




# Data Canopy

Accelerating Exploratory Statistical Analysis

Abdul Wasay | Xinding Wei | Niv Dayan | Stratos Idreos





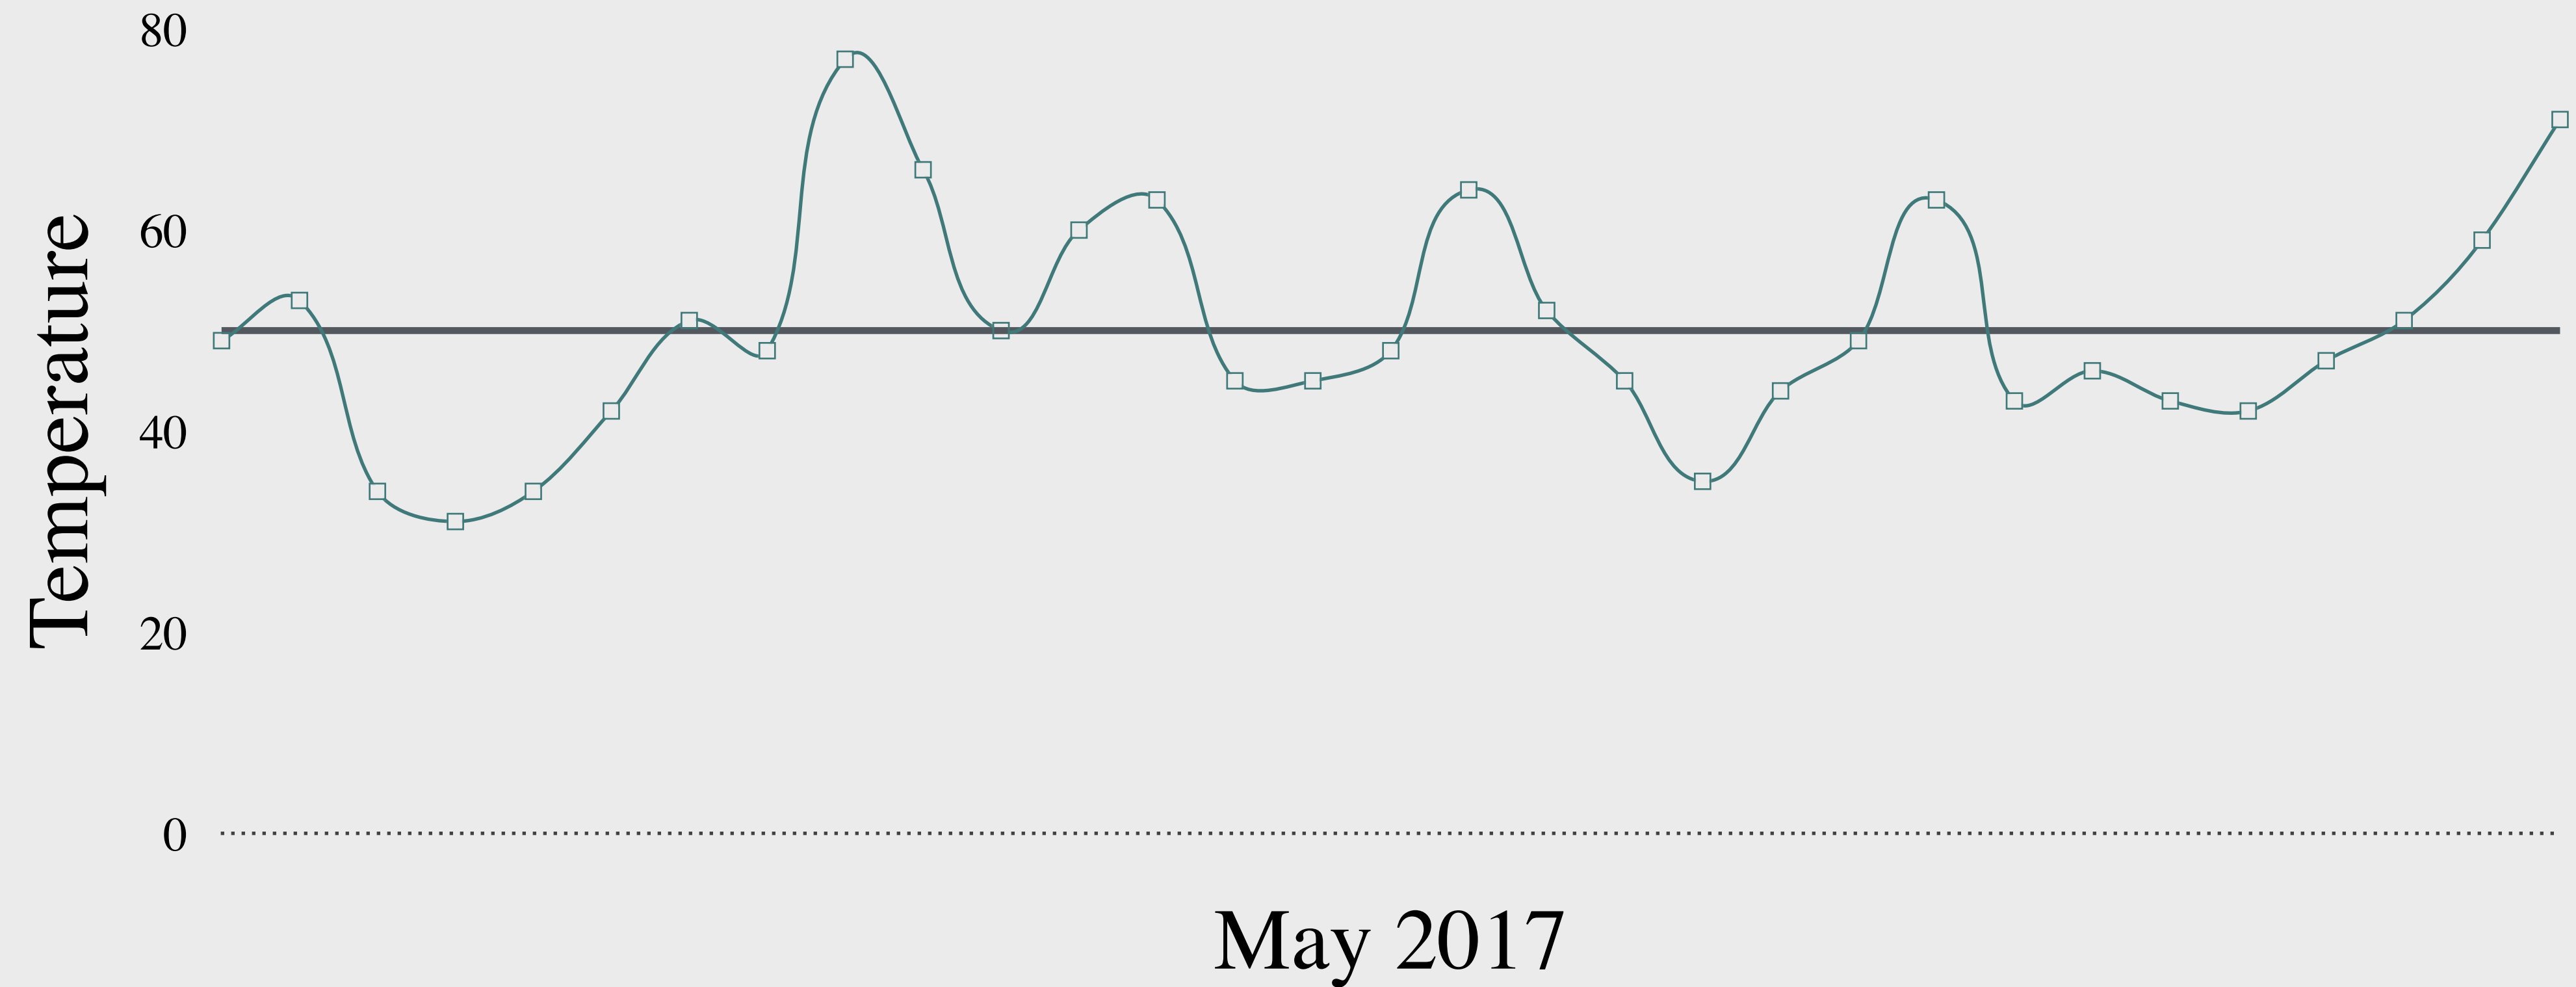

**DASlab**  
@ Harvard SEAS





# Statistics are everywhere!

Algorithms | Systems | Analytic Pipelines









# Mean



# Variance





Correlation



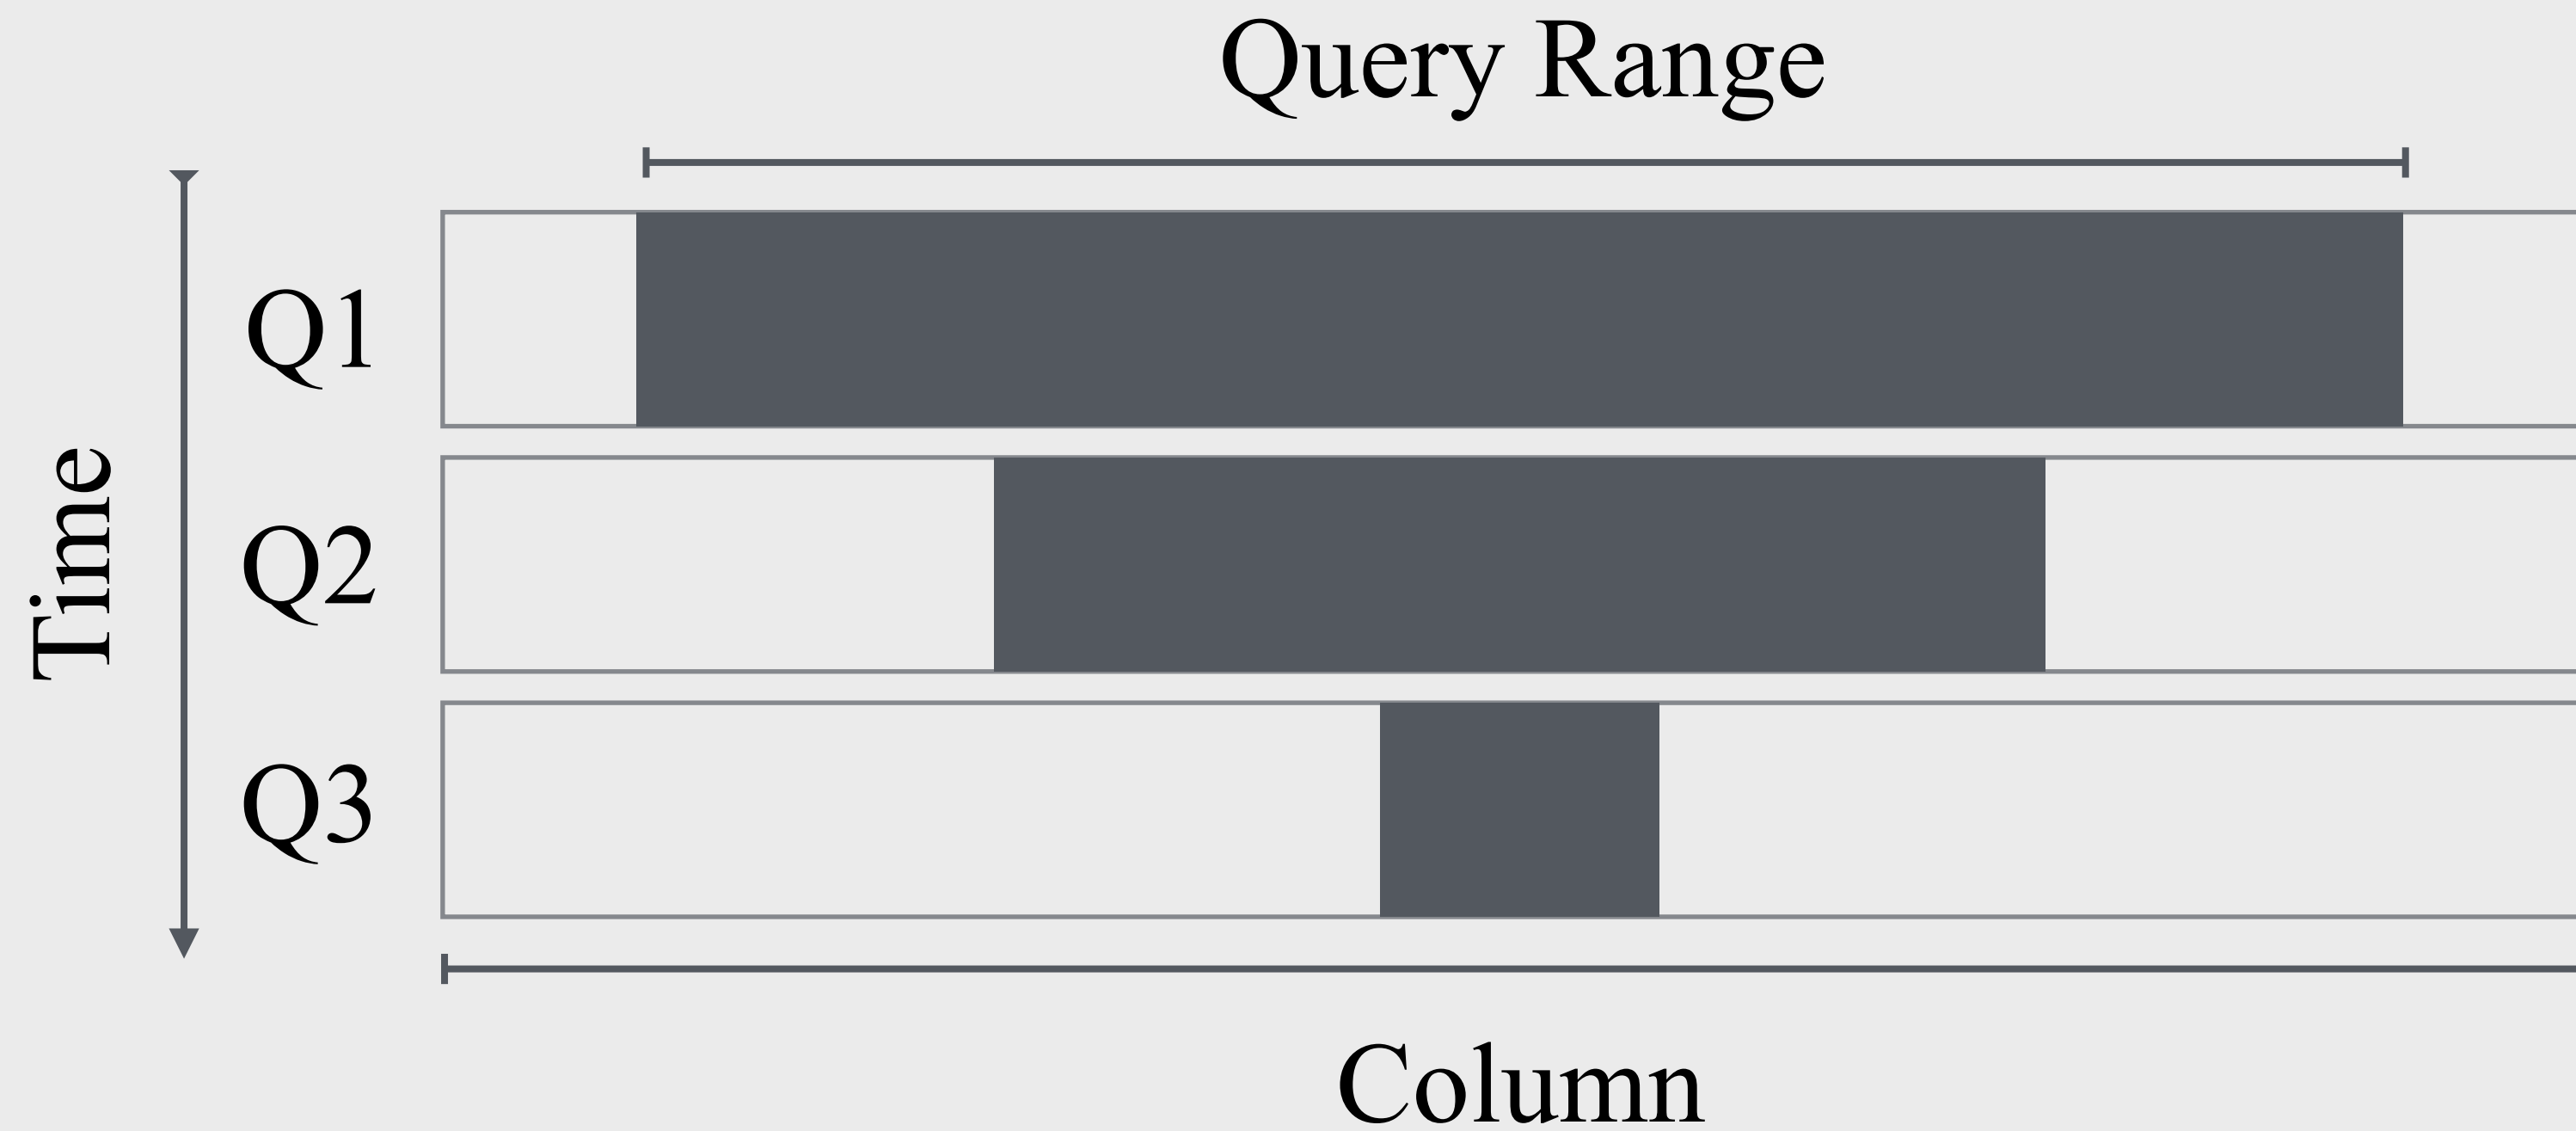
# Correlation

# Repetitive Statistics







# Repetition takes multiple forms



## Sub-range




# Repetition takes multiple forms



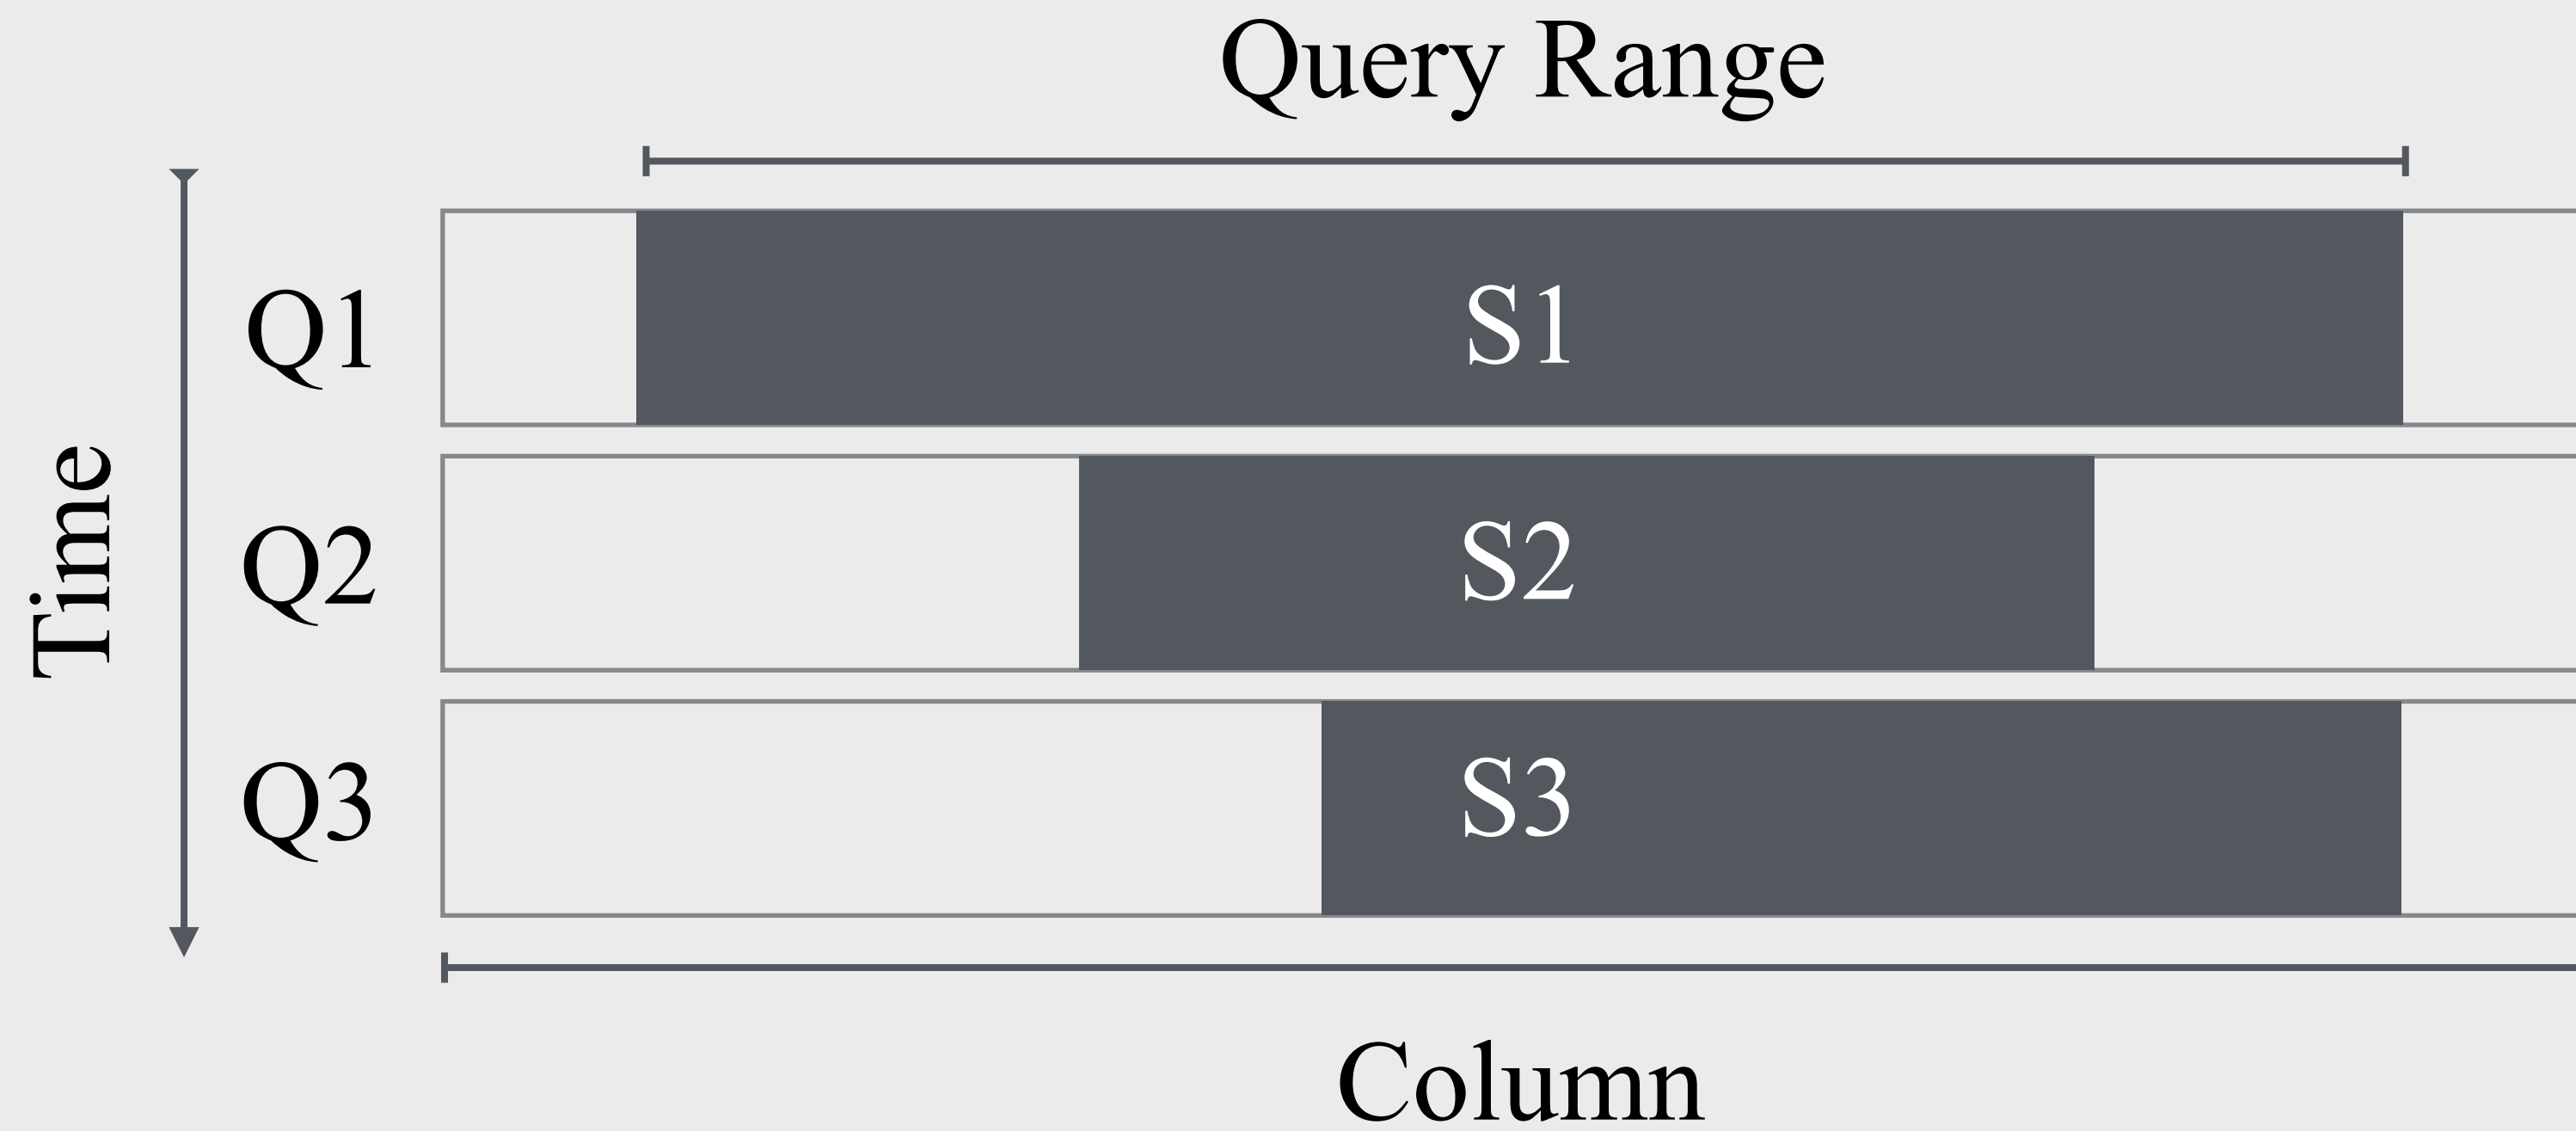
# Overlap

# Repetition takes multiple forms



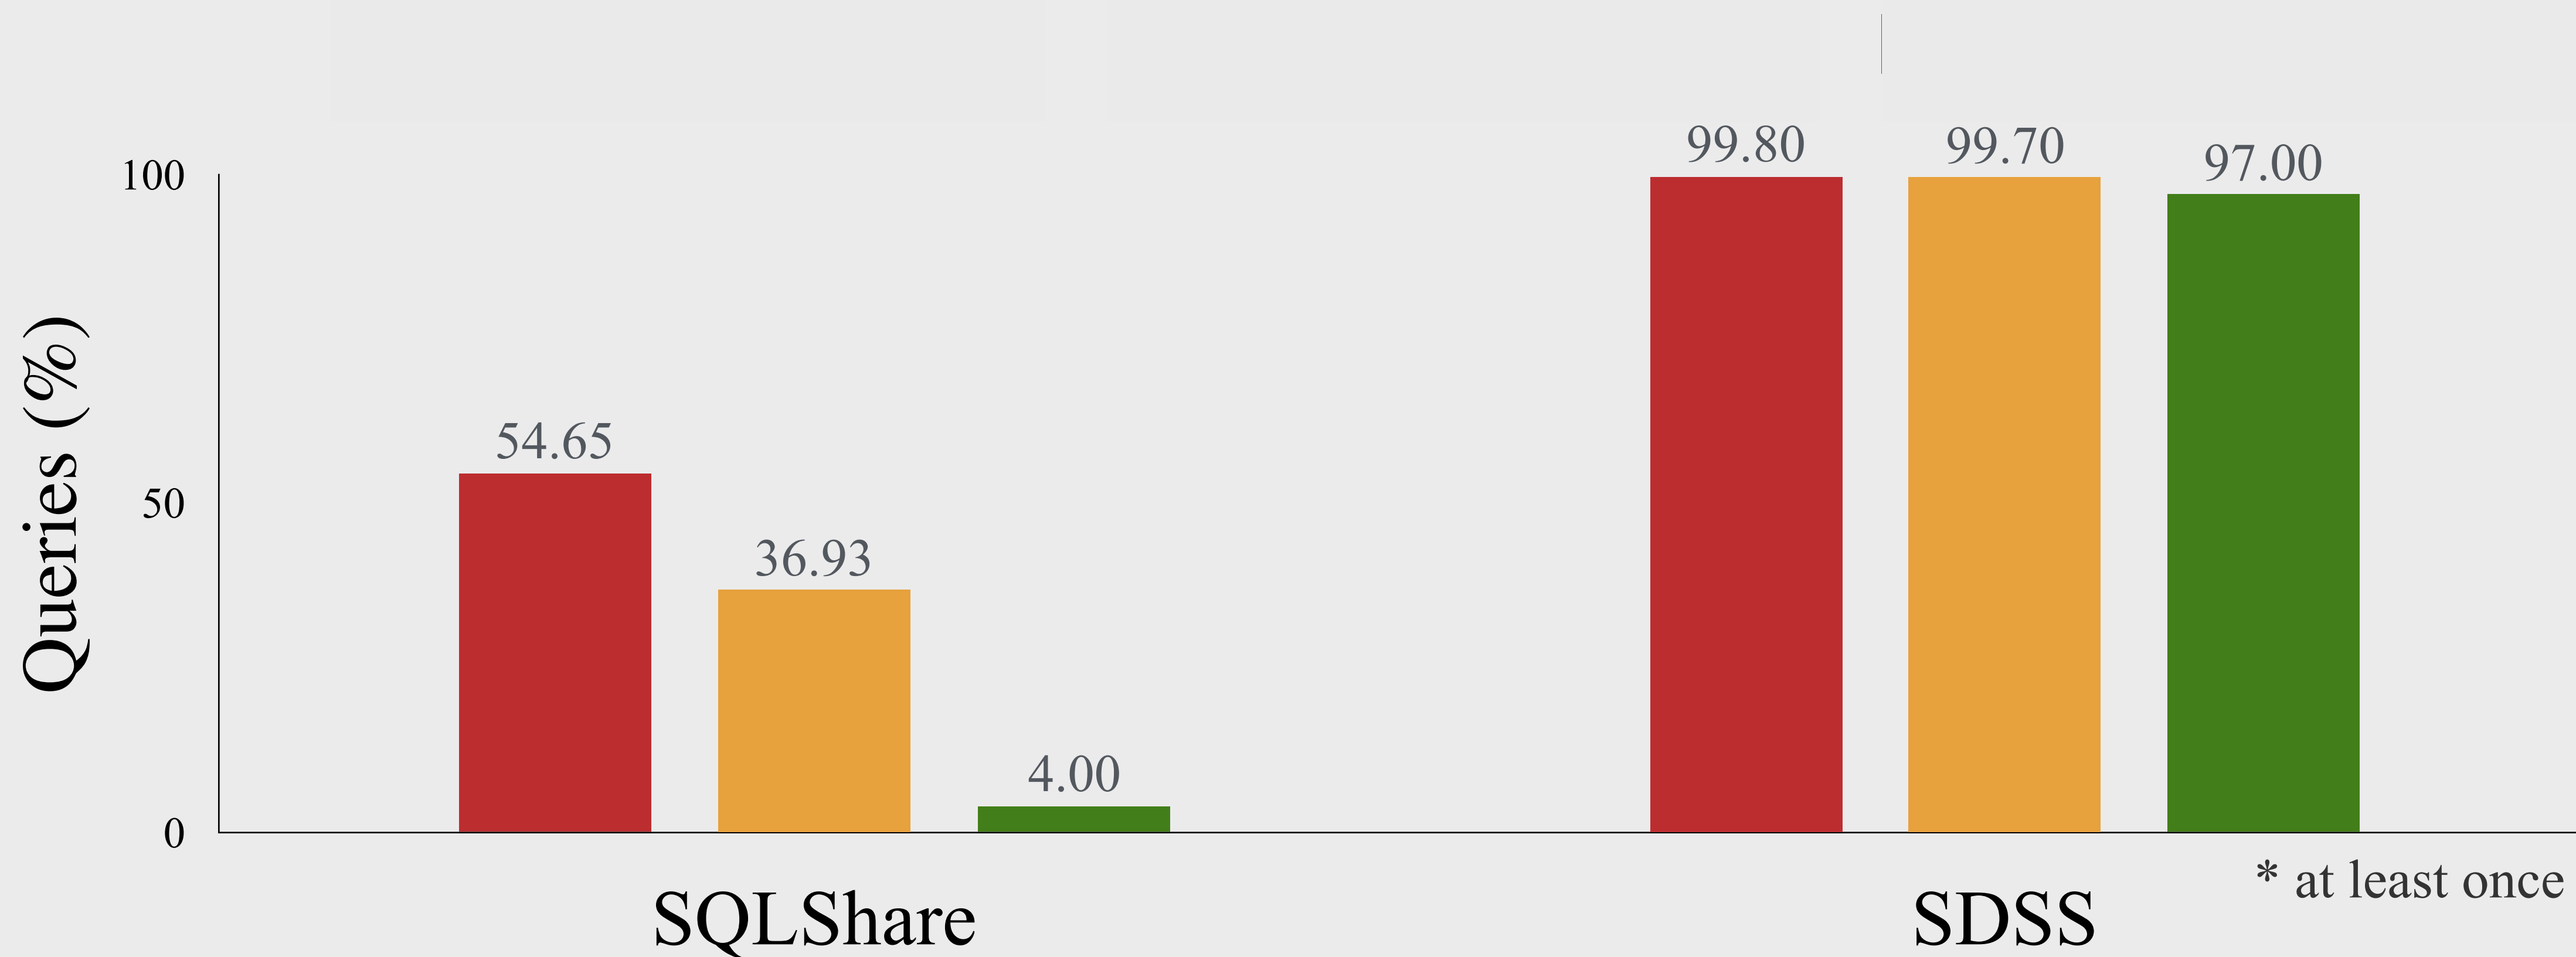
## Different Statistics

# Repetition takes multiple forms



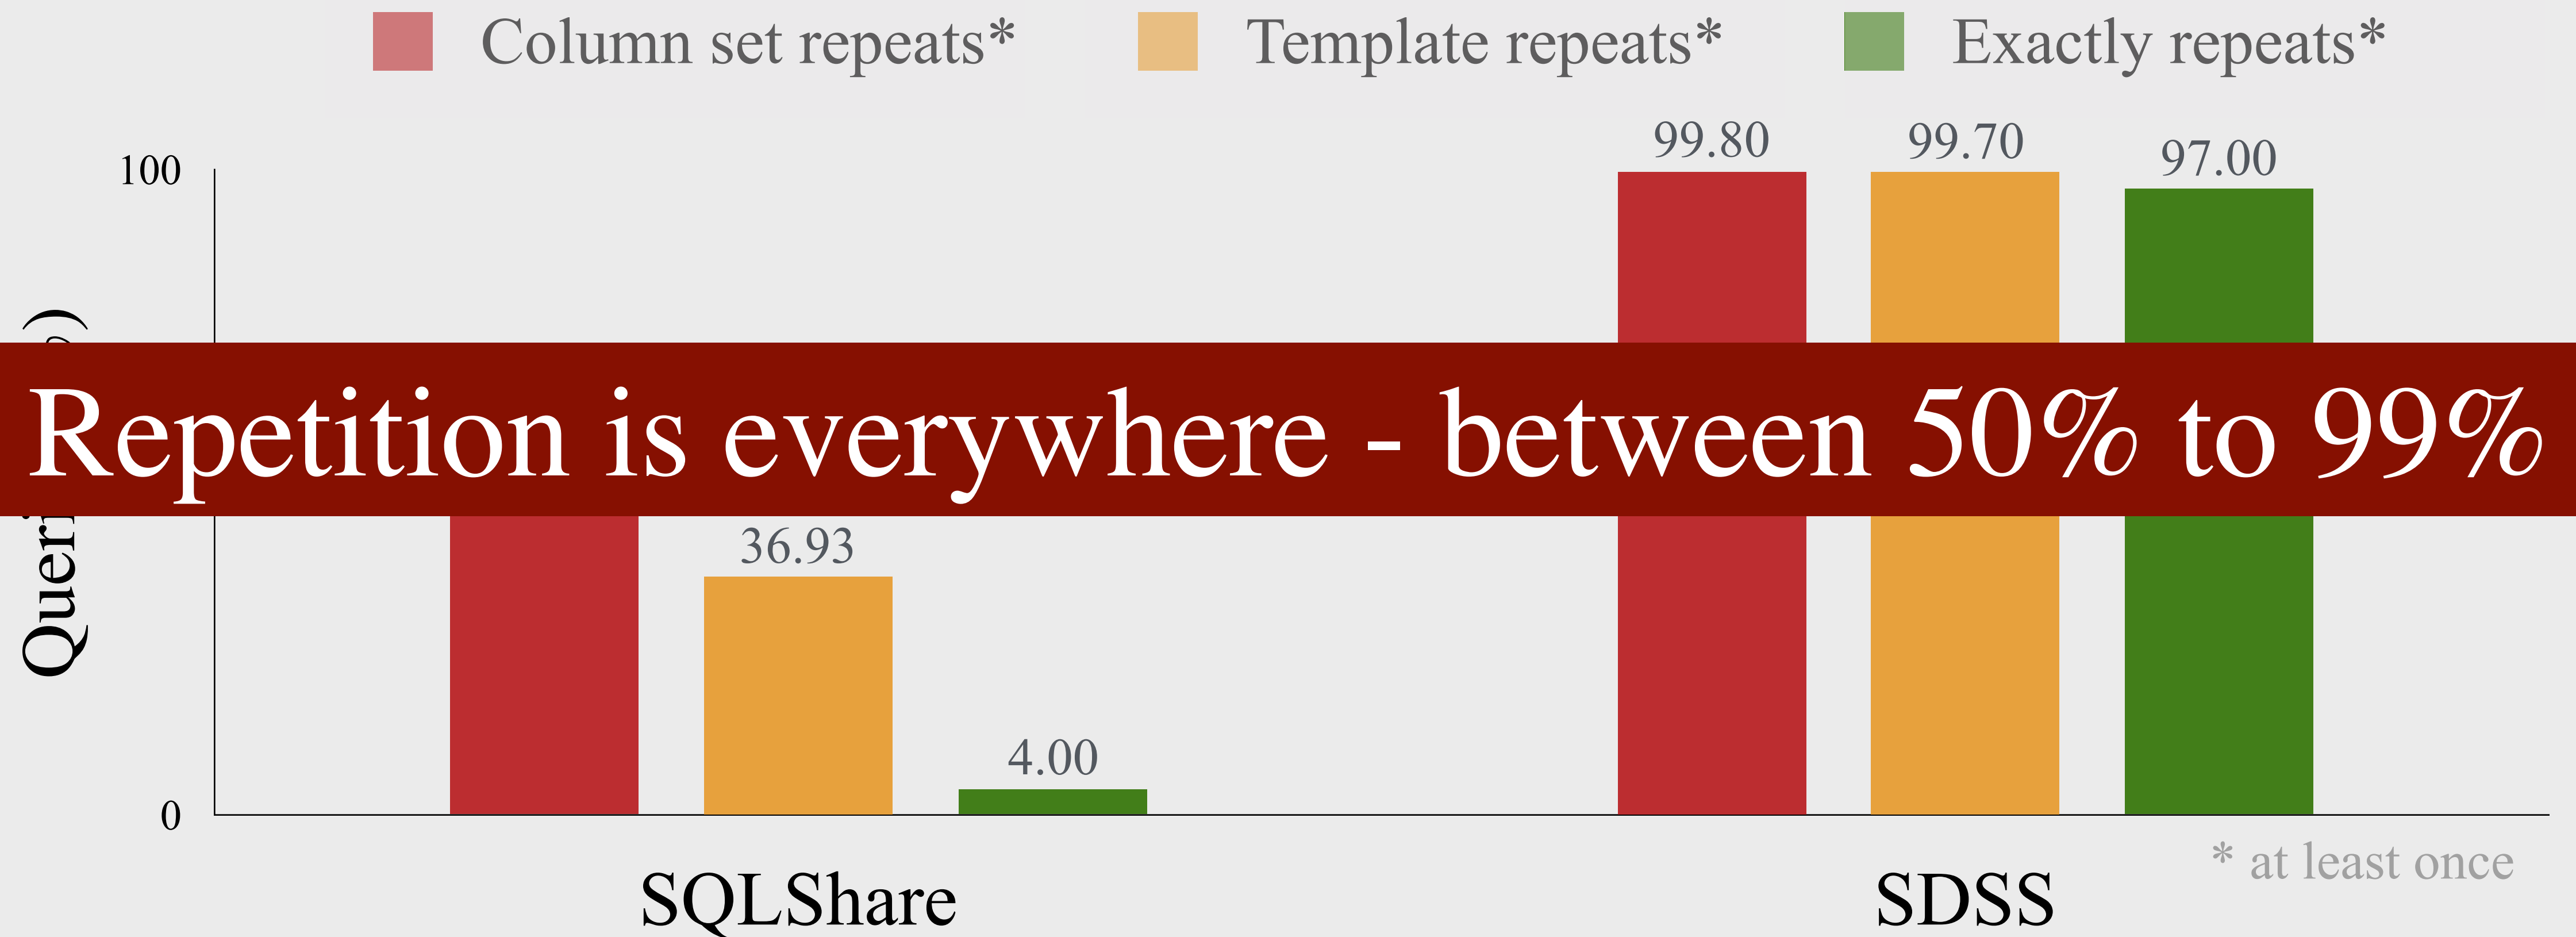
# Mixed

# Exploratory Workloads Exhibit Repetition




SQLShare: Results from a Multi-Year SQL-as-a-Service Experiment  
Shrainik Jain, Dominik Moritz, Bill Howe, Ed Lazowska. SIGMOD 2016

# Exploratory Workloads Exhibit Repetition




SQLShare: Results from a Multi-Year SQL-as-a-Service Experiment  
Shrainik Jain, Dominik Moritz, Bill Howe, Ed Lazowska. SIGMOD 2016



