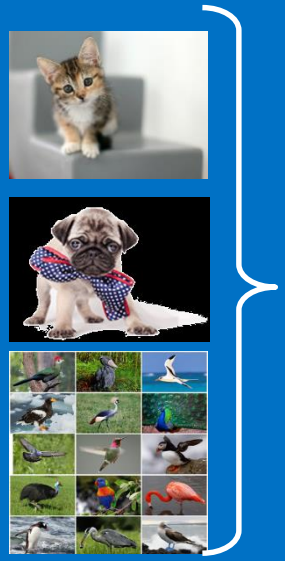

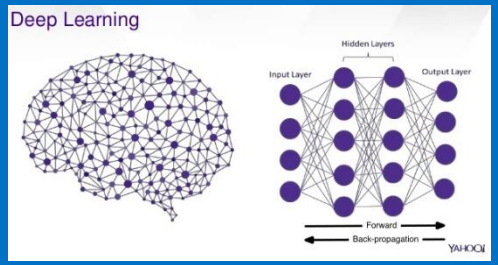

Images data and metadata

Find me cats

Visual Data Management System

Vishakha Gupta-Cledat, Luis Remis, Christina Strong, Ragaad Altarawneh, Scott Hahn

vishakha.s.gupta, luis.remis, christina.r.strong, ragaad.altarawneh, scott.hahn@intel.com

Intel Labs

What is VDMS?

A novel Visual Data Management System

- For storing, accessing and transforming visual data

What is VDMS?

A novel Visual Data Management System

- For storing, accessing and transforming visual data
- Primarily geared towards visual analytics pipelines and data science queries

What is VDMS?

A novel Visual Data Management System

- For storing, accessing and transforming visual data
- Primarily geared towards visual analytics pipelines and data science queries
- With a goal of efficiently achieving cloud scale while maintaining ease-of-use

What is VDMS?

A novel Visual Data Management System

- For storing, accessing and transforming visual data
- Primarily geared towards visual analytics pipelines and data science queries
- With a goal of efficiently achieving cloud scale while maintaining ease-of-use

Also aims to

- Exploit Intel's heterogeneous memory and storage hierarchy

What is VDMS?

A novel Visual Data Management System

- For storing, accessing and transforming visual data
- Primarily geared towards visual analytics pipelines and data science queries
- With a goal of efficiently achieving cloud scale while maintaining ease-of-use

Also aims to

- Exploit Intel's heterogeneous memory and storage hierarchy
- Be general purpose e.g. common core for medical imaging, sports, retail

Visual Data: Scale and Applications

Billions of sources

Visual Data: Scale and Applications

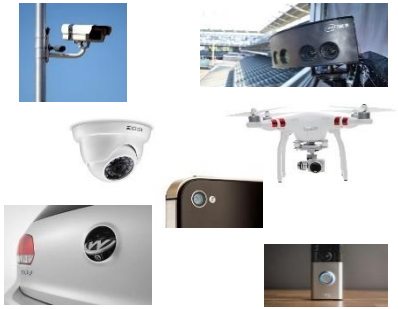

Billions of sources

Images

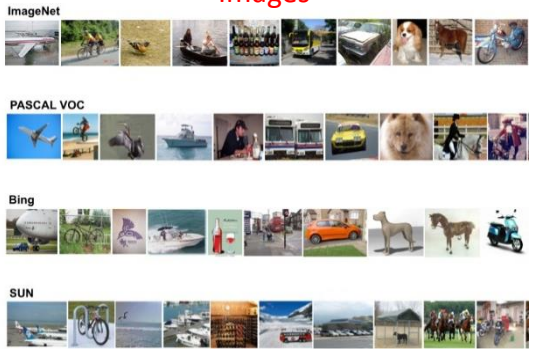

Videos

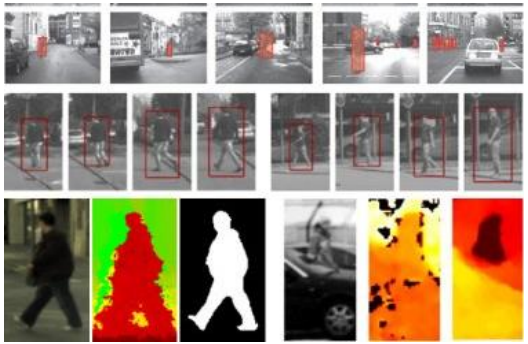

Large in size (individual object could range in size from KB to GB)
Increasingly being used for visual understanding in a range of machine learning applications

Visual Data: Scale and Applications

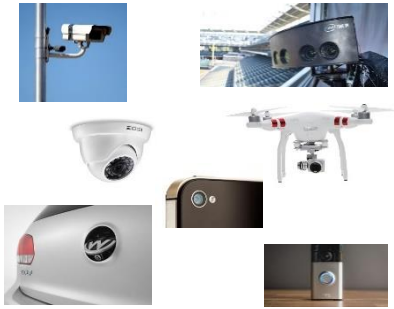

Billions of sources

Images

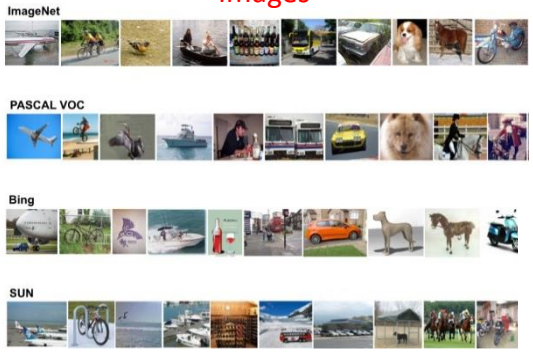

Videos

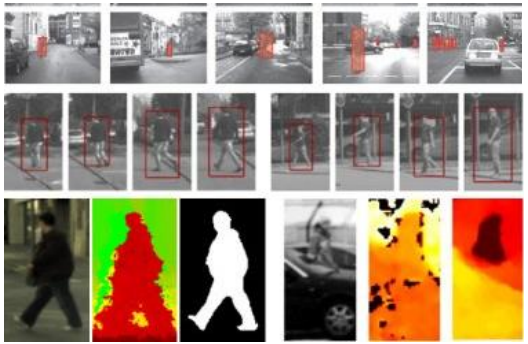

Vision
Deep CNN

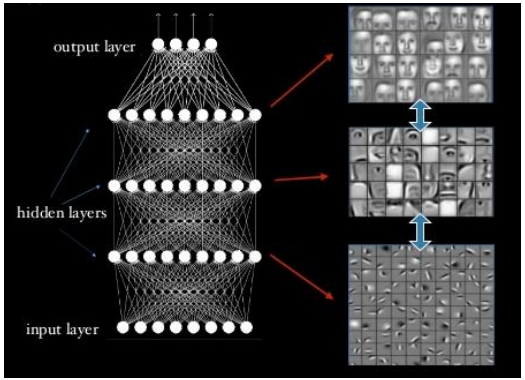

Feature Vectors / Descriptors

Large in size (individual object could range in size from KB to GB)
Increasingly being used for visual understanding in a range of machine learning applications

The Unsustainable Current Solutions

Resolve visual computing challenges and frameworks first

- Improving accuracy of algorithms on more and more complex data
- Storage has not become a bottleneck yet!

The Unsustainable Current Solutions

Resolve visual computing challenges and frameworks first

- Improving accuracy of algorithms on more and more complex data
- Storage has not become a bottleneck yet!

Application-specific solutions, if data does become a problem

- Organize media files
- Manually gather and normalize relevant metadata
- Build custom scripts to tie together many stages of complex processing

The Unsustainable Current Solutions

Resolve visual computing challenges and frameworks first

- Improving accuracy of algorithms on more and more complex data
- Storage has not become a bottleneck yet!

Application-specific solutions, if data does become a problem

- Organize media files
- Manually gather and normalize relevant metadata
- Build custom scripts to tie together many stages of complex processing

Visual data management for scale and reuse is still an open problem.

VDMS Storage Architecture

Exploding amount of visual data

- For any request, access only the required subset of data – exploit metadata

VDMS Storage Architecture

Exploding amount of visual data

- For any request, access only the required subset of data – exploit metadata

Even individual objects could be large

- Speed up access to this desired data
- Preprocess while reading where possible e.g. crop or detect edges before transferring

VDMS Storage Architecture

Exploding amount of visual data

- For any request, access only the required subset of data – exploit metadata

Even individual objects could be large

- Speed up access to this desired data
- Preprocess while reading where possible e.g. crop or detect edges before transferring

High performance as well as ease-of-use

- Suitable design choices for metadata and data, at scale
- Intel hardware optimizations e.g. 3D Xpoint, media hardware, disk offload
- Simple API and client libraries

VDMS Implementation

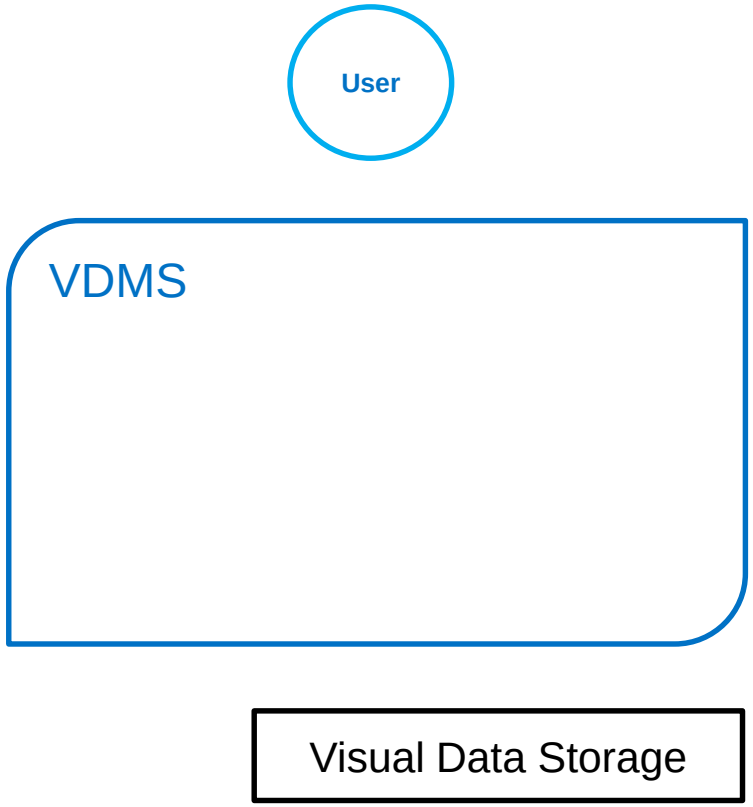

VDMS Implementation

Efficient metadata access via Persistent Memory Graph Database (PMGD) for visual data

- Optimized for metadata storage and access patterns
- Easy to evolve schema with new vision research

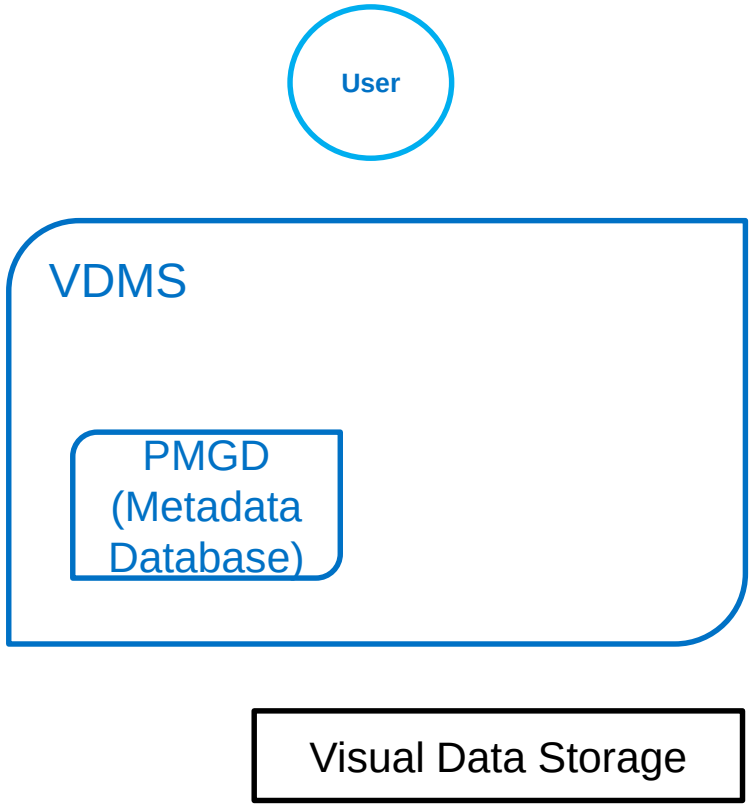

VDMS Implementation

Efficient metadata access via Persistent Memory Graph Database (PMGD) for visual data

