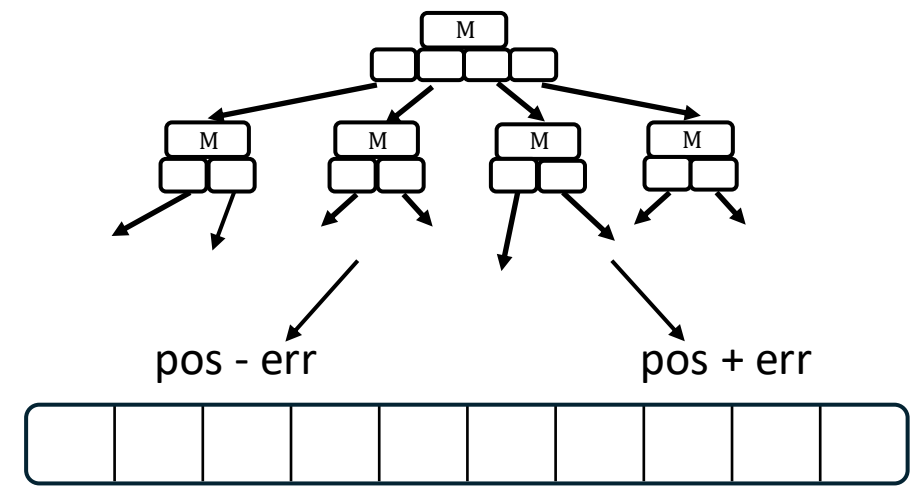


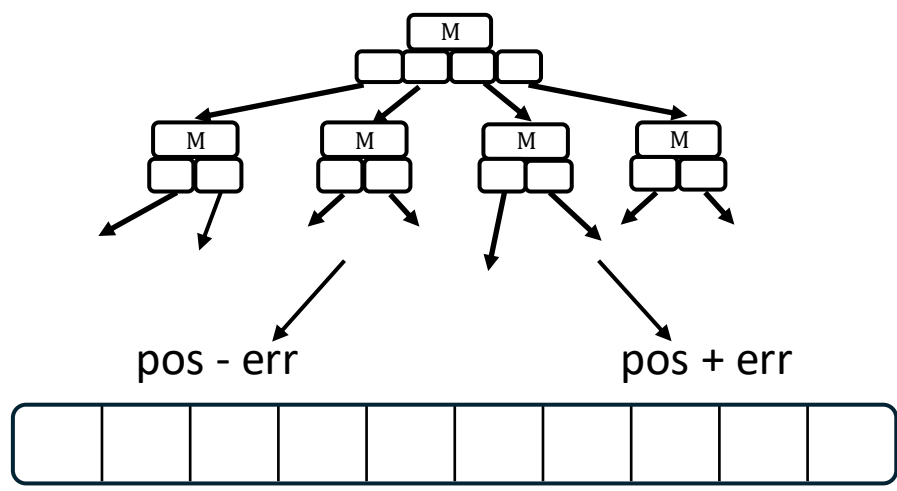
Improving the physical-to-sorted mapping for Learned Indexes

Anwesh Saha, Aneesh Raman, Ryan Marcus, Manos Athanassoulis

Learned Indexes

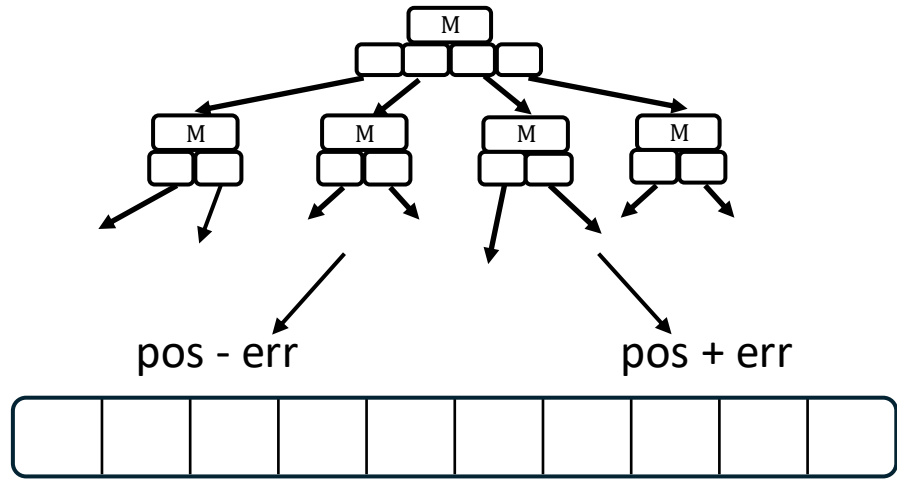


Learned Indexes



Lower Memory

Learned Indexes

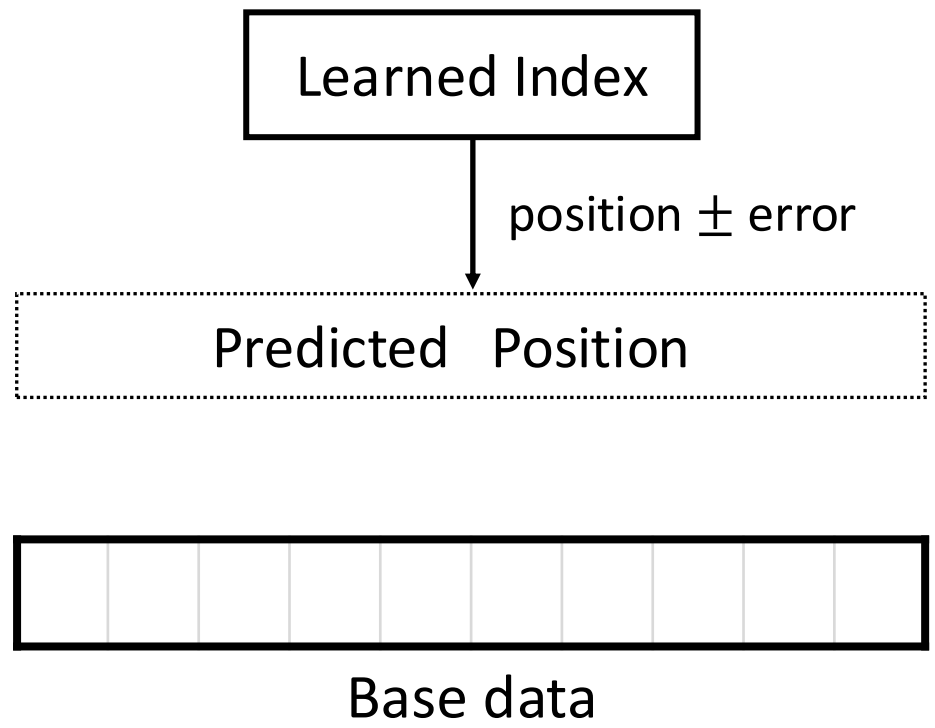


Lower Memory

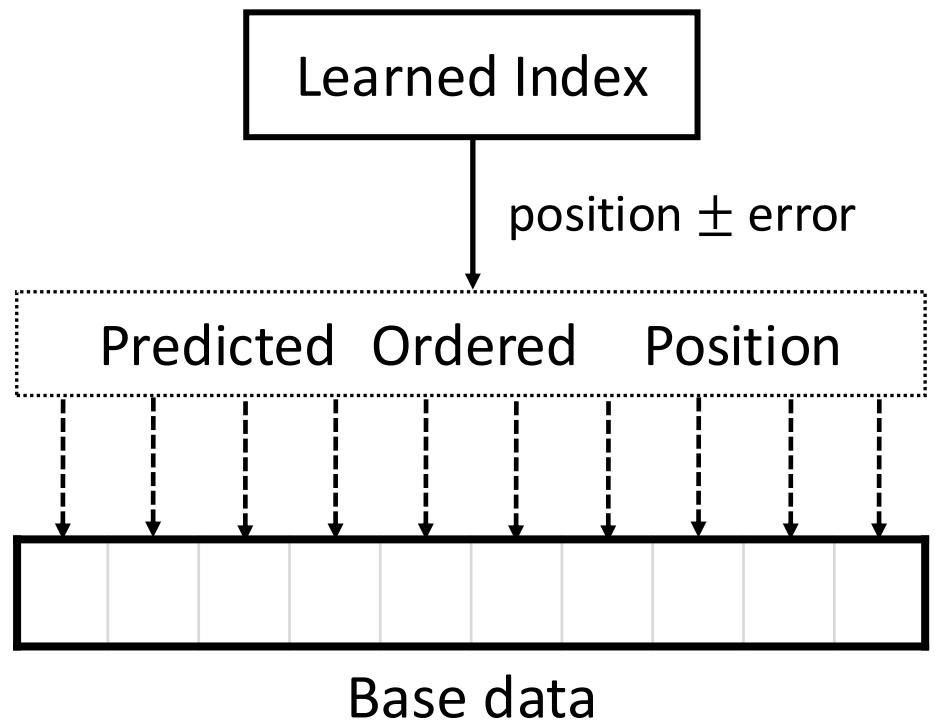


Faster Access

Learned Indexes Assumptions

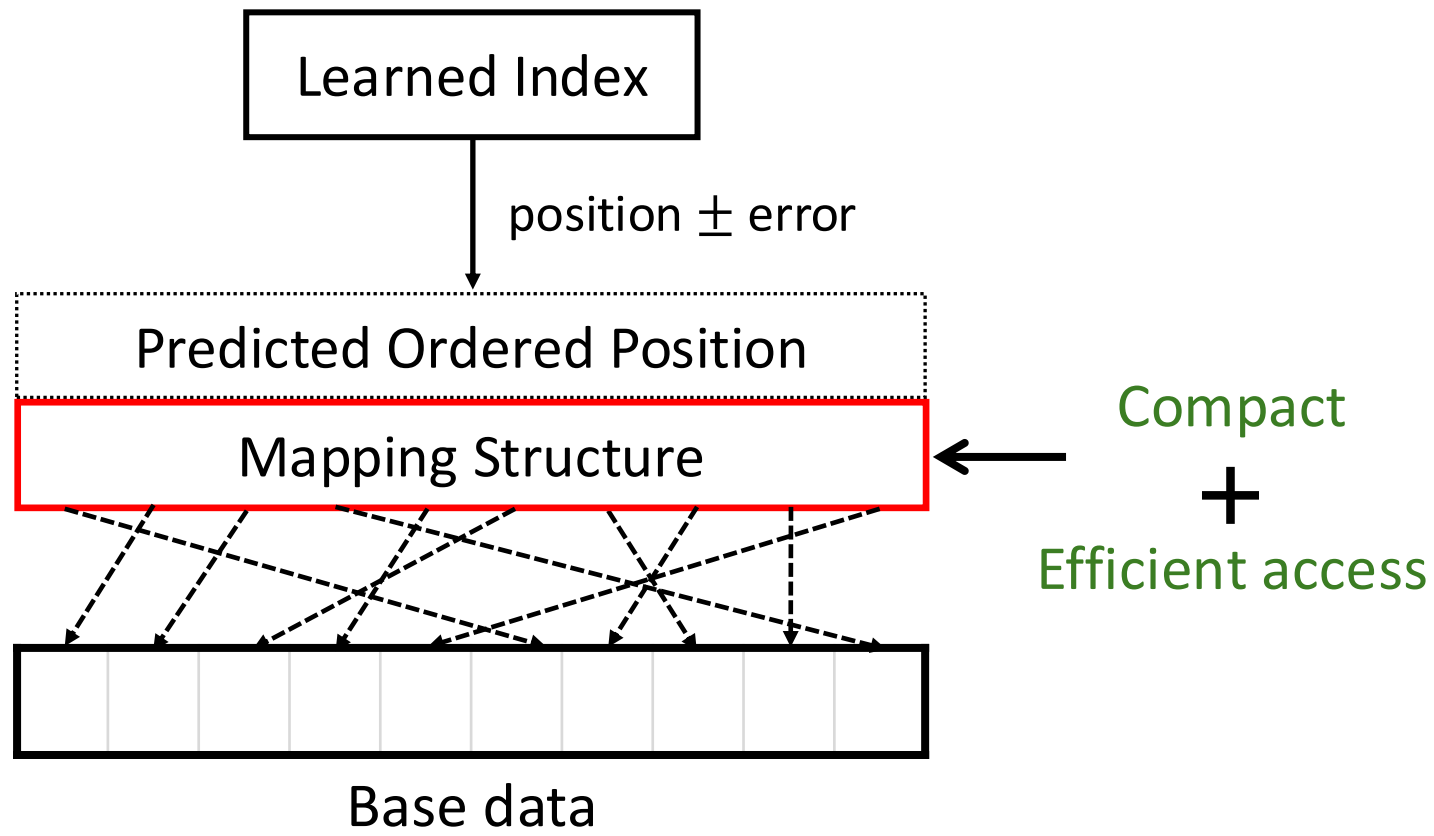


Learned Indexes Assumptions



Assumes Sorted data underneath!

Proposed Design

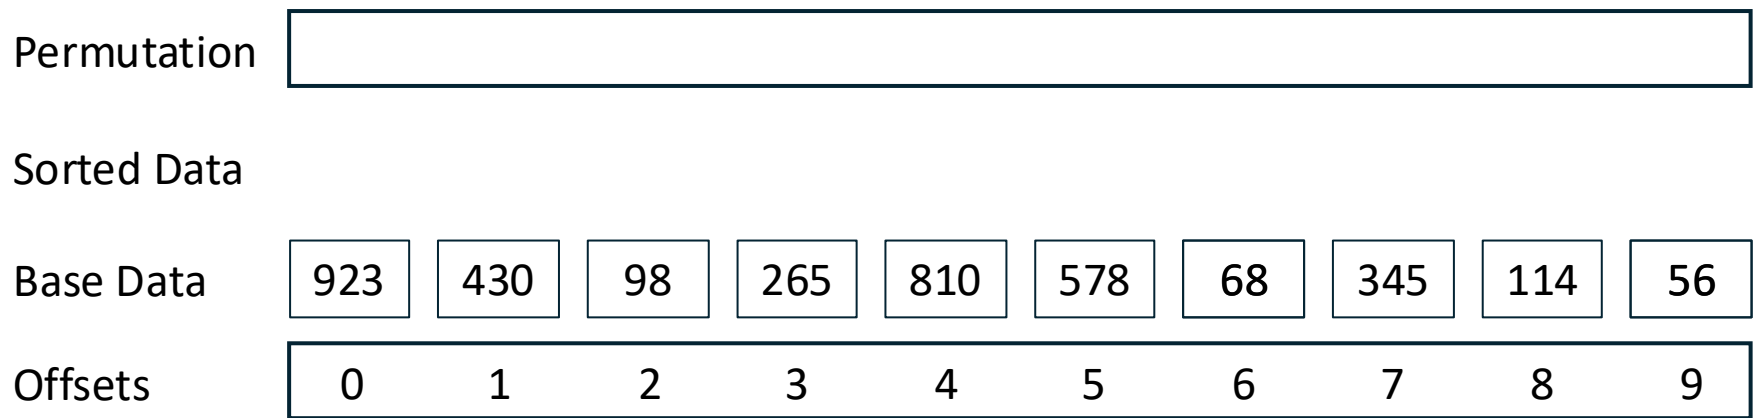


Access(Mapping structure) << Access(Traditional index)

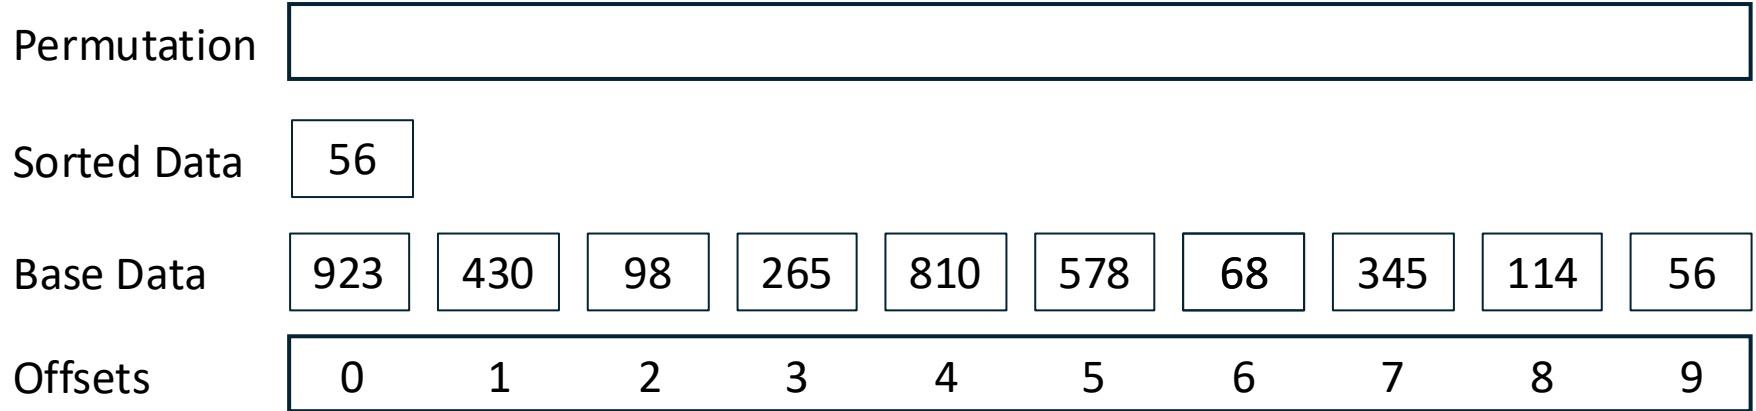
Prior Work

Base Data	923	430	98	265	810	578	68	345	114	56
Offsets	0	1	2	3	4	5	6	7	8	9

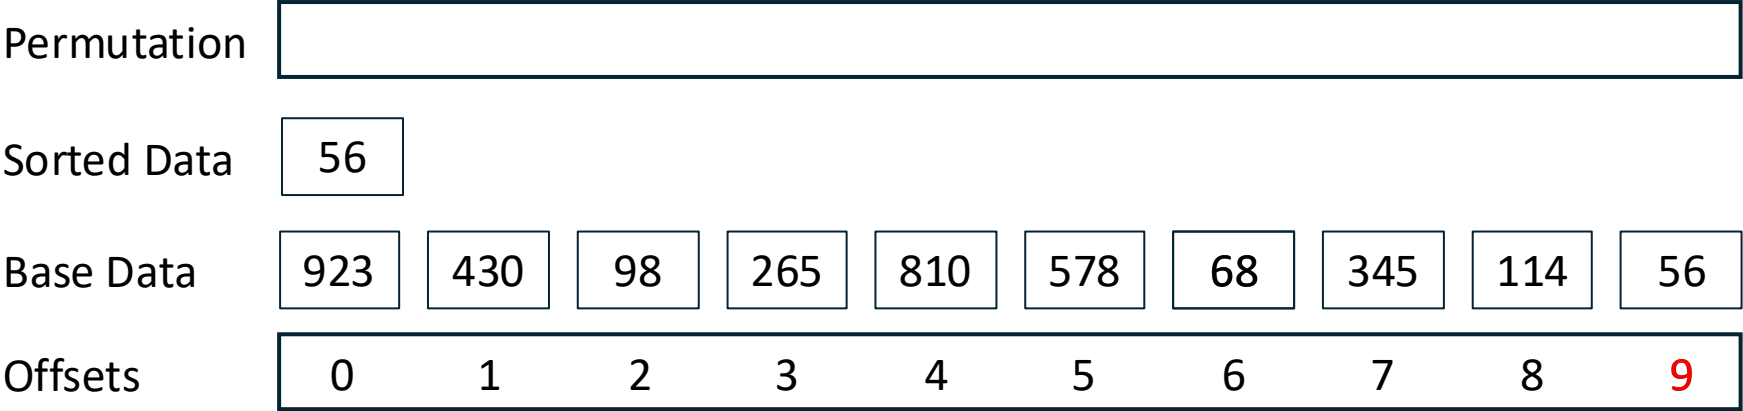
Prior Work



Prior Work



Prior Work



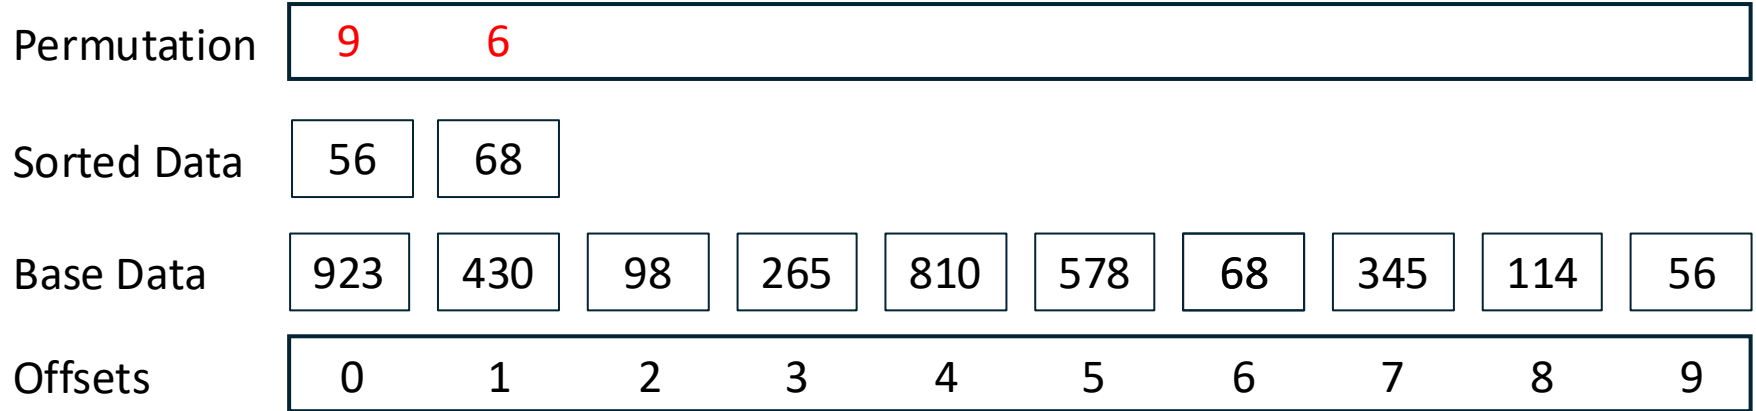
Prior Work

Permutation	9									
Sorted Data	56									
Base Data	923	430	98	265	810	578	68	345	114	56
Offsets	0	1	2	3	4	5	6	7	8	9

Prior Work

Permutation	9									
Sorted Data	56	68								
Base Data	923	430	98	265	810	578	68	345	114	56
Offsets	0	1	2	3	4	5	6	7	8	9

Prior Work



Prior Work

Permutation	9	6	2	8	3	7	1	5	4	0
Sorted Data	56	68	98	114	265	345	430	578	810	923
Base Data	923	430	98	265	810	578	68	345	114	56
Offsets	0	1	2	3	4	5	6	7	8	9

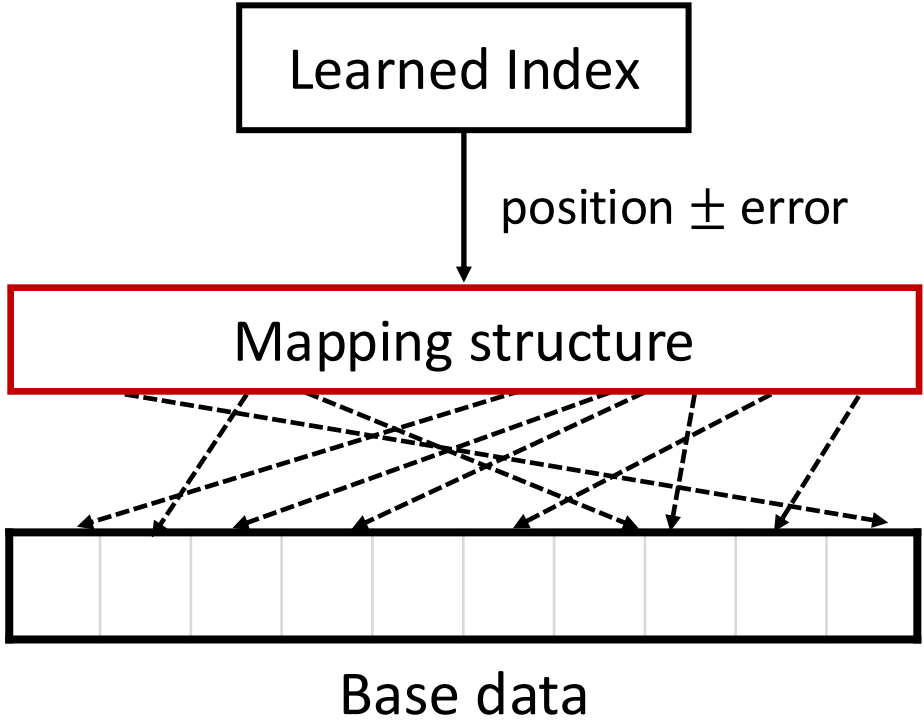
Prior Work

Mapping Structure

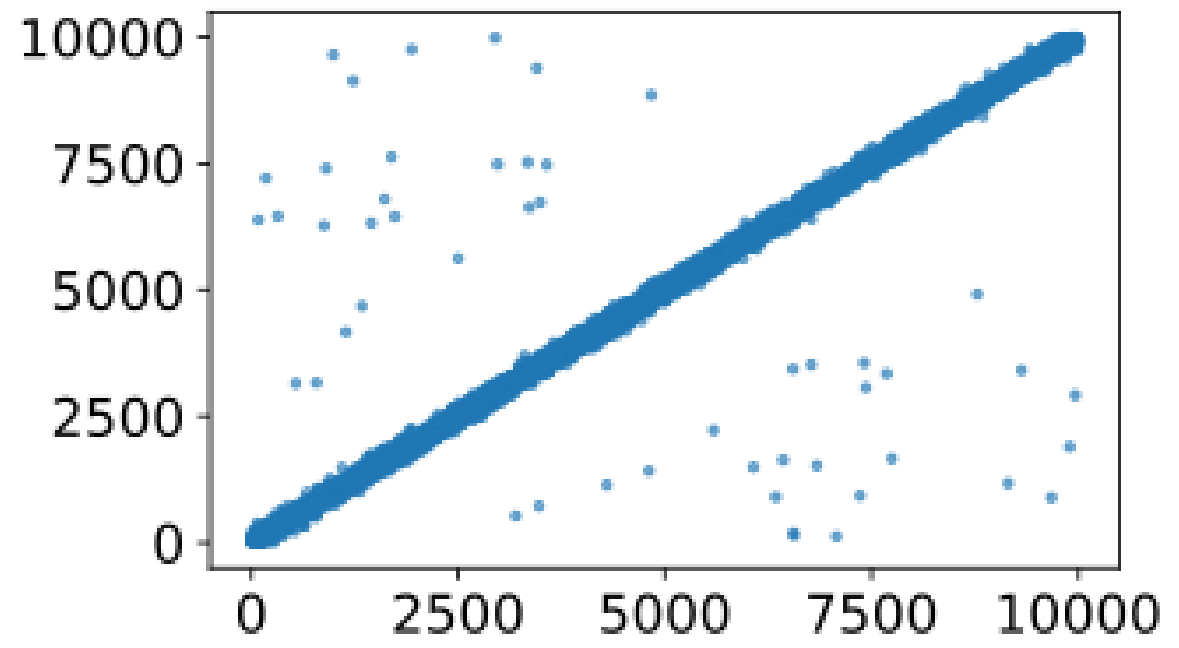
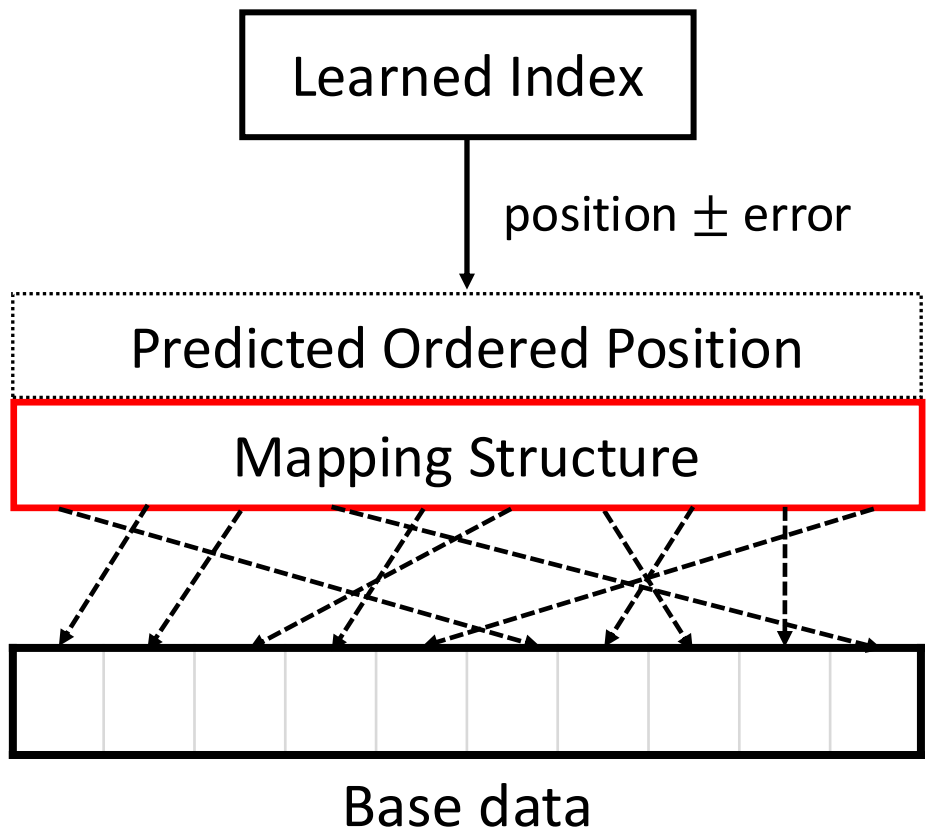
Permutation	9	6	2	8	3	7	1	5	4	0
Sorted Data	56	68	98	114	265	345	430	578	810	923
Base Data	923	430	98	265	810	578	68	345	114	56
Offsets	0	1	2	3	4	5	6	7	8	9

✗ Permutation Vector size increases with size of data

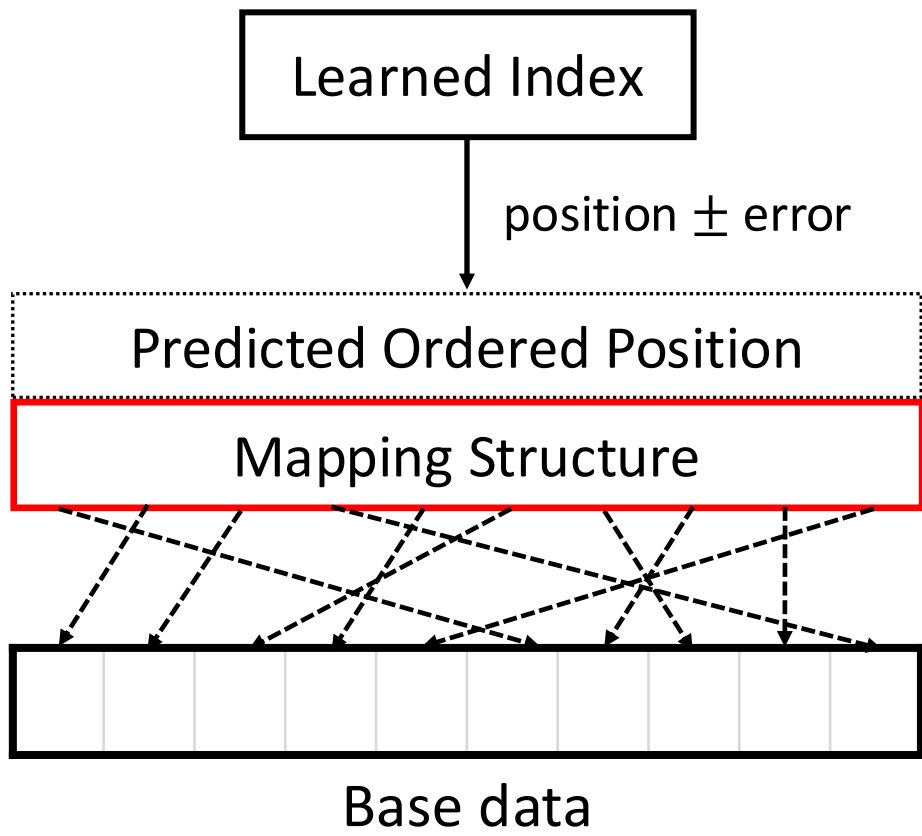
What do we Want?



