Data Management for Video Analytics

Brandon Haynes, Maureen Daum, Amrita Mazumdar,
Magdalena Balazinska, Luis Ceze, & Alvin Cheung

1

Video data is everywhere.

2

2

Existing systems treat video data like it's the
20TH CENTURY

3

3

4

Many recent video applications require jointly querying multiple cameras, reasoning about position and orientation, or querying complex metadata.

5

5

6

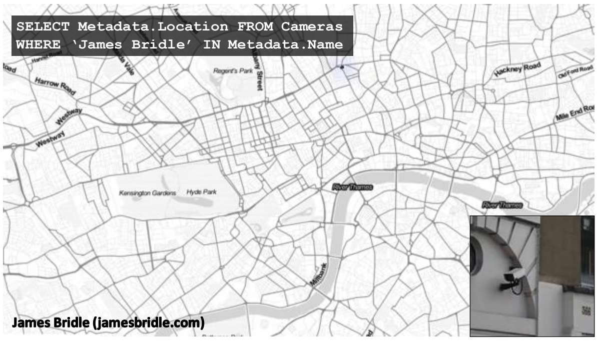
SELECT Metadata.Location FROM Cameras
WHERE 'James Bridle' IN Metadata.Name
James Bridle (jamesbridle.com)

7


500
20
National Geographic

8

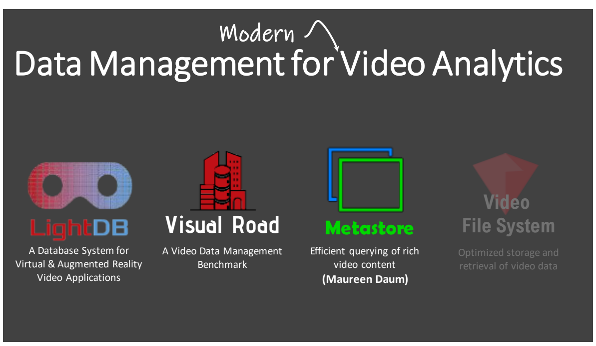
Modern
Data Management for Video Analytics
LightDB
A Database System for Virtual & Augmented Reality Video Applications
Visual Road
A Video Data Management Benchmark
Metastore
Efficient querying of rich video content
(Maureen Daum)
Video File System
Optimized storage and retrieval of video data

9

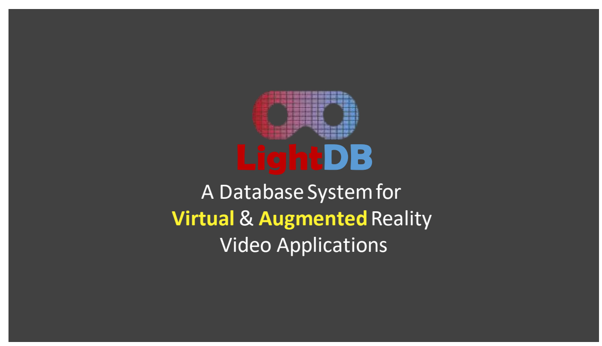
LightDB
A Database System for
Virtual & Augmented Reality
Video Applications

10

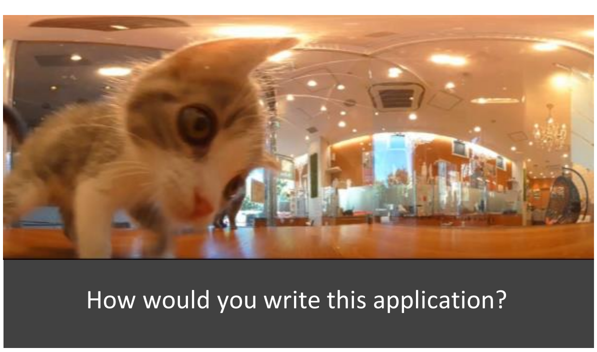
How would you write this application?