- Optimized for metadata storage and access patterns
- Easy to evolve schema with new vision research

Efficient data access via Visual Compute Library

- Enable alternate image/video analysis friendly storage formats as compared to viewer friendly ones
- Process data while accessing it

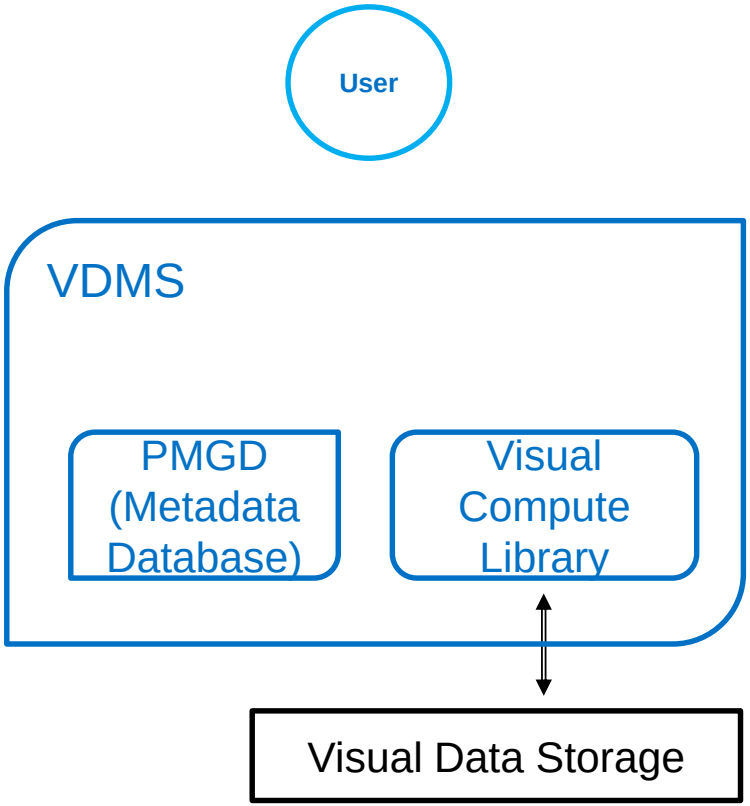

VDMS Implementation

Efficient metadata access via Persistent Memory Graph Database (PMGD) for visual data

- Optimized for metadata storage and access patterns
- Easy to evolve schema with new vision research

Efficient data access via Visual Compute Library

- Enable alternate image/video analysis friendly storage formats as compared to viewer friendly ones
- Process data while accessing it

Ease-of-use via Request Server

- Implement a unified and simple client API
- Route query (or parts) to the right components for a coherent user response

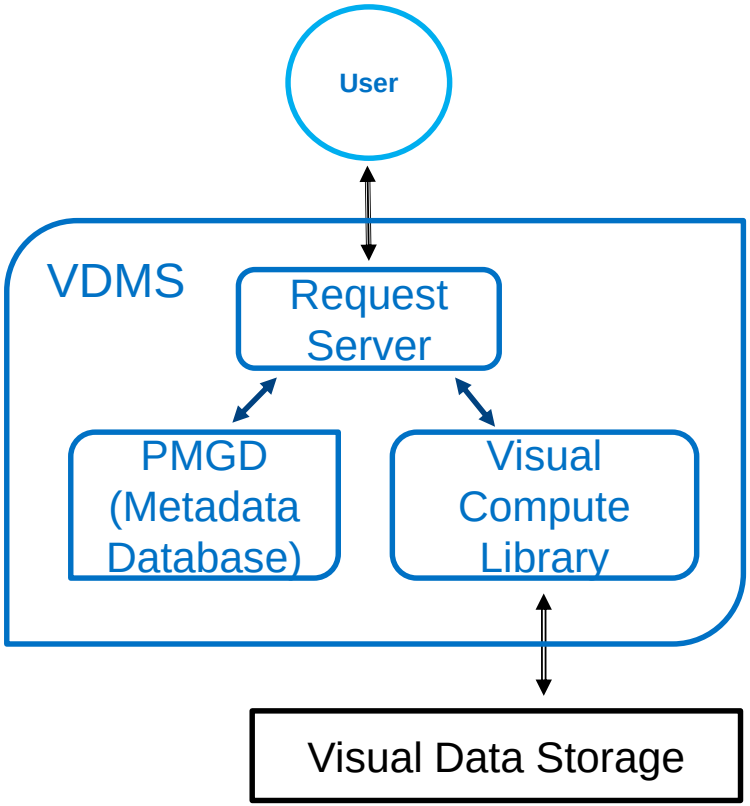

Where We Are Now

User API v1.0 defined with internal feedback

Where We Are Now

User API v1.0 defined with internal feedback

Functional one node server and client libraries

Where We Are Now

User API v1.0 defined with internal feedback

Functional one node server and client libraries

Three interesting proofs of concept at various stages of development with input from product groups

- Real data and concrete use case: medical imaging application
- Large scale, real time, intensive use case: FreeD sports storage architecture
- Integration with a larger analytic framework: Retail shopper insights application

Medical Imaging Proof of Concept on VDMS

The Cancer Image Archive: http://www.cancerimagingarchive.net/

- 60TB of medical images (Volumetric data)
- ~1000 patients metadata information (very sparse)

Medical Imaging Proof of Concept on VDMS

The Cancer Image Archive: http://www.cancerimagingarchive.net/

- 60TB of medical images (Volumetric data)
- ~1000 patients metadata information (very sparse)

For our PoC:

- 457 patients metadata, including drug and radiation treatments
- Scans for 384 patients (60K images)
- Replicated metadata x10 and x100, keeping the original distribution

Medical Imaging Proof of Concept on VDMS

The Cancer Image Archive: http://www.cancerimagingarchive.net/

- 60TB of medical images (Volumetric data)
- ~1000 patients metadata information (very sparse)

For our PoC:

- 457 patients metadata, including drug and radiation treatments
- Scans for 384 patients (60K images)
- Replicated metadata x10 and x100, keeping the original distribution