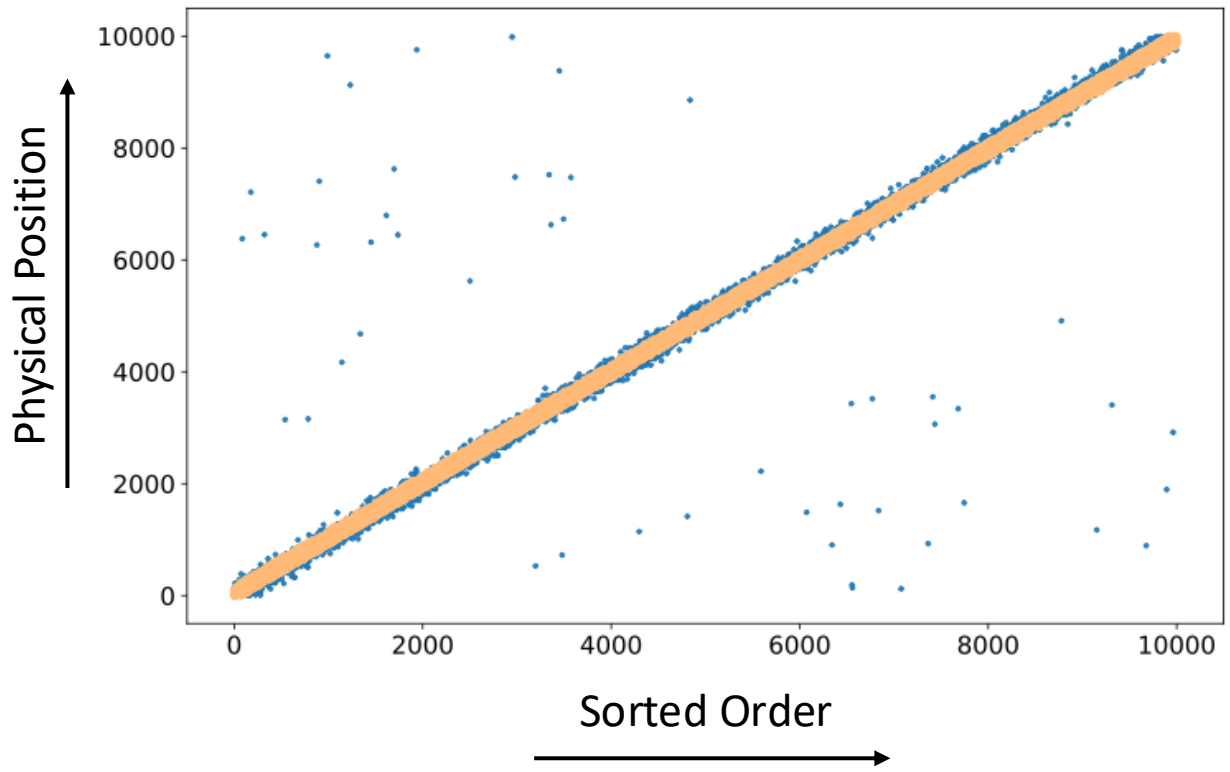
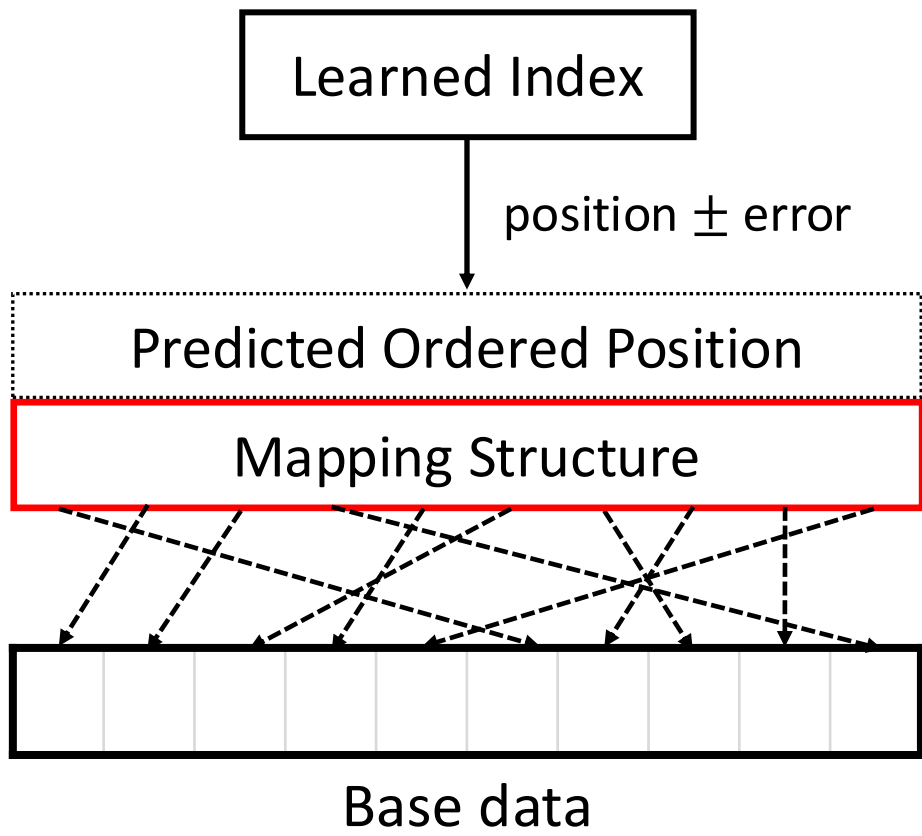
Constellation Maps as Mapping Structures



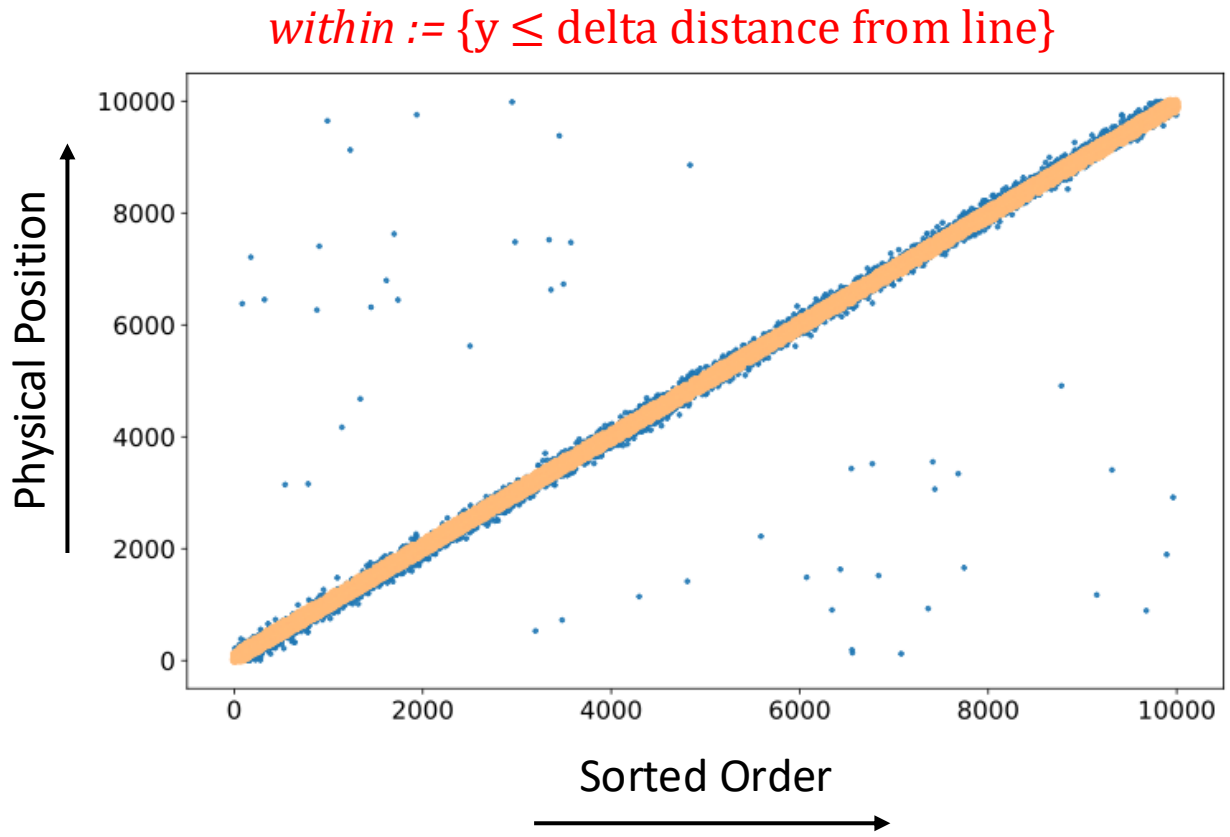
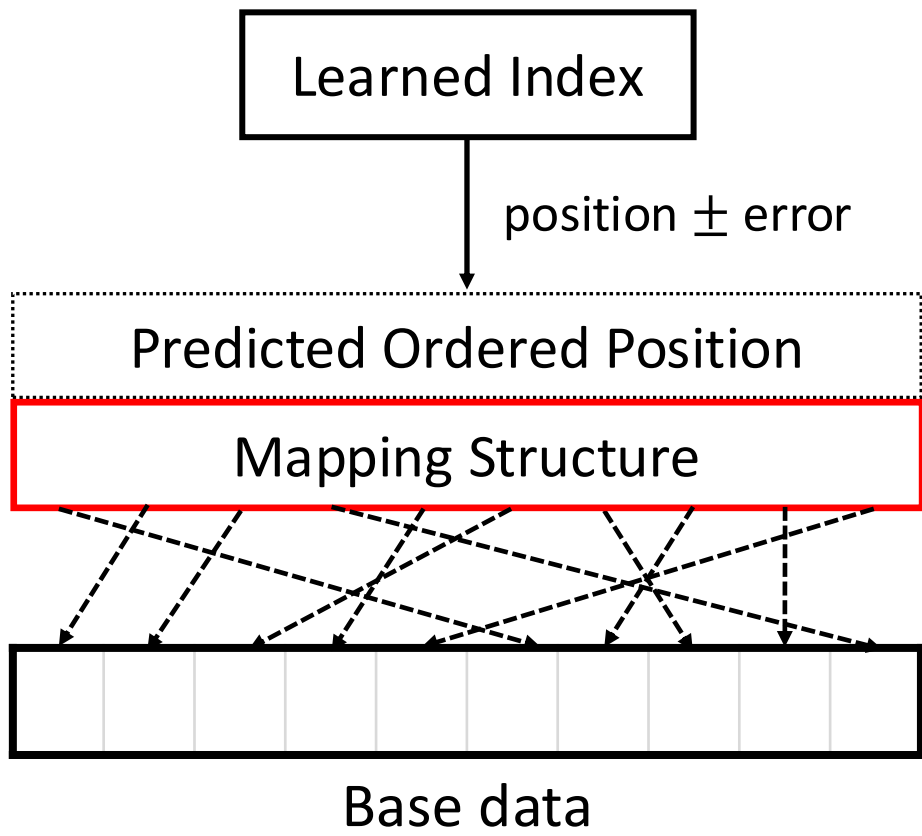
Constellation Maps as Mapping Structures



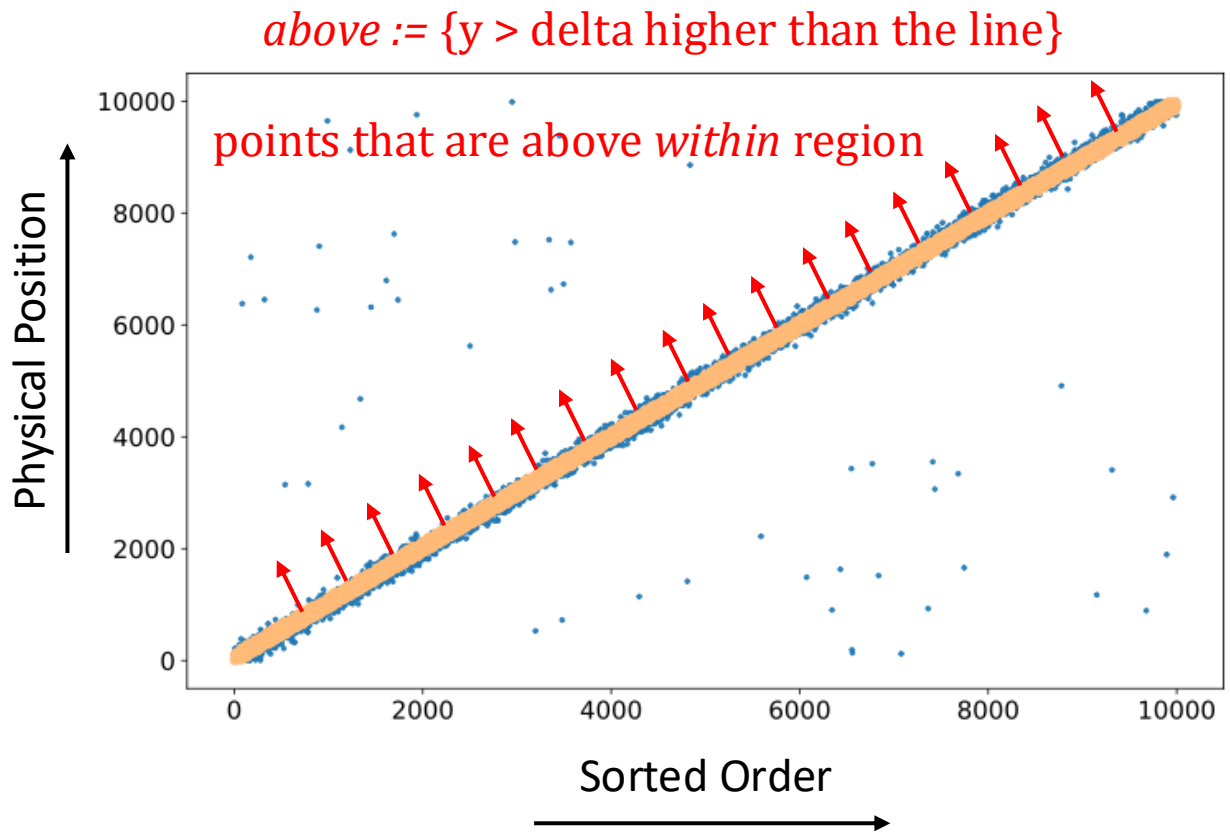
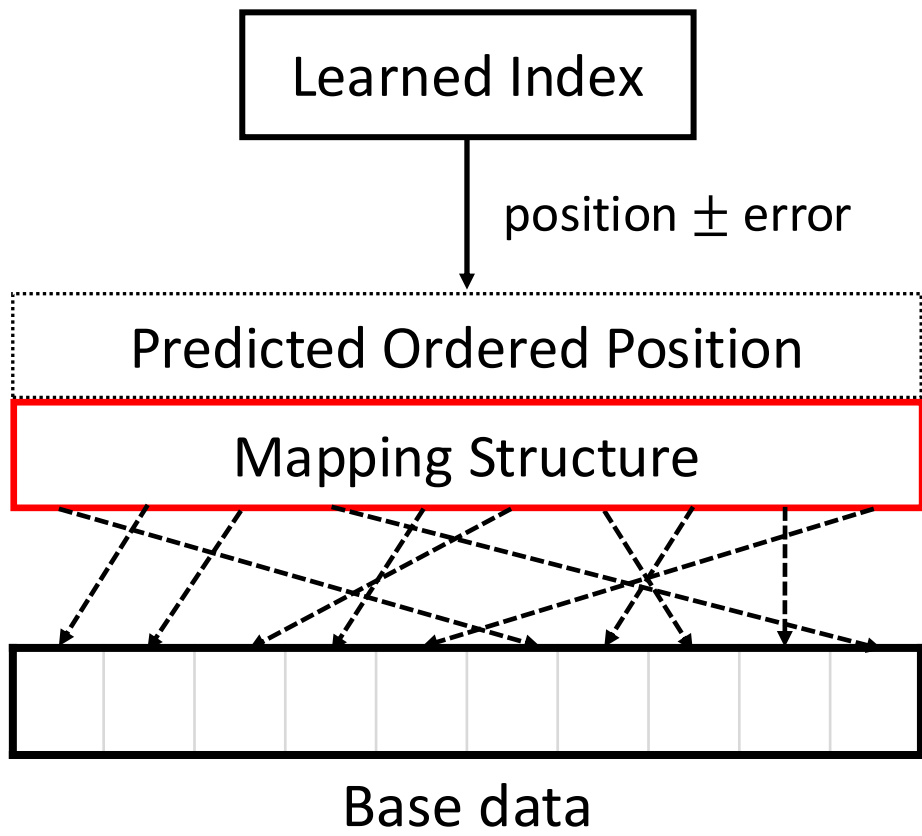
Constellation Maps as Mapping Structures



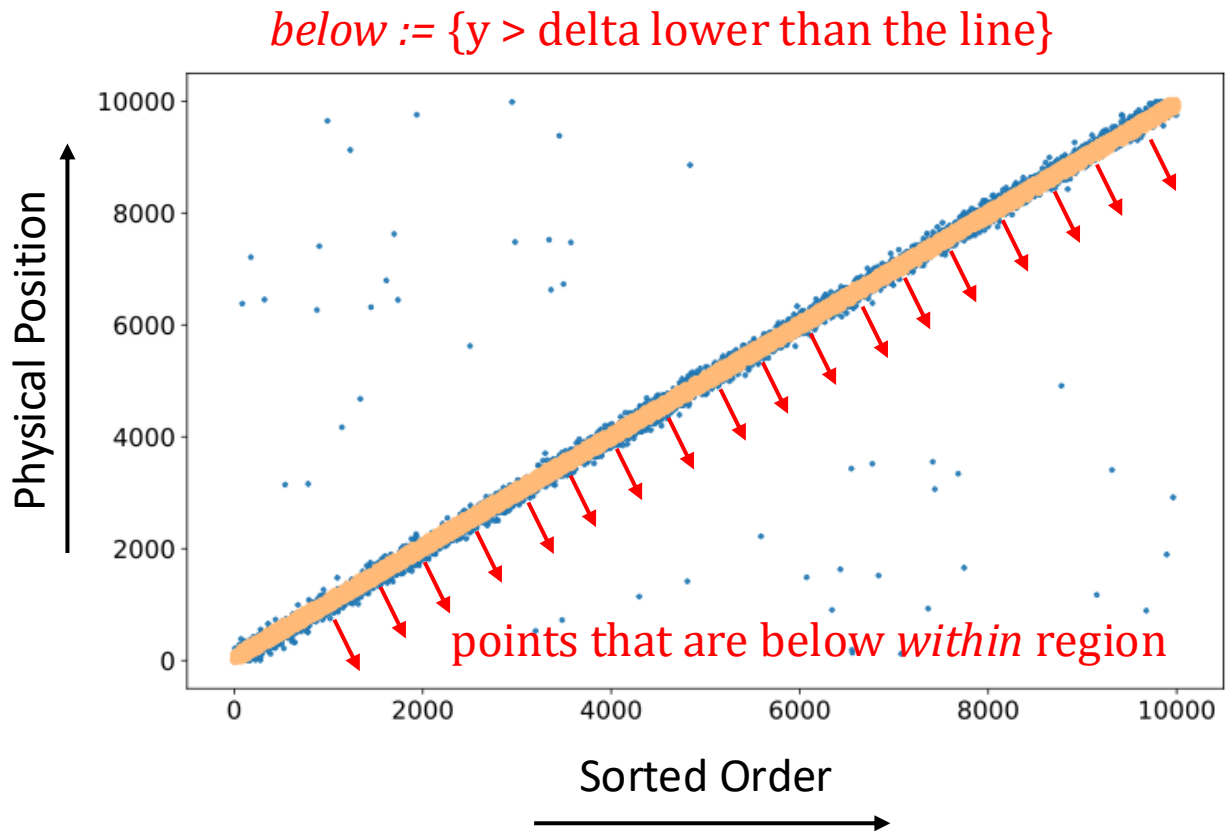
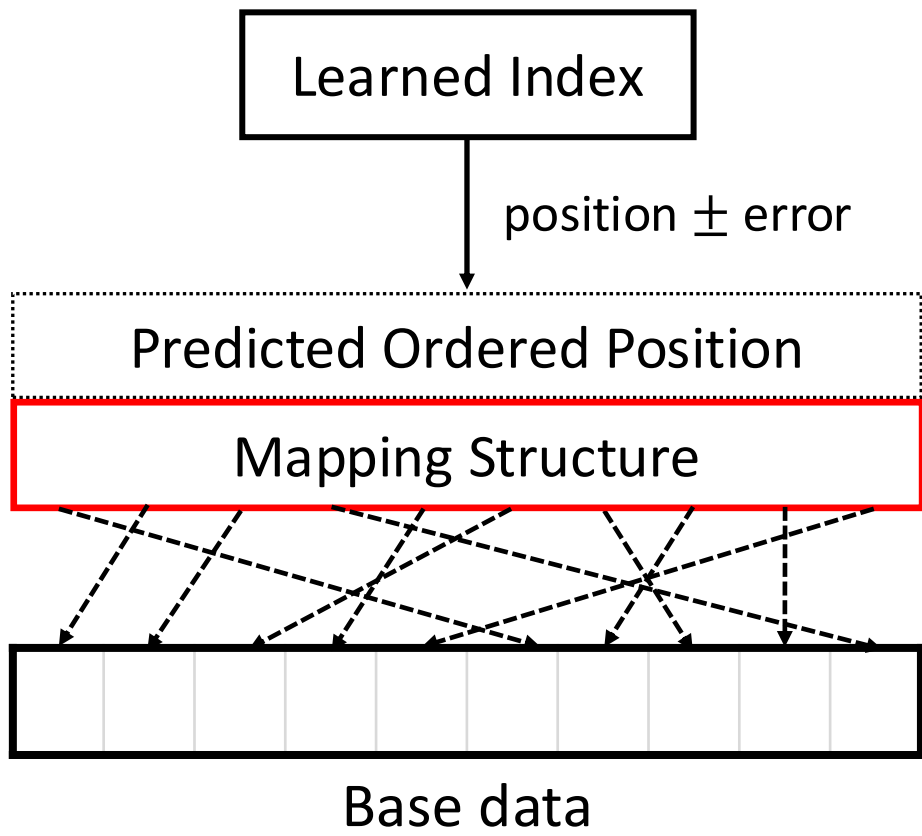
Constellation Maps as Mapping Structures



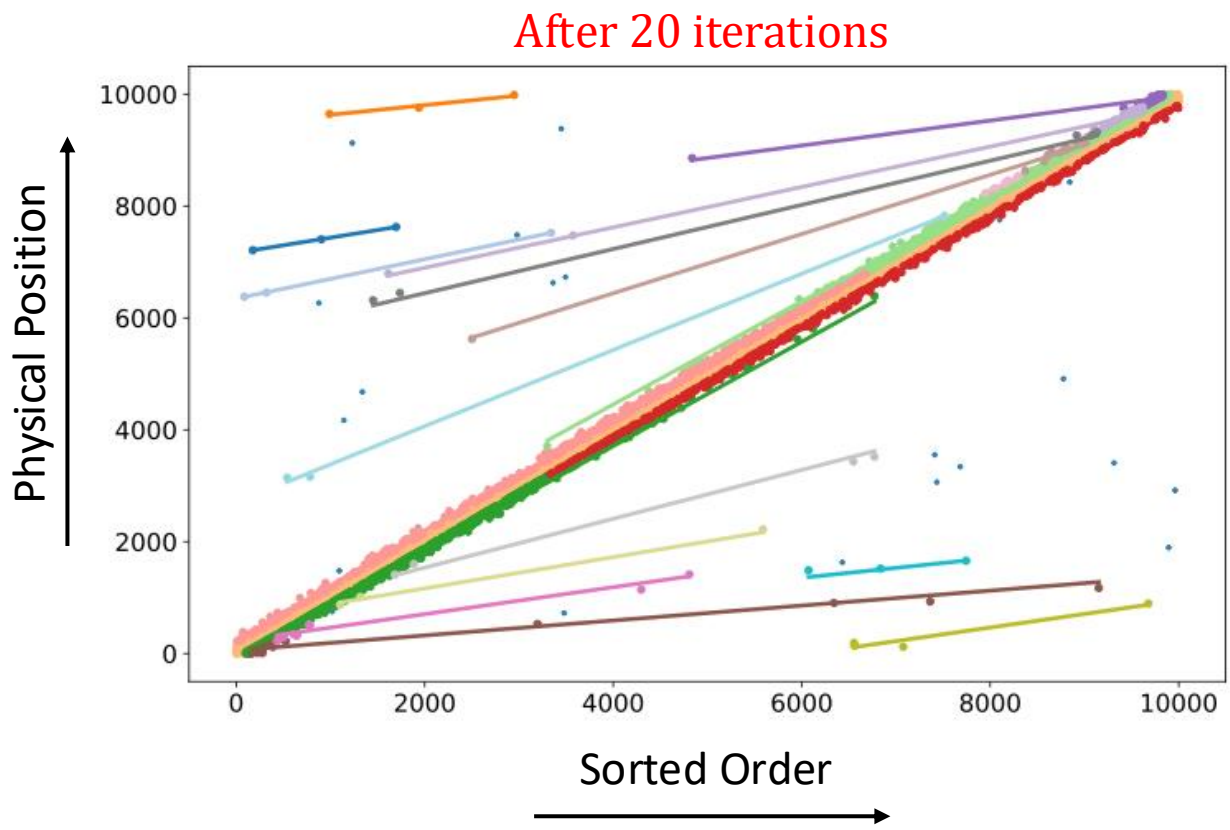
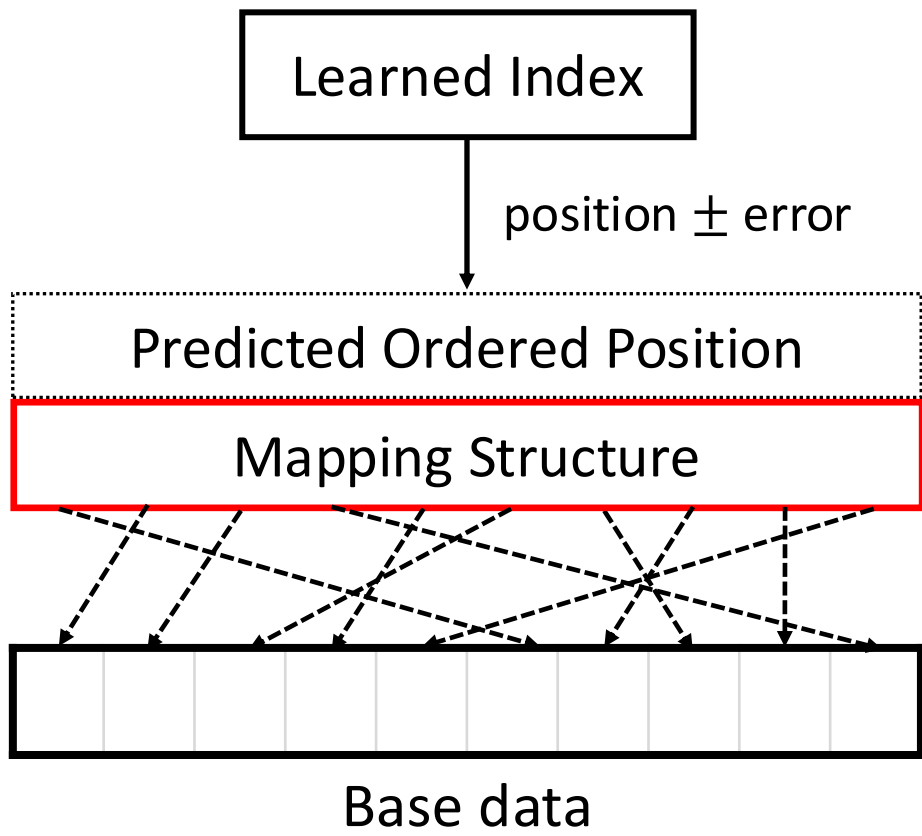
Constellation Maps as Mapping Structures



Constellation Maps as Mapping Structures

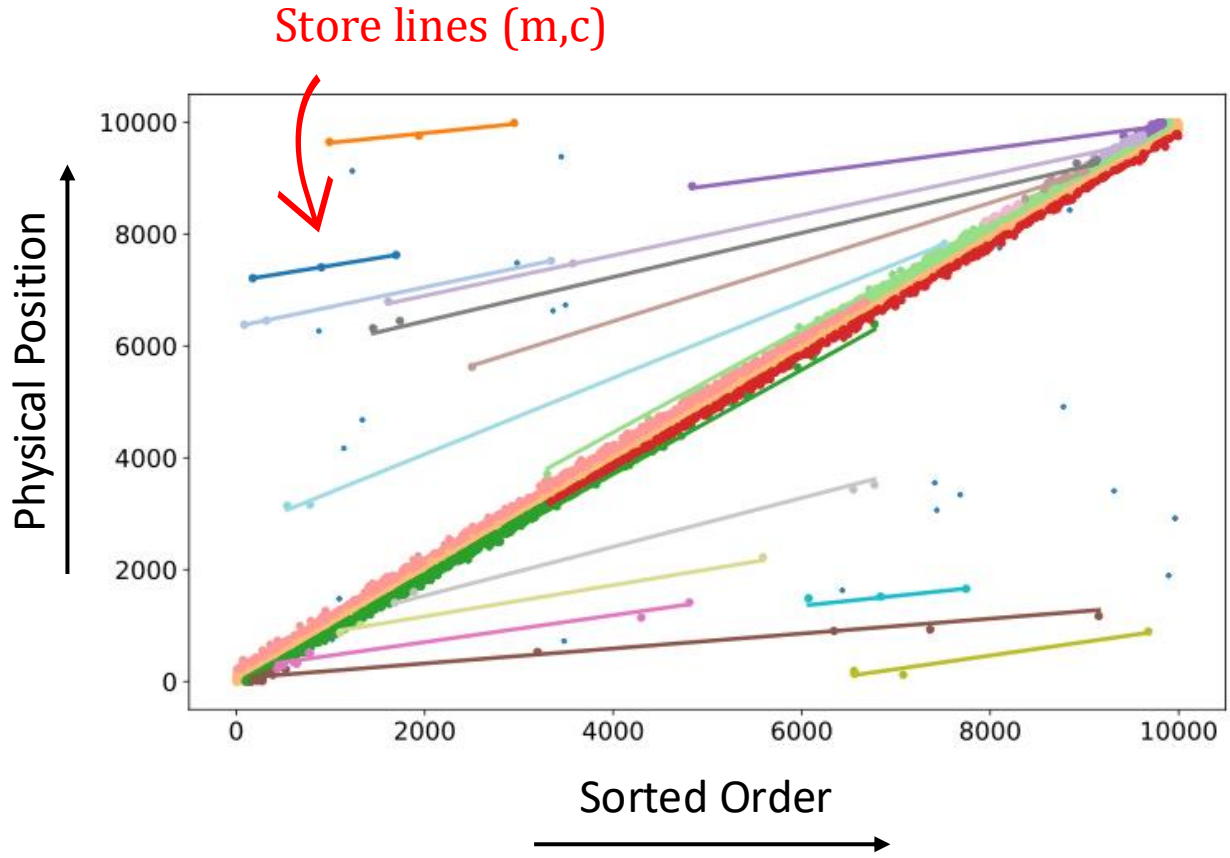
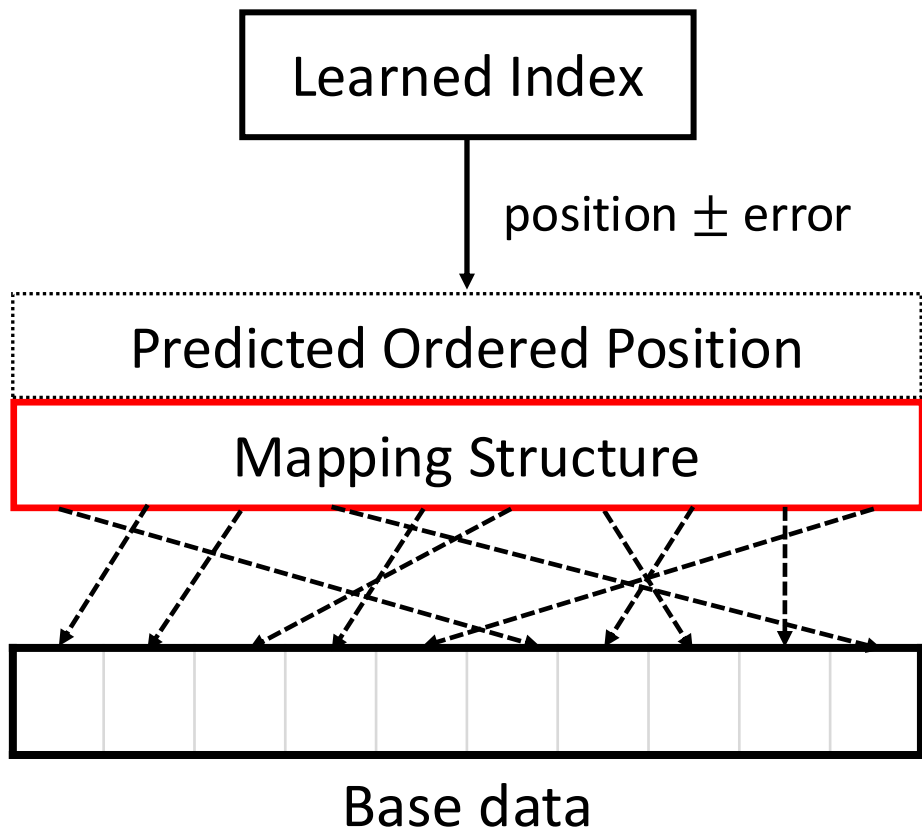


