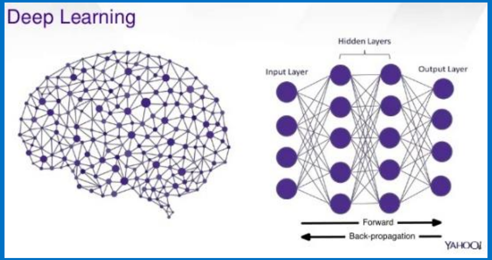

Images data and metadata

Find me cats

Addressing the Dark Side of Vision Research: Storage

Vishakha Gupta-Cledat, Luis Remis, Christina Strong

Intel Labs

Growing Image and Video Traffic

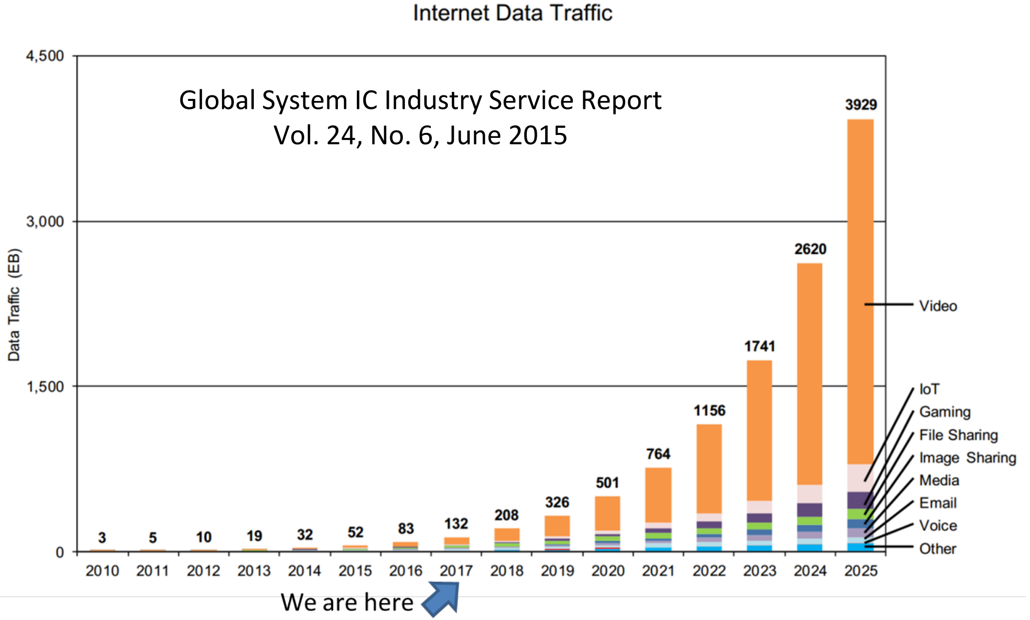

Enormous amounts of visual data, increasingly accessed and processed only by machines for a new class of vision applications