# How Do Existing Tools Perform?

NumPy (Python) | ModelTools (R) | MonetDB



# How Do Existing Tools Perform?

NumPy (Python) | ModelTools (R) | MonetDB

Sequence of Queries

**Q1-** Mean


**Q2-** Var.

**Q3-** Cov.

**Q4-** Mean

**Q5-** Var.

**Q6-** Cov.



# How Do Existing Tools Perform?

NumPy (Python) | ModelTools (R) | MonetDB

Sequence of Queries

**Q1- Mean**

**Q2- Var.**

**Q3- Cov.**

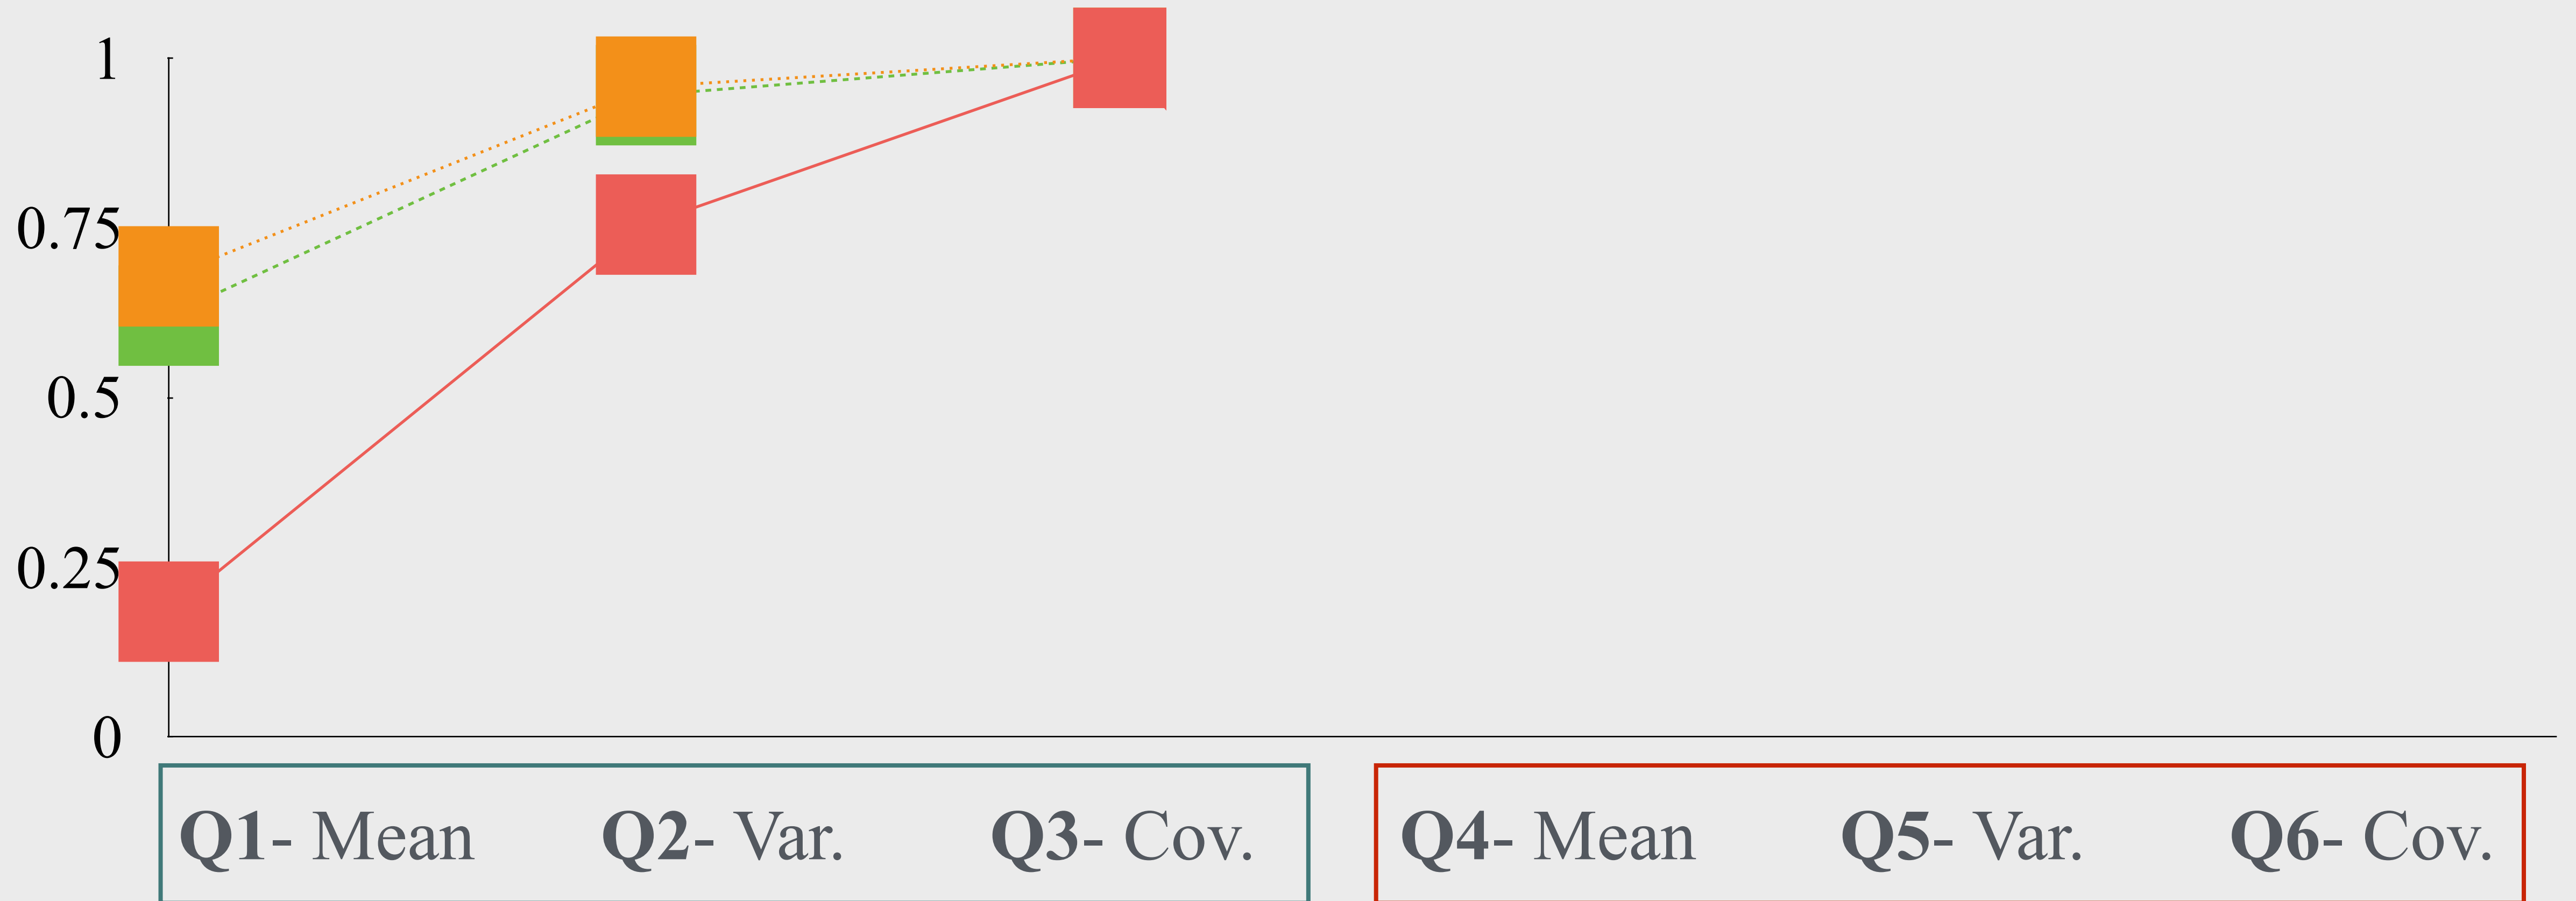
**Q4- Mean**


**Q5- Var.**

**Q6- Cov.**

# How Do Existing Tools Perform?

Normalized execution time






# How Do Existing Tools Perform?

■ NumPy      ■ Modeltools      ■ MonetDB

Normalized execution time



Q1- Mean

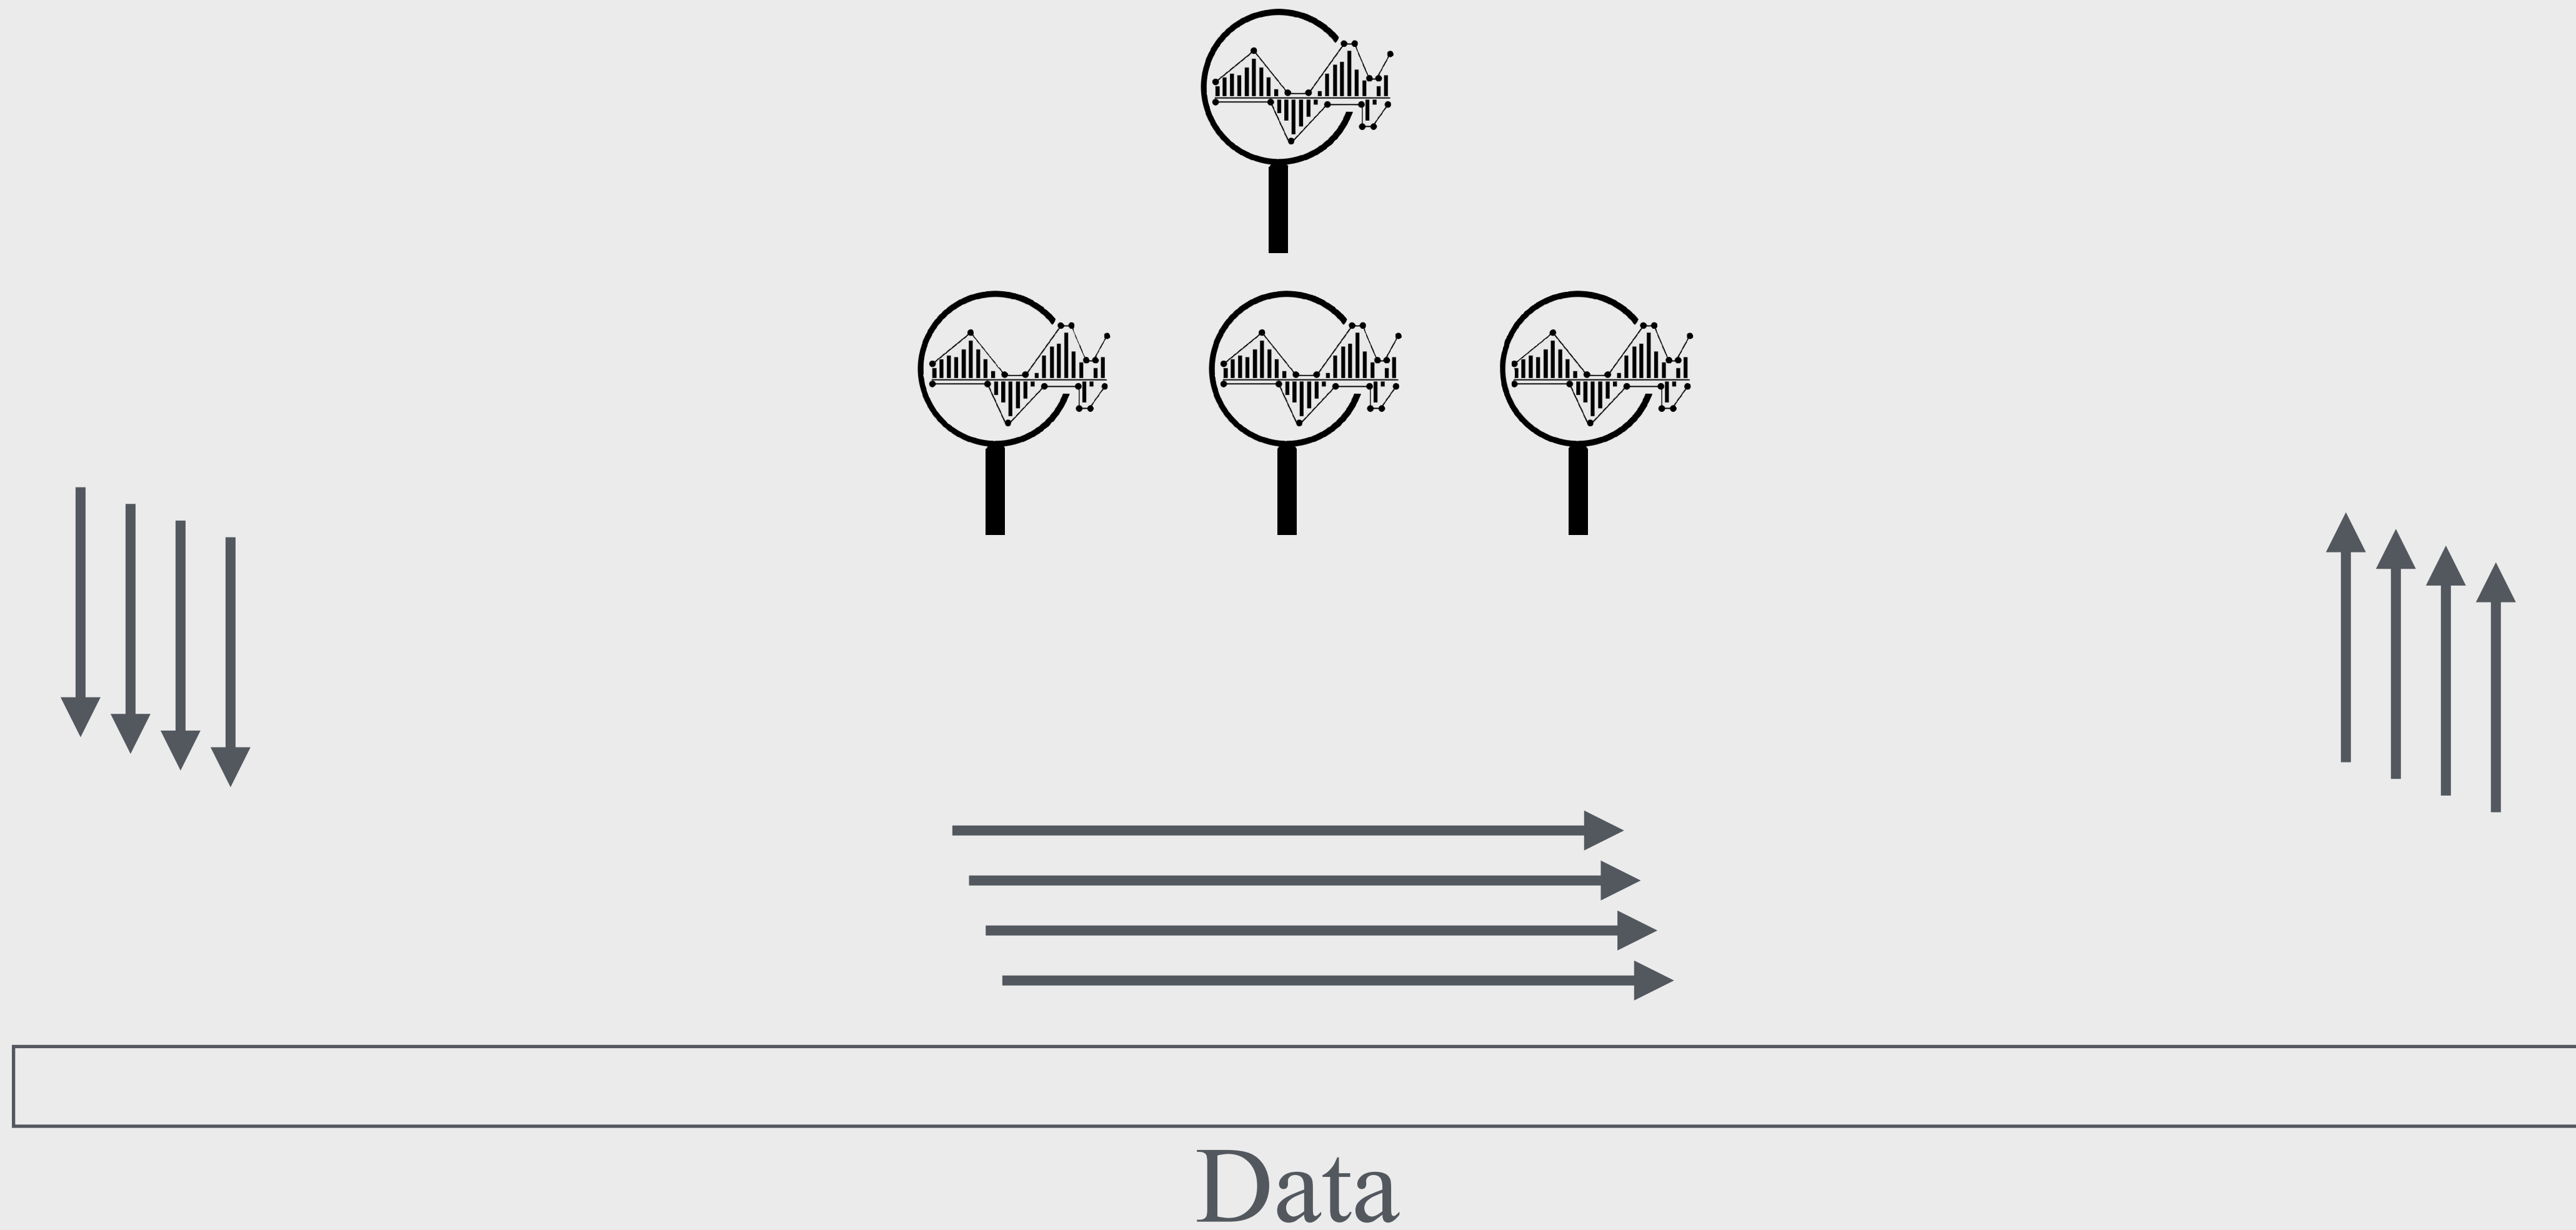

Q2- Var.



Q3- Cov.

Q4- Mean




Q5- Var.

Q6- Cov.









Existing systems always compute statistics from scratch



Data




# Data Canopy




# Data Canopy

Statistical queries




↑  
Data



# Data Canopy

Statistical queries




Library of building blocks




Data






# Data Canopy



Statistical queries



Library of building blocks

↑  
Data





# Data Canopy




Statistical queries


Avoid redundant data access to accelerate statistical analysis



↑  
Data












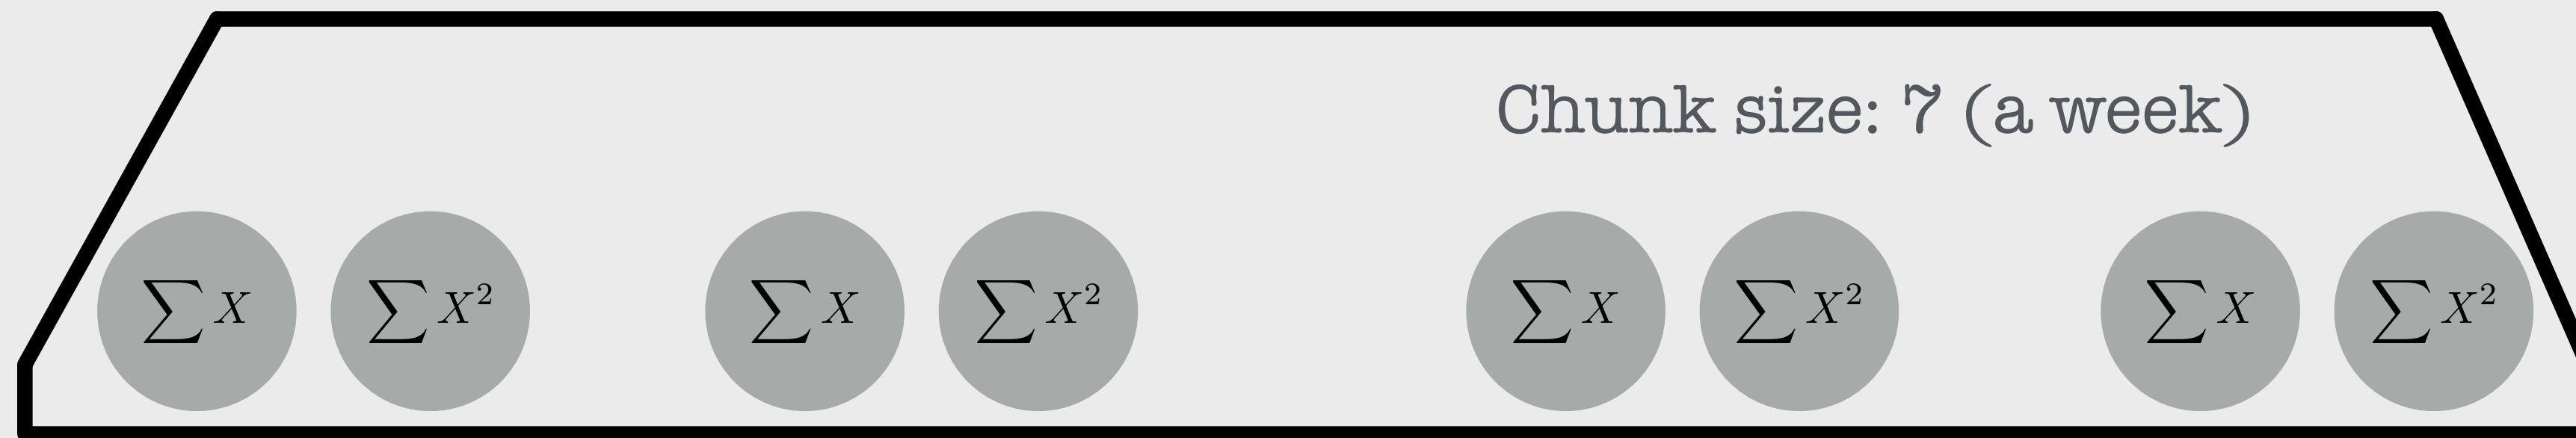
# Q: Monthly Variance


$$\text{Variance} = \left(\frac{1}{n} \sum X^2\right) - \left(\frac{1}{n} \sum X\right)^2 \longrightarrow \sum X \quad \sum X^2$$









# Q: Monthly Variance





Q: Monthly Variance  $\left(\frac{1}{n} \sum X^2\right) - \left(\frac{1}{n} \sum X\right)^2$





Var (first week)

$$\frac{1}{7} \sum X_i^2 - \left(\frac{1}{7} \sum X_i\right)^2$$

Reuse between ranges

Monthly mean


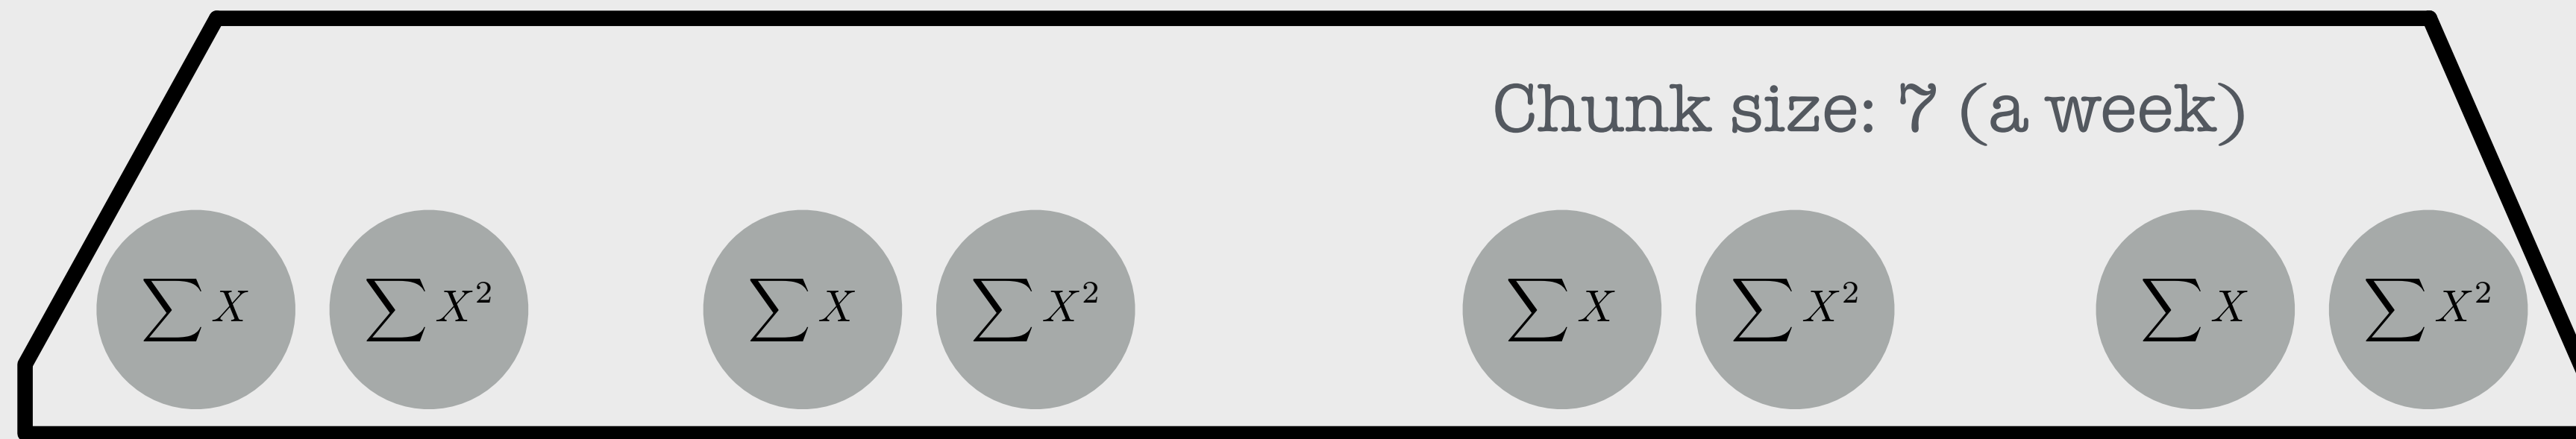
$$\frac{1}{31} \sum X_i$$


Reuse between statistics

Mean (first week)

$$\frac{1}{7} \sum X_i$$

Mixed


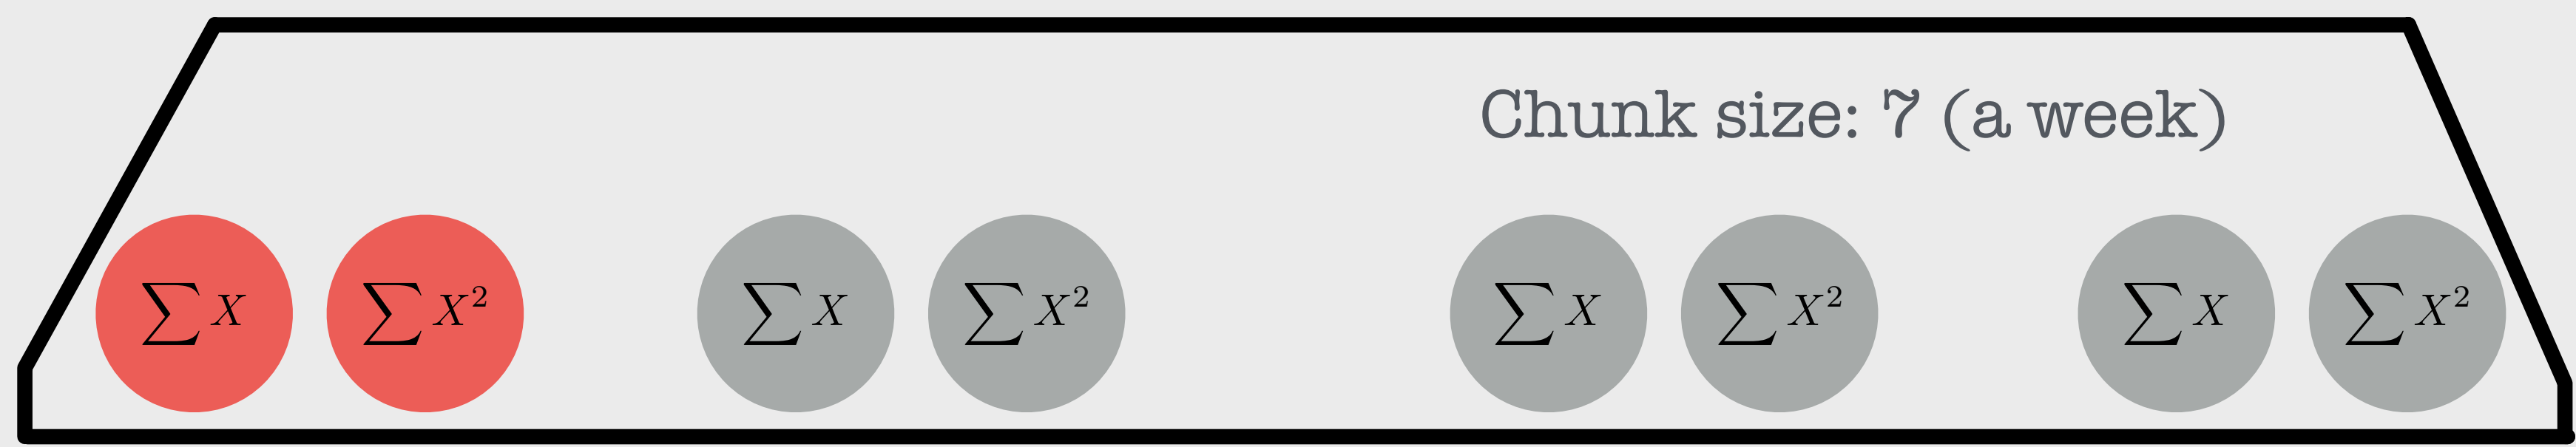





Var (first week)	$\frac{1}{7} \sum X_i^2 - (\frac{1}{7} \sum X_i)^2$	Reuse between ranges
------------------	---	----------------------

Monthly mean	$\frac{1}{31} \sum X_i$	Reuse between statistics
--------------	-------------------------	--------------------------

Mean (first week)	$\frac{1}{7} \sum X_i$	Mixed
-------------------	------------------------	-------


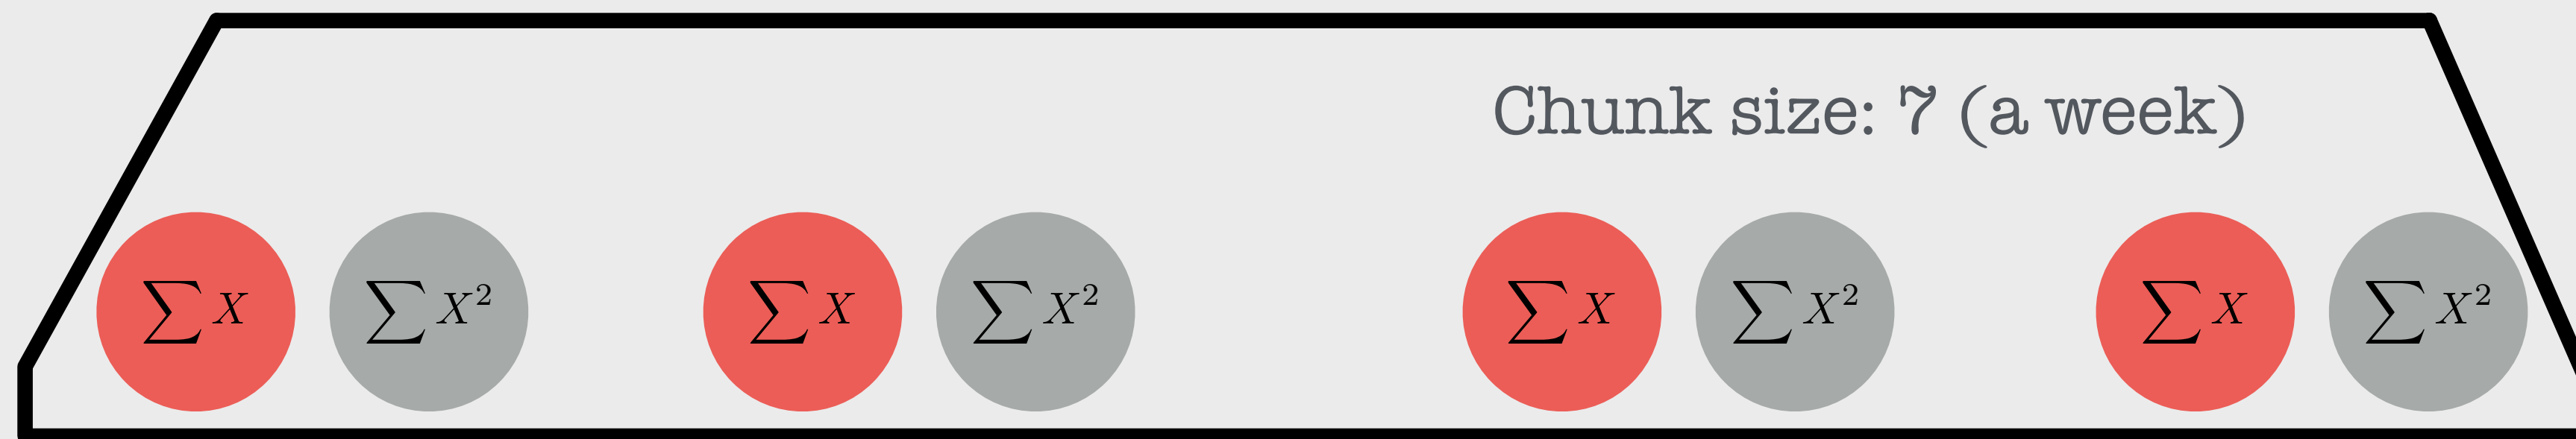





Var (first week)       $\frac{1}{7} \sum X_i^2 - (\frac{1}{7} \sum X_i)^2$       Reuse between ranges

Monthly mean       $\frac{1}{31} \sum X_i$       Reuse between statistics

Mean (first week)       $\frac{1}{7} \sum X_i$       Mixed

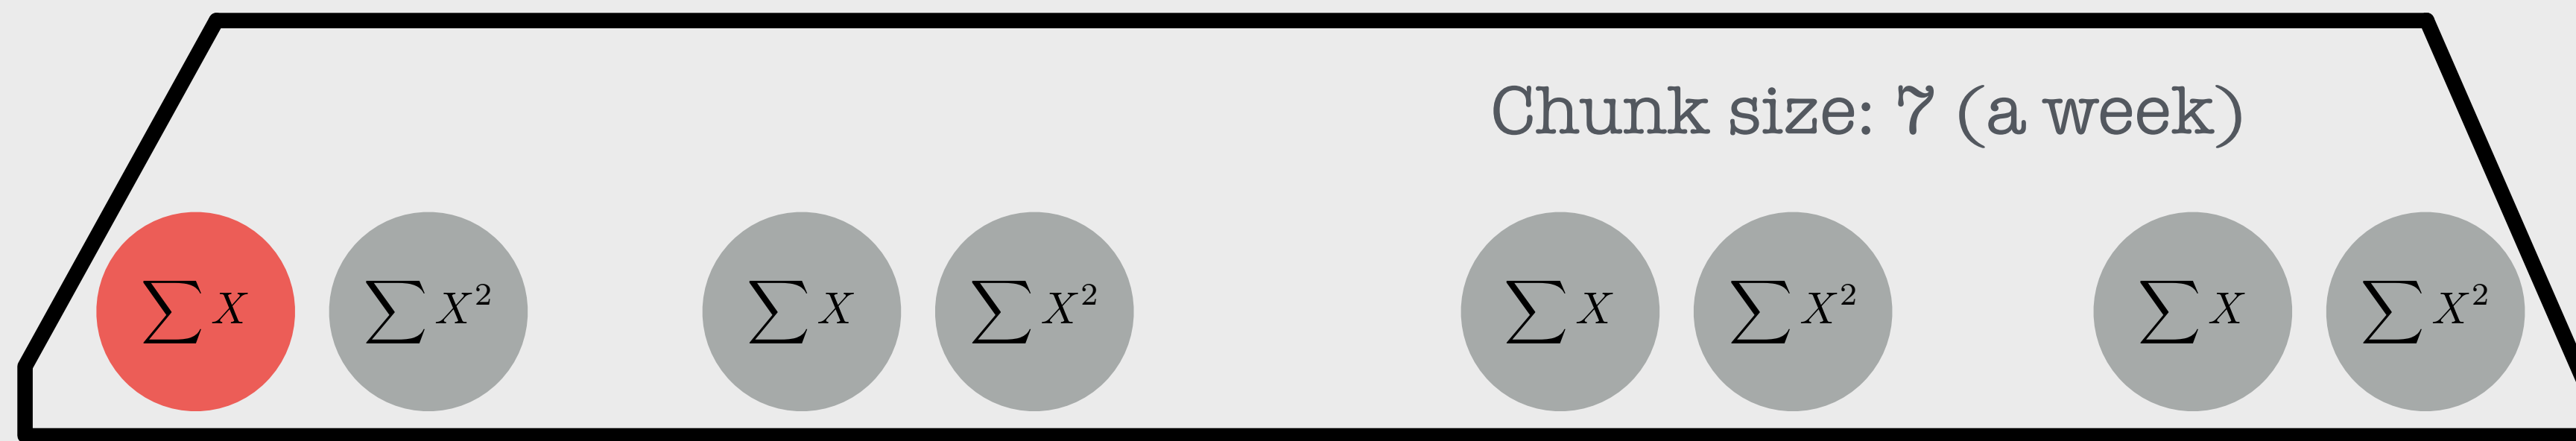





Var (first week)       $\frac{1}{7} \sum X_i^2 - (\frac{1}{7} \sum X_i)^2$       Reuse between ranges