Segmentation pipeline for demo

Segmentation Pipeline

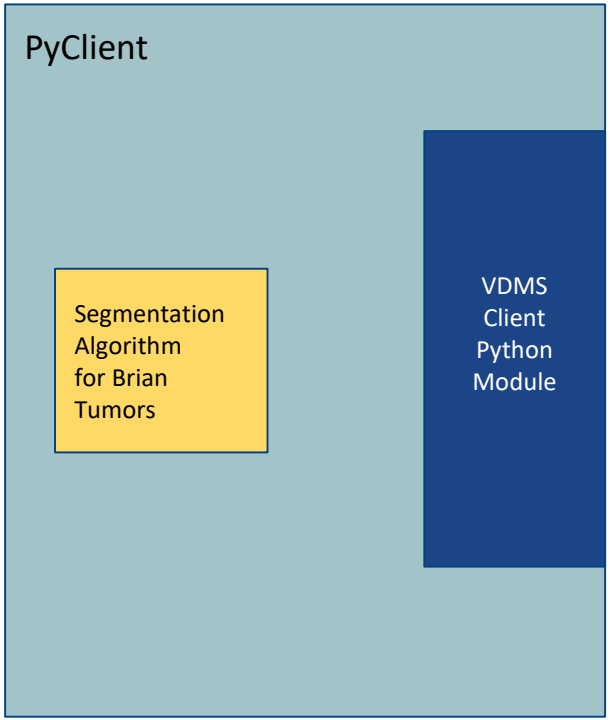

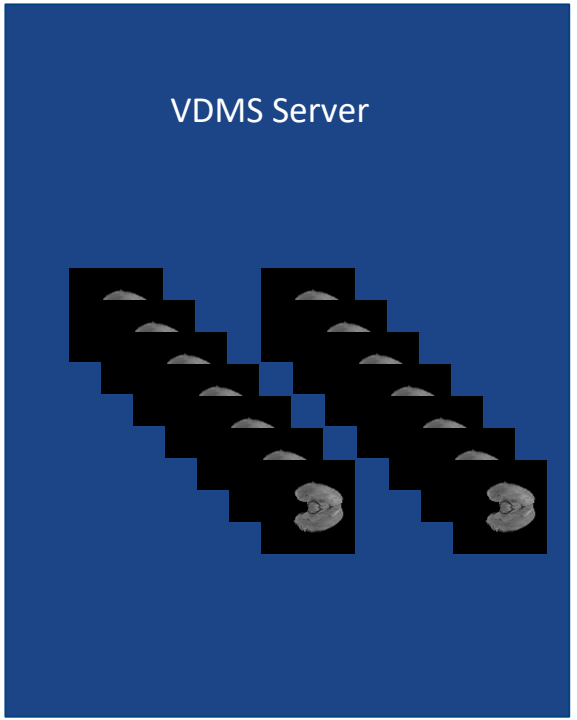

Segmentation Pipeline

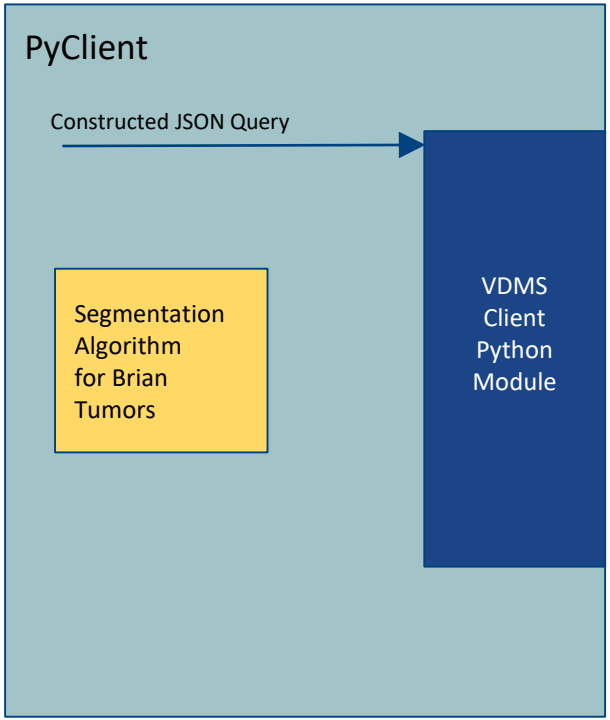

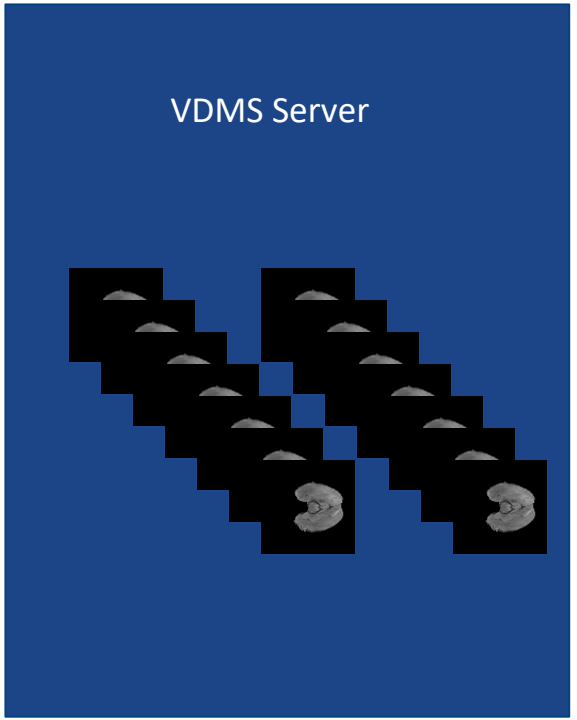

Segmentation Pipeline

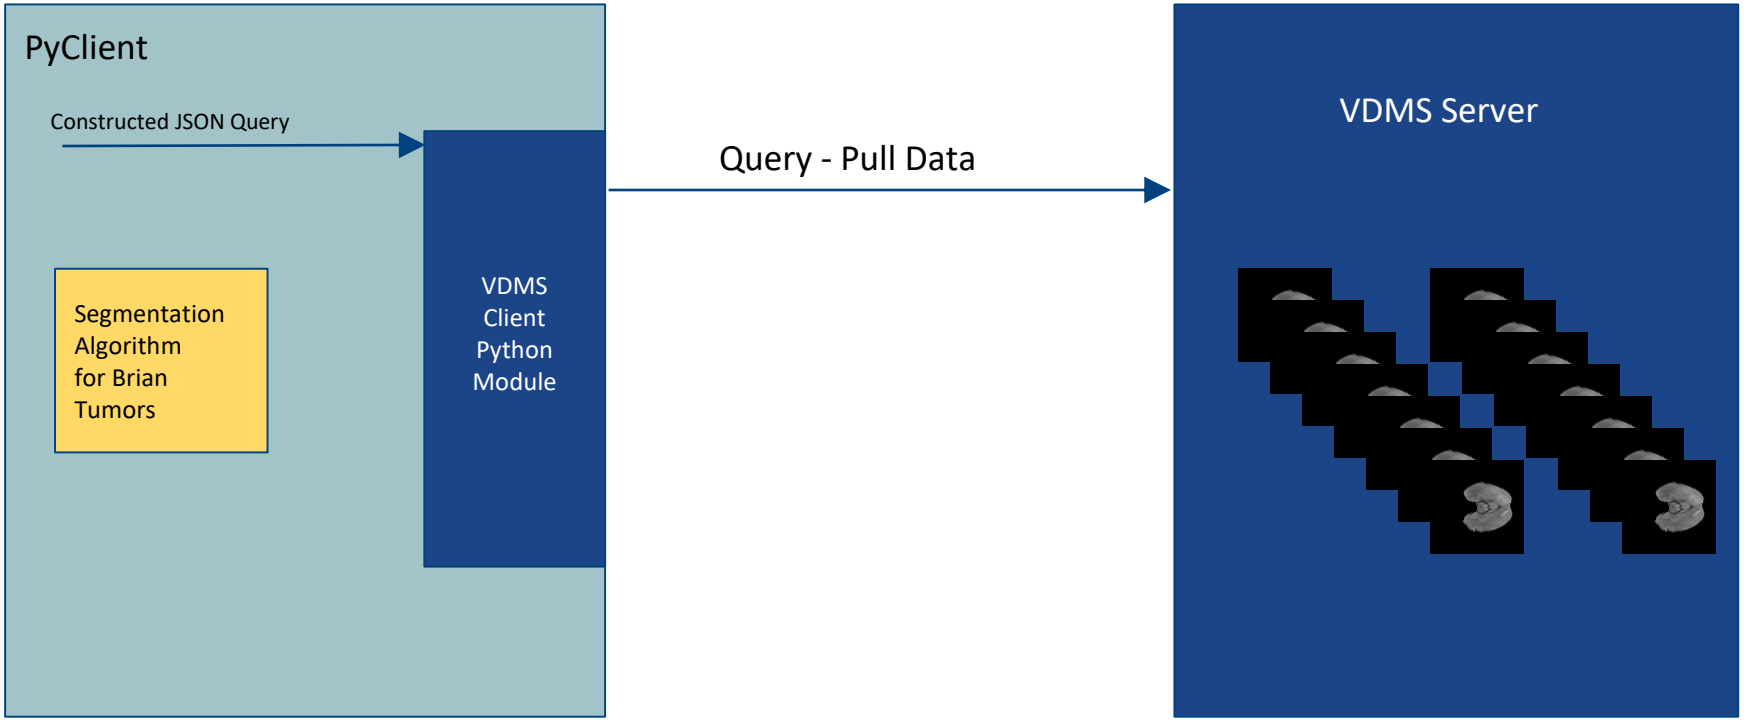

Segmentation Pipeline

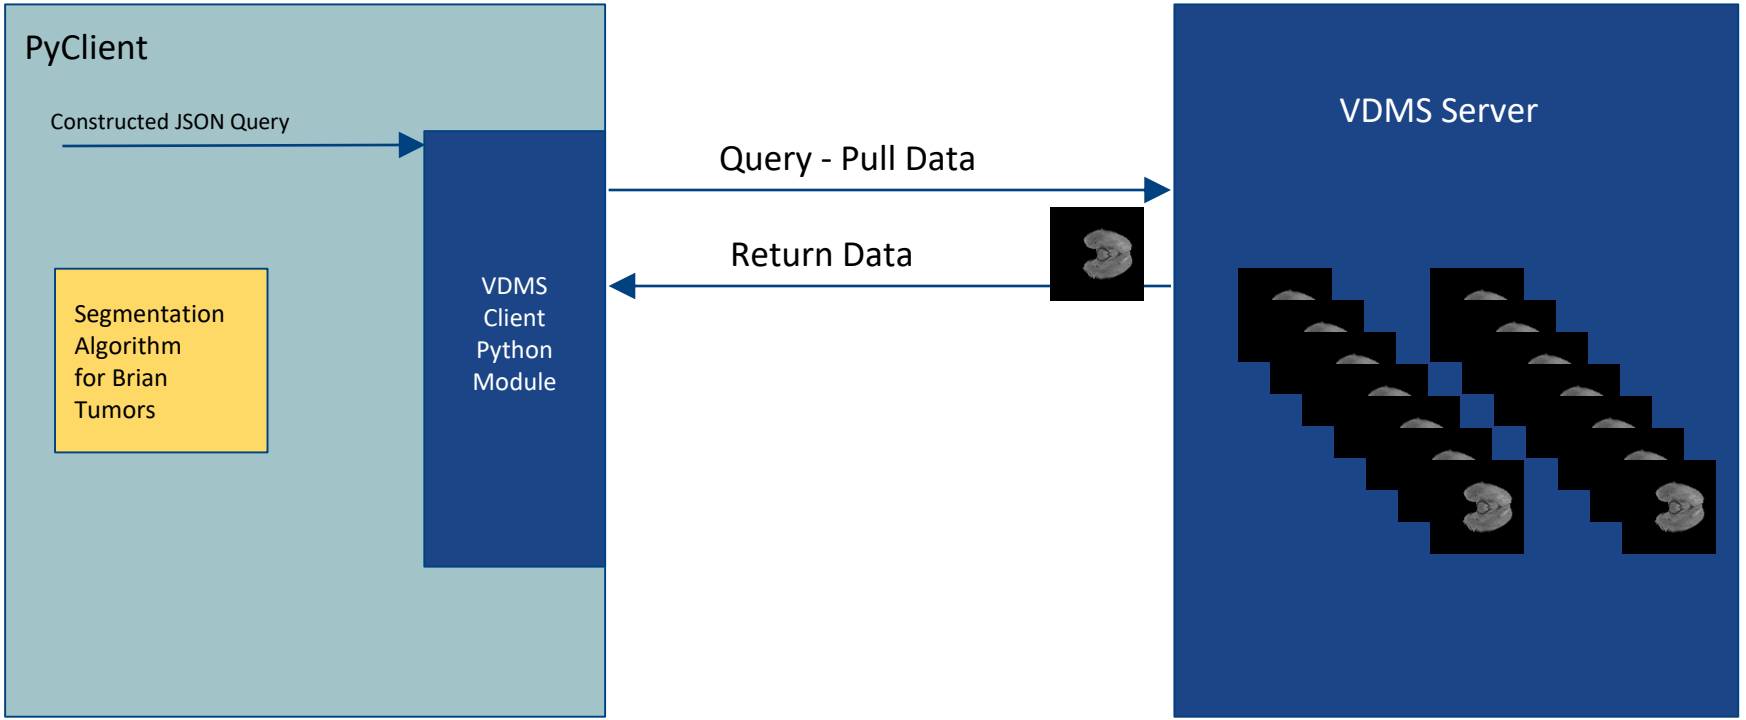