Constellation Maps as Mapping Structures

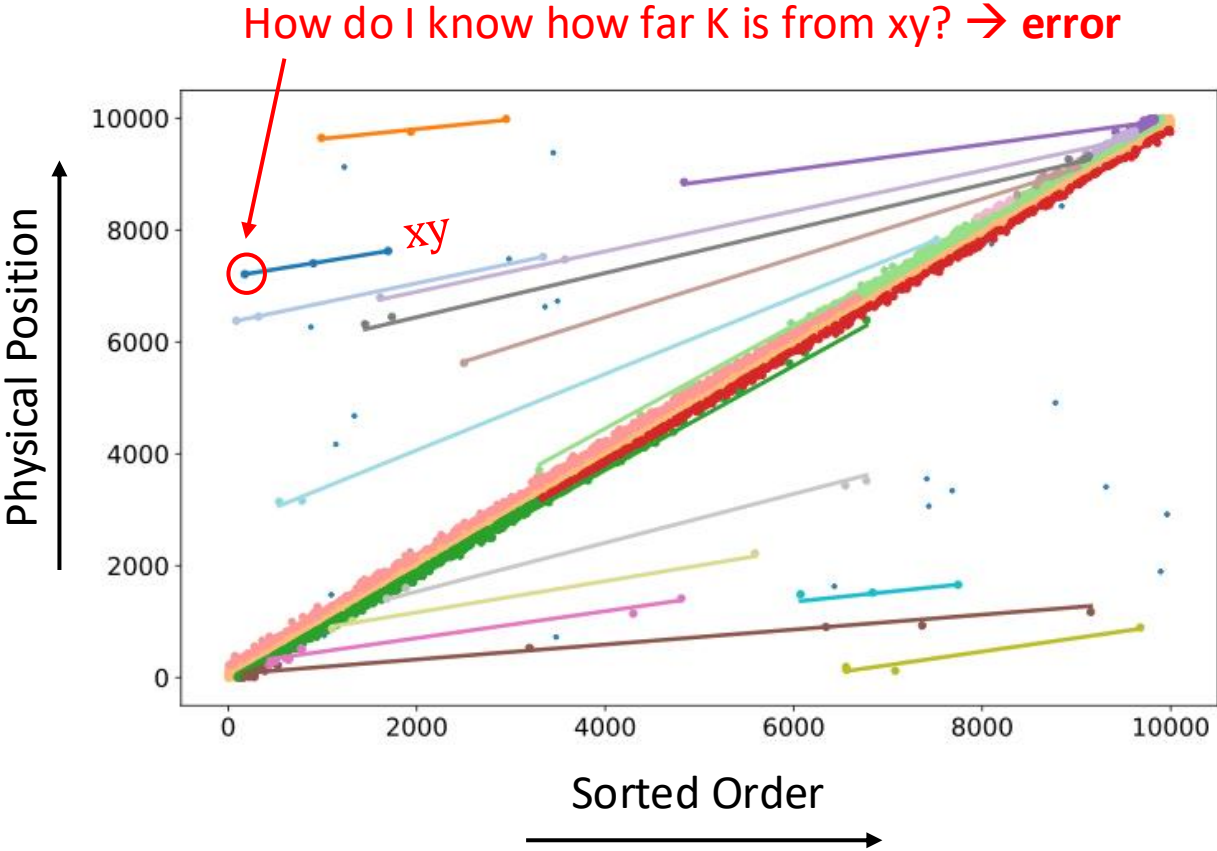
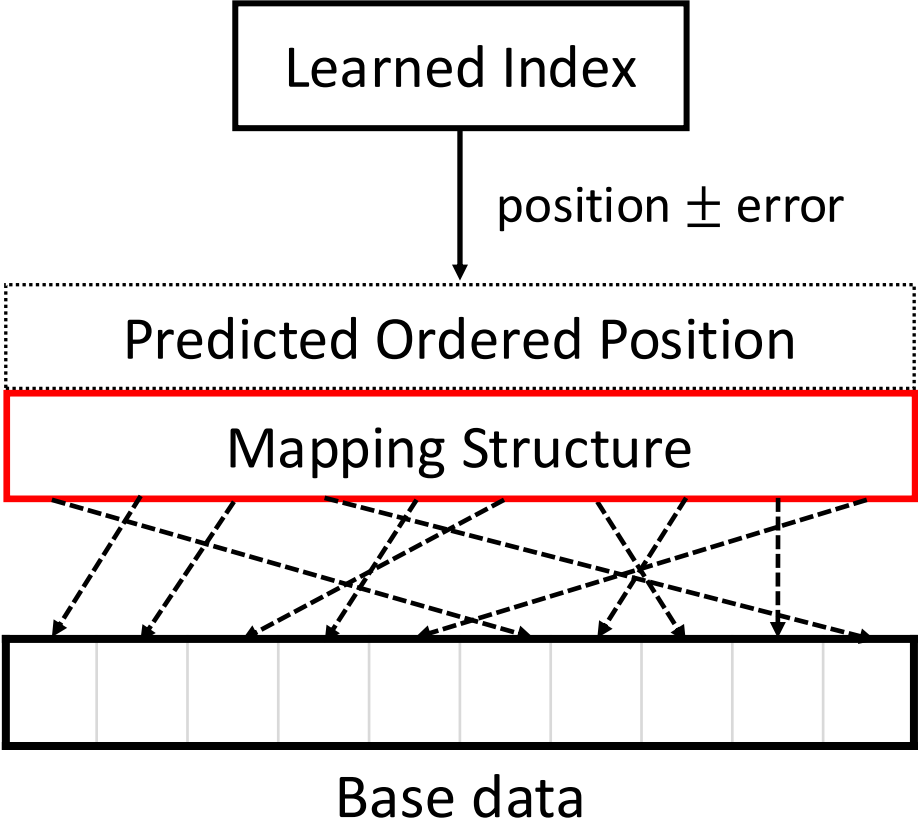


No line can cross any other line

Constellation Maps as Mapping Structures



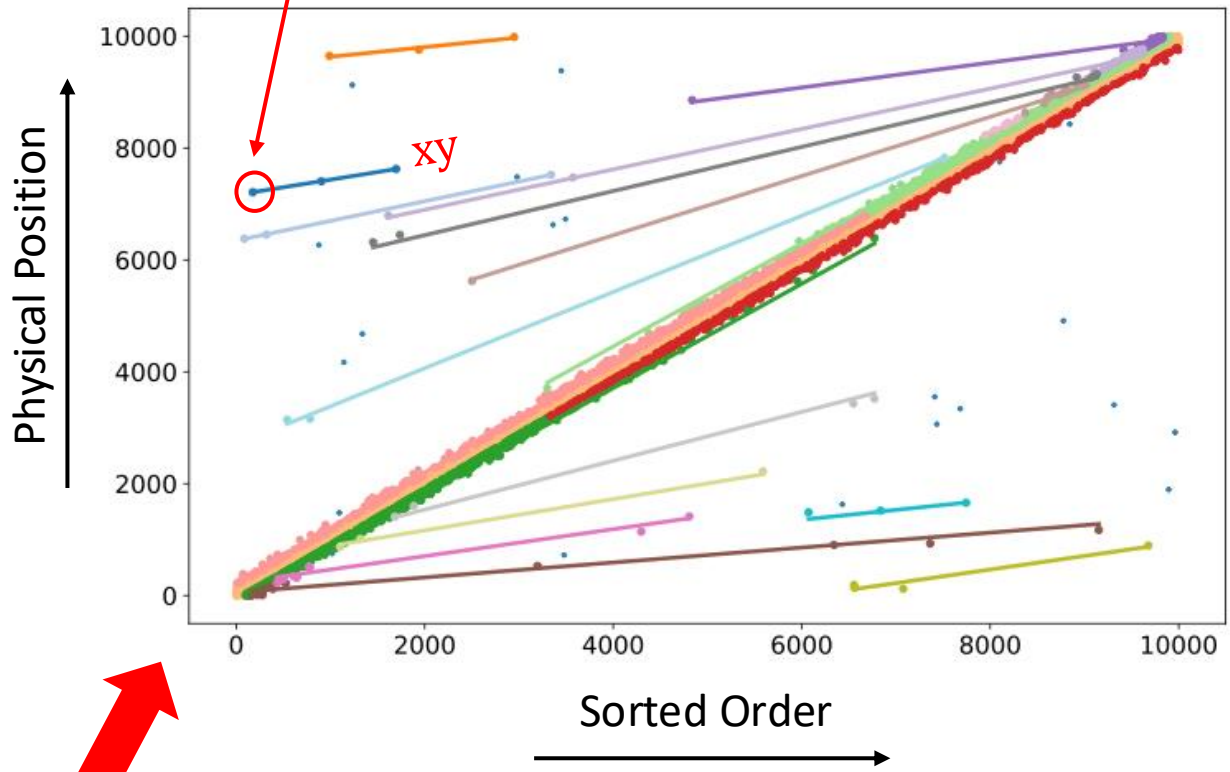
Constellation Maps as Mapping Structures



Constellation Maps as Mapping Structures

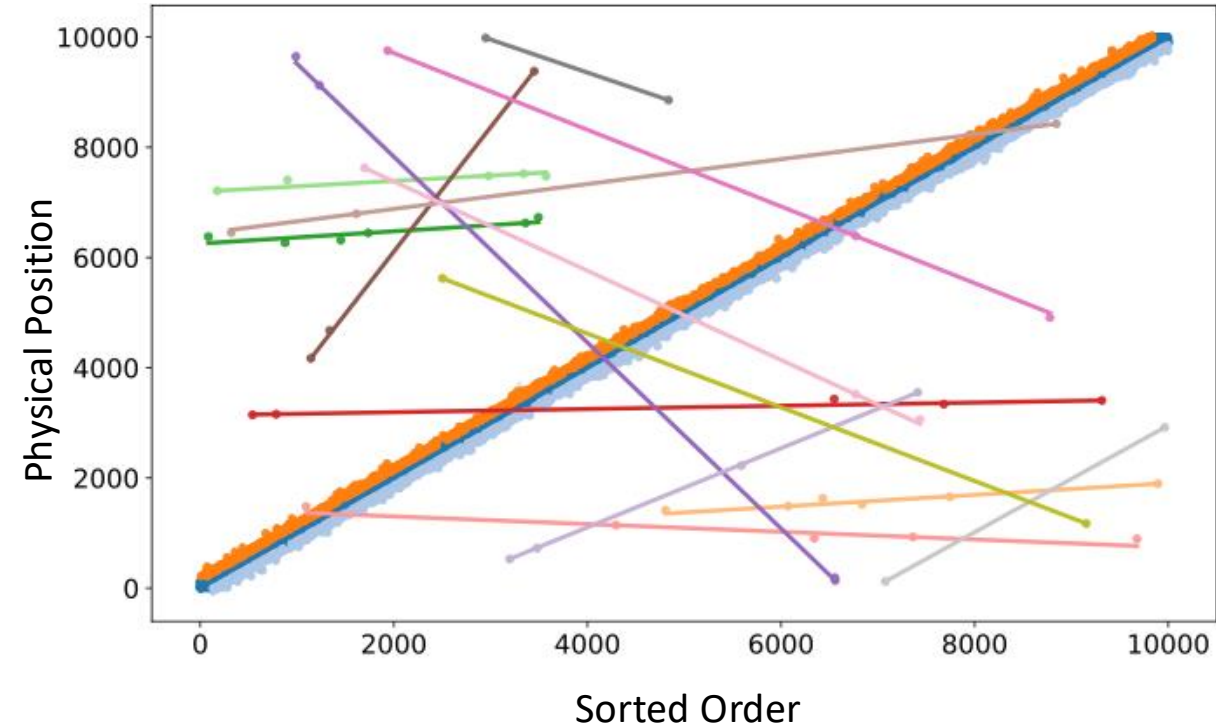


How do I know point is mapped by line xy ? \rightarrow mapping
How do I know how far K is from xy ? \rightarrow error

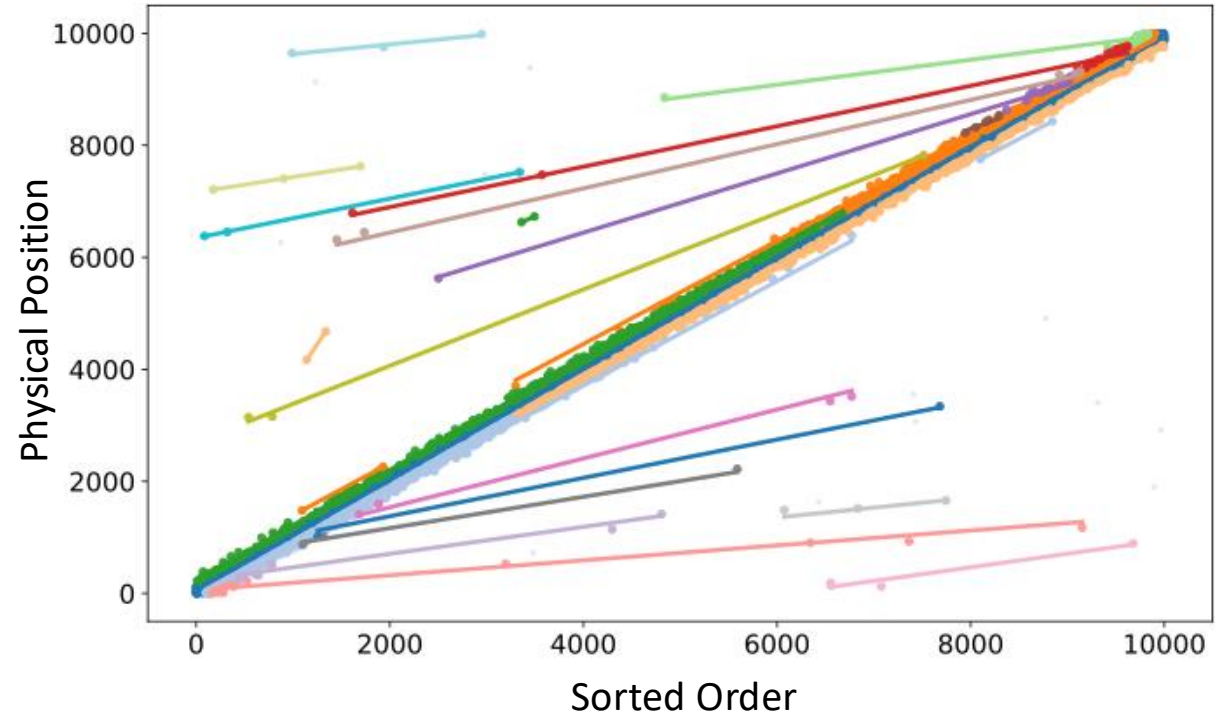


Constellation Maps!!!