Rising Tide of Vision and Machine Learning Applications

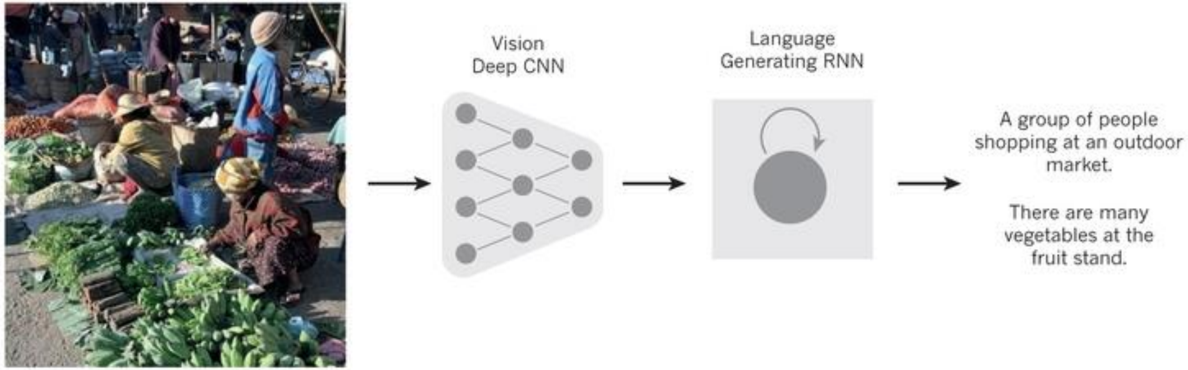

Rising Tide of Vision and Machine Learning Applications

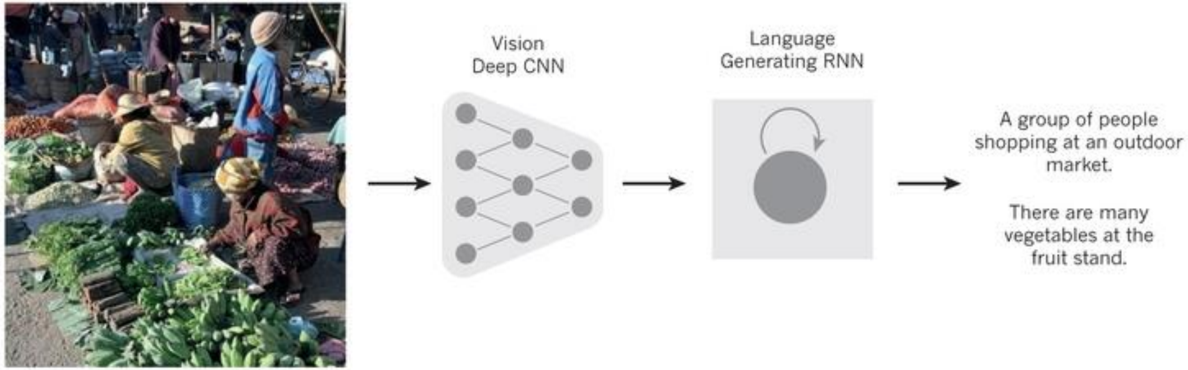

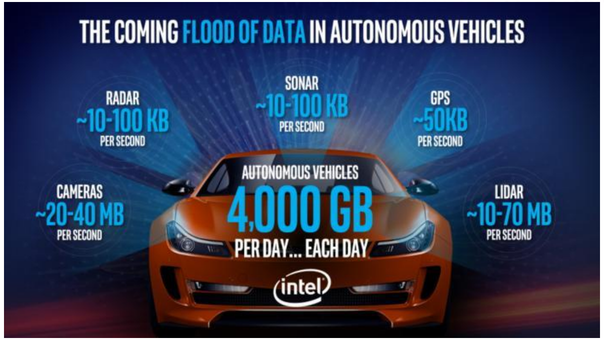

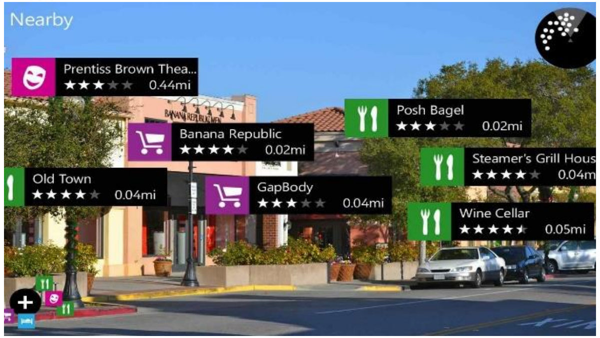

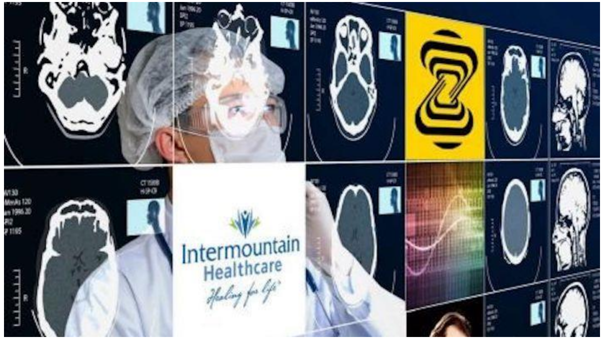

The increasing amount of information that computer vision and machine learning algorithms are identifying in images and videos not only produces more data but also enables new applications

Example Vision Workload

Stream of incoming data or data from file system → Vision processing to extract metadata e.g. image classification identifying objects




Example Vision Workload

Stream of incoming data or data from file system

Vision processing to extract metadata e.g. image classification identifying objects

