L}\}.$$

A bucket vector $\mathbf{j} = (j_1, \dots, j_k)$ defines $Q_{\mathbf{j}} = \{R_1^{(j_1)}, \dots, R_k^{(j_k)}\}$, and $\text{Join}(Q) = \bigsqcup_{\mathbf{j} \in \{0, \dots, L\}^k} \text{Join}(Q_{\mathbf{j}})$.



Two Examples



$$\mathbf{j} = (1, 0, 2) \\ Q_{\mathbf{j}} = R_1^{(1)} \bowtie R_2^{(0)} \bowtie R_3^{(2)}$$

$$\mathbf{j}' = (1, 2, 0) \\ Q_{\mathbf{j}'} = R_1^{(1)} \bowtie R_2^{(2)} \bowtie R_3^{(0)}$$

Each sub-join has nice property

For $\mathbf{u} \in \text{Join}(Q_{\mathbf{j}})$:

$$2^{-\sum_i j_i - k} < p(\mathbf{u}) \leq 2^{-\sum_i j_i}.$$

- if $\sum_i j_i < L$: $Q_{\mathbf{j}}$ is 2^k -uniform;
- if $\sum_i j_i \geq L$: $Q_{\mathbf{j}}$ is **light**.

For each sub-join, we build index to support DirectAccess [Zhao et al.'18, Carmeli et al.'20].

Feasible, but still expensive: $(L+1)^k$ sub-joins; each input tuple appears in $(L+1)^{k-1}$ sub-indexes.


Optimization: Partition the Conceptual Join Results

The first attempt classifies Q_j using only the total $\sum_i j_i$, not the whole vector \mathbf{j} .

Per-tuple score. For $\mathbf{u} \in R_i$, let $\phi(\mathbf{u}) = \lfloor -\log_2 p_i(\mathbf{u}) \rfloor$.

Join-result score. $\bar{\phi}(\mathbf{u}) = \sum_{i=1}^k \phi(\mathbf{u}[\text{schema}(R_i)])$. **Merged bucket.** $\mathcal{B}_\ell = \{\mathbf{u} \in \text{Join}(Q) : \bar{\phi}(\mathbf{u}) = \ell\}$.

For $\ell < L$, \mathcal{B}_ℓ merges all $\text{Join}(Q_j)$ with $\sum_i j_i = \ell$; for $\mathbf{u} \in \mathcal{B}_\ell$, $2^{-\ell-k} < p(\mathbf{u}) \leq 2^{-\ell}$.




Optimization: Partition the Conceptual Join Results

The first attempt classifies Q_j using only the total $\sum_i j_i$, not the whole vector \mathbf{j} .

Per-tuple score. For $\mathbf{u} \in R_i$, let $\phi(\mathbf{u}) = \lfloor -\log_2 p_i(\mathbf{u}) \rfloor$.

Join-result score. $\bar{\phi}(\mathbf{u}) = \sum_{i=1}^k \phi(\mathbf{u}[\text{schema}(R_i)])$. **Merged bucket.** $\mathcal{B}_\ell = \{\mathbf{u} \in \text{Join}(Q) : \bar{\phi}(\mathbf{u}) = \ell\}$.

For $\ell < L$, \mathcal{B}_ℓ merges all $\text{Join}(Q_j)$ with $\sum_i j_i = \ell$; for $\mathbf{u} \in \mathcal{B}_\ell$, $2^{-\ell-k} < p(\mathbf{u}) \leq 2^{-\ell}$.



Each bucket has nice property

- if $\ell < L$: \mathcal{B}_ℓ is 2^k -uniform;
- if $\ell \geq L$: \mathcal{B}_ℓ is **light**.

No known results for supporting DirectAccess in these buckets.


Optimization: Partition the Conceptual Join Results

The first attempt classifies Q_j using only the total $\sum_i j_i$, not the whole vector \mathbf{j} .

Per-tuple score. For $\mathbf{u} \in R_i$, let $\phi(\mathbf{u}) = \lfloor -\log_2 p_i(\mathbf{u}) \rfloor$.