11


20TH CENTURY Imperative Code
LightDB Query
source   = Scan("kittens")
detection = source.Map(detect)
result   = Union(source, detection)
result.Save("output.mp4")
12

12

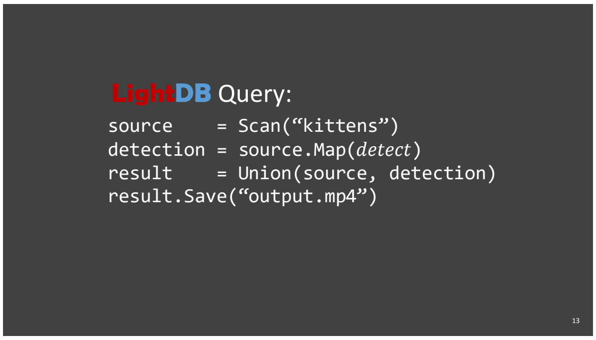

13

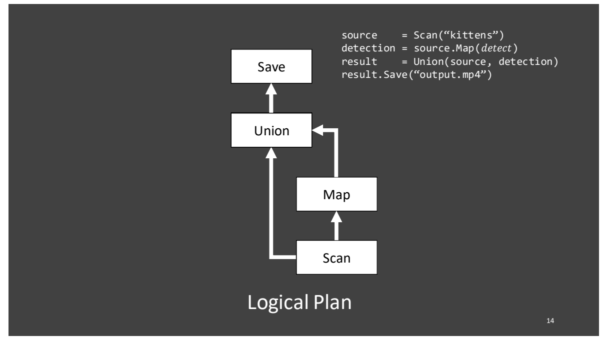

14

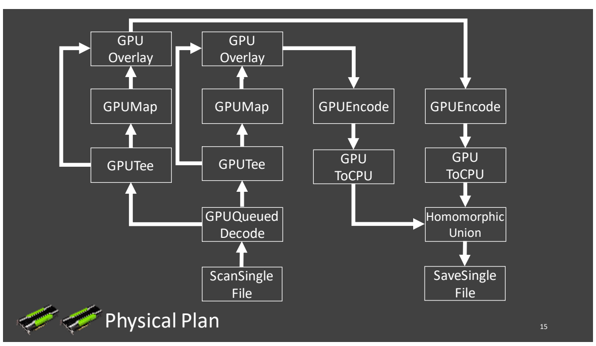

15

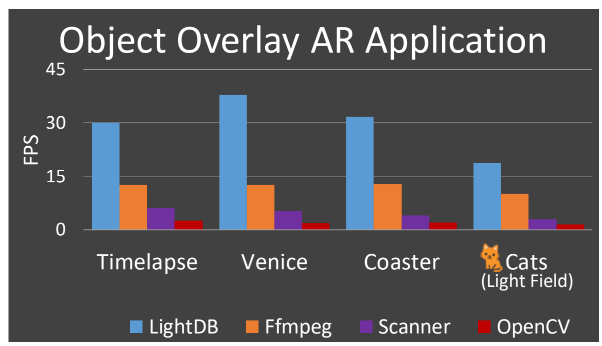

16

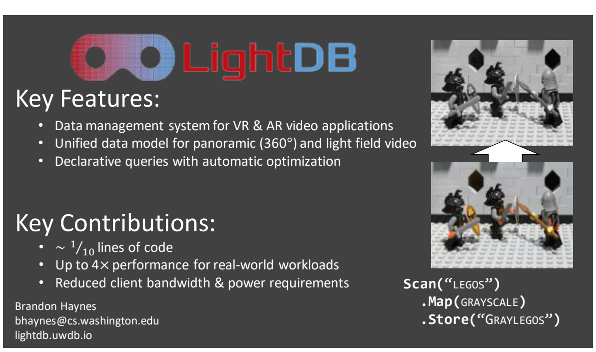

17

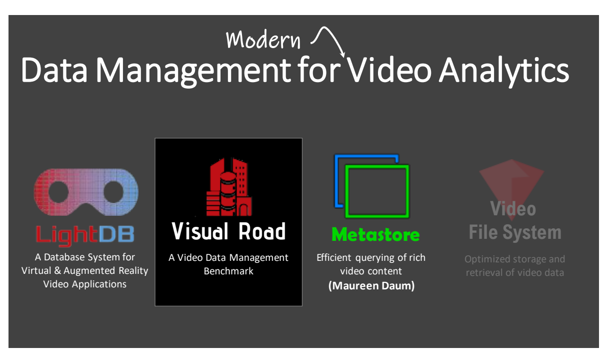

18

The performance of video systems is evaluated like it's still the 20TH CENTURY

19

Distinct videos performance tested by system

Video System	# Distinct Videos
LightDB (2018)	4
Chameleon (2018)	5
BlazeIt (2018)	6
NoScope (2017)	7
Focus (2018)	14
DeepLens (2019)	~16
Scanner (2018)	>100 (only 14 joined)