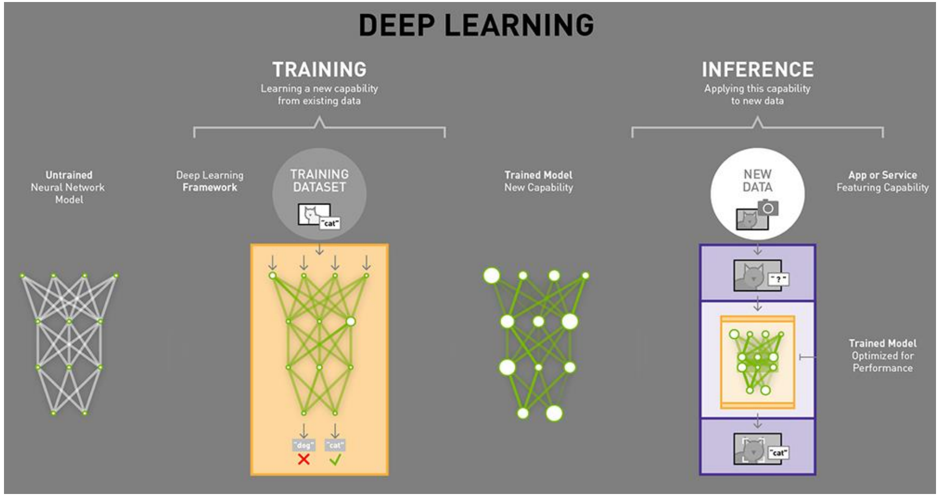

Example Vision Workload

Stream of incoming data or data from file system

Vision processing to extract metadata e.g. image classification identifying objects

Store data for future use

Queries for search of relevant data

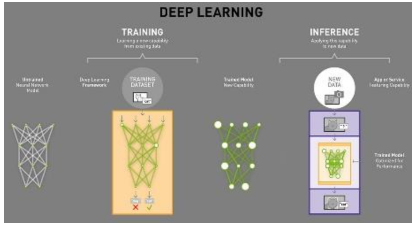

Research community focus on compute since it is expensive and algorithms still need development; storage is an after-thought

Example Vision Workload

Stream of incoming data or data from file system

Vision processing to extract metadata e.g. image classification identifying objects

Store data for future use

Queries for search of relevant data

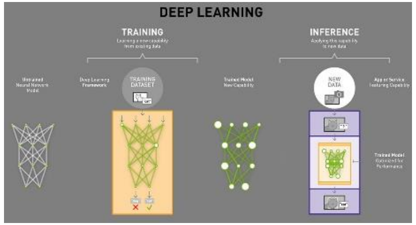

Research community focus on compute since it is expensive and algorithms still need development; storage is an after-thought

But Storage is a Problem: E.g. Image Classification using Deep Learning

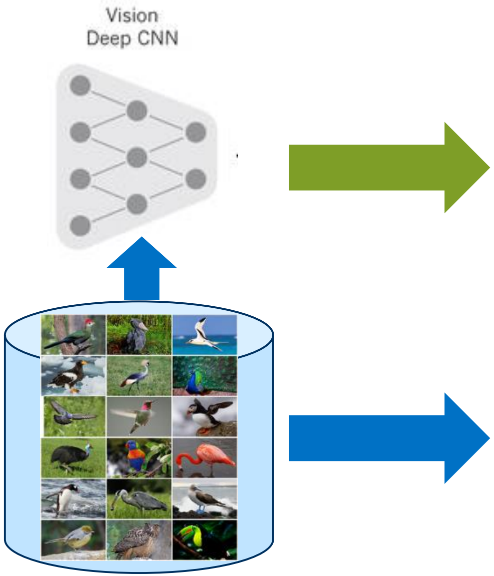

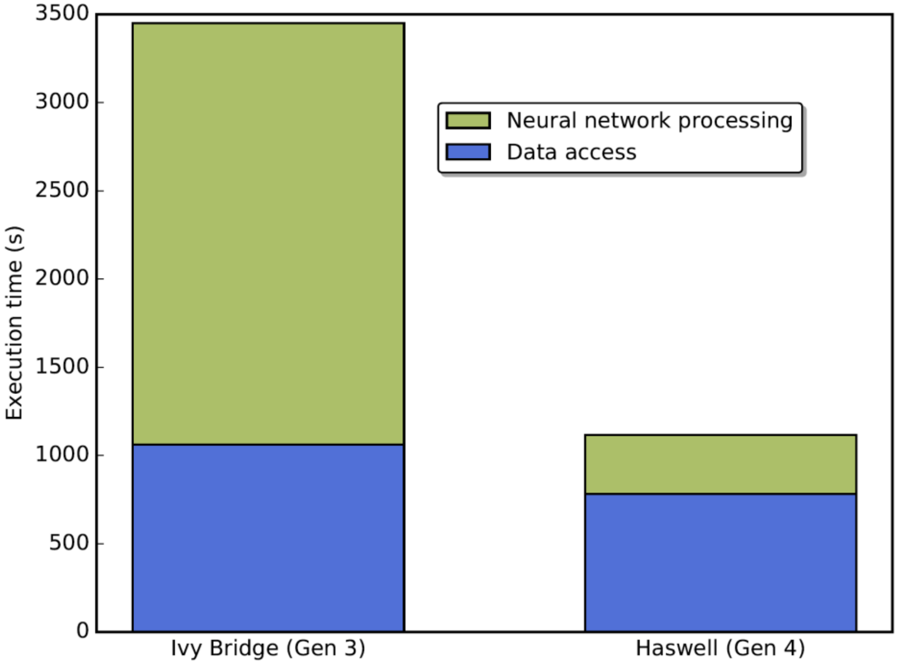

As processing capabilities and algorithms improve, amount of data increases, and data reuse becomes a possibility, storage goes from an afterthought to a real challenge

