


On the nature of Strings

A deep dive into real world strings and how LiquidCache handles them



Xiangpeng Hao

5th-year PhD@UW-Madison

Database/storage systems

Co-created LiquidCache

Funded by: InfluxData, Bauplan, SpiralDB, tax payers of Wisconsin and federal government

PublicBI Dataset

Real world business intelligence from Tableau

46

Tables

710

SQL queries analyzed

386GB

Data size

String: the most important datatype

61%

Bytes are string

40%

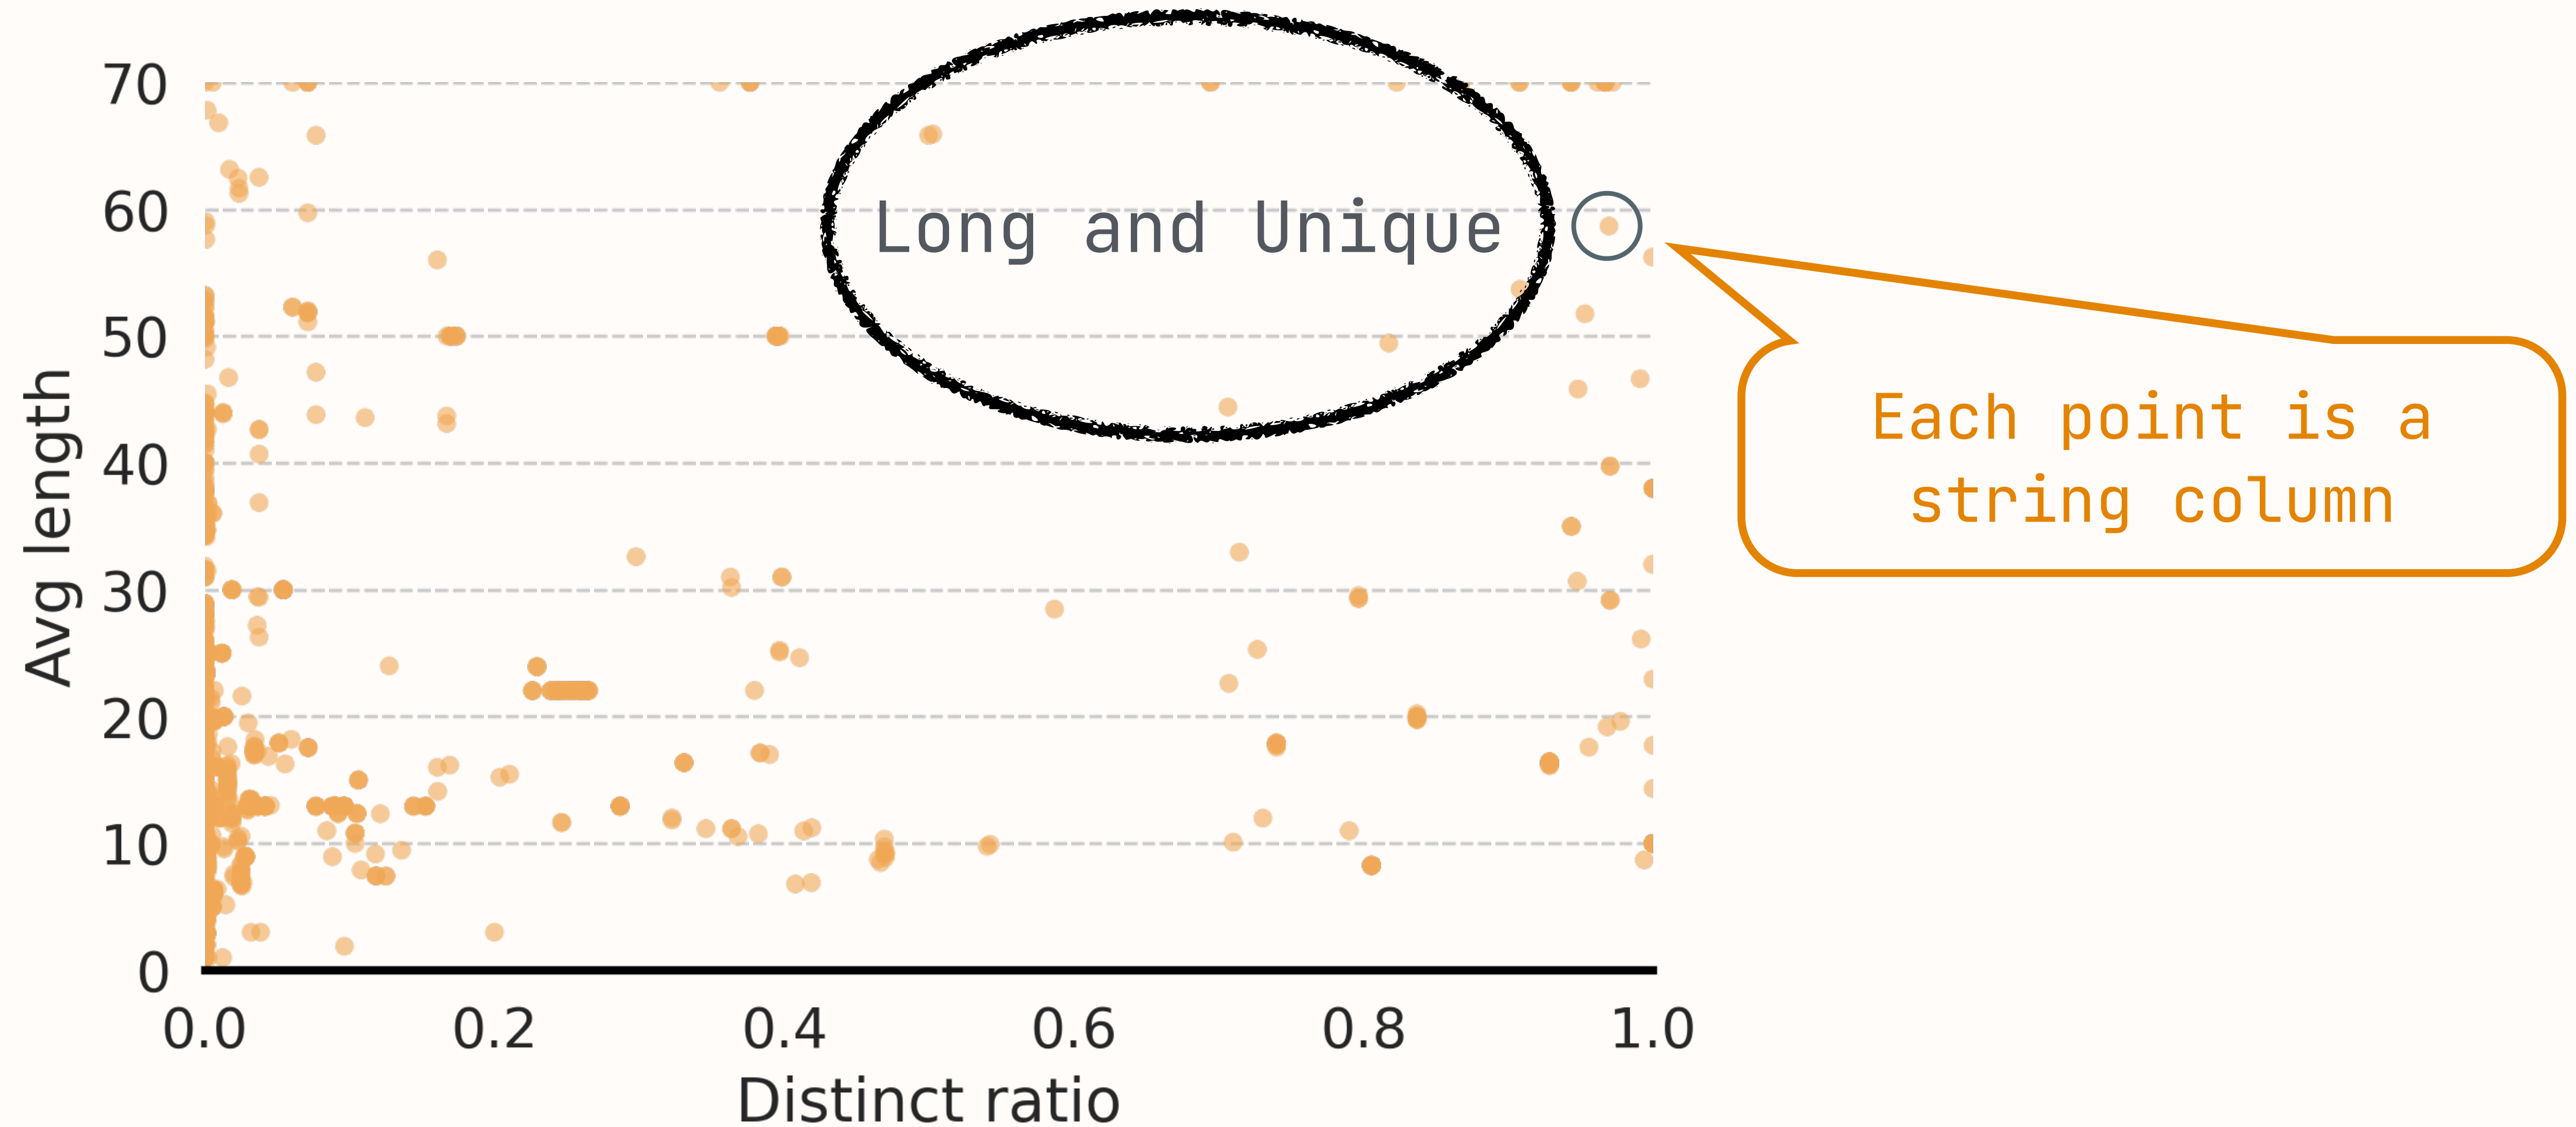
SQL projections
are string

55%

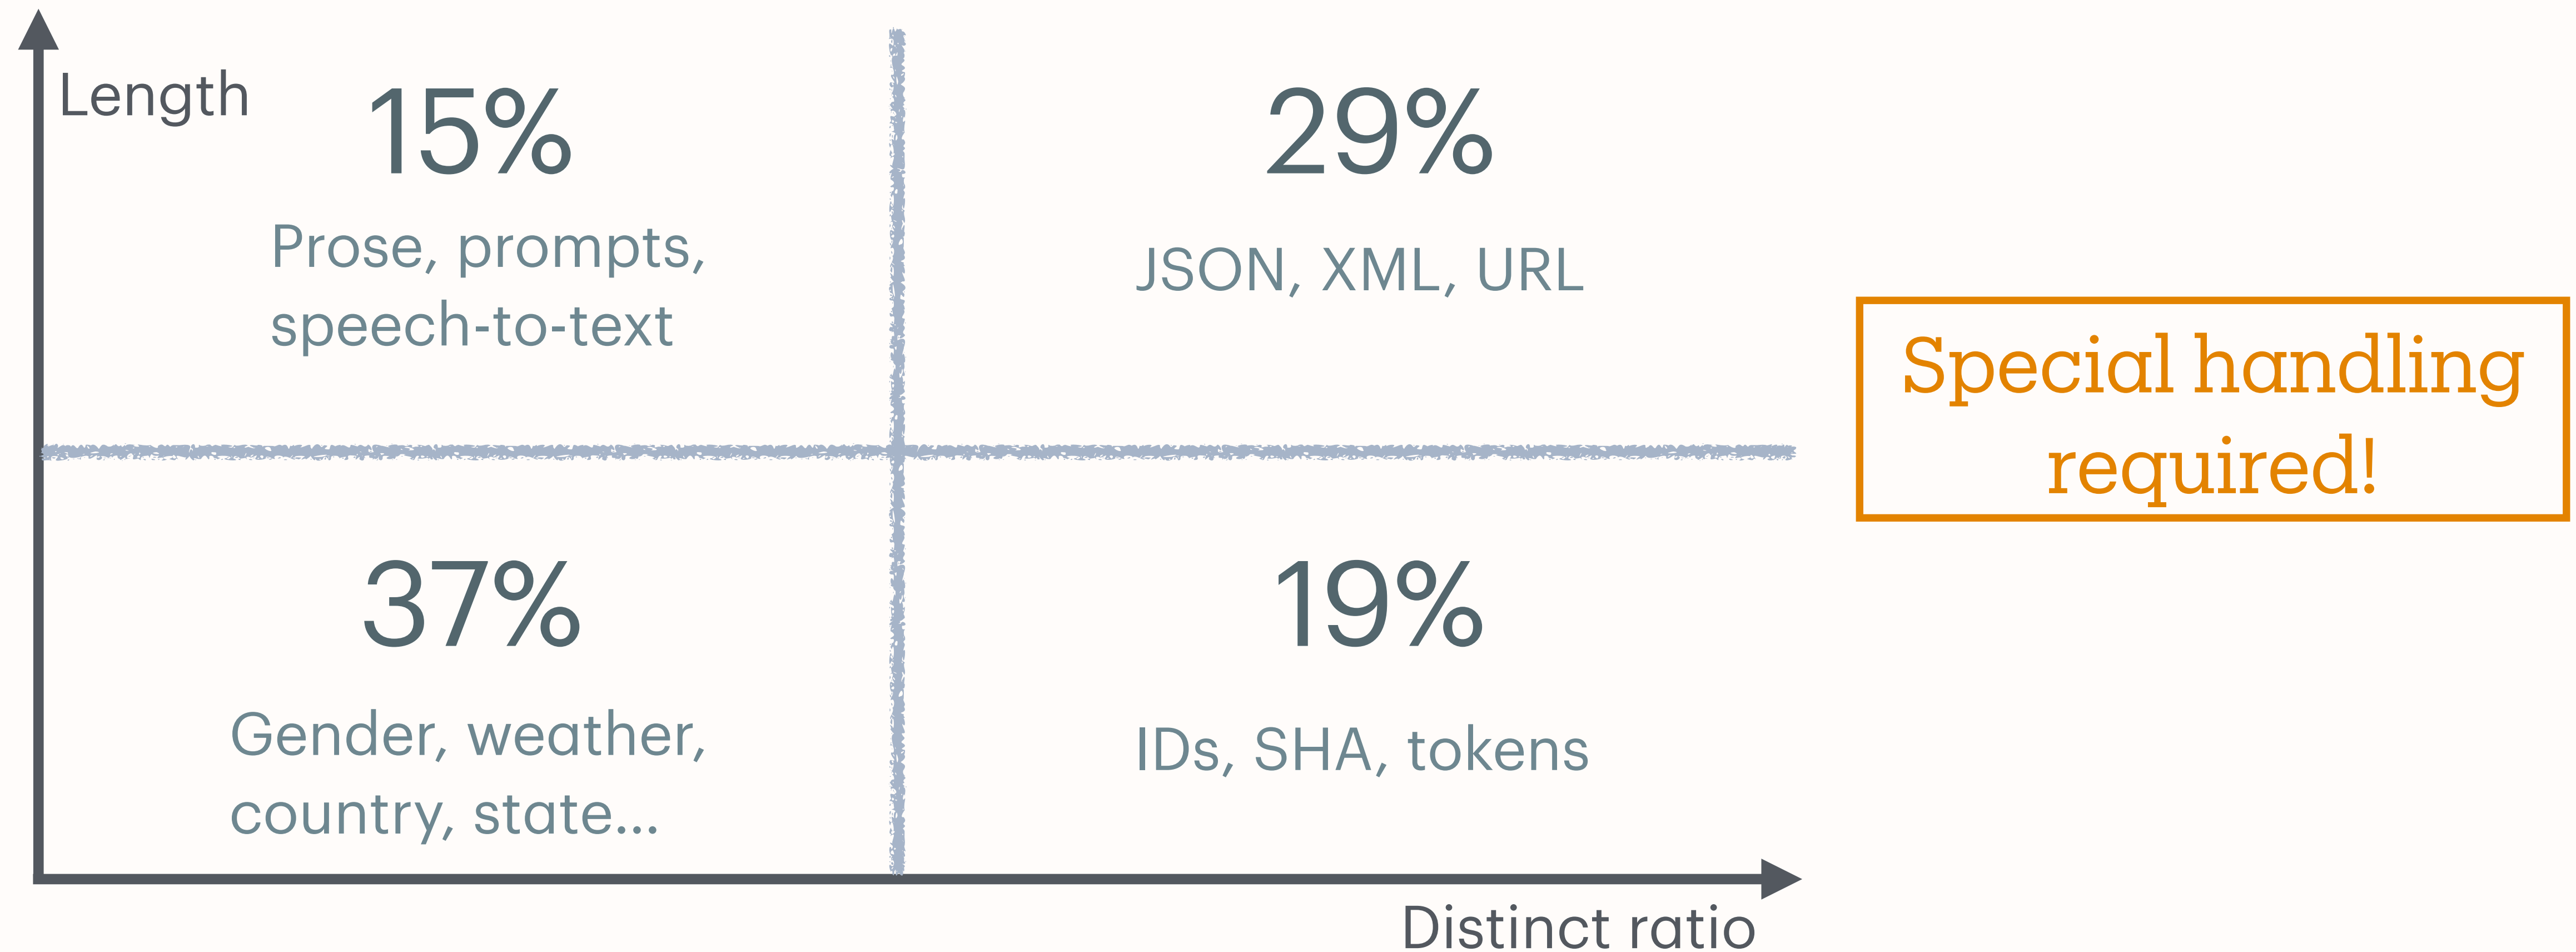
String columns read
at least once

Non-string:34%


Strings are not created equal



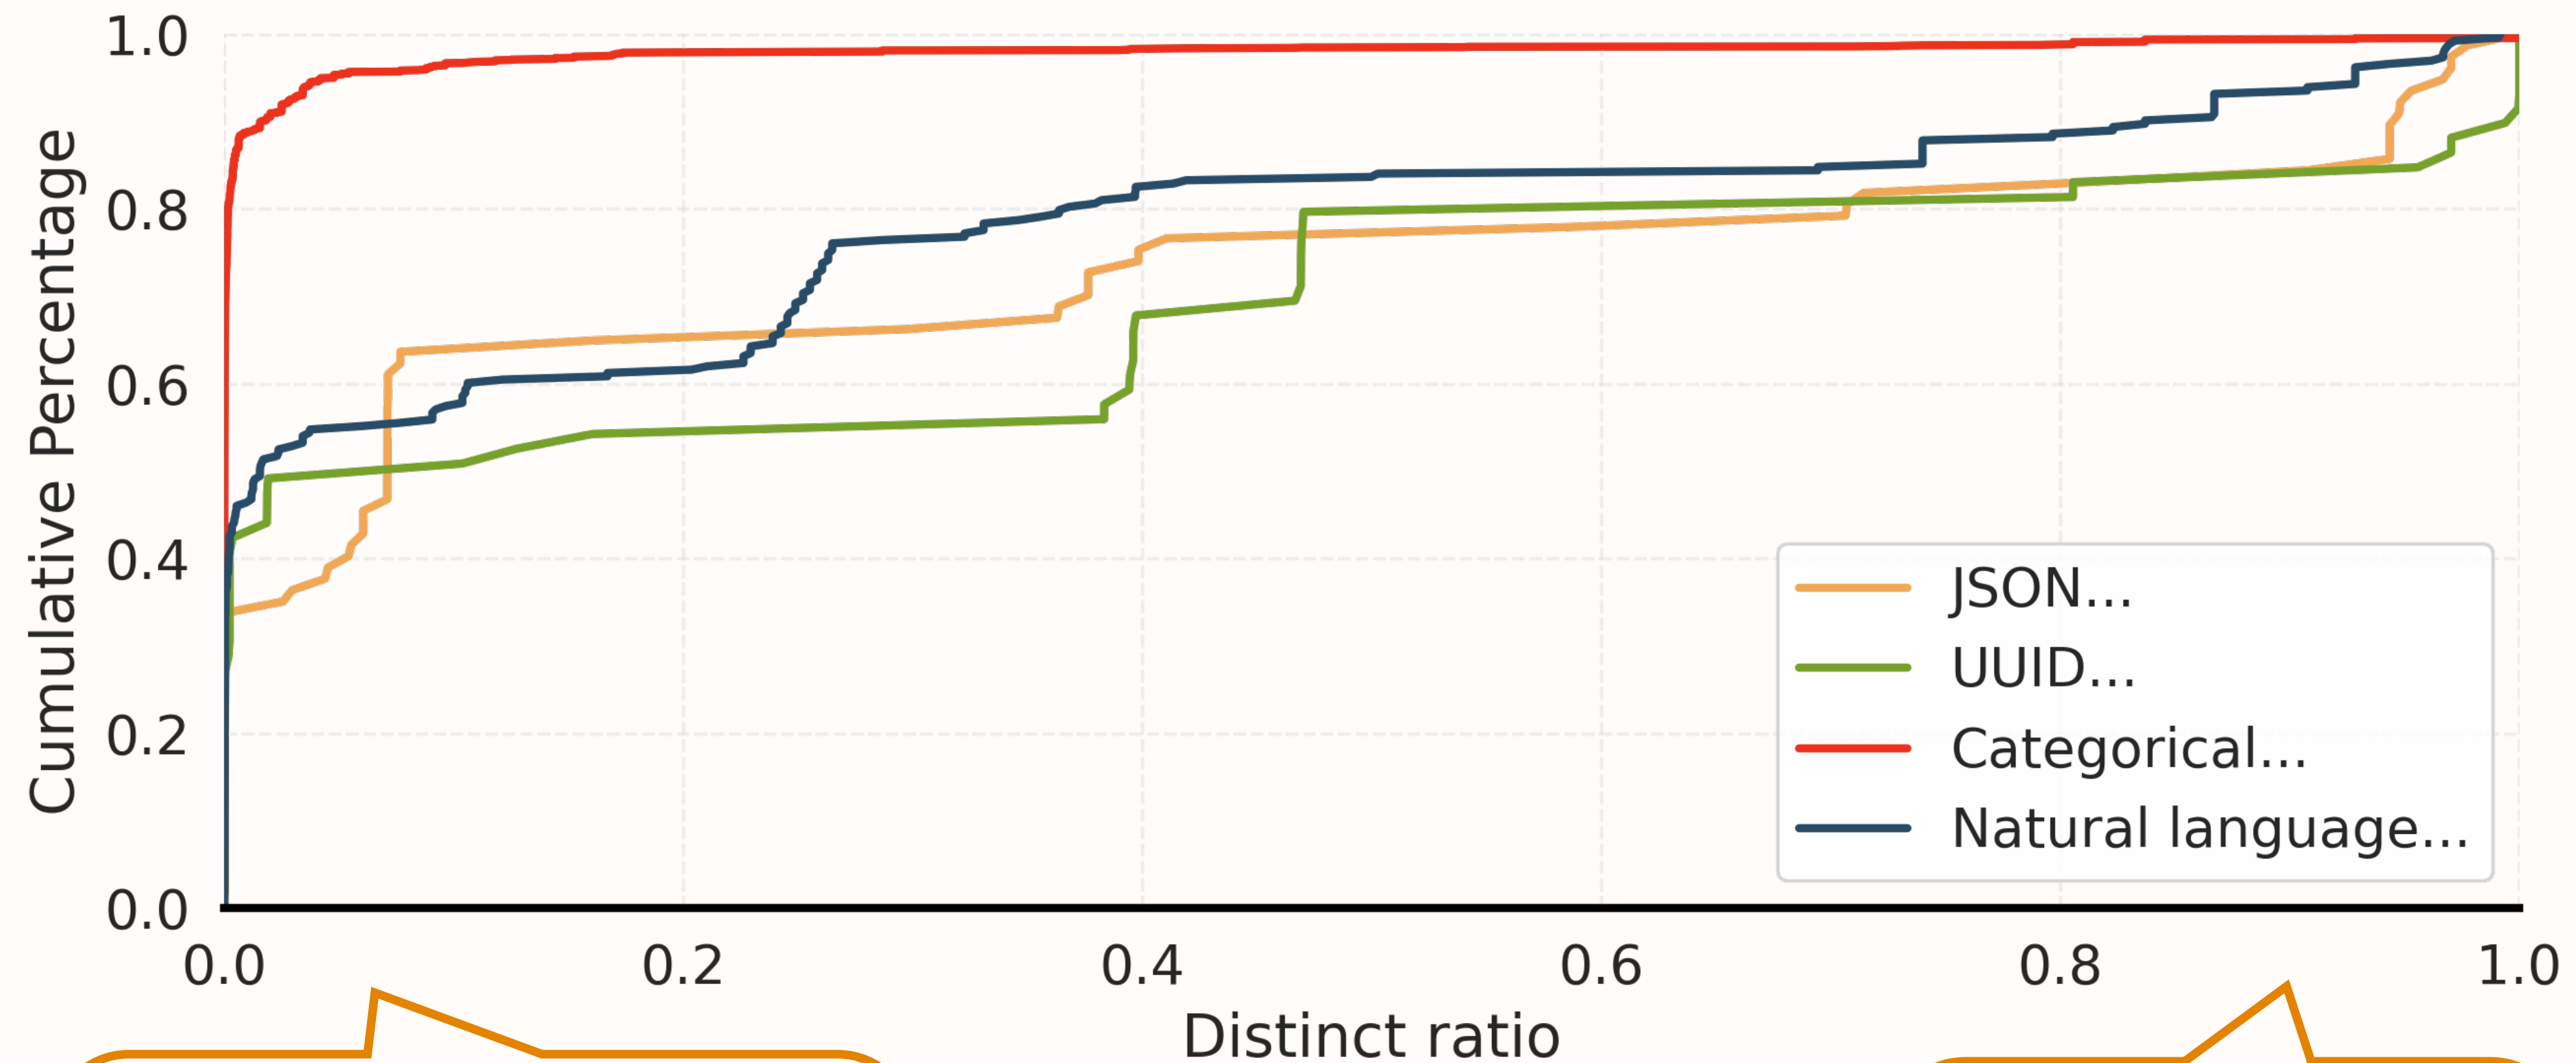
Strings can be categorized into 4 types



Distinct ratio affects encoding



Dictionary coding: optimized for repetitive data