20

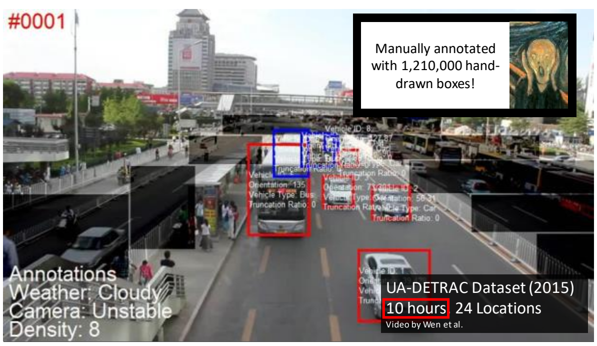

21

Type	Benchmarks	Test Data
Video	(Visual Road)	~~Ad Hoc~~ Synthetic
OLTP	(TPC-H)	Synthetic
OLAP	(SSB, DWEB)	Synthetic
Streaming	(Linear Road, DTDW)	Synthetic
NoSQL	(YCSB)	Synthetic
Graph	(LDBC)	Synthetic
Privacy	(SDV)	Synthetic

22

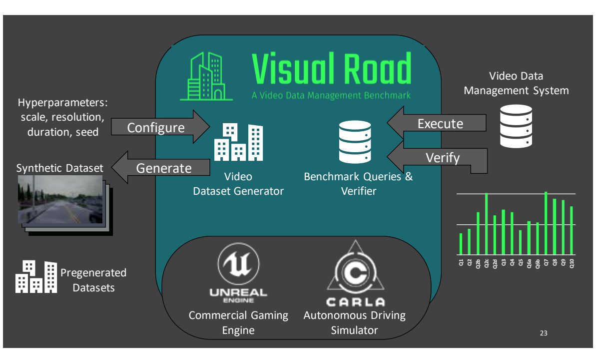

23

24

25

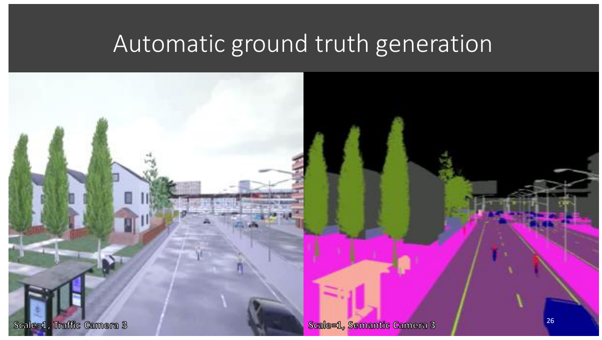

26

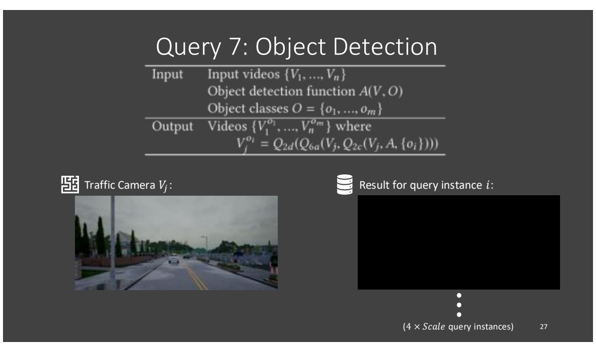

27

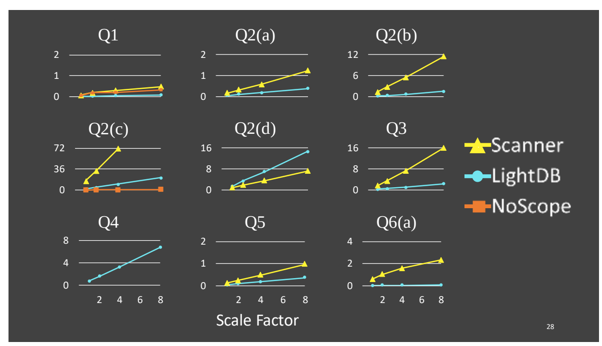