Join-result score. $\bar{\phi}(\mathbf{u}) = \sum_{i=1}^k \phi(\mathbf{u}[\text{schema}(R_i)])$. **Merged bucket.** $\mathcal{B}_\ell = \{\mathbf{u} \in \text{Join}(Q) : \bar{\phi}(\mathbf{u}) = \ell\}$.

For $\ell < L$, \mathcal{B}_ℓ merges all $\text{Join}(Q_j)$ with $\sum_i j_i = \ell$; for $\mathbf{u} \in \mathcal{B}_\ell$, $2^{-\ell-k} < p(\mathbf{u}) \leq 2^{-\ell}$.



Each bucket has nice property

- if $\ell < L$: \mathcal{B}_ℓ is 2^k -uniform;
- if $\ell \geq L$: \mathcal{B}_ℓ is light.

No known results for supporting DirectAccess in these buckets.

Main remaining question: support DirectAccess in each **conceptual** bucket \mathcal{B}_ℓ .
(What we are familiar with: only one bucket.)


Preprocessing: Compute Weights of L Buckets by DP

Fix a join tree \mathcal{T} of Q . For node R_i , child $R_j \in \mathcal{C}_i$, tuple $\mathbf{u} \in R_i$, score ℓ :

$$W_{i,\mathbf{u}}^j(\ell) = \# \{ \text{tuples in } \bowtie_{h \in \mathcal{T}_i^j} R_h \text{ joinable with } \mathbf{u}, \text{ score } \ell \}; \quad M_{i,\mathbf{v}}(\ell) = \sum_{\mathbf{u} \in R_i \times \mathbf{v}} W_{i,\mathbf{u}}^\emptyset(\ell).$$

Recursive structure ($R_{\text{next}(j)} = \text{next sibling of } R_j$):

$$W_{i,\mathbf{u}}^j(\ell) = \sum_{\ell_1 + \ell_2 = \ell - \phi(\mathbf{u})} M_{j,\mathbf{u}[\text{key}(j)]}(\ell_1) \cdot W_{i,\mathbf{u}}^{\text{next}(j)}(\ell_2).$$



Bottom-up: $W_{i,\mathbf{u}}^j$ over children in decreasing order; $M_{i,\mathbf{v}}$ aggregates siblings.

This is a convolution \Rightarrow use FFT

For each (R_i, \mathbf{u}, R_j) , all L values $W_{i,\mathbf{u}}^j(\cdot)$ in $O(L \log L)$ — not $O(L^2)$.

Preprocessing Cost

Index built in $O(N \log N \log \log N)$ time, $O(N \log N)$ space.

Direct Access: Retrieve the τ -th Tuple in \mathcal{B}_ℓ in $O(\log N)$

Goal. Given $\ell < L$ and rank $\tau \in [|\mathcal{B}_\ell|]$, return the τ -th tuple in \mathcal{B}_ℓ *without listing* \mathcal{B}_ℓ .

RecursiveAccess($i, j, \mathbf{v}, \ell, \tau$) — τ -th tuple in $\mathcal{T}_i^j \bowtie \mathbf{v}$, score ℓ

1. **Locate** $\mathbf{u} \in R_i \bowtie \mathbf{v}$: binary search on prefix sums of $W_{i,\mathbf{u}'}^j(\ell)$. $O(\log N)$
2. **Locate score pair** (ℓ_1, ℓ_2) with $\ell_1 + \ell_2 = \ell - \phi(\mathbf{u})$ absorbing τ . $O(\log N)$
3. **Compute sub-ranks** $\tau_1 = \lceil \tau / W_{i,\mathbf{u}}^{\text{next}(j)}(\ell_2) \rceil$, $\tau_2 = ((\tau - 1) \bmod \dots) + 1$.
4. **Recurse** on \mathcal{T}_j with (ℓ_1, τ_1) and $\mathcal{T}_i^{\text{next}(j)}$ with (ℓ_2, τ_2) ; join the answers.