Monthly mean       $\frac{1}{31} \sum X_i$       Reuse between statistics

Mean (first week)       $\frac{1}{7} \sum X_i$       Mixed






**? Decompose**



Statistic





Basic Aggregates



**? Synthesize**

**? Store**

**? Chunk size**



a chunk




**? Updates**


**? Build**

Data


**? Memory Pressure**



**? Decompose**



Statistic





Basic Aggregates



**? Synthesize**

**? Store**

**? Chunk size**



a chunk




**? Updates**

**? Build**

Data


**? Memory Pressure**



**? Decompose**



Statistic




Basic Aggregates


**? Synthesize**

**? Store**


**? Chunk size**




a chunk




Data




Statistic → Basic Aggregates



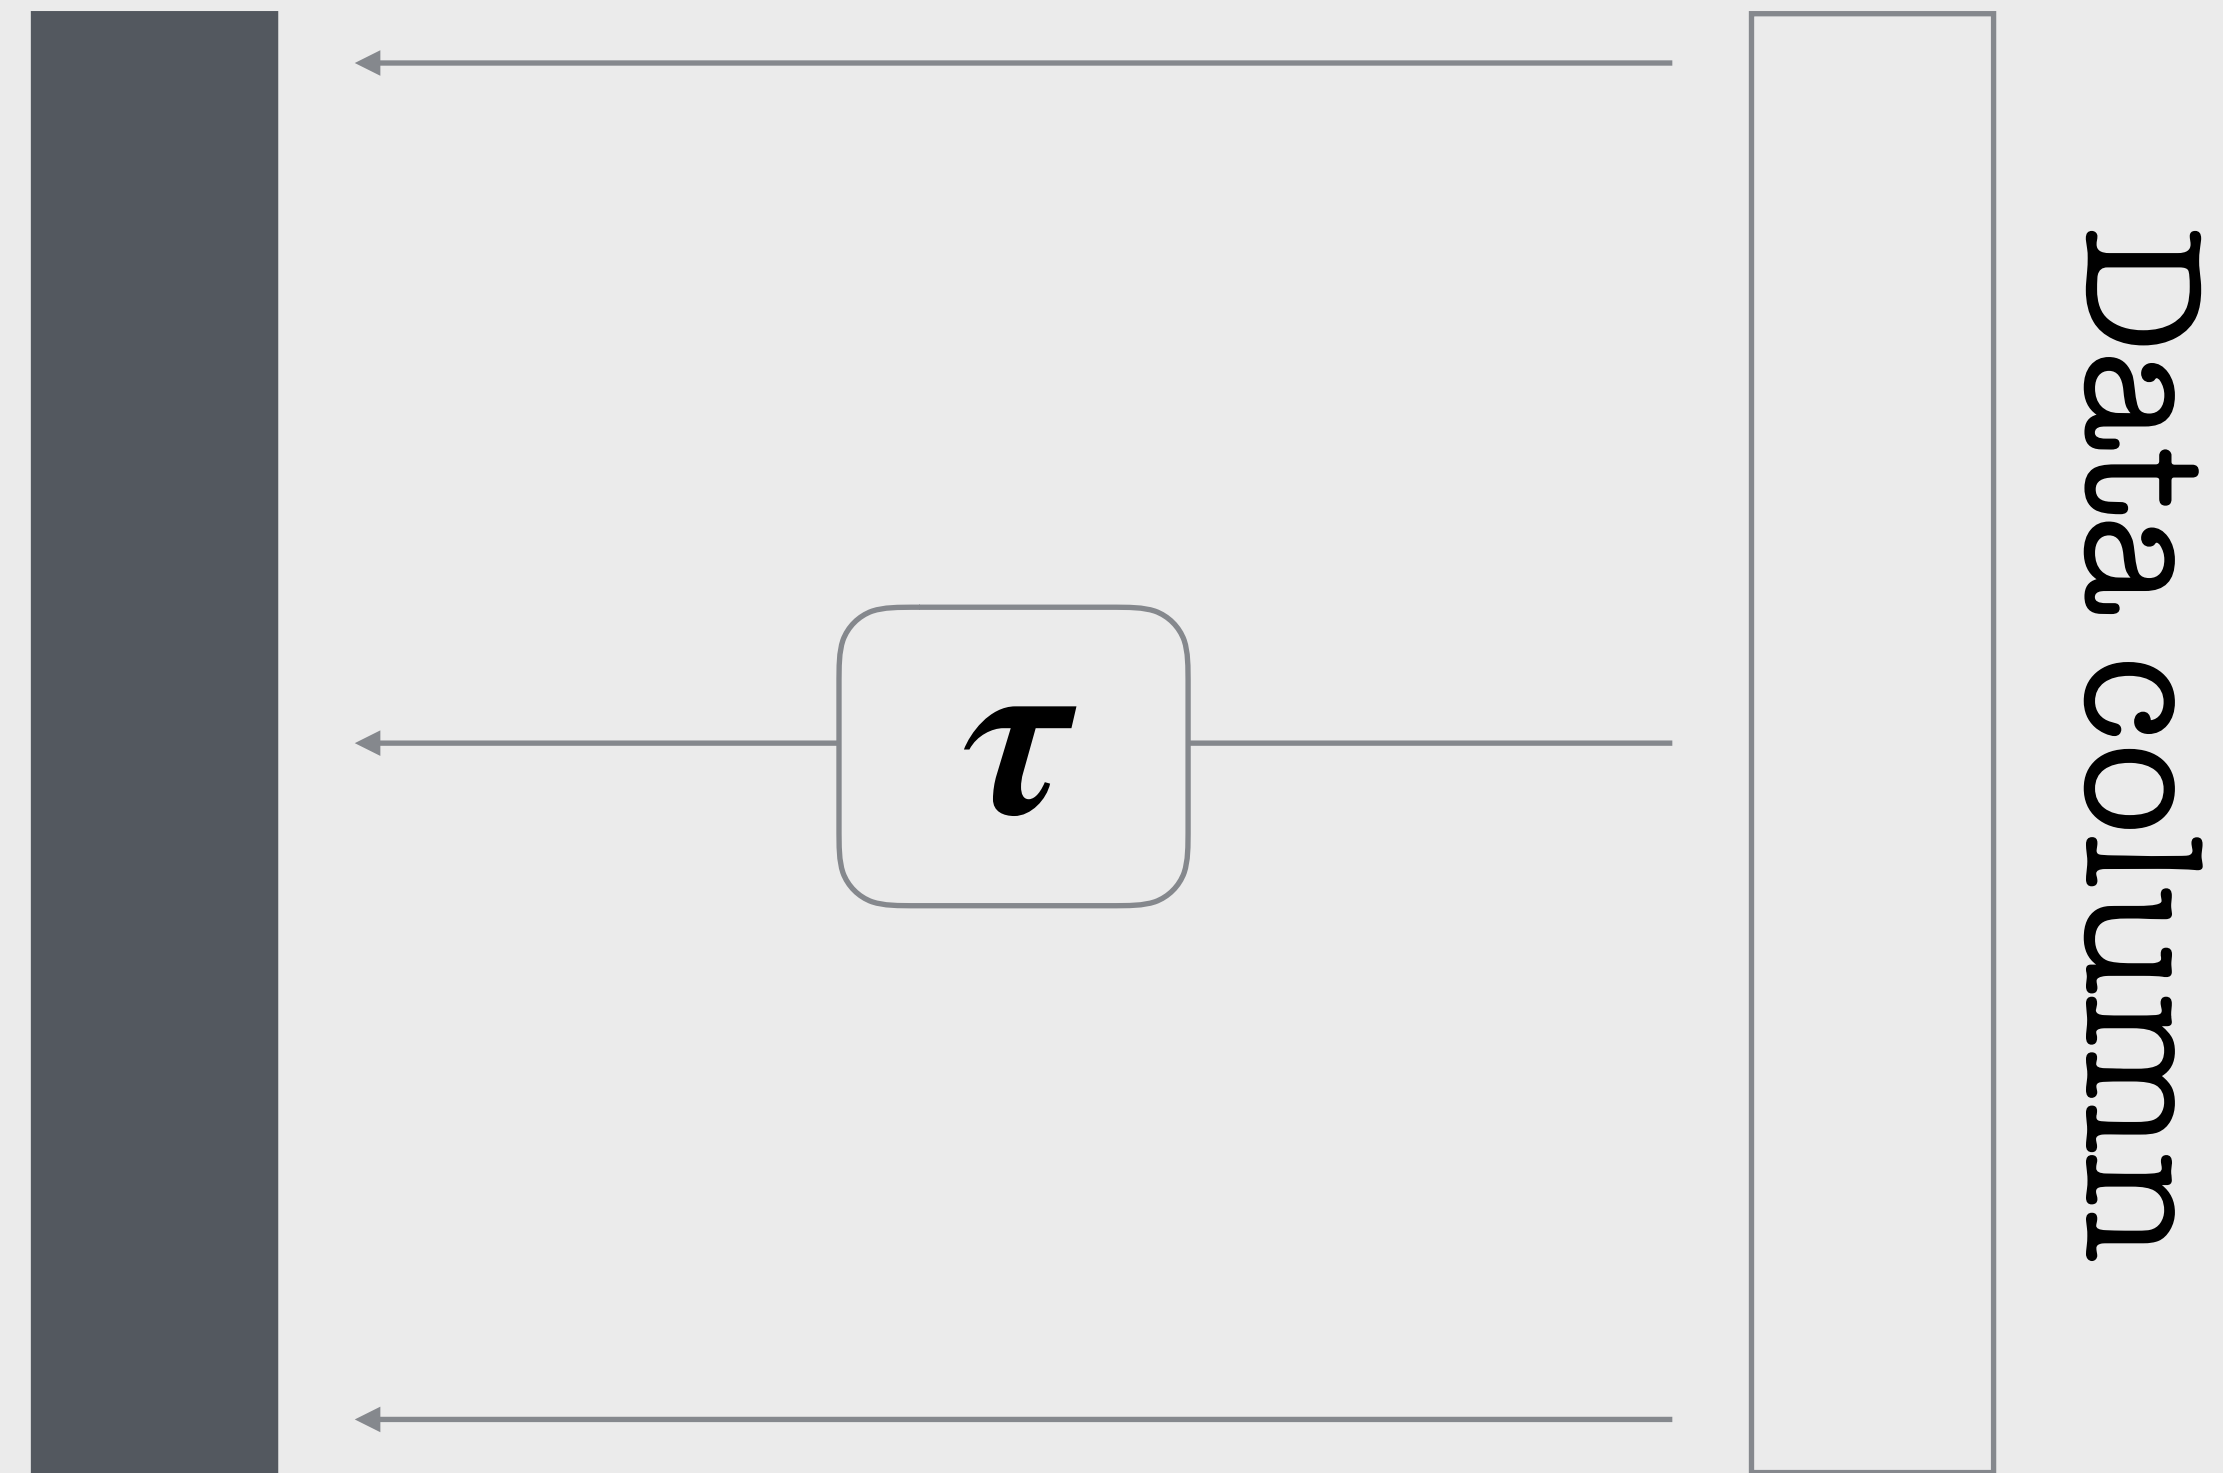
# A basic aggregate




Data column

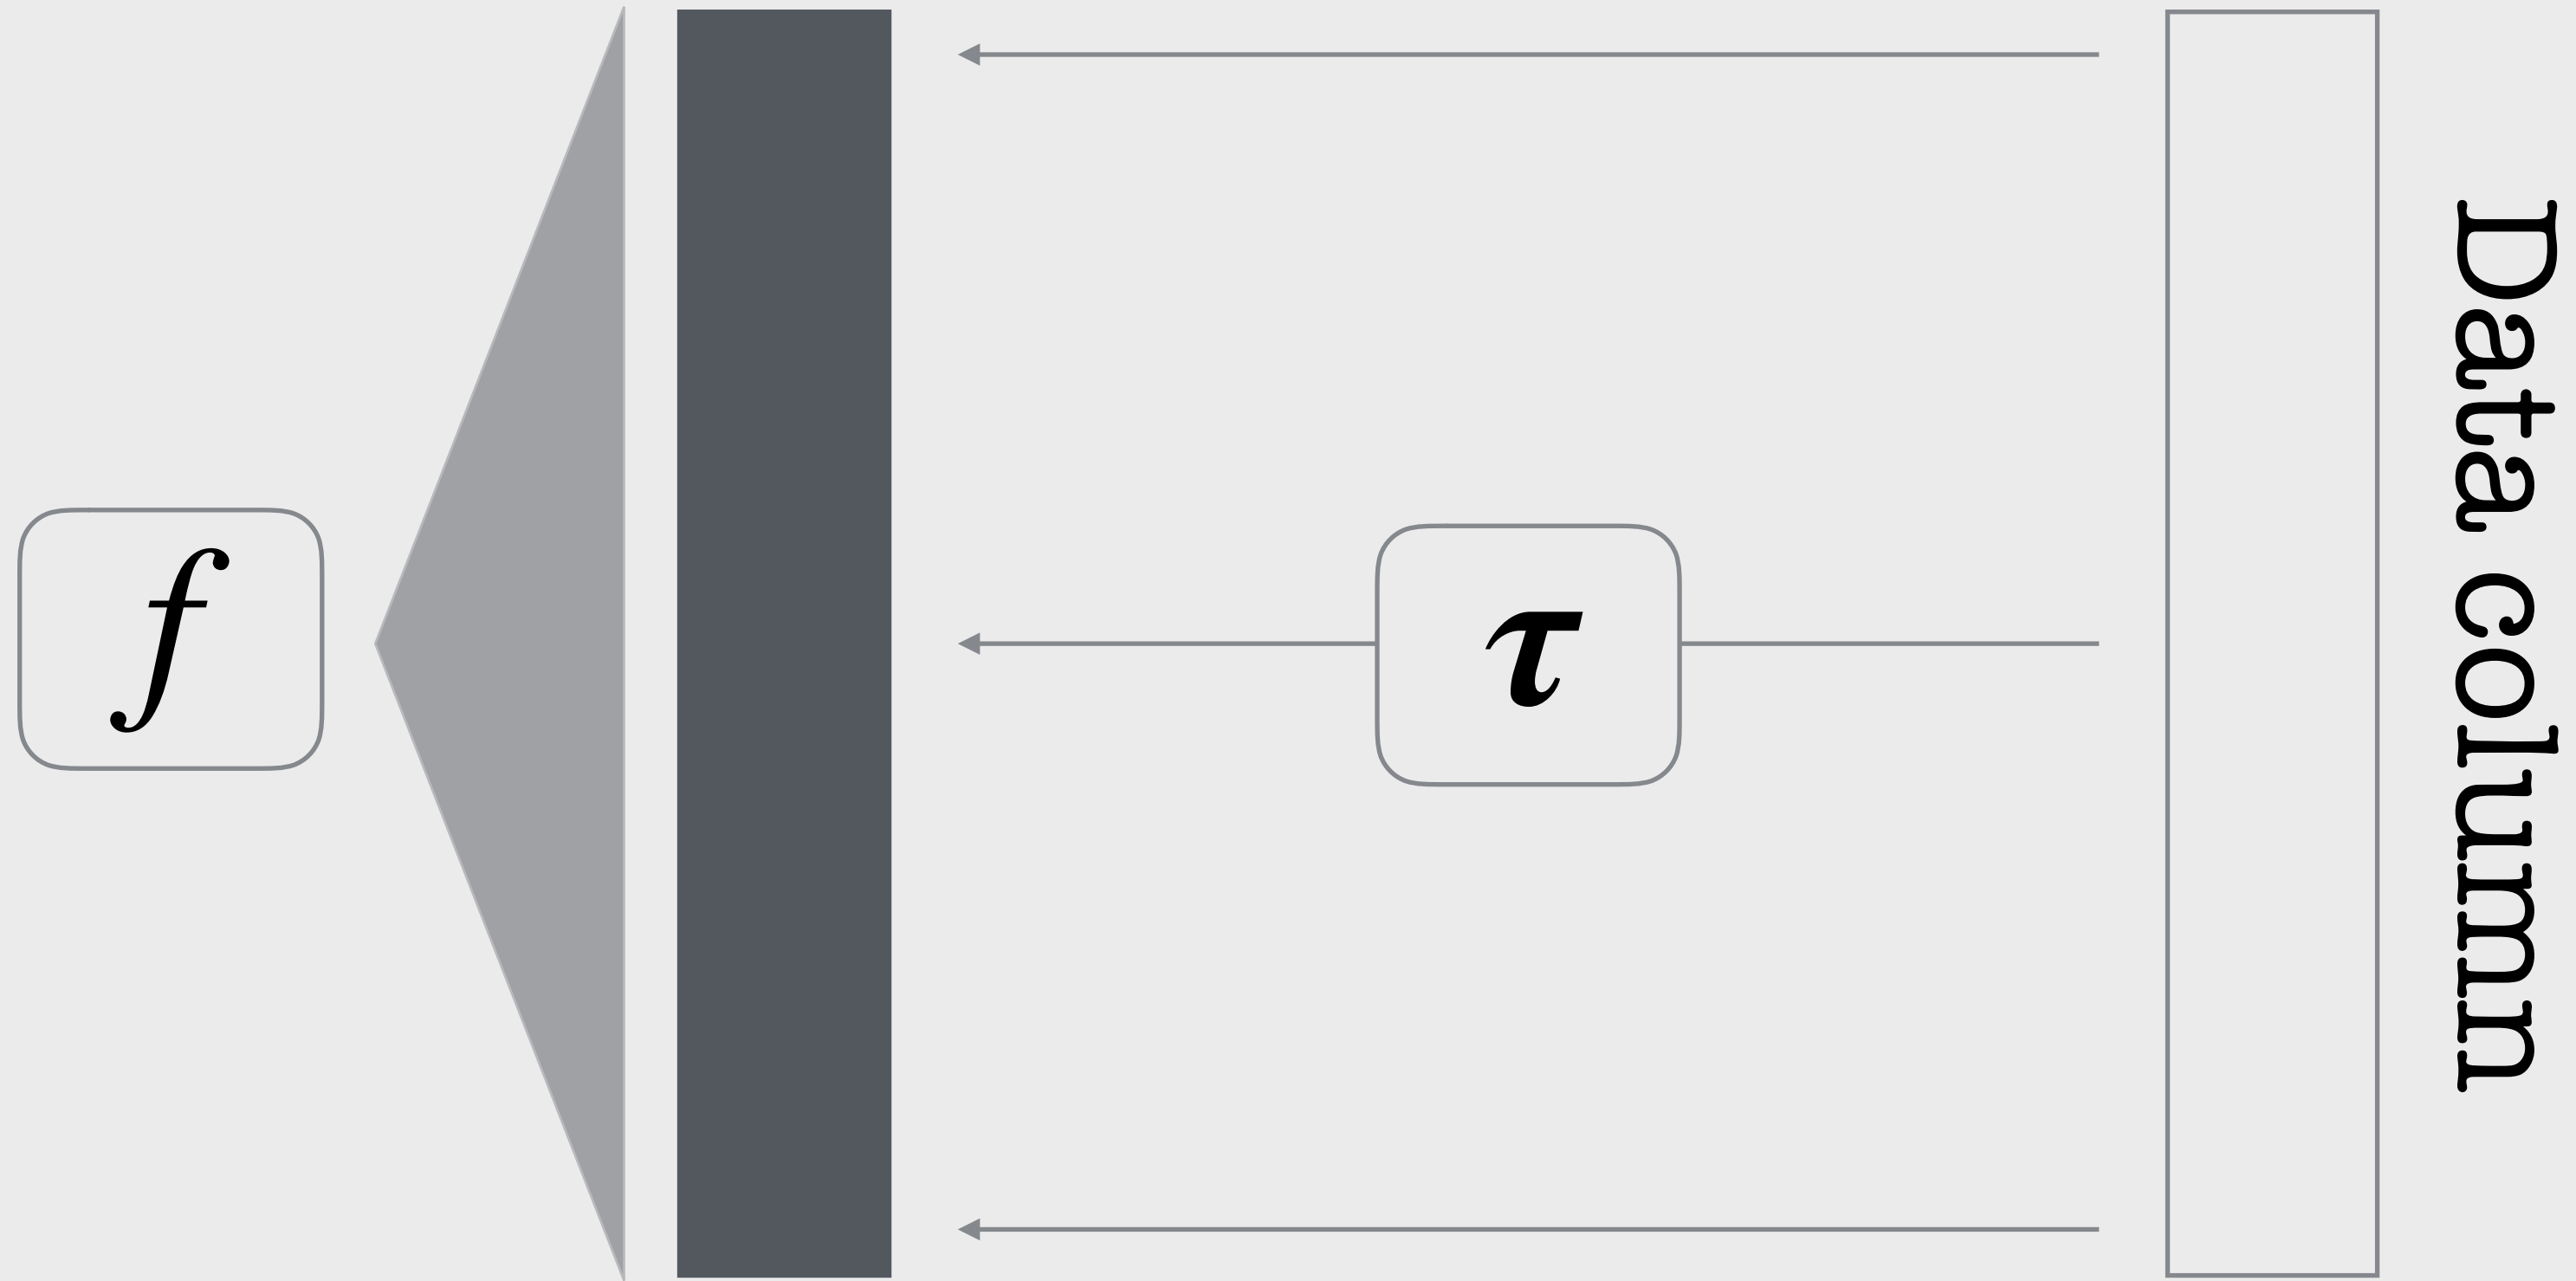



# A basic aggregate






# A basic aggregate







# A basic aggregate




$$\tau(x) = x \mid \tau(x) = x^2$$




# A basic aggregate



$$f(\text{[long bar]}) \rightarrow f(\text{[short bar]}) f(\text{[short bar]}) f(\text{[short bar]})$$

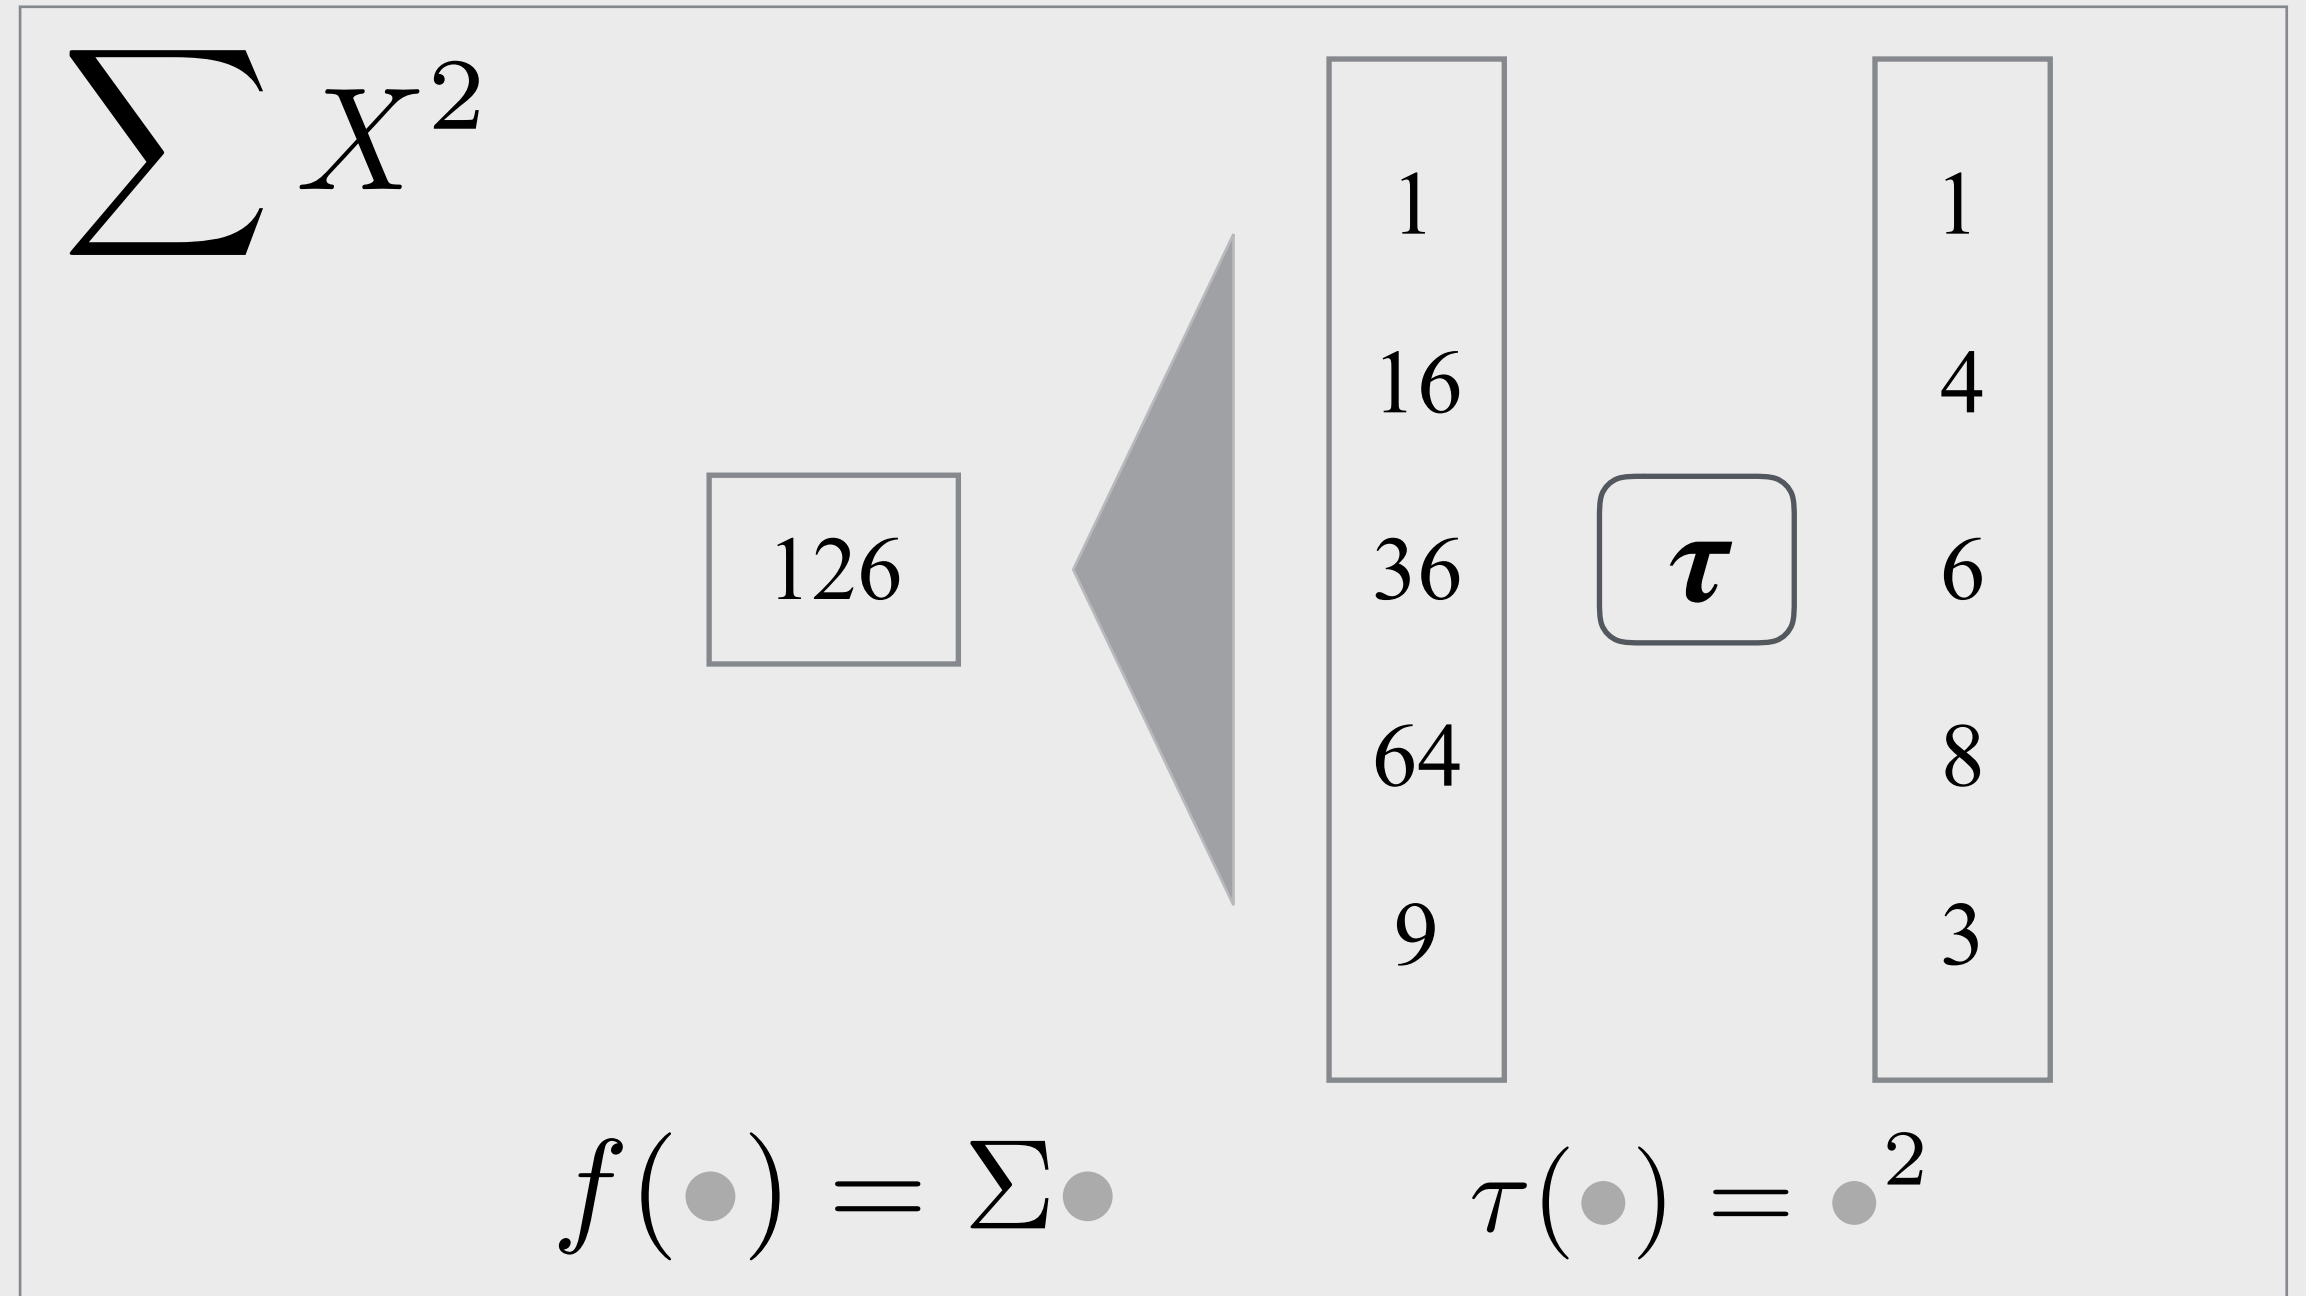




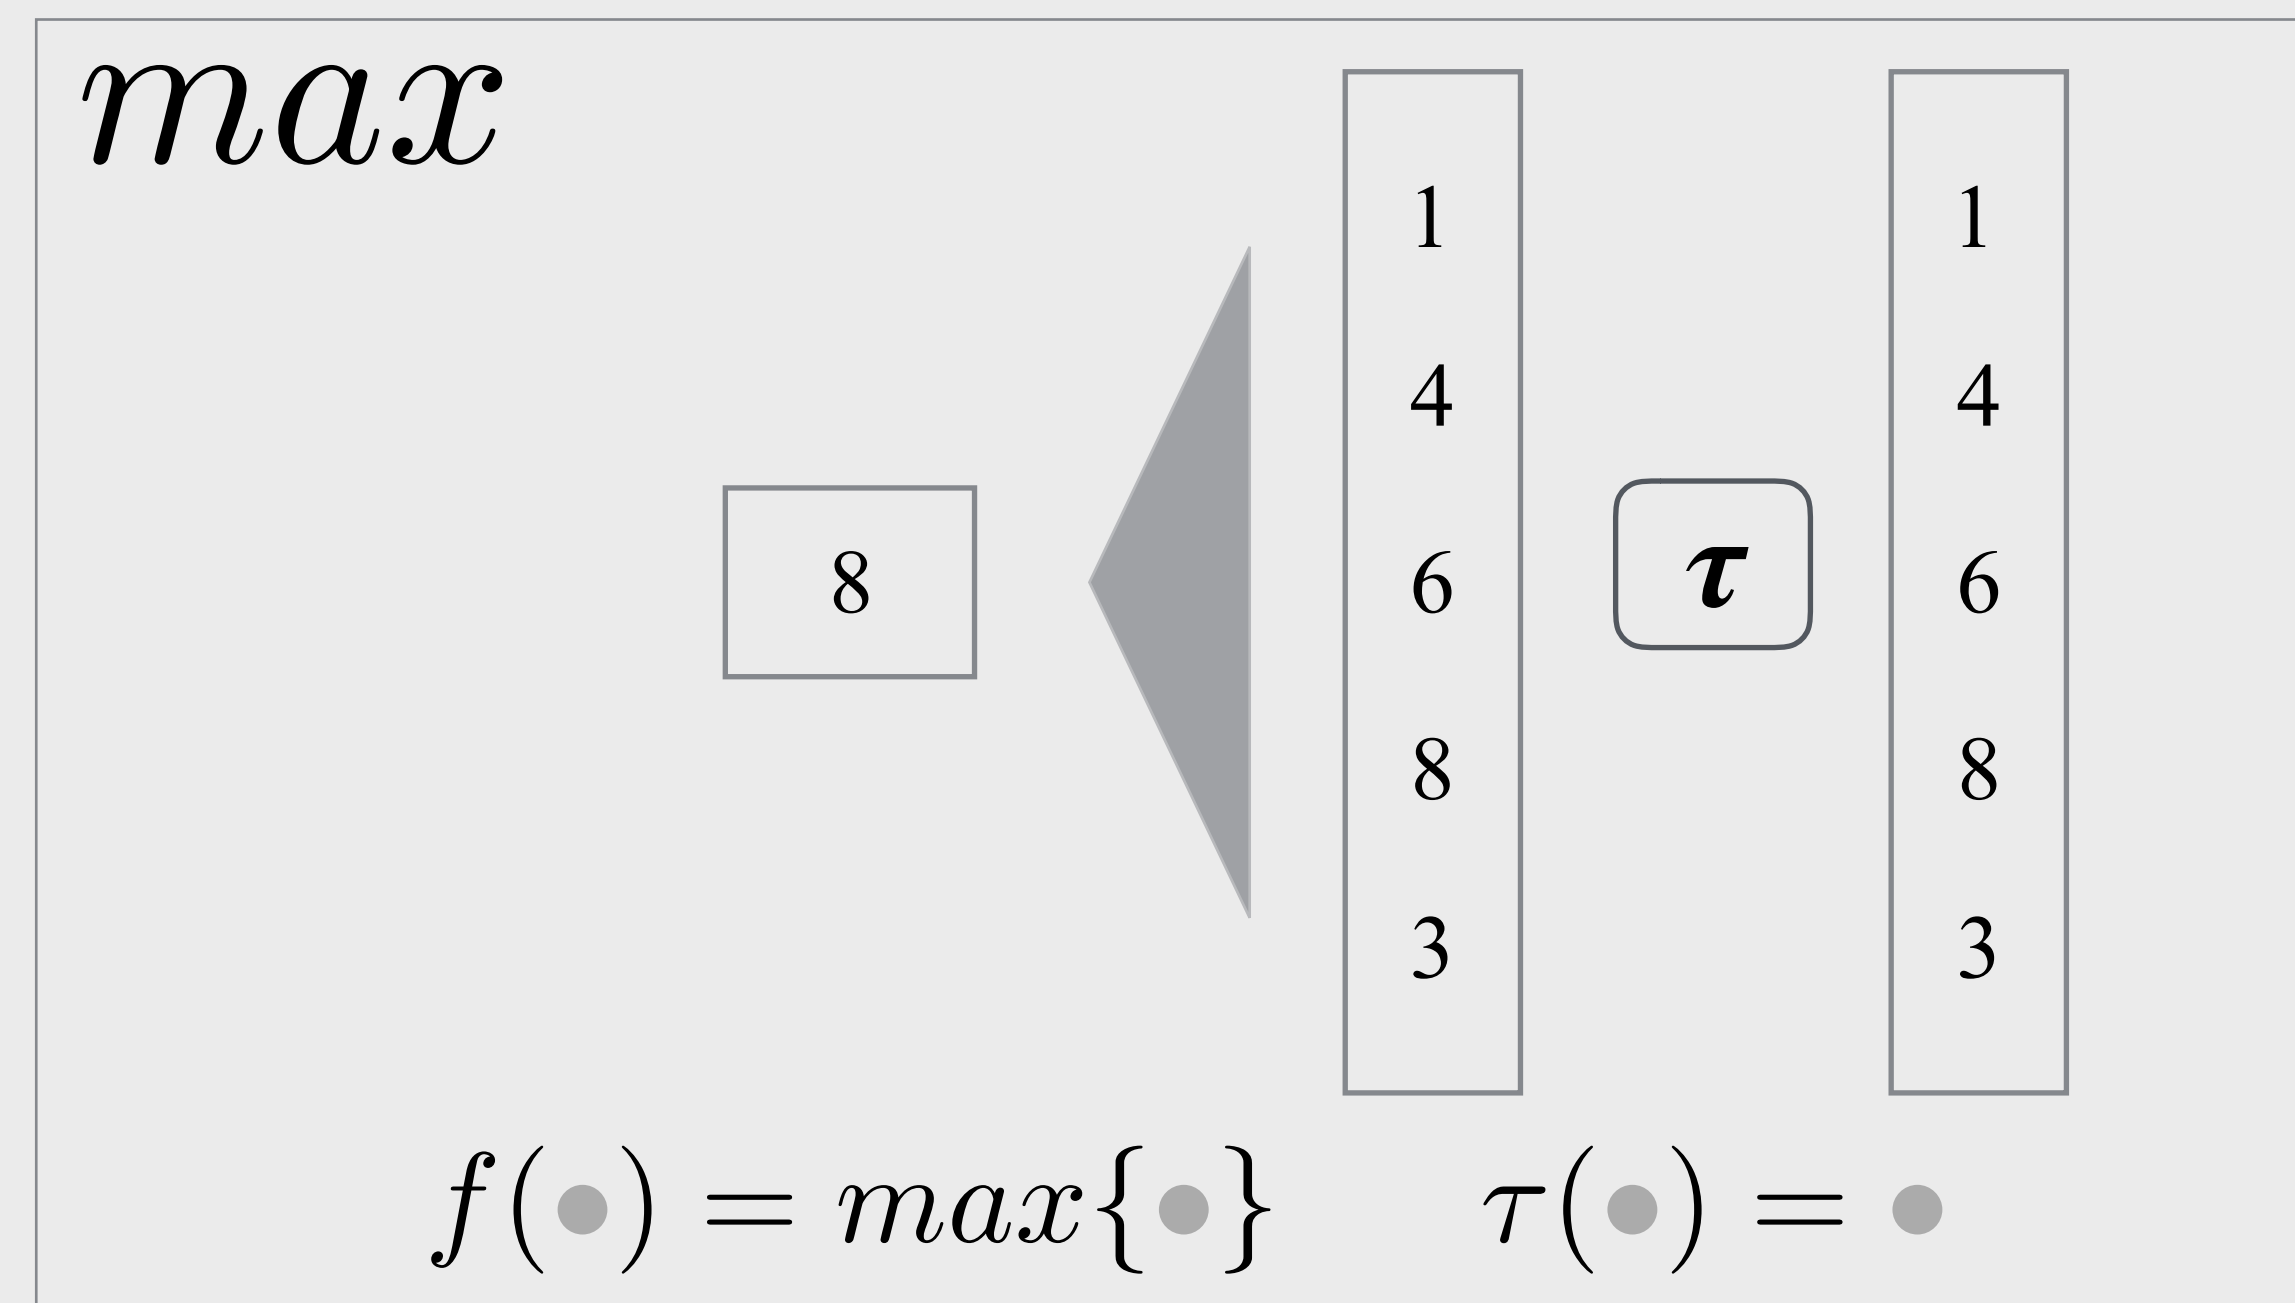


# A basic aggregate



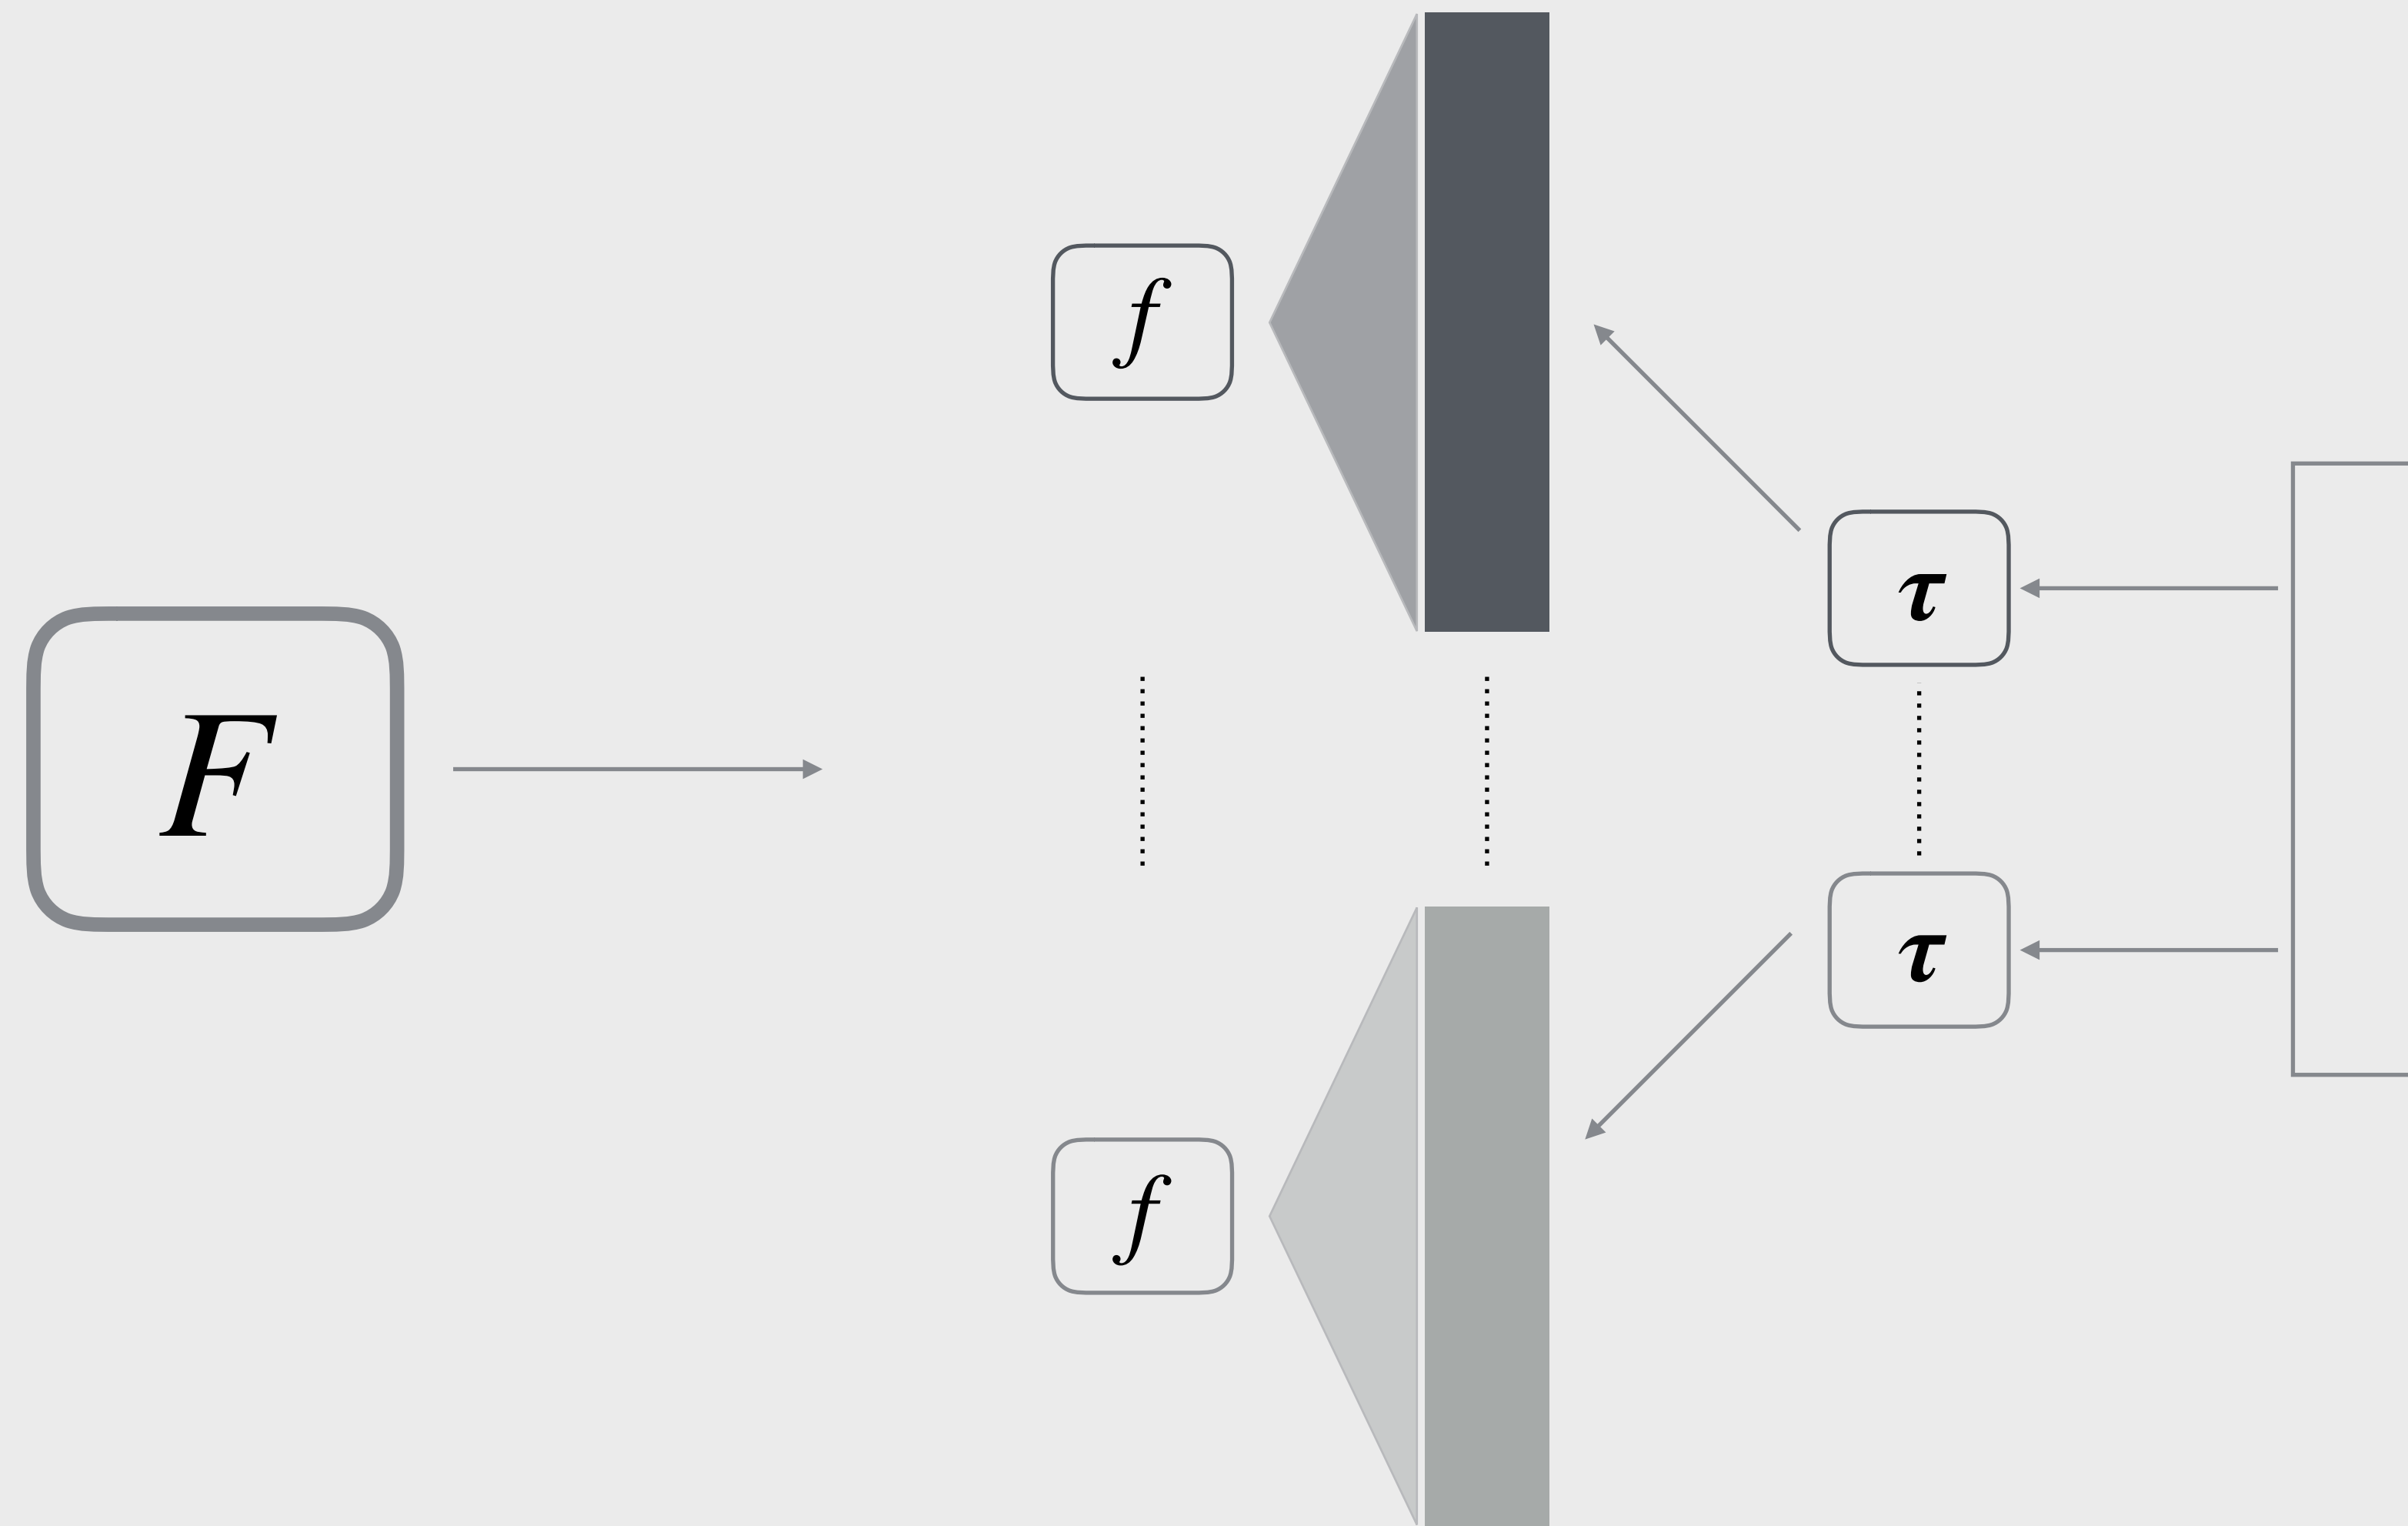
$$f(\text{[Data Column]}) \rightarrow f(\text{[Partition 1]}) \quad f(\text{[Partition 2]}) \quad f(\text{[Partition 3]})$$