28

29

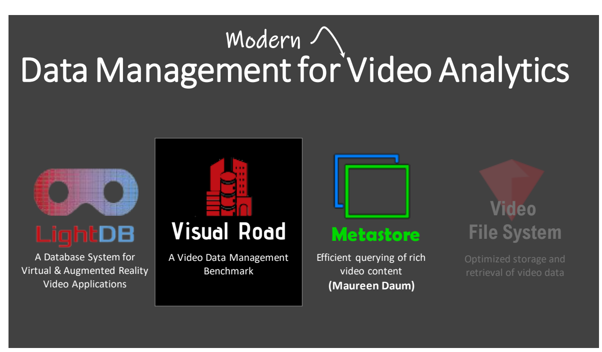

30

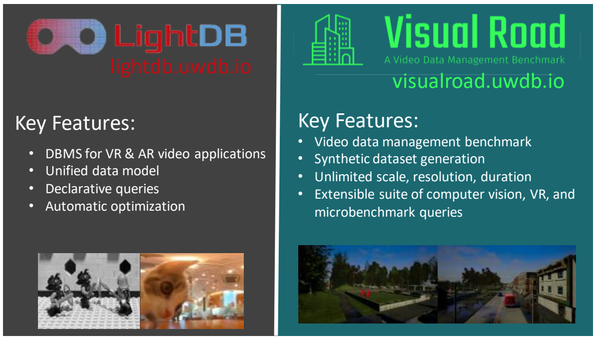

31

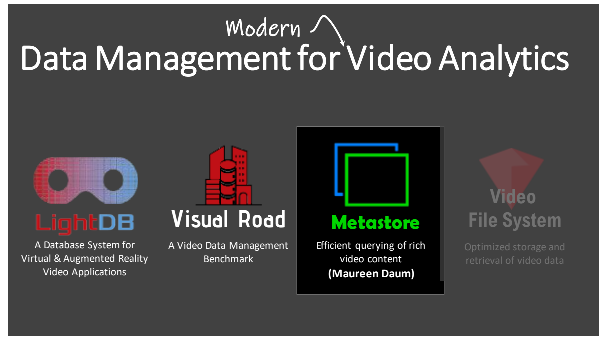

32

33

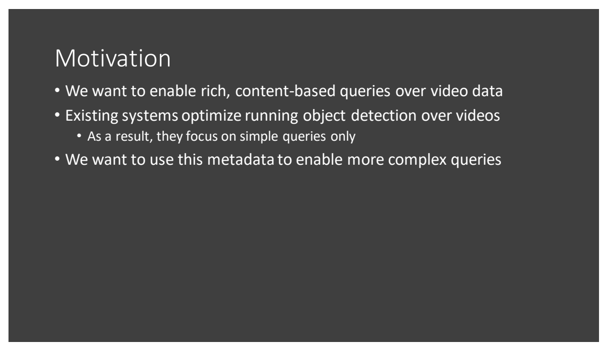

34

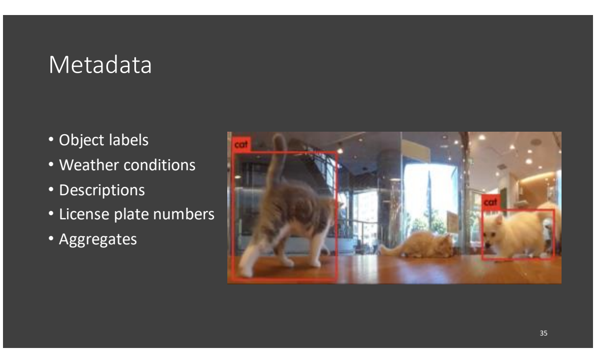

35

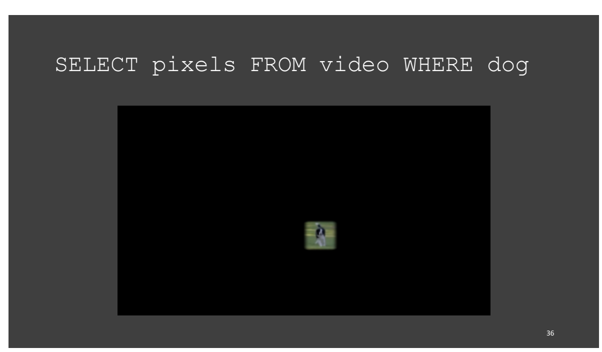

36

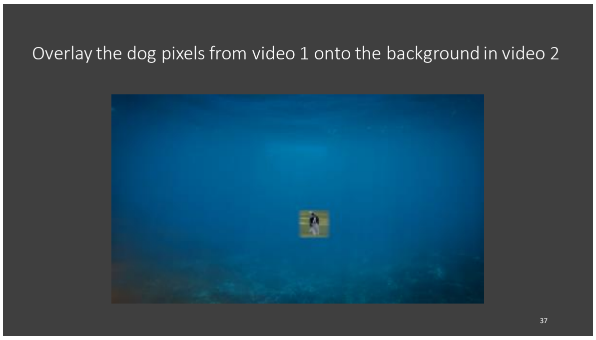

37

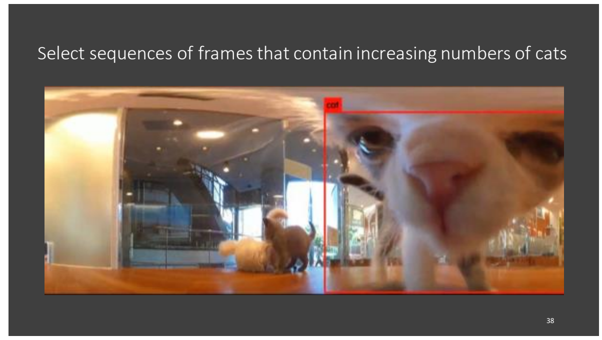

38

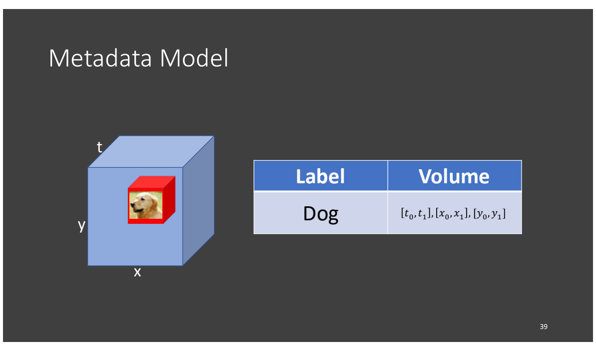

39