Complexity. $k^2 = O(1)$ recursive calls, each $O(\log N)$ work $\Rightarrow O(\log N)$ total. Distinct calls return distinct tuples (induced lex order on \mathbf{u} and (ℓ_1, ℓ_2)).

Theorem (Static Index)

For acyclic Q with PRODUCT weights: index of size $O(N \log N)$ built in $O(N \log N \log \log N)$; each subset-sampling query in expected $O(1 + \mu_\Psi \log N)$.

One-Shot Subset Sampling — Batching DirectAccess Queries

Baseline (static-index + one query): $O(N \log N \log \log N + \mu_\Psi \log N)$. The $\mu_\Psi \log N$ hurts when $\mu_\Psi \gg N$.

Where the $\log N$ comes from. Two binary searches inside RECURSIVEACCESS: locating $\mathbf{u} \in R_i \times \mathbf{v}$, and locating the score pair (ℓ_1, ℓ_2) .

Batch all rank queries together — BatchRecursiveAccess($i, j, \mathcal{P}_{i,j}$)

1. Precompute prefix arrays $X_{i,j,\mathbf{v},\ell}$ and score-pair arrays $Y_{i,j,\mathbf{u},\ell}$.
2. **Radix-sort** the rank set $\mathcal{P}_{i,j}$ by τ ; group by (i, j, \mathbf{v}, ℓ) .
3. *Sweep once* through each array — no per-query binary search.
4. Recurse on the two children with the propagated rank sets.

Theorem (One-shot)

One subset sample in expected time $O(\min\{N \log N \log \log N + \mu_\Psi \log N, N \log^2 N + \mu_\Psi\})$. Faster in the high-output regime where the $\mu_\Psi \log N$ term dominates.

Supporting Insertion Stream

Why exact maintenance is hopeless. Even maintaining the join size on a 3-chain join takes $\Omega(\sqrt{N})$ per insertion — far too slow.

Idea: maintain power-of-2 aggregate upper bounds

- Store $\hat{M}_{i,v}(\ell) = \sum_{\mathbf{u}} \tilde{W}_{i,\mathbf{u}}^\theta(\ell)$, then $\tilde{M}_{i,v}(\ell) = 2^{\lceil \log \hat{M}_{i,v}(\ell) \rceil}$.
- Insertion-time recursion: only re-convolve at $R_{\text{parent}(i)}$ when \tilde{M} crosses a power of 2.
- $\tilde{M}_{i,v}(\ell)$ doubles at most $O(\log N)$ times — charges only logarithmically.

Querying with approximate stats. The implicit array contains *dummy* tuples. Rejection sampling on the fly: constant-factor approximation \Rightarrow each draw is real with probability $\geq 2^{-c}$.

Theorem (Dynamic, optimized)

Under tuple insertions, an index on acyclic Q is maintained with

amortized update $O(\log^3 N \log \log N)$, sample $O(1 + \mu_\Psi \log N)$, space $O(N \log N)$.

Stream one-shot: $O(N \log^3 N \log \log N + \mu_\Psi \log N)$. *Full dynamics:* non-hierarchical joins have OMv lower bound $\Omega(N^{0.5-\epsilon})$ for Boolean maintenance.

Takeaways & Open Directions

What we showed

Formulate **subset sampling over joins** and design an efficient framework. Near-optimal in N and μ_Ψ :

- Static: $O(N \log N \log \log N)$ build, $O(1 + \mu_\Psi \log N)$ query.
- One-shot: $O(\min\{N \log N \log \log N + \mu_\Psi \log N, N \log^2 N + \mu_\Psi\})$.
- Dynamic: $O(\log^3 N \log \log N)$ update.

Reusable techniques

- **Database Partitioning** — make each subproblem easy to handle.
- **DP + FFT** — exploit convolutional structure.
- **Batching DirectAccess** — share computation among DirectAccess queries.
- **Power-of-2 approximations** — charge only when a counter doubles.

Open problems. Non-decomposable / holistic aggregations; practical engineering [BNPV'26].

Thank you — questions?