There is an early opportunity to influence the way visual data is treated in the storage system

Visual Metadata and New Formats

Exploit rich metadata

- Media data easily enhanced by rich metadata computed in advance or on the fly
- Metadata much smaller and can be used to zoom in, on only the desired raw data

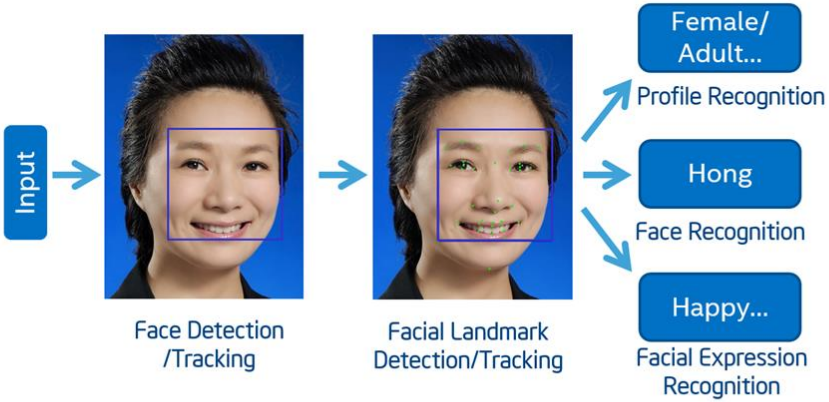

Search photos by faces, scenes, objects, and actions/events

Source: Yurong Chen, Intel Labs China

Visual Metadata and New Formats

Exploit rich metadata

- Media data easily enhanced by rich metadata computed in advance or on the fly
- Metadata much smaller and can be used to zoom in, on only the desired raw data

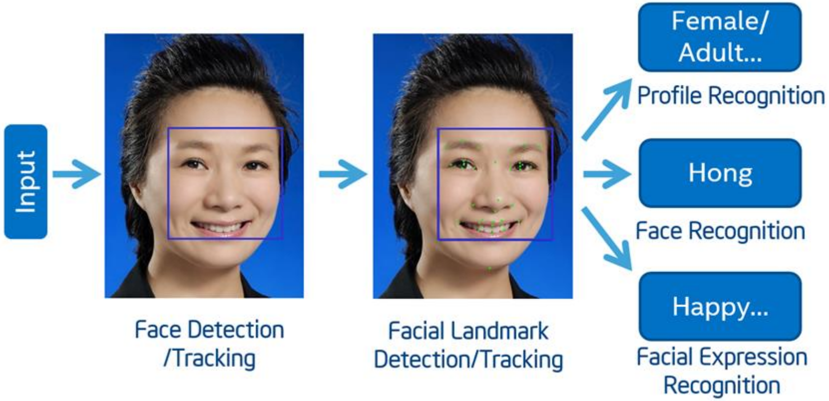

Search photos by faces, scenes, objects, and actions/events

Source: Yurong Chen, Intel Labs China

Speed up access to desired data

- Analysis friendly formats

Visual Storage Architecture

Persistent Memory Graph Database (PMGD)

- Support efficient query of metadata via a graph database developed for new non-volatile memory (Persistent Memory)

Visual Compute Library (VCL)