*Example: max, min, sum, product*





# Statistics are decomposed into building blocks

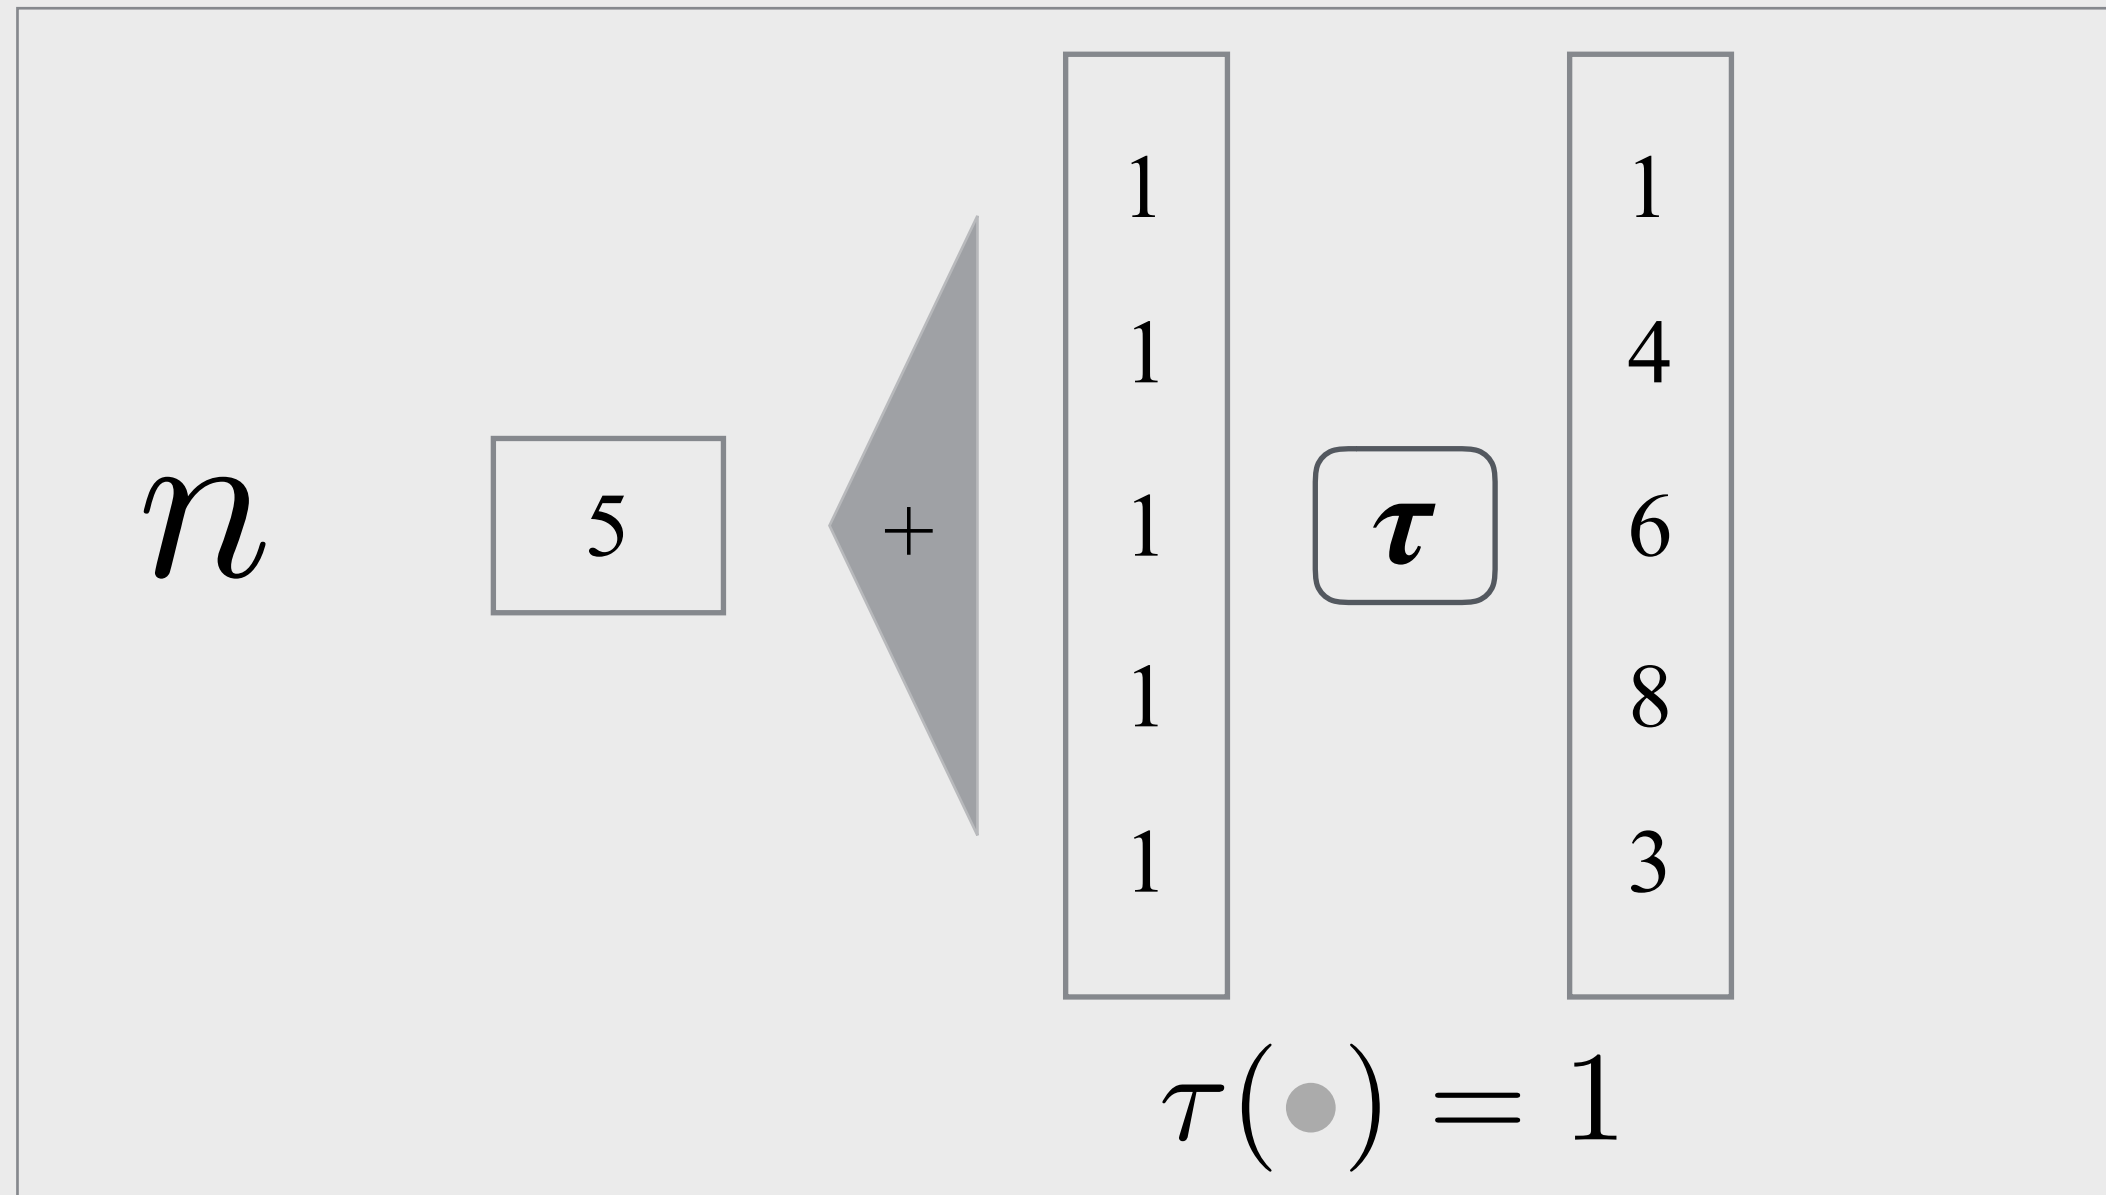
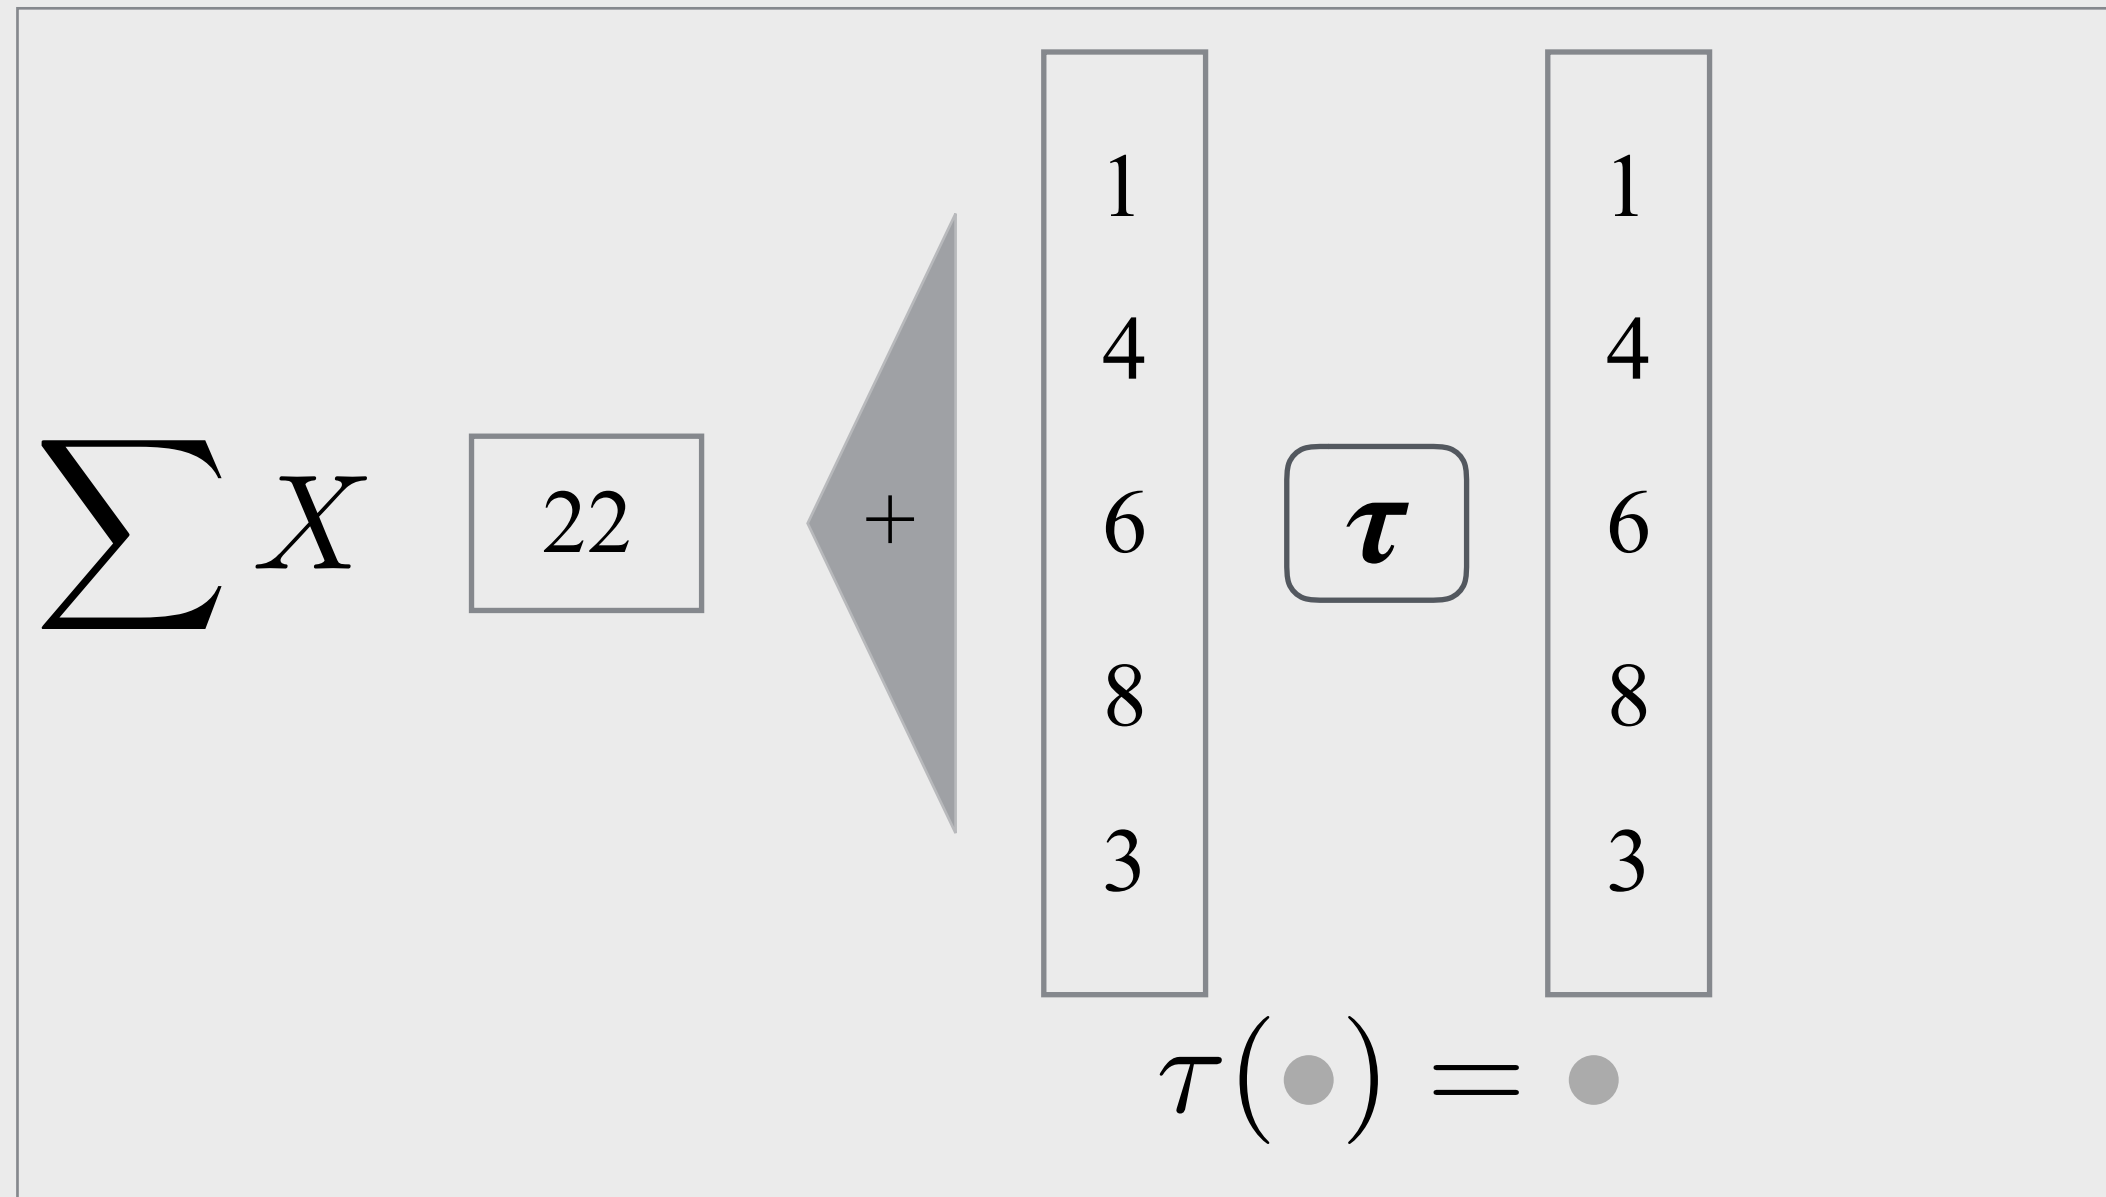