Constellation Maps in Action

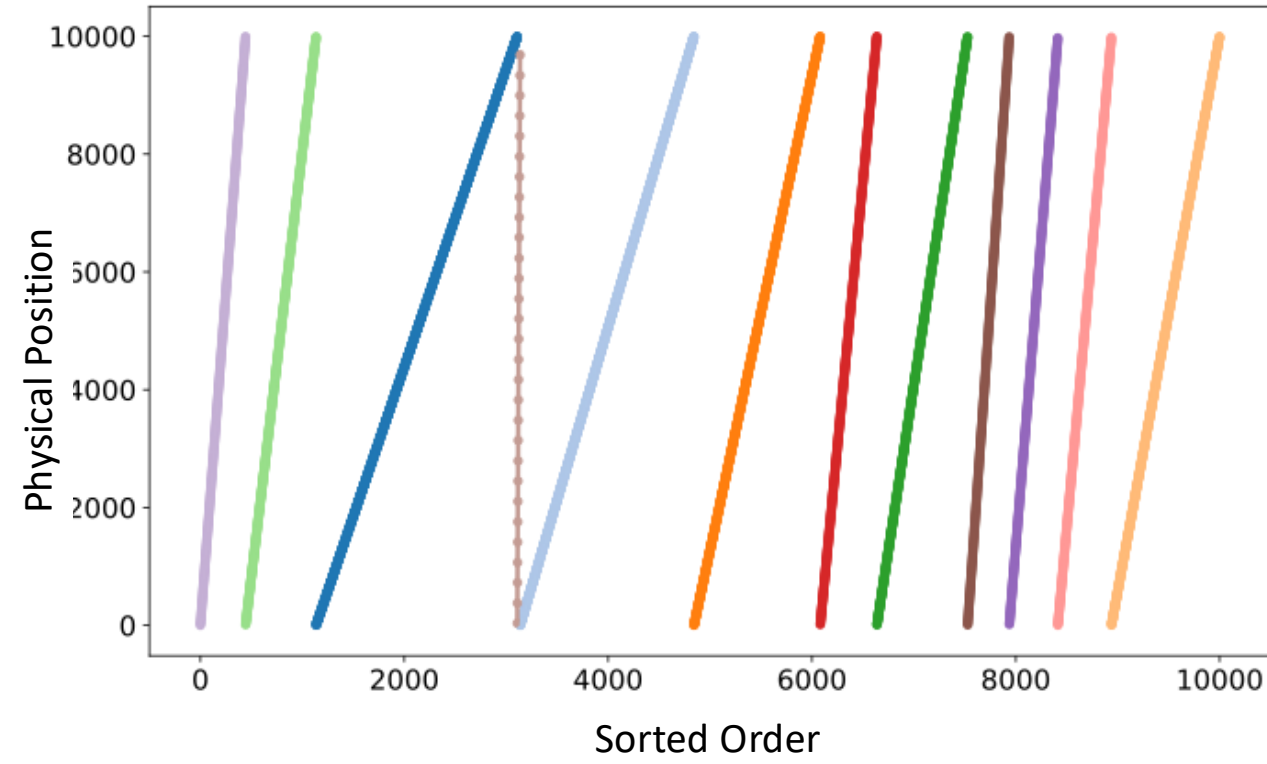


RANSAC = 17 lines

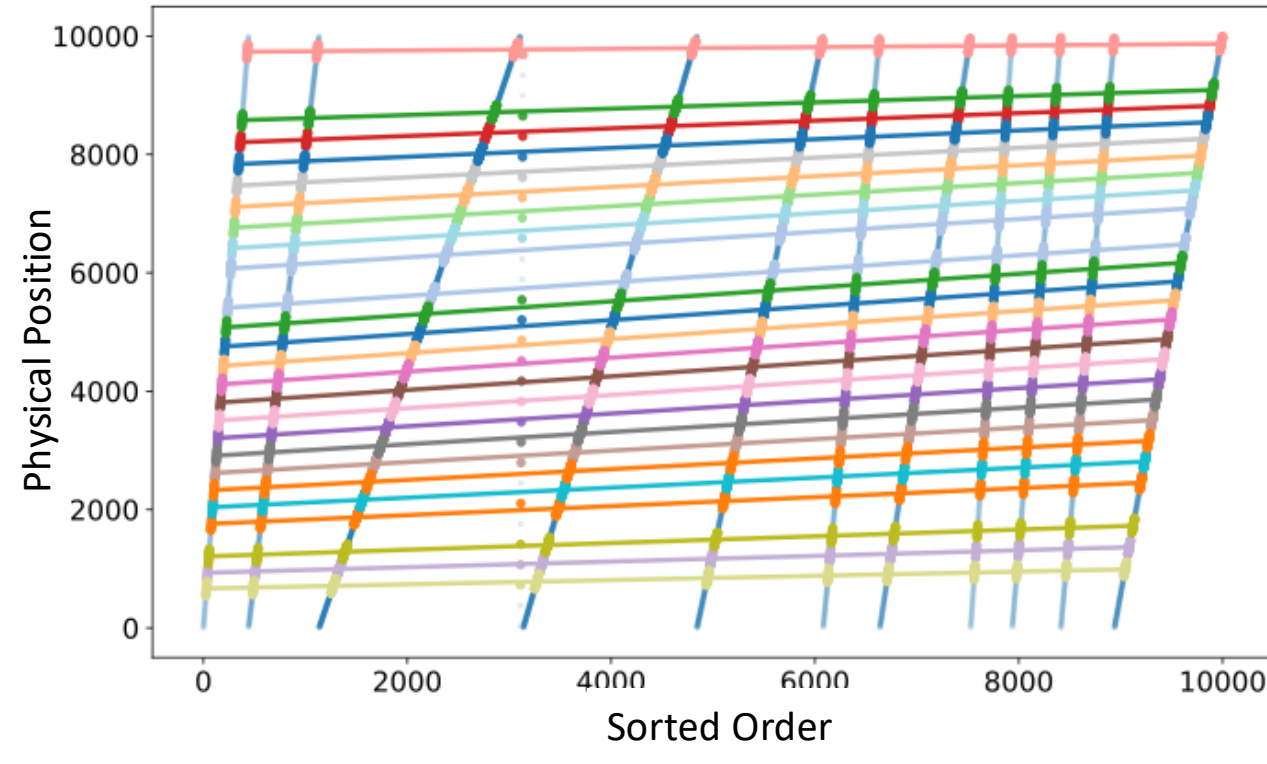


LR = 29 lines

Constellation Maps in Action

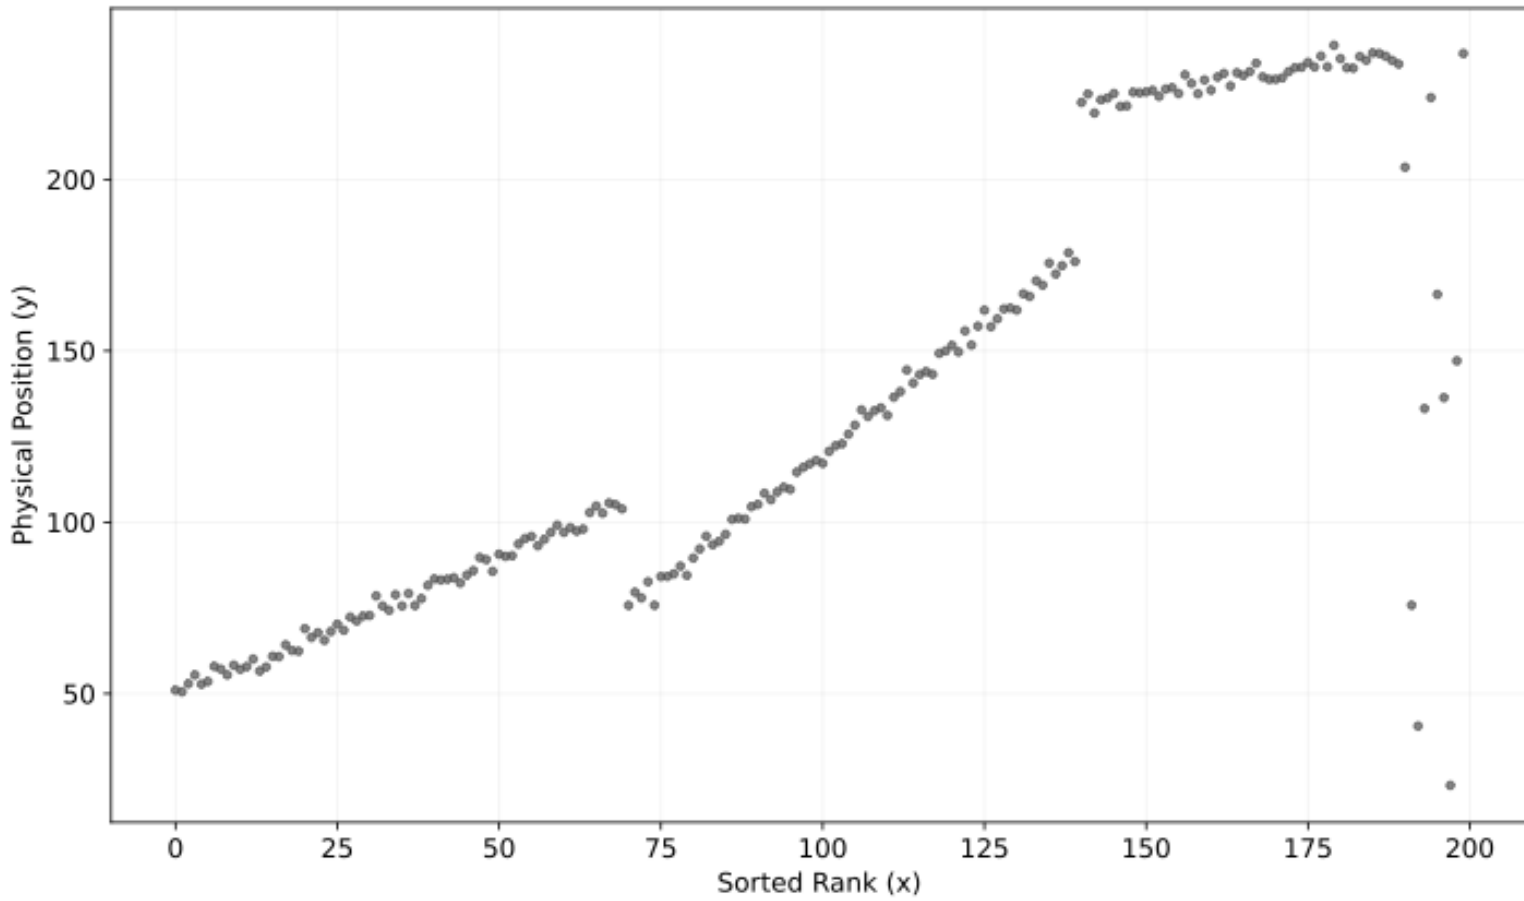


RANSAC = 12 lines

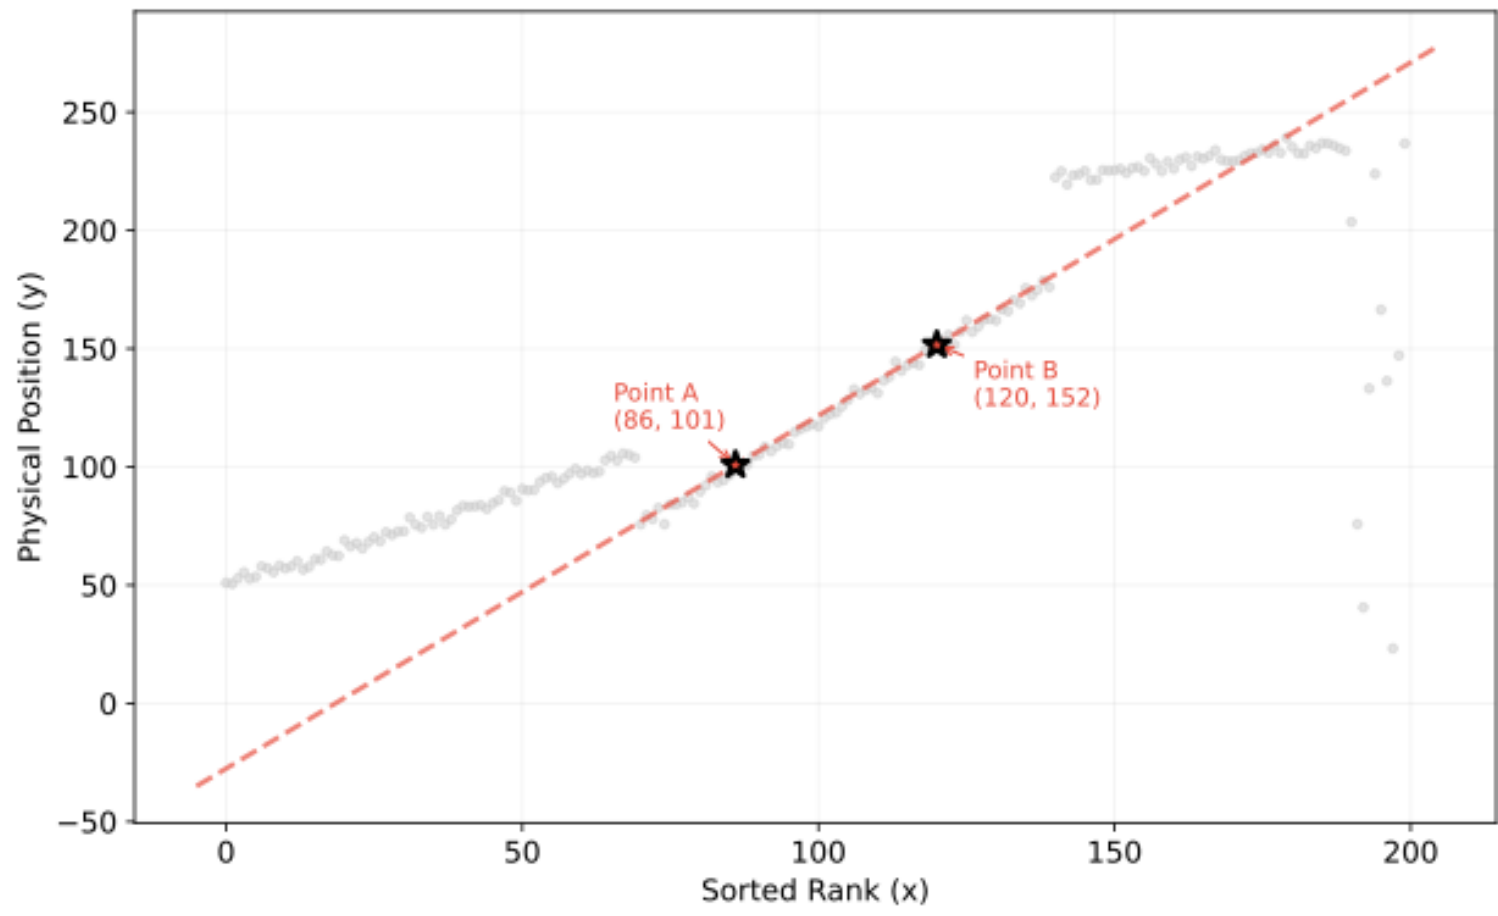


LR = 63 lines

Problems with Default Algorithm

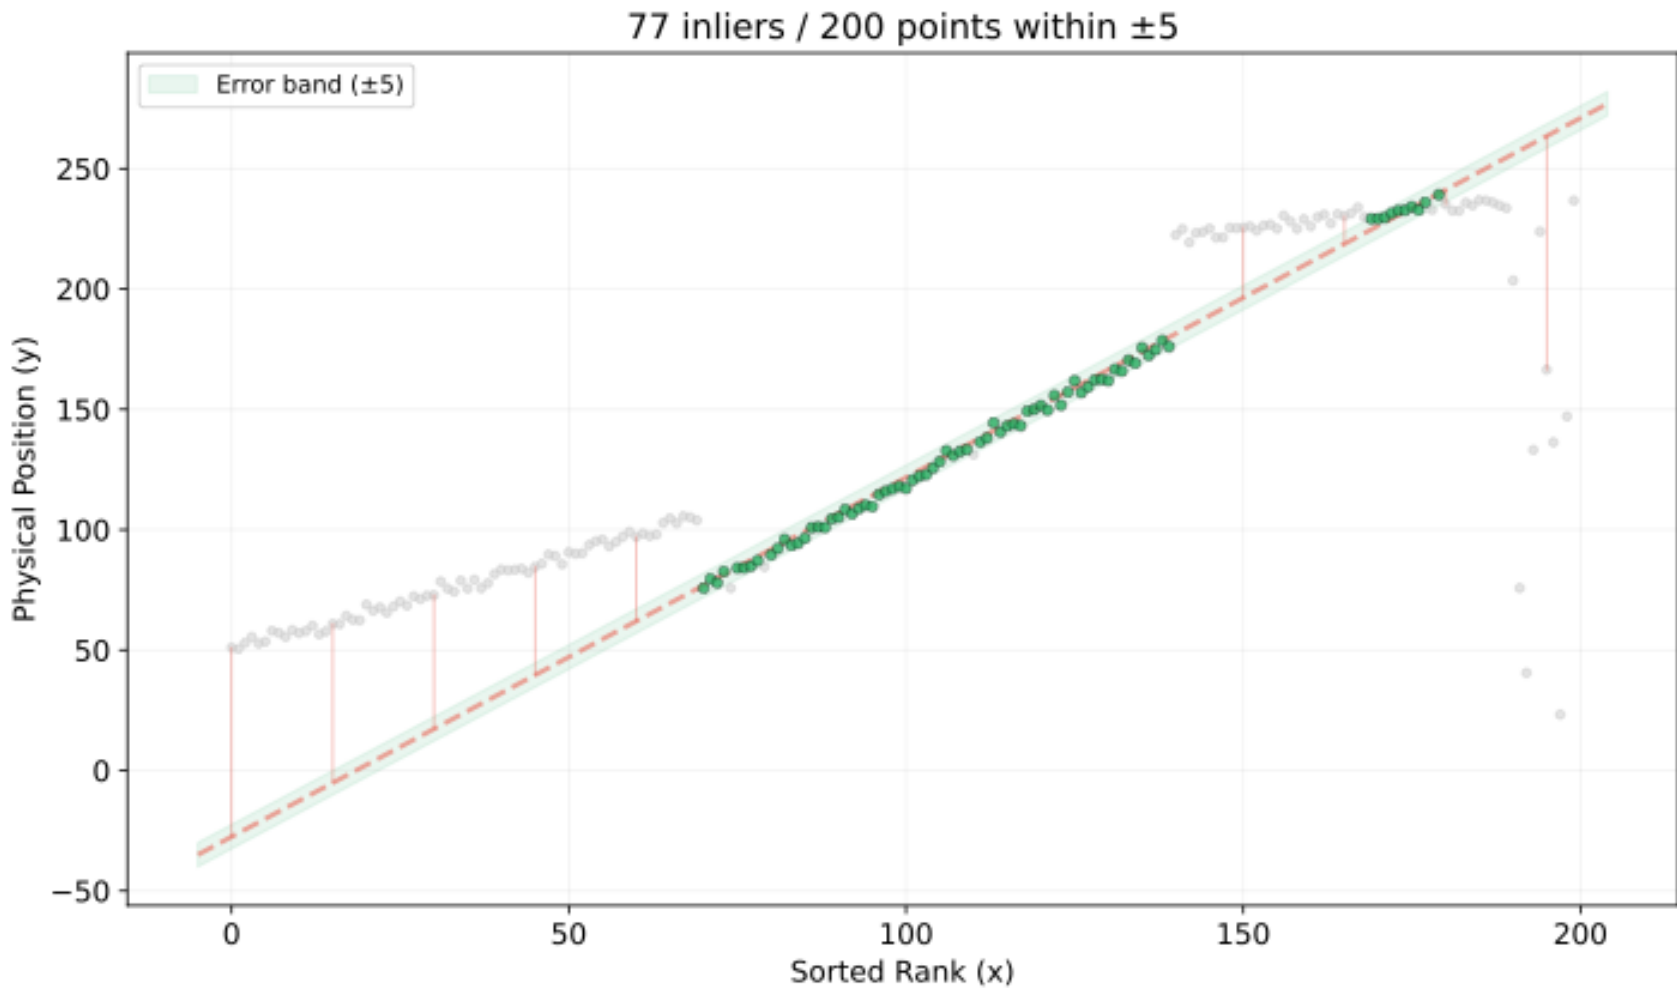


Problems with Default Algorithm

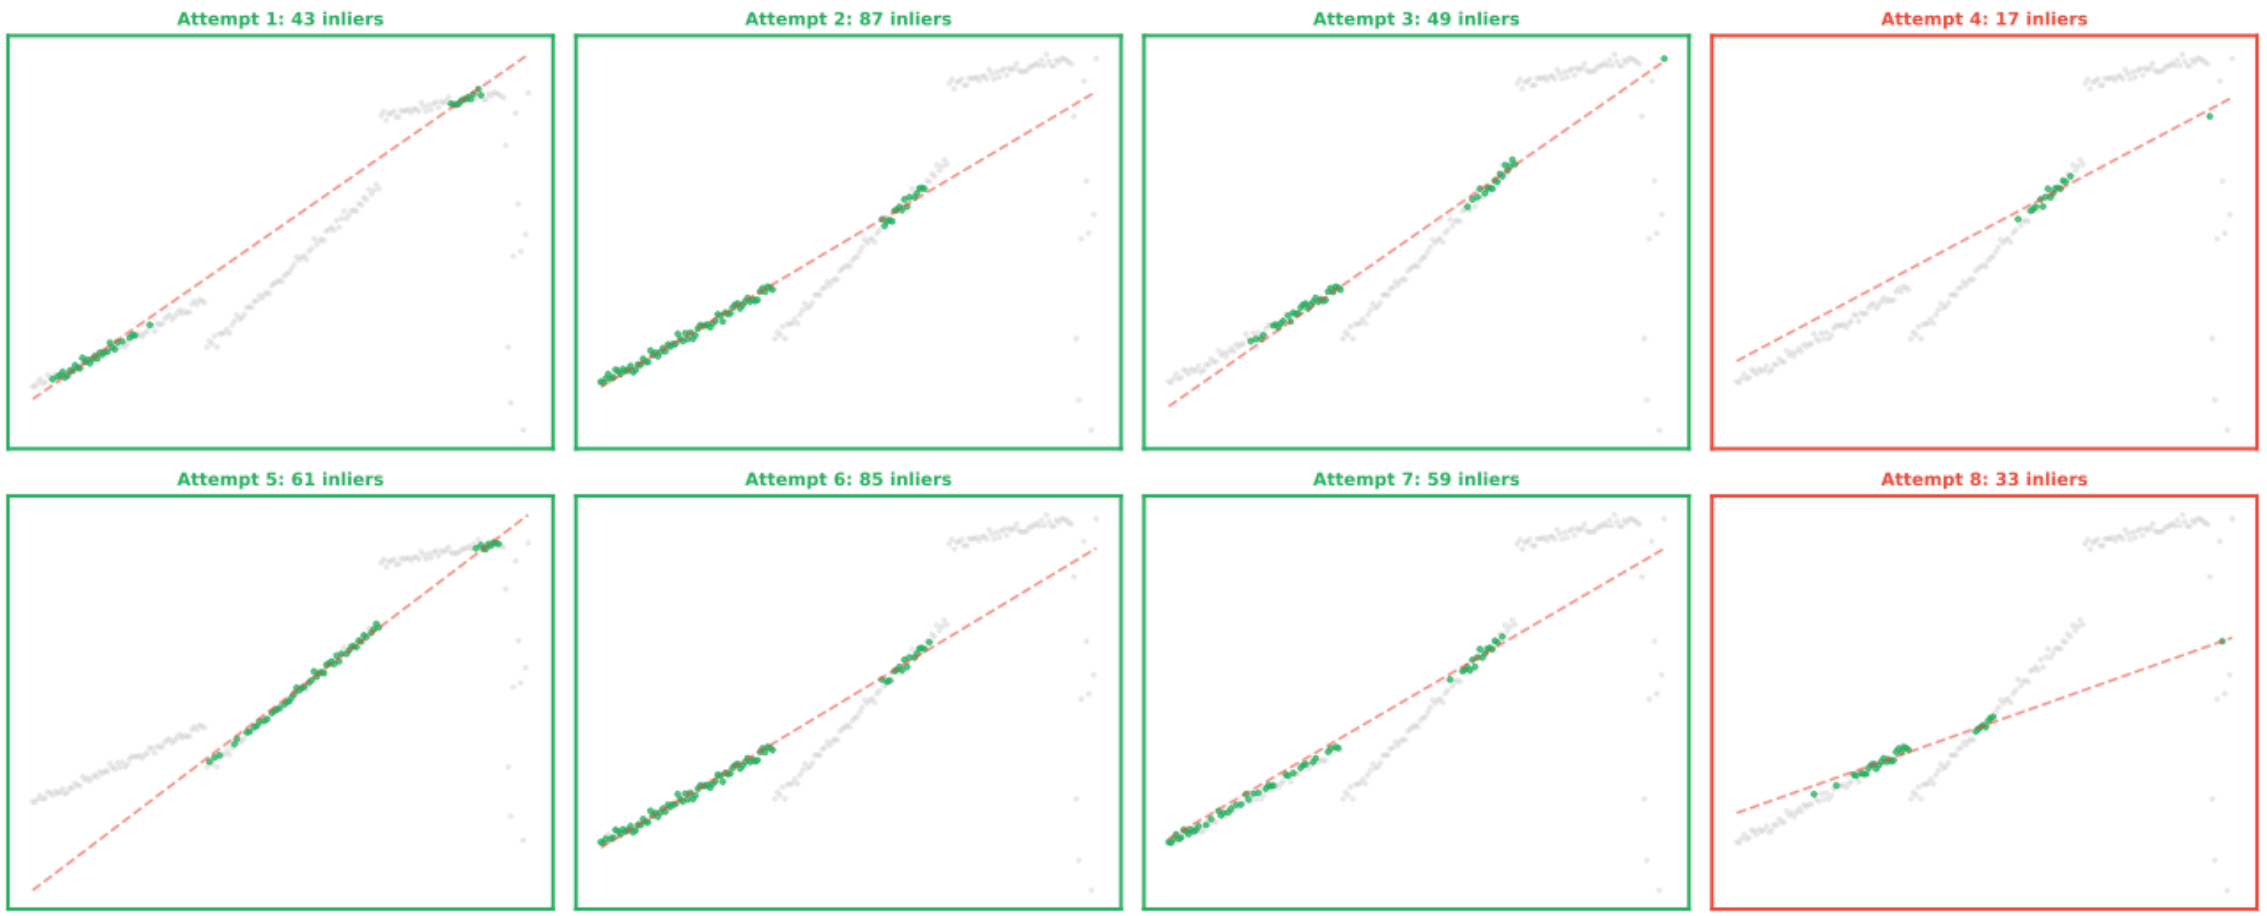


2 random points selected to form a line

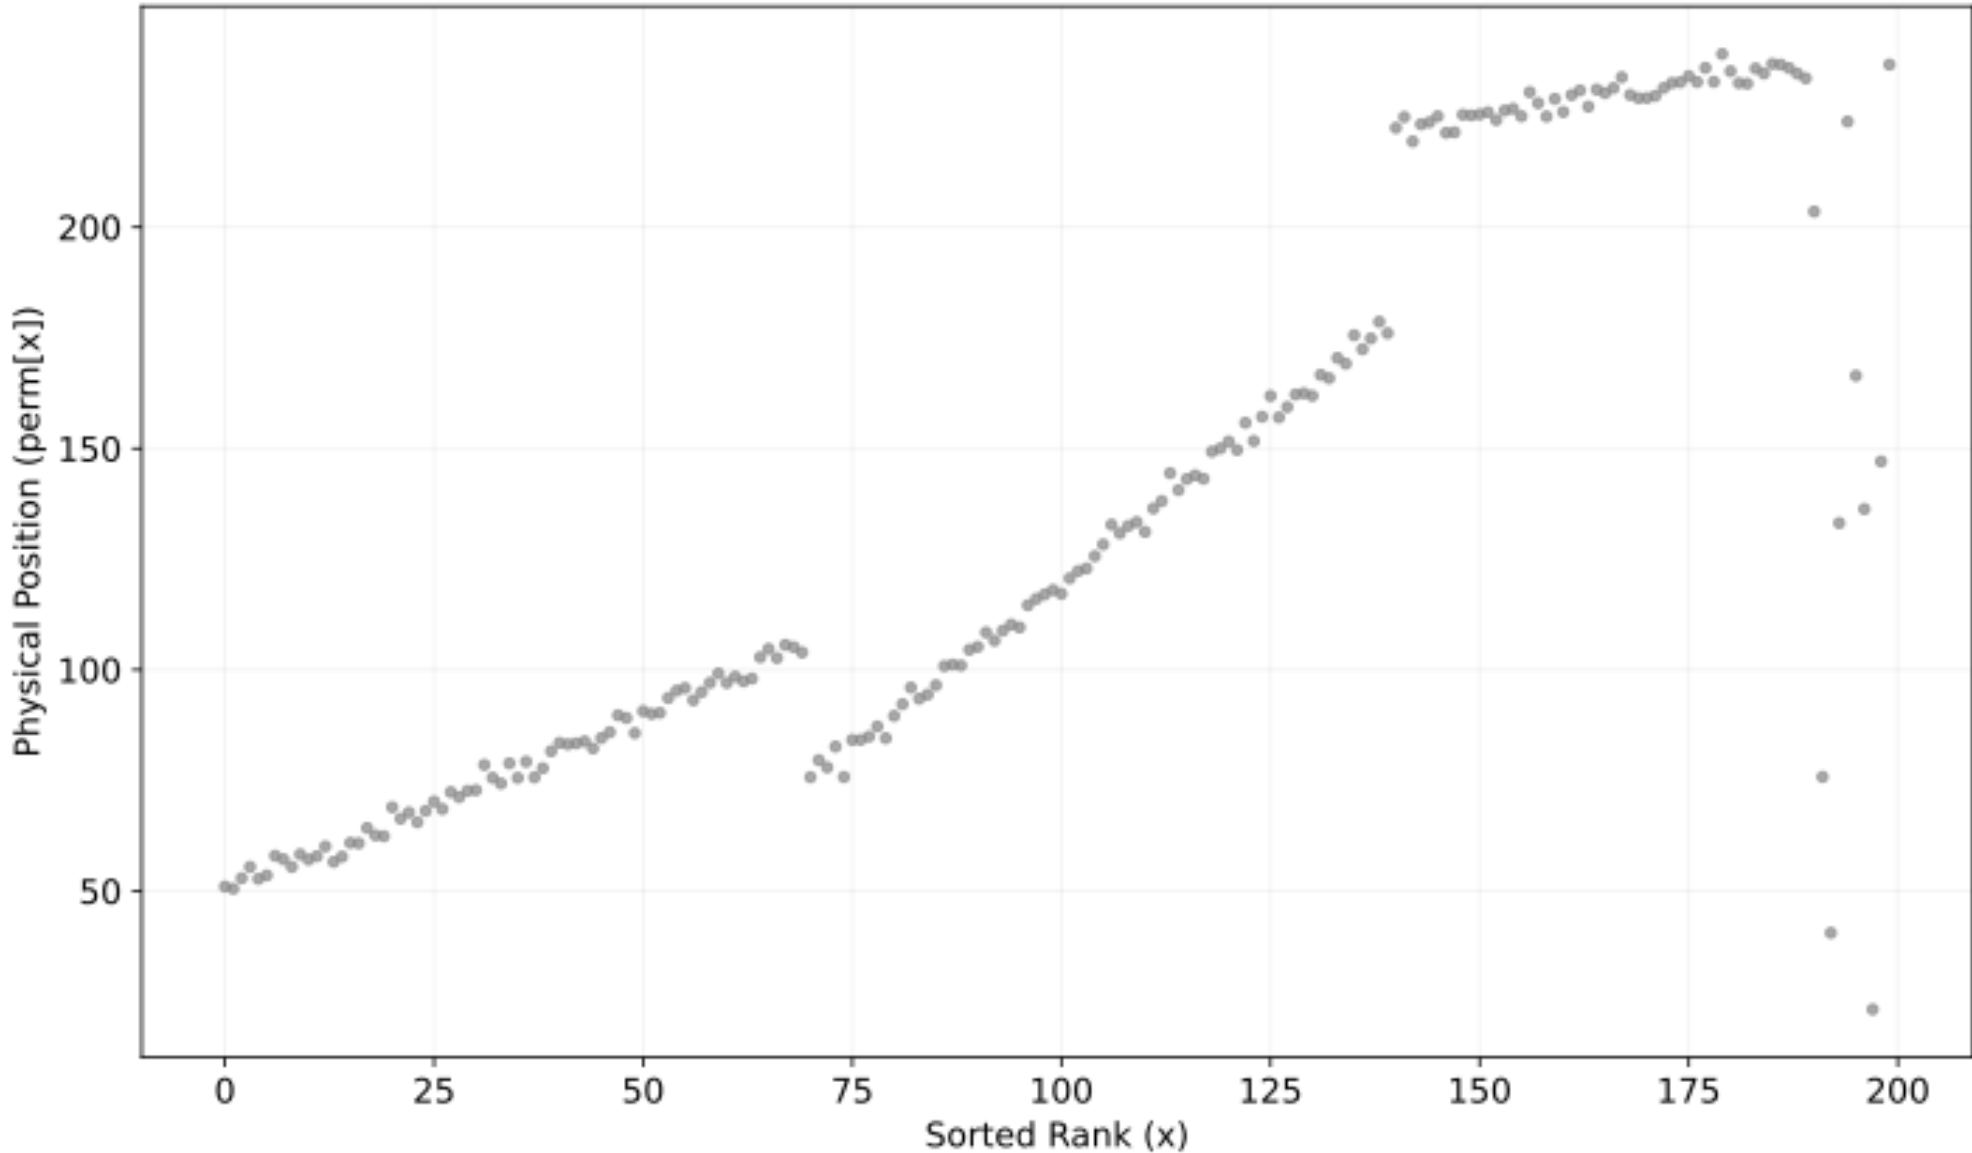
Problems with Default Algorithm



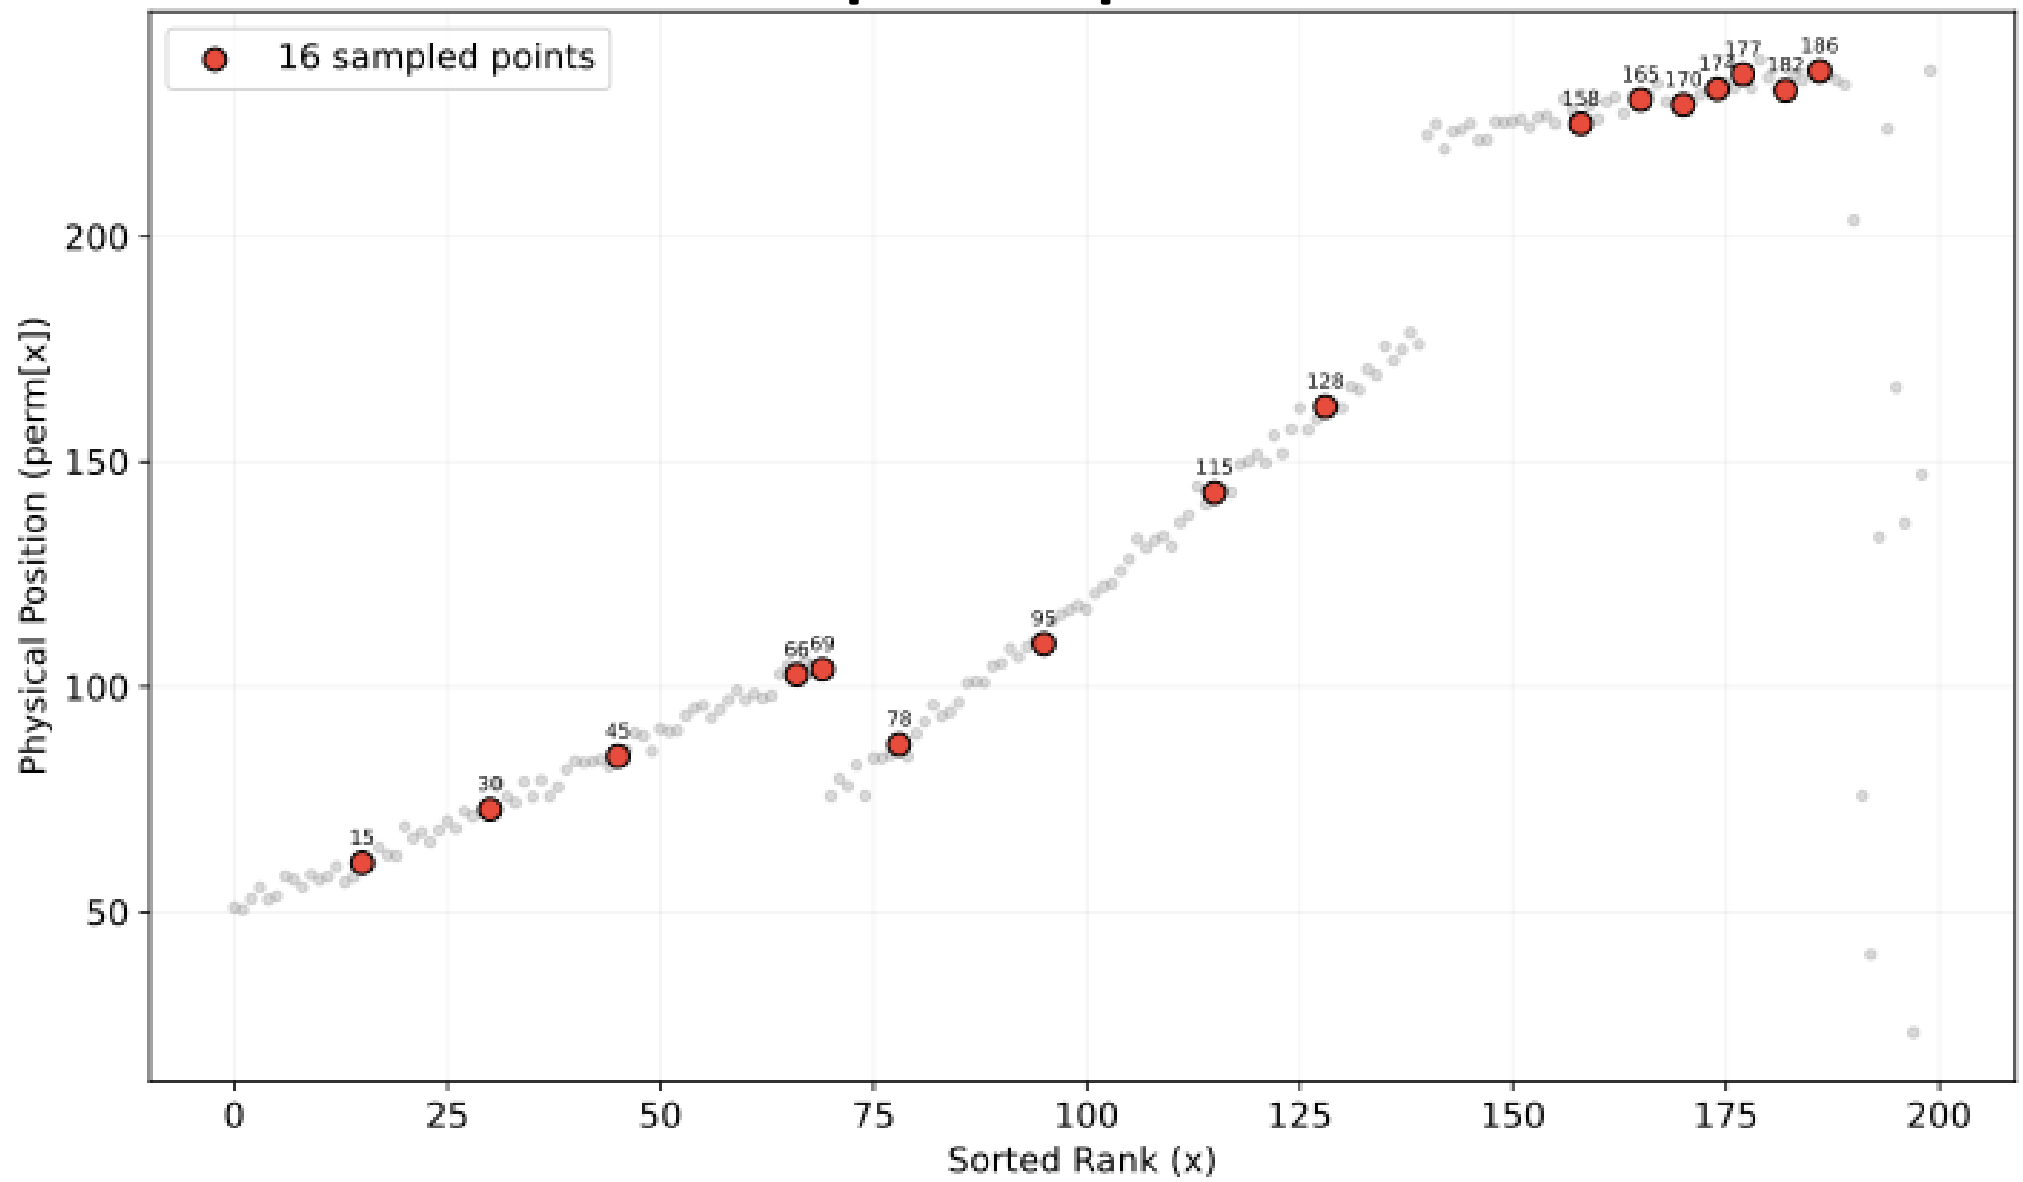
Problems with Default Algorithm



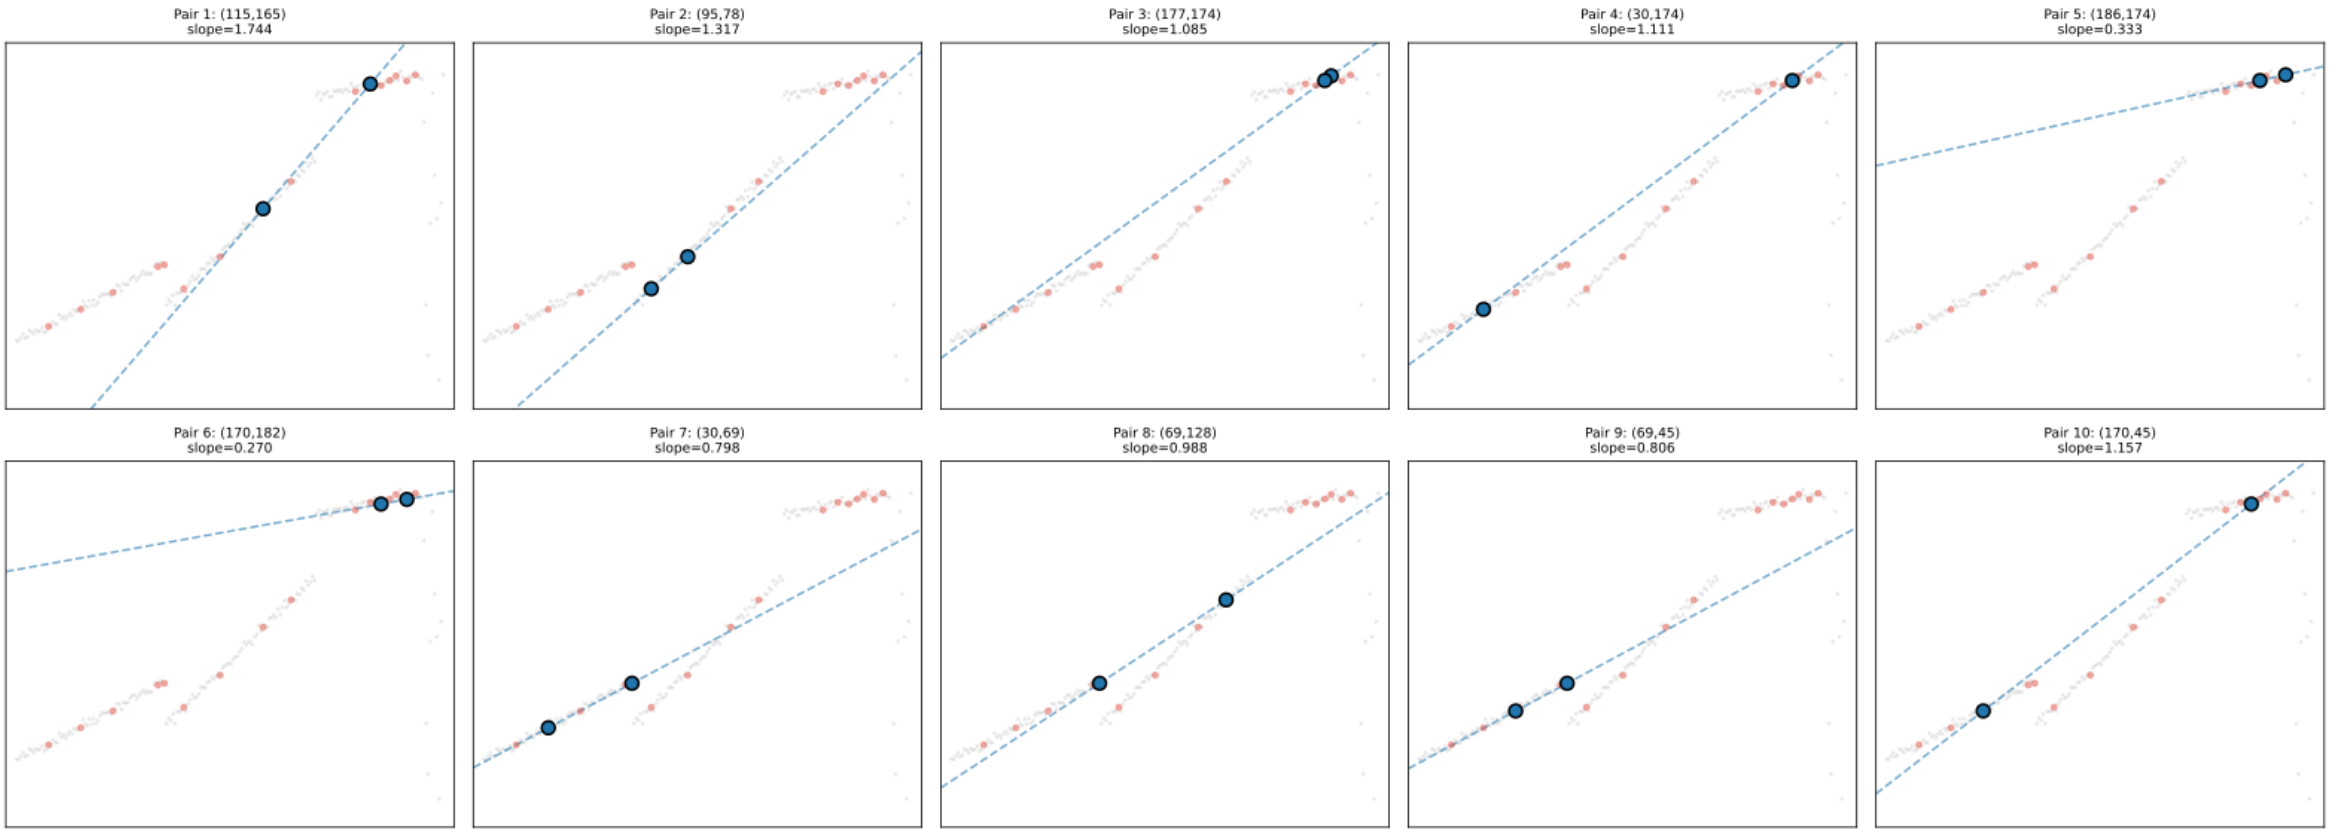
Building a Better Algorithm



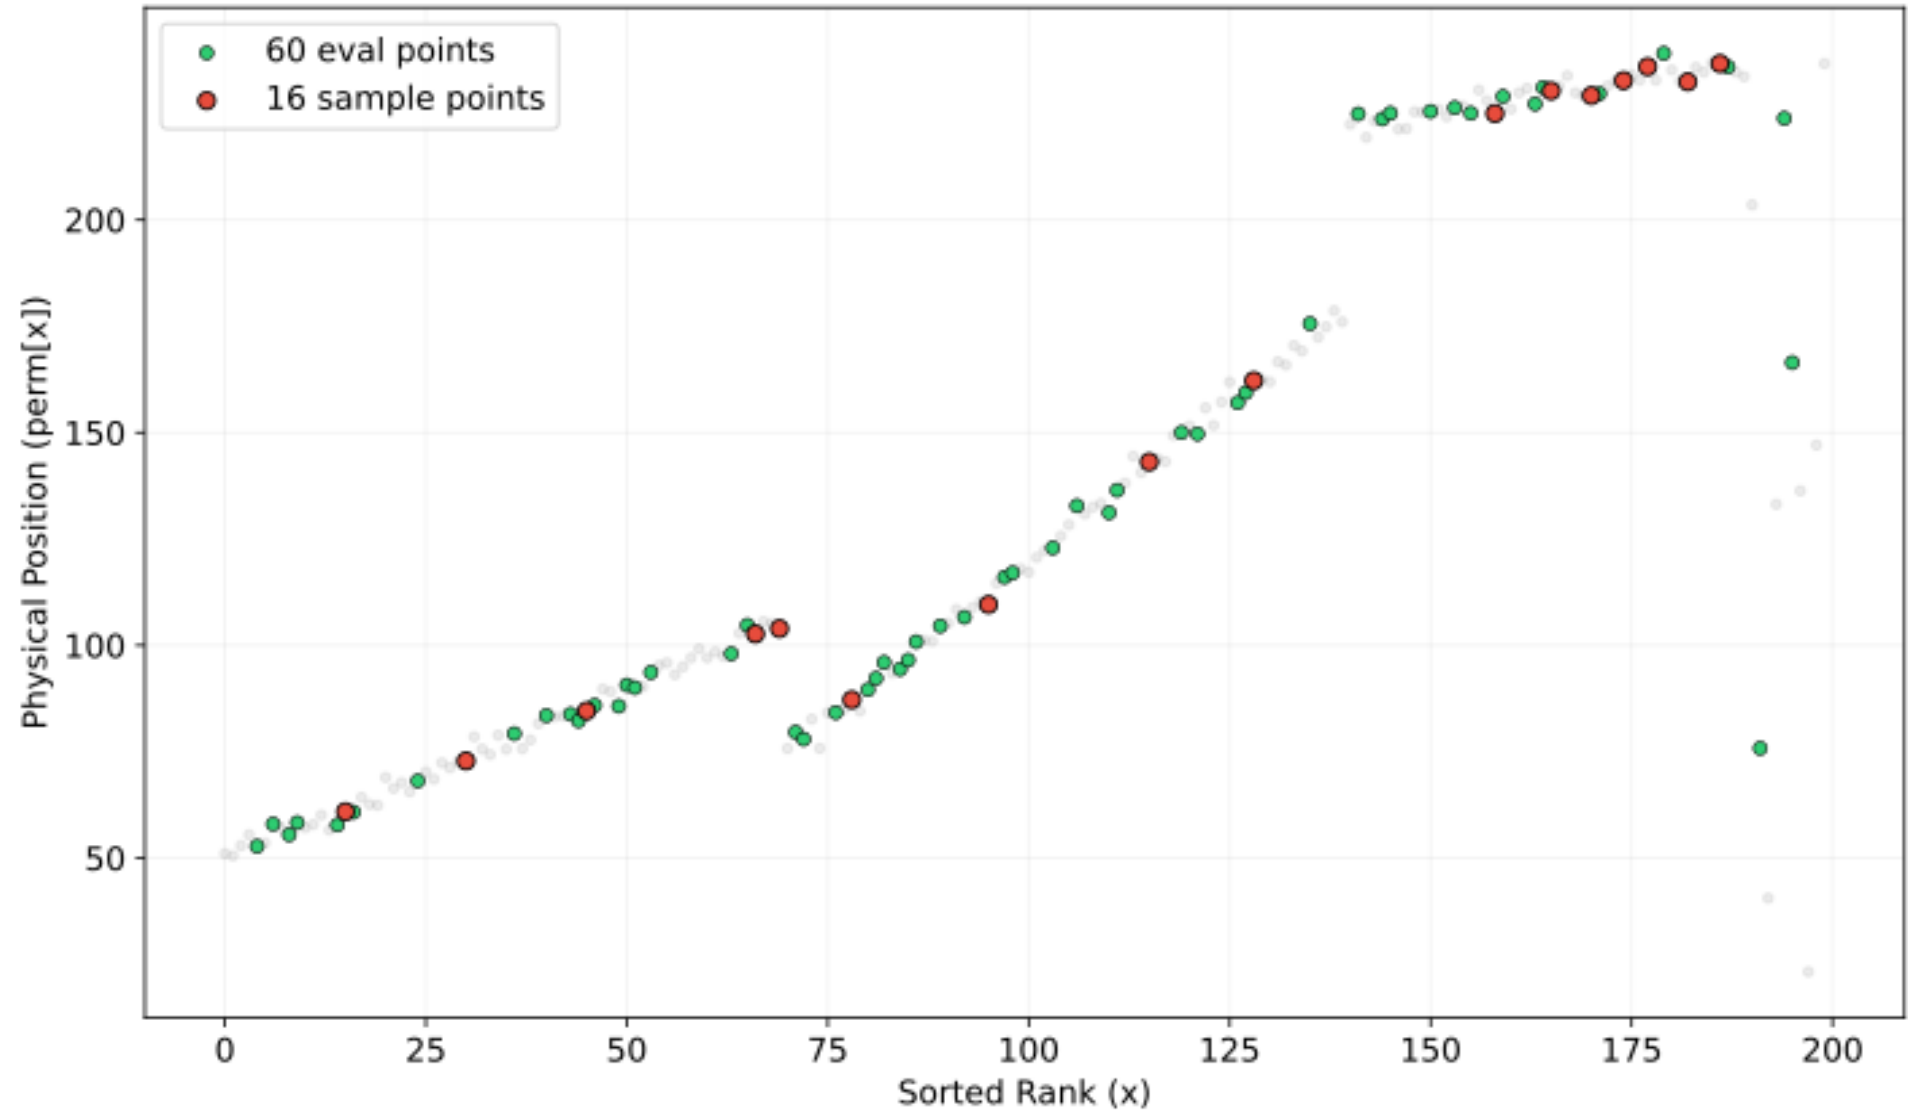