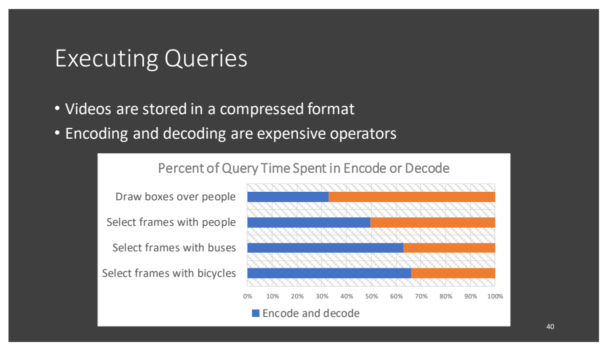

40

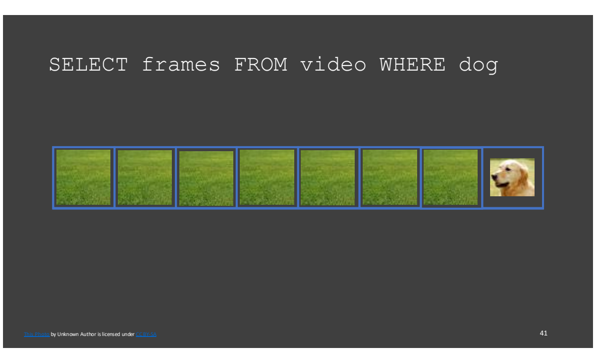

41

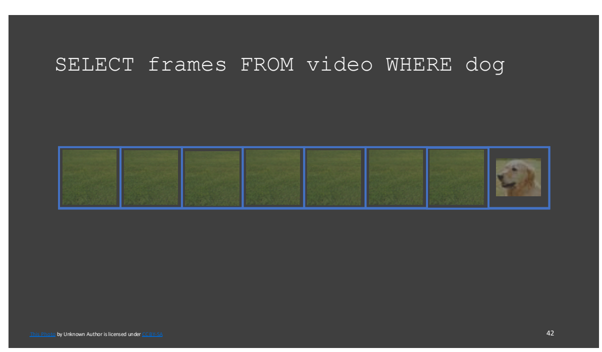

42

43

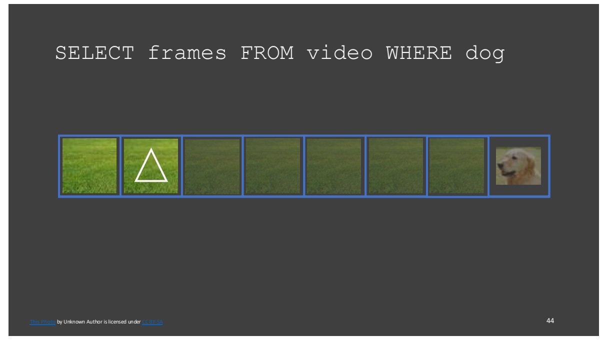

44

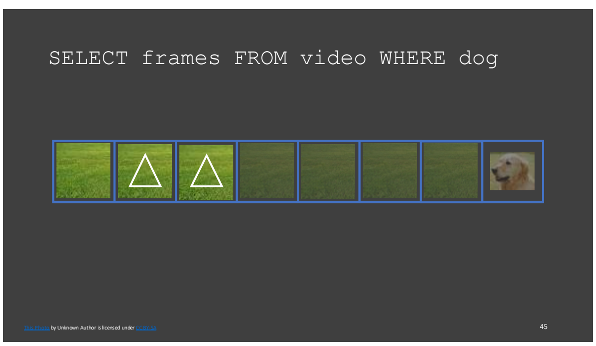

45

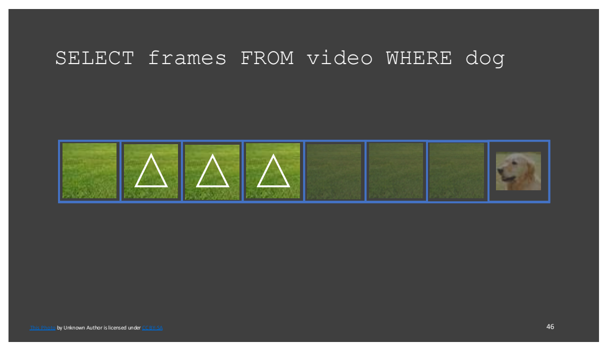

46

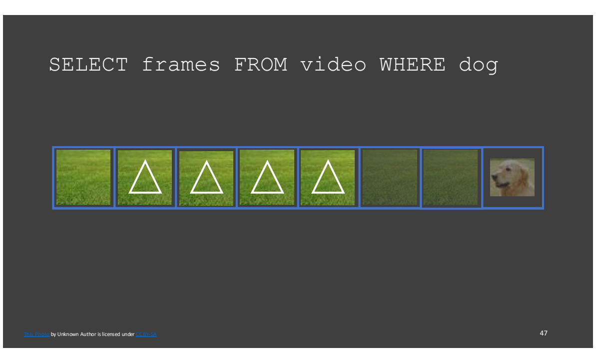

47

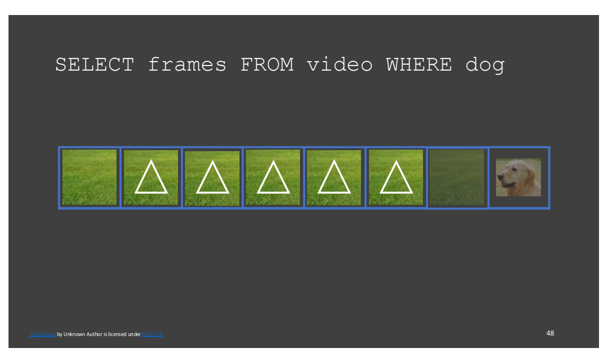

48

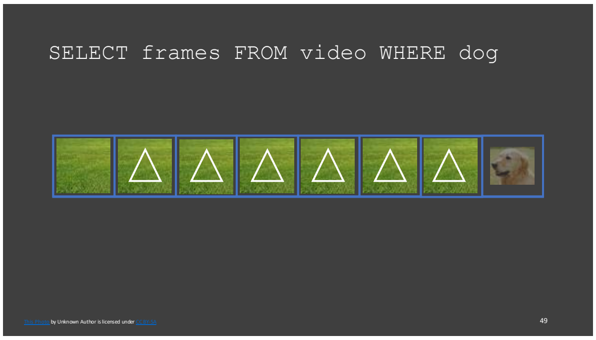

49

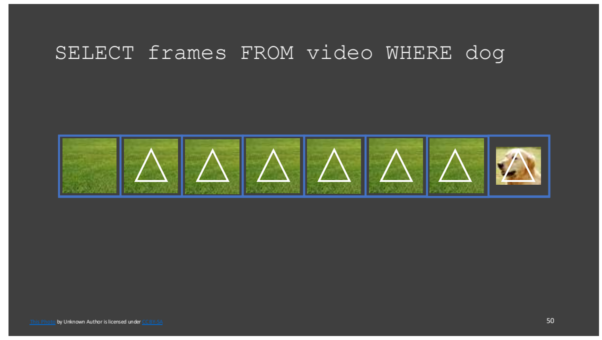