- Enable alternate image/video analysis-friendly storage formats

Request Server

- Coordinate requests and responses PMGD, VCL and client
- Implement client API

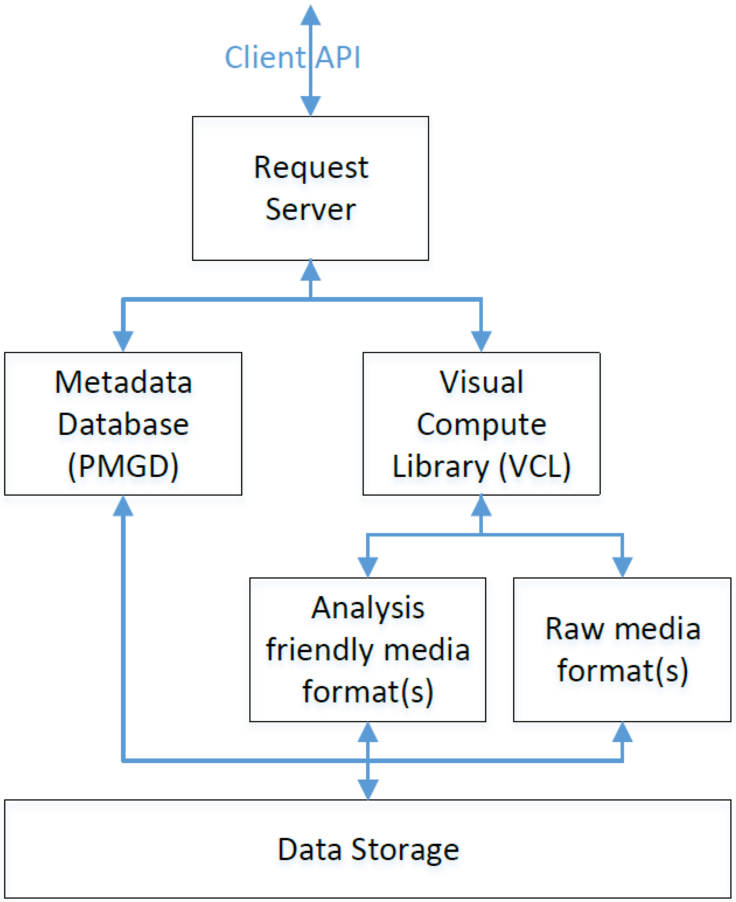

Visual Storage Architecture

Persistent Memory Graph Database (PMGD)

- Support efficient query of metadata via a graph database developed for new non-volatile memory (Persistent Memory)

Visual Compute Library (VCL)