Segmentation Pipeline

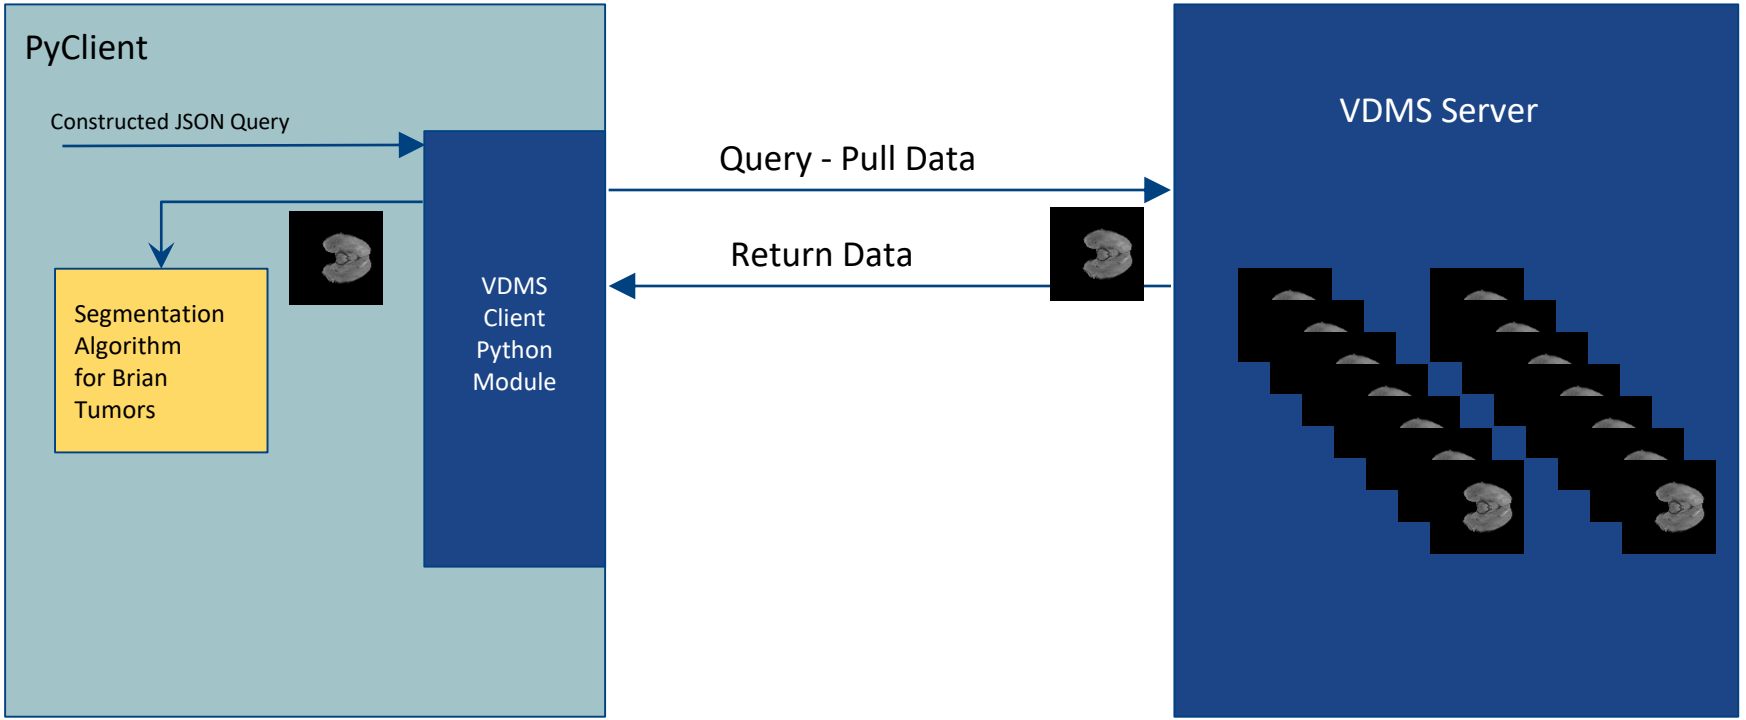

Segmentation Pipeline

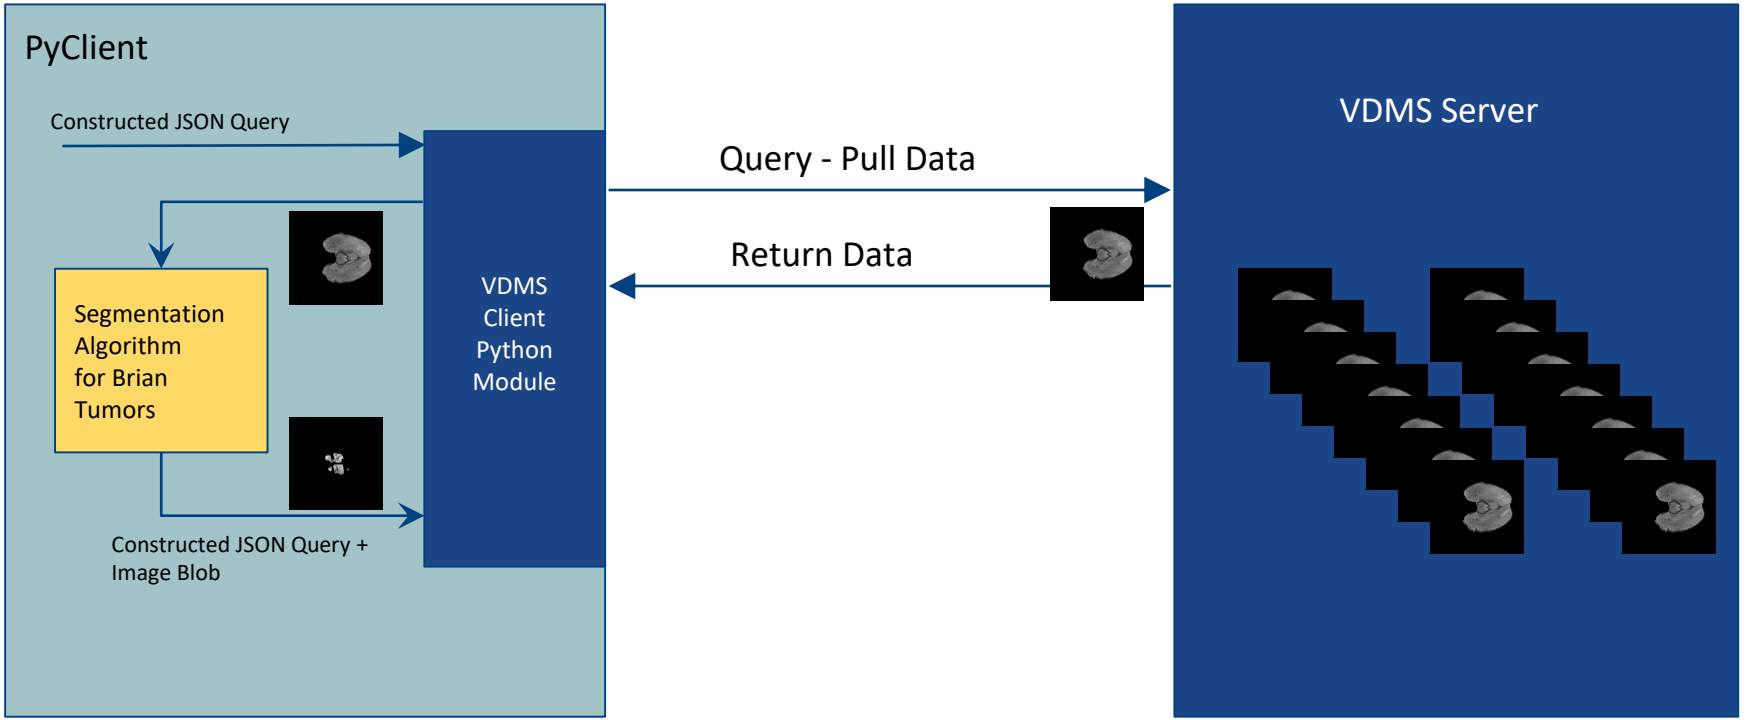

Segmentation Pipeline

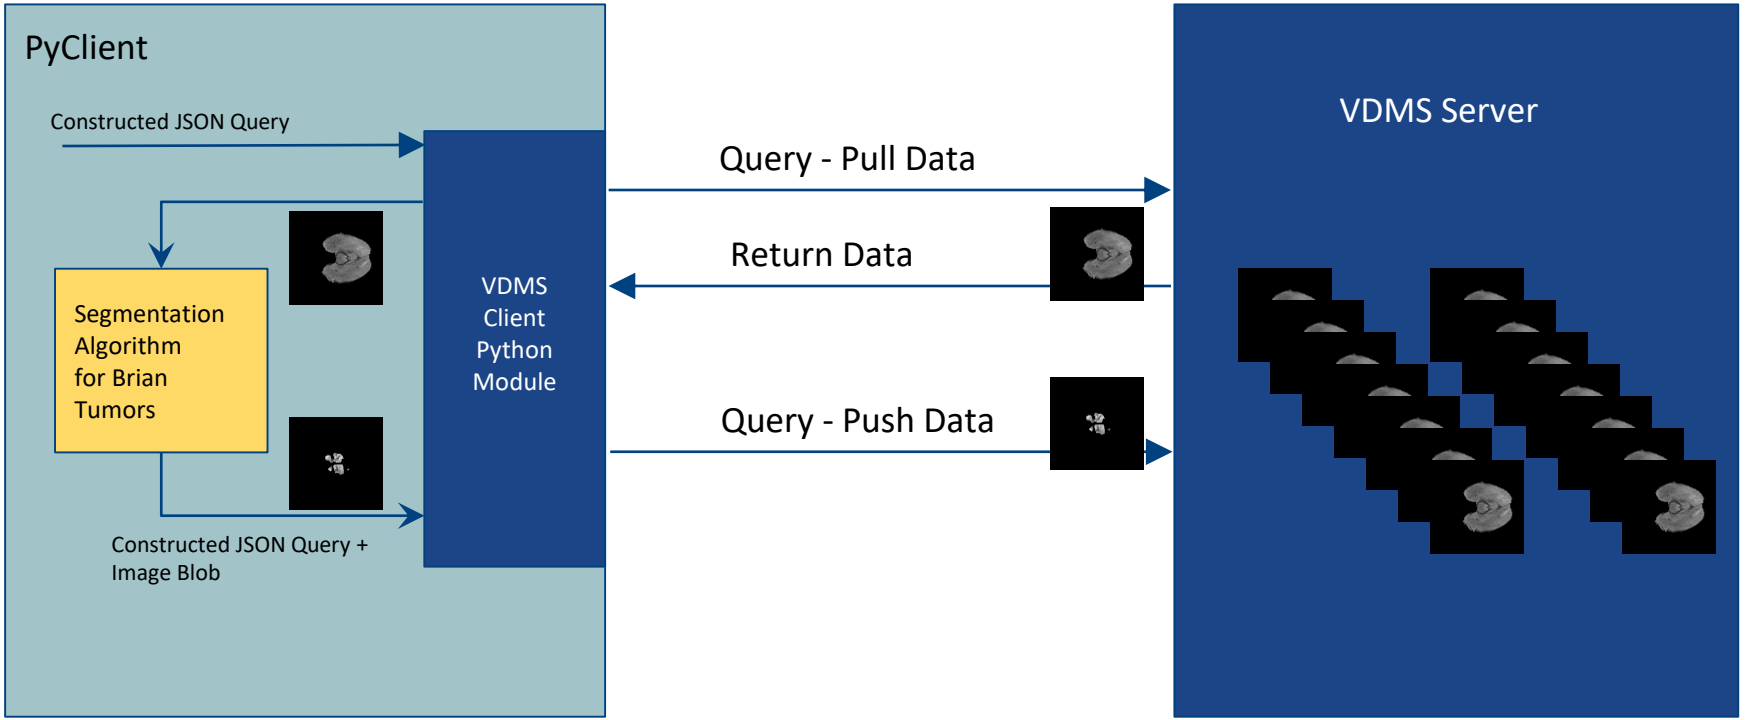

Segmentation Pipeline

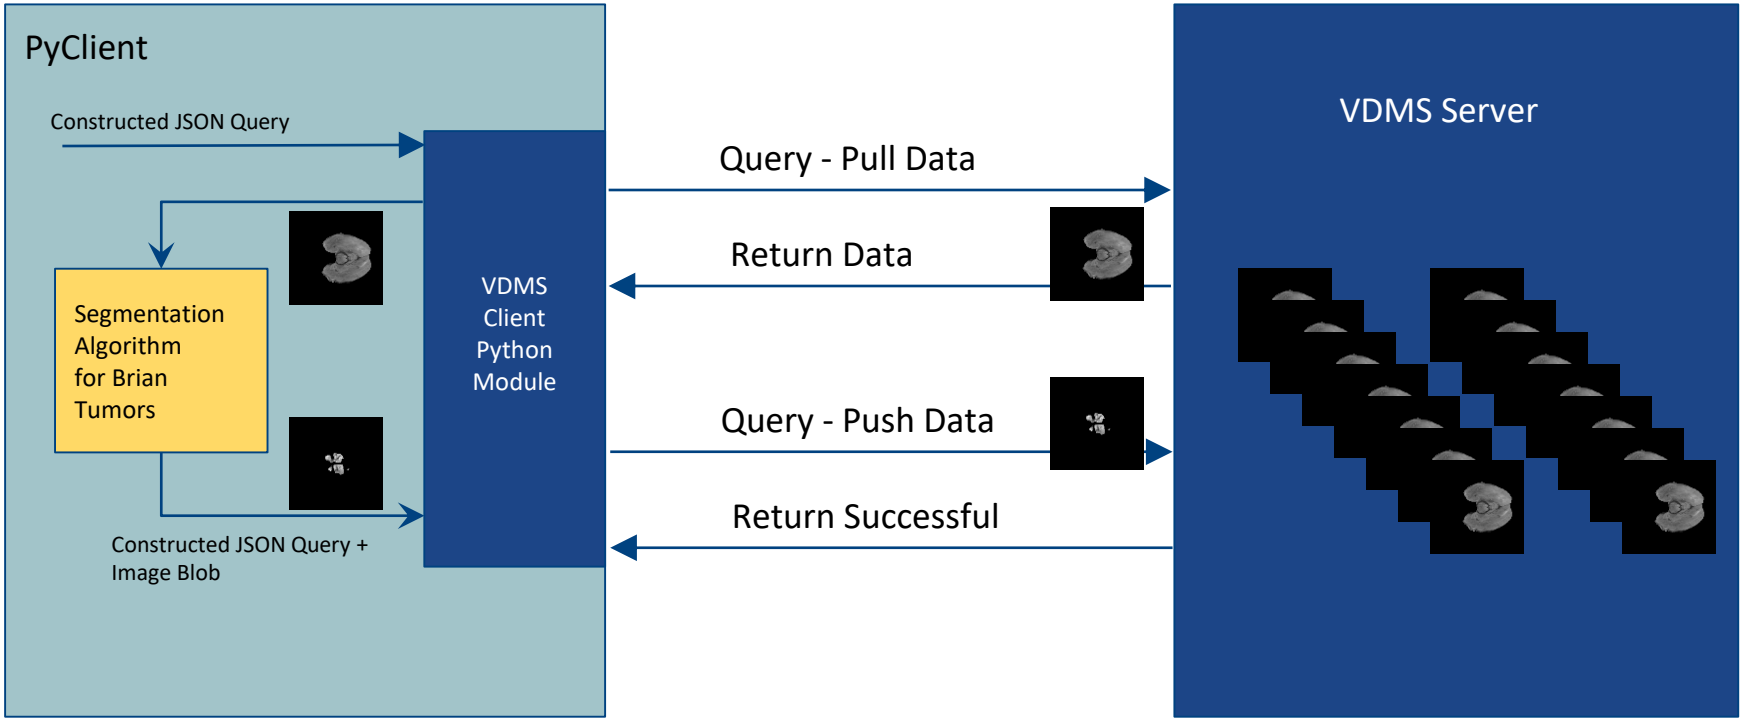

Domain Specific Queries - Some Examples

Query 1: Retrieve a single image (200x200), searching by its unique name.