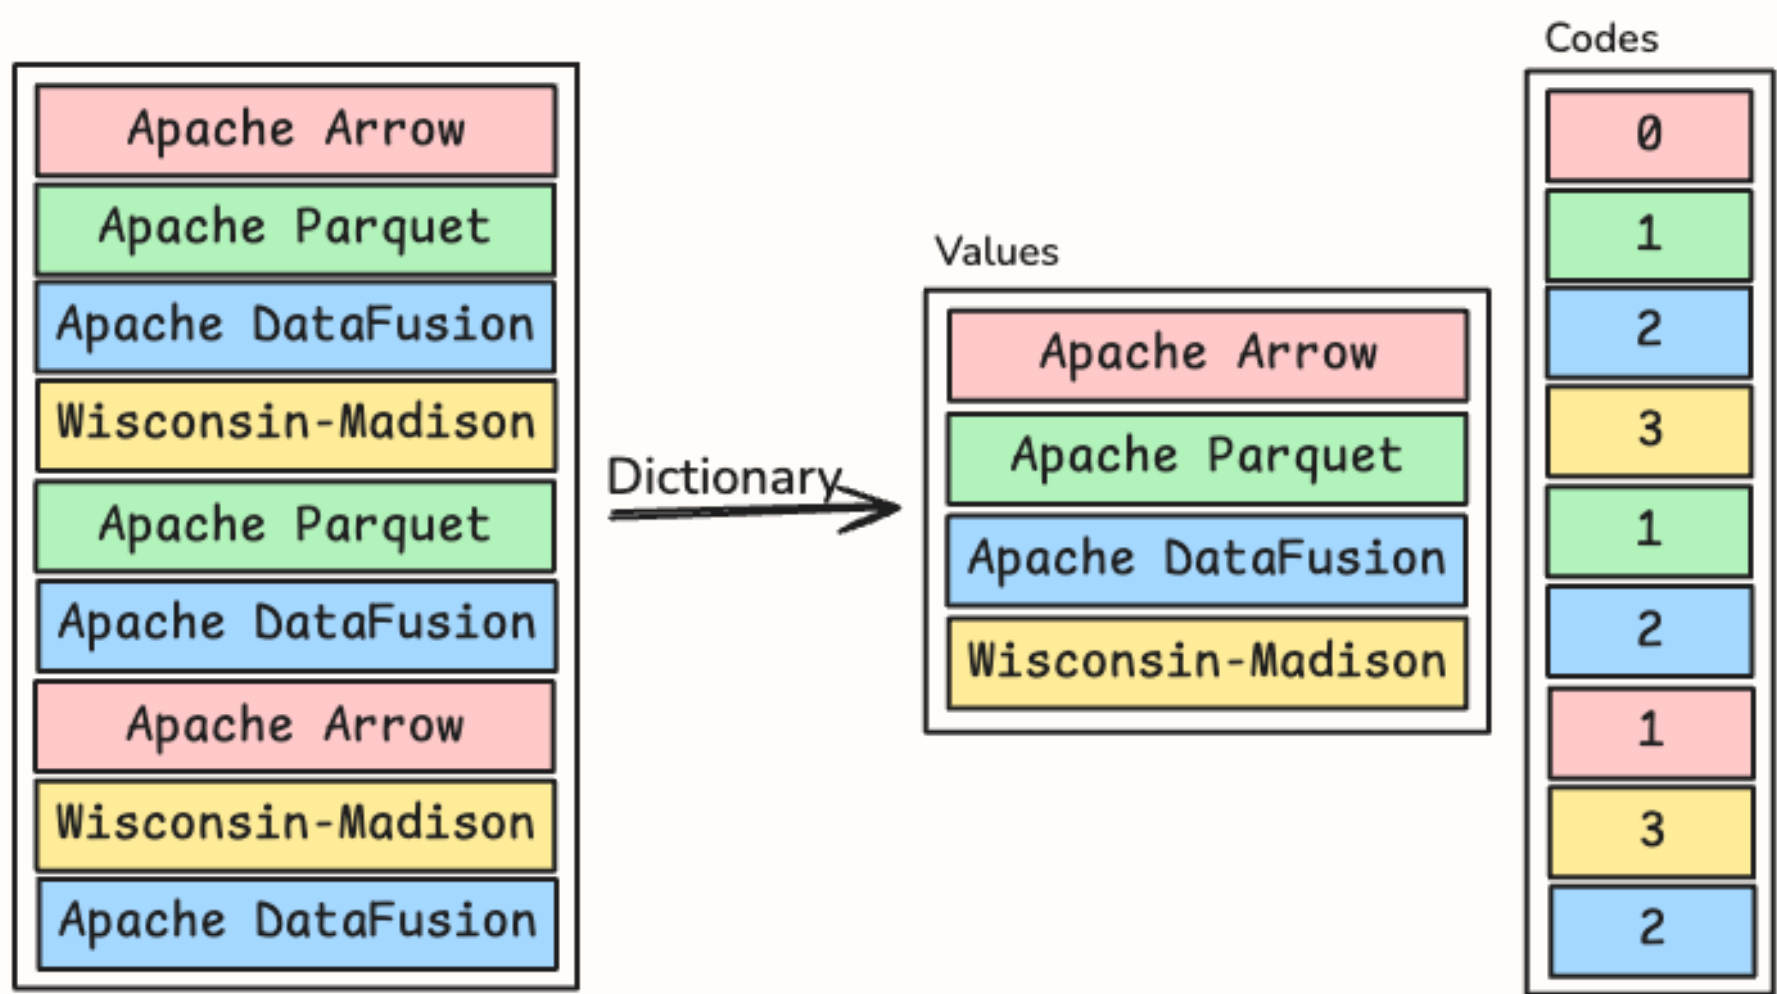
Dictionary encoding

Plain encoding

Dictionary encoding in practice

Distinct ratio

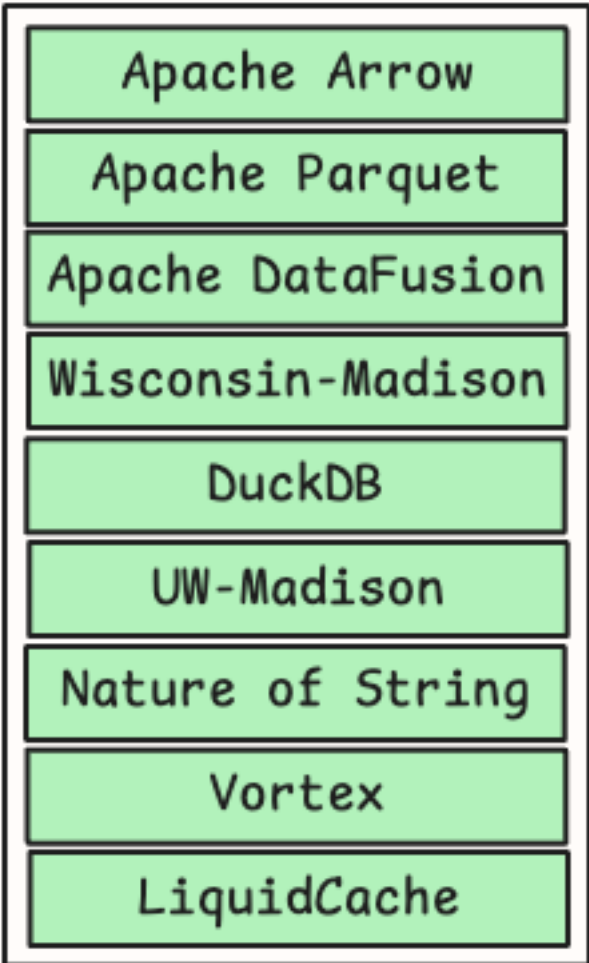
Dictionary encoding



Lower memory footprint

Faster sorting/ comparisons

Plain encoding



Parquet: starts with dictionary encoding, fallback to plain encoding

Dictionary encoding is crucial for string operations

Dictionary encoded data

26%

Went through a filter