Sample x points



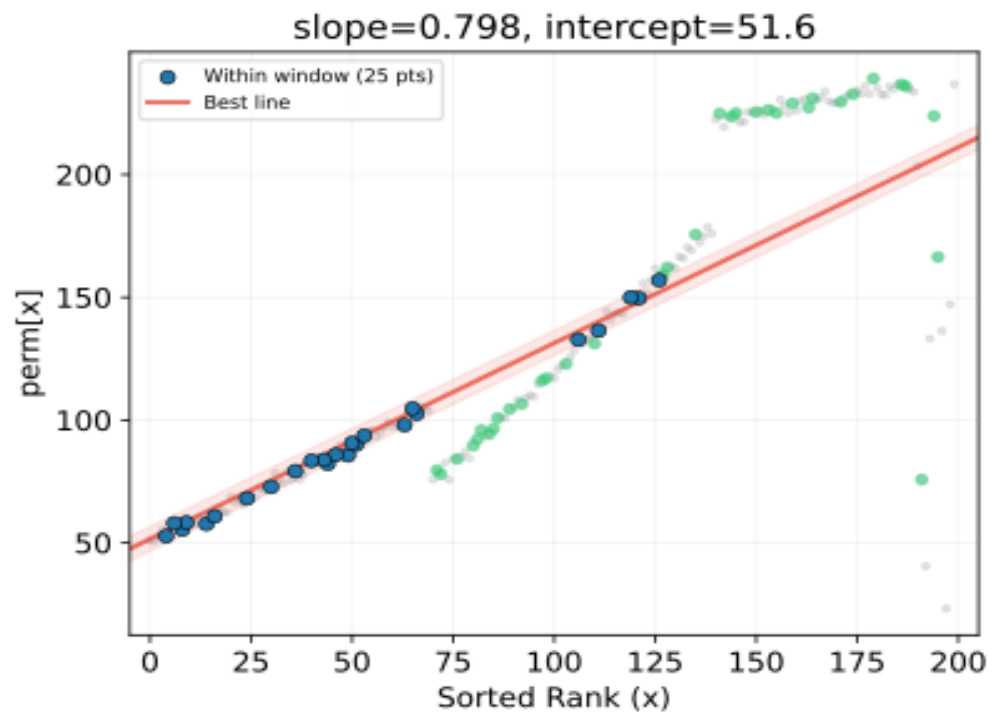
Compute slope for x points



Sample eval points



Find intercepts

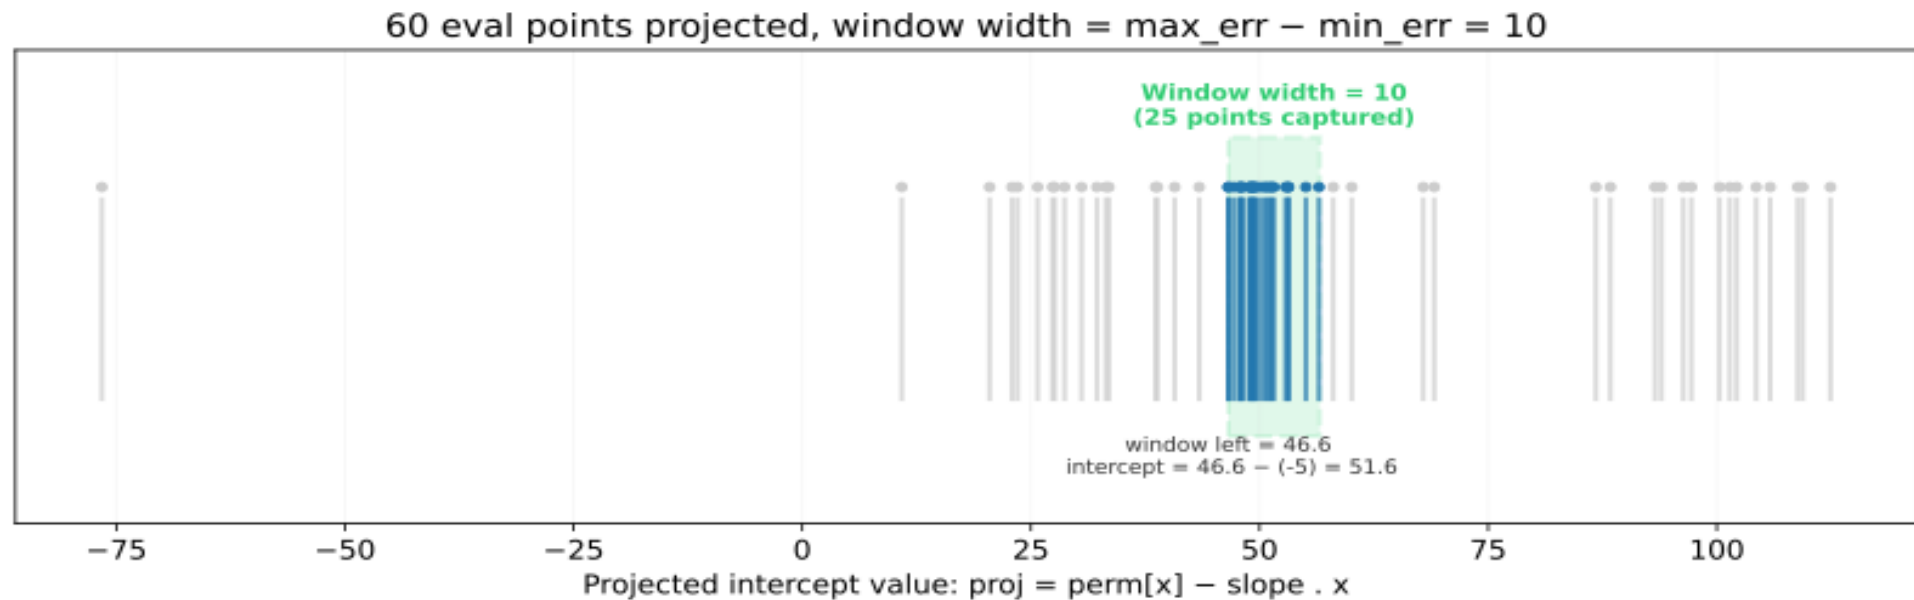


Find c for each eval point:
 $c = y - m.x$

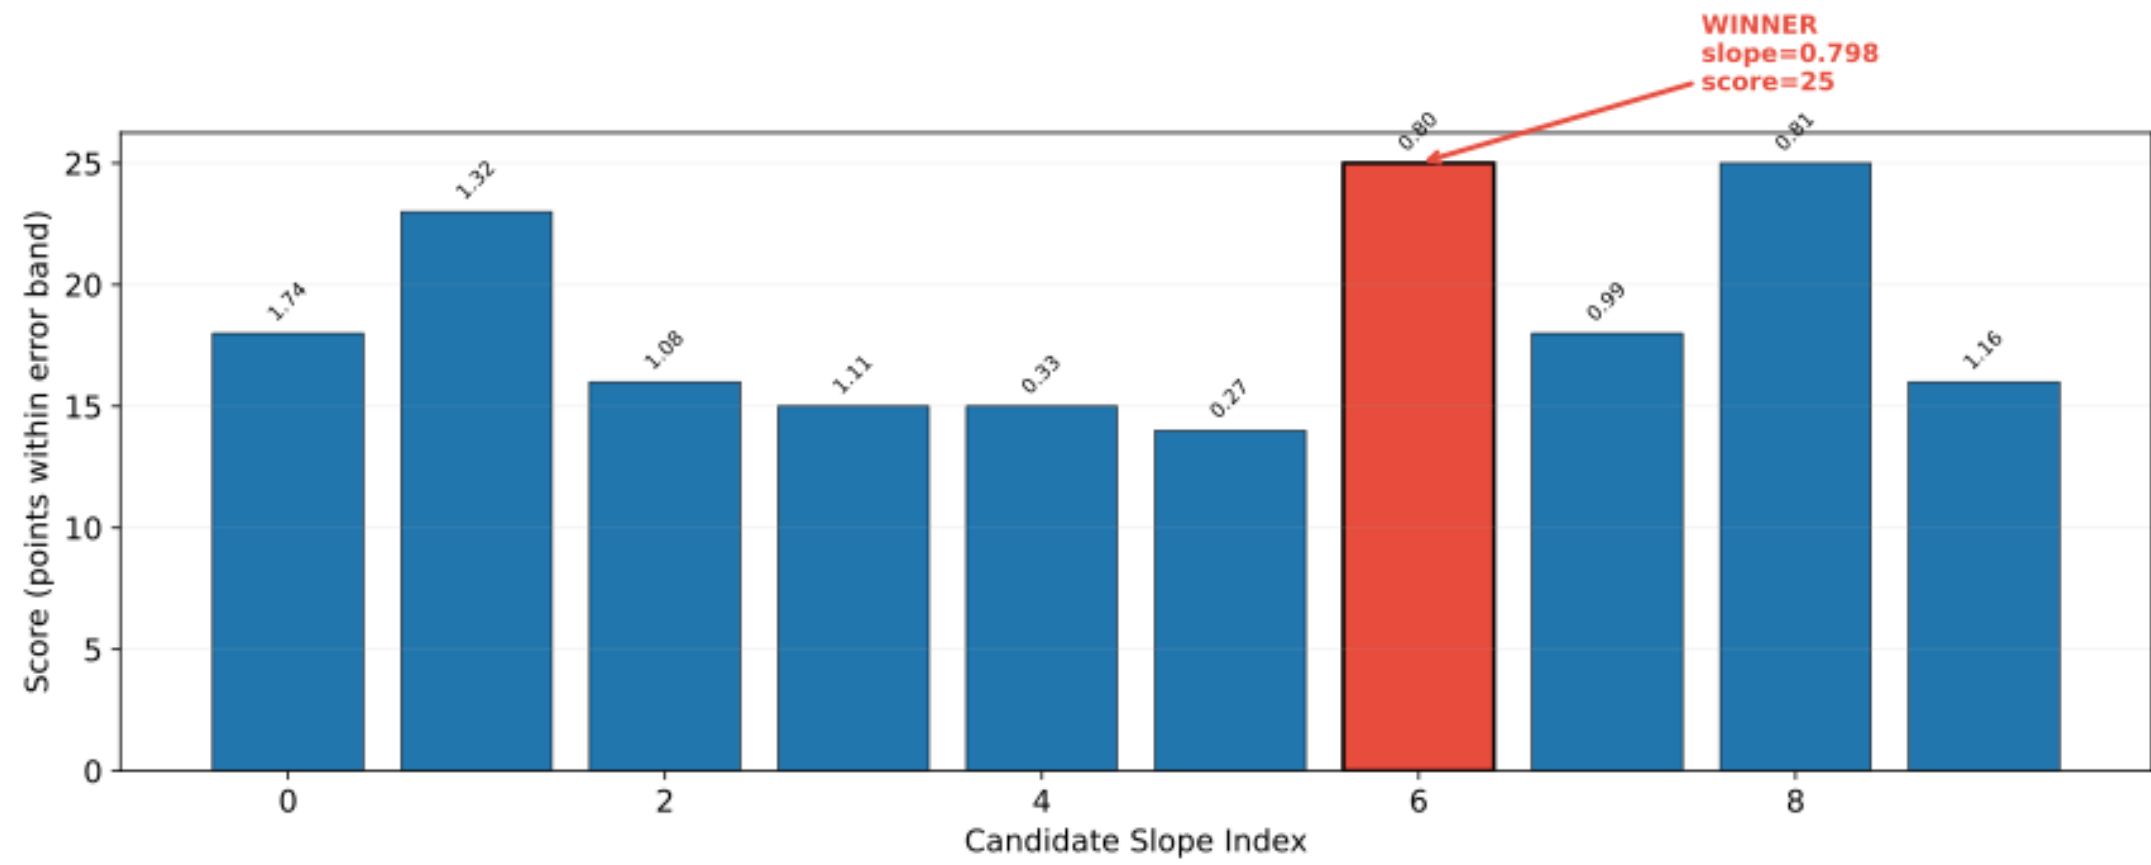
Aim: Find best slope

Find intercepts fitting in window

Sort the intercepts and slide a window to find the intercept fitting max points

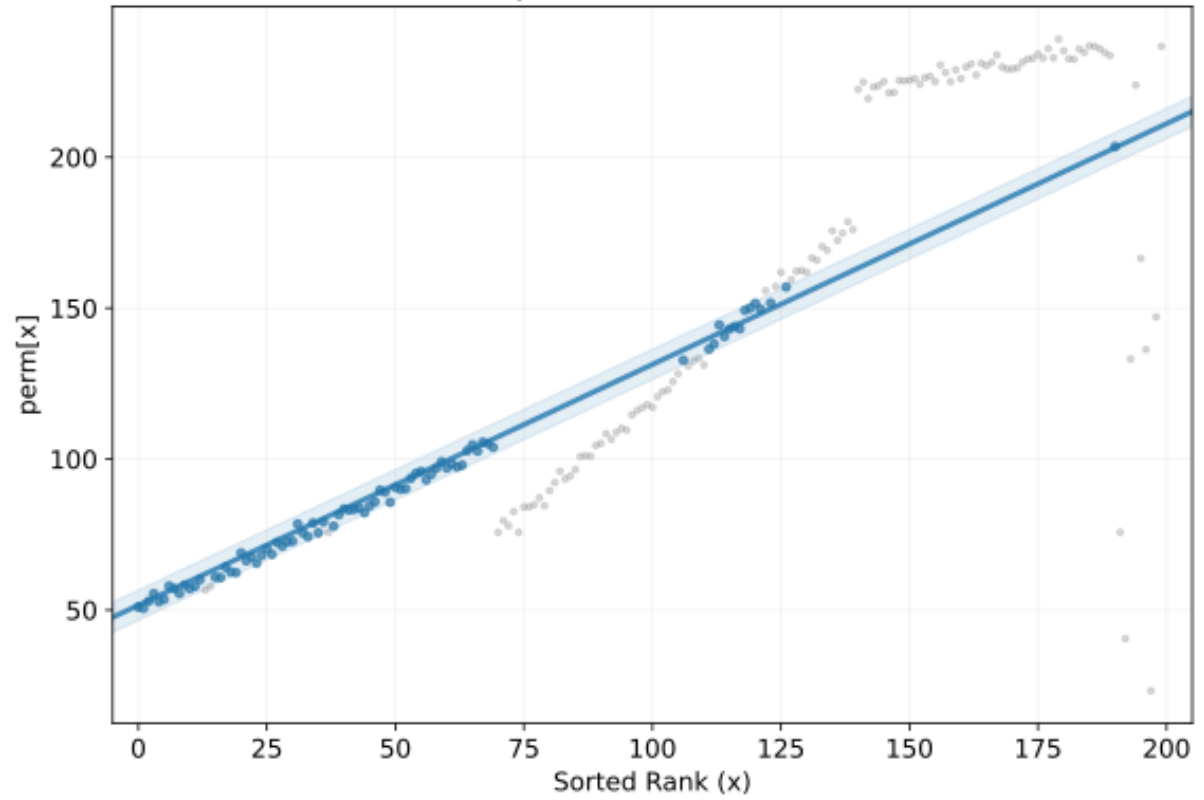


Find best slope

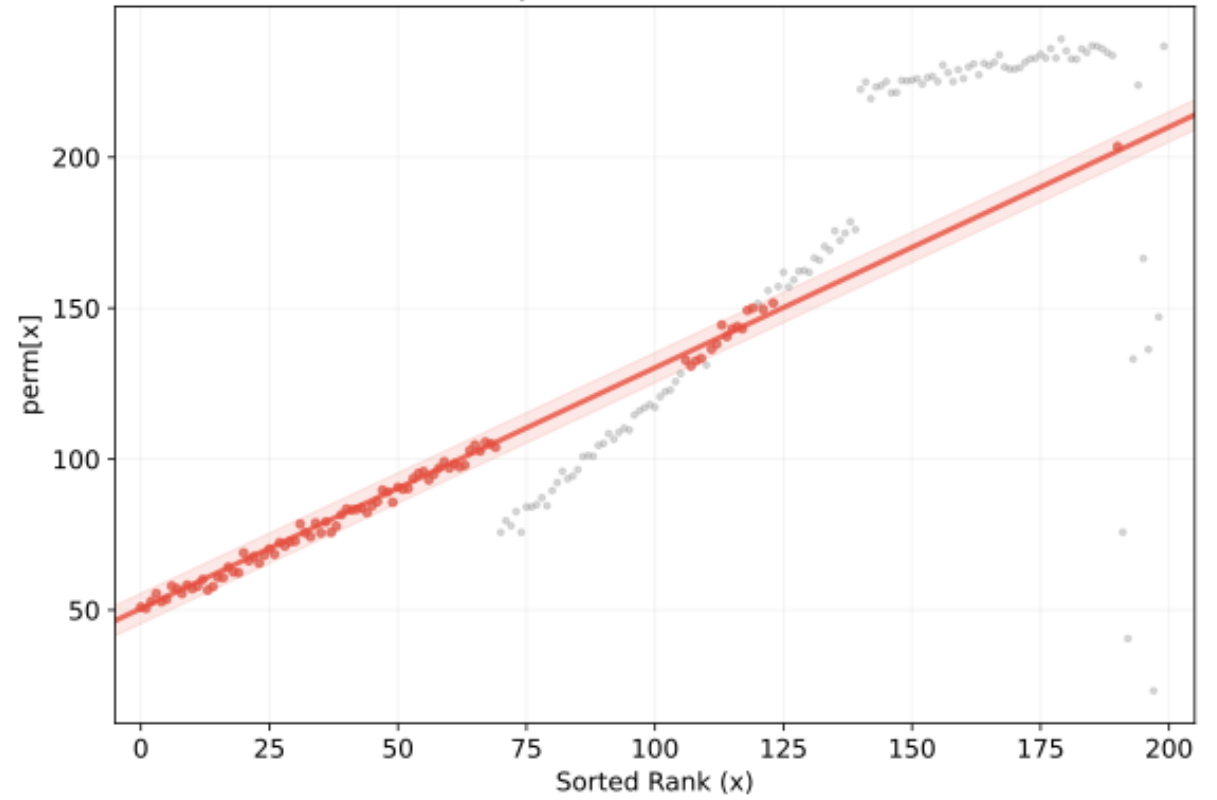


Refine intercept of best slope

Small Eval (1024 pts)
intercept=51.58, covers 82/200

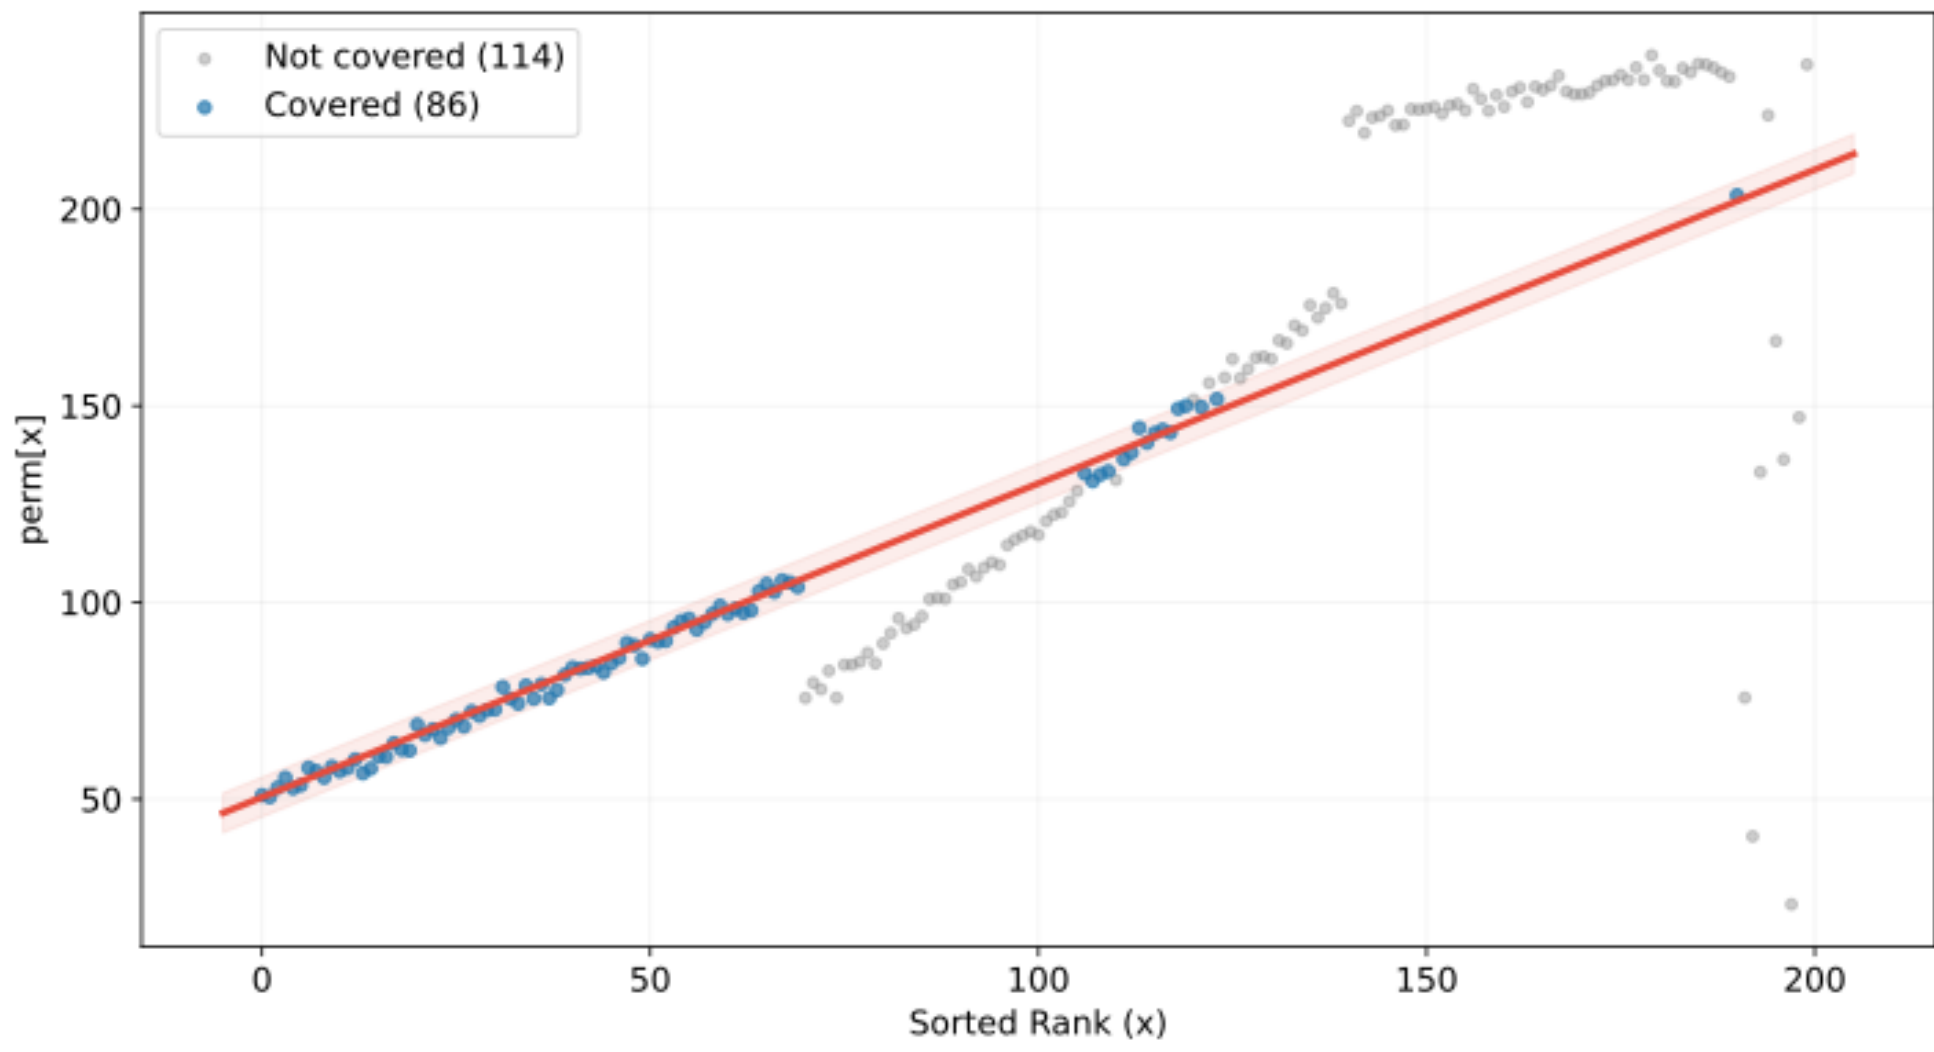


Big Eval (65536 pts)
intercept=50.49, covers 86/200



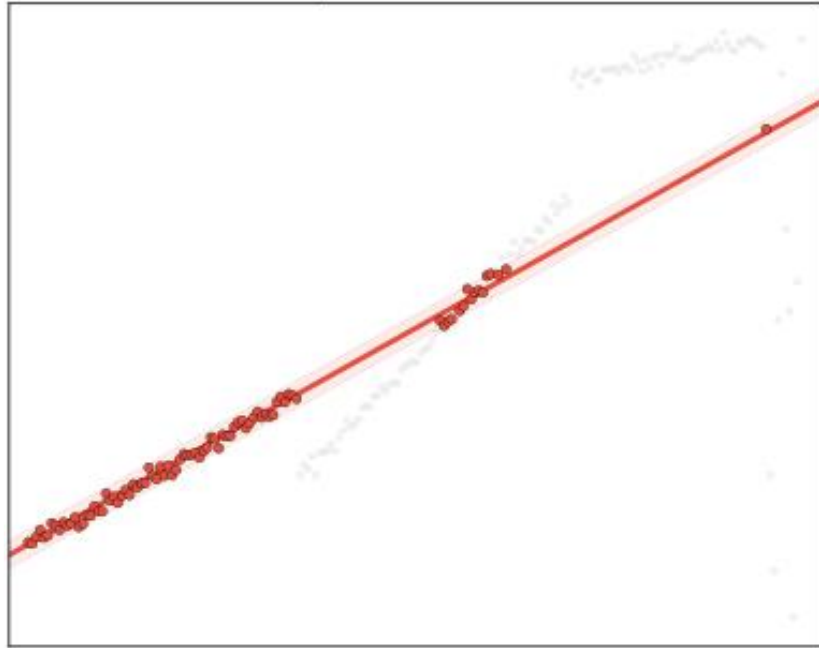
Found first line!