→ Retrieve single image

Domain Specific Queries - Some Examples

Query 1: Retrieve a single image (200x200), searching by its unique name.

→ Retrieve single image

Query 2: Retrieve a complete brain scan (155 images) from a particular patient.

→ Retrieve 155 images

Domain Specific Queries - Some Examples

Query 1: Retrieve a single image (200x200), searching by its unique name.

→ Retrieve single image

Query 2: Retrieve a complete brain scan (155 images) from a particular patient.

→ Retrieve 155 images

Query 3: Retrieve all brain scans corresponding to people over 75 who had a chemotherapy using the drug “Temodar”.

→ Retrieve 1600 images after 3 neighbor hops

Comparison Baseline

No single solution to compare

Comparison Baseline

No single solution to compare

Create “likely” combination of well-known options

- MemSQL for storing metadata
- Apache HTTP server for requesting images via http
- OpenCV for performing preprocessing

Performance Improvements - Metadata

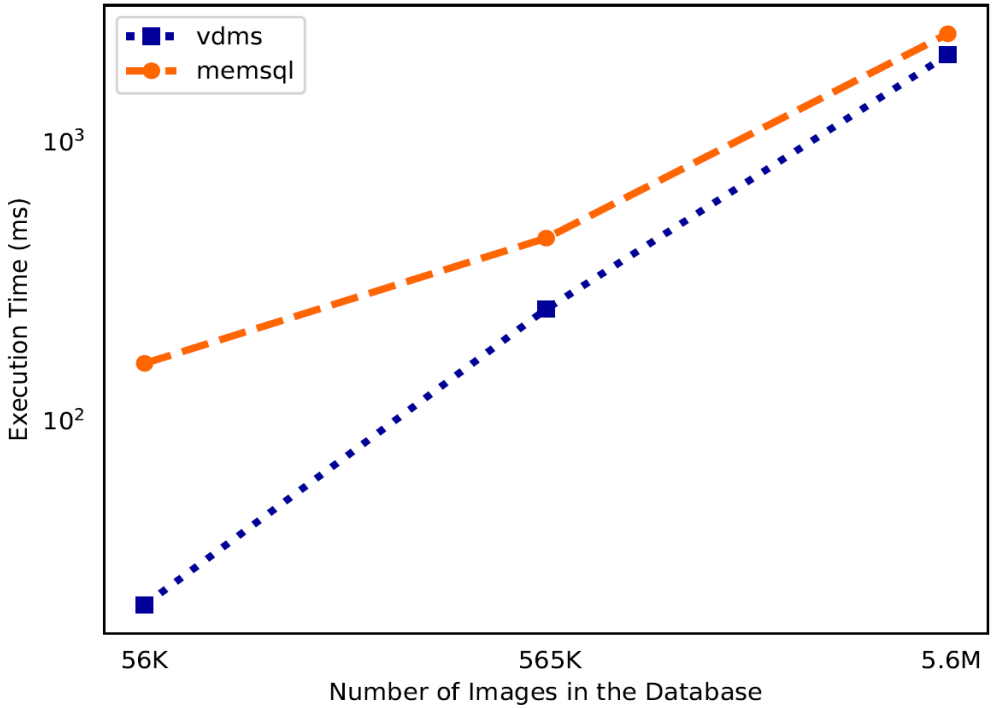

Query 3: Retrieve 1600 image names after 3 neighbor hops

Performance Improvements - Metadata

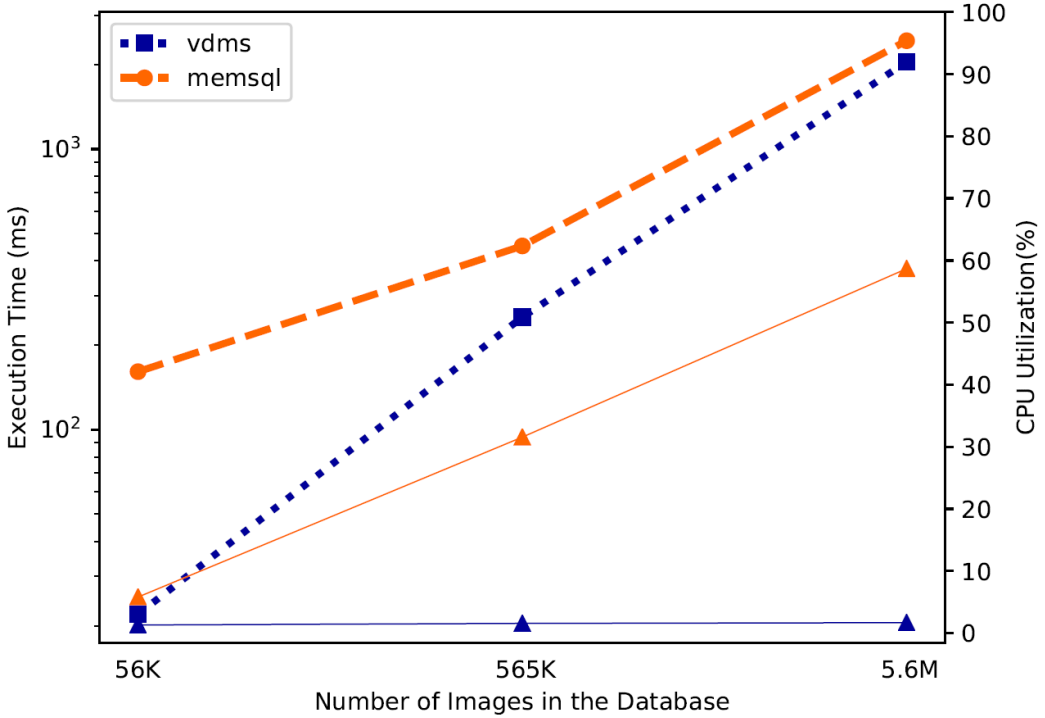

Query 3: Retrieve 1600 image names after 3 neighbor hops

Performance Improvements - Metadata

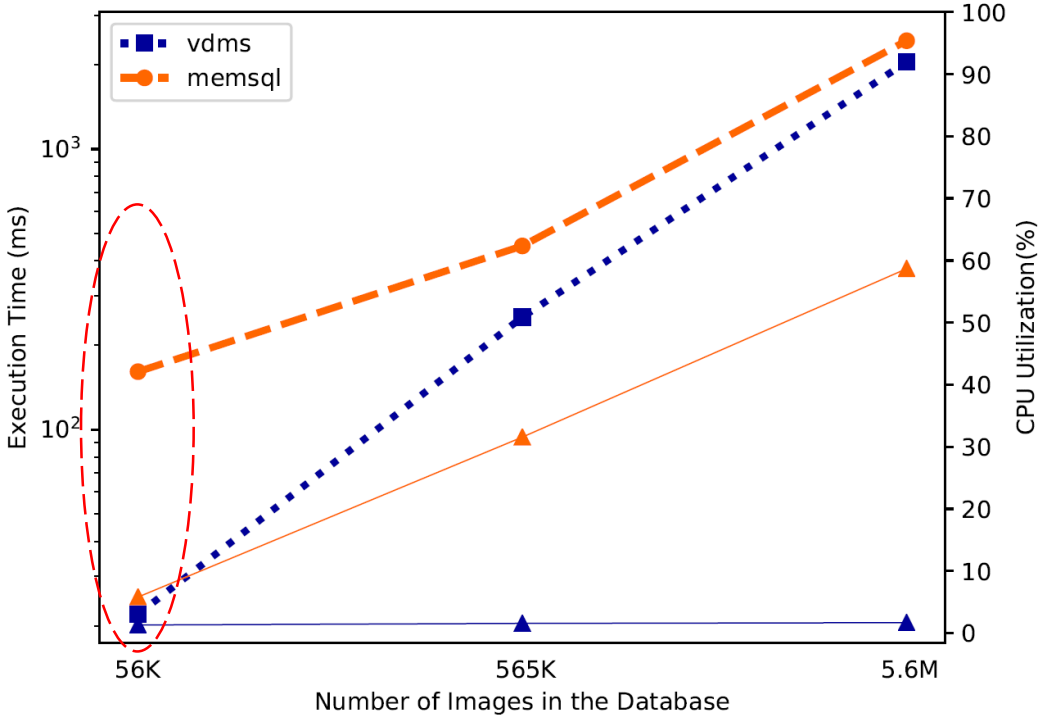

Query 3: Retrieve 1600 image names after 3 neighbor hops

Performance Improvements - Metadata

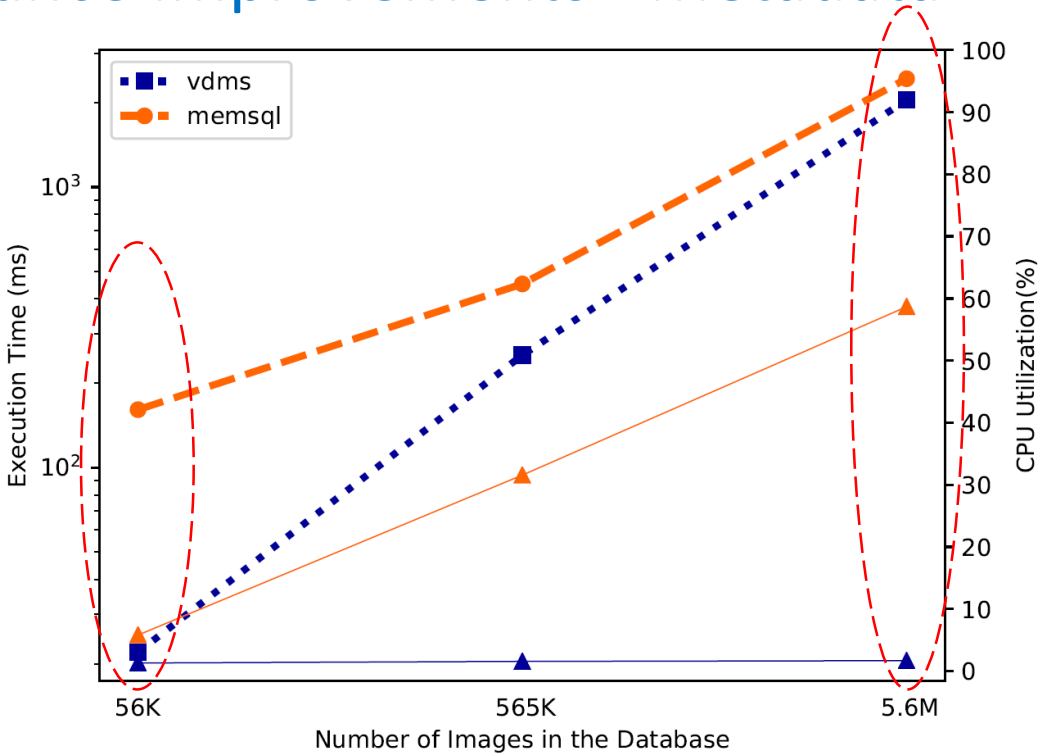

Query 3: Retrieve 1600 image names after 3 neighbor hops

VDMS performs up to one order of magnitude better compared to MemSQL
A Graph Database is a logical choice for visual metadata.