50

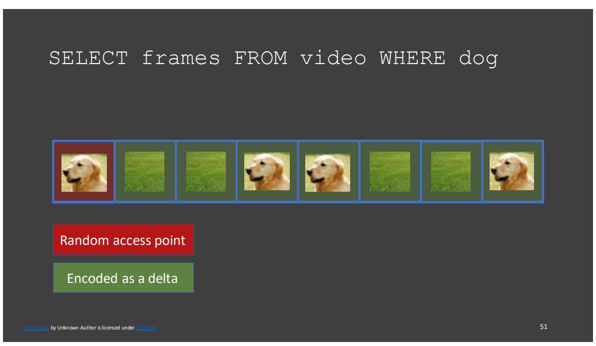

51

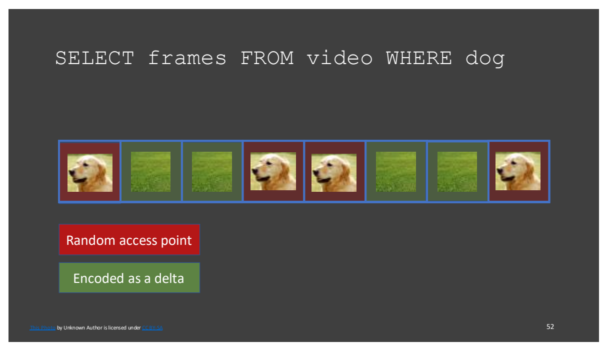

52

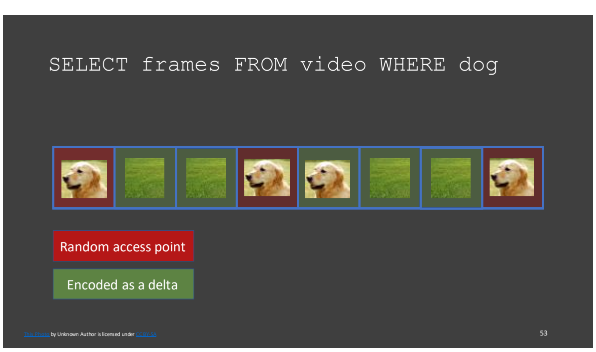

53

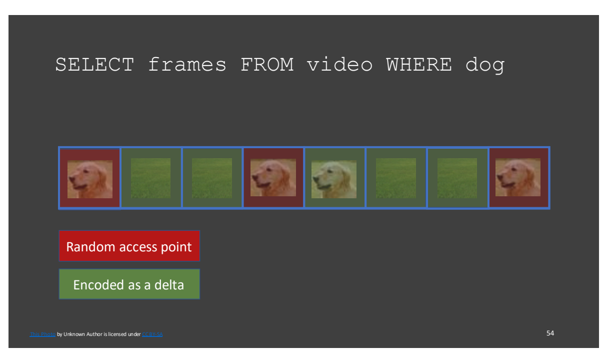

54

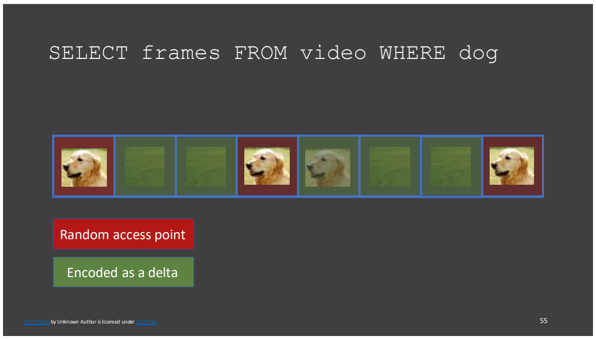

55

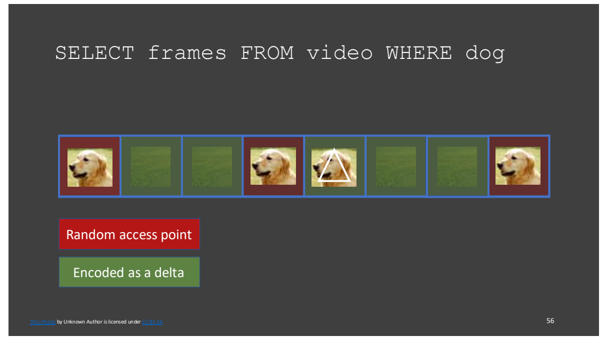

56

57

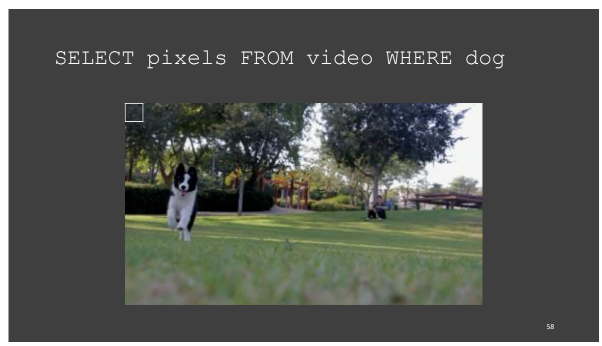

58

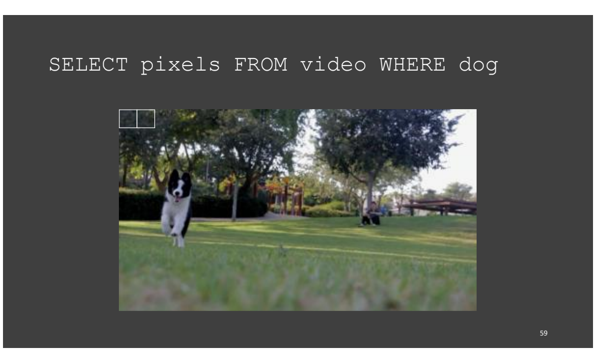

59

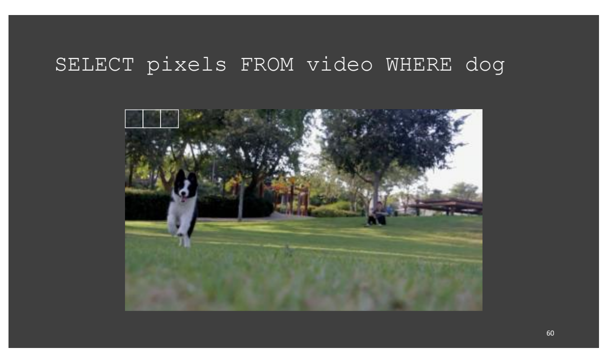

60

SELECT pixels FROM video WHERE dog