Filter on dictionary, reduce
repetitive filter evaluation

19%

Went through a hash
aggregation

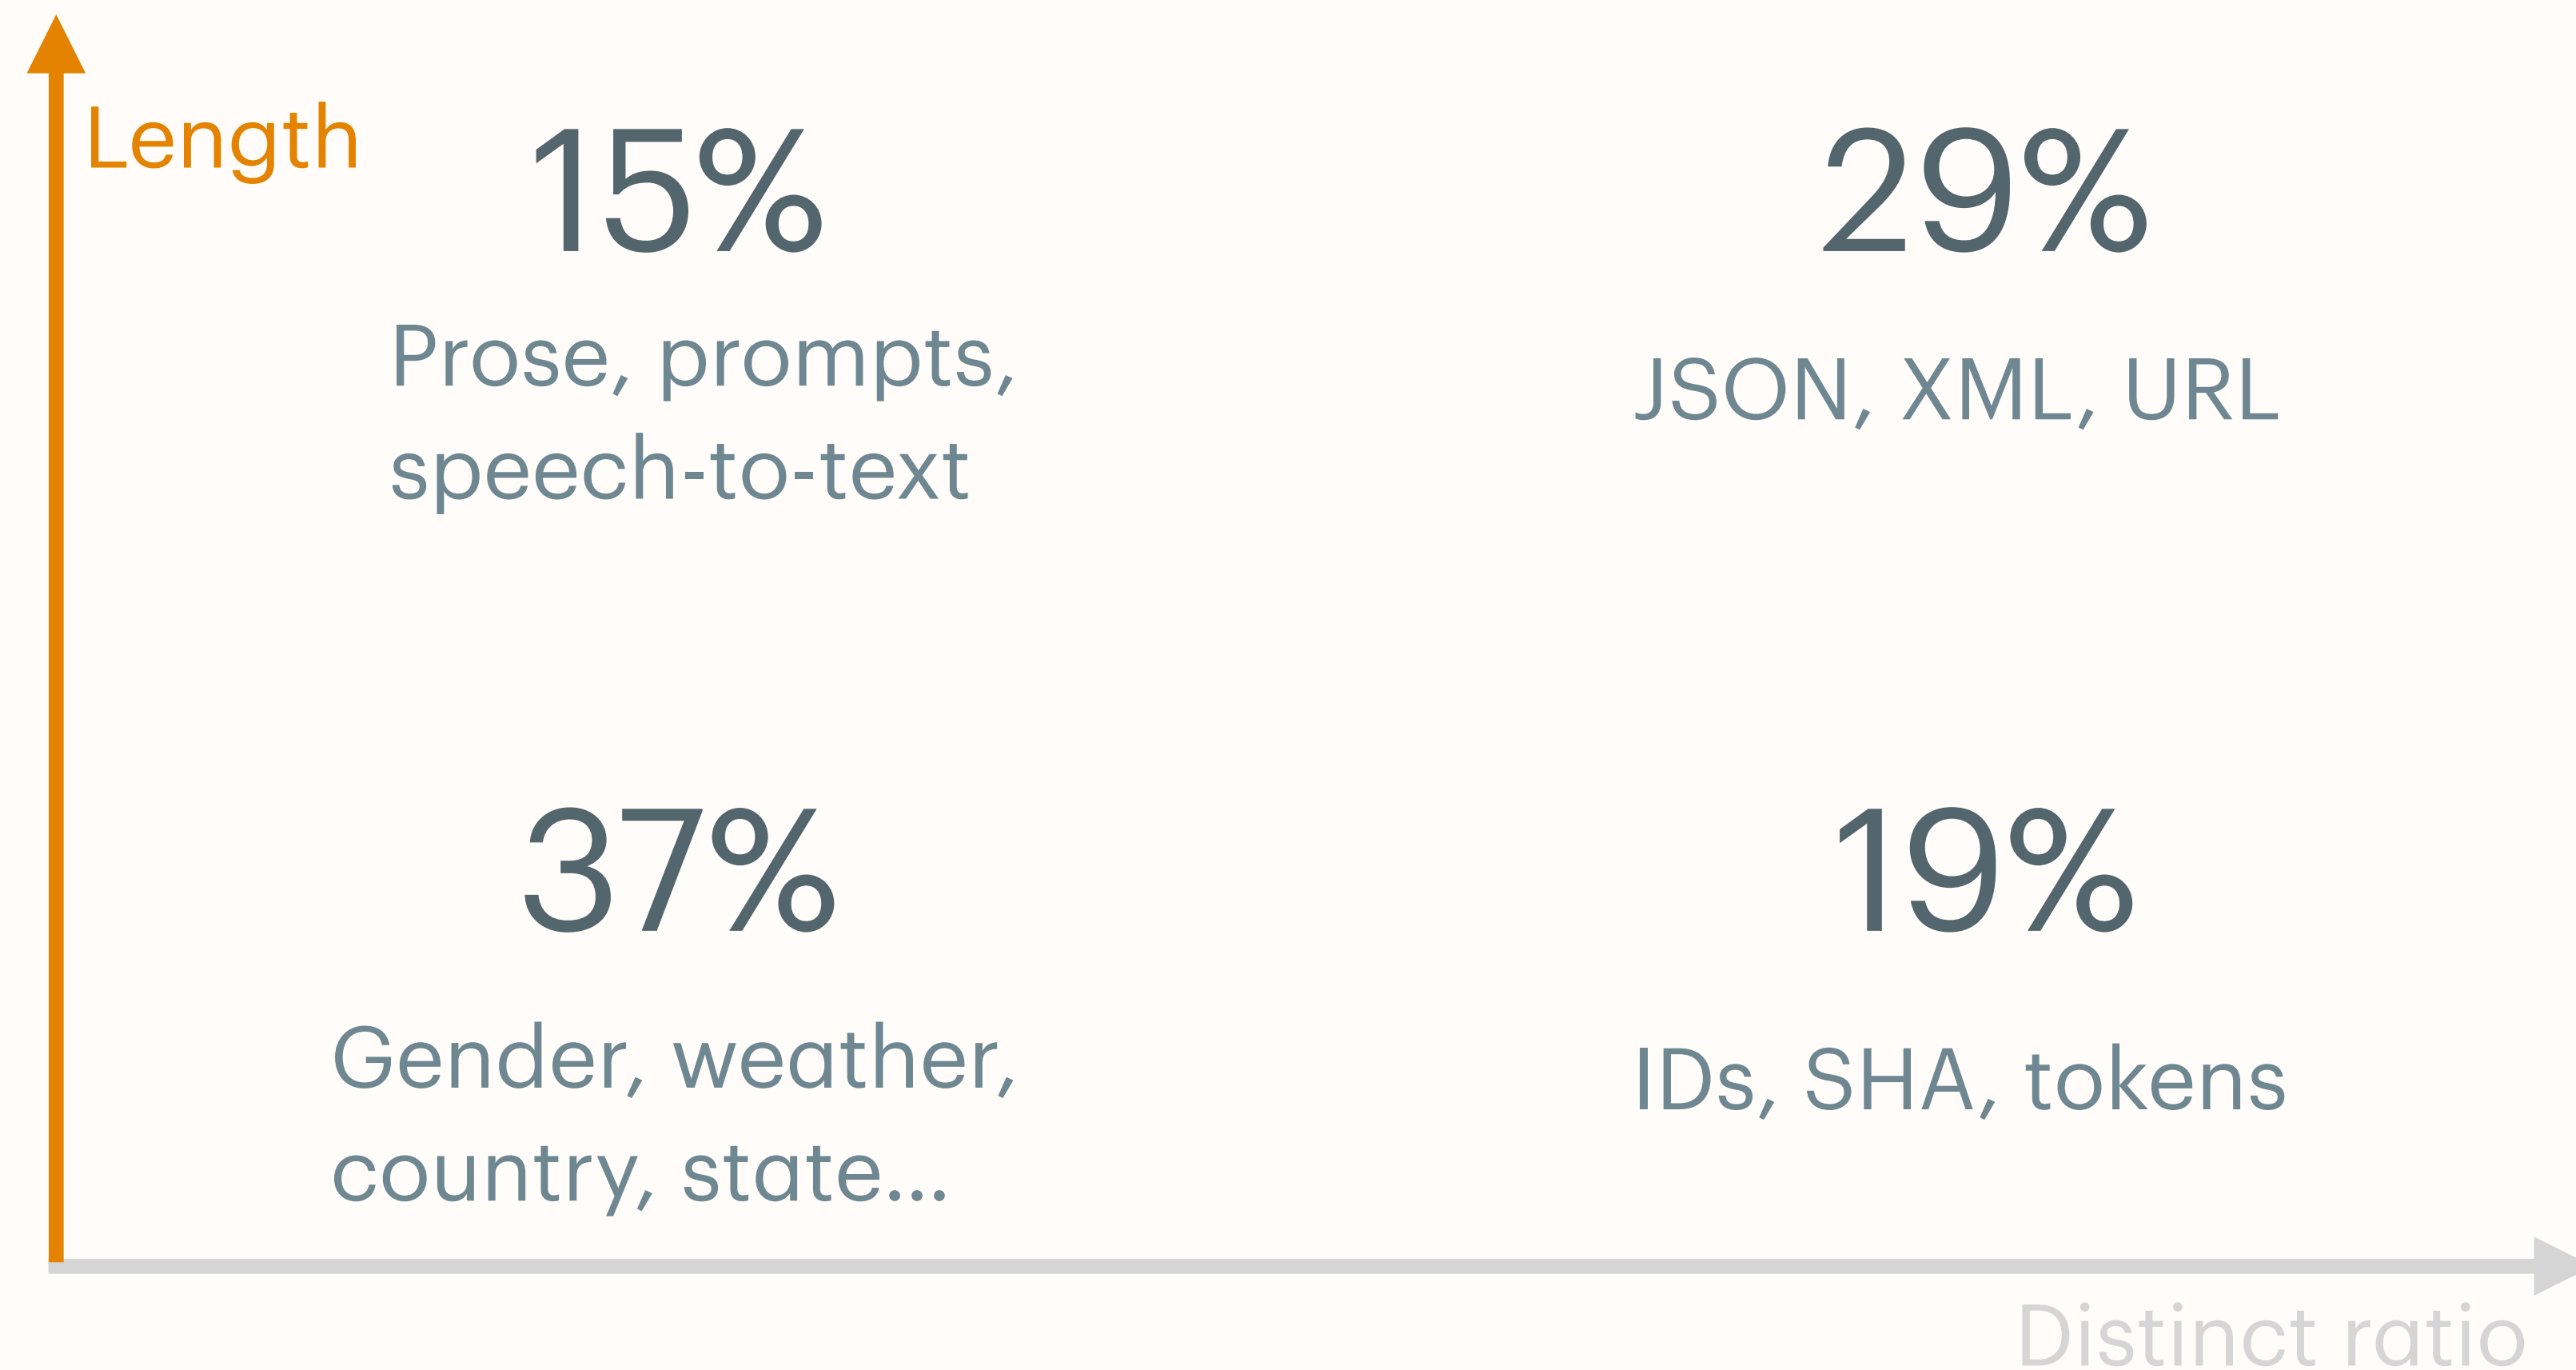
Directly merge distinct values

55%

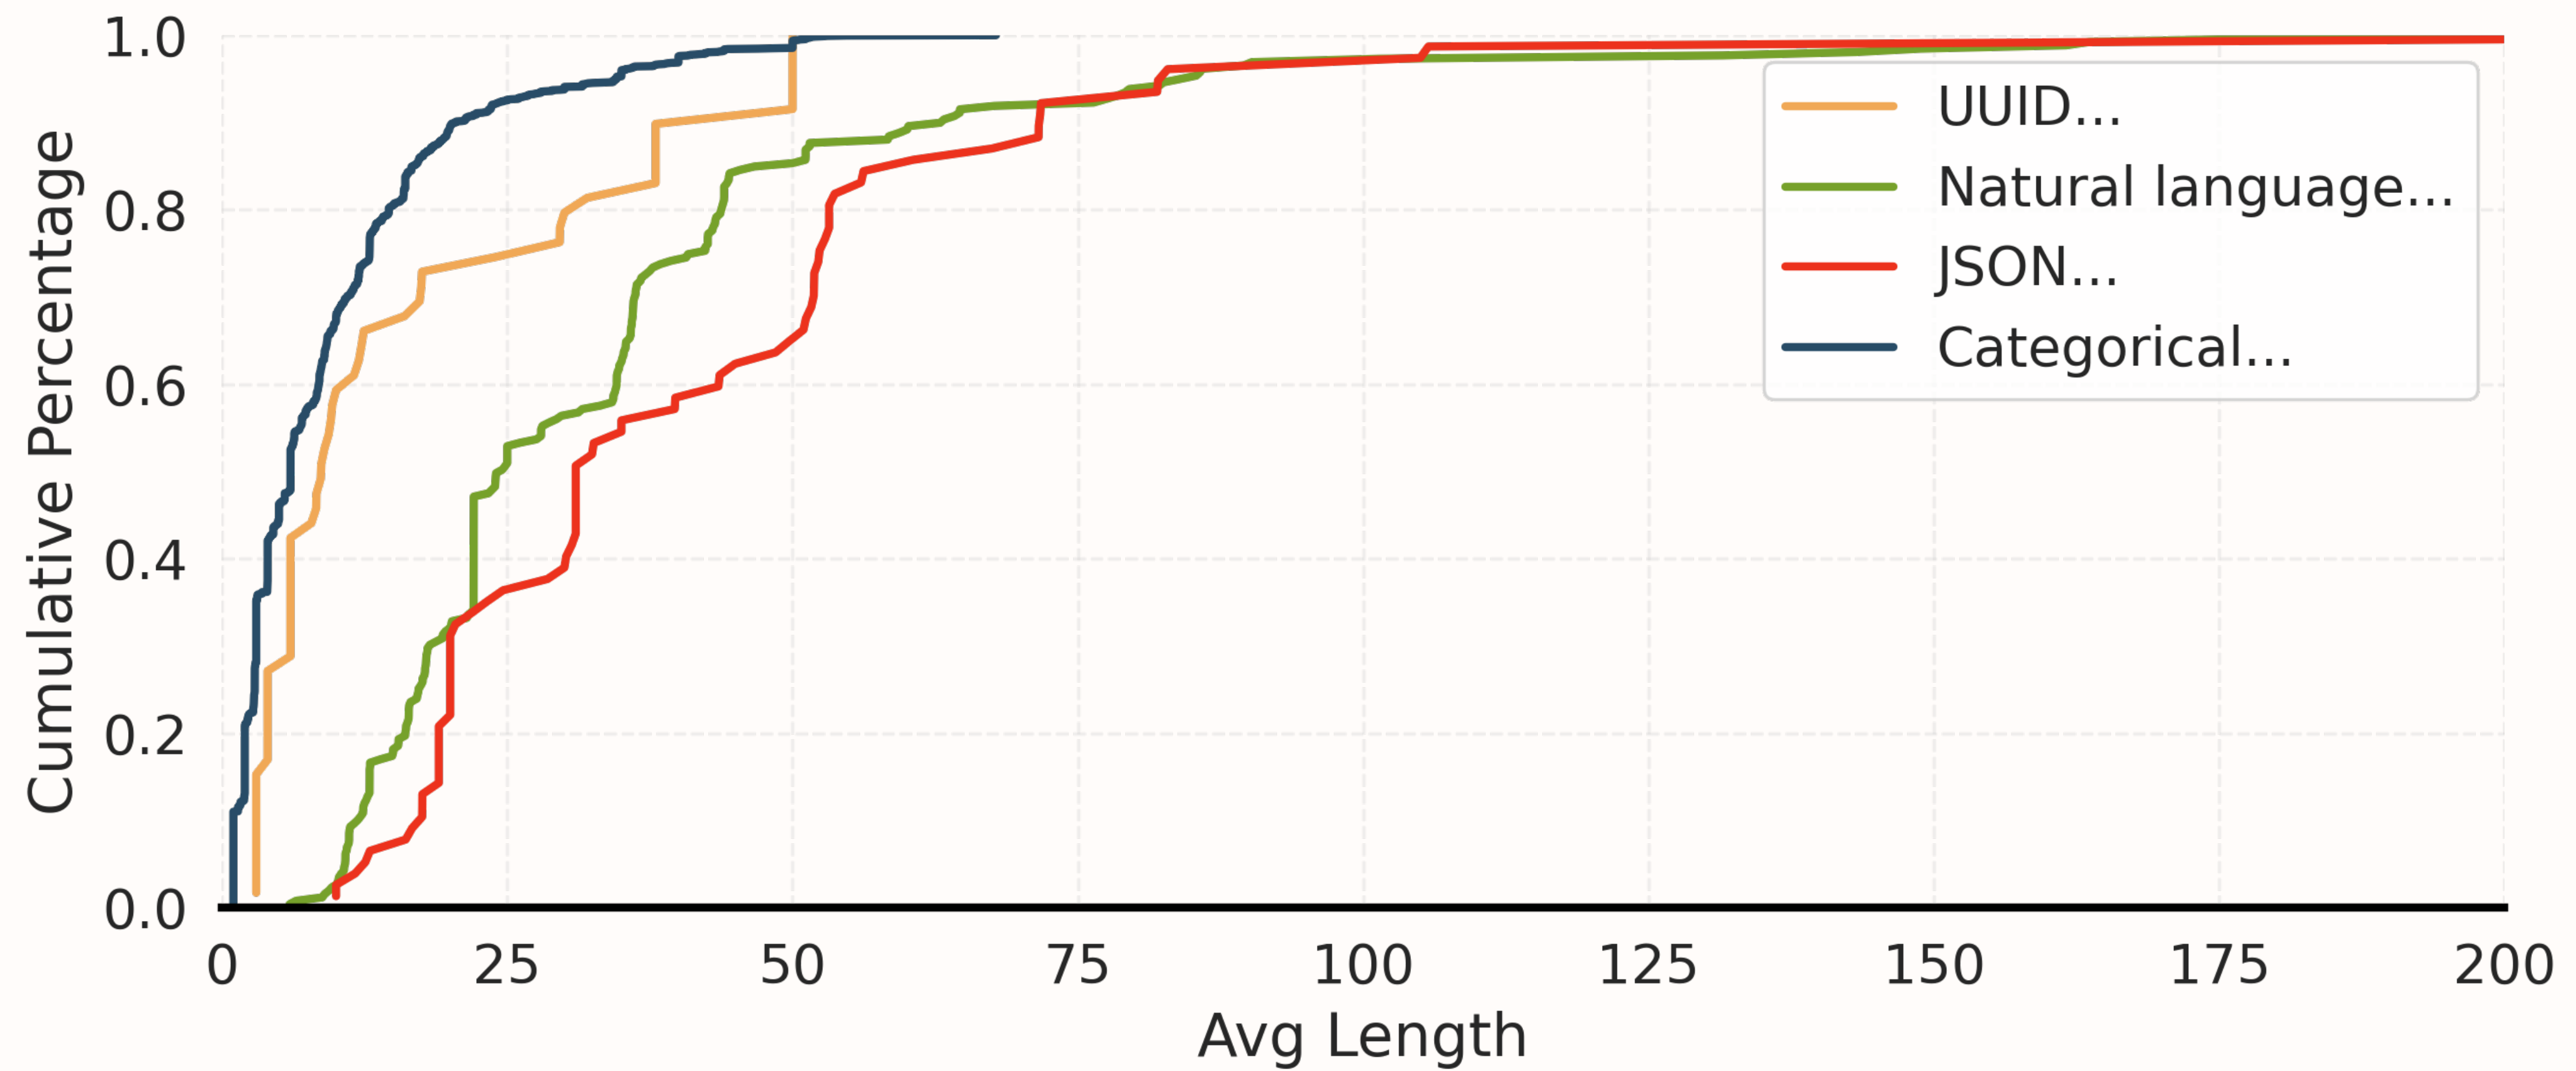
Projection + select

Memory saving, zero-copy of
actual string data

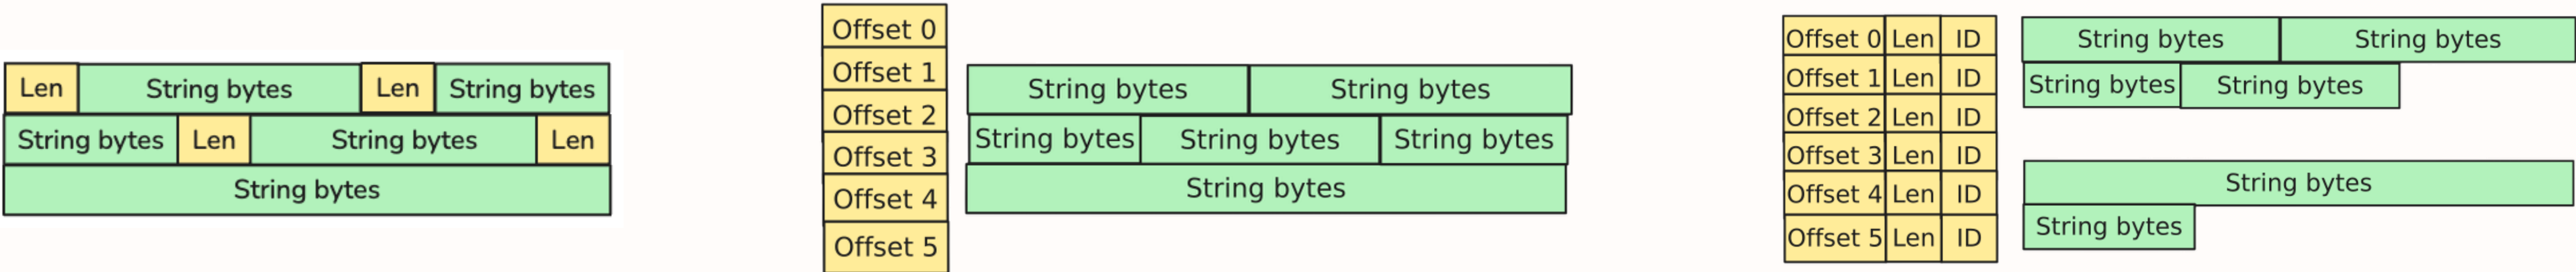
Variable string lengths complicate storage



Variable string lengths complicate storage