Visual Compute Library: E.g. Transformation Operations

Images in Analytics-friendly TDB Format (uses TileDB)

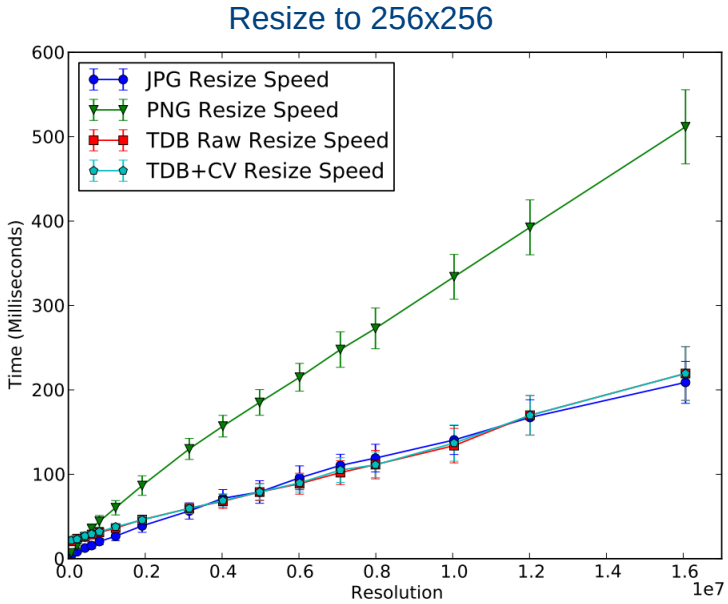

Visual Compute Library: E.g. Transformation Operations

Images in Analytics-friendly TDB Format (uses TileDB)

Resize to 256x256

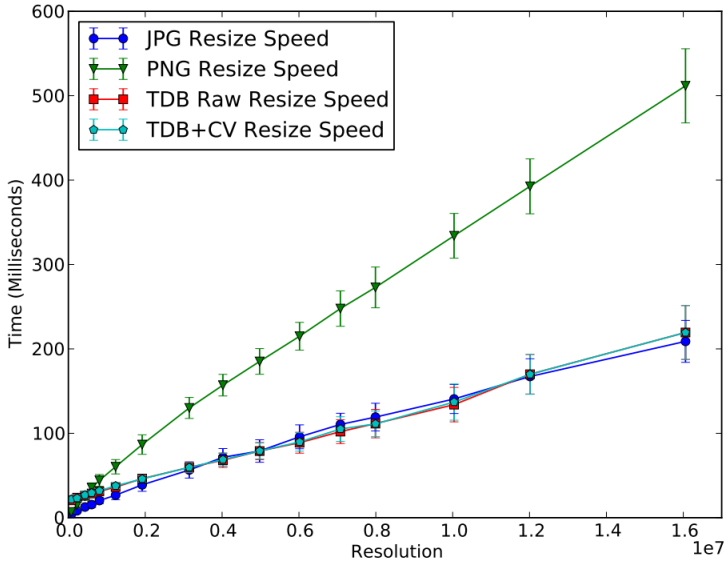

Crop to one-sixth the size

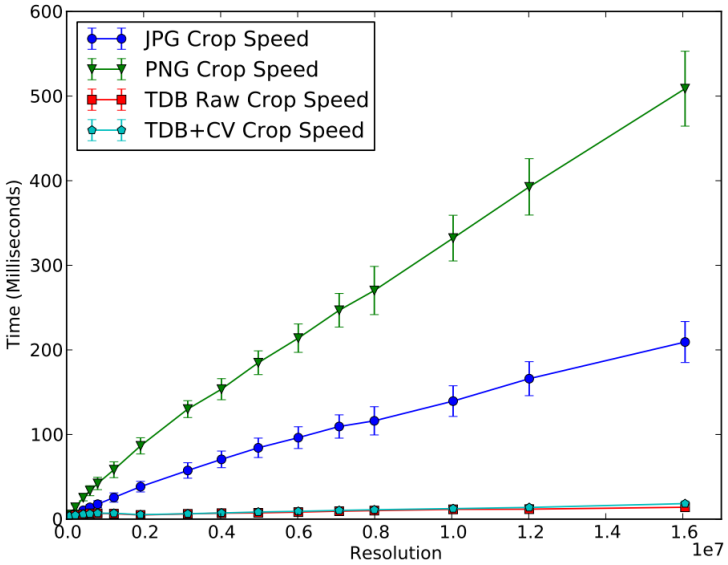

Visual Compute Library: E.g. Transformation Operations

Images in Analytics-friendly TDB Format (uses TileDB)

Resize to 256x256

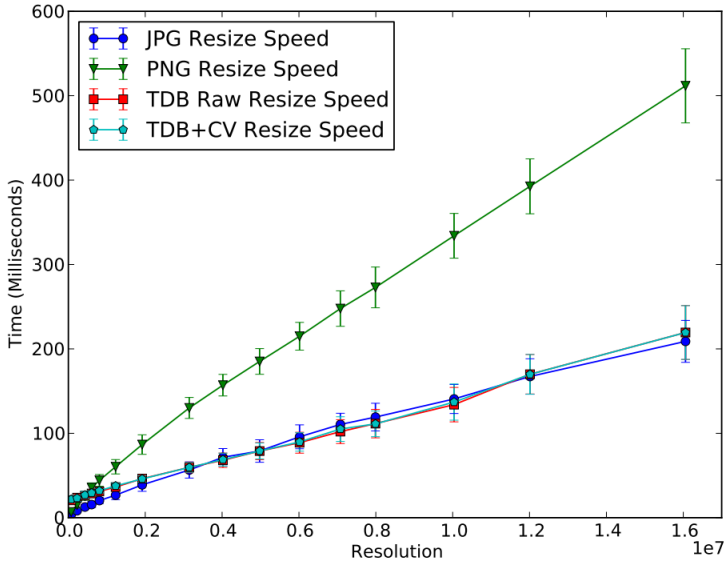

Crop to one-sixth the size

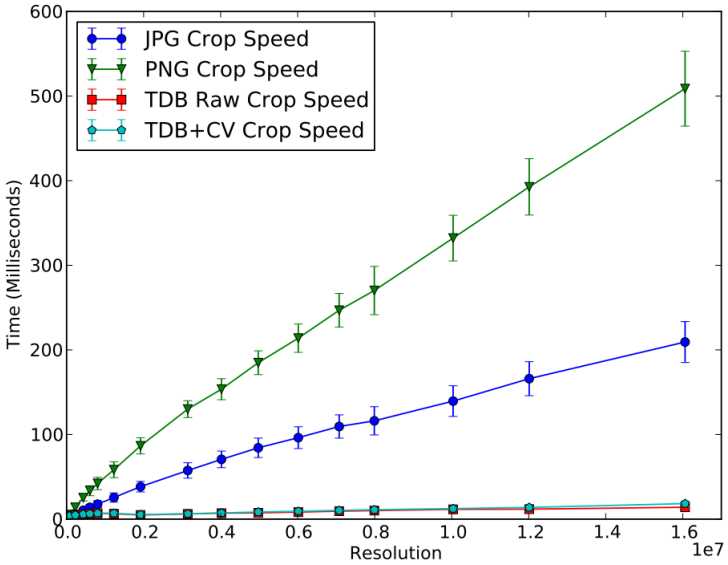

Images stored in the TDB format provide faster access and processing, thus making it a great format for visual analytics pipelines, specially for large images.

Overall Improvements

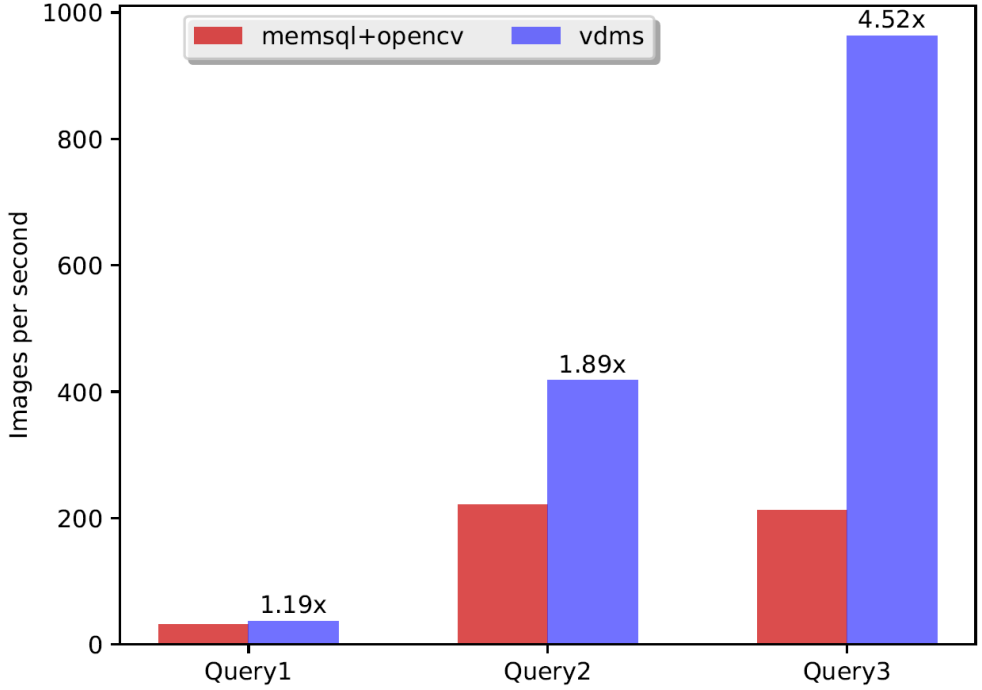

Query 1: Retrieve single image

Query 2: Retrieve 155 images for a patient

Query 3: Retrieve 1600 images after 3 neighbor hops

Overall Improvements

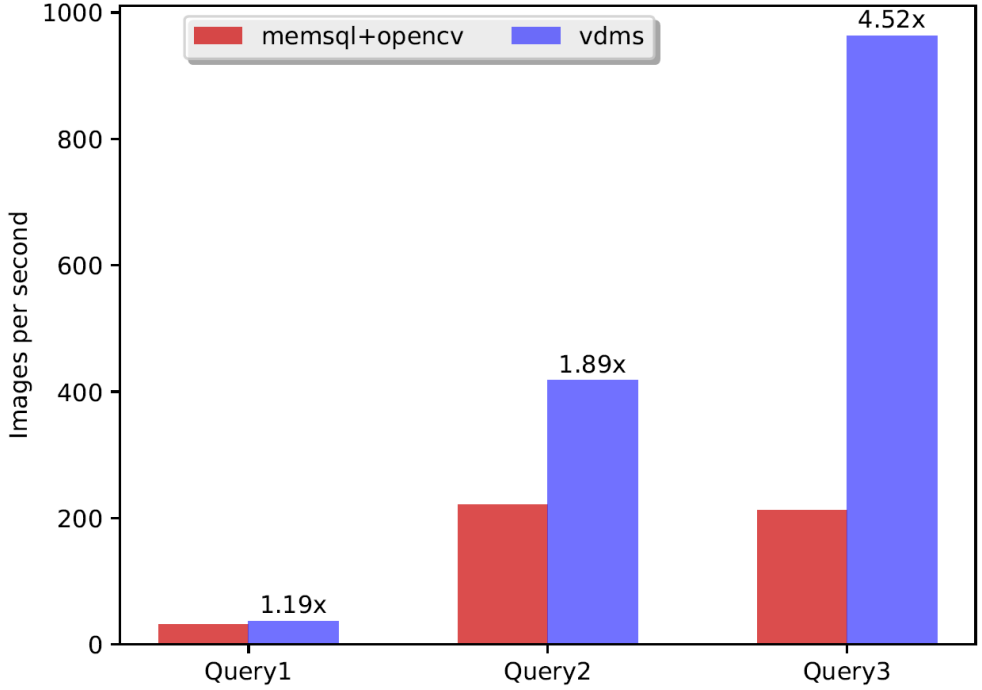

Query 1: Retrieve single image

Query 2: Retrieve 155 images for a patient

Query 3: Retrieve 1600 images after 3 neighbor hops

VDMS performs significantly better when dealing with more complex queries, without incurring in overhead in more simple tasks