$$\frac{1}{n} \sum X$$

Arithmetic Mean

4.4


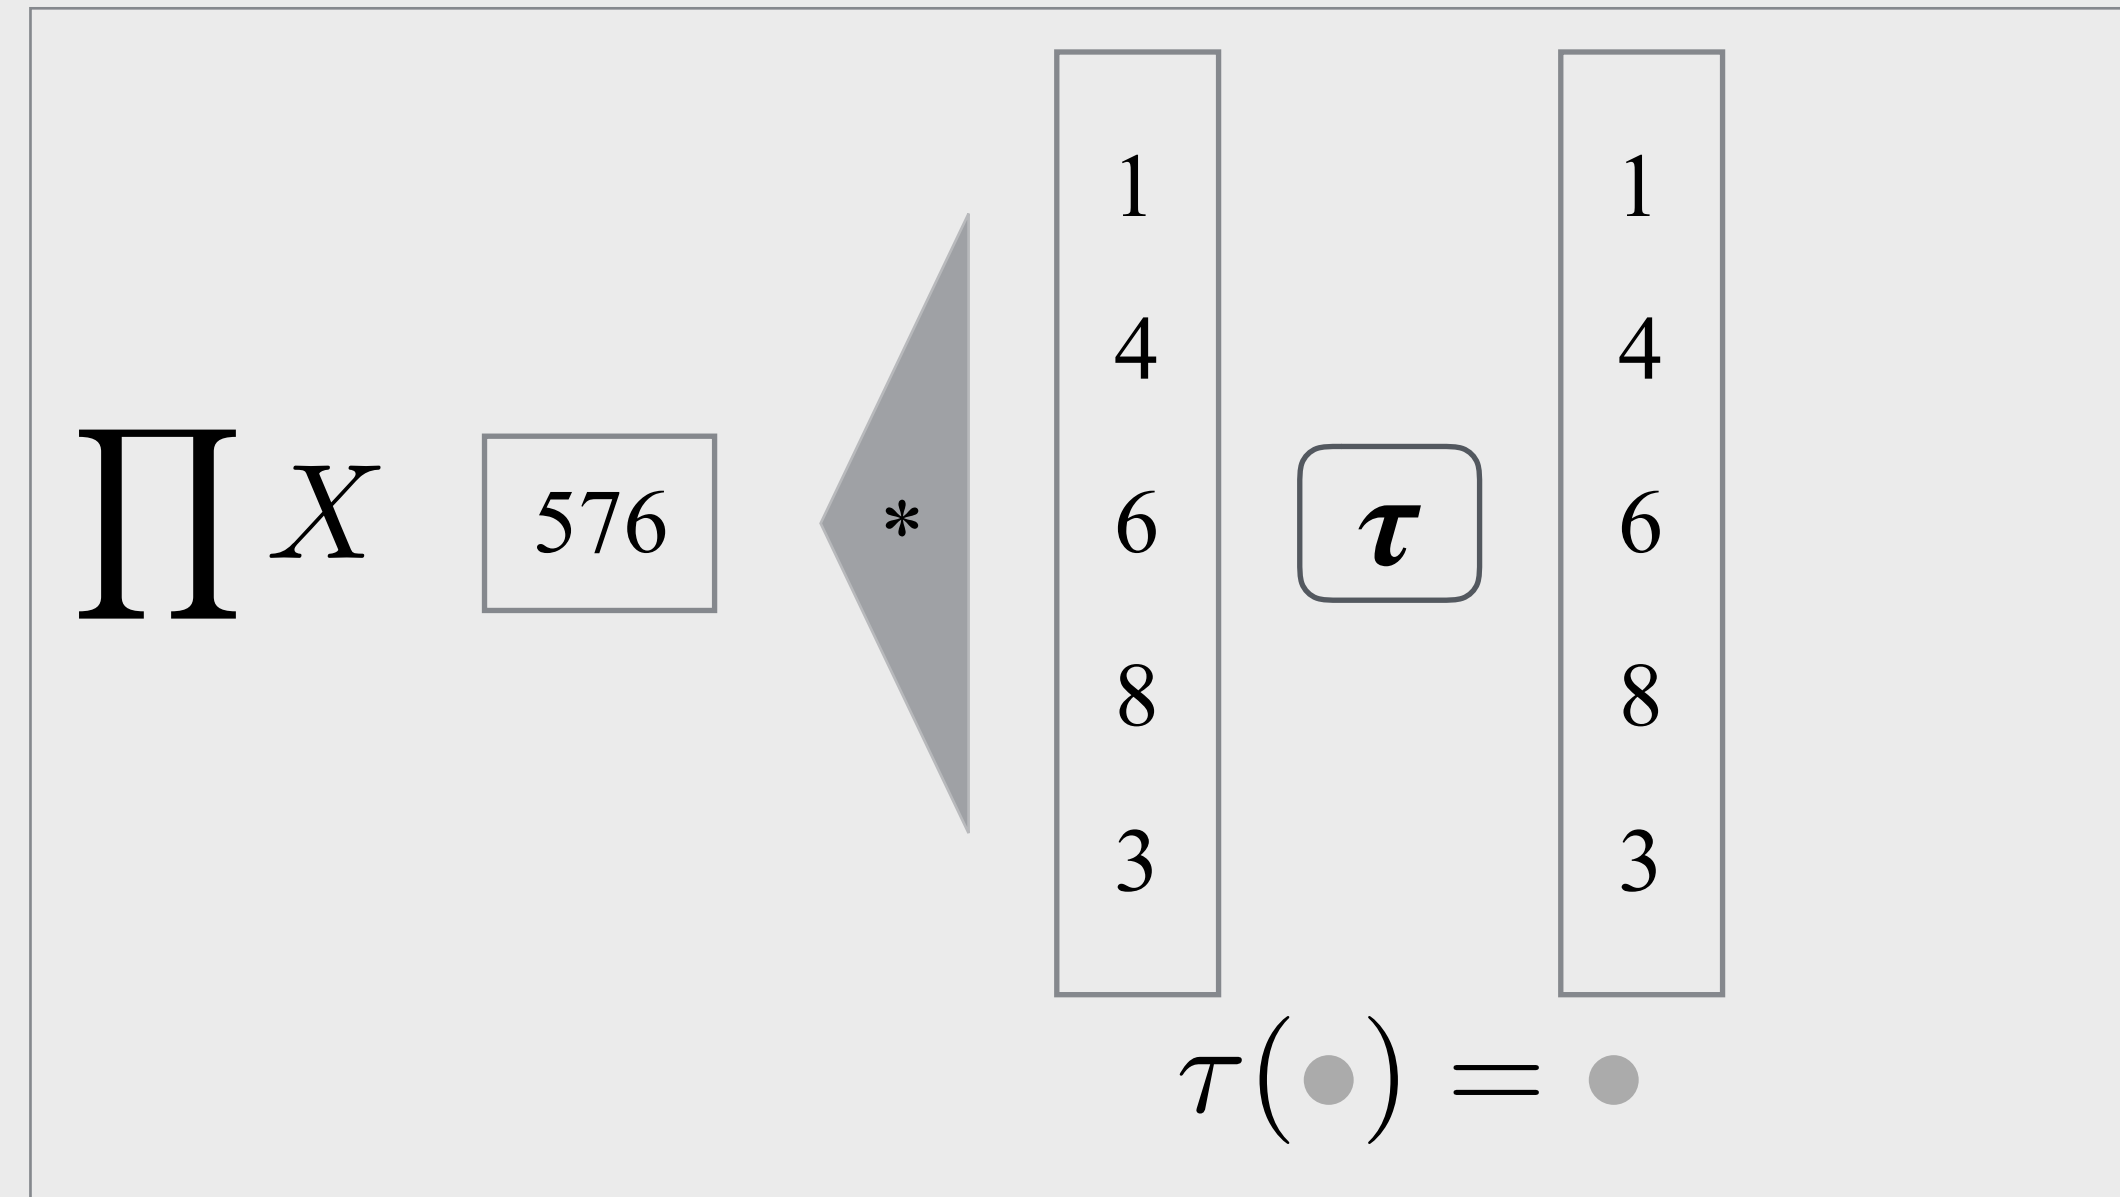


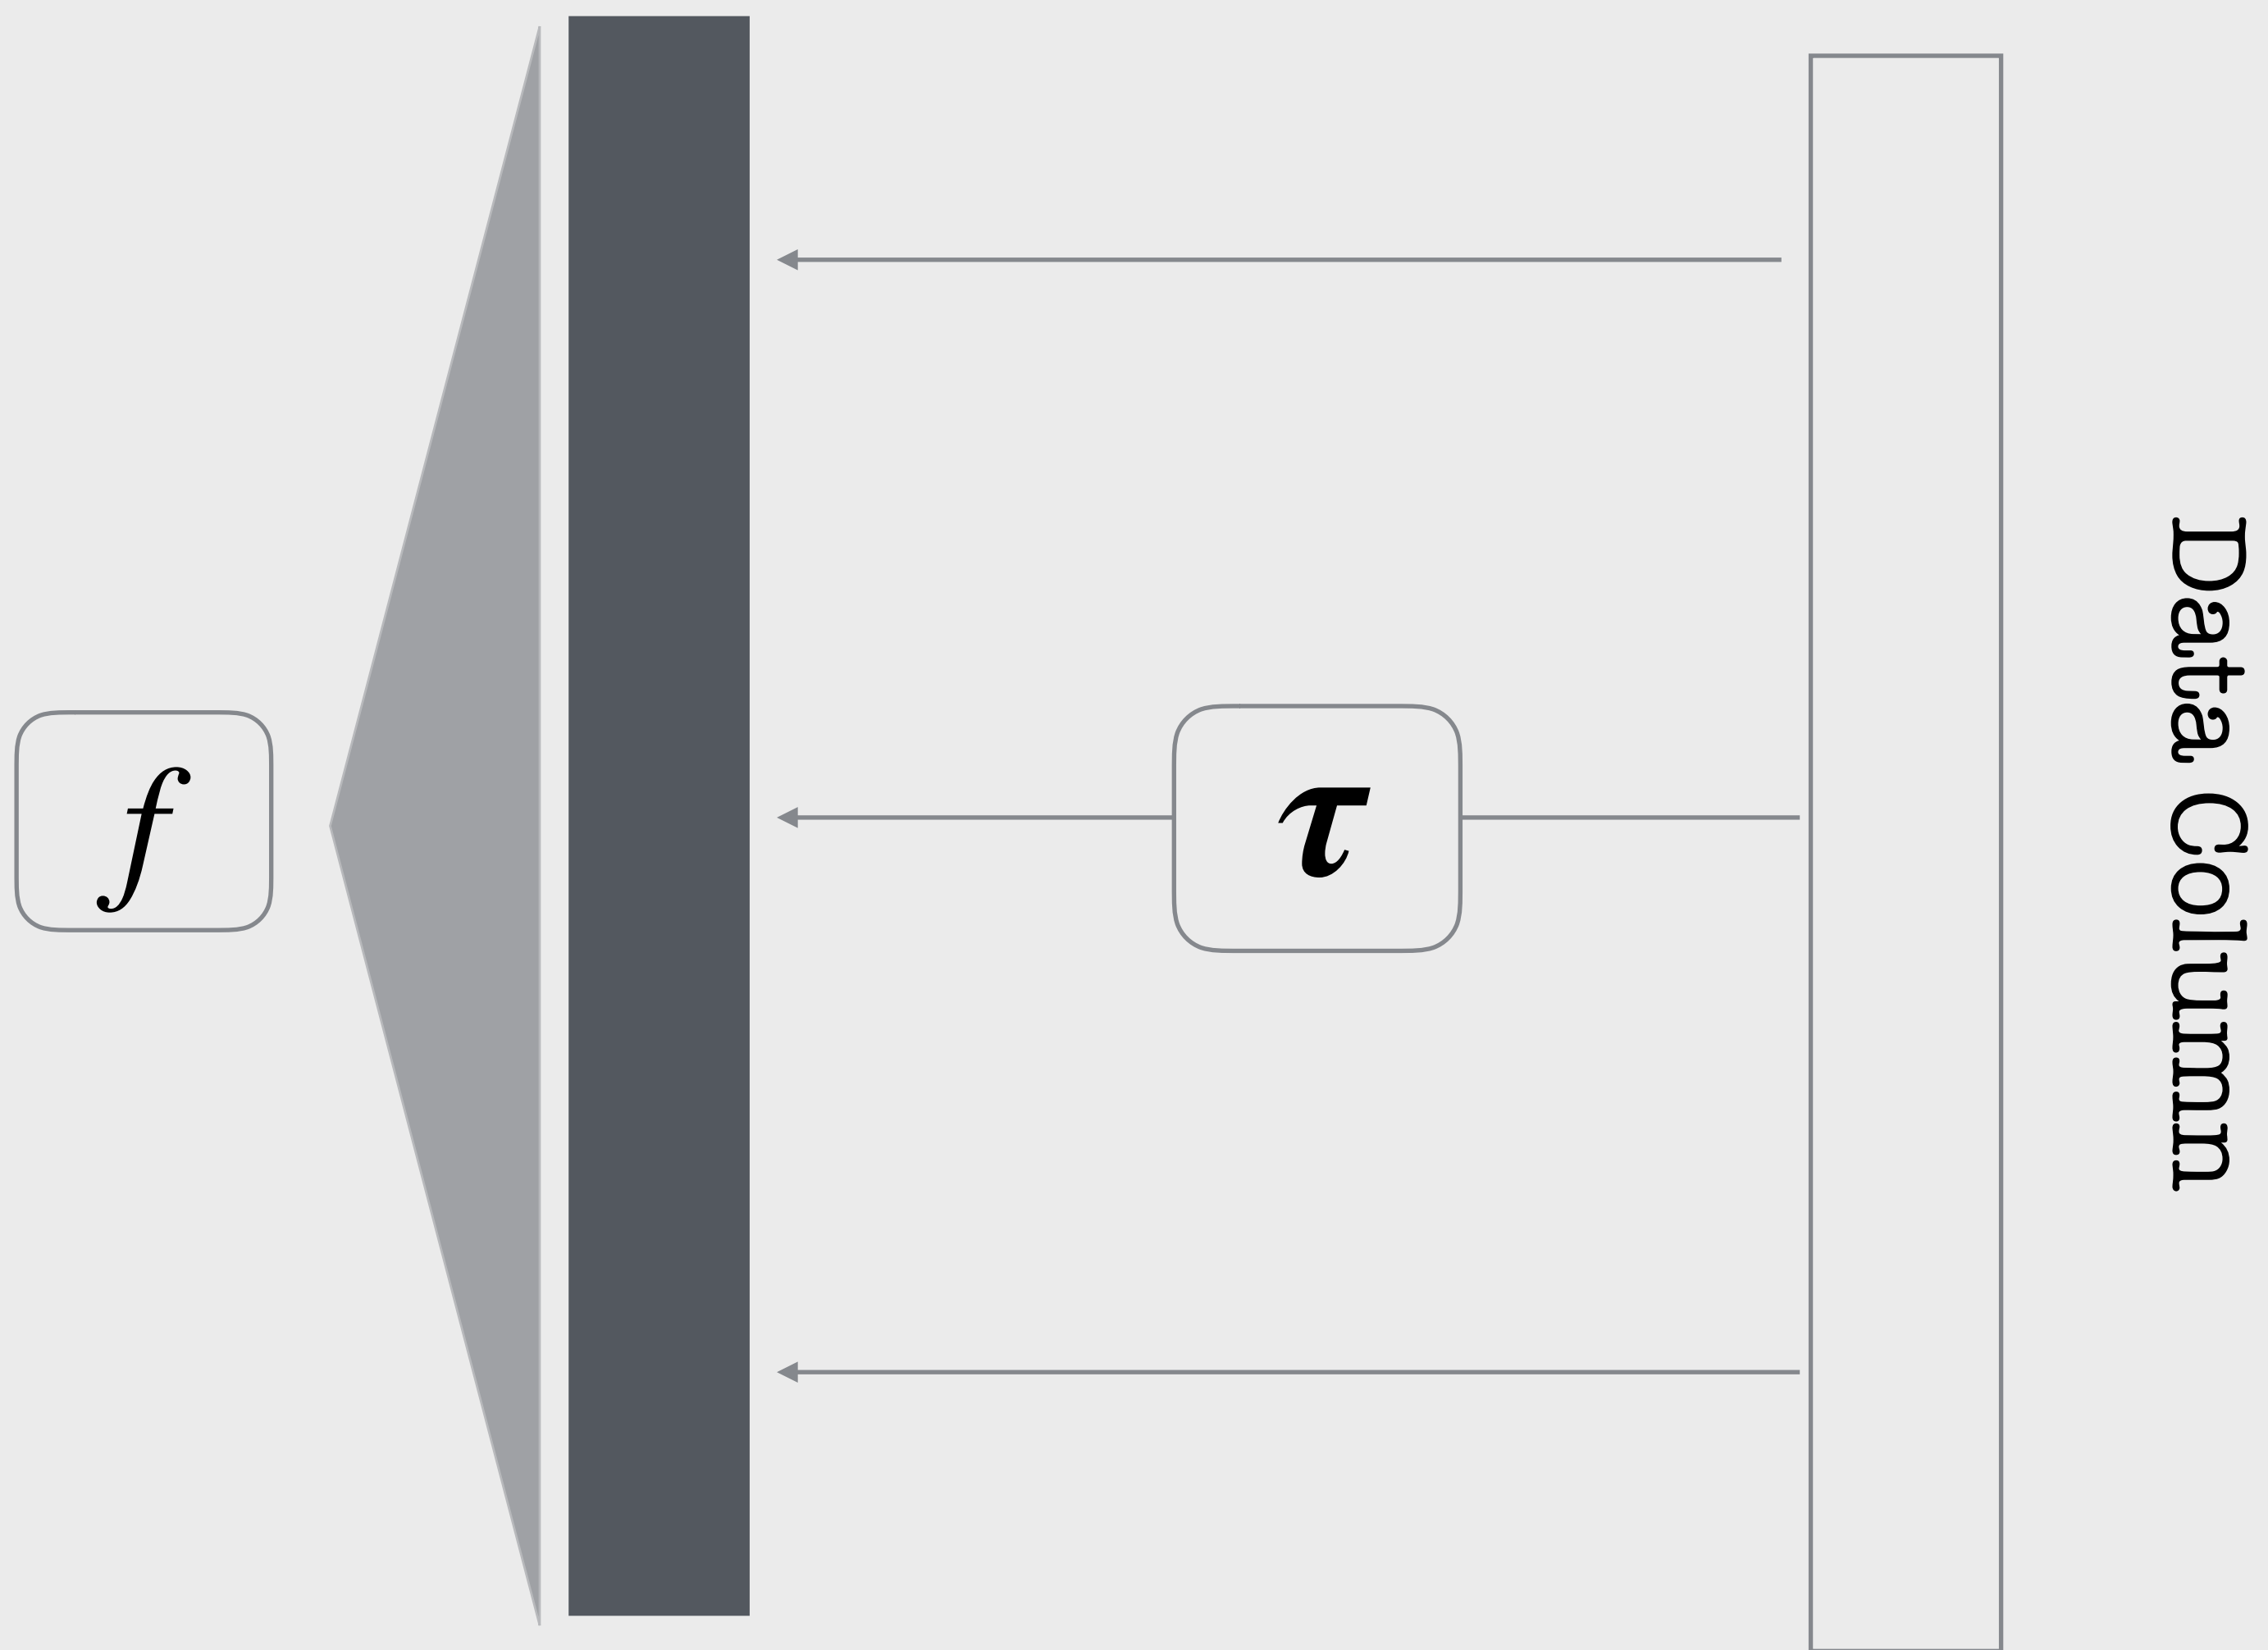






$$\left(\prod X\right)^{\frac{1}{n}}$$

Geometric Mean

3.6

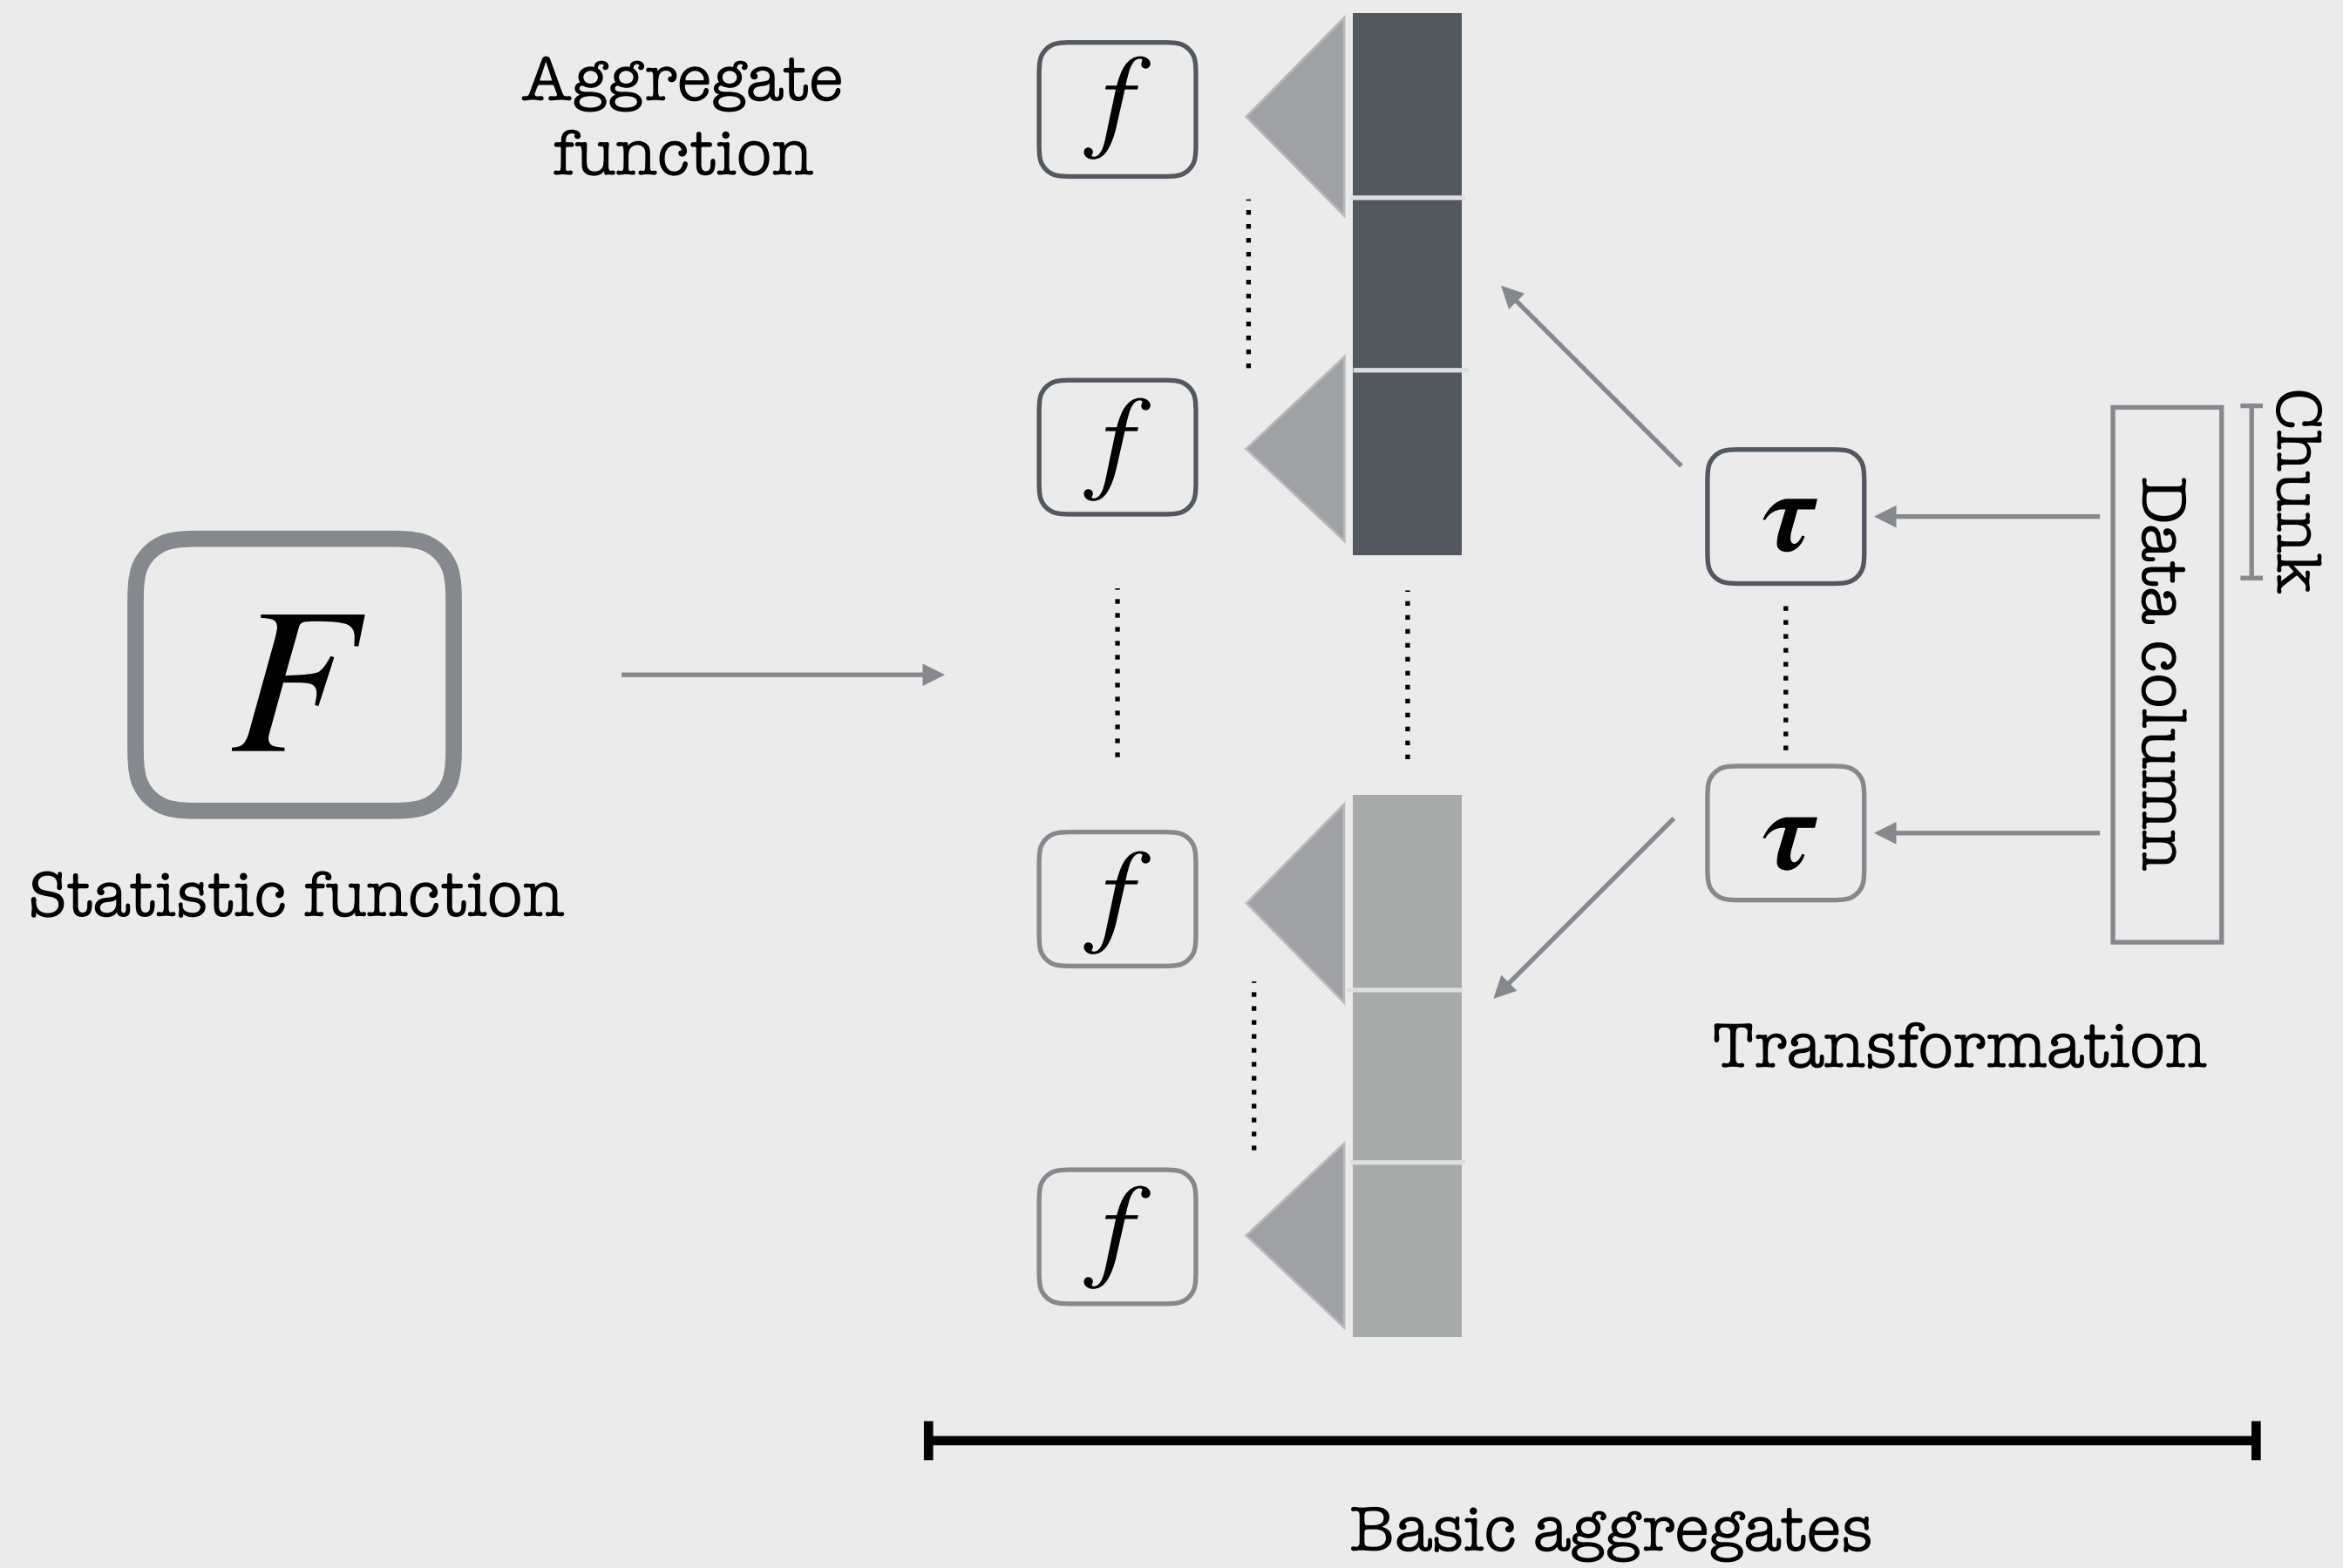







# Basic Aggregates

# Chunks



# Basic Aggregates

# Chunks

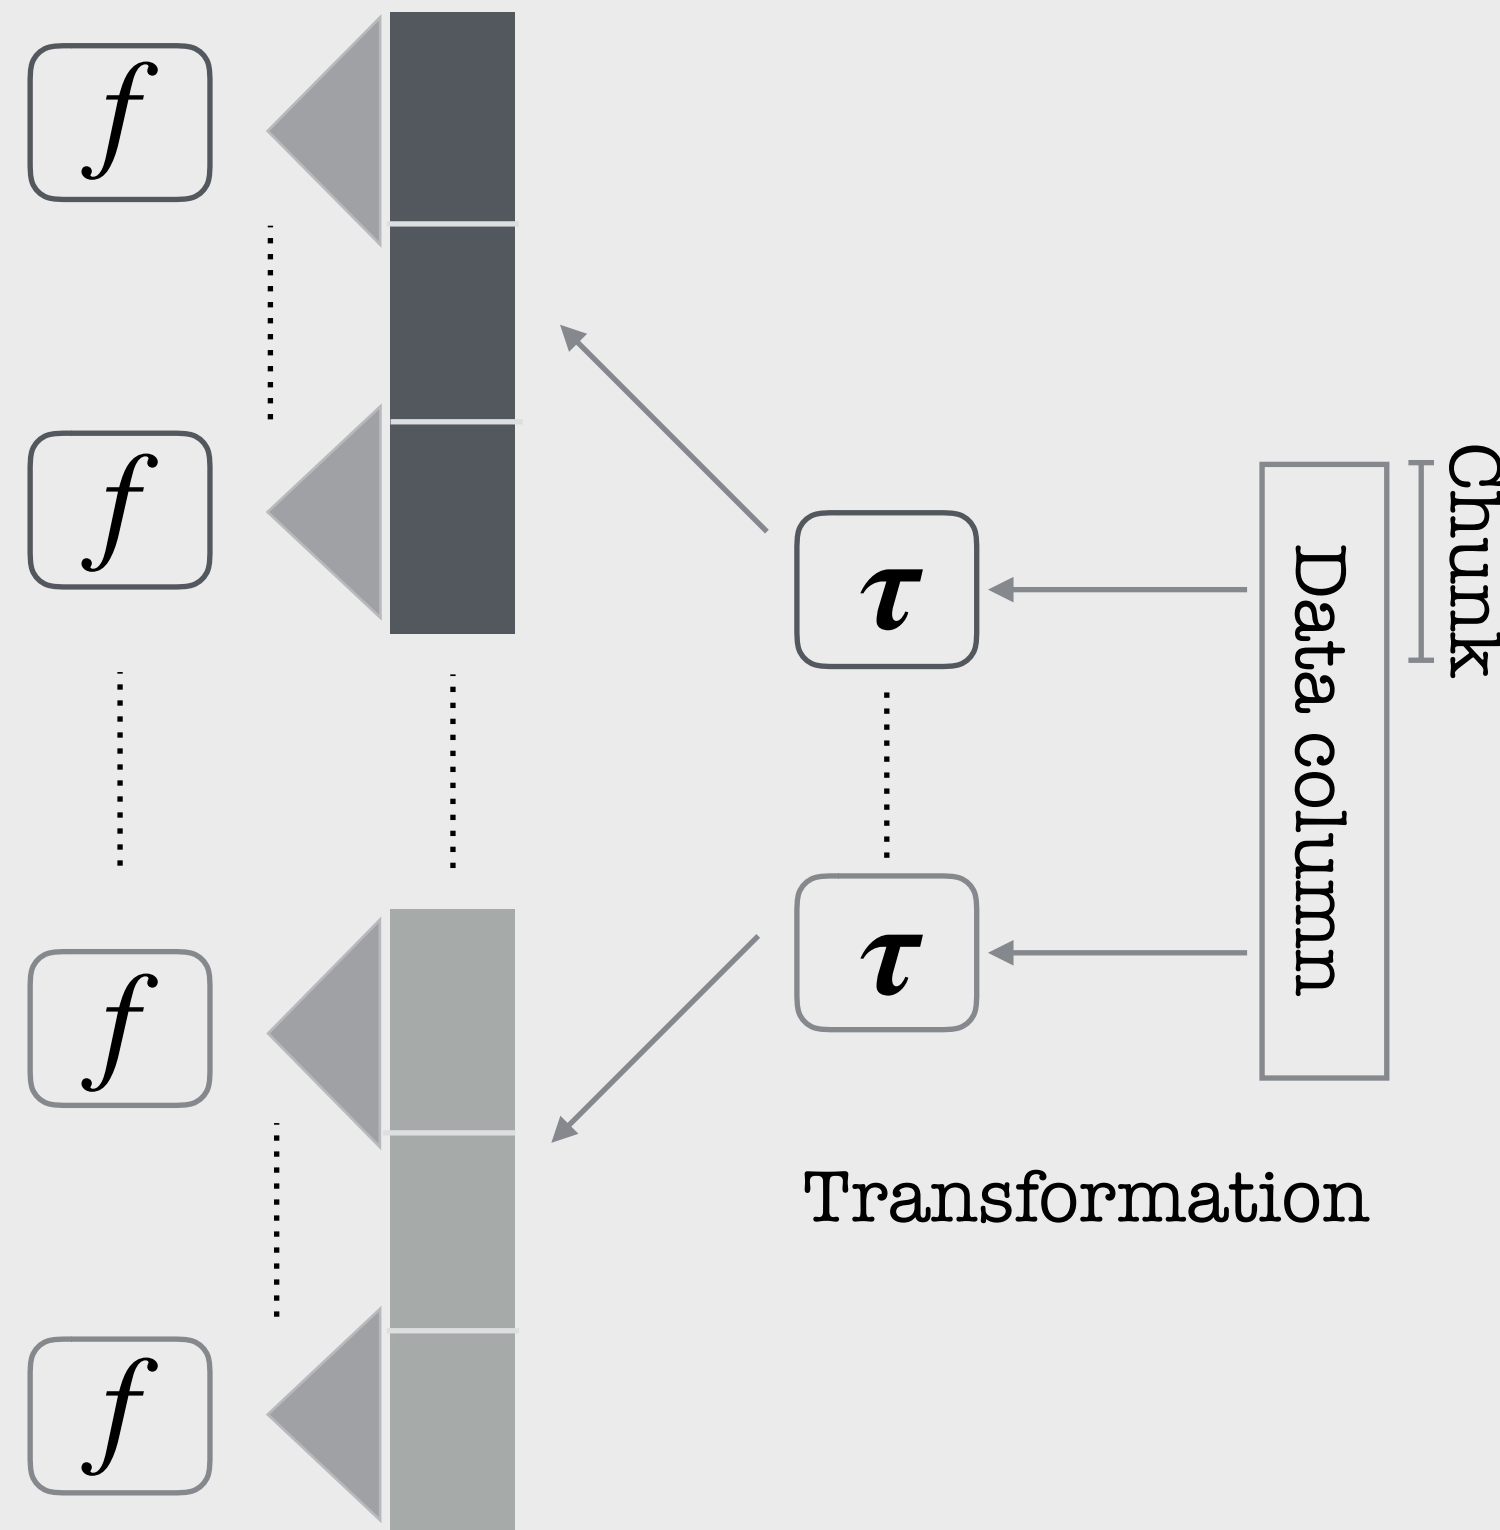


Statistic type



Statistic function

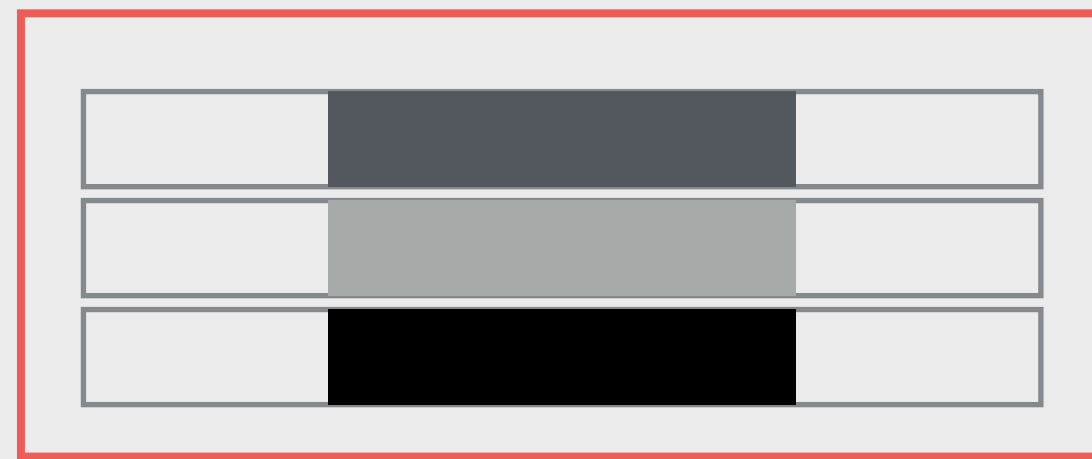
Aggregate function



Basic aggregates

# Basic Aggregates

# Chunks


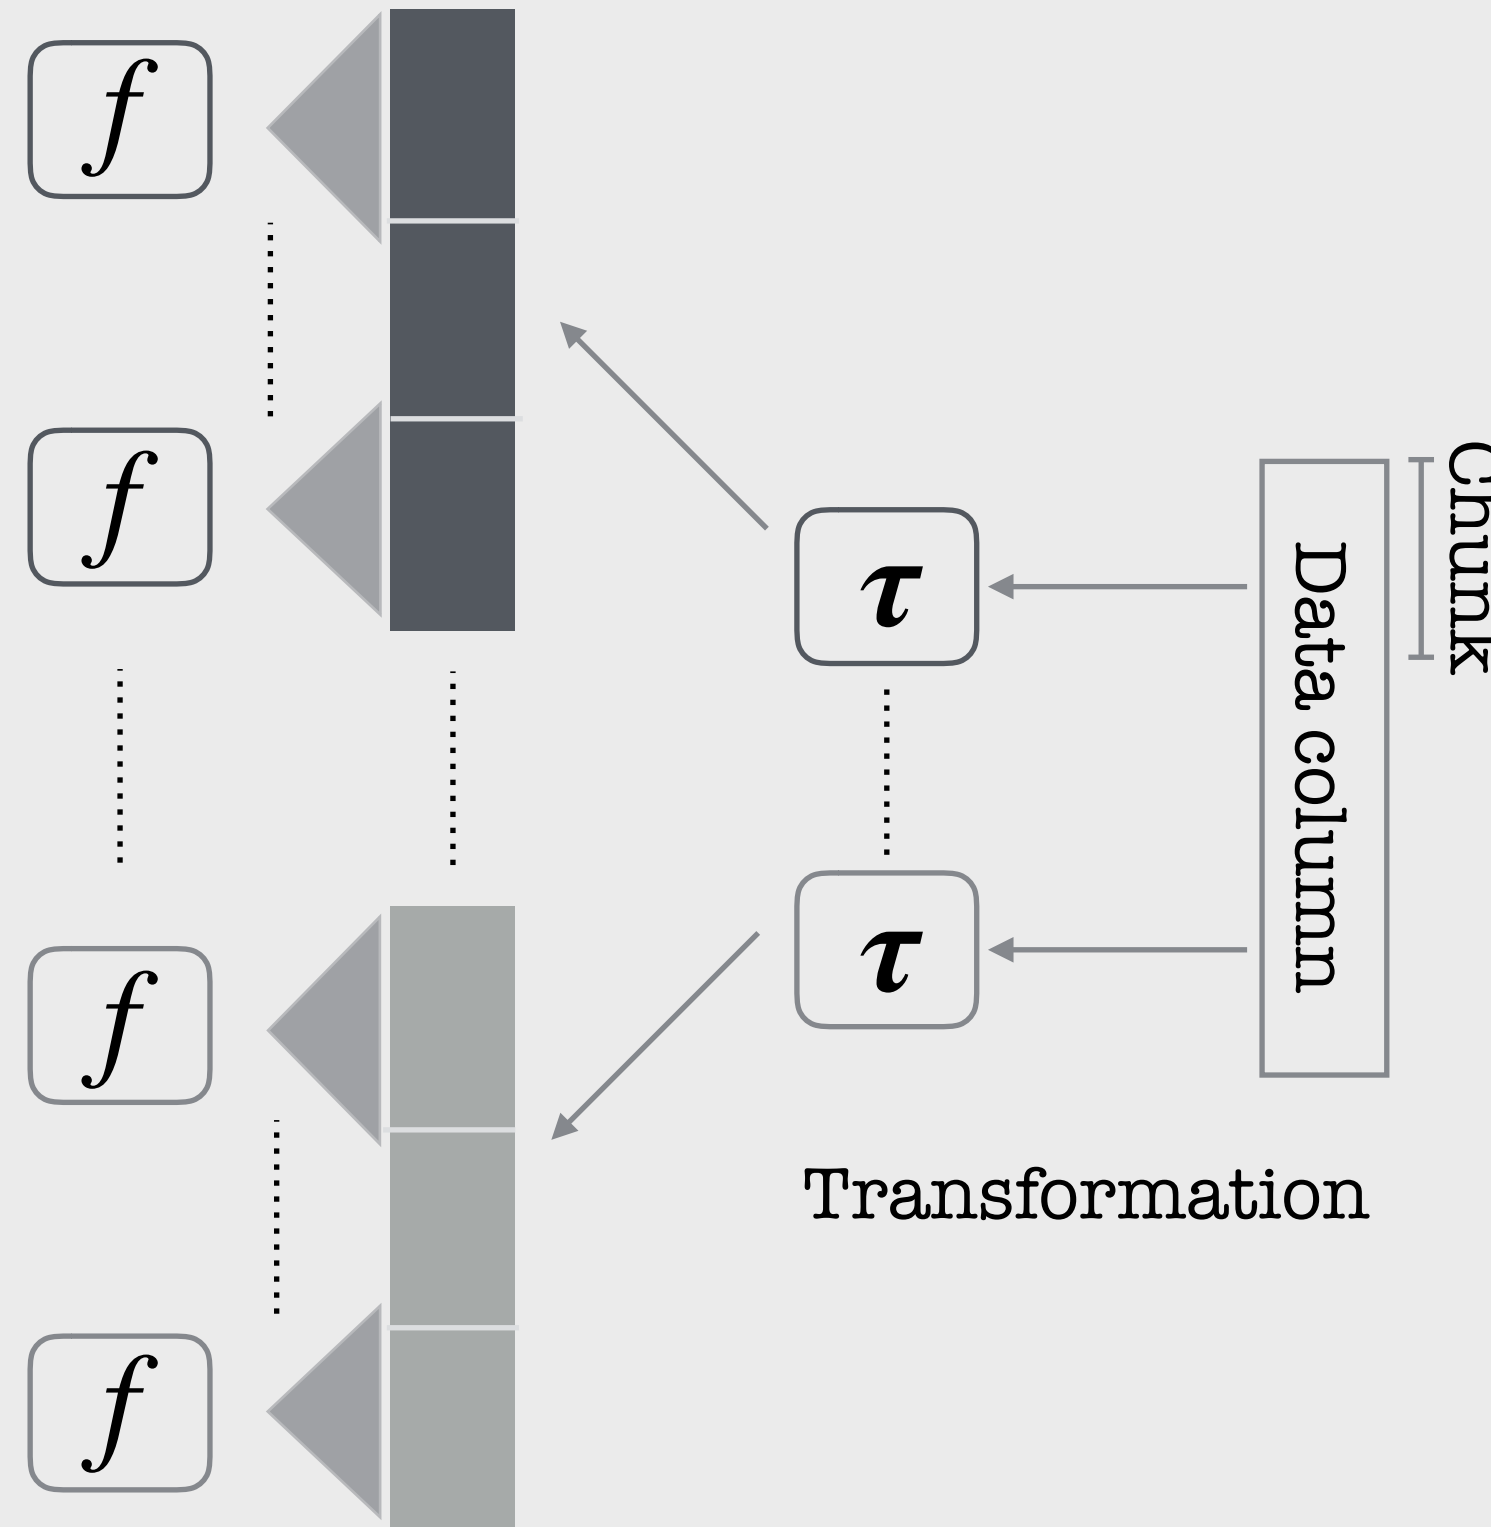


Statistic type




Statistic function

Aggregate function



Overlapping ranges

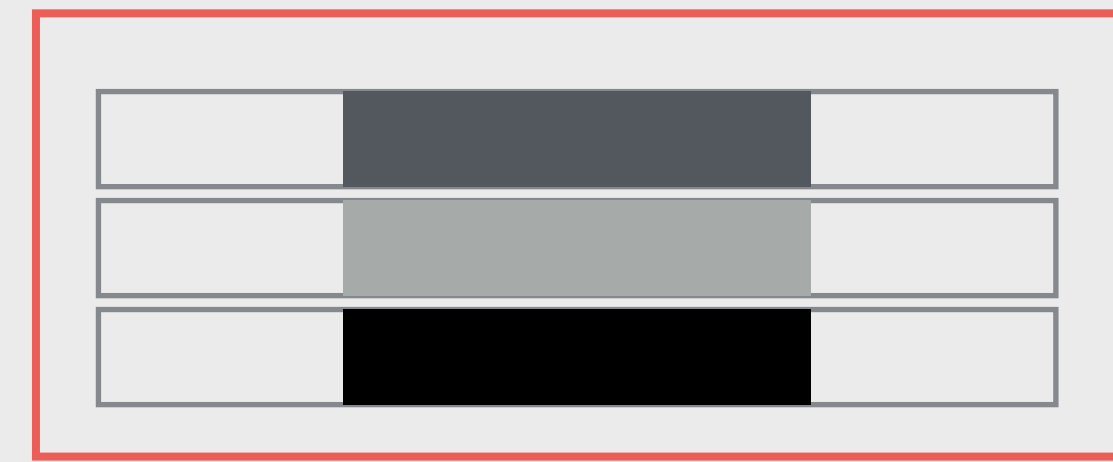


Sub-ranges


Basic aggregates

# Basic Aggregates

# Chunks



Statistic type




Mixed




Statistic function

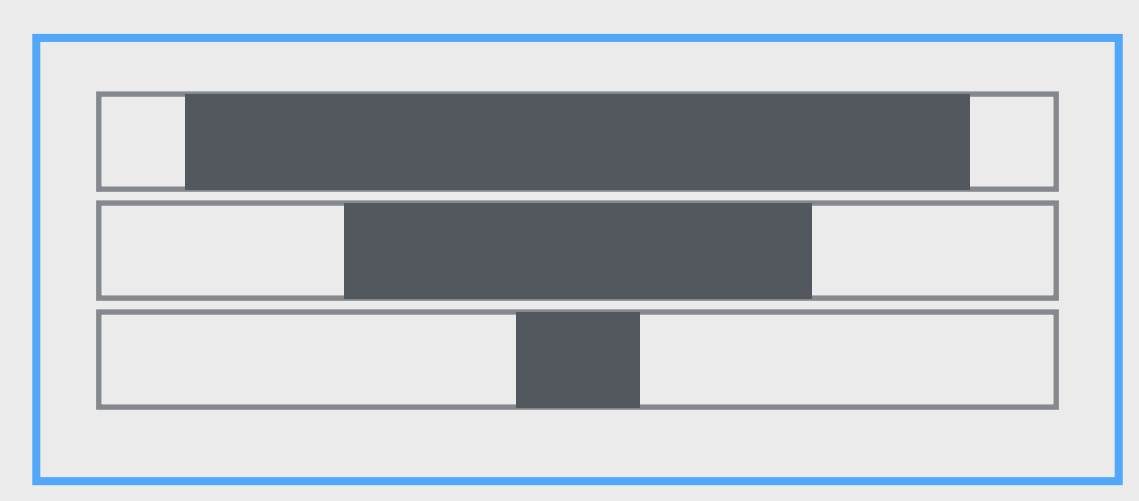
Aggregate function




Basic aggregates




Overlapping ranges




Sub-ranges



 Decompose



Statistic





Basic Aggregates

 Synthesize

 Store


 Chunk size




a chunk




Data



 Decompose



Statistic





Basic Aggregates


 Synthesize

 Store

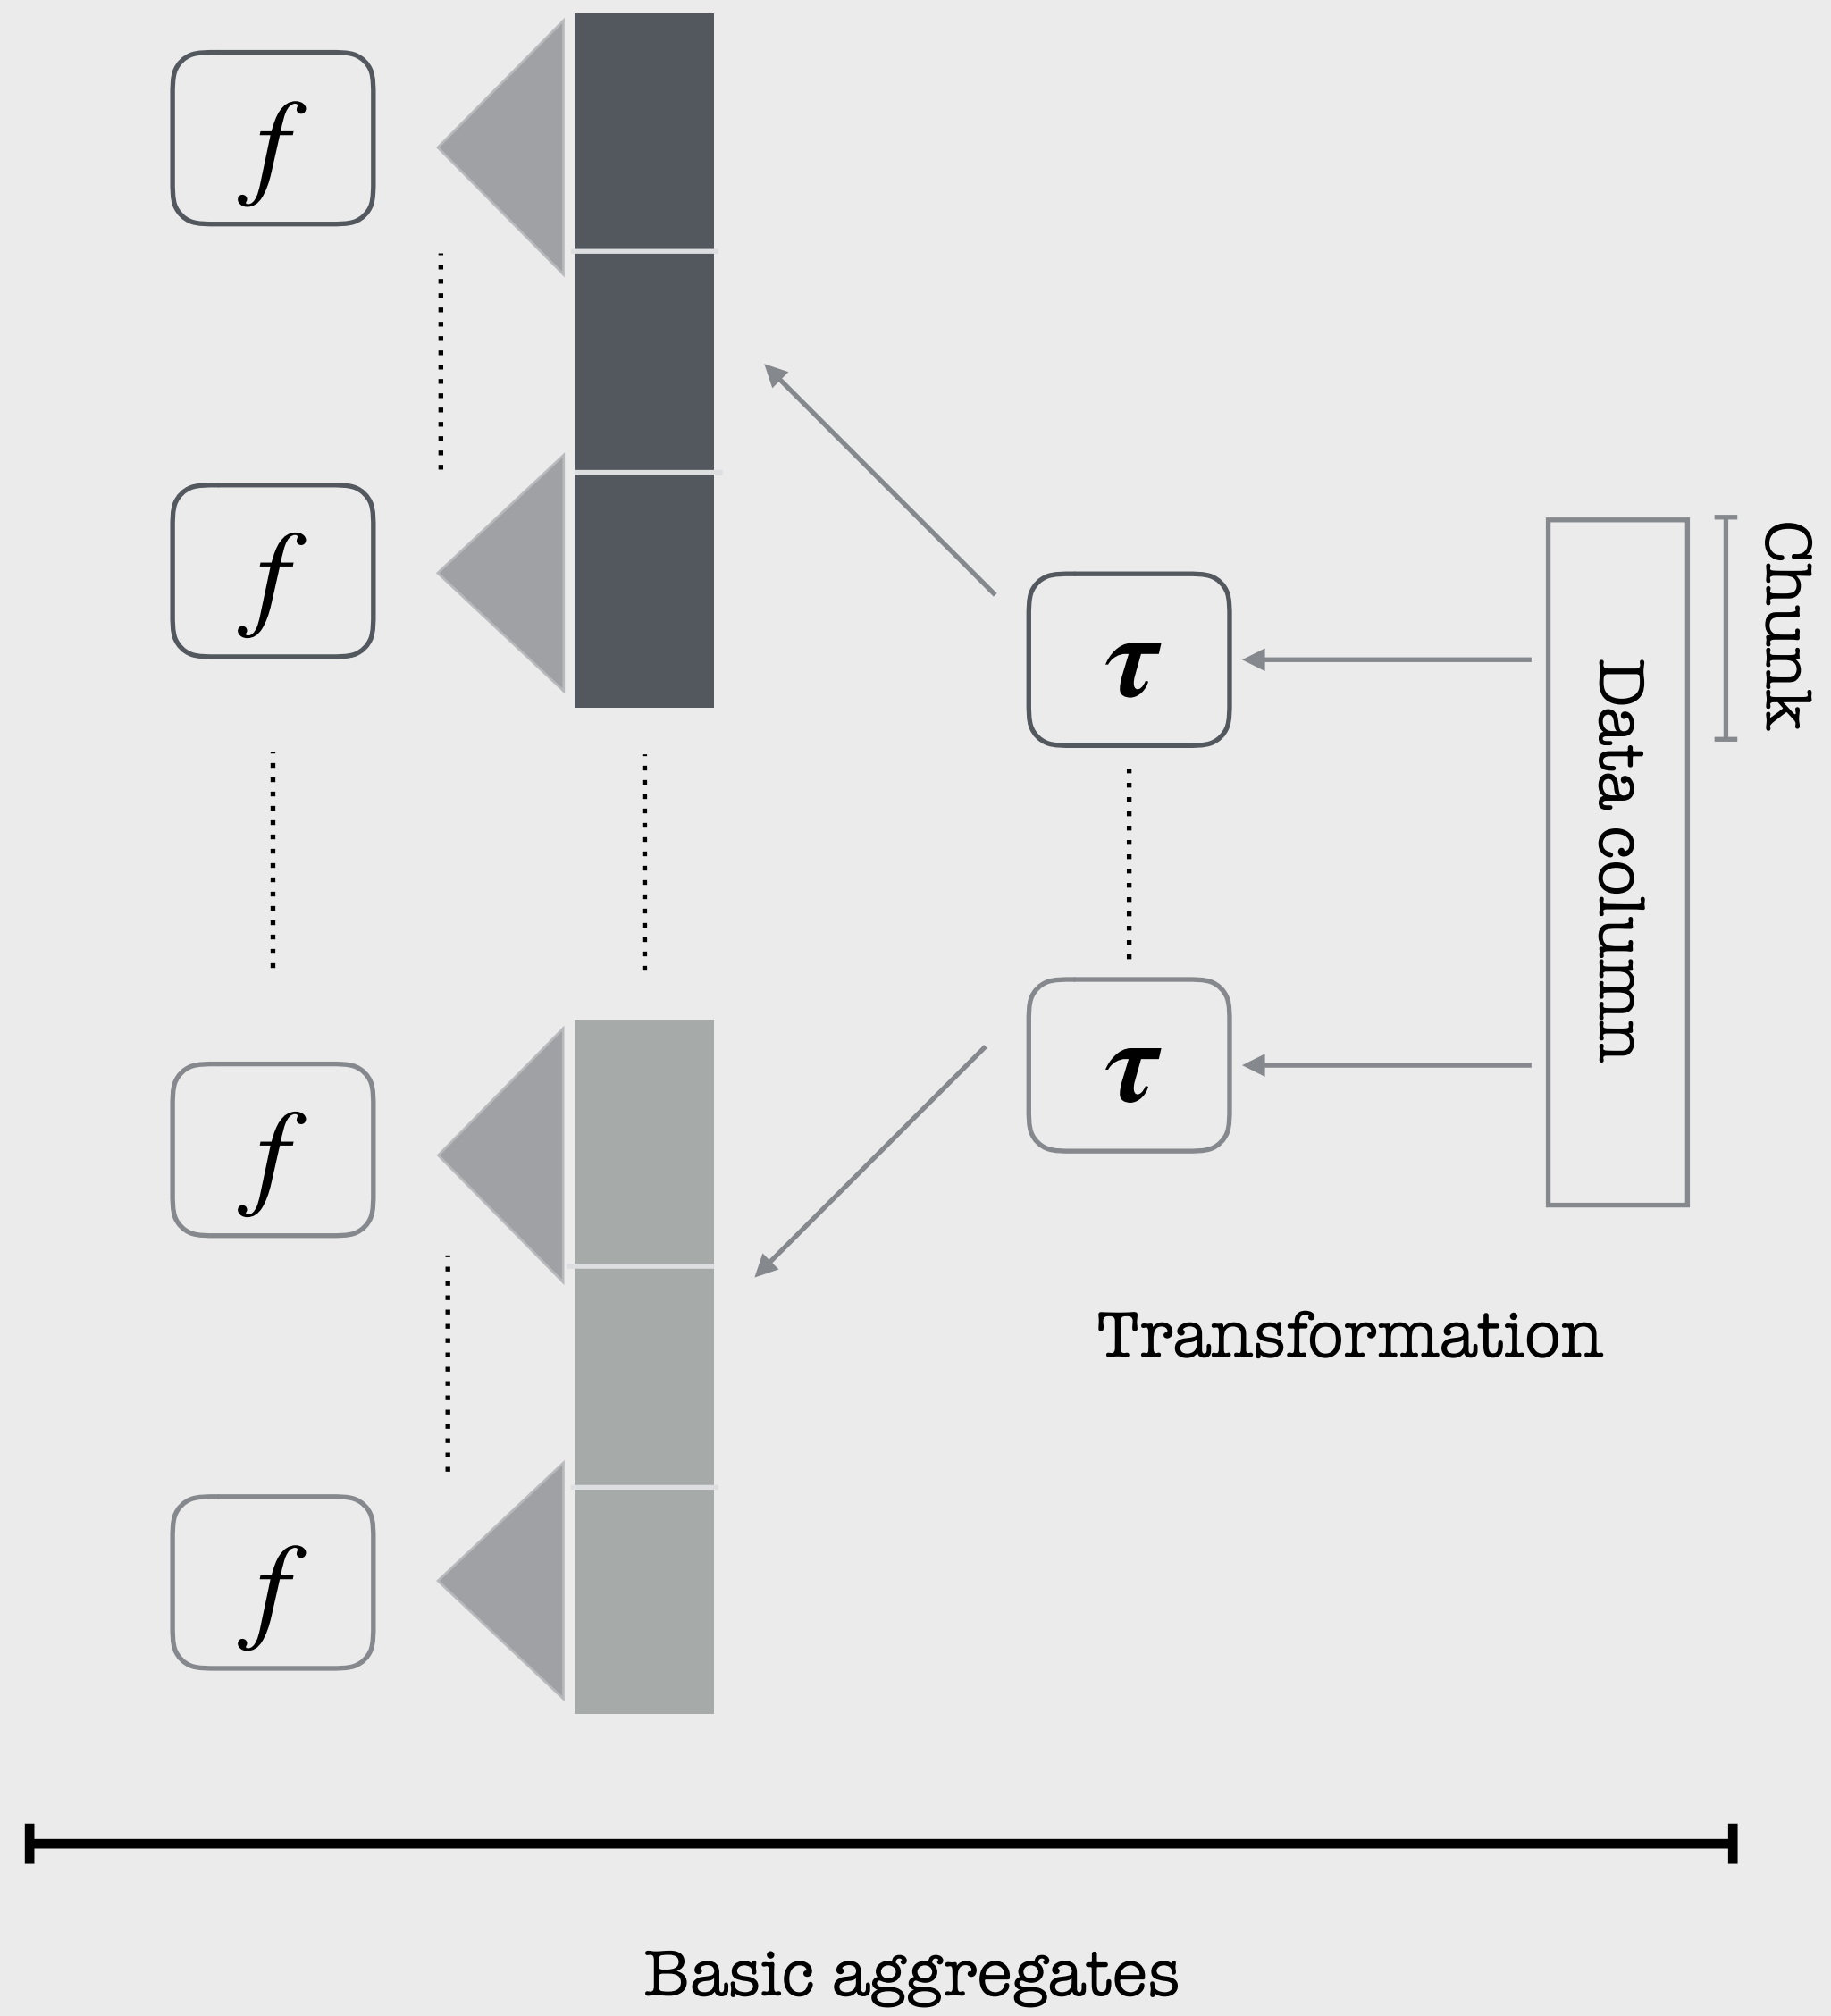

 Chunk size




a chunk

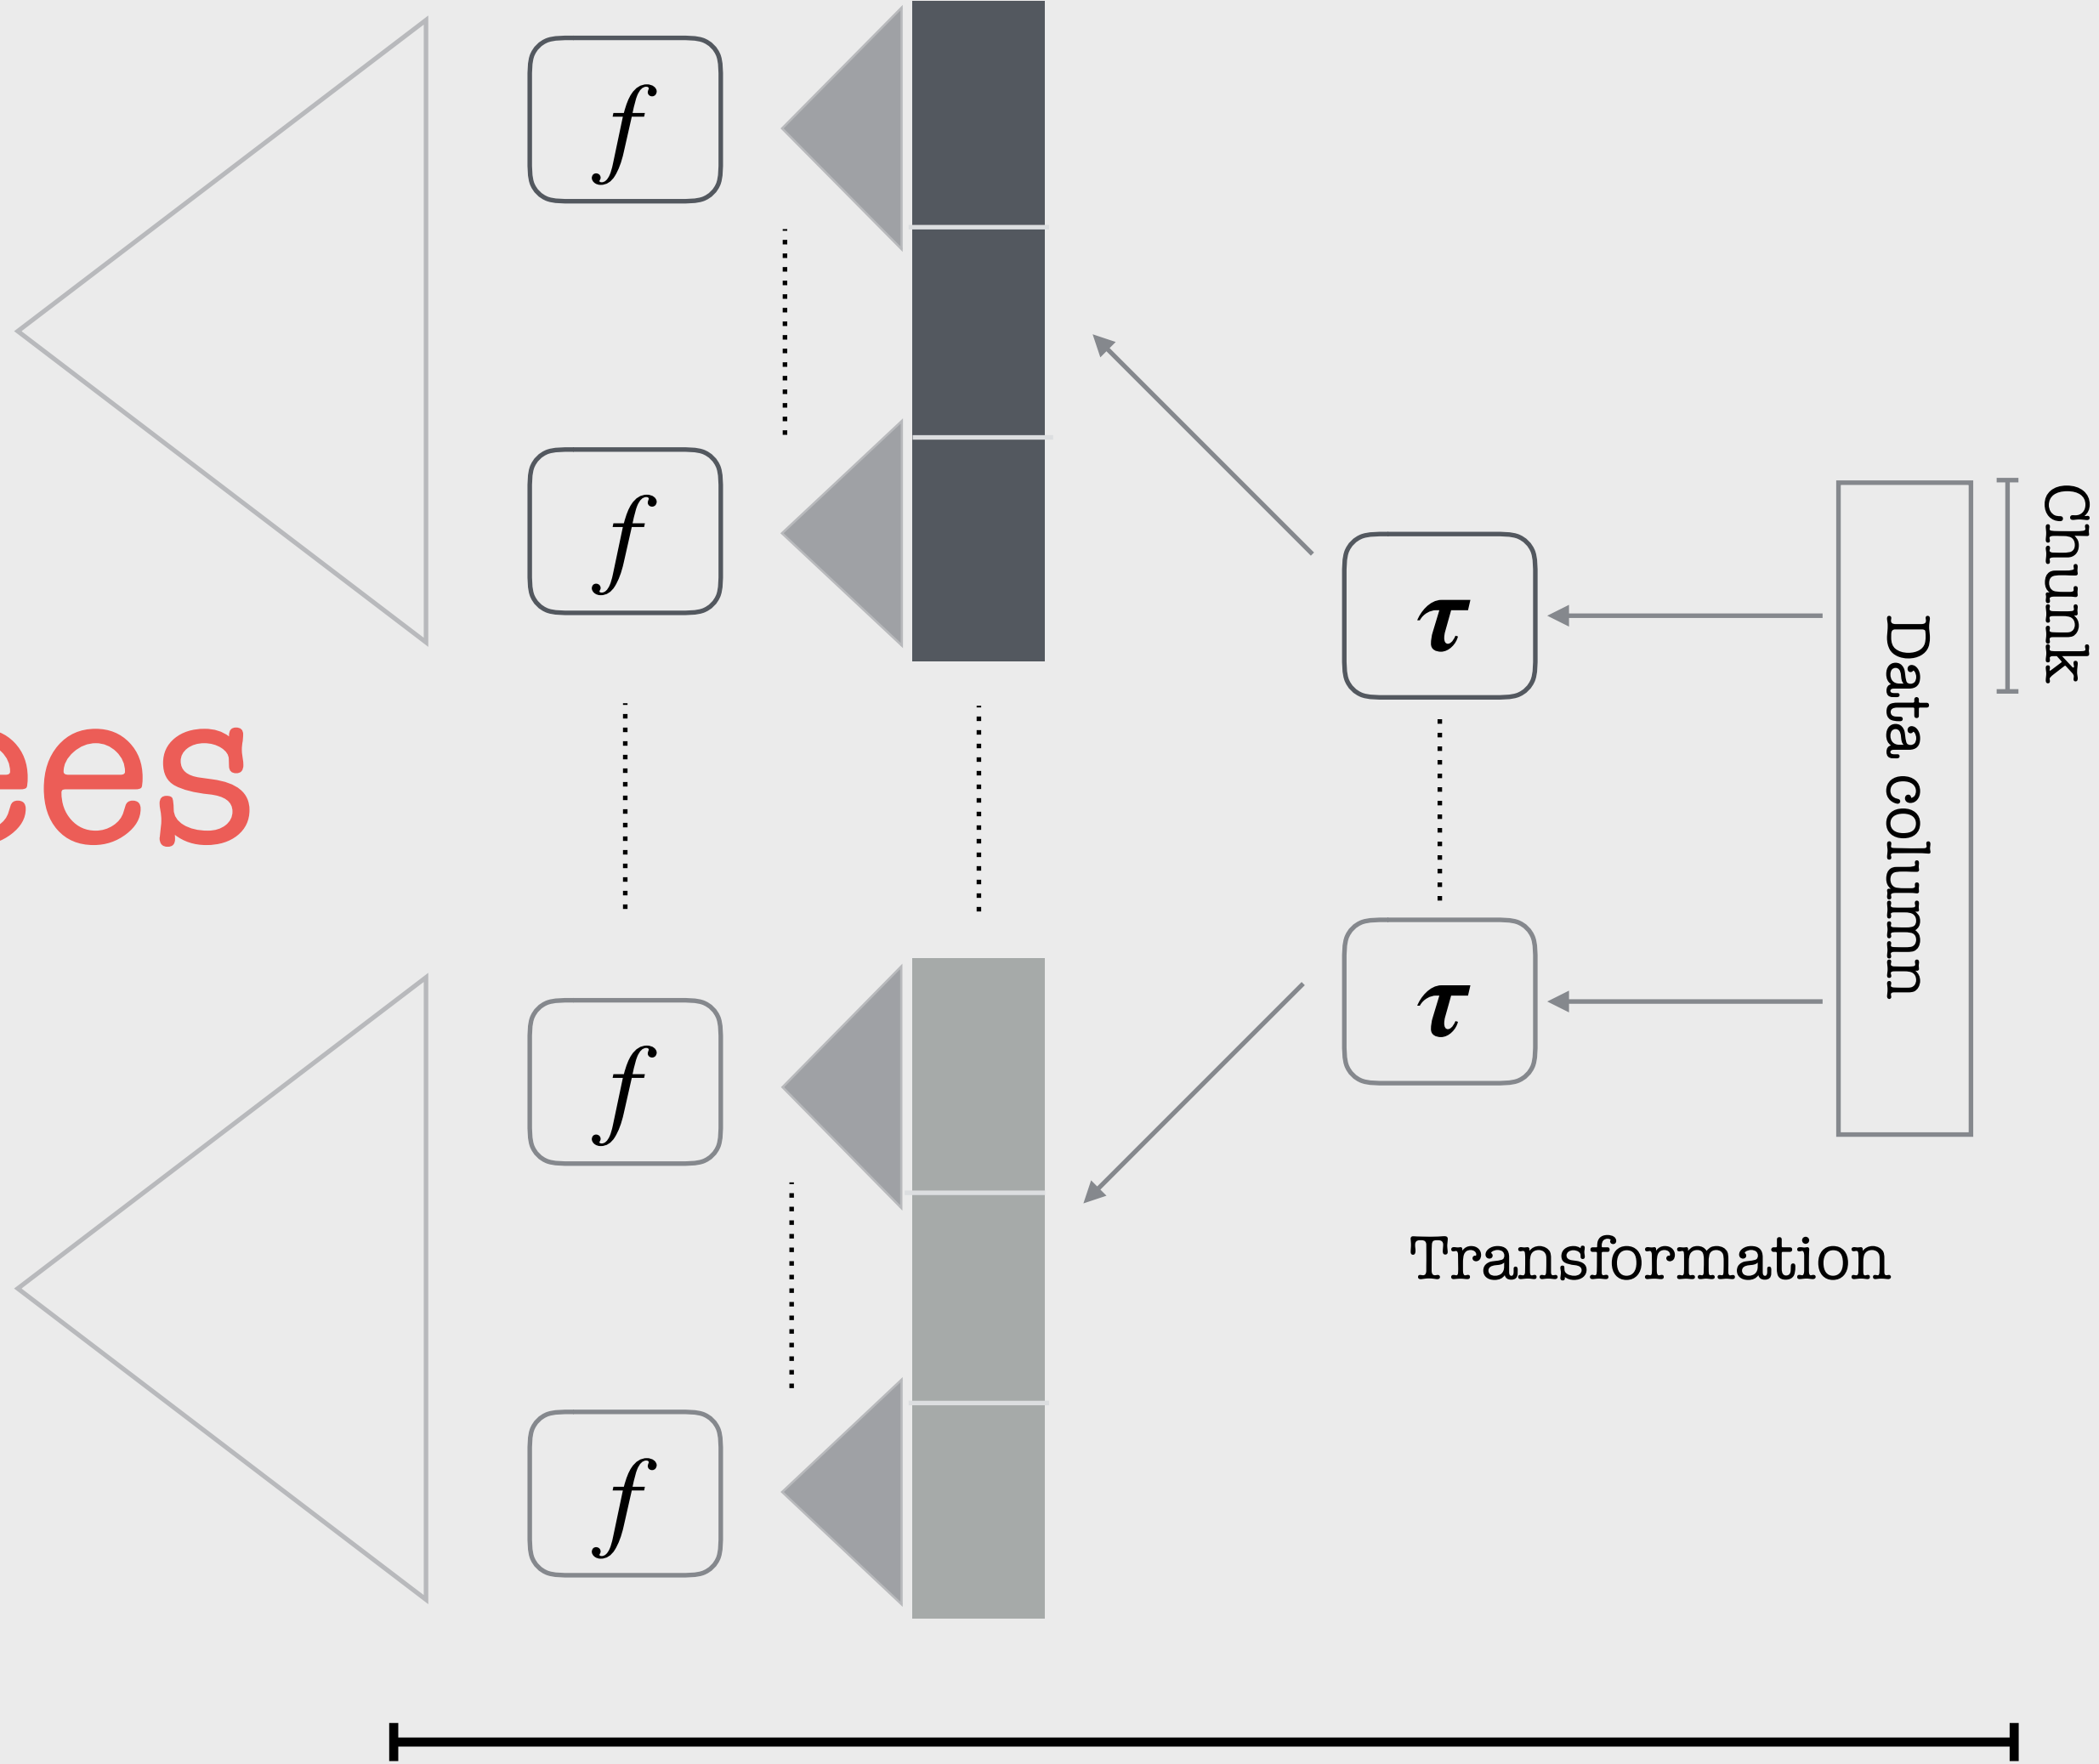



Data





# Segment trees






5 4 2 5 4 4 6 8 0 1 7 2 2 1 3 5

# Data

Logical  
chunk





## Basic aggregates

Leaves



# Data

Logical  
chunk





Leaves




# Data

Logical  
chunk





Leaves




Data




Logical chunk



Leaves




# Data



Leaves


5	4	2	5	4	4	6	8	0	1	7	2	2	1	3	5
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