slope=0.798, intercept=50.5, covered=86/200

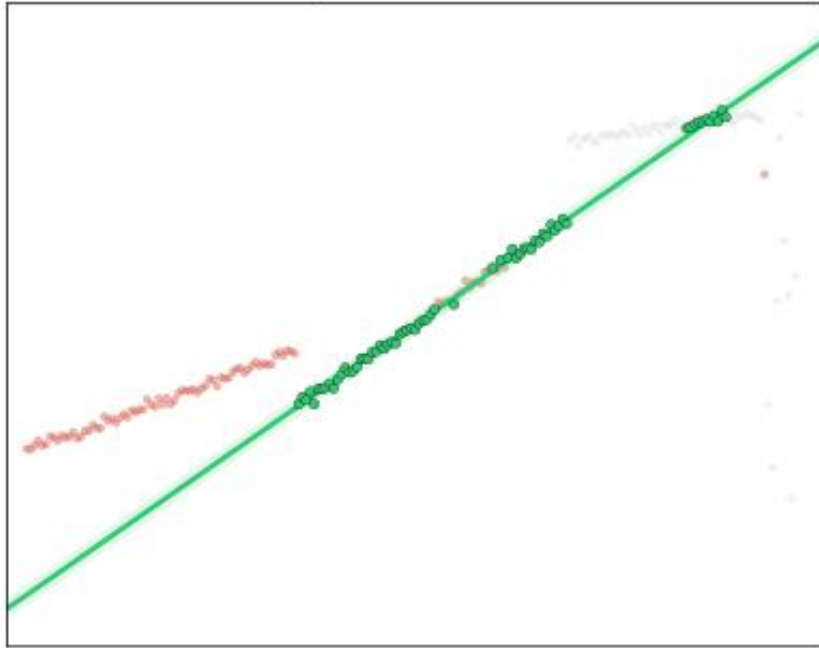


Incremental line cover

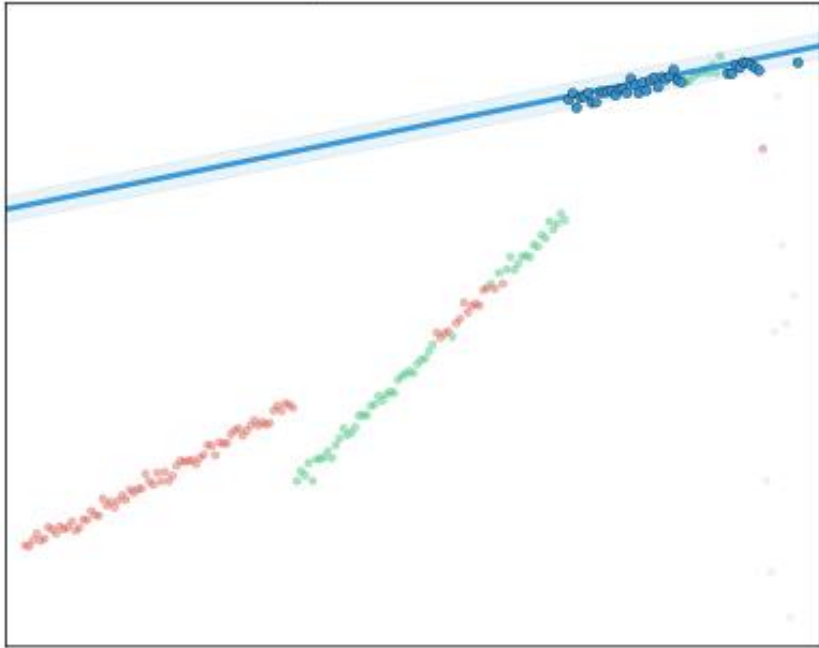
Proposal 1
slope=0.80, covers=86



Proposal 2
slope=1.50, covers=66

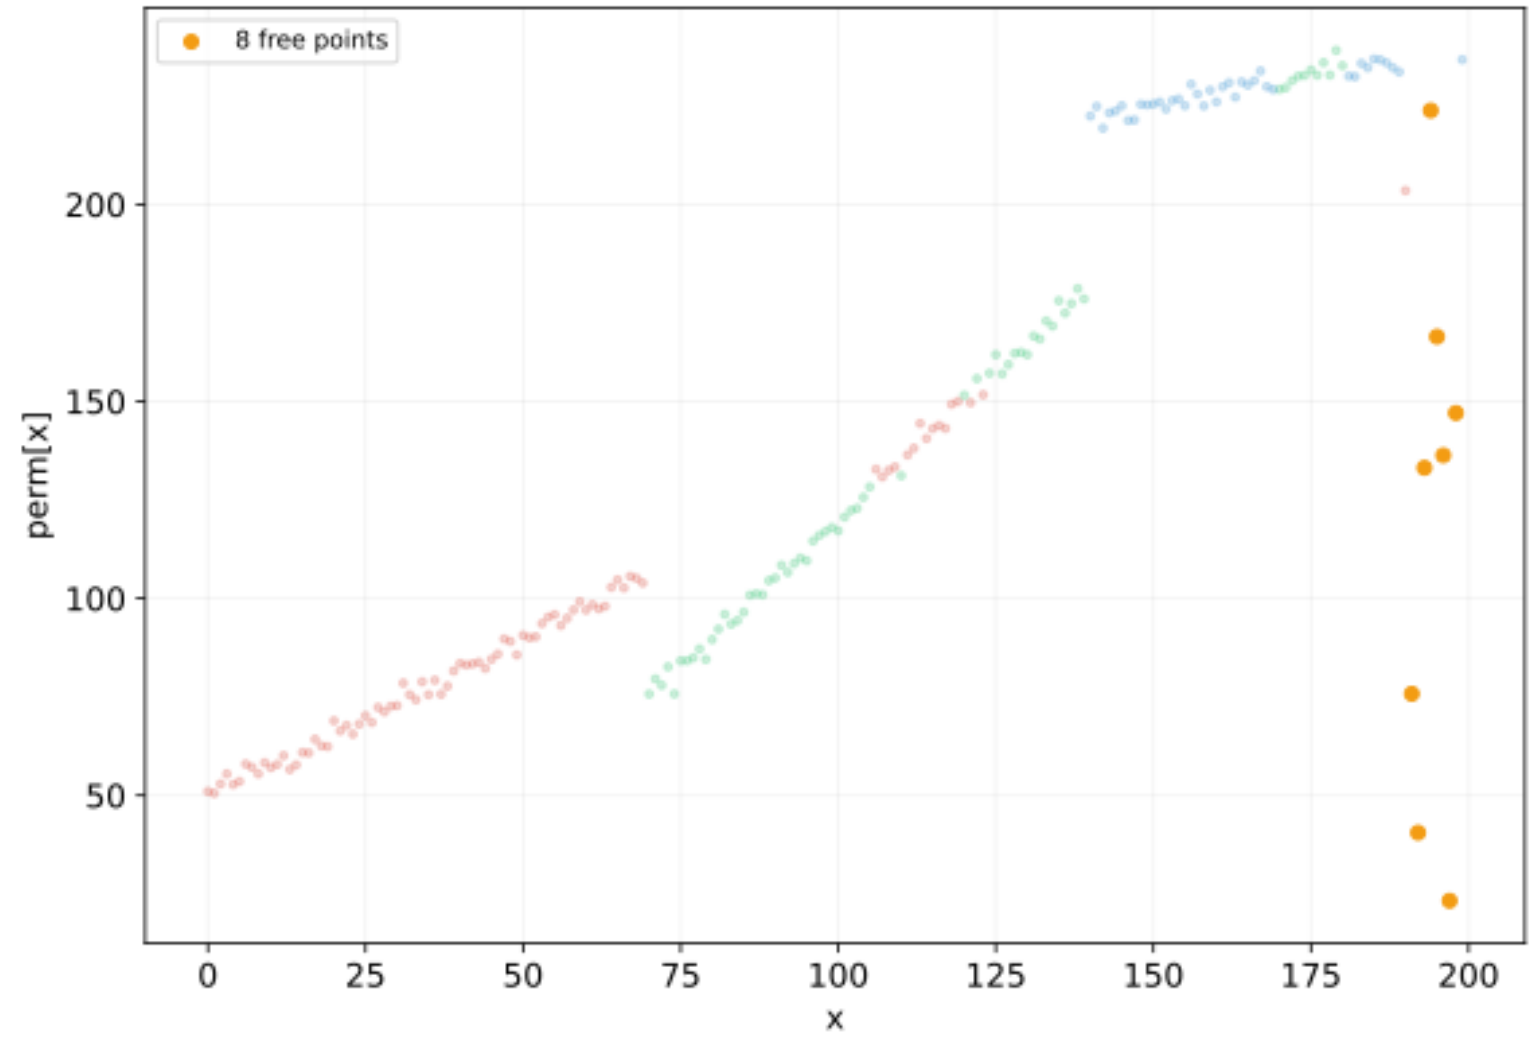


Proposal 3
slope=0.30, covers=40

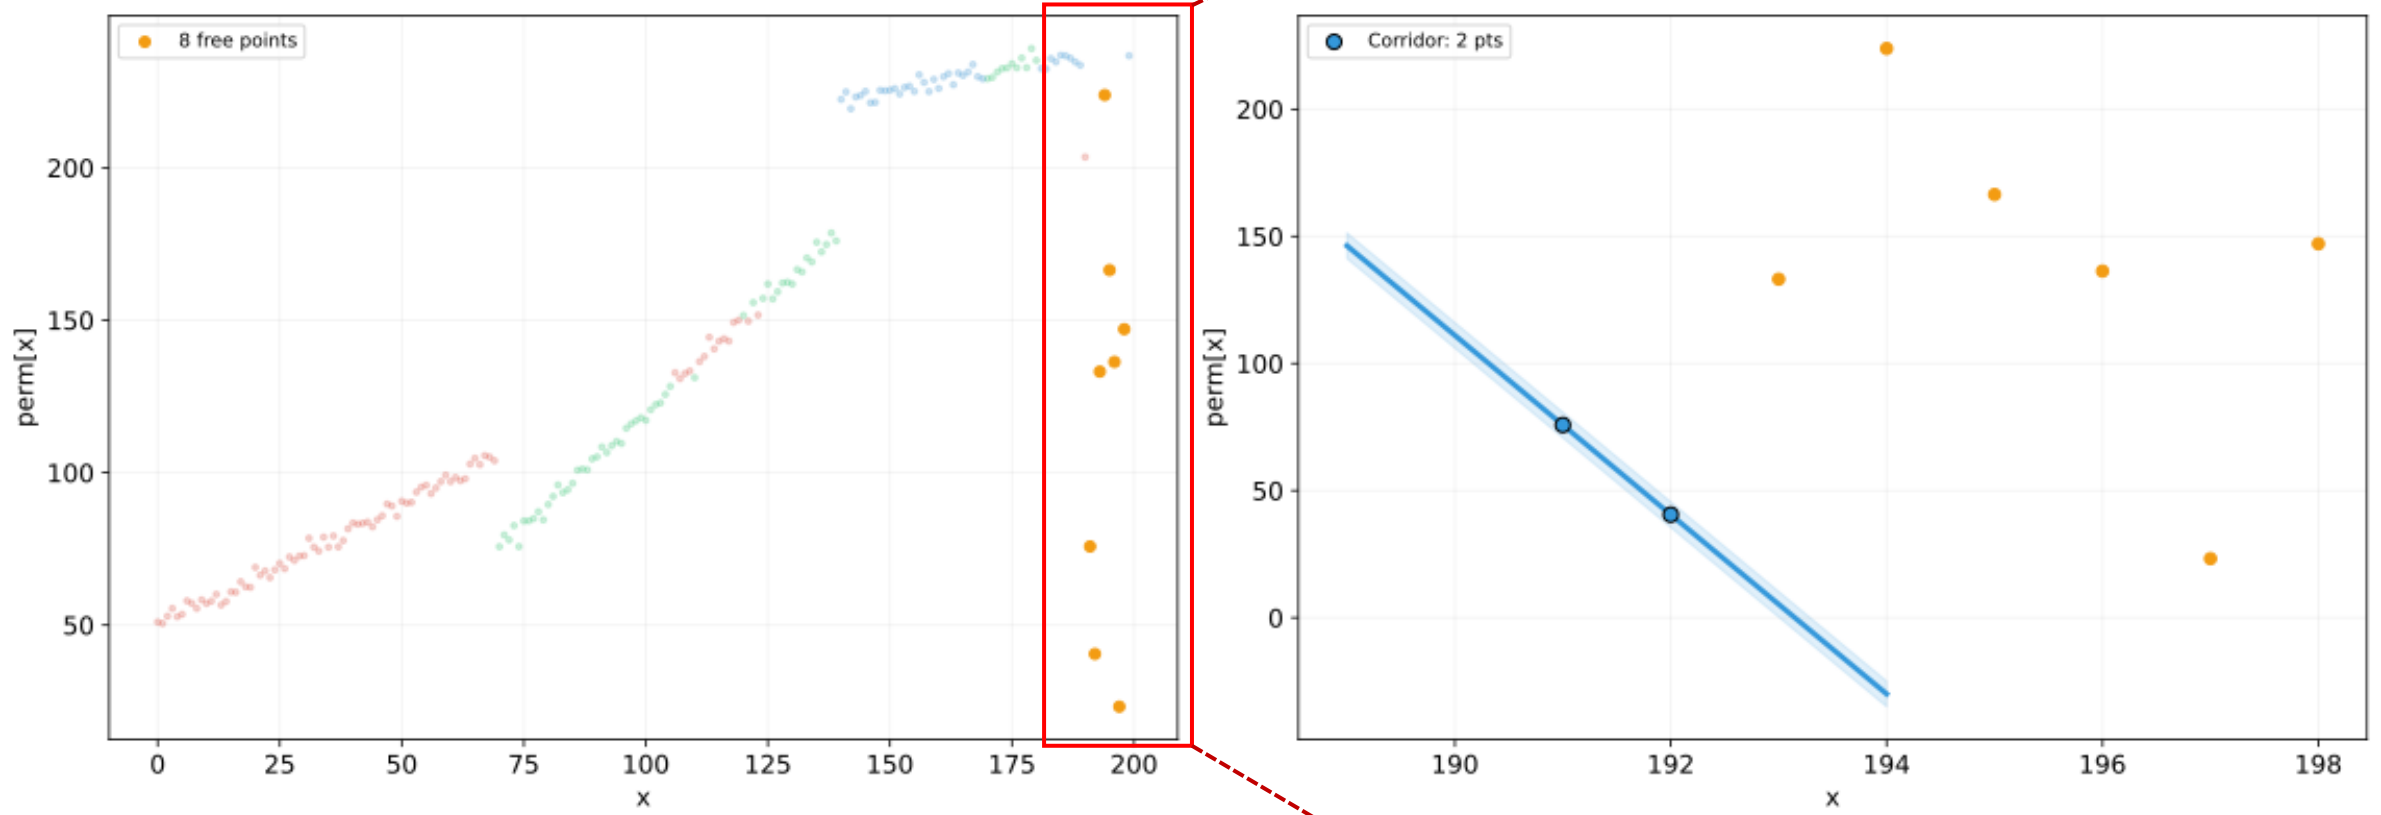


What about the uncovered points?

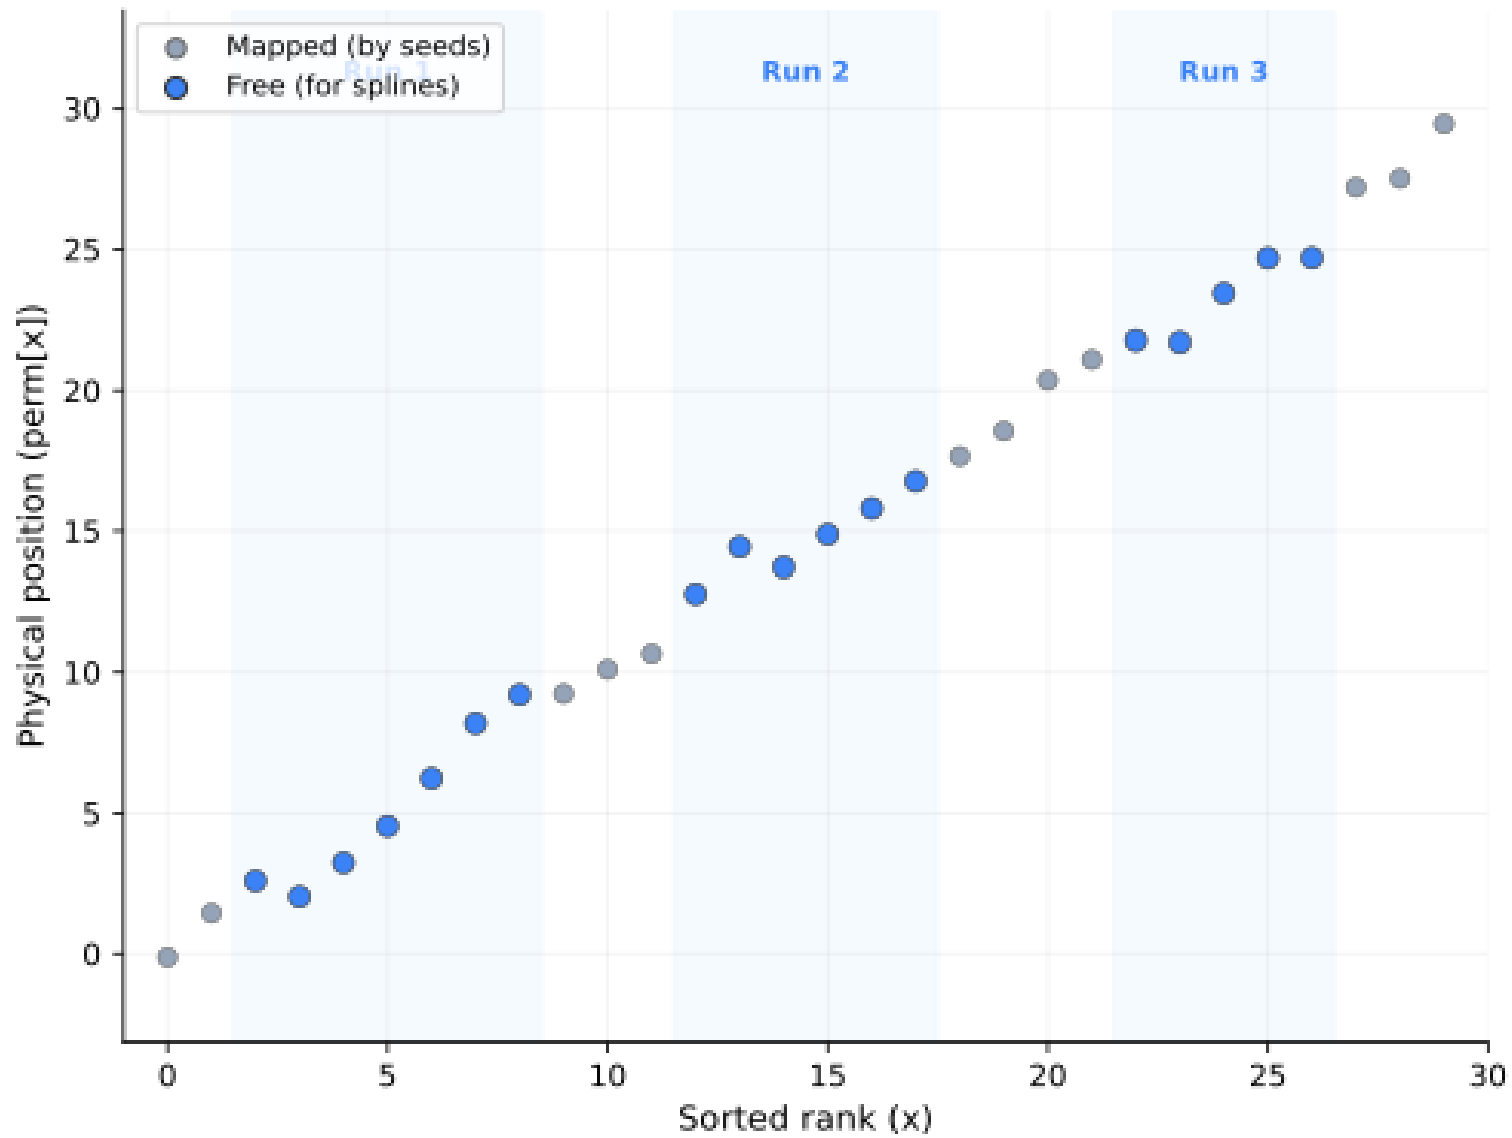
Find Local Pattern



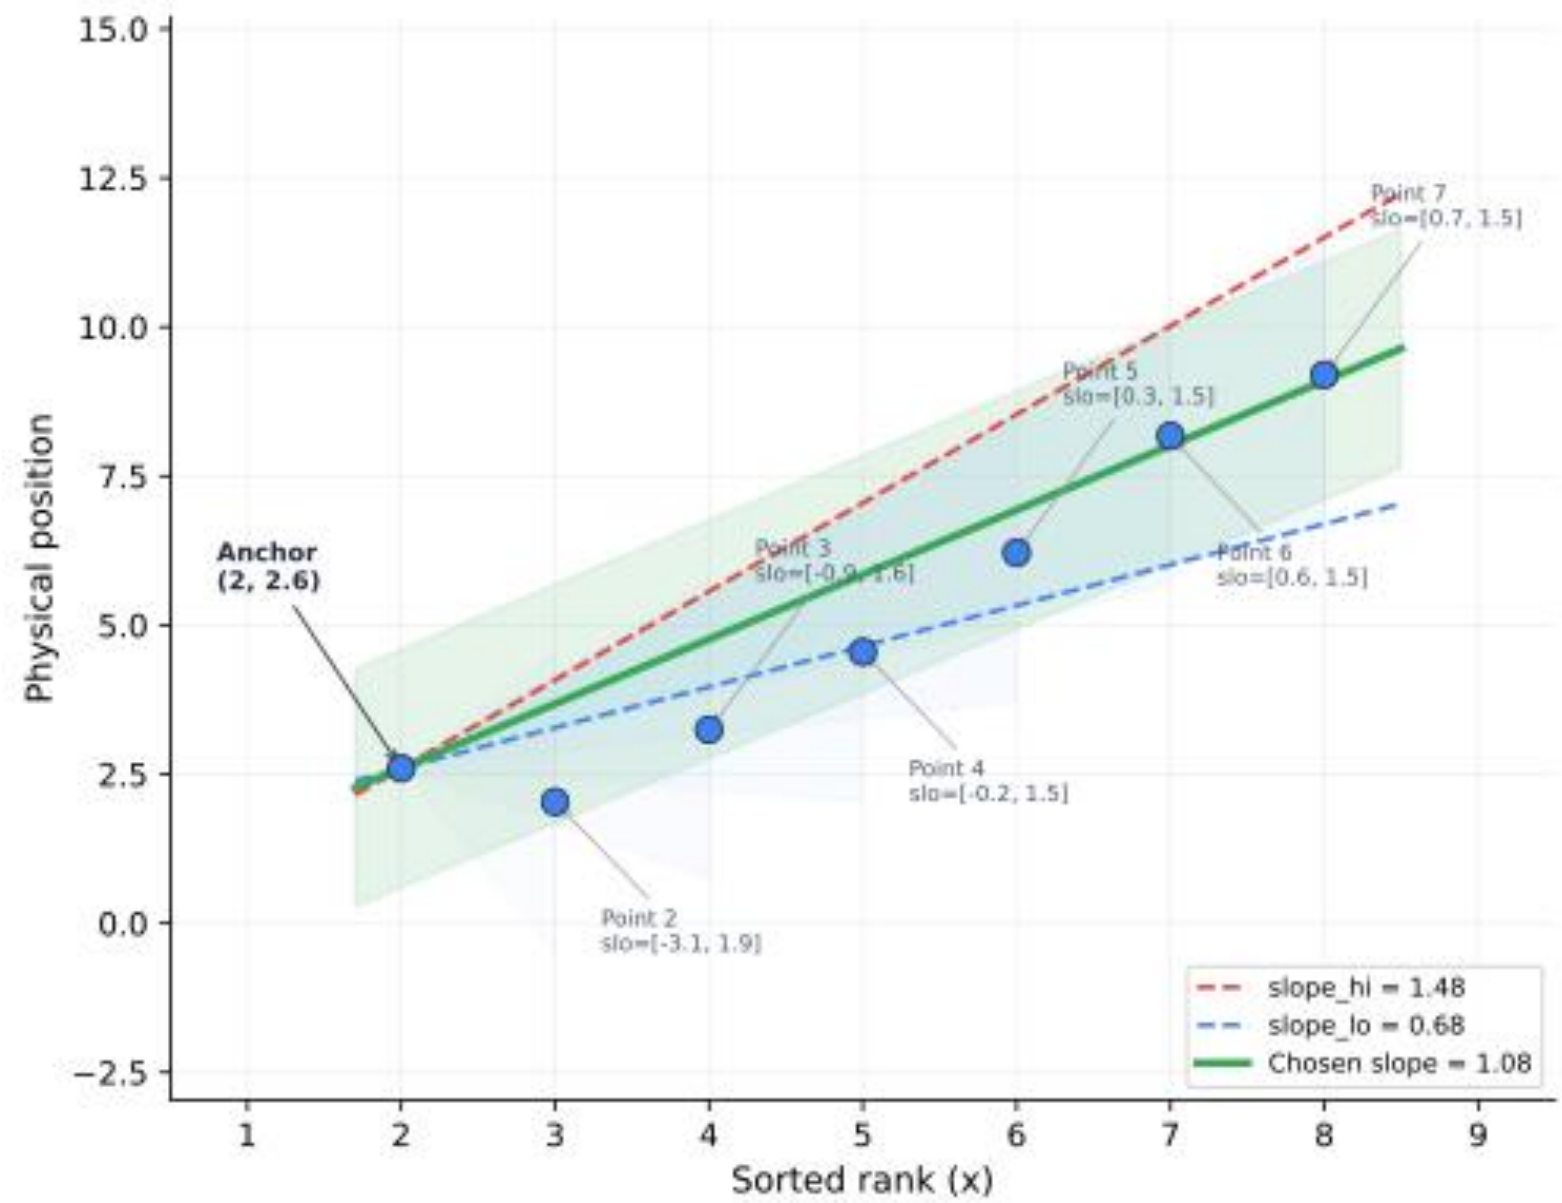
Find Local Pattern



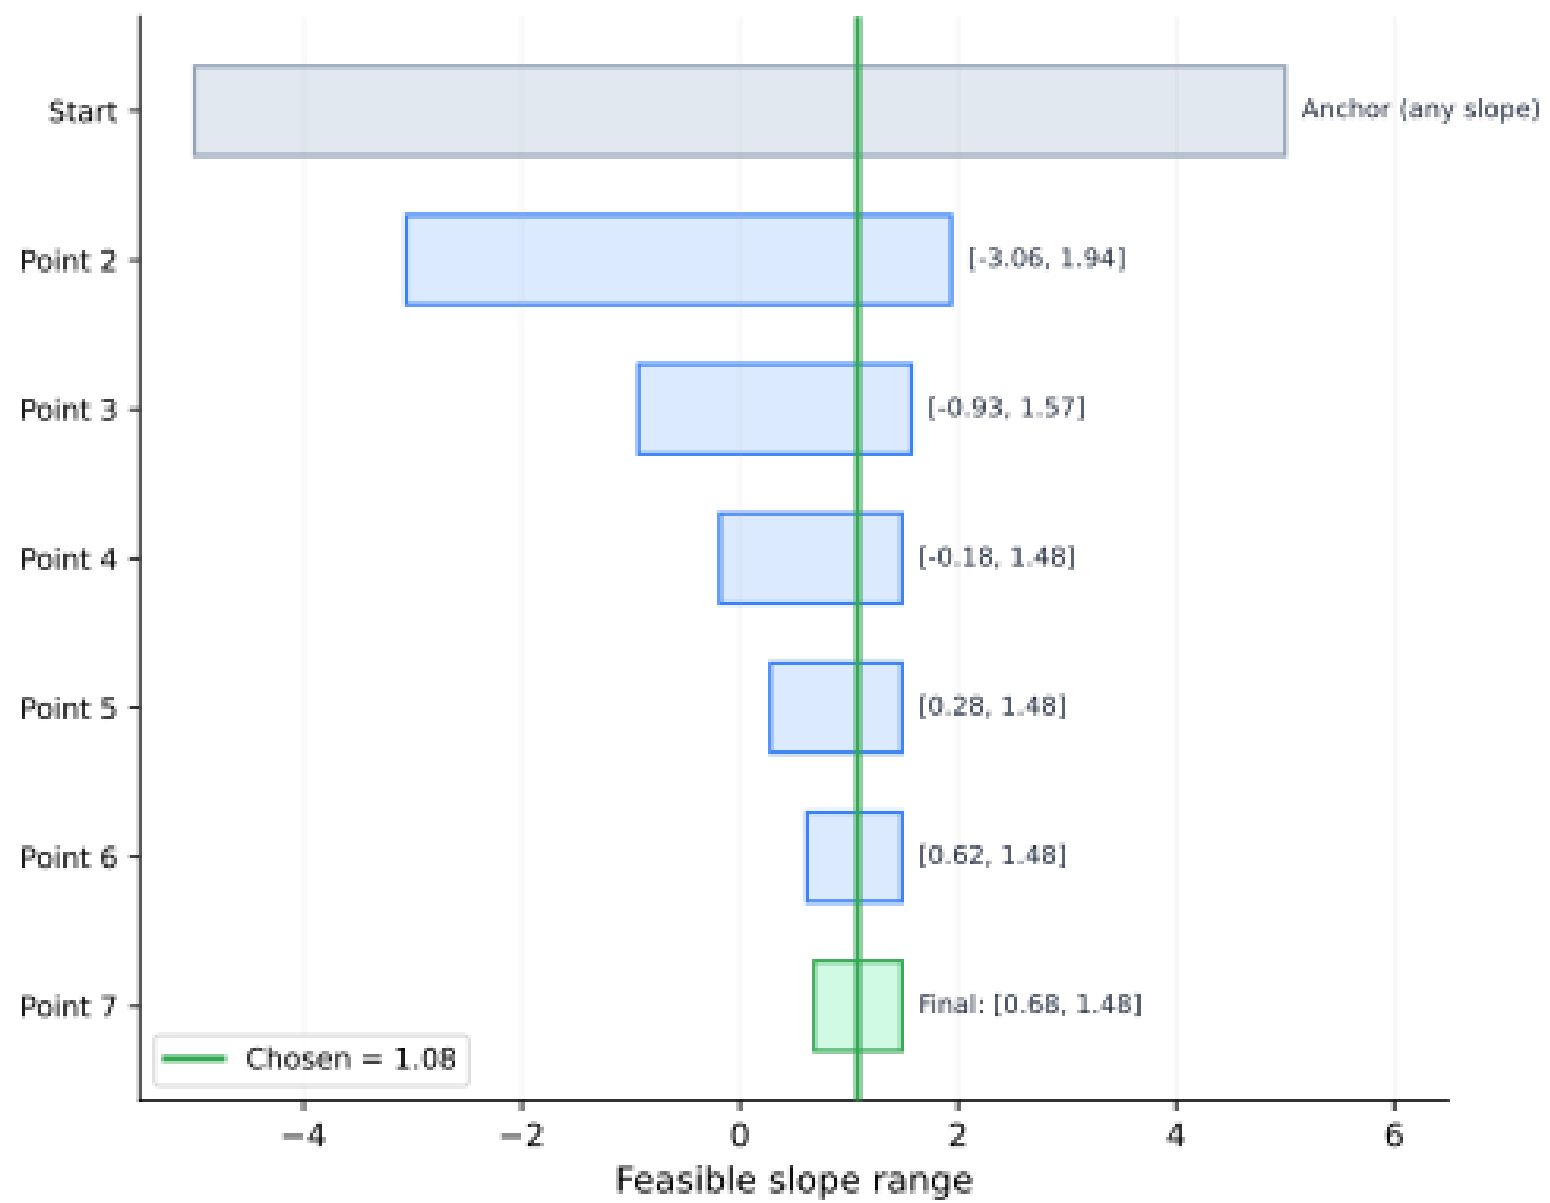
Find Local Pattern



Find Local Pattern

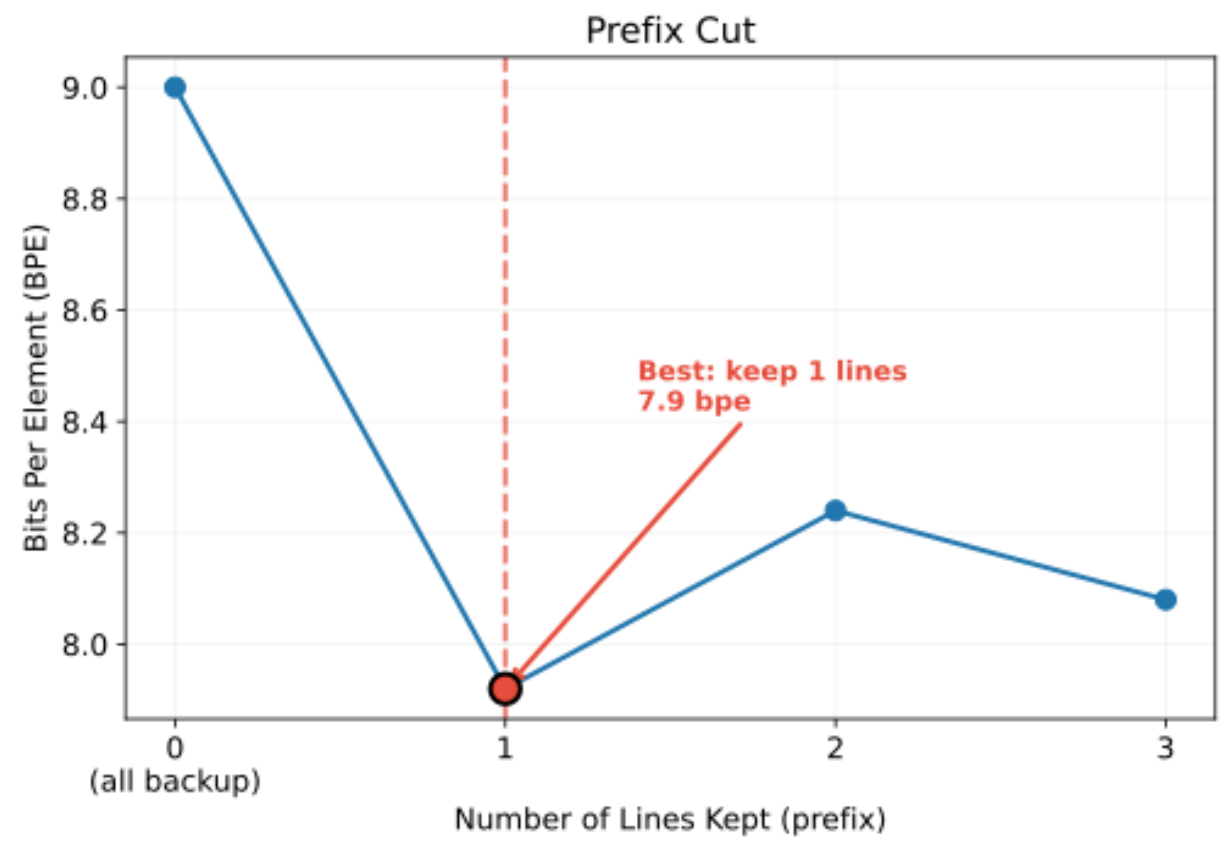
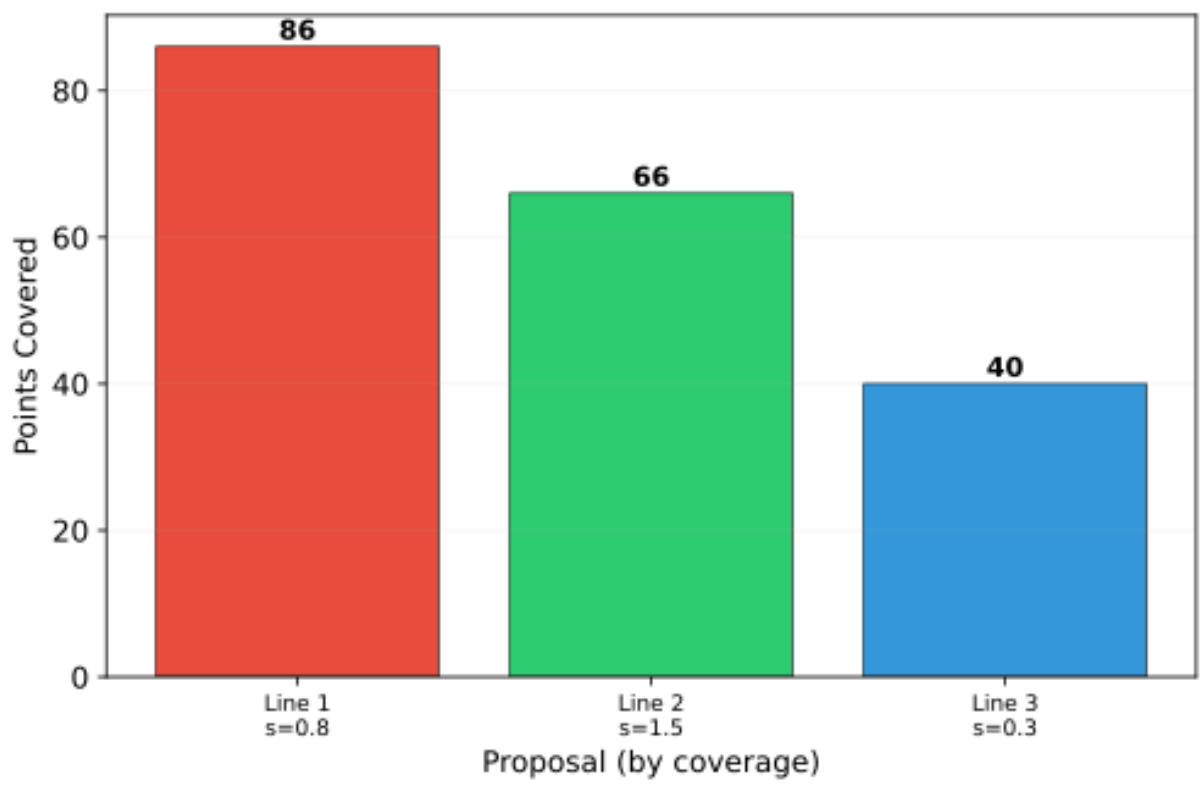


Convergence of slope, intercept



Is this method always good?

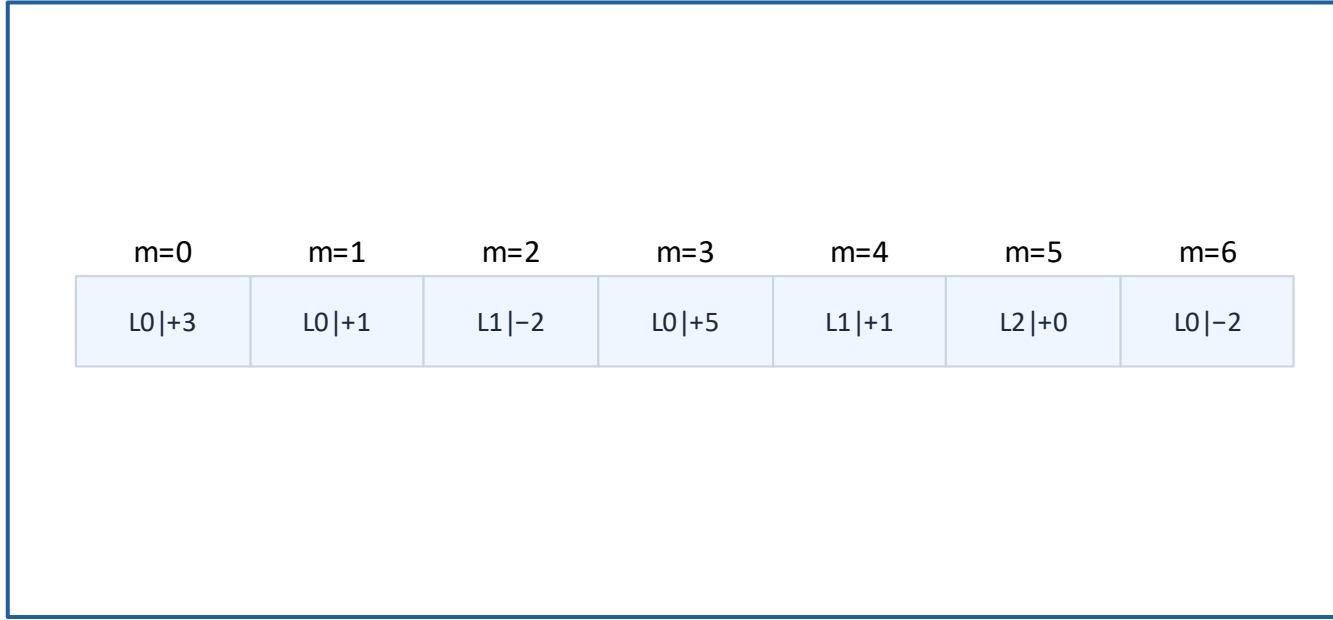
Lines are not always beneficial!



$$\text{bits} = (\text{n_lines} \cdot 64 \text{ bits per line}) + (\text{mapped} \times (\text{lid_bits} + \text{err_bits})) + (\text{backup} \times \text{ceil_log2}(\text{N})) + (\text{rank overhead})$$

What does the entire system look like?