- Enable alternate image/video analysis-friendly storage formats

Request Server

- Coordinate requests and responses PMGD, VCL and client
- Implement client API

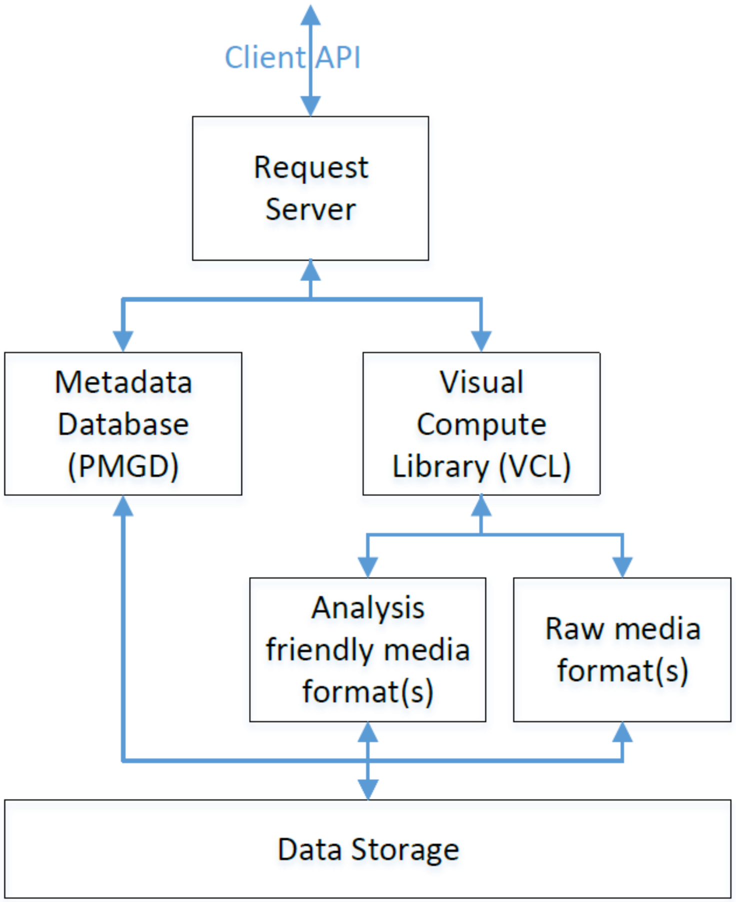

Representing Media Metadata

Photo

Name: Hawaii1.jpg

Date: 4/15/14

Size: 2MB

Photo

Name: Hawaii2.jpg

Date: 4/16/14

Size: 2.5MB

Find all photos of Alice from Hawaii

Representing Media Metadata

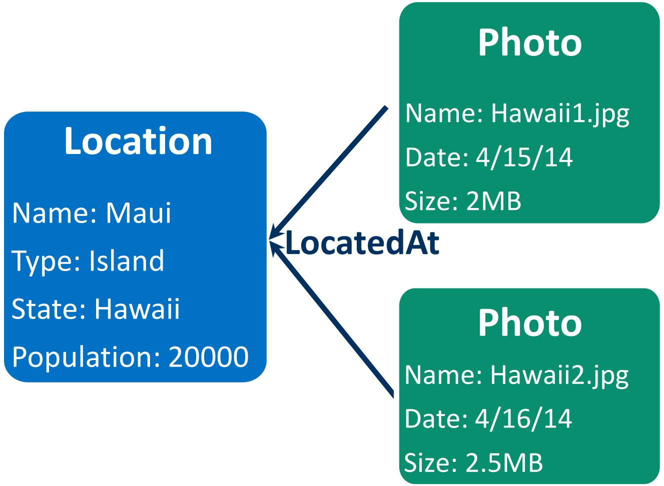

Find all photos of Alice from Hawaii

Representing Media Metadata

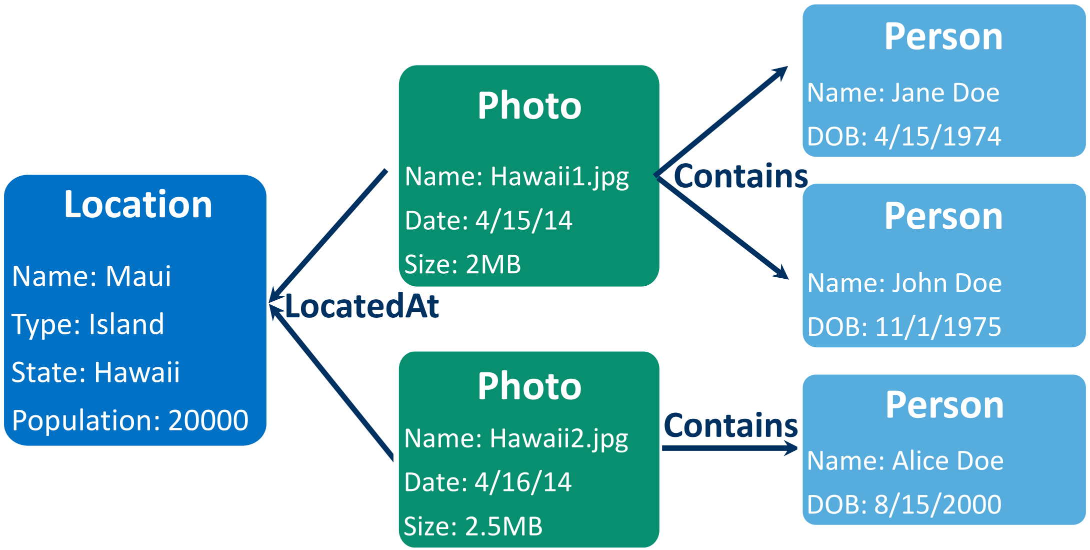

Find all photos of Alice from Hawaii

Representing Media Metadata

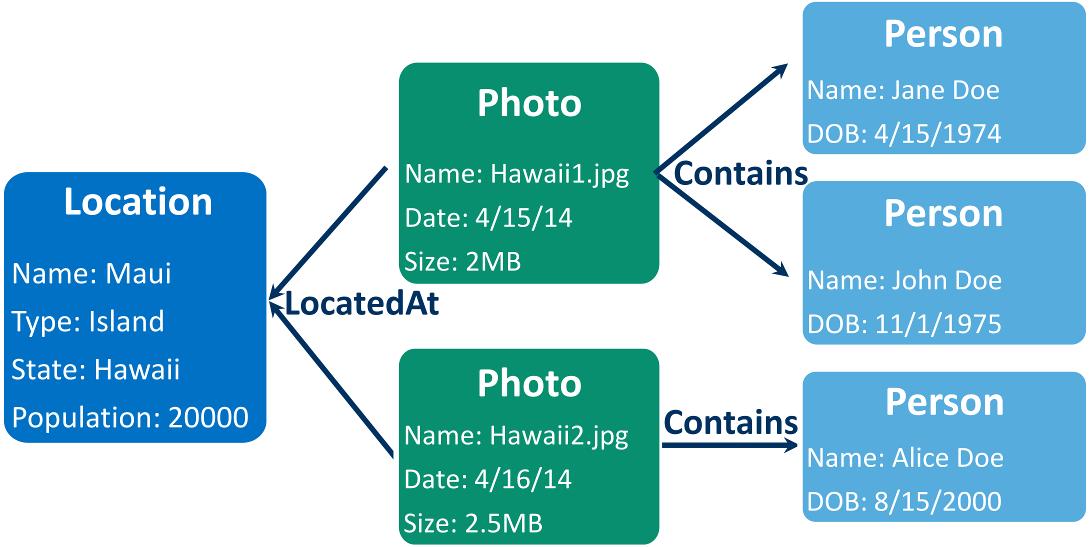

While this metadata schema will be application-specific, it looks like a property graph:

- *Nodes* connected with *Edges*
- *Properties* on nodes/edges
- (optional) Group by *tags*

Find all photos of Alice from Hawaii

Representing Media Metadata

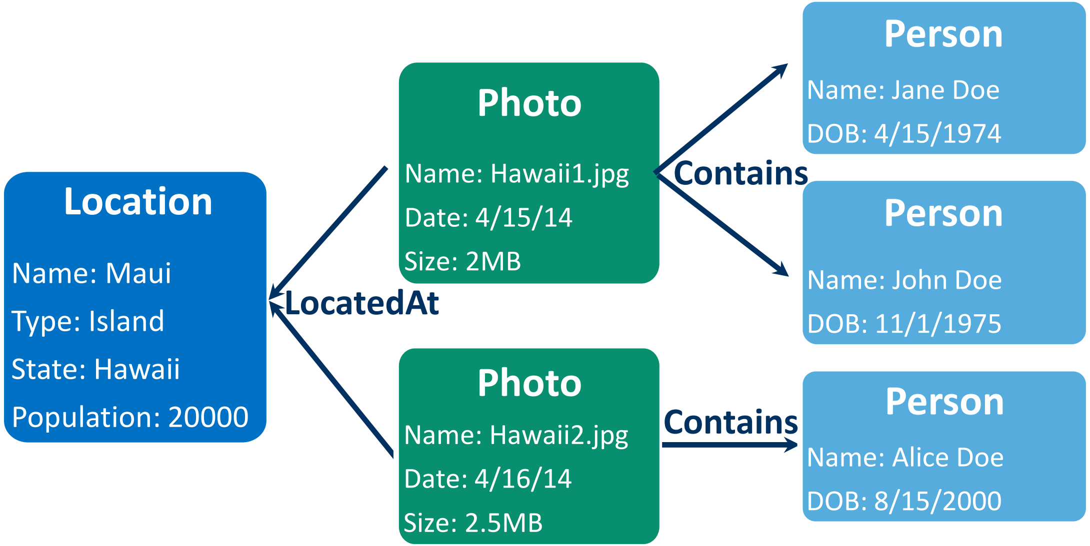

Find all photos of Alice from Hawaii

While this metadata schema will be application-specific, it looks like a property graph:

- *Nodes* connected with *Edges*
- *Properties* on nodes/edges
- (optional) Group by *tags*

Support evolving schema

Variety of indexes

Persistent Memory Graph Database (PMGD)

Traditional property graph databases plagued by disk latencies

New non-volatile memory technology (e.g. 3D Xpoint) with performance close to DRAM

Opportunity to avoid a lot of legacy software → PMGD

- Graph database implementation targeting persistent memory

Data Storage via the Visual Compute Library

More and more machine consumption of data for processing

- Think beyond standard formats for visual data

Visual Compute Library

- Enable alternate image/video storage formats (e.g. TileDB [1])
- Perform common operations closer to the data

[1] Stavros Papadopoulos, Kushal Datta, Samuel Madden, and Timothy Mattson. 2016. The TileDB array data storage manager. Proc. VLDB Endow. 10, 4 (November 2016), 349-360. DOI: https://doi.org/10.14778/3025111.3025117

Example Visual Pipeline(s)

"Find images with cars on beaches to train a new model for recognizing their make"

Images from YFCC100M → Image Classification (Places) → Filter (beach images) → Image Classification (Objects) → Filter (cars on beaches) → Images with cars on beaches

Using Yahoo Flickr dataset, create a biased set (700K images)

Image classification using Caffe deep learning framework (different weights and machine learning models for locations and objects)

Metadata Query Comparison

t(“images with beaches after place classification”) + t(“images with cars and beaches after object classification”)