# Data



7 22 10 3

42



One segment tree per **basic aggregate** per **column**



Column a

Column b

# One segment tree per basic aggregate per column

$$\sum X$$

$$\sum X^2$$


$$\sum X$$


$$\sum X^2$$

Column a


Column b

# One segment tree per basic aggregate per column






 Decompose



Statistic





Basic Aggregates


 Synthesize

 Store


 Chunk size




a chunk




Data



 Decompose



Statistic





Basic Aggregates




 Synthesize

 Store


 Chunk size




a chunk



Data



Calculate variance in temperature between May 15 and Oct. 15




Calculate `variance` in `temperature` between `May 15 and Oct. 15`

`variance`

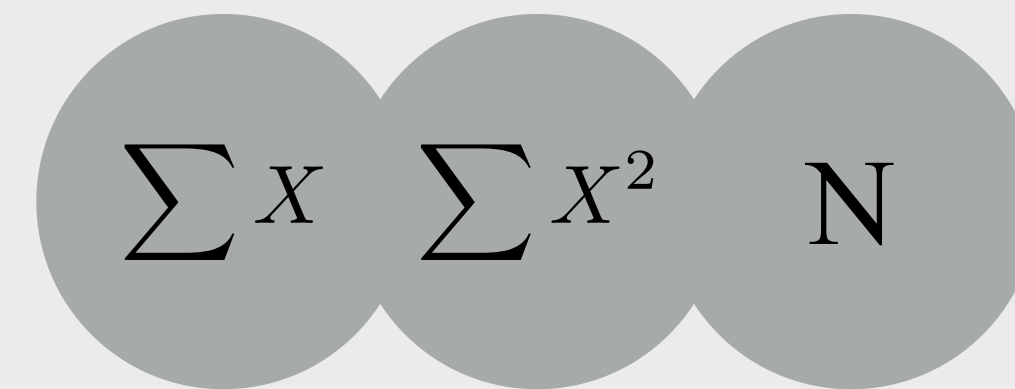
`temperature`

`May 15 and Oct. 15`



Calculate variance in temperature between May 15 and Oct. 15


variance



Recipe


temperature

May 15 and Oct. 15



Calculate variance in temperature between May 15 and Oct. 15


variance




Recipe




Chunk size = 1 month



Calculate variance in temperature between May 15 and Oct. 15



Plan



Recipe


temperature



base data

segment trees


base data



✓ Decompose



Statistic





Basic Aggregates

✓ Synthesize

✓ Store


? Chunk size



a chunk




Data



✓ Decompose



Statistic





Basic Aggregates


✓ Synthesize

✓ Store


? Chunk size





a chunk




Data









## Segment tree traversal




$n$  | column size;  $c$  | chunk size





## Segment tree traversal



## Residual range scan




$n$  | column size;  $c$  | chunk size

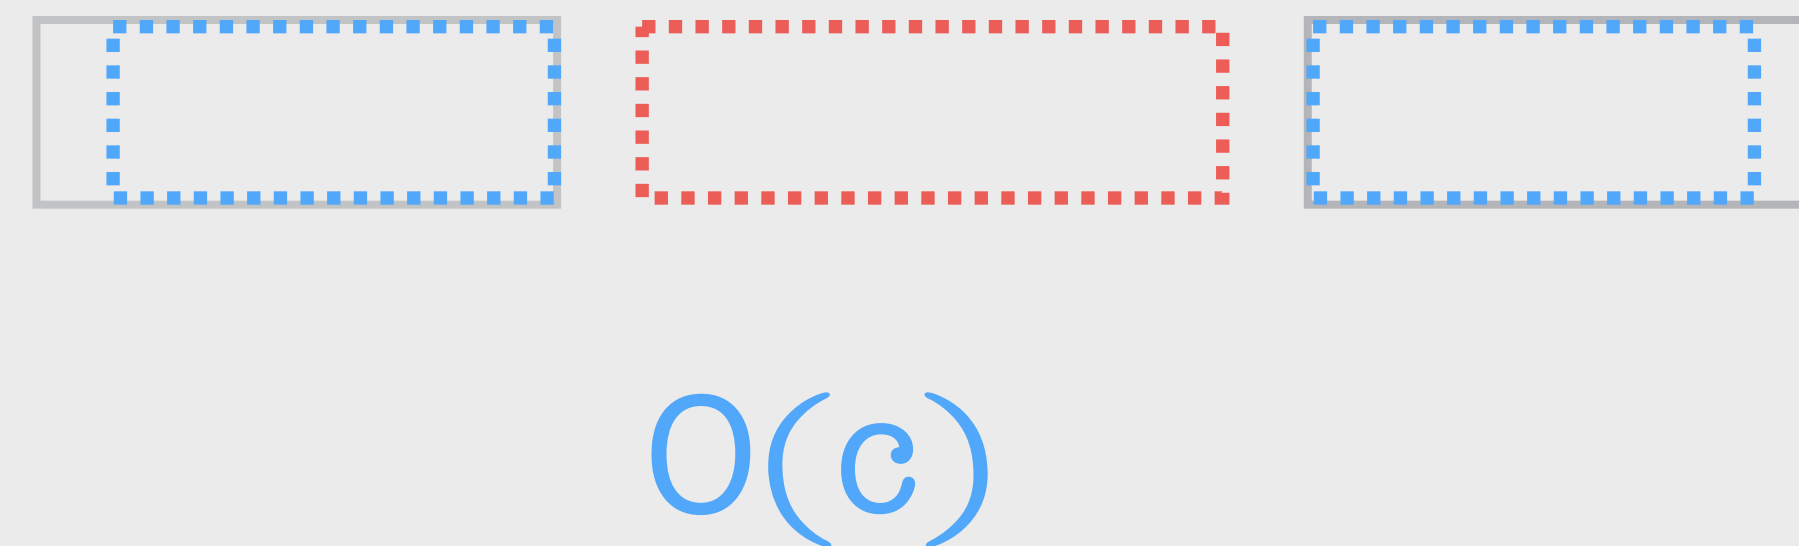


$$\text{Total cost} = O(\log(n/c) + c)$$


Segment tree traversal




Residual range scan




$n$  | column size;  $c$  | chunk size




$$\text{Total cost} = O(\log(n/c) + c)$$




$n$  | column size;  $c$  | chunk size




$$\text{Total cost} = O(\log(n/c) + c)$$




$n$  | column size;  $c$  | chunk size



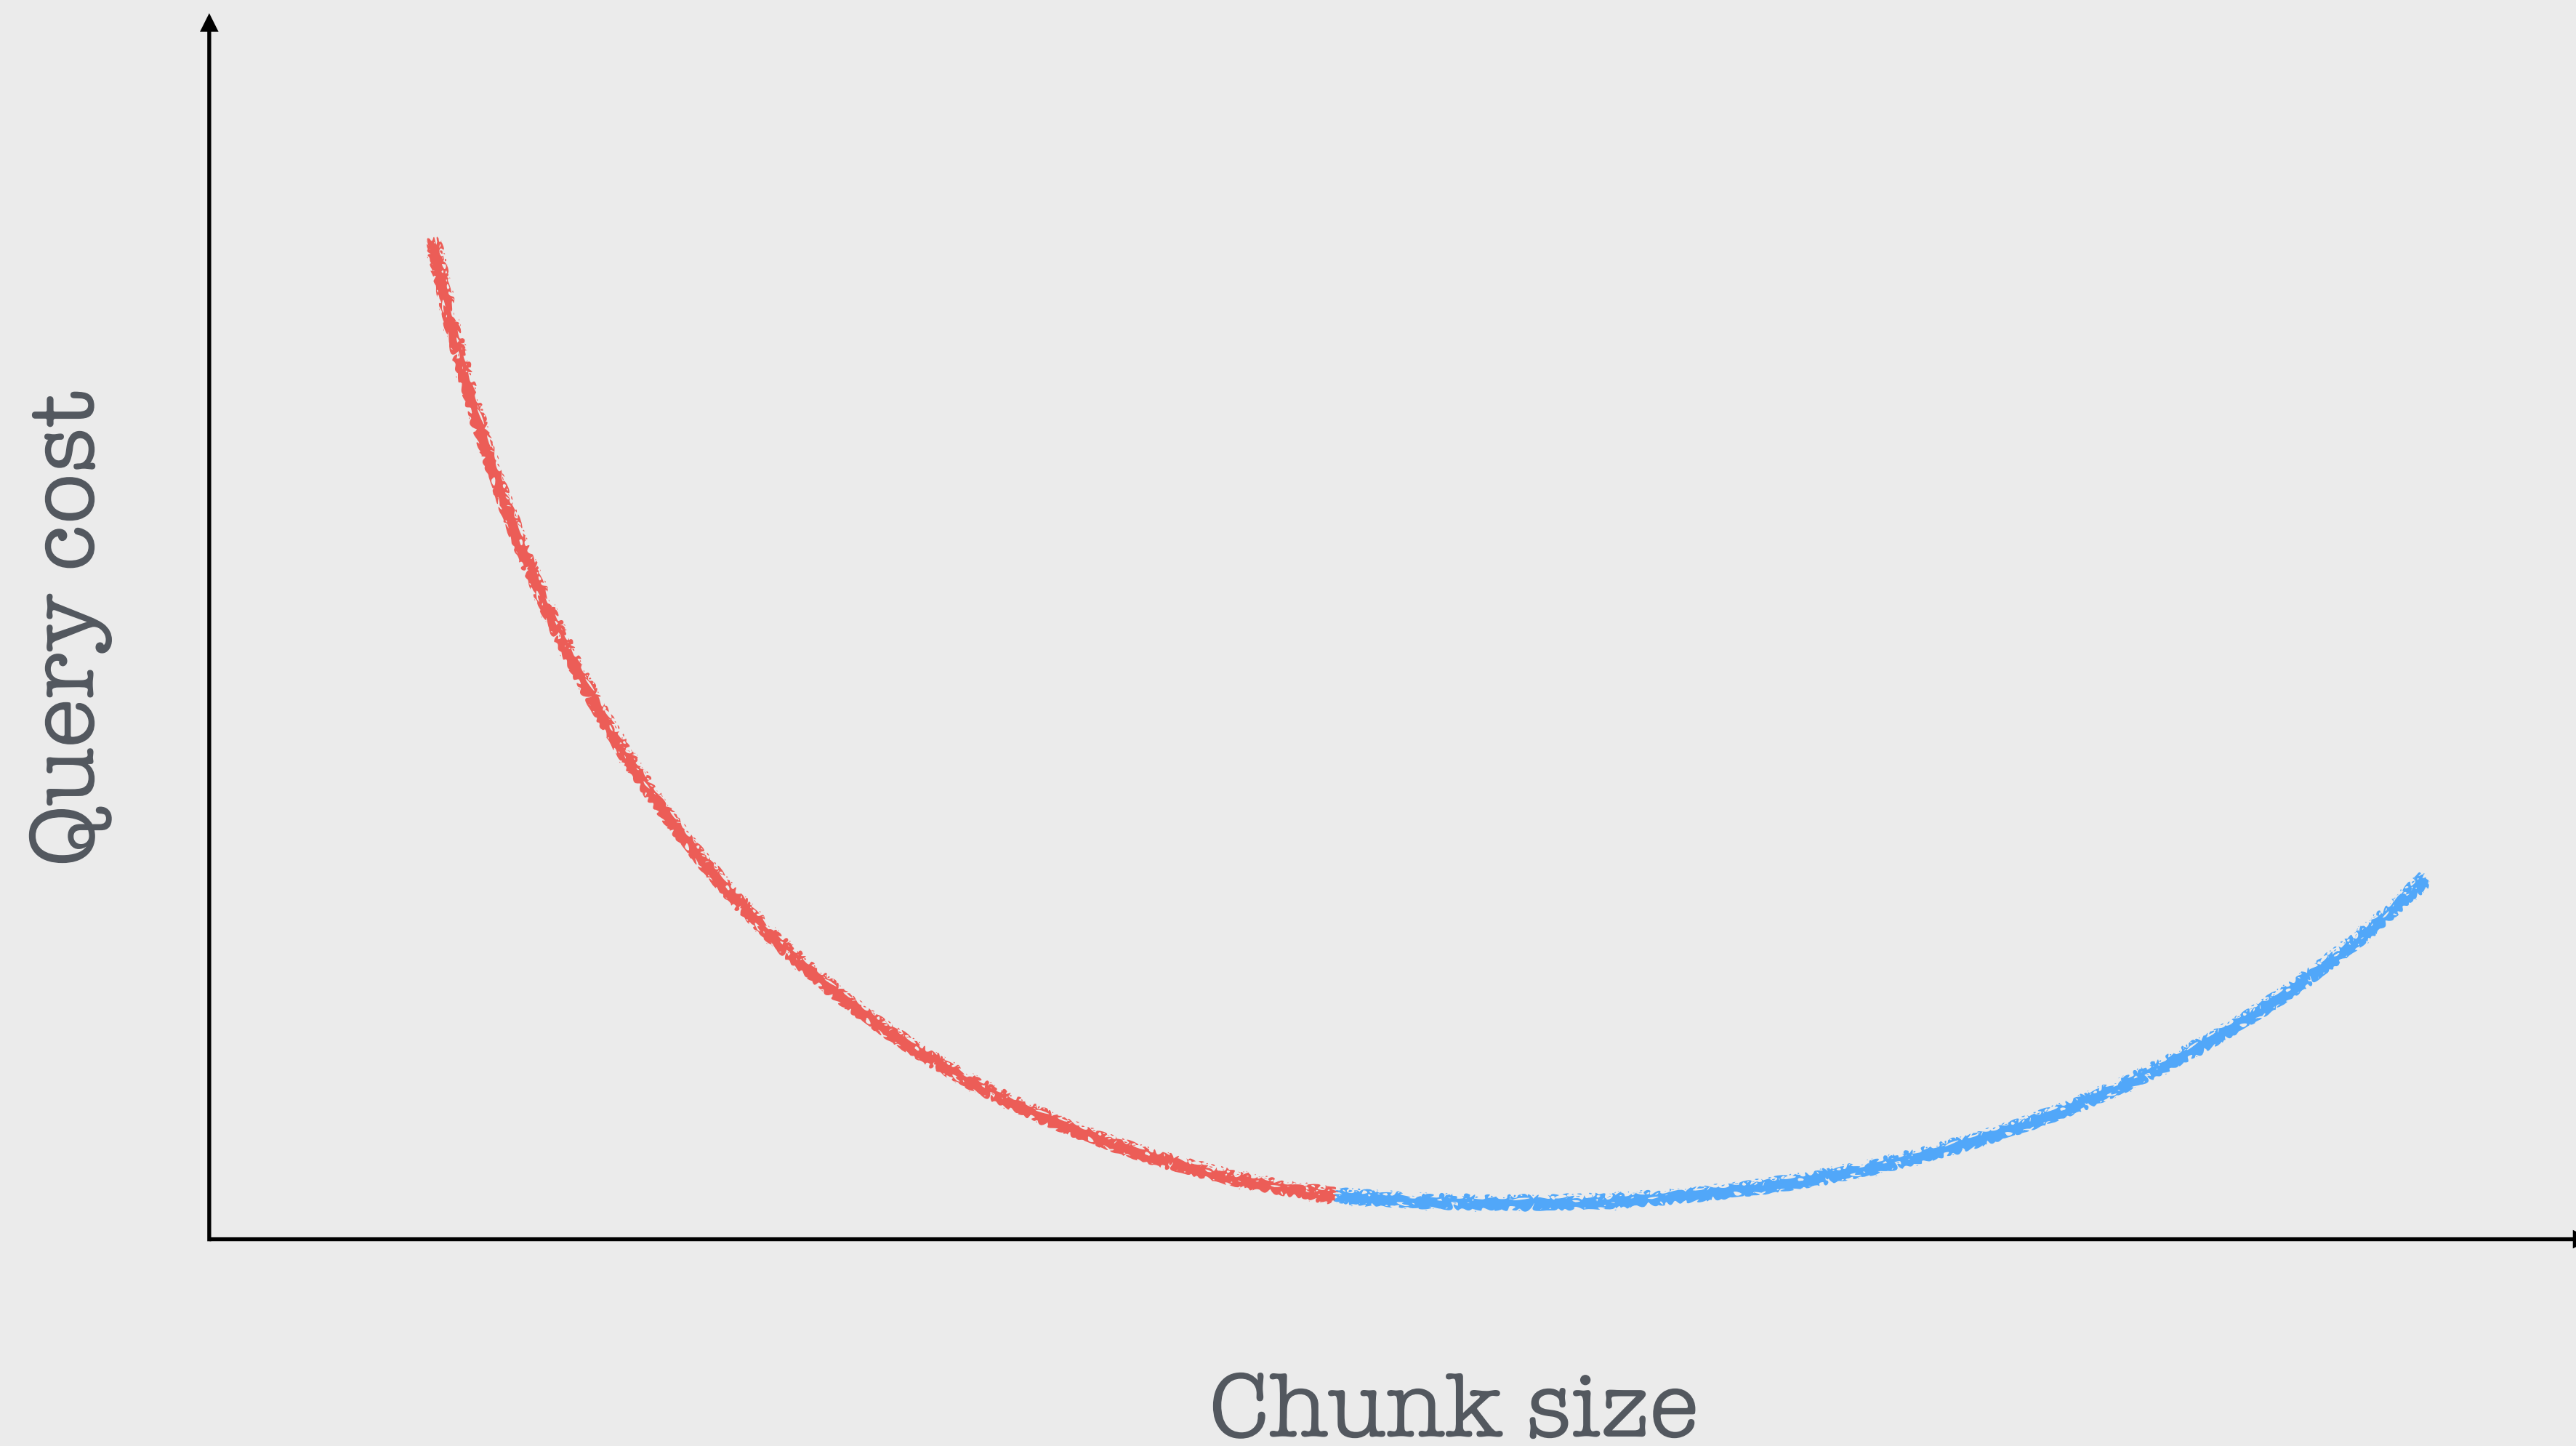
$$\text{Total cost} = O(\log(n/c) + c)$$