How to handle variable-length strings?



Parquet

- + High data locality
- Can't random lookup
- Slow UTF-8 validation

Arrow (StringArray)

- + Simple, efficient slice
- + Random lookup
- Single string buffer

DataFusion/DuckDB

- + Multiple buffer, reduced copy
- Extra memory overhead

Prefix matters

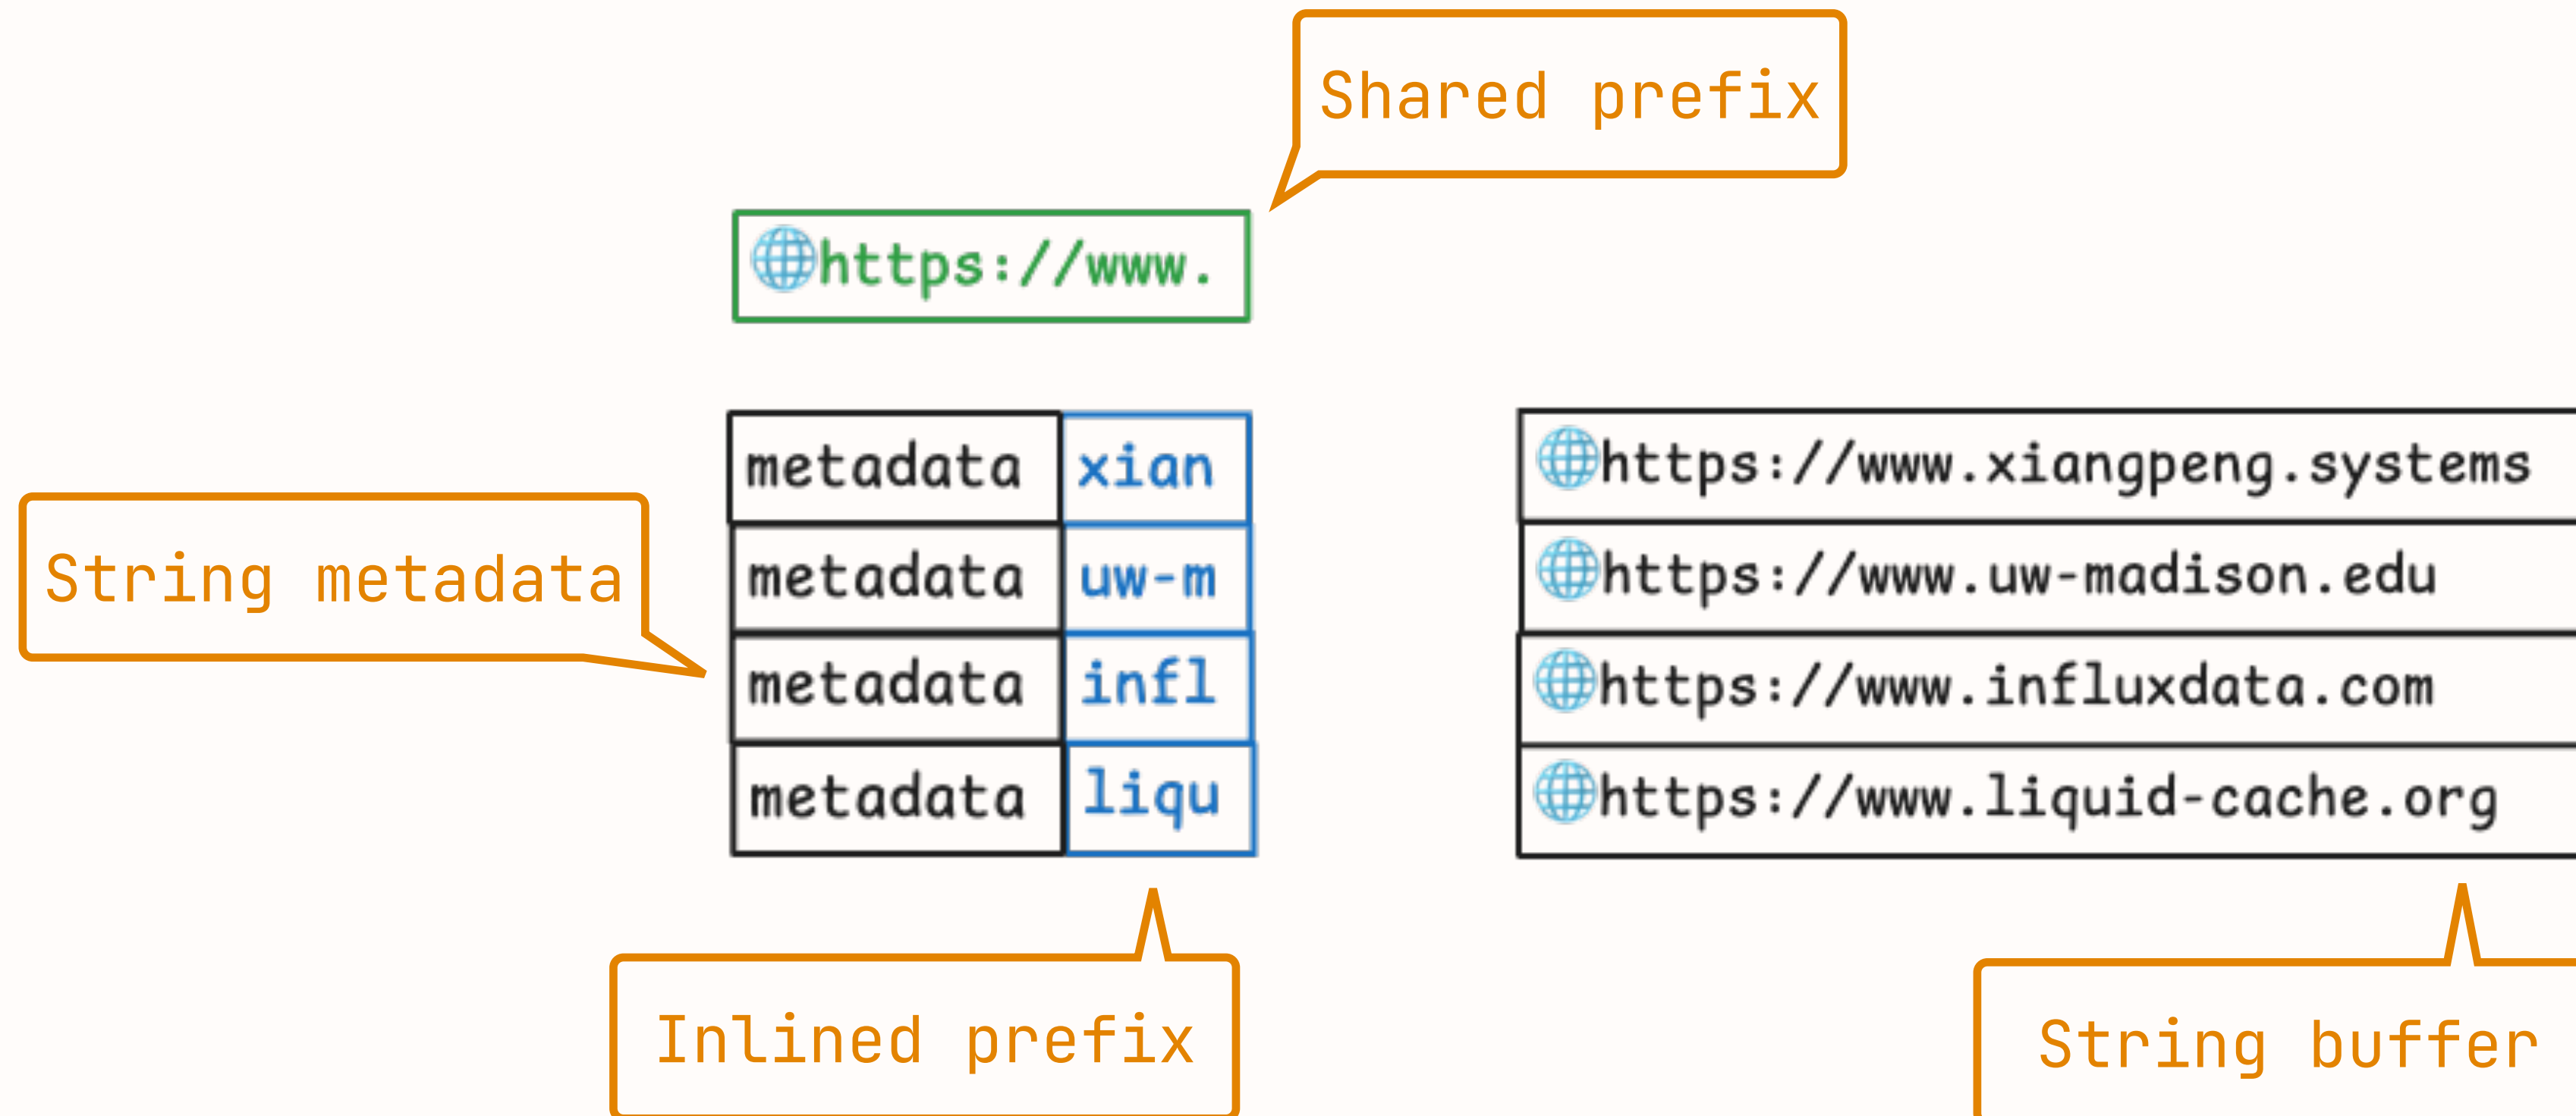
```
#breakforcist
#hasicontent
#monstercakes
$1 Ride
'70s Robot Anime: Geppy-X
'83
'90s Football Stars
'Bhop' Game
'I Want It All' - Queen
'Project: F'
'Saudade'
'Tis the Season...
'n Verlore Verstand
( ° ͜ ° )
($SRCROOT)\dungeon
(04/12) The Life of Erich Zann
(ASMR) Vin Diesel DMing a Game o...
(DEMO) Cupcake: an Apartment Adv...
(DEMO) Project Rose
(DEMO) Star Seeker in: the Secre...
(DEMO) The Art Of Humanizing A P...
(Don't) Make Me Socialize
(Don't) Open Your Eyes
(FREE) Sci-Fi Sex Game - NSFW
```