Hermes Peak: A Framework for Ad-hoc Video Analytics

Framework for processing visual data from the edge to cloud with four focus areas within the Intel Science and Technology Center for Visual Cloud Systems

In-line Processing

- Video processing with real time turnaround
- Support arbitrary number of streams
- Programmable events
- Optimized resource utilization

Optimized Storage and Retrieval

- Optimized metadata DB
- Analysis friendly media formats
- Distributed for cloud scale
- Tiered storage for hot and cold data

Offline Processing

- Query and analytic on historic (stored) data
- Processing of large (cloud) scale video or image libraries
- Optimized resource utilization

Query Processing and Configuration

- Tools to configure pipeline and answer queries
- Visual query compiler
- Visual kernel repository

Bigger Picture: Visual Cloud Inferencing Flow

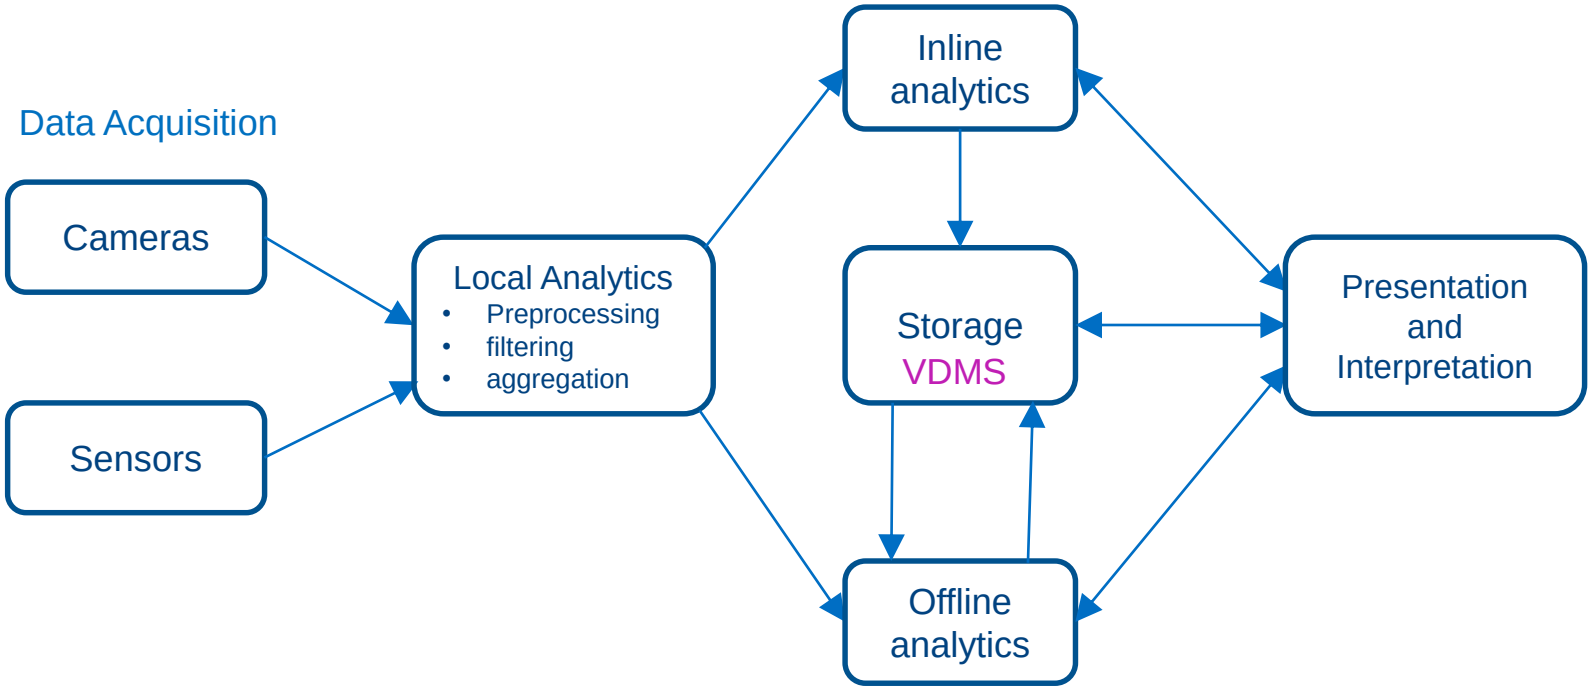

Bigger Picture: Visual Cloud Inferencing Flow

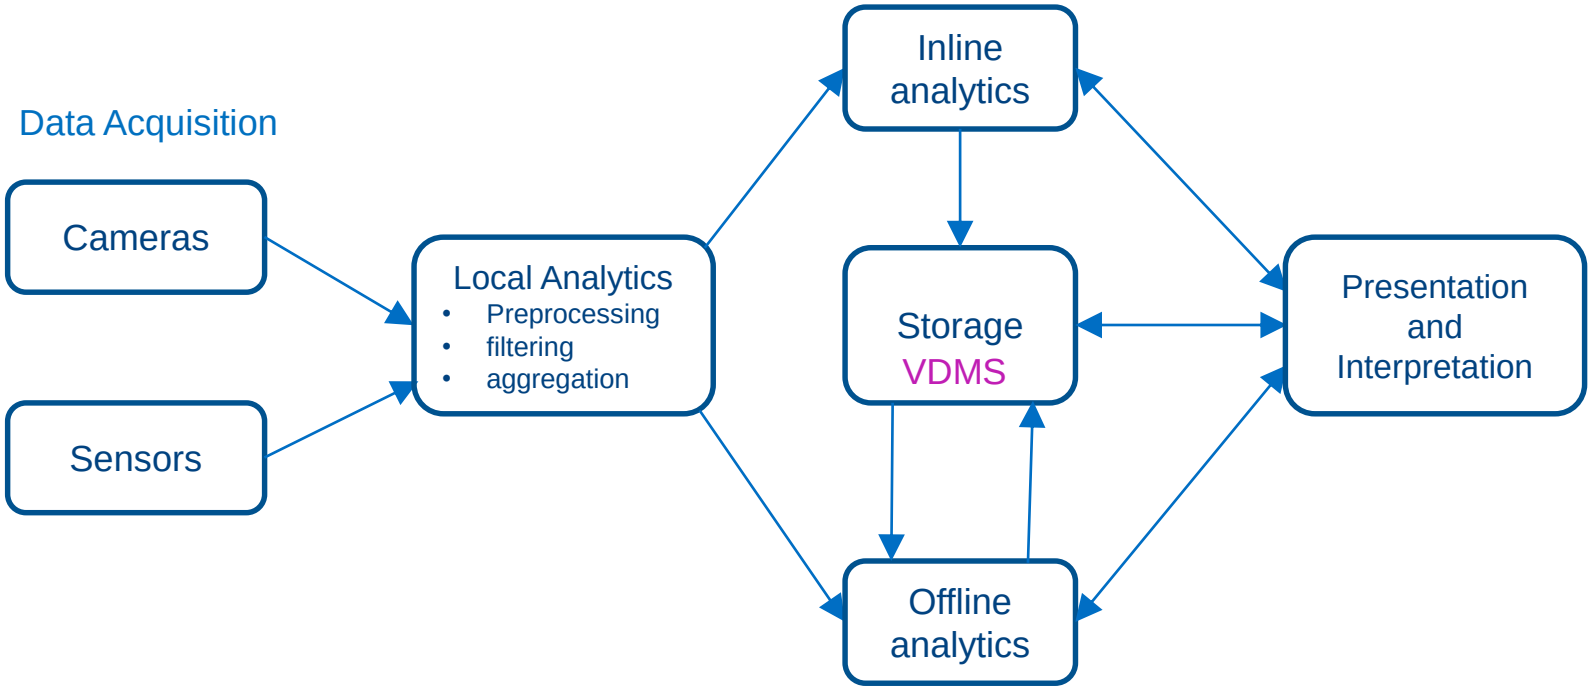

Along with our academic partners, Intel Labs is looking at the entire flow of visual data and processing from edge to cloud

Hermes Peak: A Framework for Ad-hoc Video Analytics

In-line Processing

- Video processing with real time turnaround
- Support arbitrary number of streams
- Programmable events
- Optimized resource utilization

E.g. Streamer (https://github.com/viscloud/streamer)

Optimized Storage and Retrieval

- Optimized metadata DB
- Analysis friendly media formats
- Distributed for cloud scale
- Tiered storage for hot and cold data

E.g. VDMS

Offline Processing

- Query and analytic on historic (stored) data
- Processing of large (cloud) scale video or image libraries
- Optimized resource utilization

E.g. Scanner (https://github.com/scanner-research/scanner)

Query Processing and Configuration

- Tools to configure pipeline and answer queries
- Visual query compiler
- Visual kernel repository

TBD

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

Address feature vector and video storage and search

Scale out to sustain large amount of data and high rates

- Also integrate with pub/sub model (Kafka) and evaluate

Next version of the API and open source code

Hermes Peak integration to complete a visual pipeline

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

Address feature vector and video storage and search

Scale out to sustain large amount of data and high rates

- Also integrate with pub/sub model (Kafka) and evaluate

Next version of the API and open source code

Hermes Peak integration to complete a visual pipeline

Backup

Intel Labs

Extracting Value from Visual Data – Machine Learning

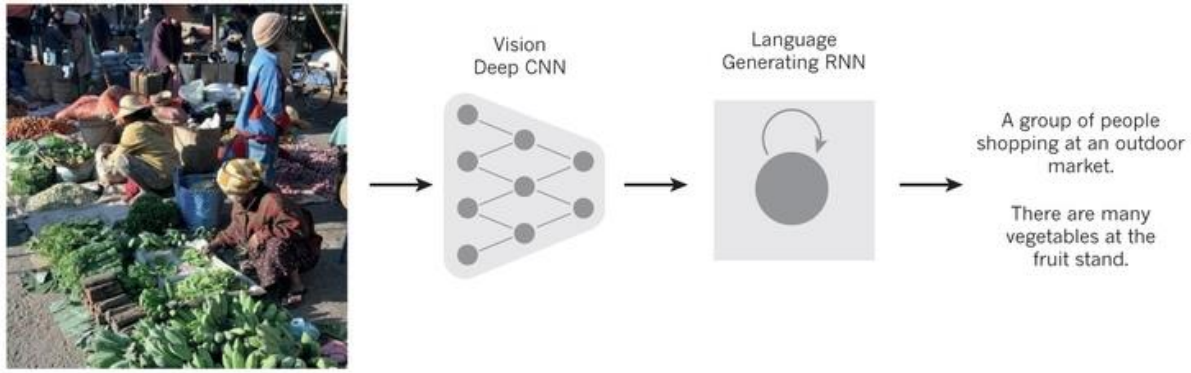

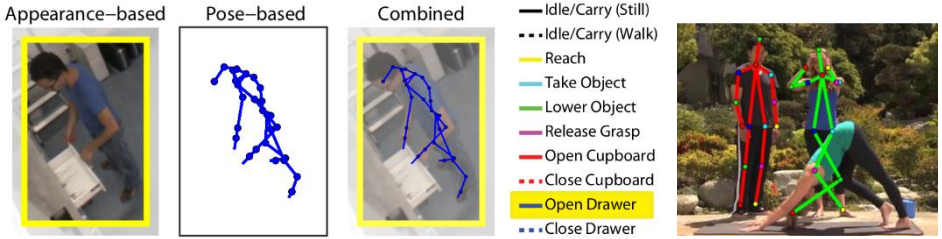

Scale - Ubiquitous Cameras, New Applications

5K cameras surround the venue

With more than thirty 5K cameras surrounding each venue, Intel® freeD™ technology delivers the clearest high-def images possible. This technology uses voxels (pixels with volume) to render replays in spectacular 3D, creating a multi-perspective view of key moments.