$n$  | column size;  $c$  | chunk size




$$\text{Total cost} = O(\log(n/c) + c)$$




$n$  | column size;  $c$  | chunk size




$$\text{Total cost} = O(\log(n/c) + c)$$




$n$  | column size;  $c$  | chunk size




$$\text{Total cost} = O(\log(n/c) + c)$$



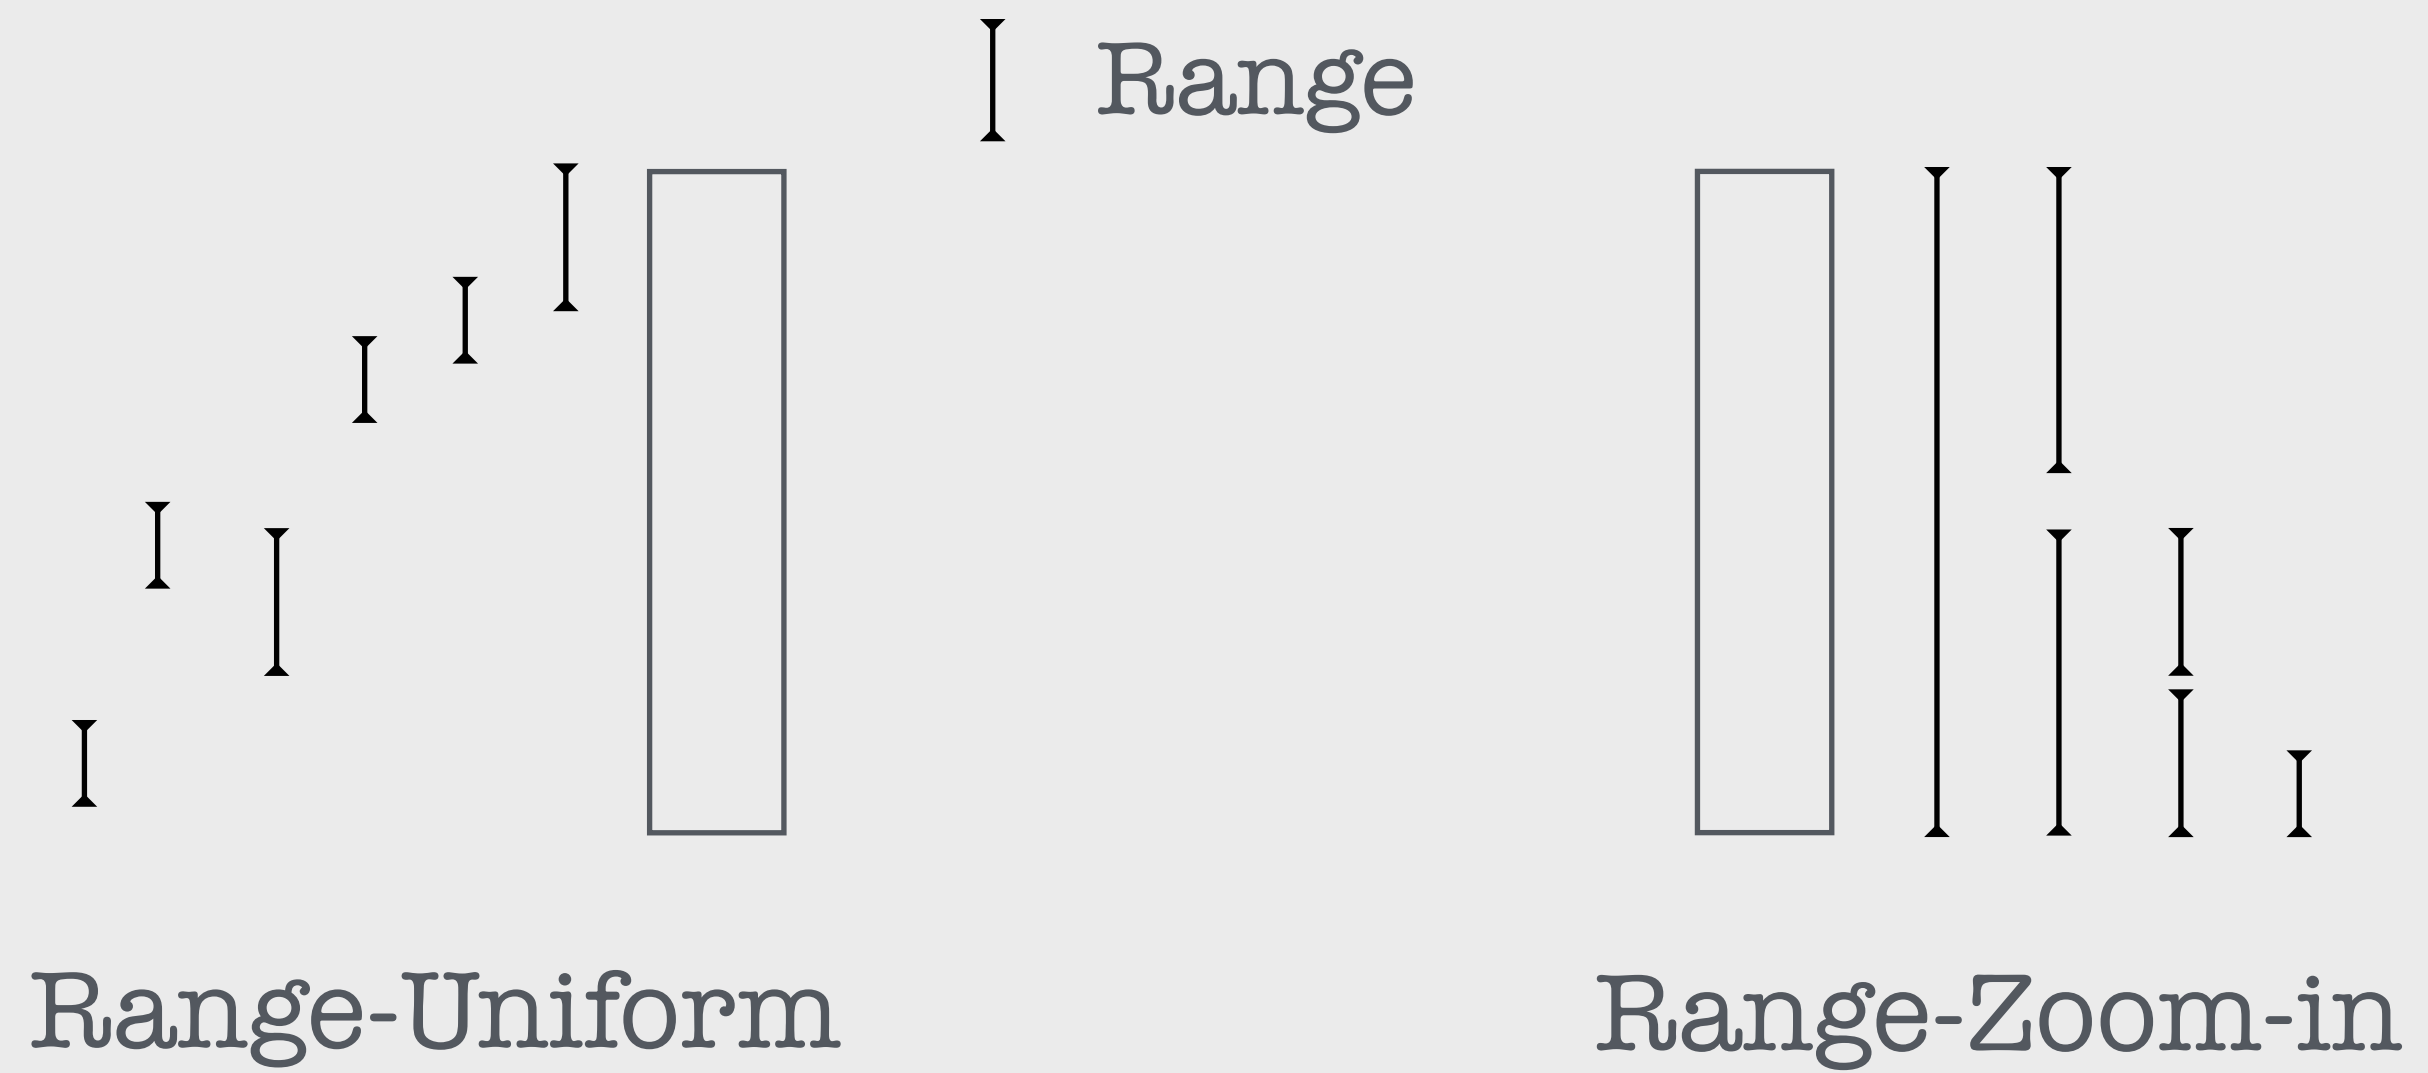
$n$  | column size;  $c$  | chunk size





# Experimental Analysis





Query range distribution






# Online Data Canopy




Online Data Canopy




Offline Data Canopy




# Range Uniform Workload

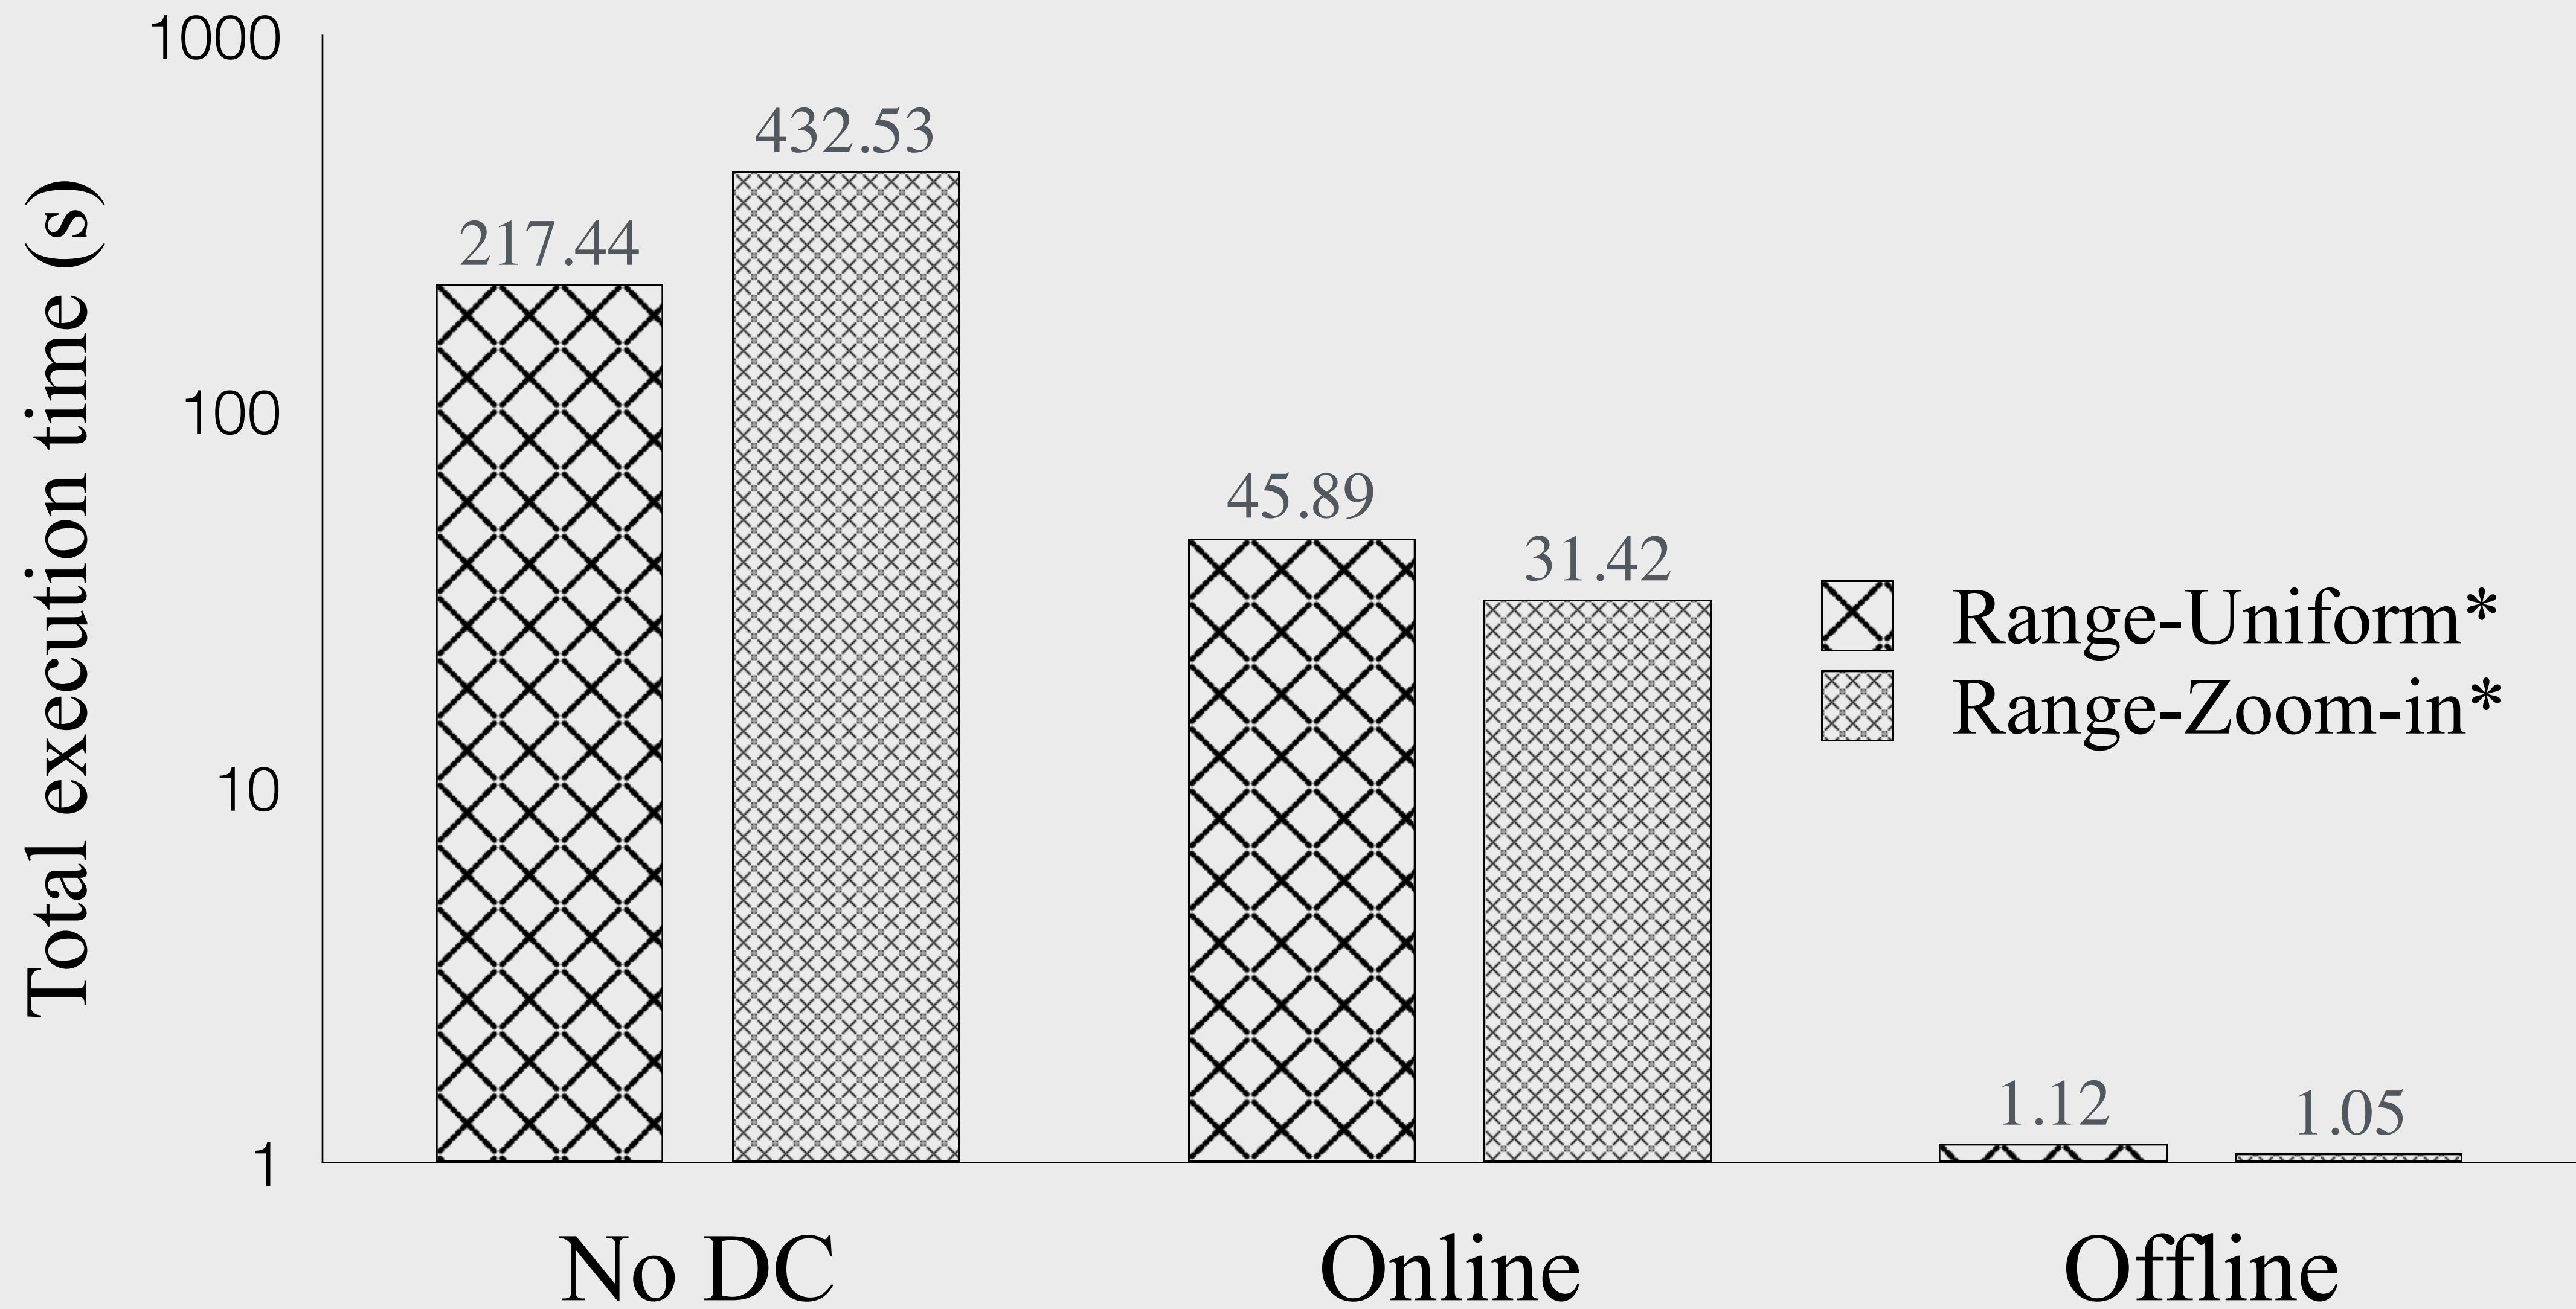







# Range Uniform Workload




DC is up to 2 orders of magnitude faster

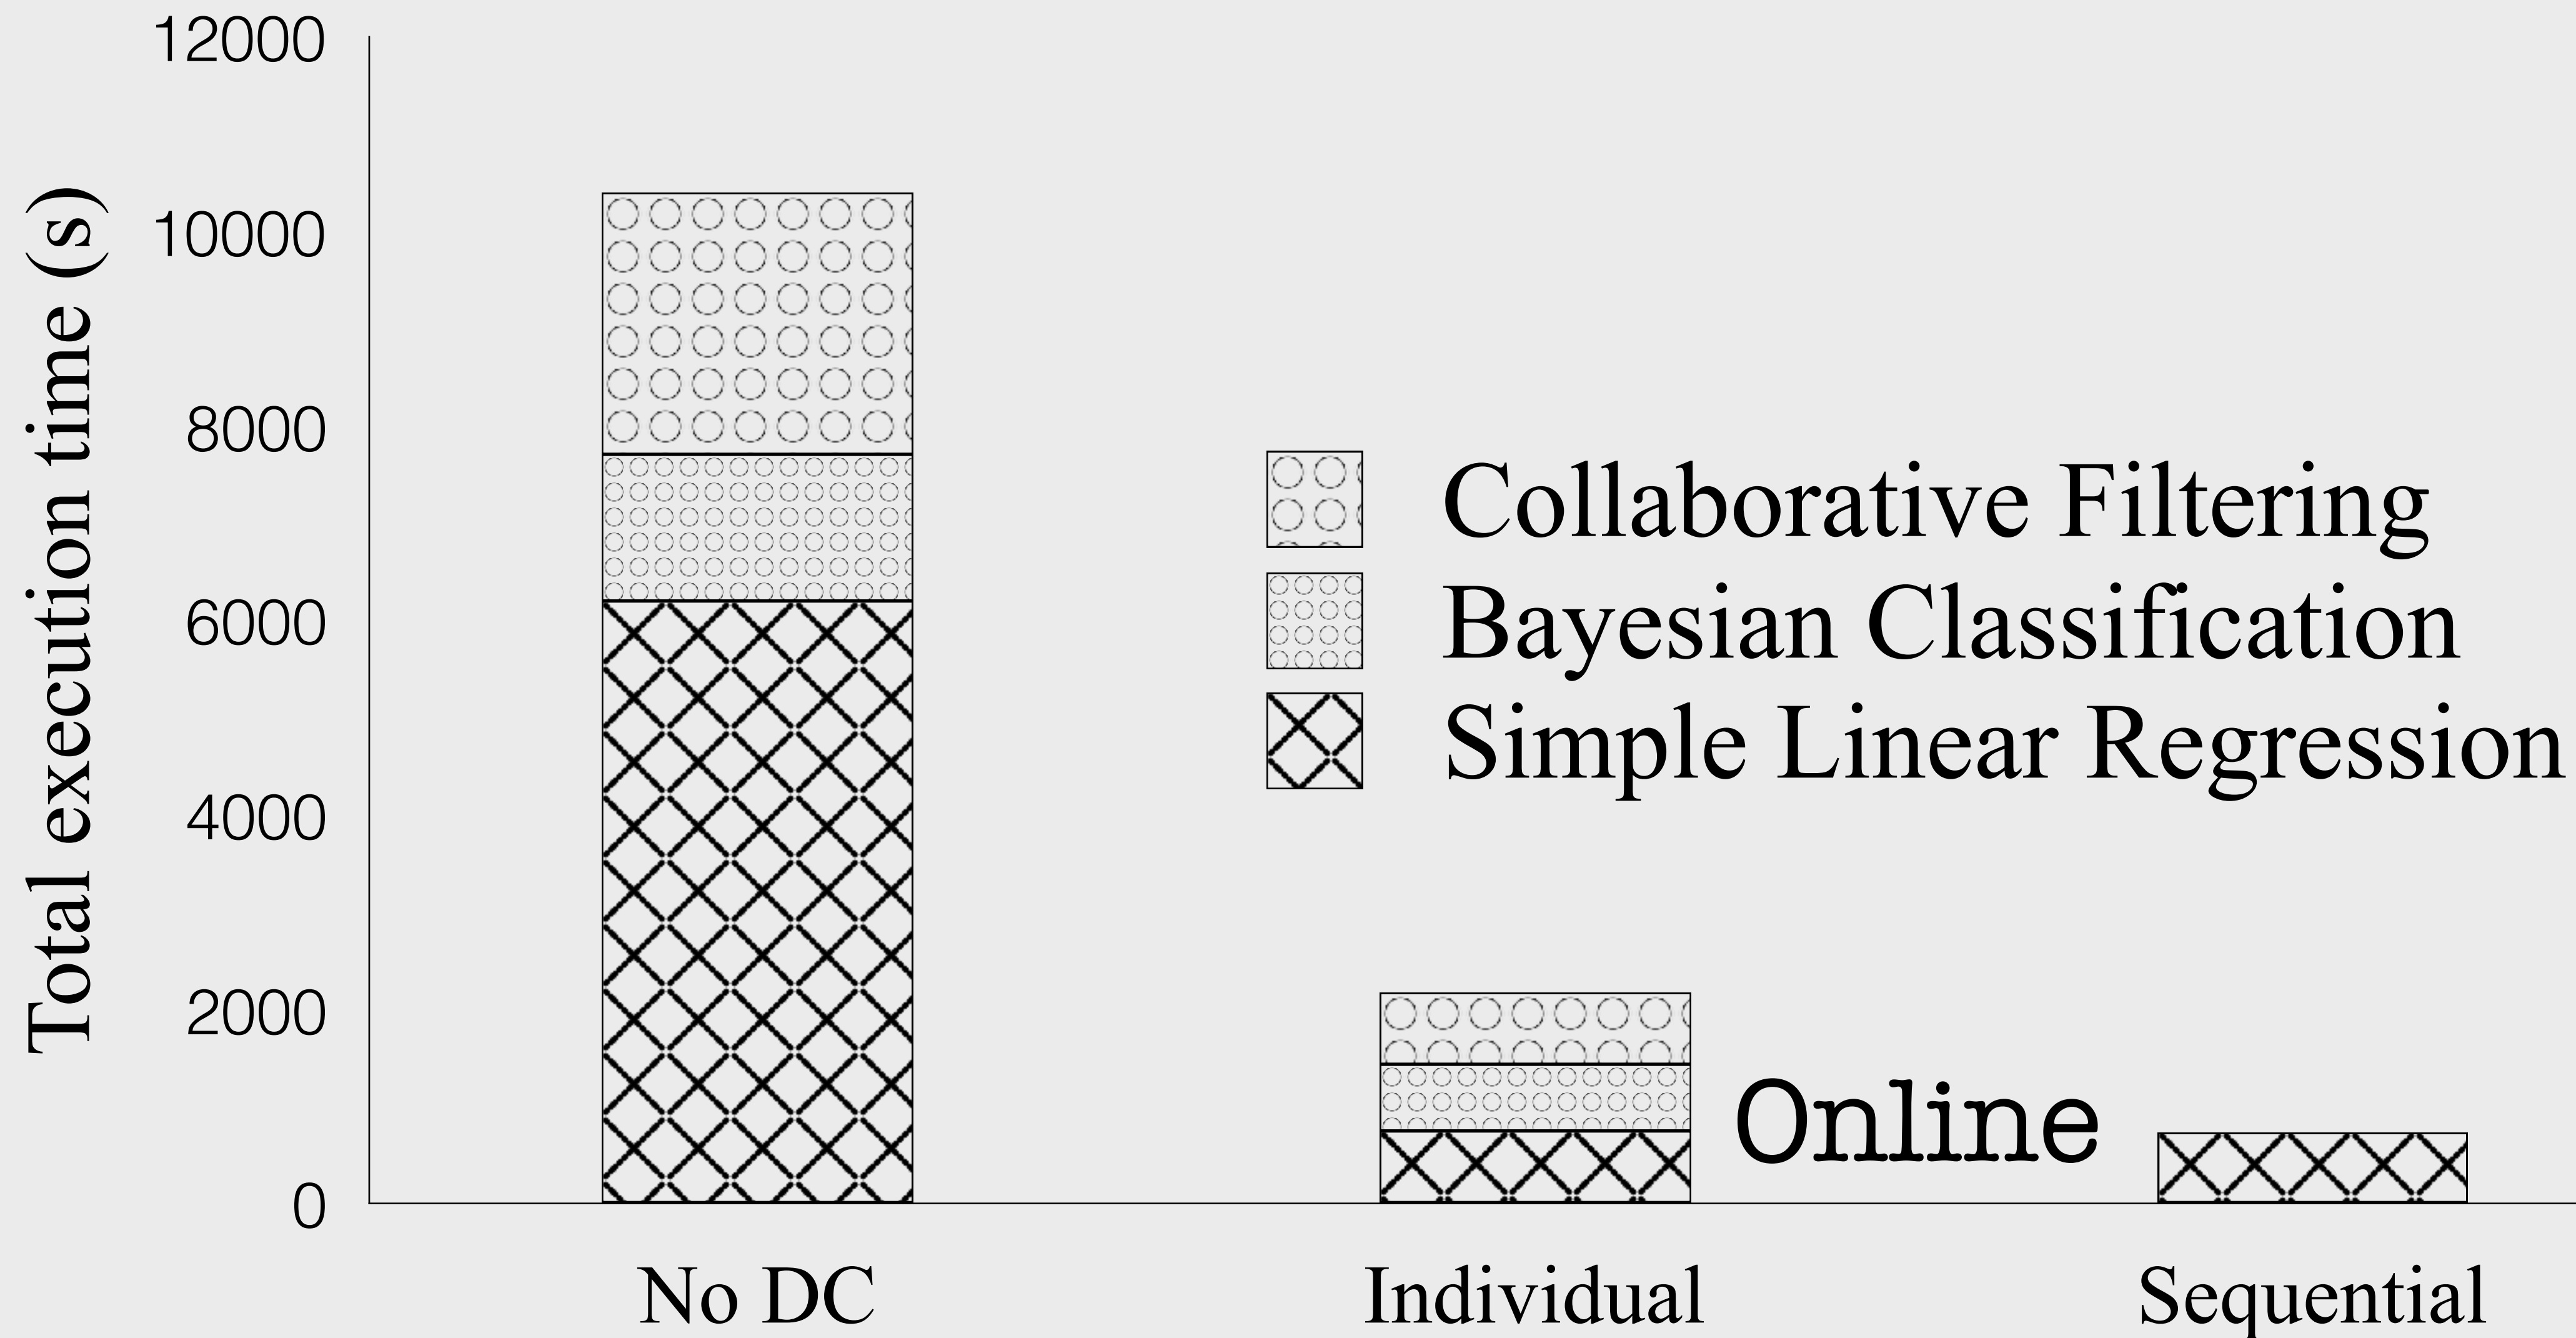




\* 2000 Queries





# DC benefits regardless of the exploration scenario








# Data Canopy: Accelerating Exploratory Statistical Analysis



**STATISTICS ARE EVERYWHERE!**

**REPETITIVE STATISTICS AND DATA ACCESS**




**DATA CANOPY SYNTHESIZES STATISTICS FROM BASIC INGREDIENTS**

# QUERIOSITY



Accelerate  
by synthesis

Provide  
hints




# DATA SYSTEMS. LABORATORY


@ Harvard School of Engineering and Applied Sciences

*Designing data systems for the big data era*

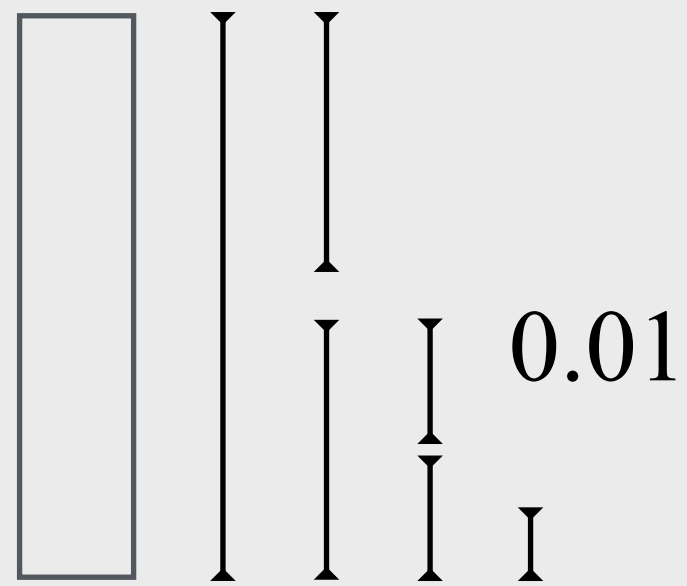
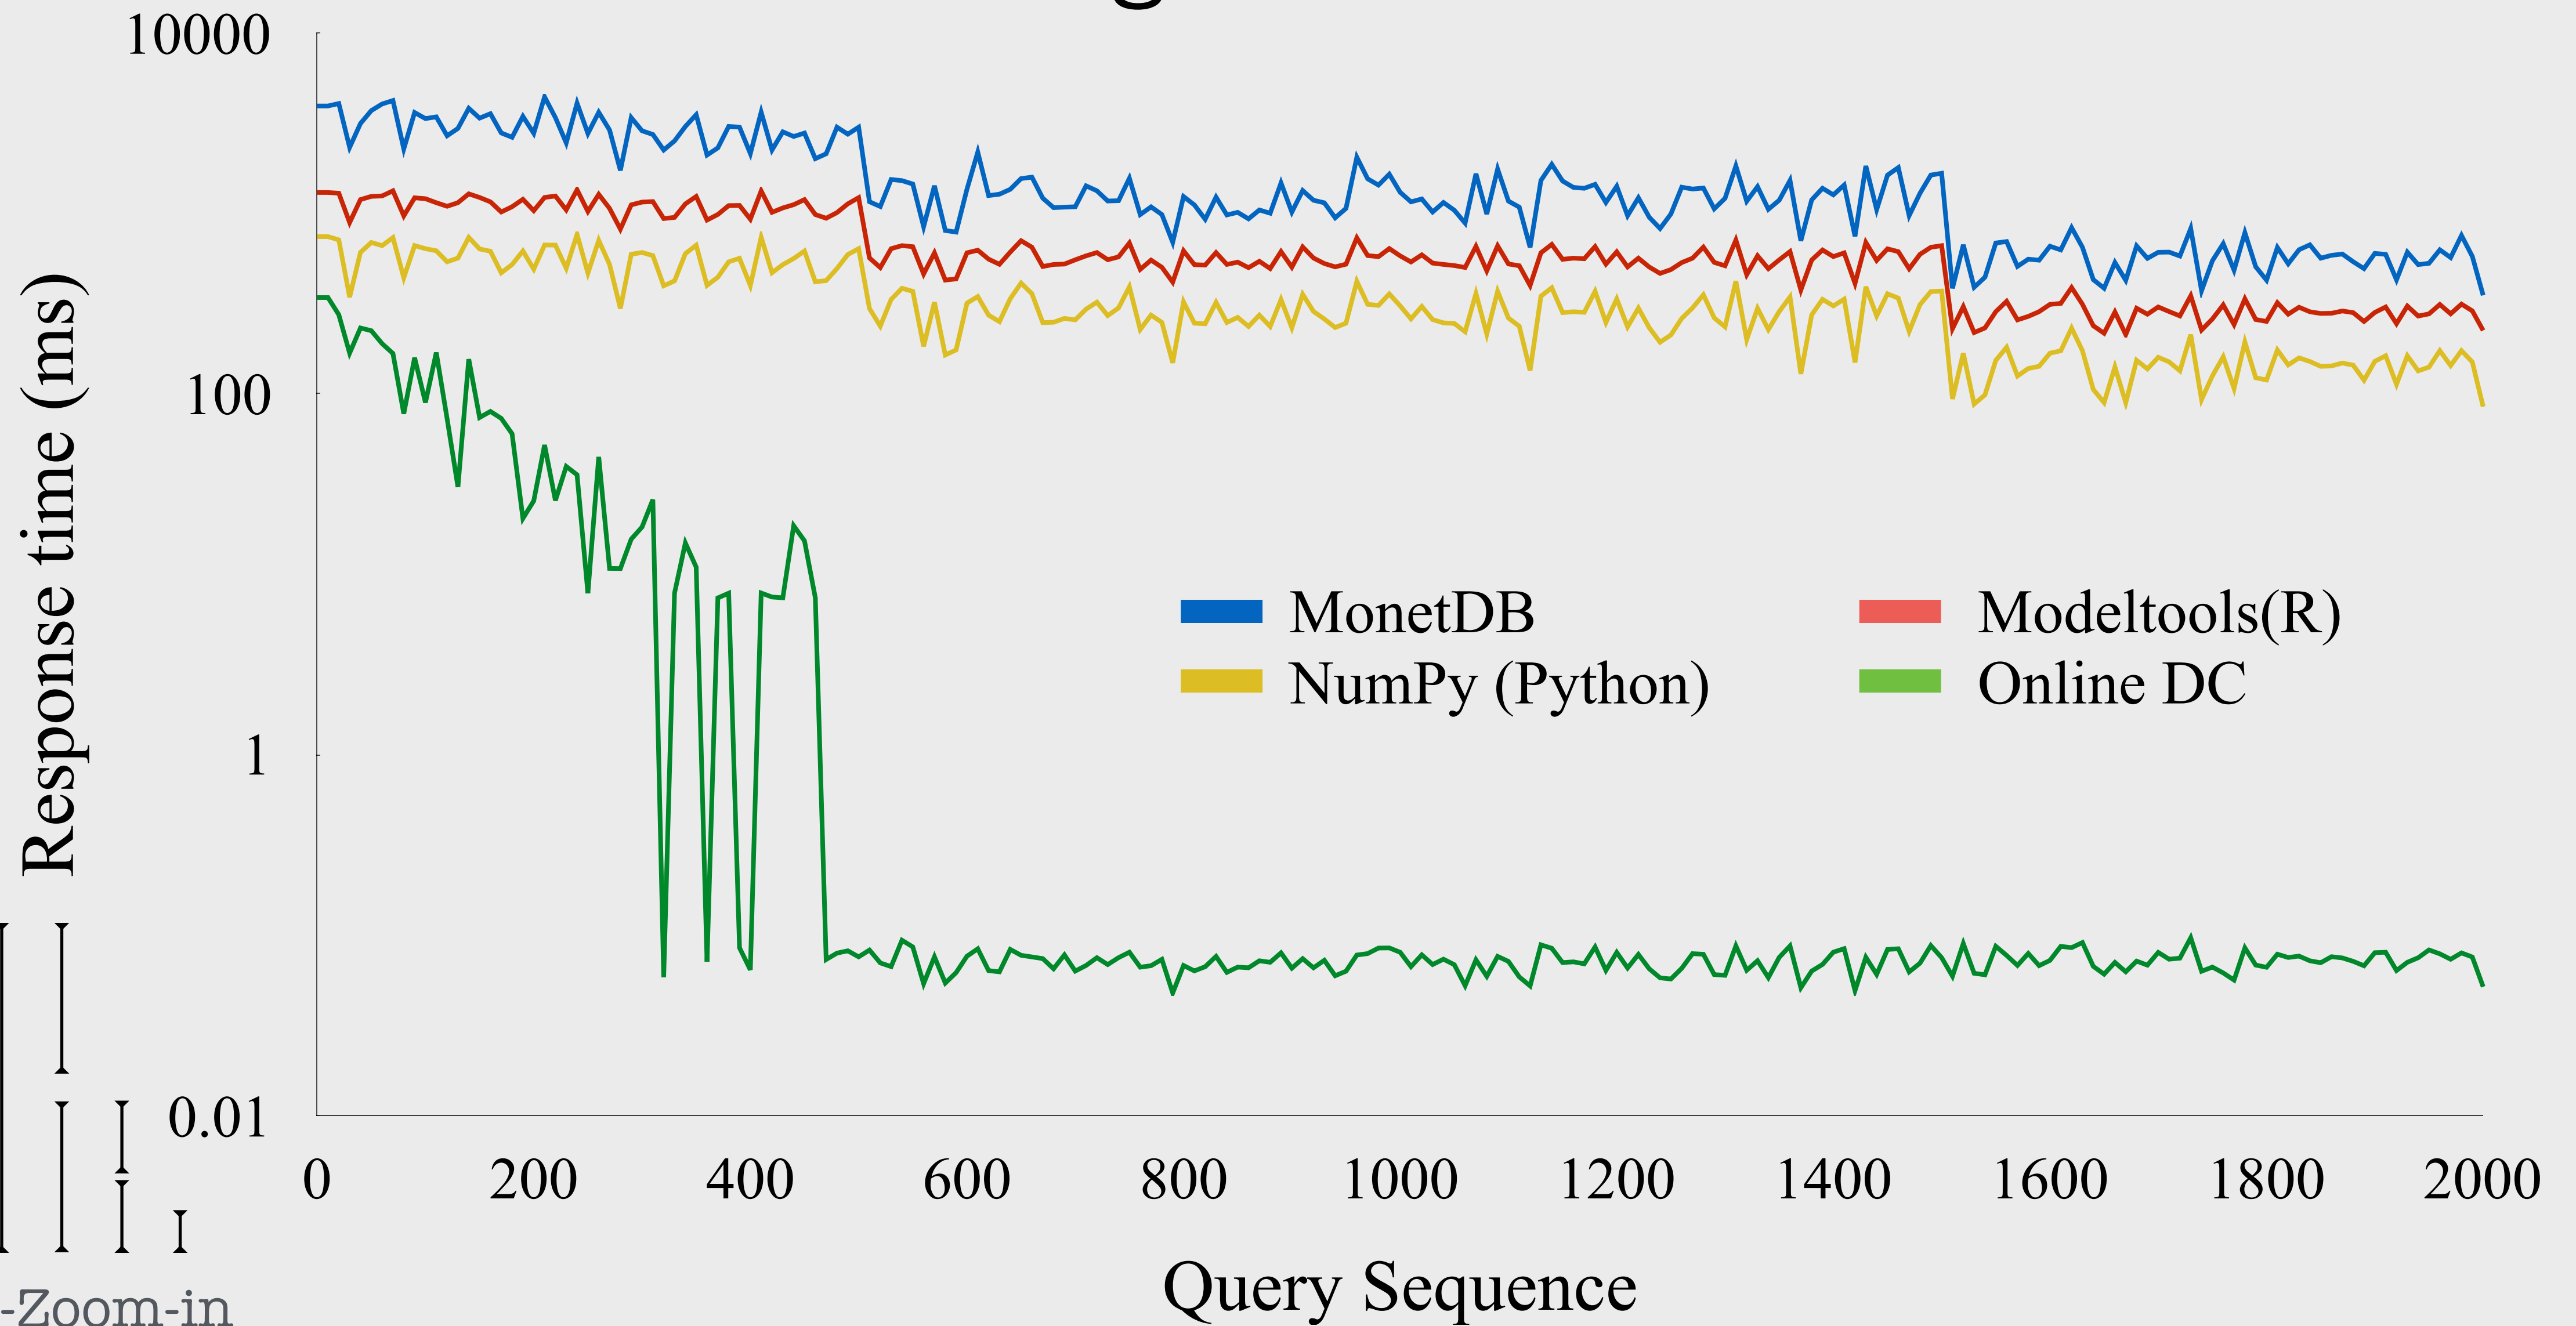
[daslab.seas.harvard.edu/](https://daslab.seas.harvard.edu/)  
[data-canopy](#)  
[queriosity](#)

Thank You !





# Range-Zoom-in




Range-Zoom-in

Data Canopy

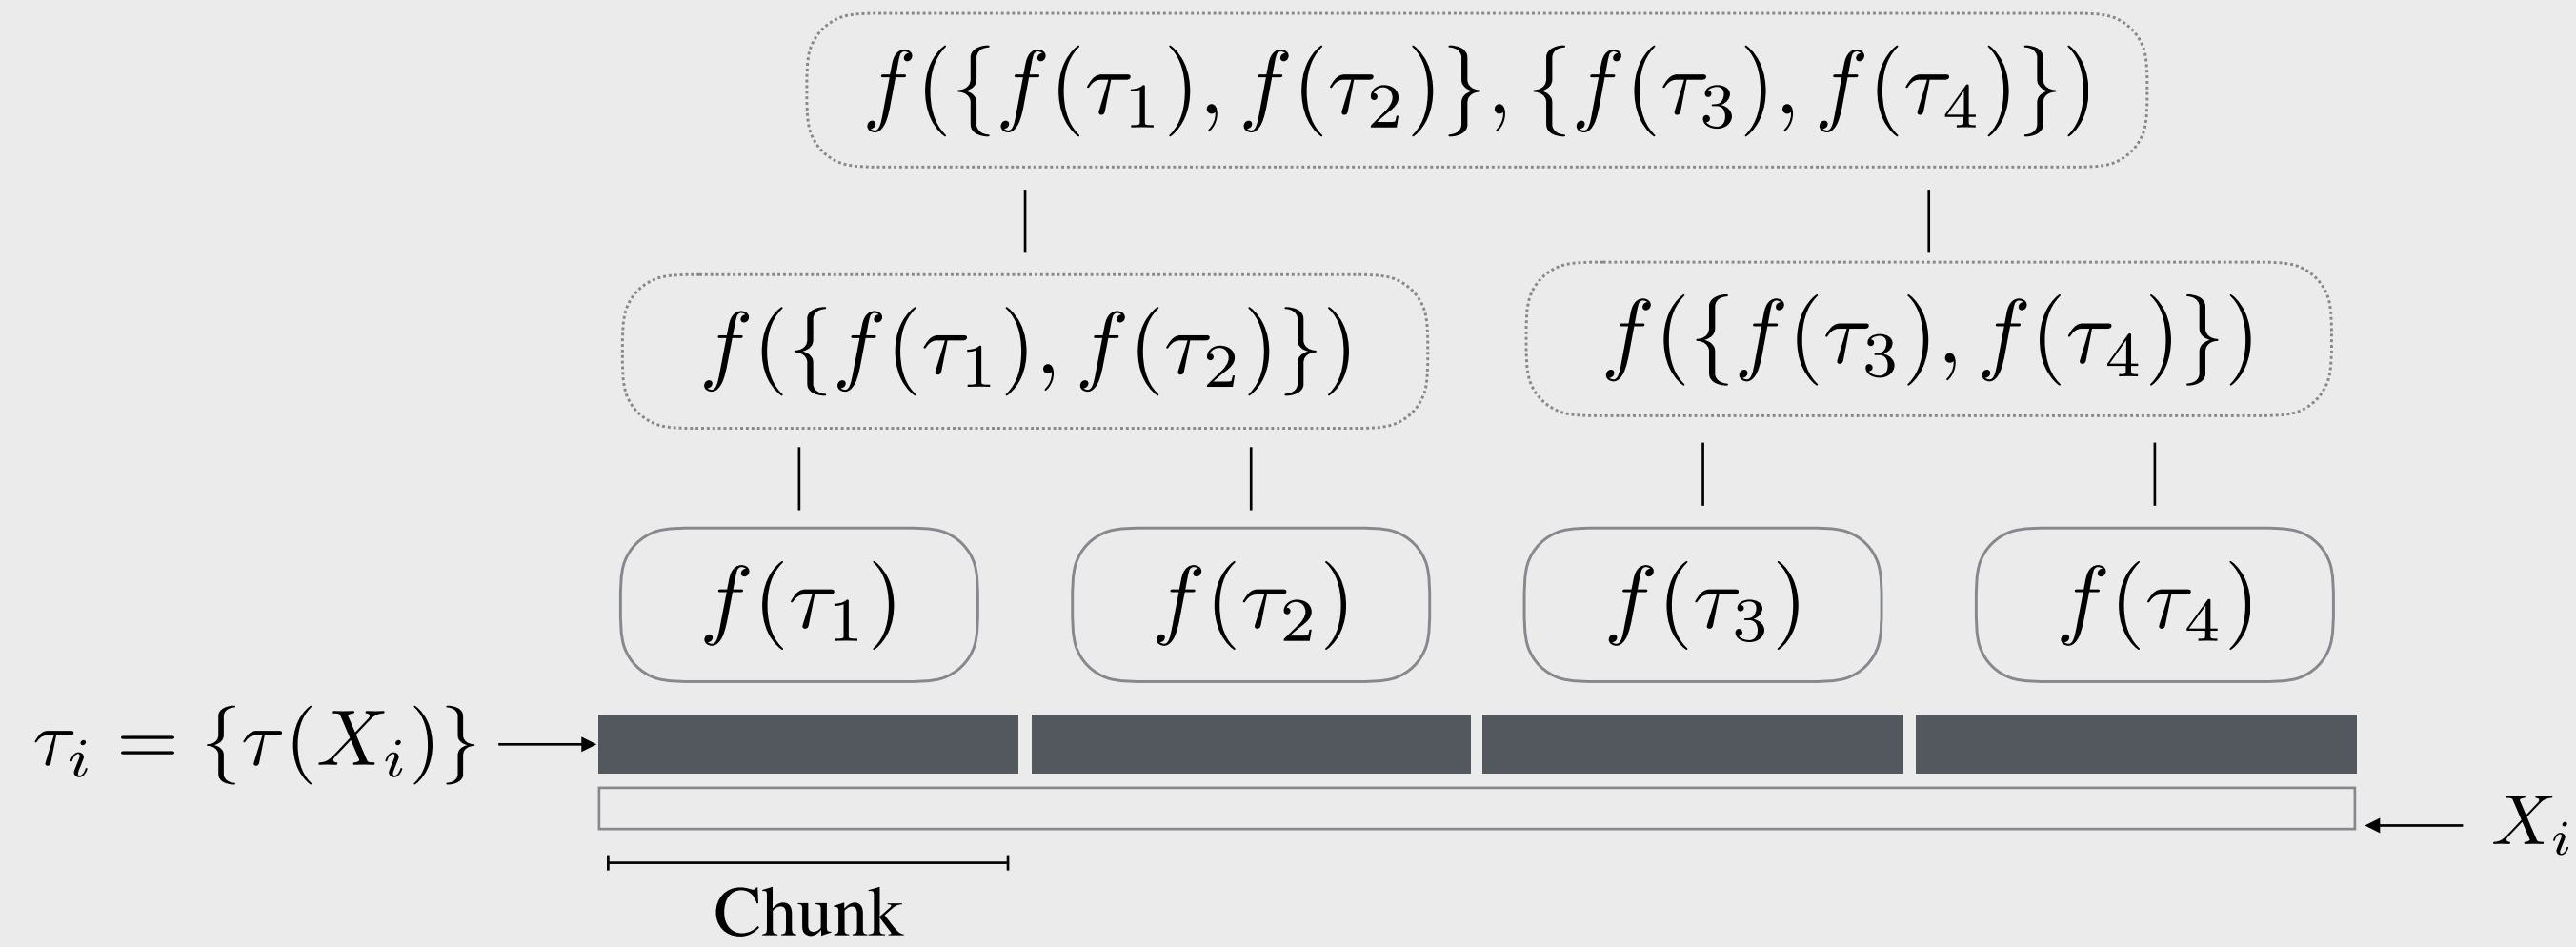
1 thread

In-memory

40M rows, 100 columns, 5 statistic types




# Segment Tree in Data Canopy



Leaves: basic aggregate on every chunk

Internal node: Aggregate function applied to its two children

$$*f(X) = f(\{ f(X_1), f(X_2), \dots f(X_n) \})$$



# Operation Modes


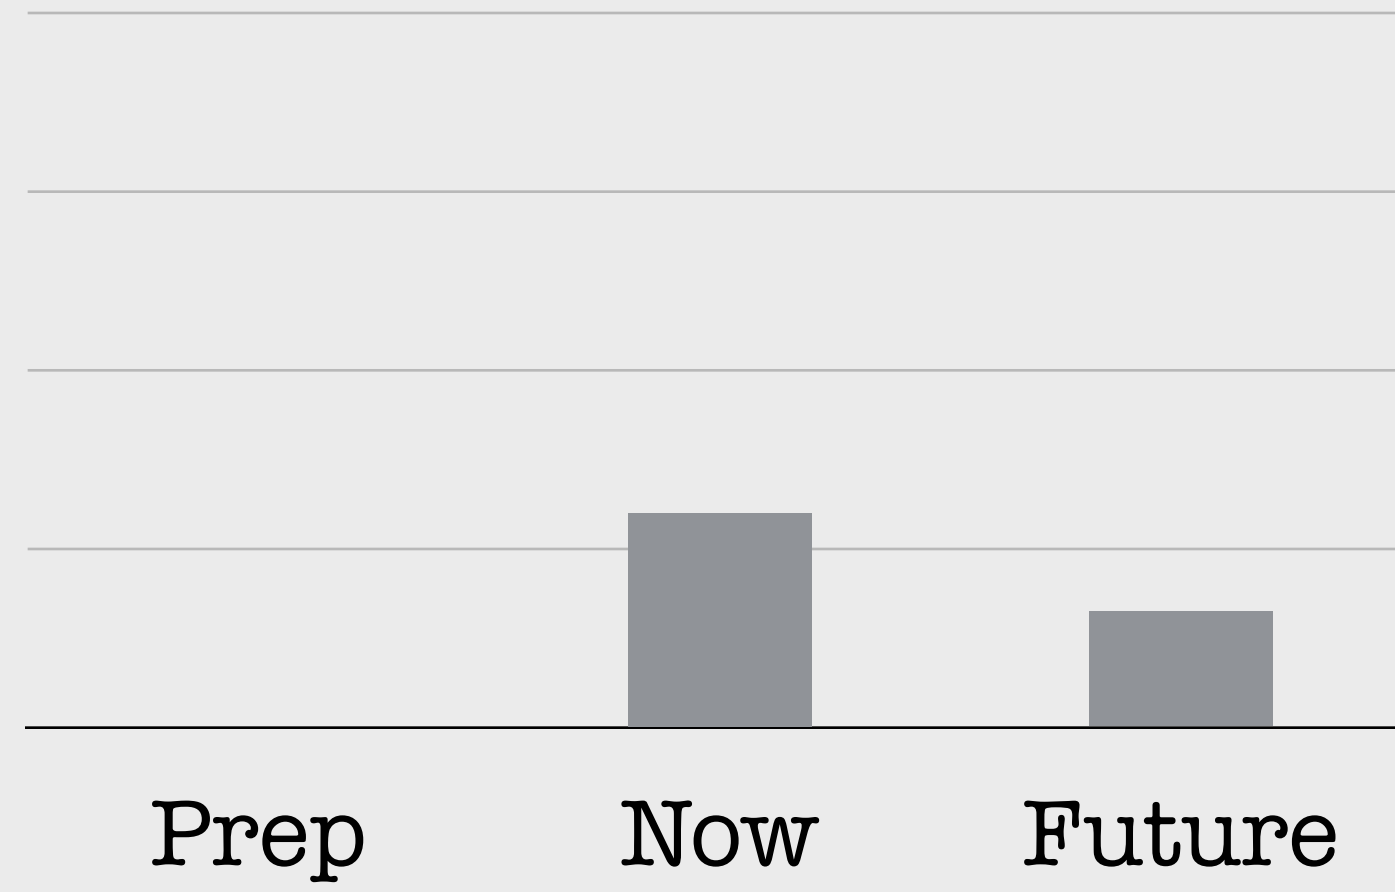
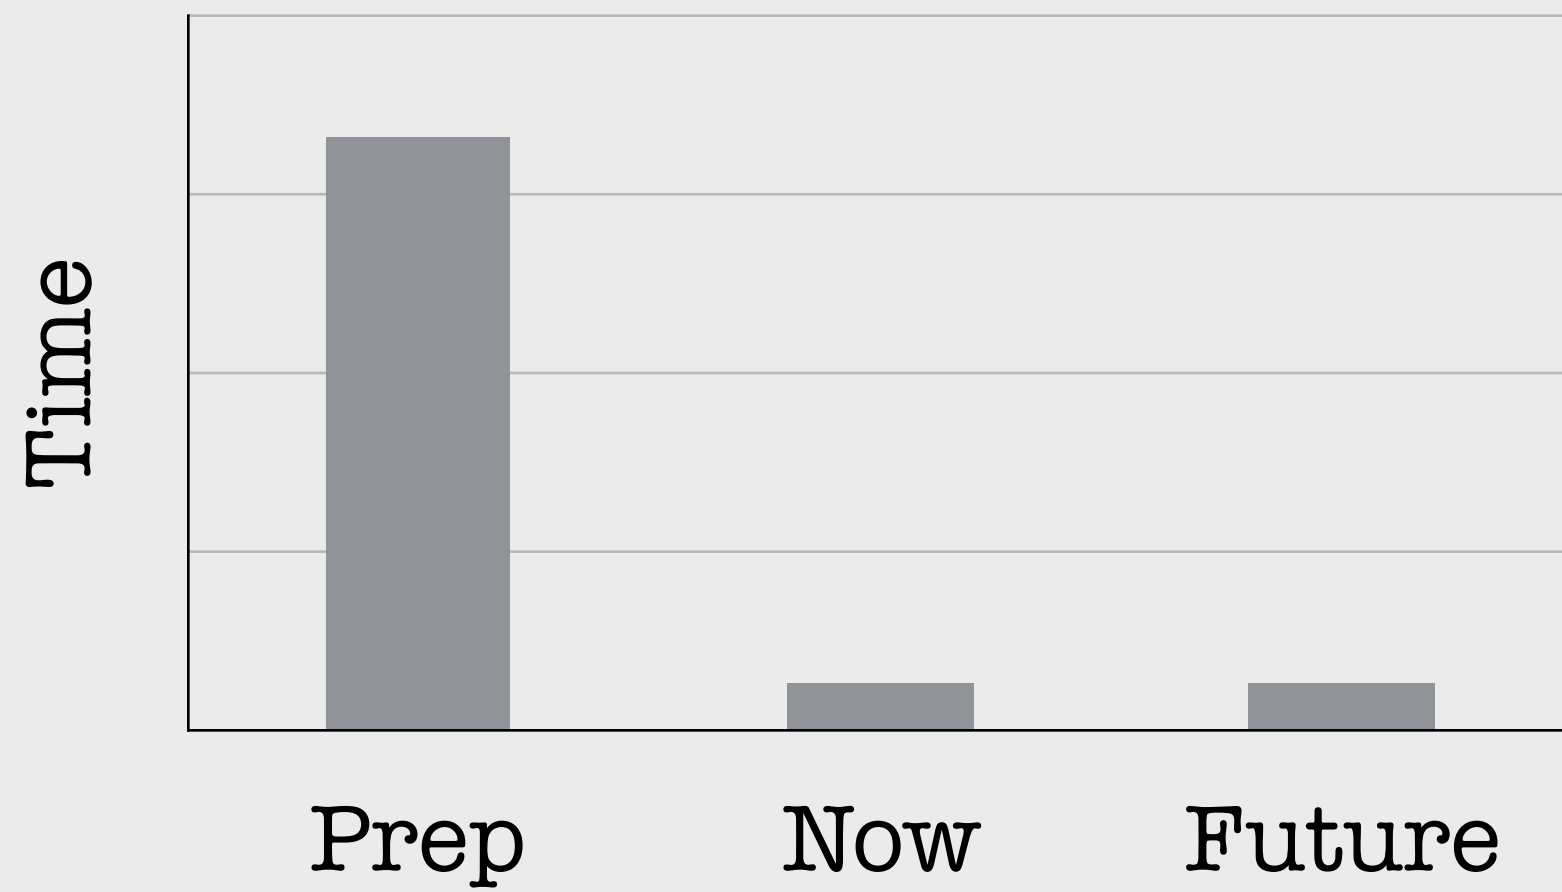
Offline