CMap Storage Structure



CMap Storage Structure

Mapped Array (map_pack_) - bit-packed [line_id | error]

m=0	m=1	m=2	m=3	m=4	m=5	m=6
L0 +3	L0 +1	L1 -2	L0 +5	L1 +1	L2 +0	L0 -2

CMap Storage Structure

Mapped Array (map_pack_) - bit-packed [line_id | error]

m=0	m=1	m=2	m=3	m=4	m=5	m=6
L0 +3	L0 +1	L1 -2	L0 +5	L1 +1	L2 +0	L0 -2

lid	slope	intercept
L0	0.82	51.3
L1	1.47	-28.1
L2	0.31	182.0

CMap Storage Structure

Mapped Array (map_pack_) - bit-packed [line_id | error]

m=0	m=1	m=2	m=3	m=4	m=5	m=6
L0 +3	L0 +1	L1 -2	L0 +5	L1 +1	L2 +0	L0 -2

Lines Table (lines_[lid])

lid	slope	intercept
L0	0.82	51.3
L1	1.47	-28.1
L2	0.31	182.0

CMap Storage Structure

Mapped Array (map_pack_) - bit-packed [line_id | error]

m=0	m=1	m=2	m=3	m=4	m=5	m=6
L0 +3	L0 +1	L1 -2	L0 +5	L1 +1	L2 +0	L0 -2

Lines Table (lines_[lid])

lid	slope	intercept
L0	0.82	51.3
L1	1.47	-28.1
L2	0.31	182.0

Decode:
 $pos = \text{round}(\text{slope} \cdot x + \text{intercept}) + \text{error}$

CMap Storage Structure

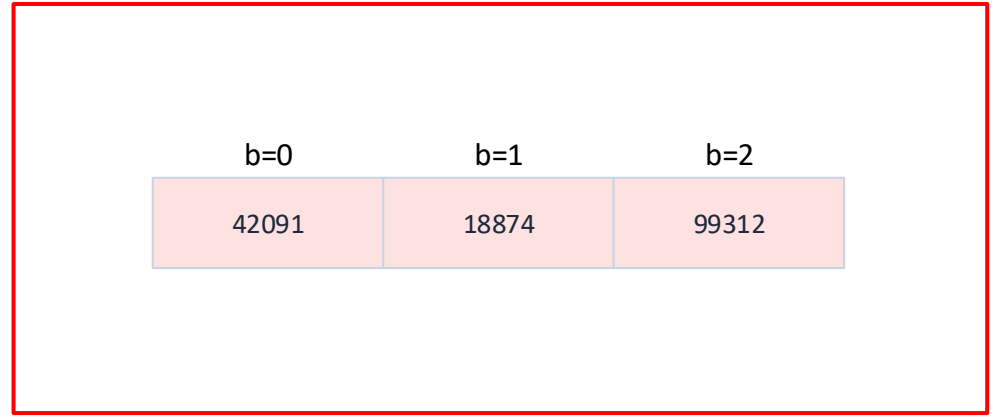
Mapped Array (map_pack_) - bit-packed [line_id | error]

m=0	m=1	m=2	m=3	m=4	m=5	m=6
L0 +3	L0 +1	L1 -2	L0 +5	L1 +1	L2 +0	L0 -2

Lines Table (lines_[lid])

lid	slope	intercept
L0	0.82	51.3
L1	1.47	-28.1
L2	0.31	182.0

Decode:
 $pos = round(s \times x + c) + err$



CMap Storage Structure

Mapped Array (map_pack_) - bit-packed [line_id | error]

m=0	m=1	m=2	m=3	m=4	m=5	m=6
L0 +3	L0 +1	L1 -2	L0 +5	L1 +1	L2 +0	L0 -2

Lines Table (lines_[lid])

lid	slope	intercept
L0	0.82	51.3
L1	1.47	-28.1
L2	0.31	182.0

Decode:

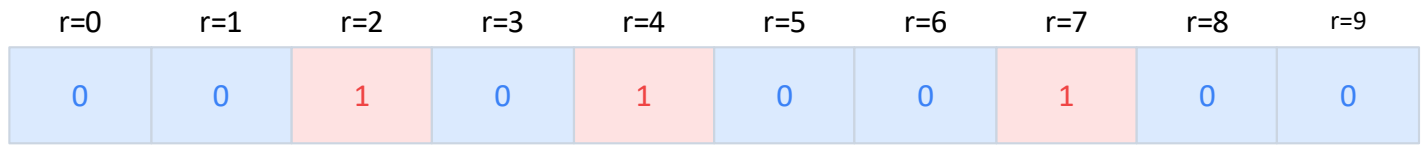
$$\text{pos} = \text{round}(s \times x + c) + \text{err}$$

Backup Array (bak_pack_) - raw positions

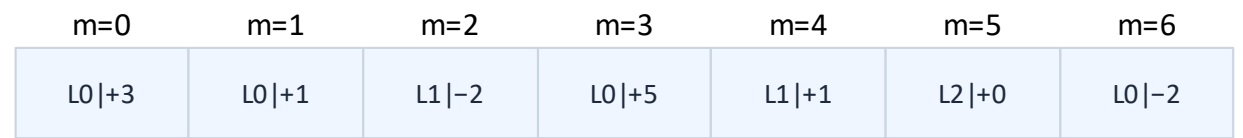
b=0	b=1	b=2
42091	18874	99312

No line decode needed

CMap Storage Structure



Mapped Array (map_pack_) - bit-packed [line_id | error]

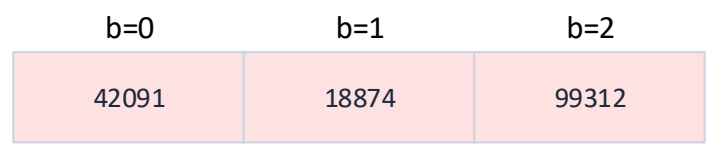


Lines Table (lines_[lid])

lid	slope	intercept
L0	0.82	51.3
L1	1.47	-28.1
L2	0.31	182.0

Decode:
 $pos = \text{round}(s \times x + c) + \text{err}$

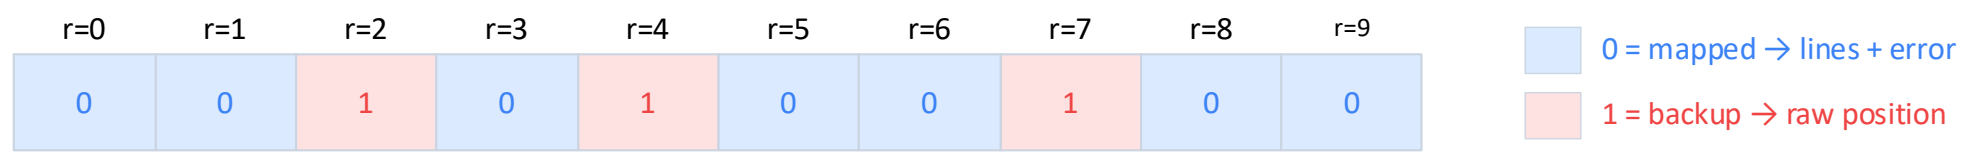
Backup Array (bak_pack_) - raw positions



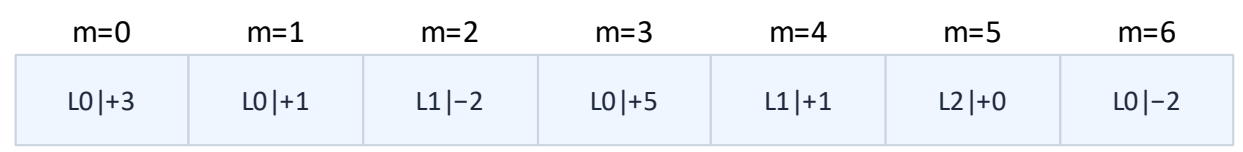
No line decode needed

CMap Storage Structure

Bitvector (1 bit/entry, SDSL rank support)



Mapped Array (map_pack_) - bit-packed [line_id | error]

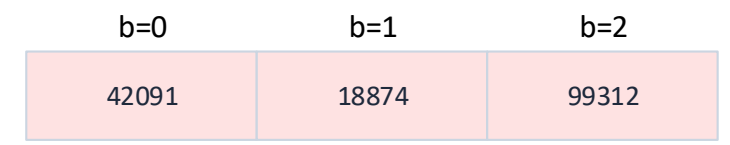


Lines Table (lines_[lid])

lid	slope	intercept
L0	0.82	51.3
L1	1.47	-28.1
L2	0.31	182.0

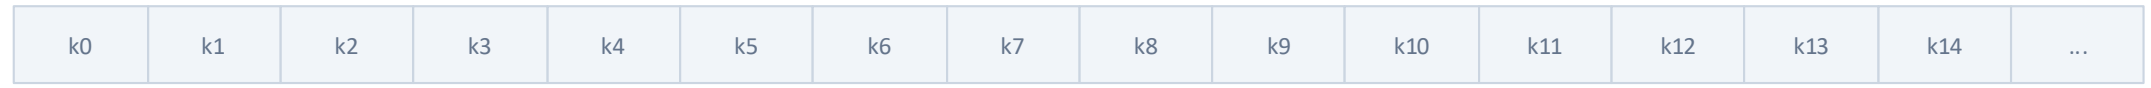
Decode:
 $pos = \text{round}(s \times x + c) + \text{err}$

Backup Array (bak_pack_) - raw positions



No line decode needed.
 Direct physical position.

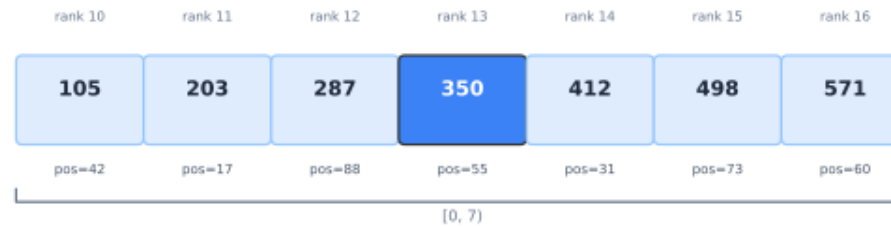
data_[pos]



Single Query: Binary Search Path

Binary search over a 7-entry window (query key = 412)

Step 1: mid=3 → key=350 < 412 → go right

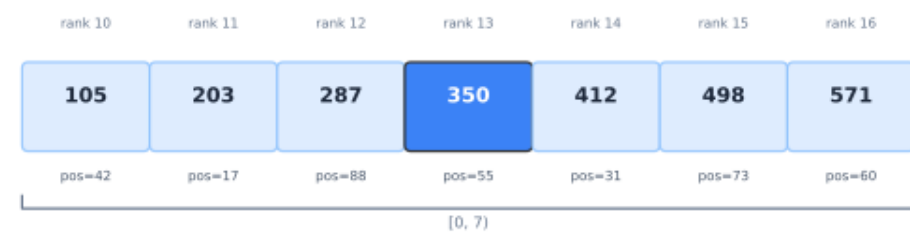


Step 2: mid=5 → key=498 > 412 → go left

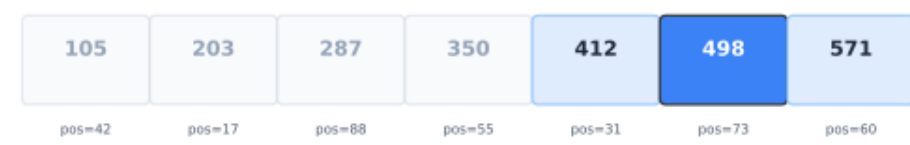
Single Query: Binary Search Path

Binary search over a 7-entry window (query key = 412)

Step 1: mid=3 → key=350 < 412 → go right



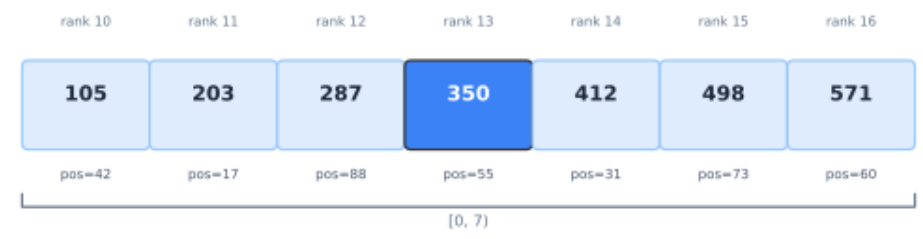
Step 2: mid=5 → key=498 > 412 → go left



Single Query: Binary Search Path

Binary search over a 7-entry window (query key = 412)

Step 1: mid=3 → key=350 < 412 → go right



Step 2: mid=5 → key=498 > 412 → go left



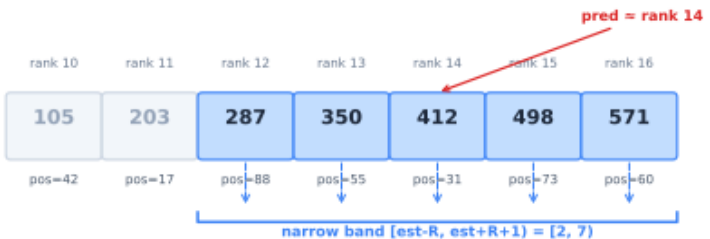
Step 3: mid=4 → key=412 = 412



Single Query: Restricted Search

Restricted-window decode (R=2, query key = 412)

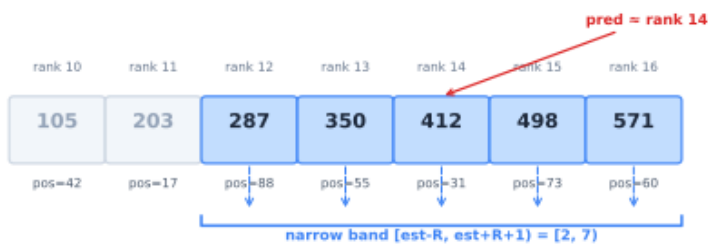
Phase 1: Decode narrow band [2..6], prefetch data[]



Single Query: Restricted Search

Restricted-window decode (R=2, query key = 412)

Phase 1: Decode narrow band [2..6], prefetch data[]



Phase 2: Scan — index 2: 287 < 412, continue



Single Query: Restricted Search

Restricted-window decode (R=2, query key = 412)

Phase 1: Decode narrow band [2..6], prefetch data[]



Phase 2: Scan — index 2: 287 < 412, continue



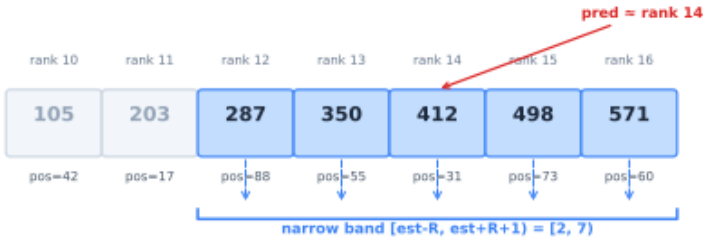
Phase 3: Scan — index 3: 350 < 412, continue



Single Query: Restricted Search

Restricted-window decode (R=2, query key = 412)

Phase 1: Decode narrow band [2..6], prefetch data[]



Phase 2: Scan — index 2: 287 < 412, continue



Phase 3: Scan — index 3: 350 < 412, continue

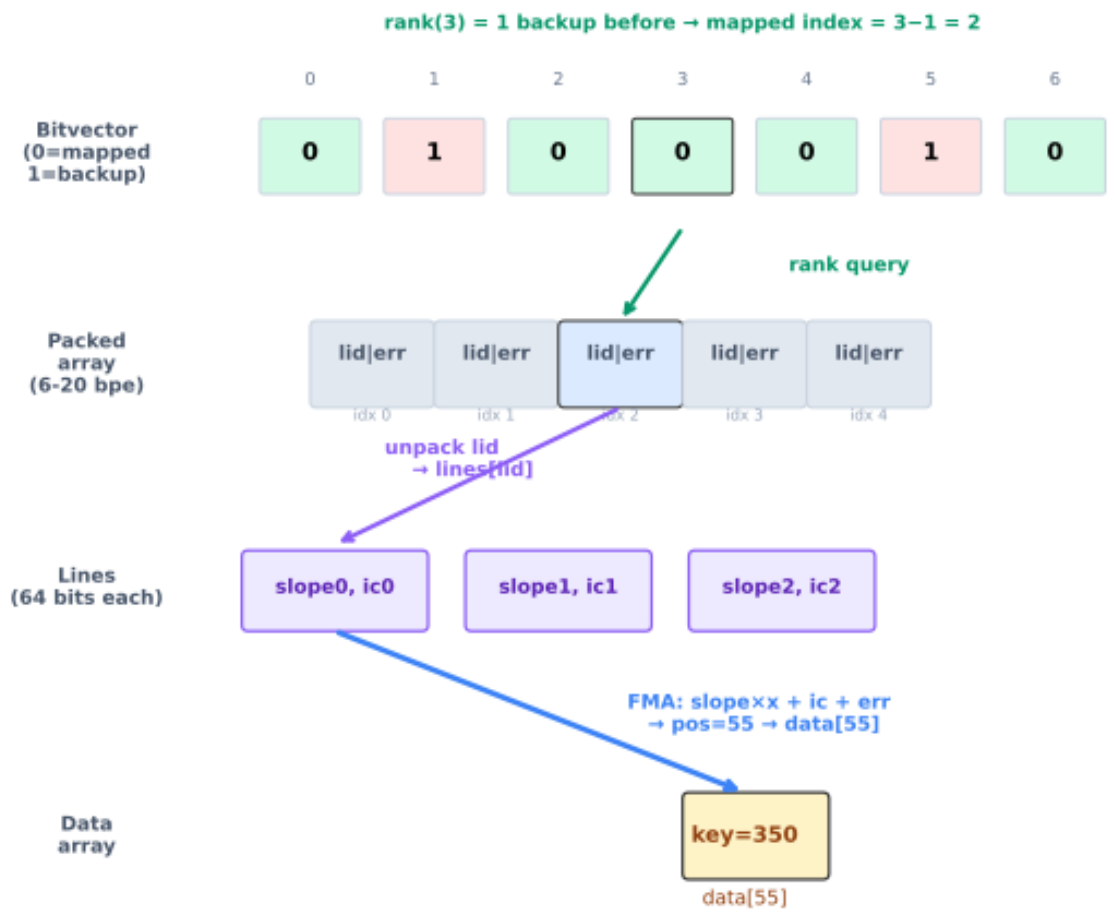


Phase 4: Scan — index 4: 412 = 412, found!



Single Query: Decode

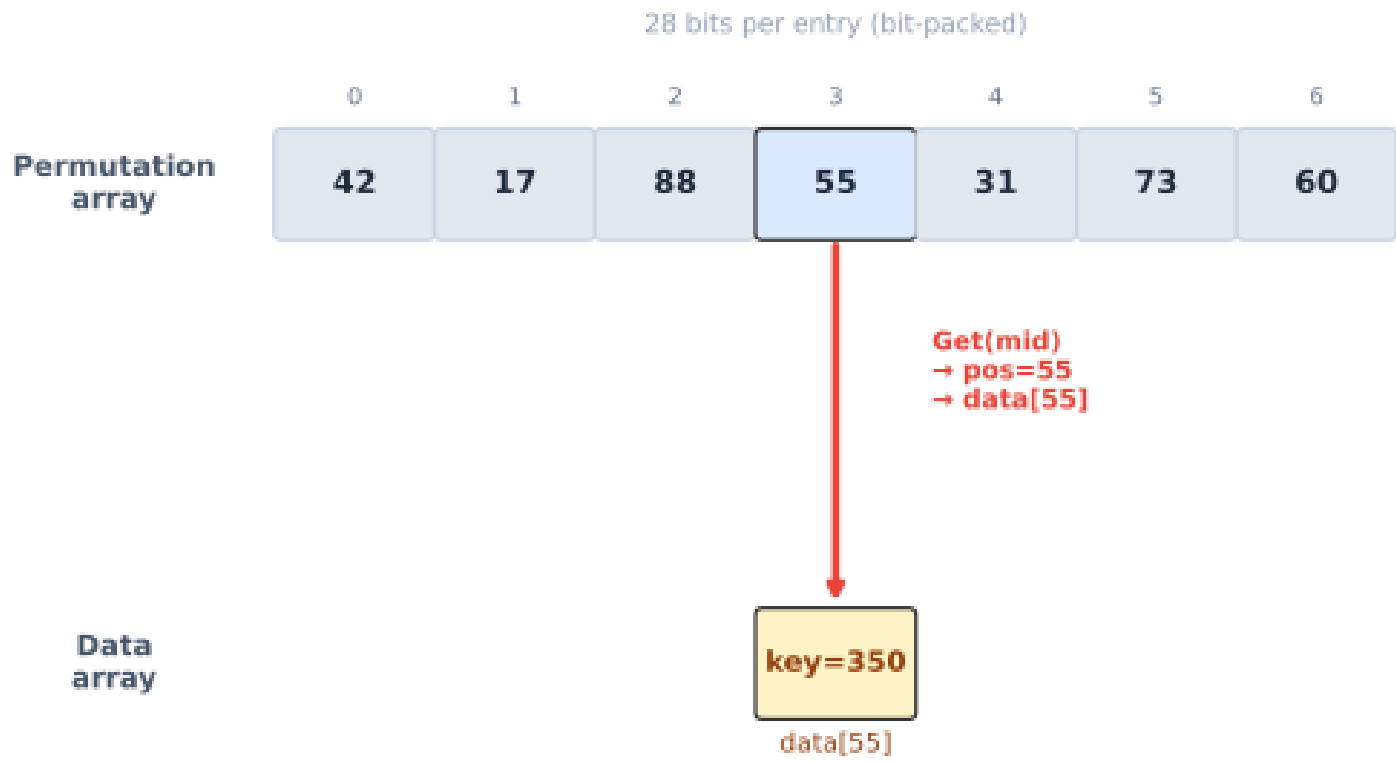
CMap decode: rank + unpack + FMA per step



rank + bit-packed load + line access + FMA + data access
= 3-4 dependent cache accesses per step

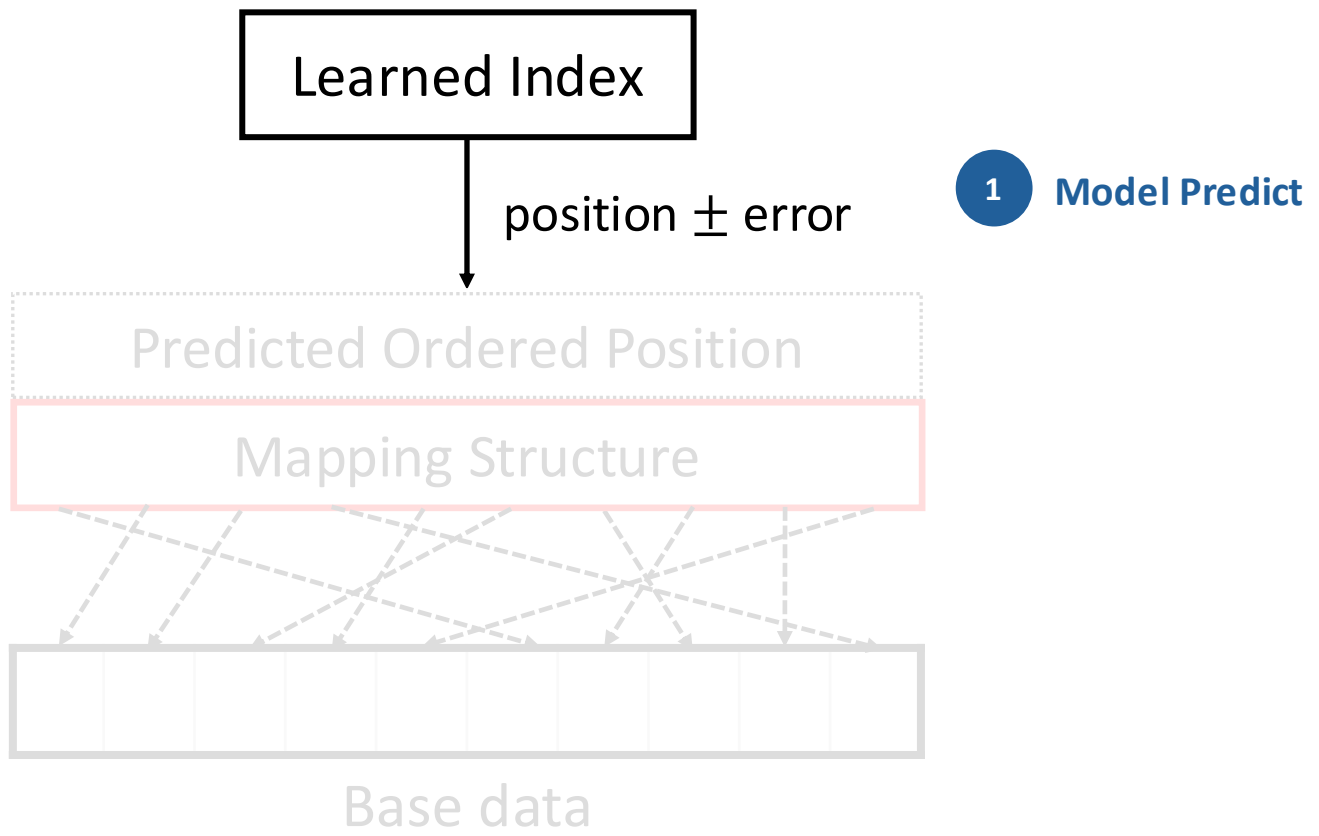
Single Query: LSI Access Path

LSI decode: one bit-packed load per step

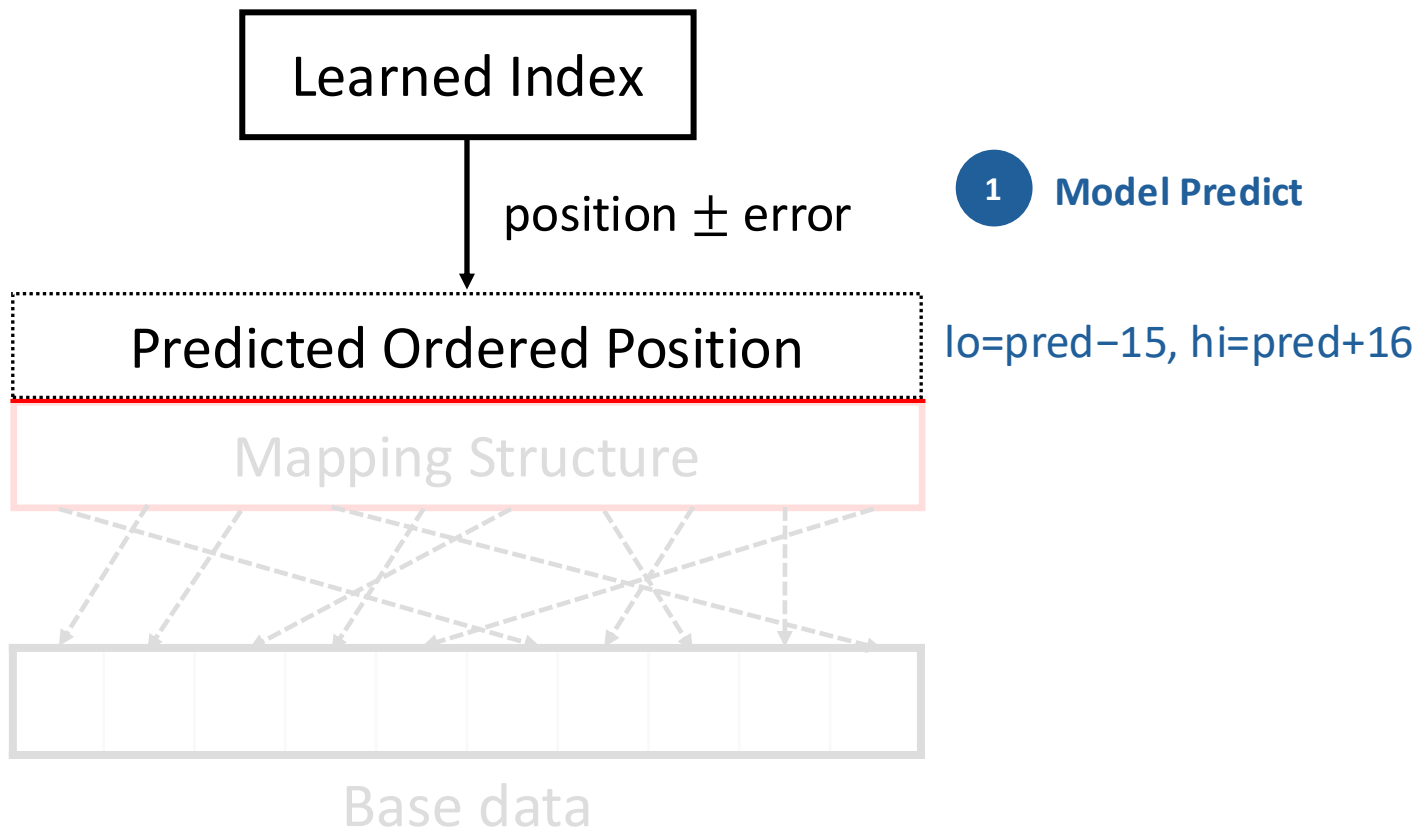


1 bit-packed load + 1 data access per step
= 2 dependent cache accesses

Single Query: Access Path



Single Query: Access Path



Single Query: Access Path



Single Query: Access Path

1 Model Predict
 lo=pred-15, hi=pred+16

2 rank_window(lo) → base_rank + 32-bit window

0	0	1	0	0	0	0	1	0	0
---	---	---	---	---	---	---	---	---	---

↑ est

3a Backup (bit=1)

bak_pack_ → raw pos

42091

3b Mapped (bit=0)

map_pack_read:

L0 +3	L1 -2	L0 +5
-------	-------	-------

↓ unpack

lines_[lid]:

L0	0.82	51.3
----	------	------

↓ FMA + SSE round

pos = round(slope × rank + intercept) + error

Single Query: Access Path

1 Model Predict
 $lo = pred - 15, hi = pred + 16$



2 rank_window(lo) → base_rank + 32-bit window

0	0	1	0	0	0	0	1	0	0
---	---	---	---	---	---	---	---	---	---

↑ est

3a Backup (bit=1)
 bak_pack_ → raw pos

42091

prefetch(data_[pos])

3b Mapped (bit=0)

map_pack_read:

L0 +3	L1 -2	L0 +5
-------	-------	-------

↓ unpack

lines_[lid]:

L0	0.82	51.3
----	------	------

↓ FMA + SSE round

prefetch(data_[pos])

$pos = \text{round}(\text{slope} \times \text{rank} + \text{intercept}) + \text{error}$

Single Query: Access Path

1 Model Predict
 lo=pred-15, hi=pred+16

2 rank_window(lo) → base_rank + 32-bit window

0	0	1	0	0	0	0	1	0	0
---	---	---	---	---	---	---	---	---	---

↑est

3a Backup (bit=1)
 bak_pack_ → raw pos
 42091
 prefetch(data_[pos])

3b Mapped (bit=0)
 map_pack_read:

L0 +3	L1 -2	L0 +5
-------	-------	-------

 ↓ unpack
 lines_[lid]:

L0	0.82	51.3
----	------	------

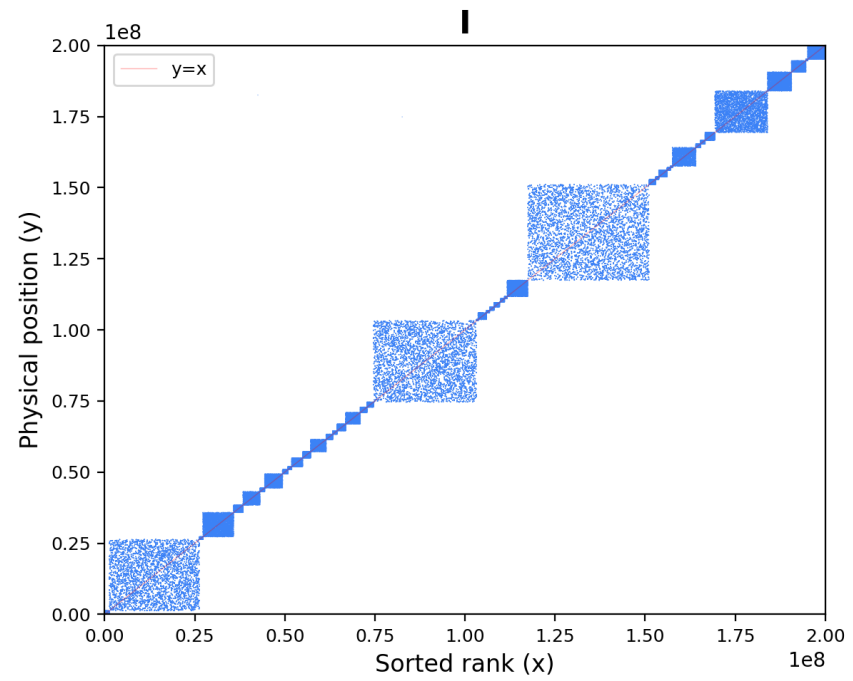
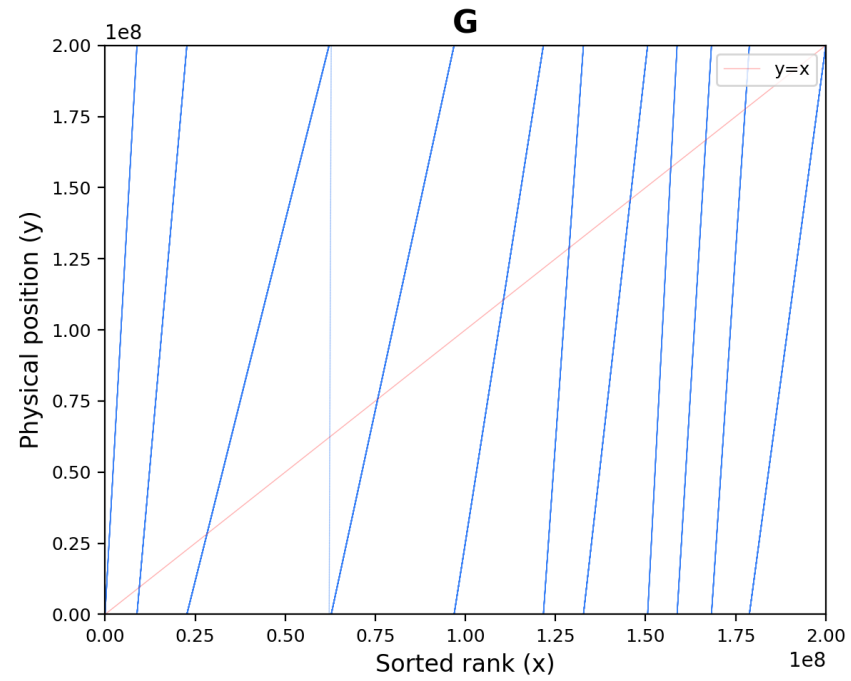
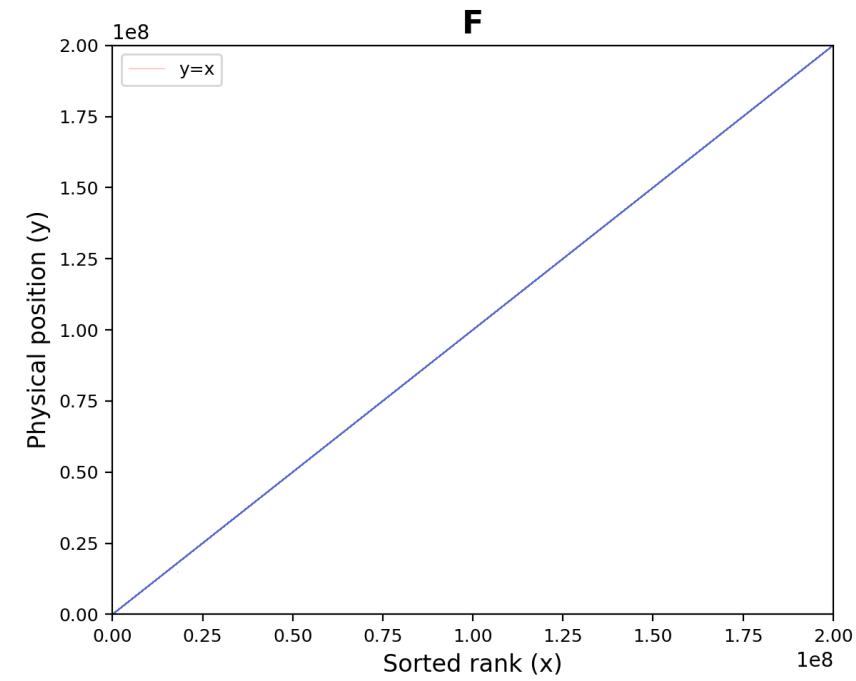
 ↓ FMA + SSE round
 prefetch(data_[pos])

pos = round(slope × rank + intercept) + error

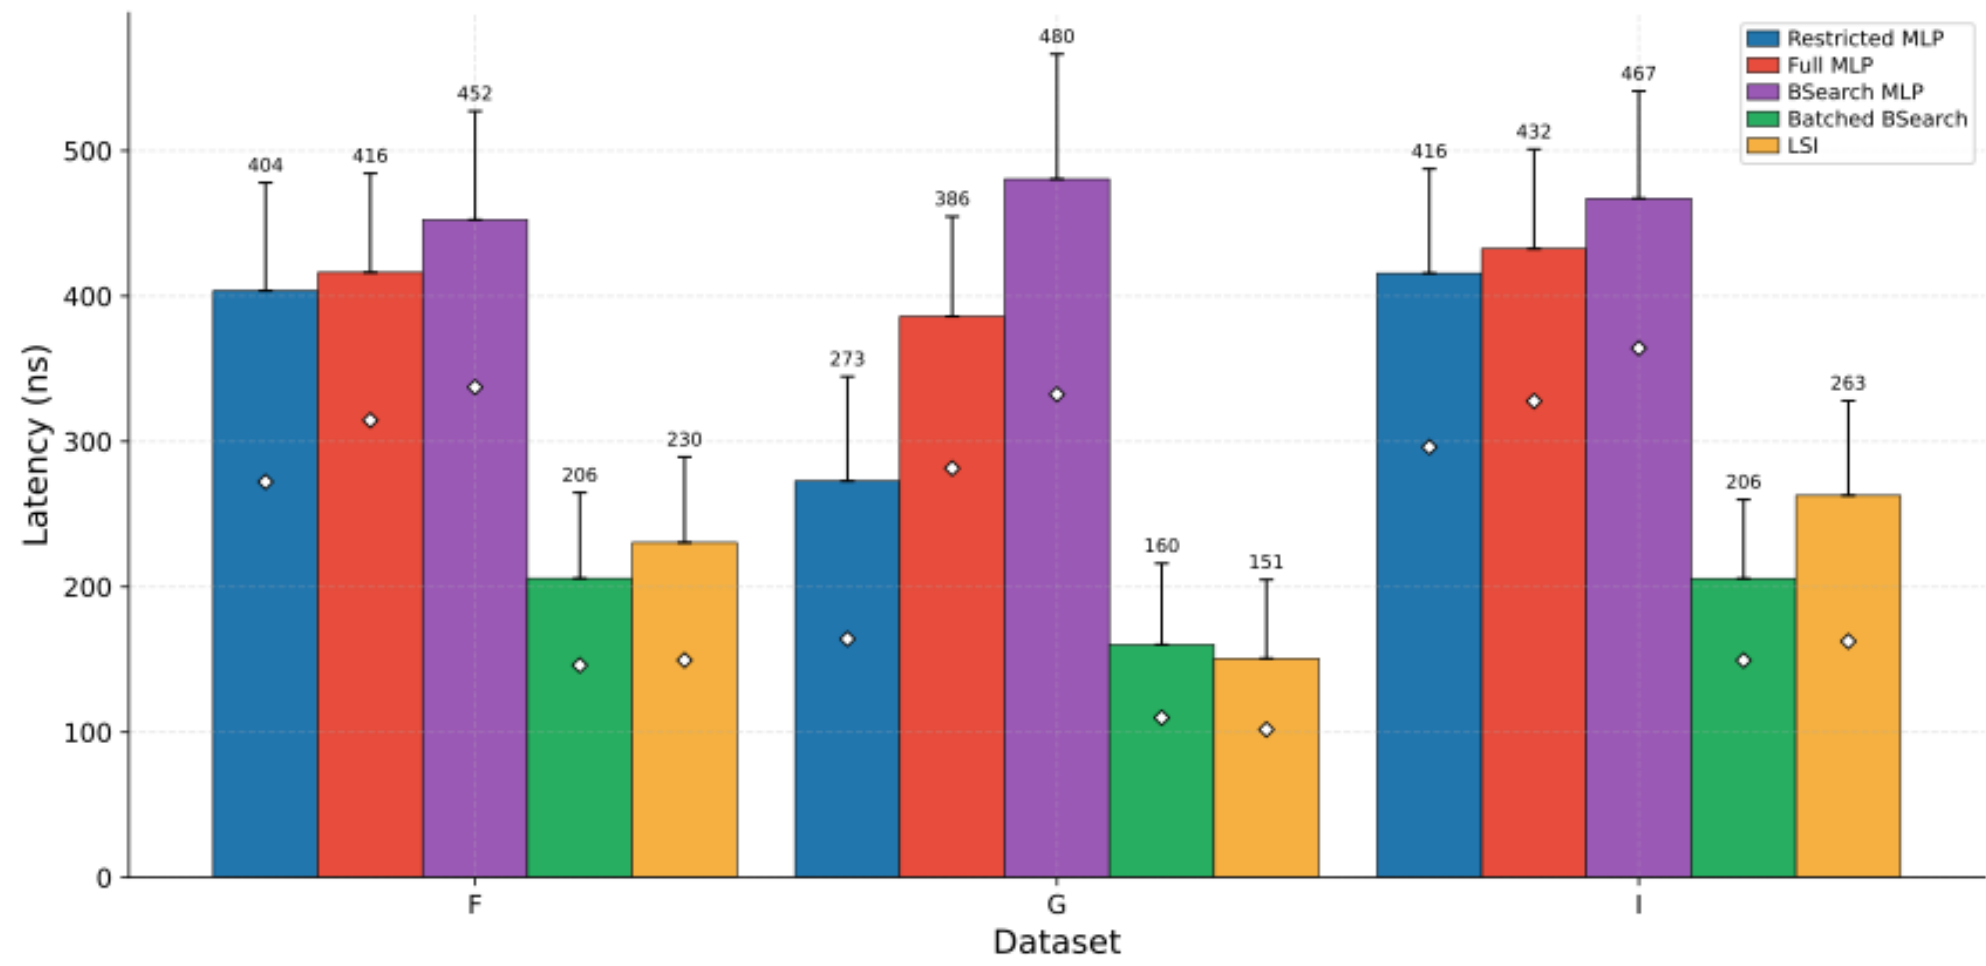
4 Scan: compare data_[pos] vs key
 == key → return pos
 > key → decode left, scan
 < key → decode right, scan

Does this Work?

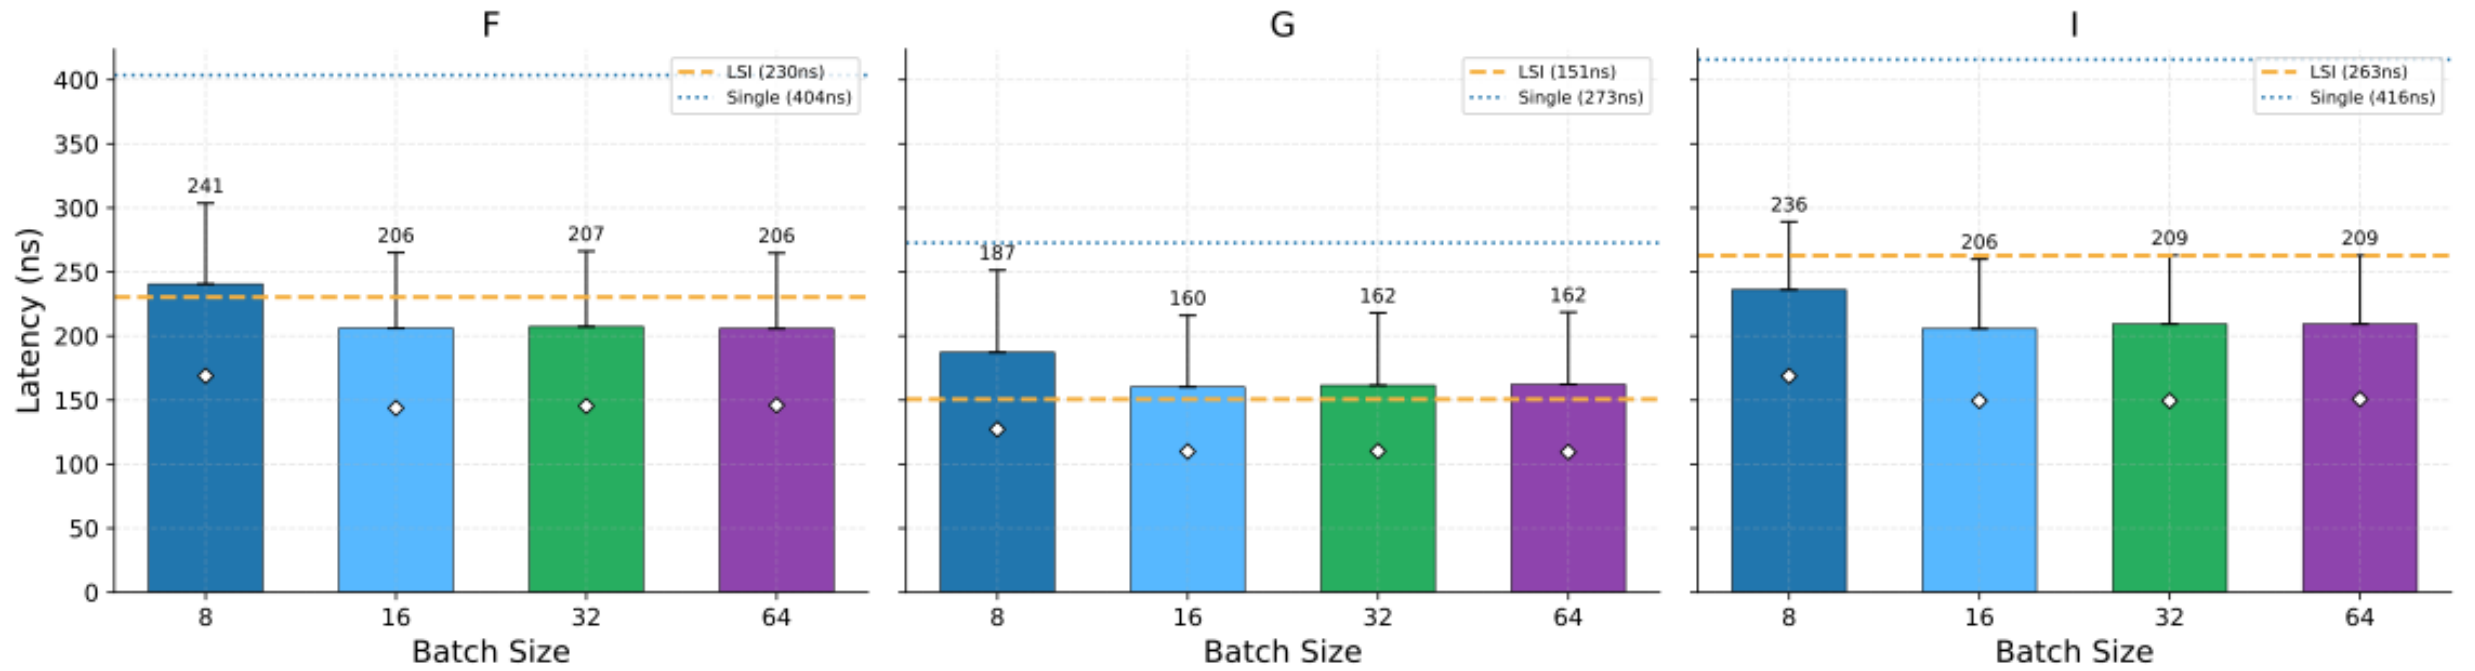
Evaluation Data Set



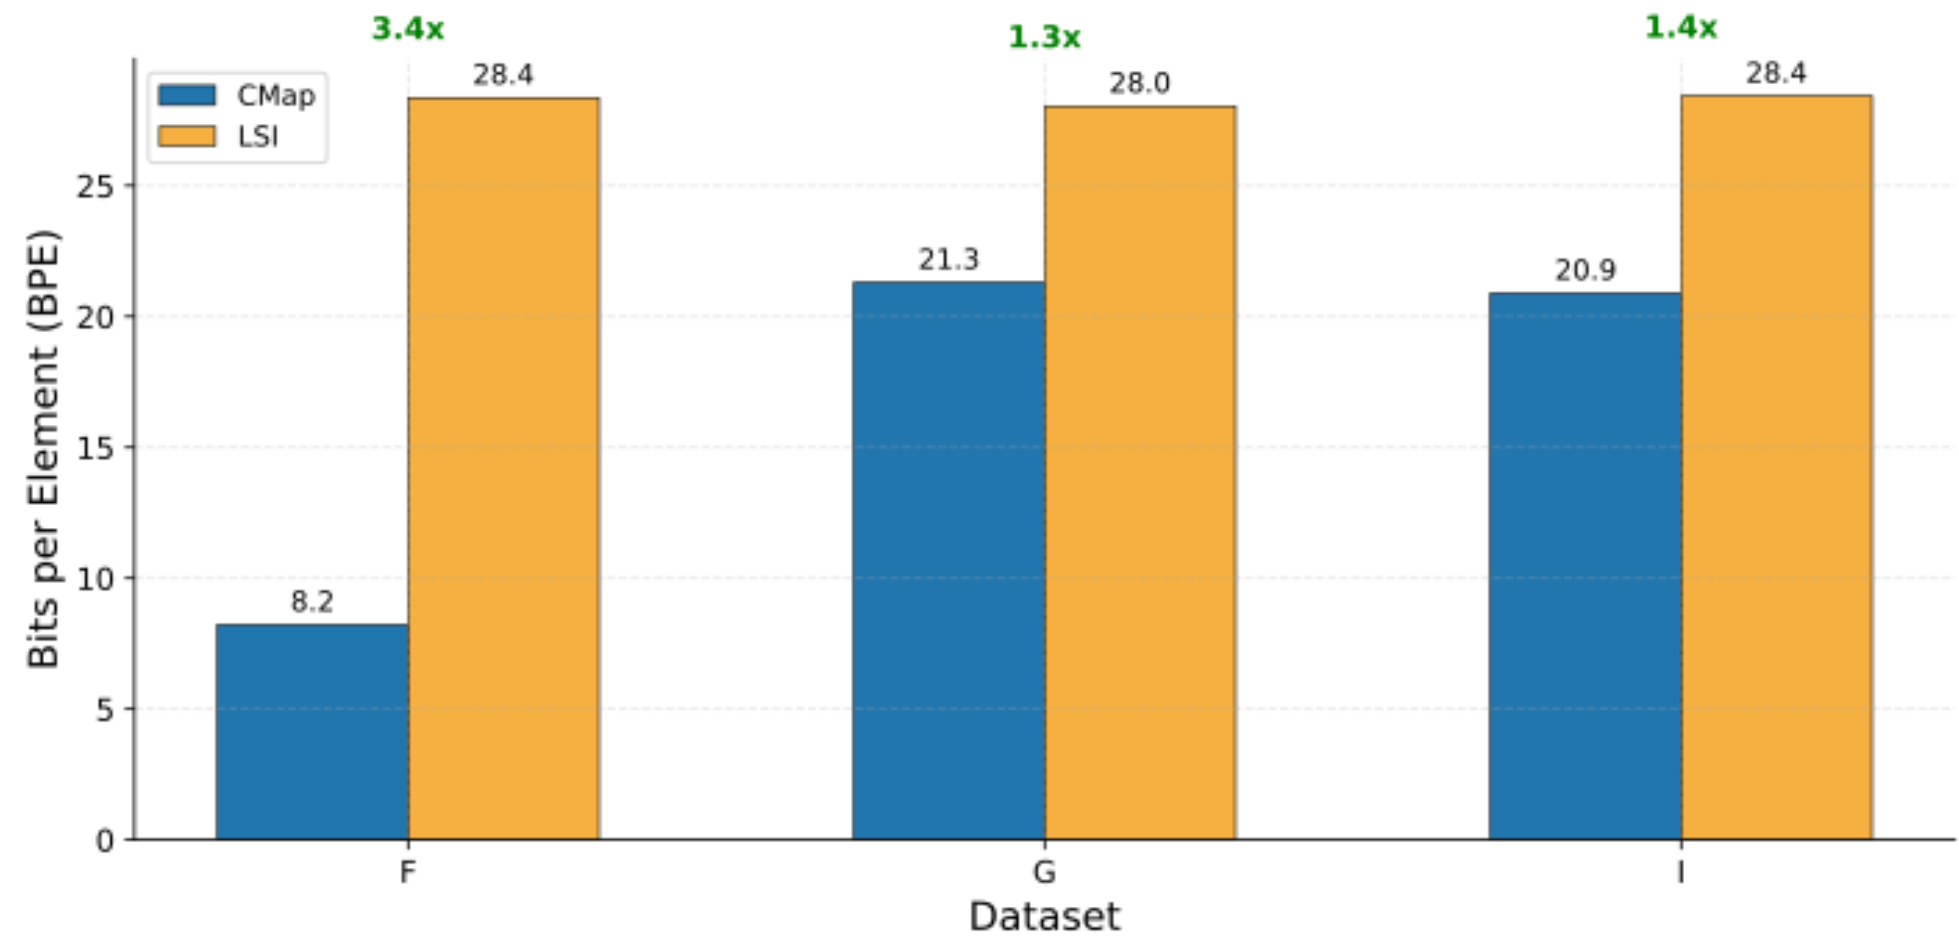
All Methods Combined



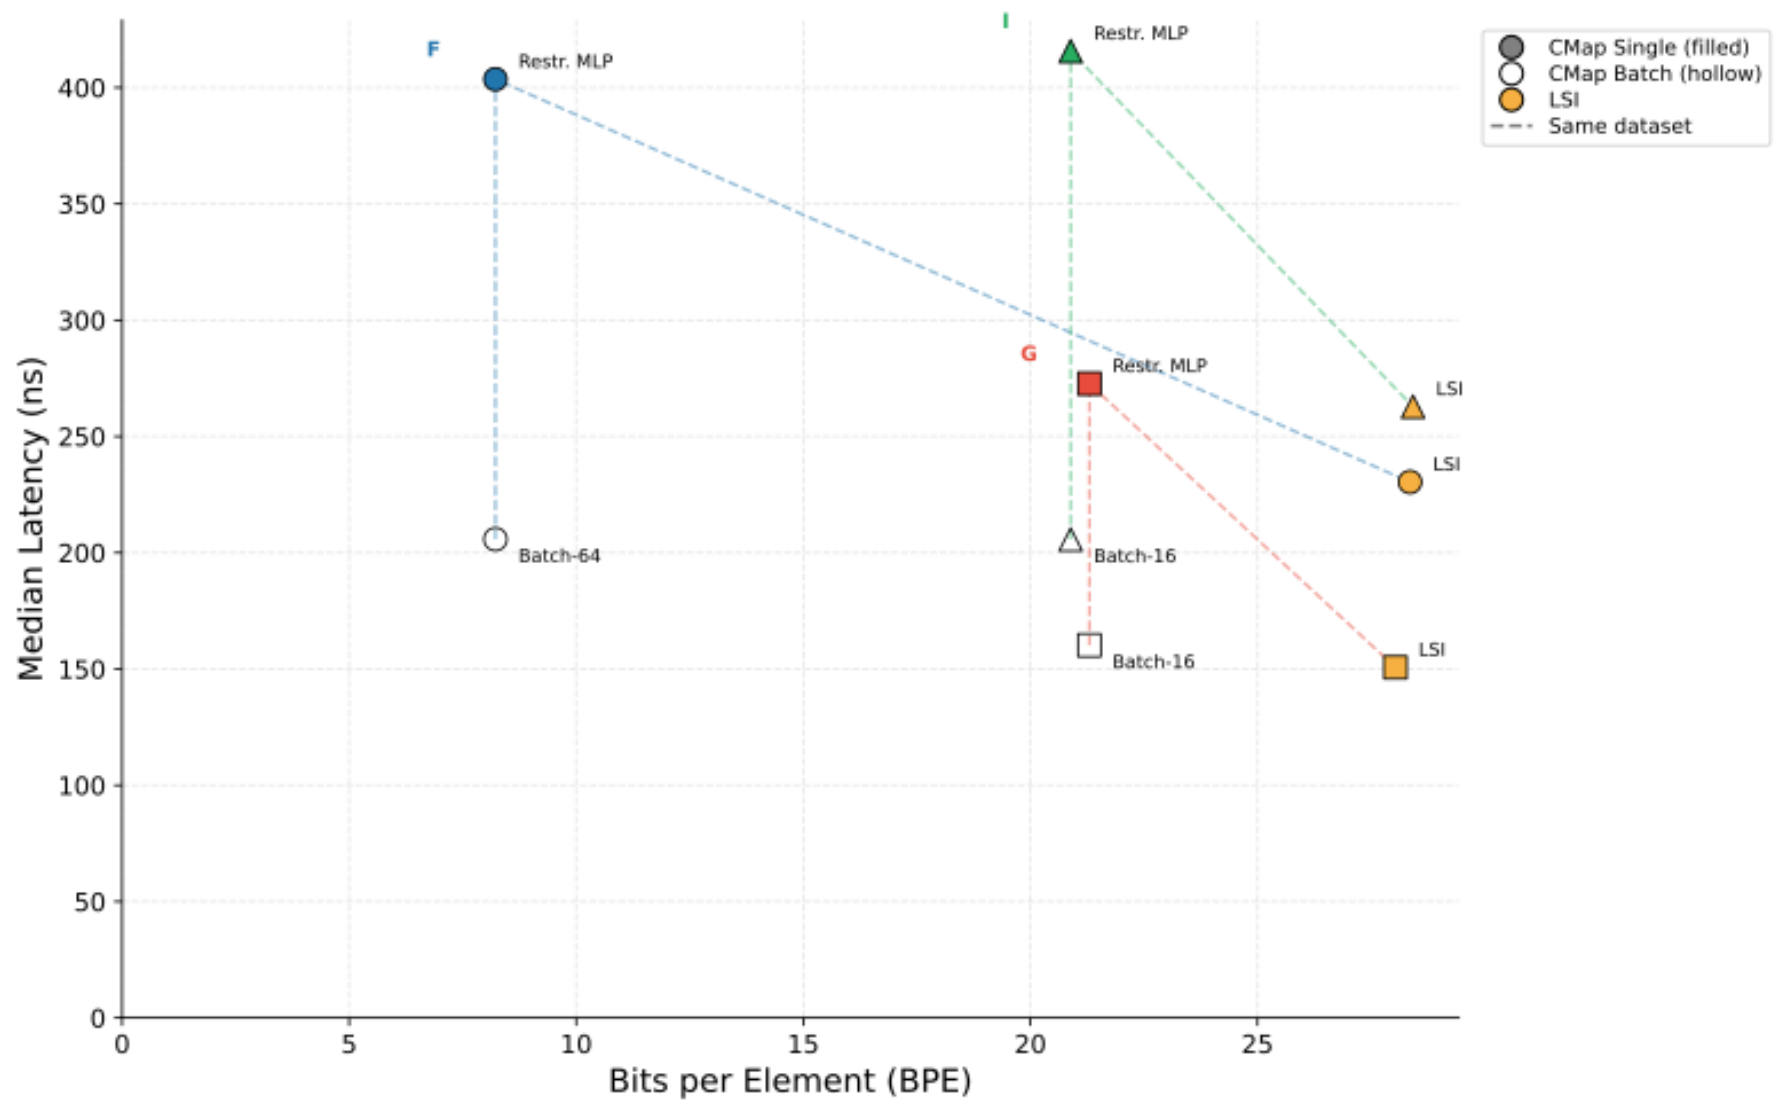
Batched Queries



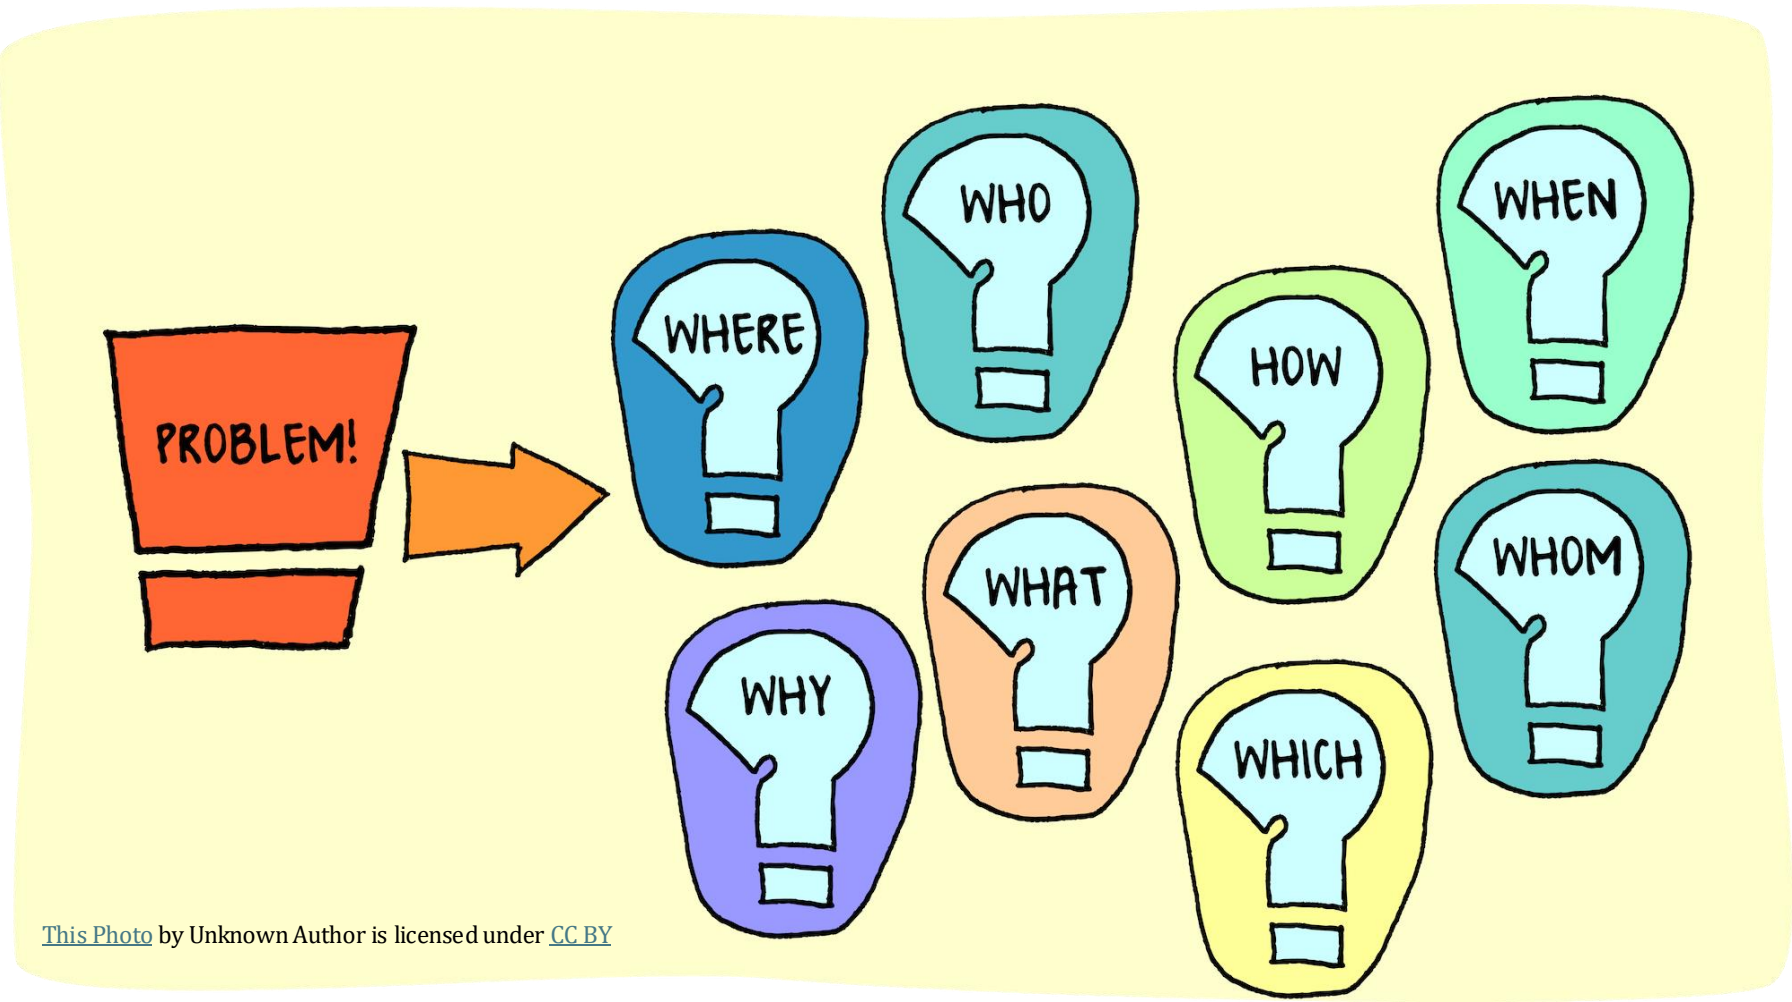
Space Savings



Space-Time Tradeoff



Questions



[This Photo](#) by Unknown Author is licensed under [CC BY](#)