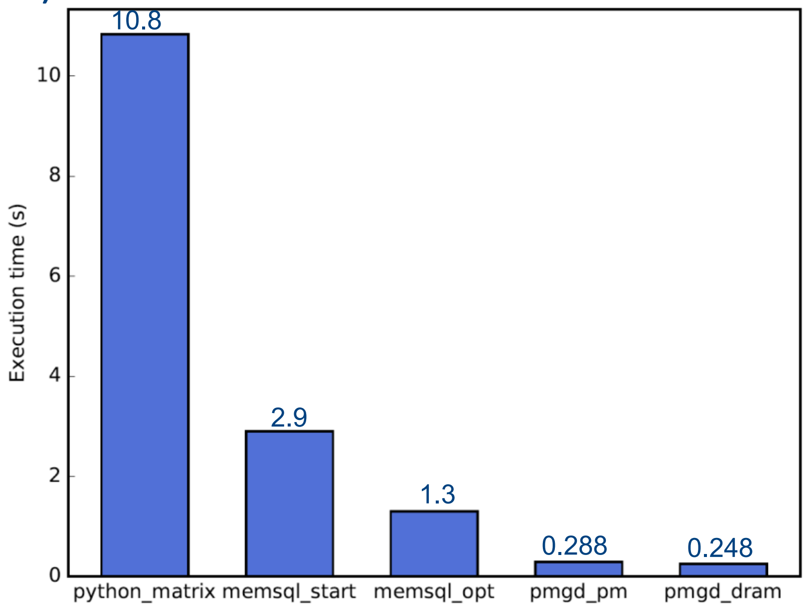

PMGD supports queries faster than an in-memory relational database while providing better schema support.

Query Scaling

t("images with beaches after place classification")

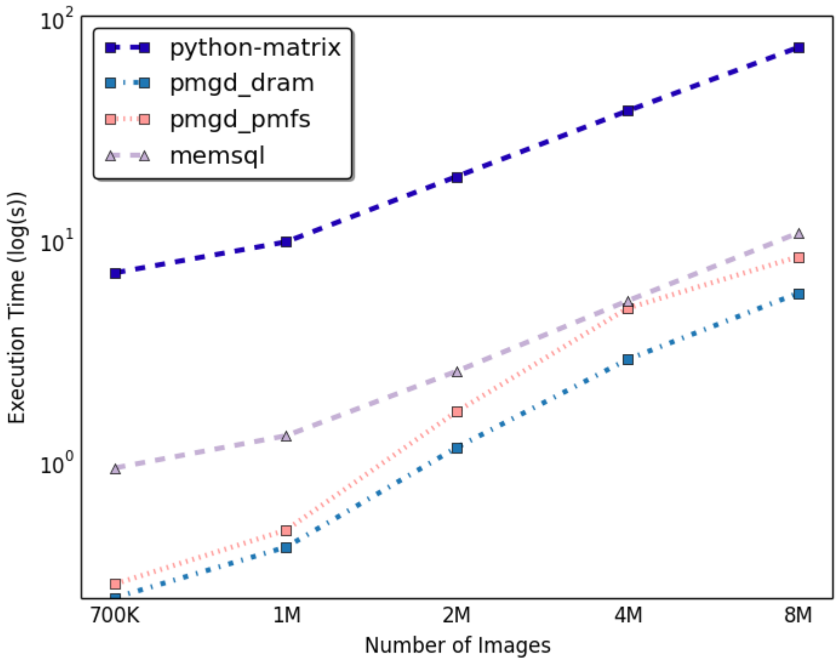

PMGD continues to perform better than MemSQL at larger scales despite a lack of special query caching provided in MemSQL

Analysis friendly processing in VCL – e.g. resize

t(“read images and resize”)

Image Size	Image Format	Read/decode Implementation	Resize Implementation	Read + Resize Time (ms)
500x291	JPEG	OpenCV	OpenCV	2.85
500x291	TileDB	VCL	OpenCV	0.37
500x291	TileDB	VCL	VCL native	1.36
1024x2048	PNG	OpenCV	OpenCV	68.04
1024x2048	TileDB	VCL	OpenCV	4.63
1024x2048	TileDB	VCL	VCL native	2.89

Resize 400 images from original size to 256x256. Average time per image

Larger the image, better the performance of VCL, more benefits for high-def images

Conclusions and Future Work

Room and need for novel storage methods in vision pipelines

Graph database, made efficient with new technology, a good option for metadata

Analysis friendly data storage a worthwhile research direction

Conclusions and Future Work

Room and need for novel storage methods in vision pipelines

Graph database, made efficient with new technology, a good option for metadata

Analysis friendly data storage a worthwhile research direction

Complete the architecture and evaluate with real applications

Address video storage and streaming use cases

Scale out

Conclusions and Future Work

Room and need for novel storage methods in vision pipelines

Graph database, made efficient with new technology, a good option for metadata

Analysis friendly data storage a worthwhile research direction

Complete the architecture and evaluate with real applications

Address video storage and streaming use cases

Scale out