Online

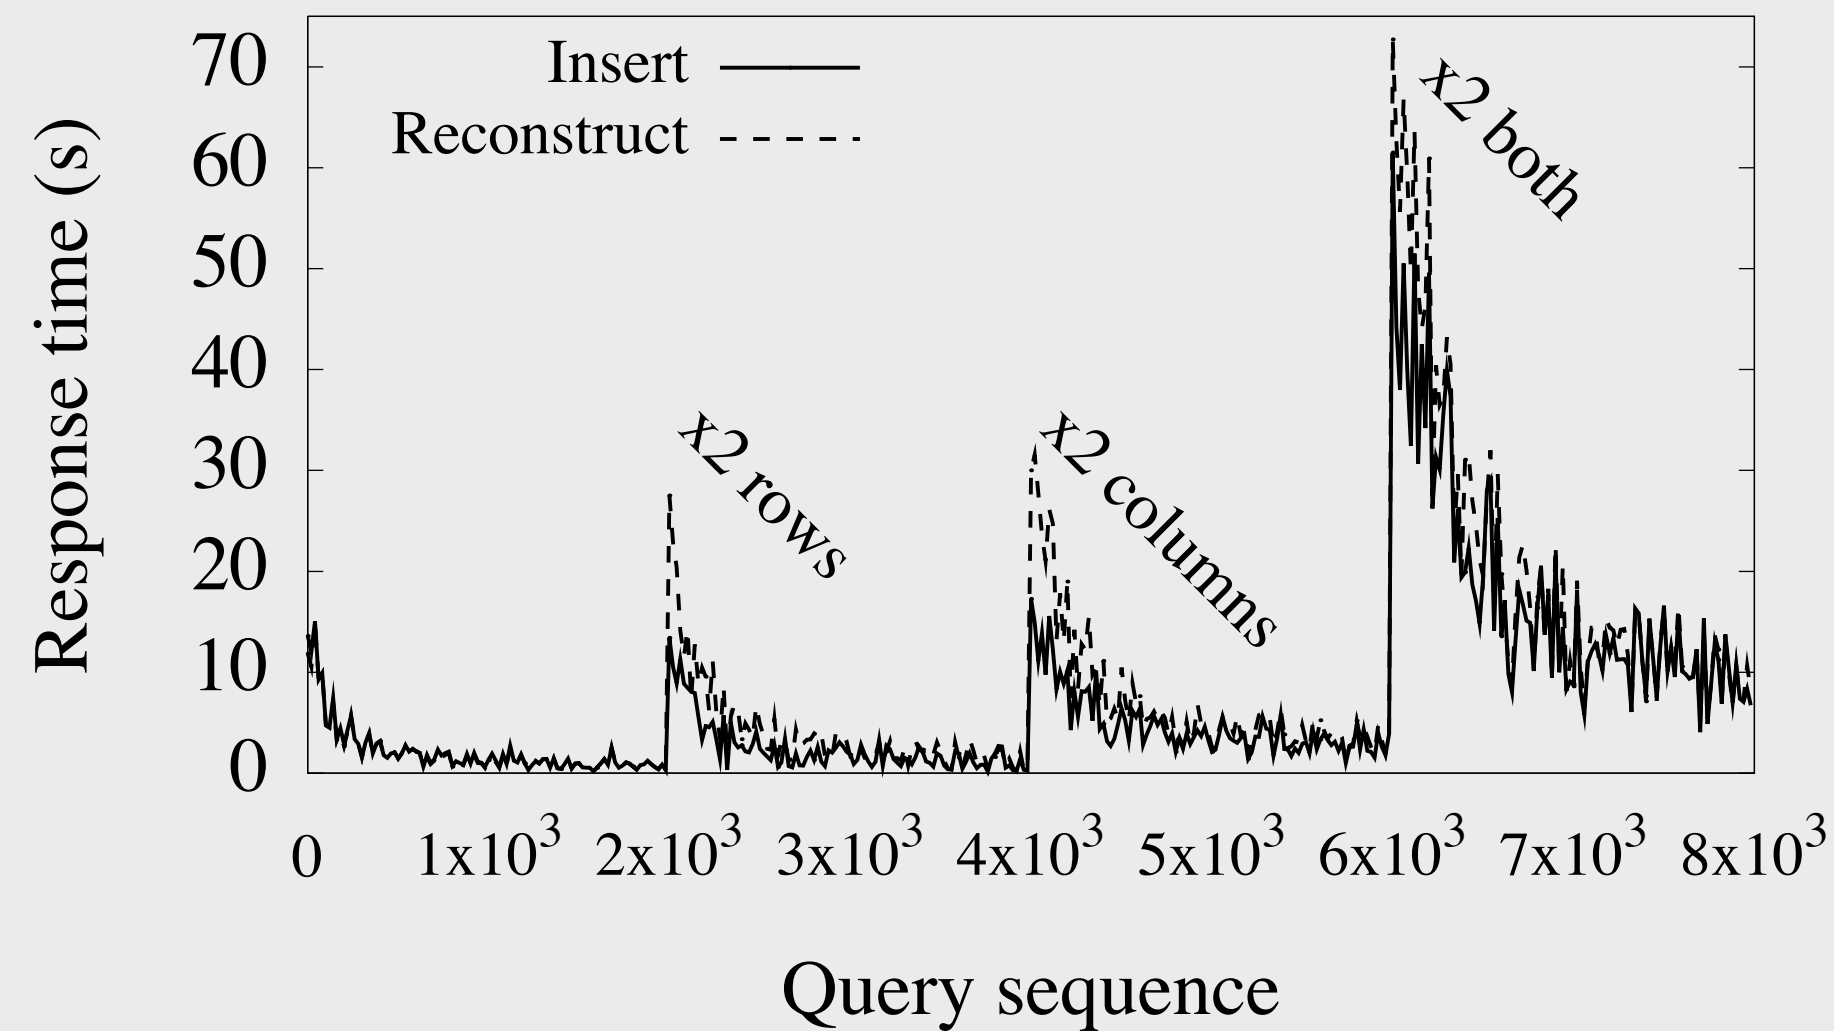


Speculative

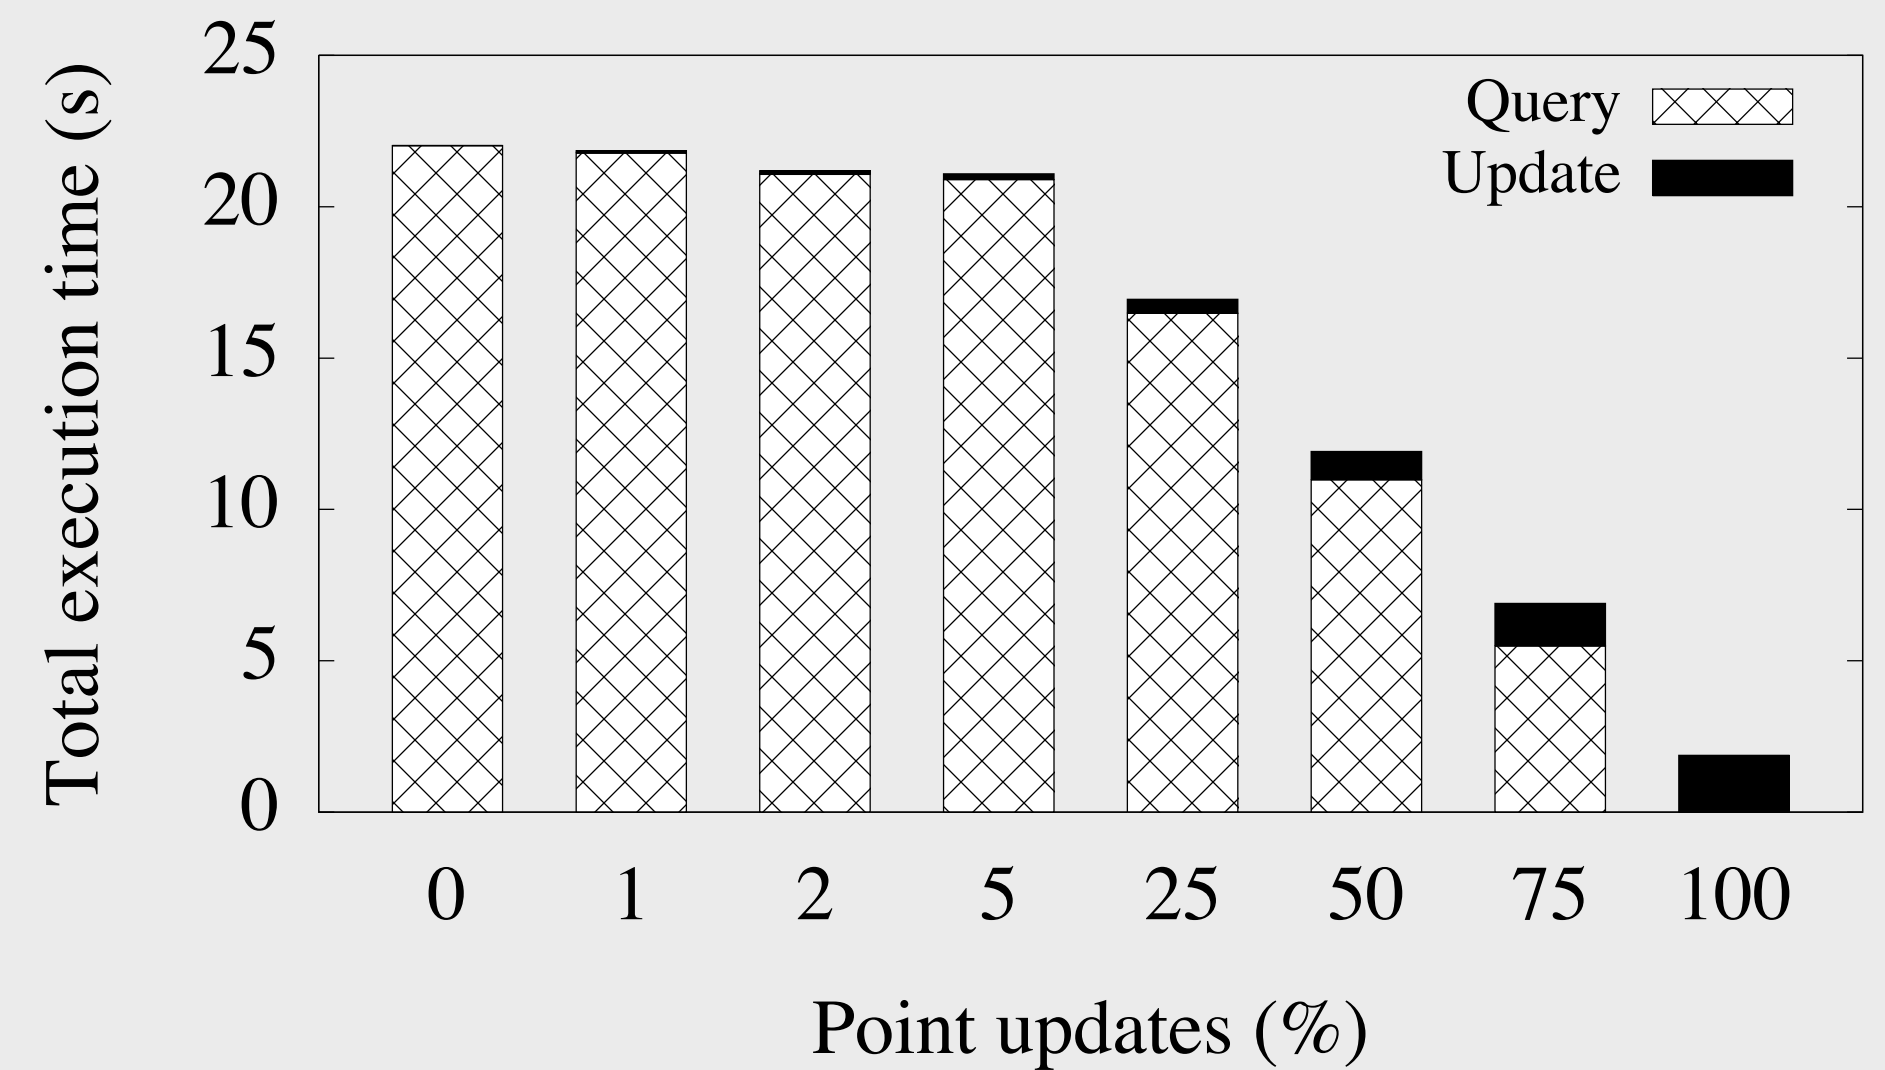




Extend segment trees as needed

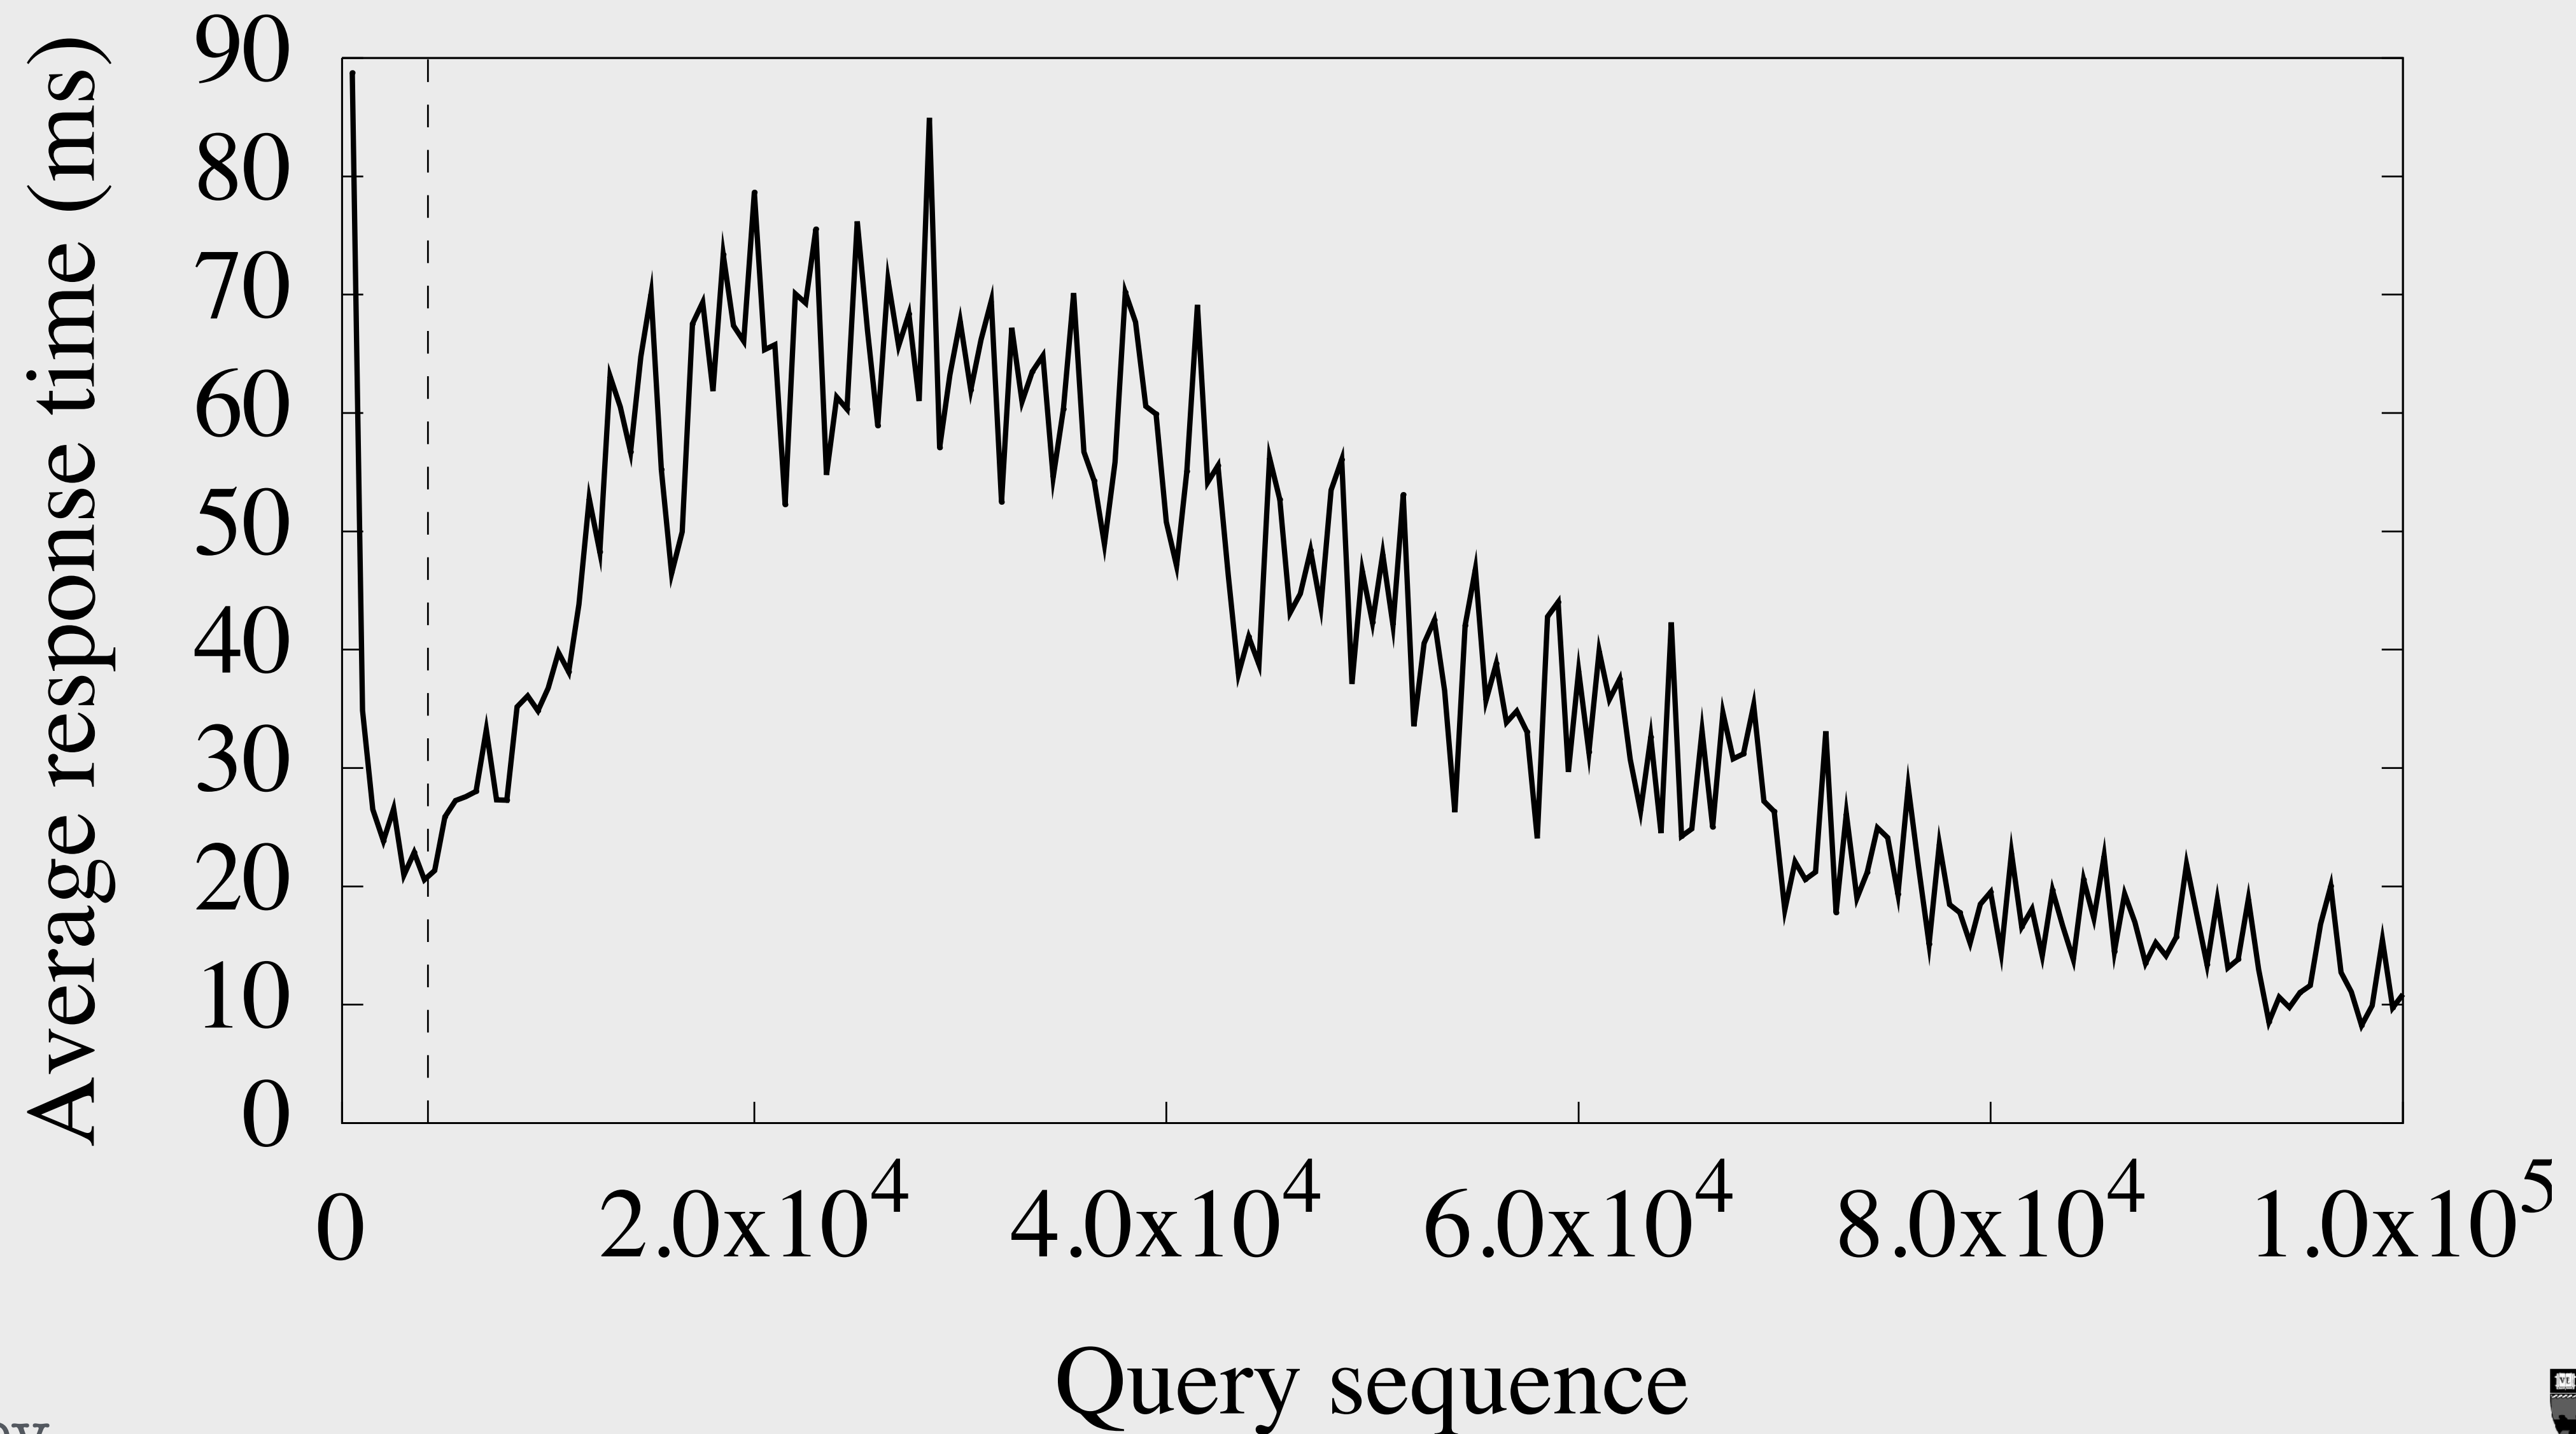


Inserts

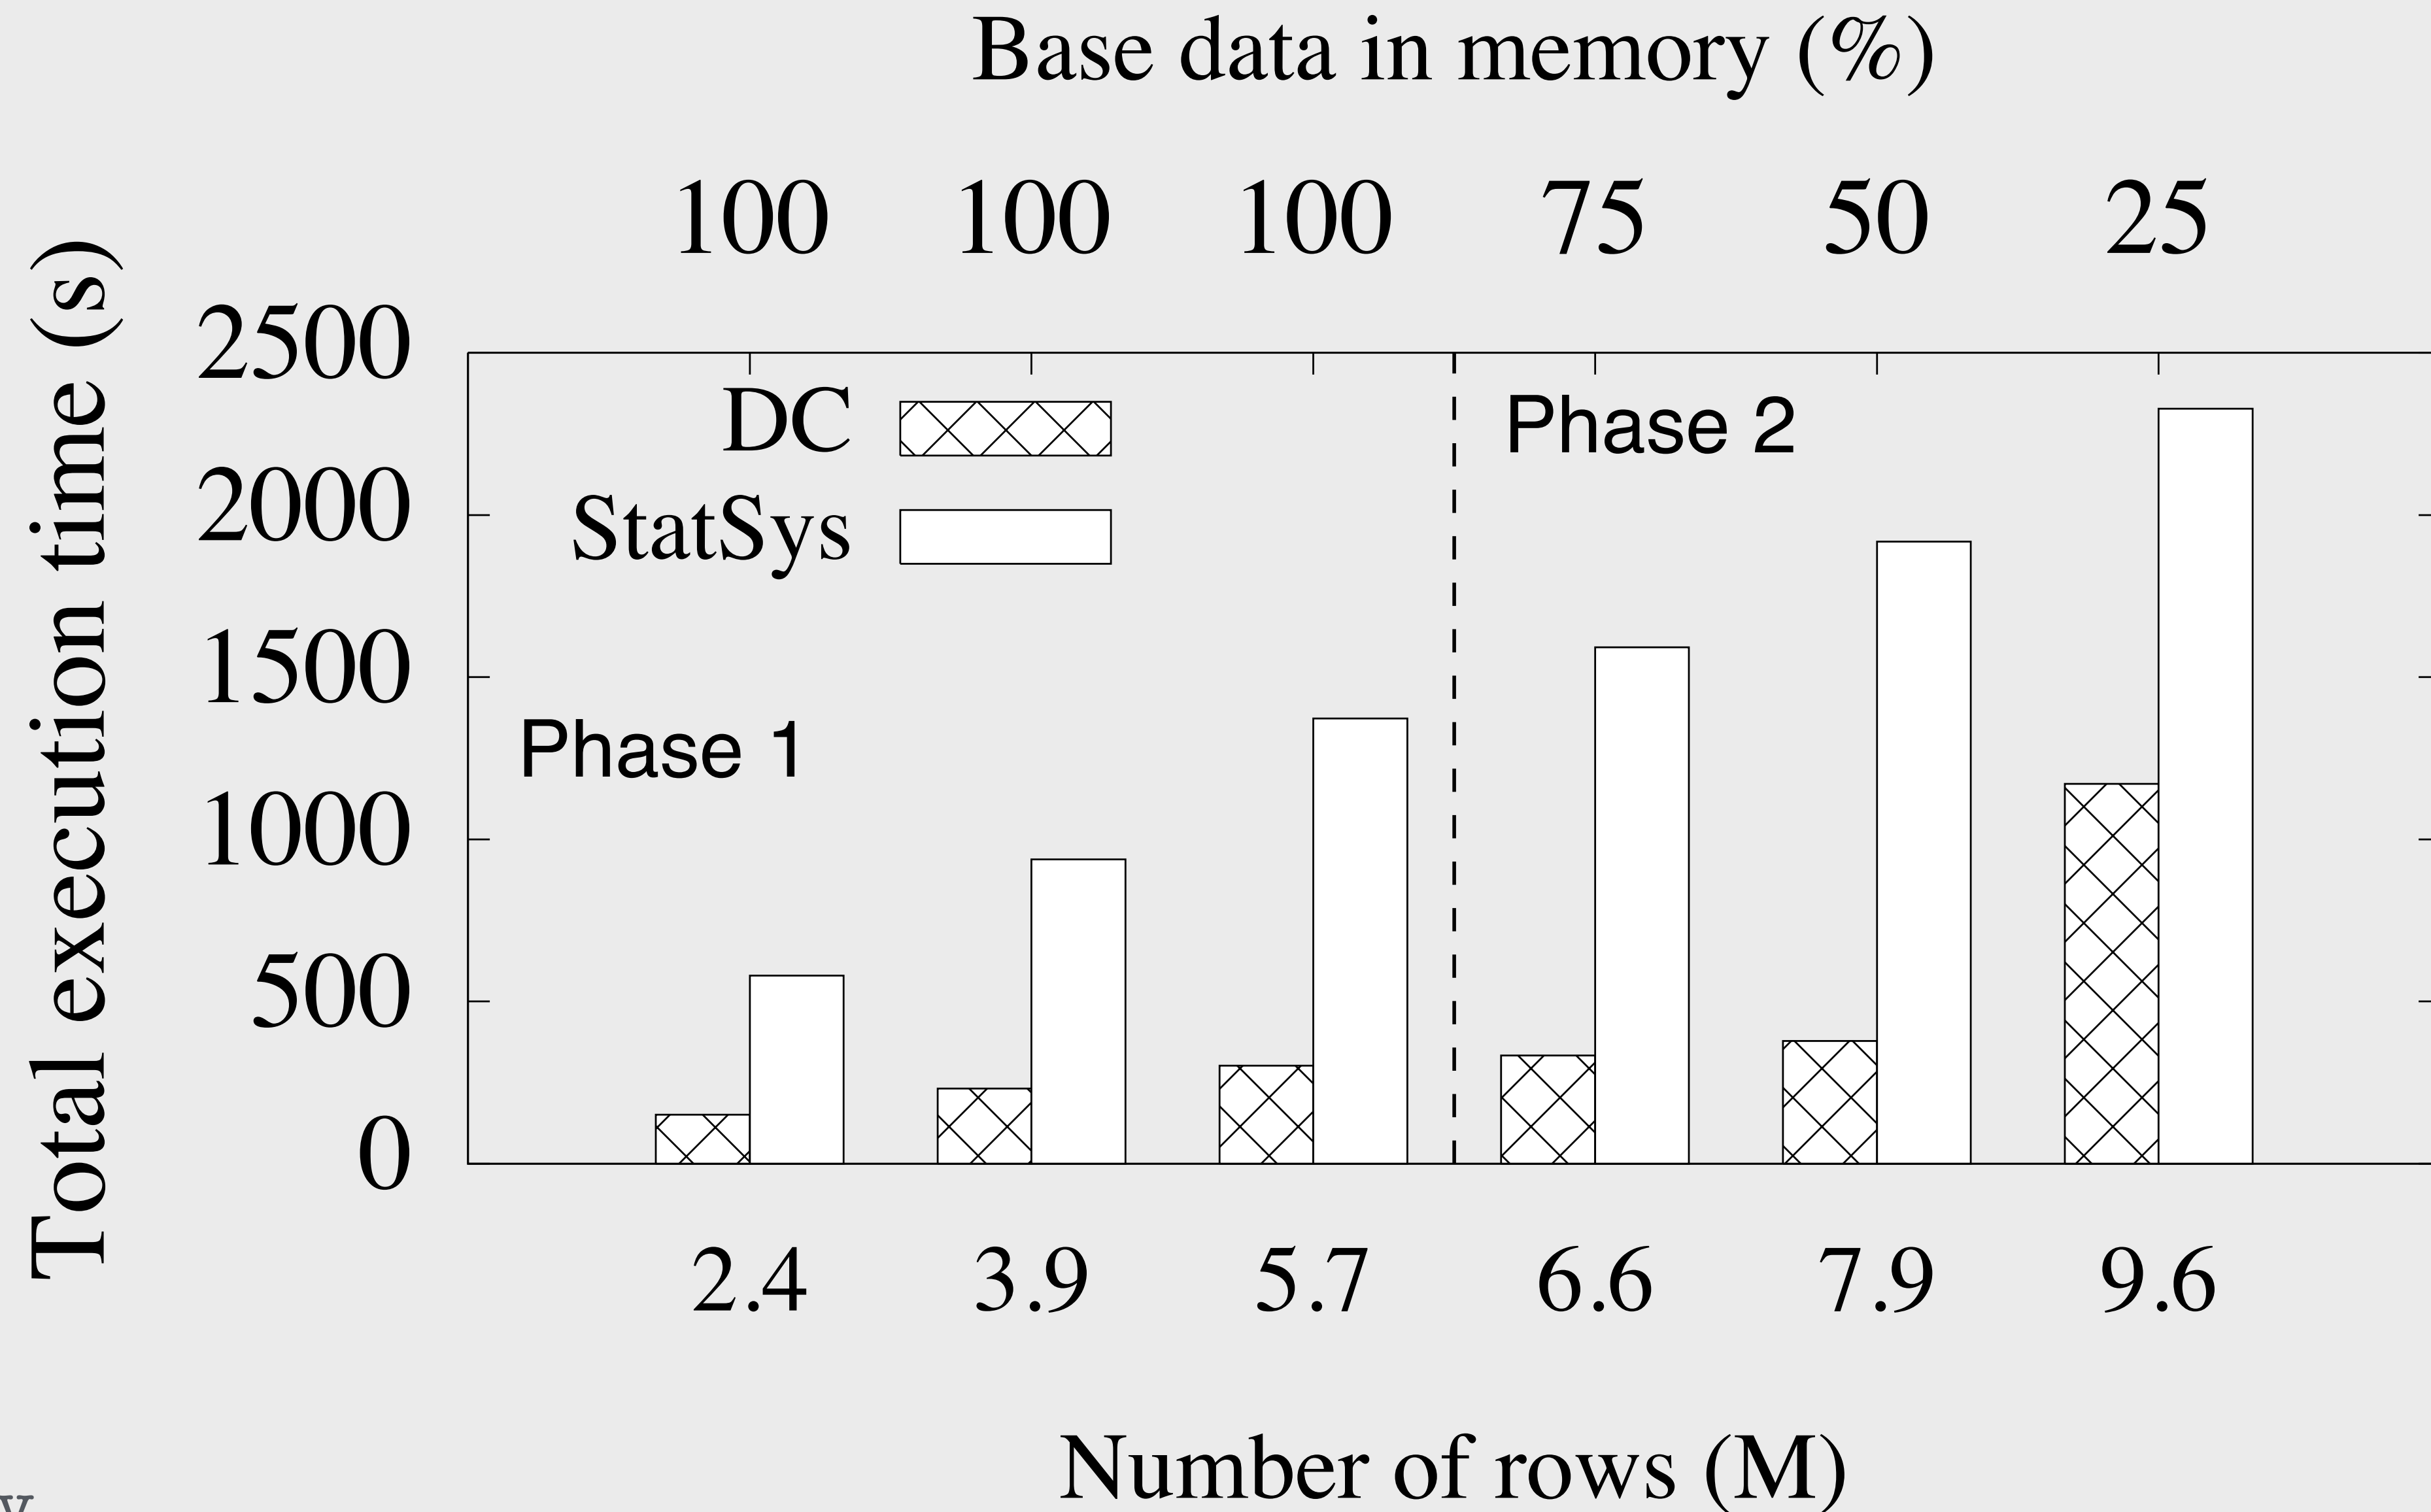


Updates


# Performance under Memory Pressure




# Handling Memory Pressure



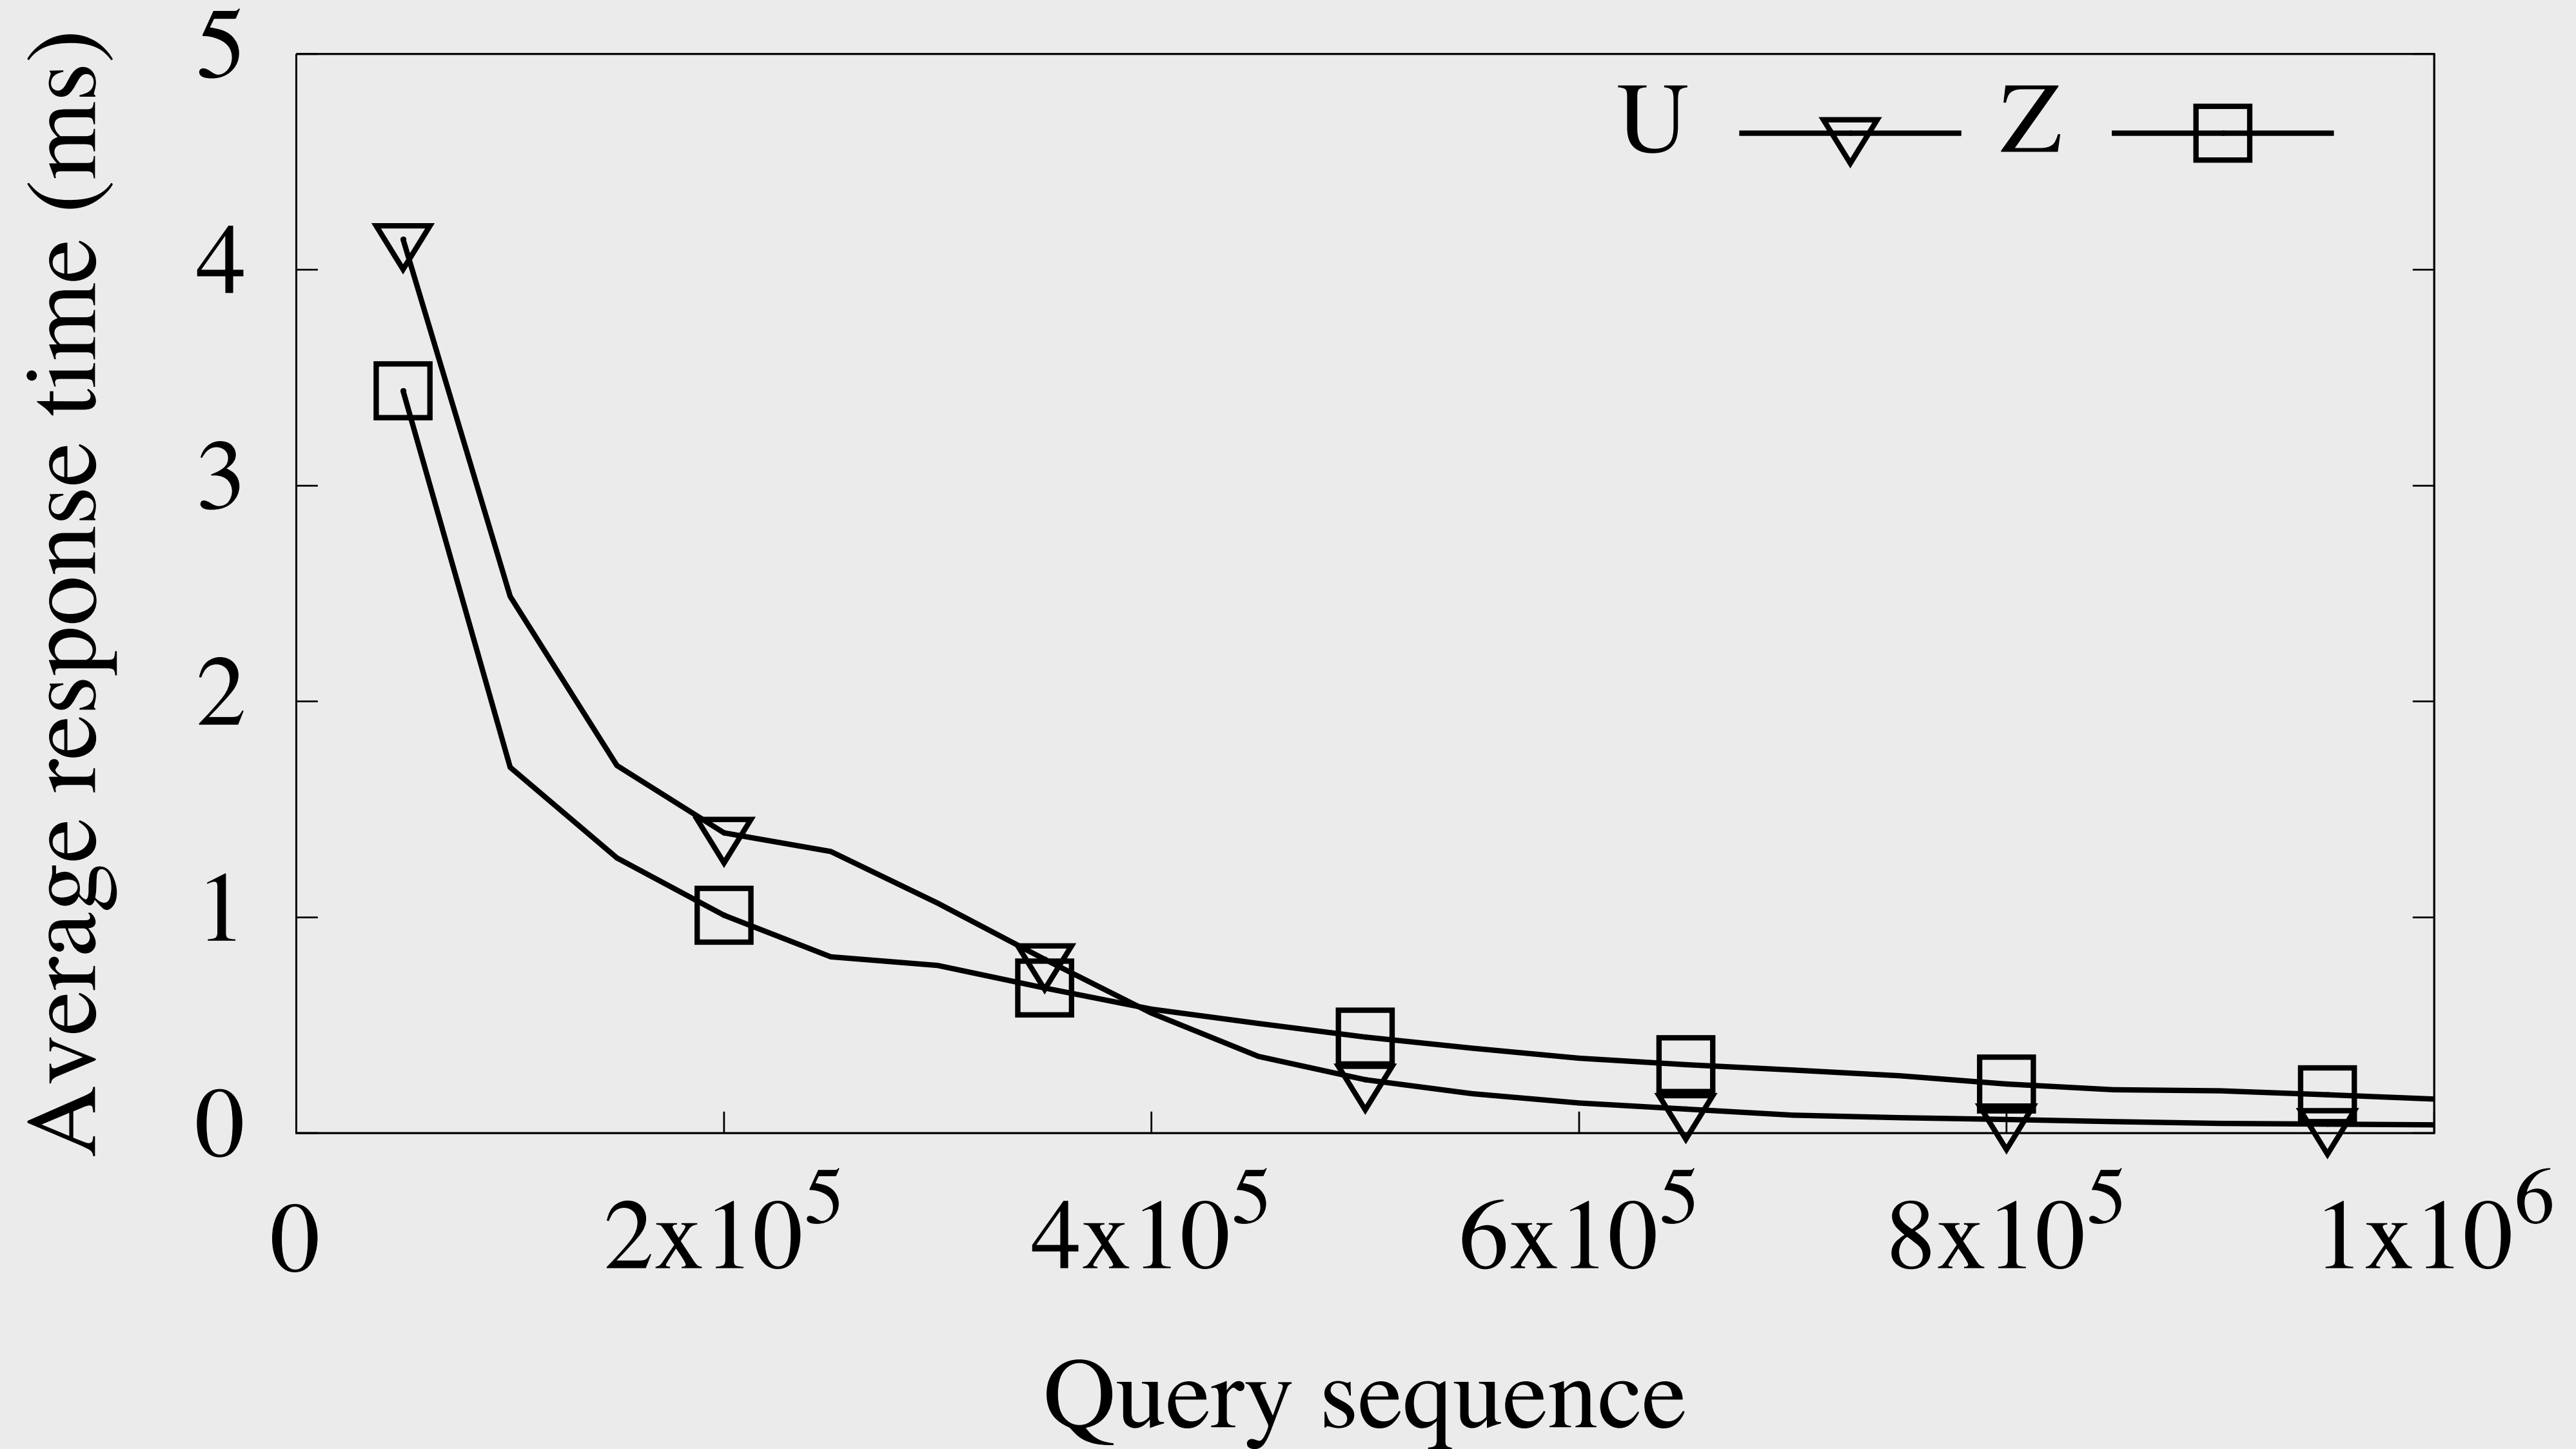
# Memory Feasibility – 8GB



# Memory Footprint



# Scaling with Queries



# Selecting the Chunk Size