Distinct prefixes
accelerates
comparisons

Common prefixes
improve
compression

```
https://img.itch.zone/aW1nLzI0OT...
https://img.itch.zone/aW1nLzIyOD...
https://media.rawg.io/media/games/000/0002...
https://media.rawg.io/media/games/000/0005...
https://media.rawg.io/media/games/000/000db...
https://media.rawg.io/media/games/000/000dc...
https://media.rawg.io/media/games/001...
https://media.rawg.io/media/games/002/00242...
https://media.rawg.io/media/games/002/00247...
https://media.rawg.io/media/games/003/0031...
https://media.rawg.io/media/games/003/0033...
https://media.rawg.io/media/games/003/0034...
https://media.rawg.io/media/games/004/0040...
https://media.rawg.io/media/games/004/0041...
https://media.rawg.io/media/games/004/0042...
https://media.rawg.io/media/games/004/004c...
https://media.rawg.io/media/games/006/0069...
https://media.rawg.io/media/games/006/006b...
https://media.rawg.io/media/games/006/006d...
https://media.rawg.io/media/games/007/0073...
https://media.rawg.io/media/games/007/0079...
https://media.rawg.io/media/games/007/007c2...
https://media.rawg.io/media/games/007/007cf...
https://media.rawg.io/media/games/007/007d...
```





String representation in action




How many bytes to inline?

 https://www.

metadata	xian
metadata	uw-m
metadata	infl
metadata	liqu

 https://www.xiangpeng.systems
 https://www.uw-madison.edu
 https://www.influxdata.com
 https://www.liquid-cache.org

Inlined prefix







DataFusion/DuckDB: 4 bytes

How many bytes to inline?

 https://www.

metadata	xian
metadata	uw-m
metadata	infl
metadata	liqu

 https://www.xiangpeng.systems
 https://www.uw-madison.edu
 https://www.influxdata.com
 https://www.liquid-cache.org

Inlined prefix

37%

bytes are needed for a
string sort
rest of the bytes are never used!

Long strings have low entropy (needs compression)

31%

String compression ratio

Overall: 66%

```
PlayStation 4|PlayStation 2
PlayStation 4|PlayStation 3
PlayStation 4|PlayStation 3|PC
PlayStation 4|PlayStation 3|PC|X...
PlayStation 4|PlayStation 5
PlayStation 4|PlayStation 5|PC|X...
PlayStation 4|Xbox 360
PlayStation 4|Xbox One
PlayStation 4|Xbox One|Nintendo Switch
PlayStation 4|Xbox One|Nintendo Switch|...
PlayStation 4|Xbox One|PC
PlayStation 4|Xbox One|PC|PlaySt...
PlayStation 4|Xbox One|PlayStati...
PlayStation 4|Xbox Series S/X|Ni...
PlayStation 4|Xbox Series S/X|Pl...
PlayStation 4|iOS
PlayStation 4|iOS|Nintendo Switch
PlayStation 4|iOS|PC|PlayStation...
PlayStation 4|iOS|Xbox One
PlayStation 5
PlayStation 5|Nintendo Switch
PlayStation 5|Nintendo Switch|PC...
PlayStation 5|PC
```

LZ4 / ZSTD

Option 1

General purpose compression

FSST

Option 2

String-specific compression

Long strings have low entropy (needs compression)

```
PlayStation 4|PlayStation 2
PlayStation 4|PlayStation 3
PlayStation 4|PlayStation 3|PC
PlayStation 4|PlayStation 3|PC|X...
PlayStation 4|PlayStation 5
PlayStation 4|PlayStation 5|PC|X...
PlayStation 4|Xbox 360
PlayStation 4|Xbox One
PlayStation 4|Xbox One|Nintendo Switch
PlayStation 4|Xbox One|Nintendo Switch|...
PlayStation 4|Xbox One|PC
PlayStation 4|Xbox One|PC|PlaySt...
PlayStation 4|Xbox One|PlayStati...
PlayStation 4|Xbox Series S/X|Ni...
PlayStation 4|Xbox Series S/X|Pl...
PlayStation 4|iOS
PlayStation 4|iOS|Nintendo Switch
PlayStation 4|iOS|PC|PlayStation...
PlayStation 4|iOS|Xbox One
PlayStation 5
PlayStation 5|Nintendo Switch
PlayStation 5|Nintendo Switch|PC...
PlayStation 5|PC
```