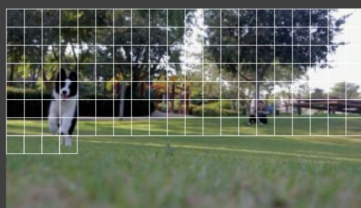

61

61

SELECT pixels FROM video WHERE dog

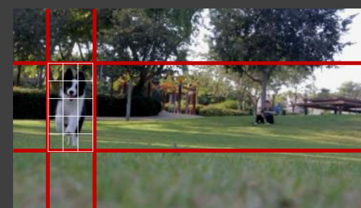

62

62

Putting Things Together

Similar to DB cracking, incrementally partition the video and add indices

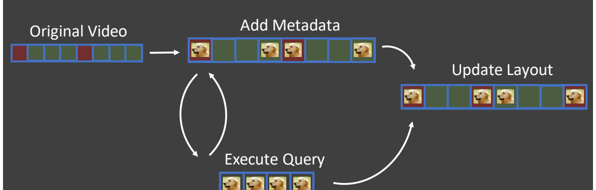

63

63

Ongoing Work

- Measure the effectiveness of selection optimization techniques
- Investigate optimization techniques for more compute-heavy queries
- Determine how to effectively layout videos with a lot of metadata
  - Possibly store multiple versions of a video with different layouts

64

64

Conclusion

- Deep learning opens the door to rich queries over video data
- Videos are large and slow to process
- Database techniques can accelerate such queries
  - Partitioning
  - Indexing
  - Incremental physical tuning
- Indexing must be balanced with maintaining reasonable storage sizes

65