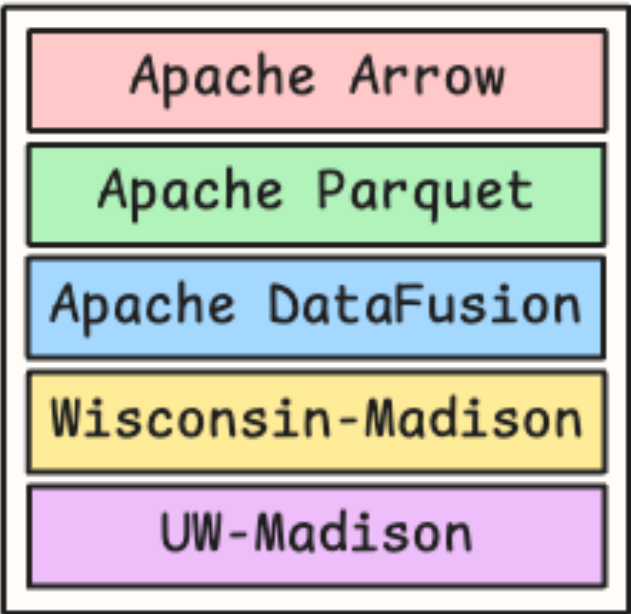

LZ4 / ZSTD

Option 1
General purpose compression

FSST

Option 2
String-specific compression

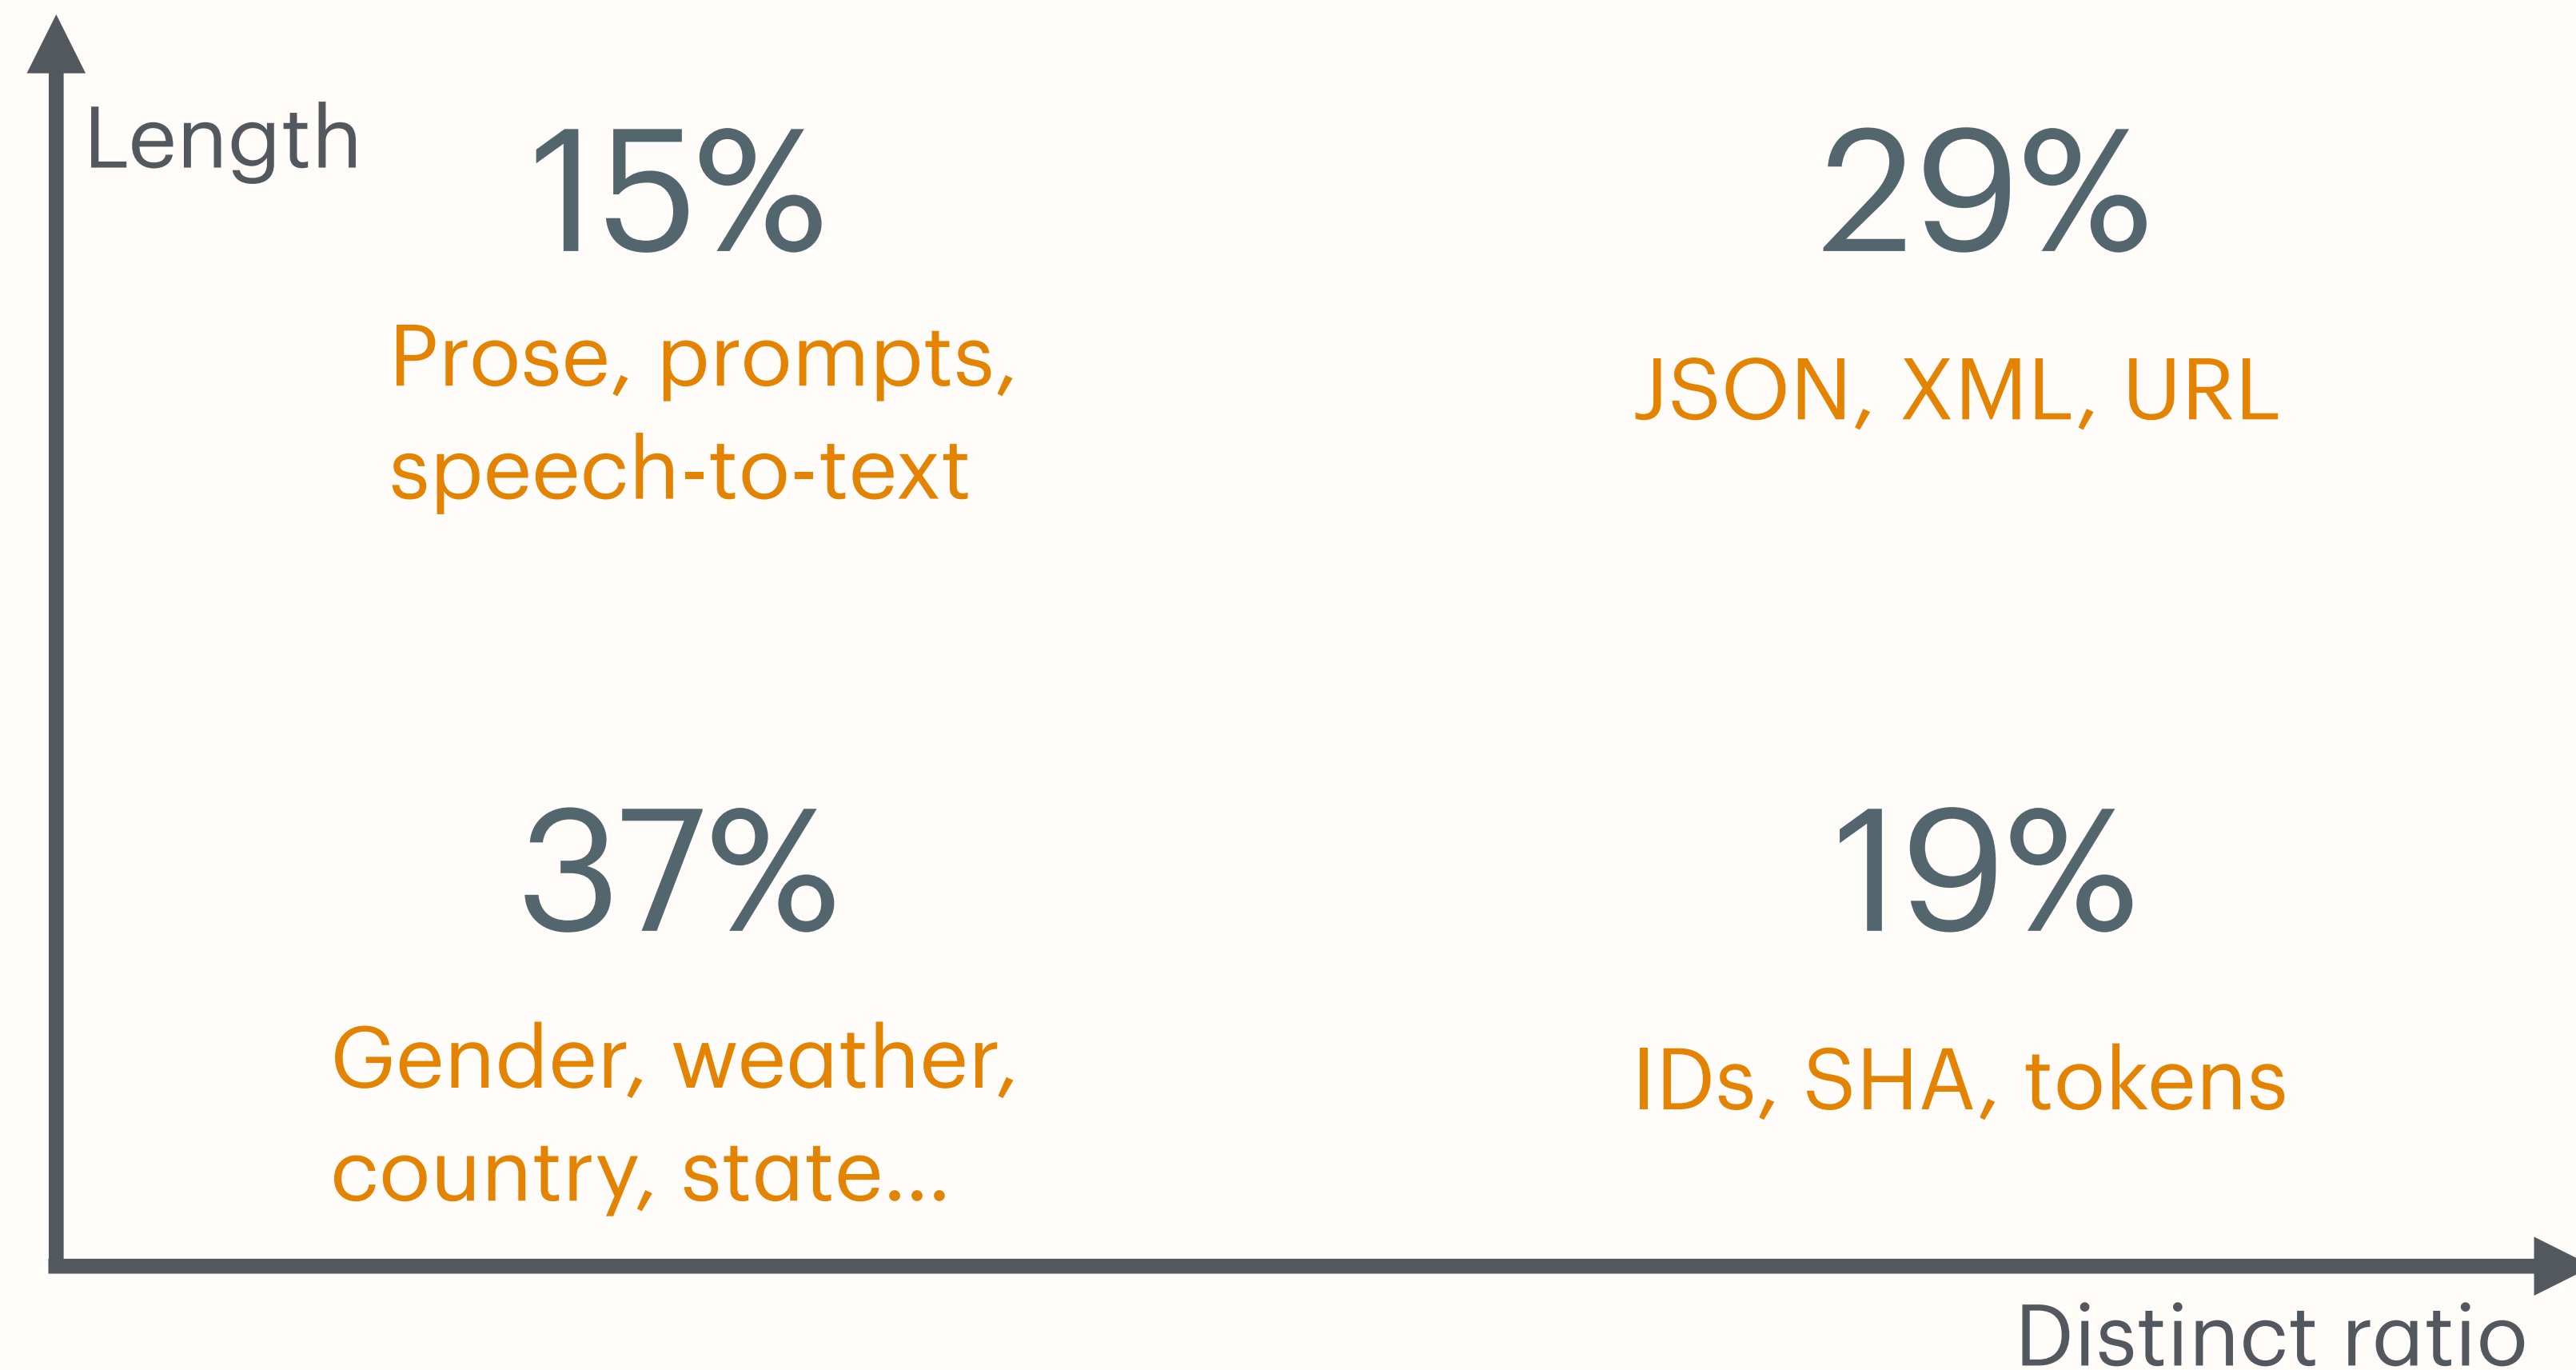
FSST is independently decodable!



Symbol table

Apache	-	a
Arrow	-	b
Parquet	-	c
DataFusion	-	d

Different strings are all called “strings”



Most common string: Null

Null

Proper null

“Null” | “null”

“Stringified” null

“” | “[]”

Empty means null

“NaN” | “nan”

Not a number is null

“None” | “none”

Rust/Python also has null

Null handling is important

Size

193 GB -> 308 GB

If “null”s are not handled properly:

- Force a column to be string
- Repetitive “null” strings

Sorting

“null”s are everywhere

Expected semantic: nulls
before or after sorted string

Skipping

Cause unnecessary scans

“null” strings breaks:

- min/max values
- common prefix

Putting everything together: string in LiquidCache