Despite Computing Challenges, Data Access Can't be Ignored E.g. Image Classification using Deep Learning

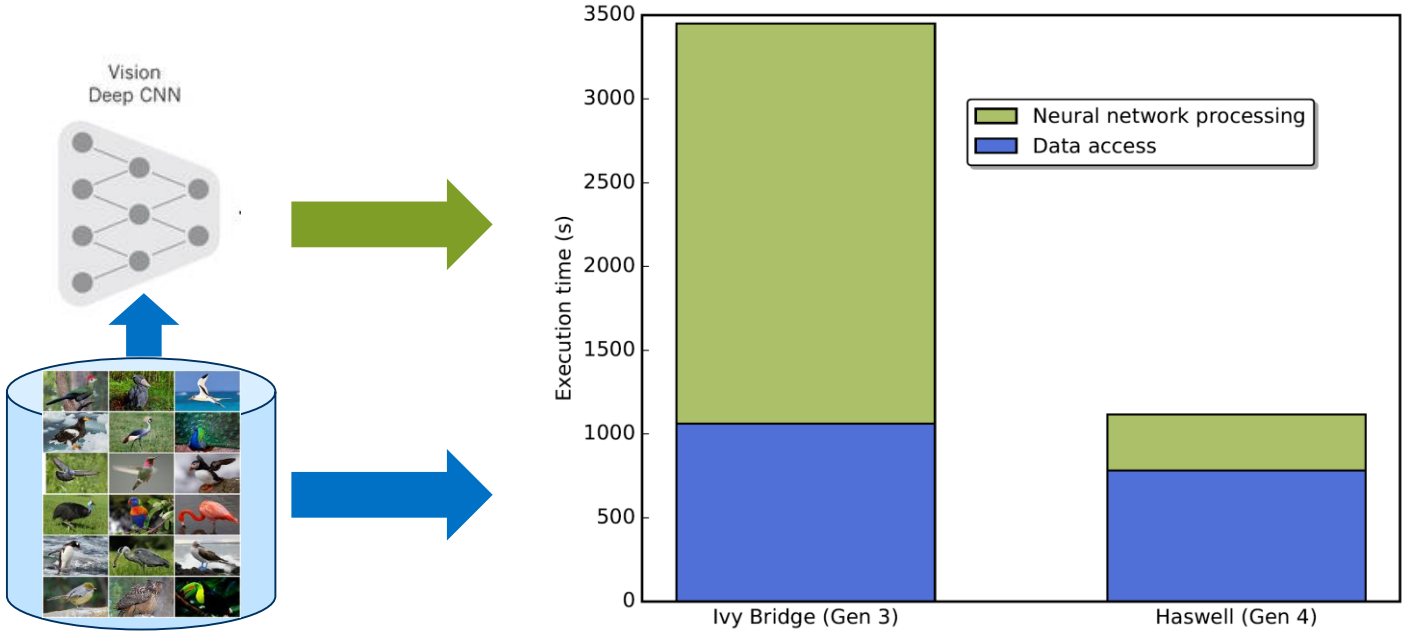

As processing capabilities and algorithms improve, amount of data increases, and data reuse becomes a possibility, data access goes from an afterthought to a real challenge

Exploit Rich Visual Metadata

Media data easily leads to 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

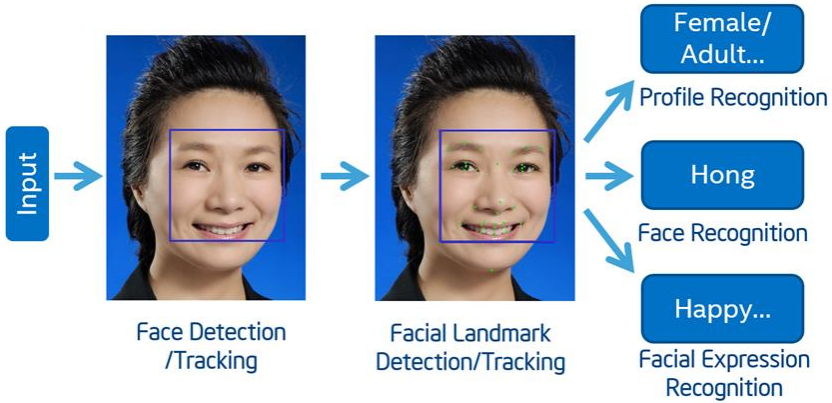

Search photos by faces, scenes, objects, and actions/events

Source: Yurong Chen, Intel Labs China

Representing Media Metadata

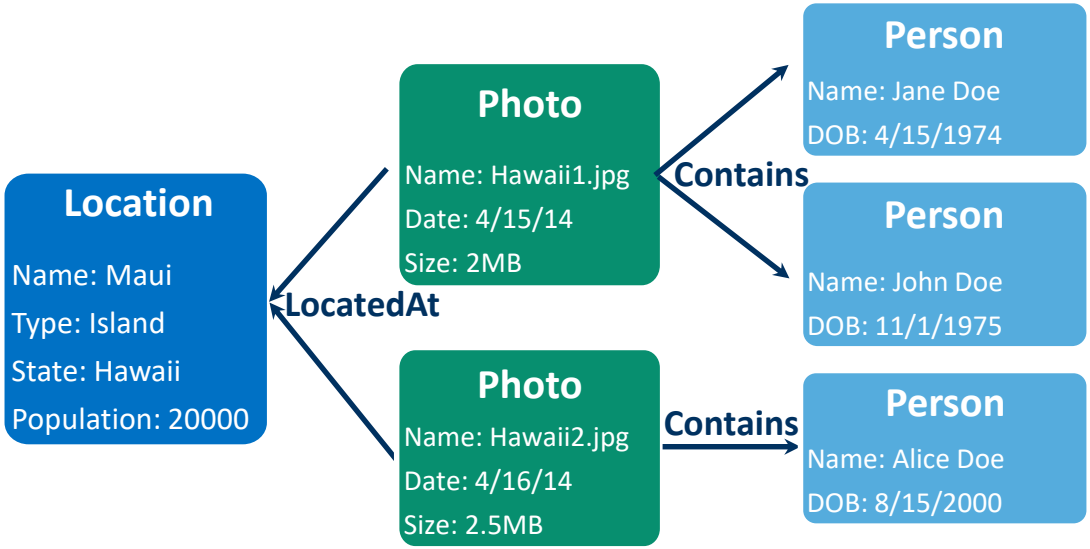

Find all photos of Alice from Hawaii

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

PMGD Comparison to Neo4j

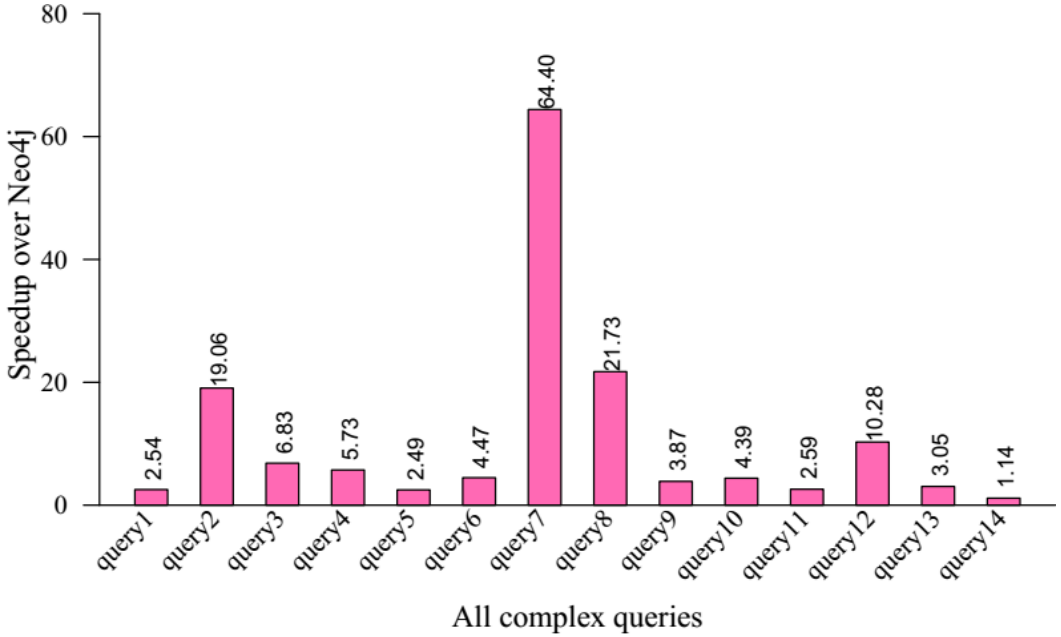

Queries taken from the LDBC social network benchmark

Bars show speedup over Neo4j

The more graph traversals, the better PMGD does

Speeding up Access to Desired Data

More and more machine consumption of data for processing

- Think beyond standard formats for visual data
- Create formats better suited for processing

Visual Compute Library (VCL)

- Explore alternate formats for images, videos and feature vectors
- Implement suitable processing on traditional and new formats

VCL::Image

Implement alternate image storage formats to use when beneficial

- TDB format, based on TileDB [1]

Higher level interaction with images in traditional or TDB format

- Perform processing such as crop, resize, threshold, ROI access, as data is read

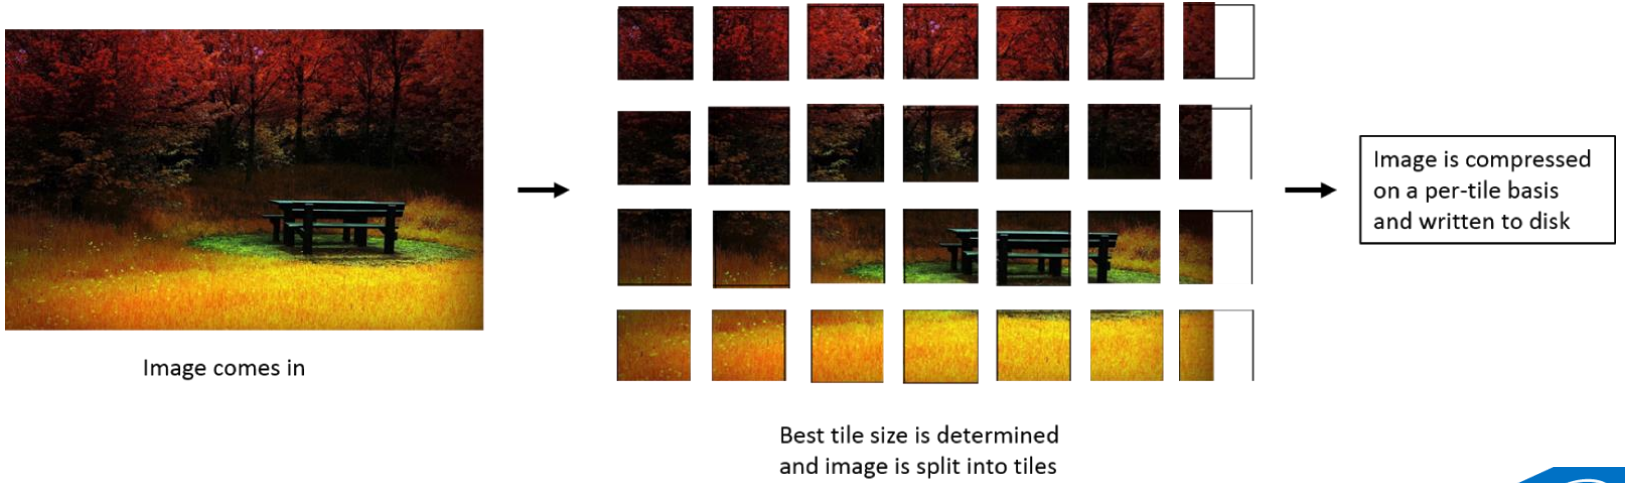

[1] Stavros Papadopoulos, Kushal Datta, et. al. 2016. The TileDB array data storage manager. VLDB 2016

TDB Performance

Write Performance

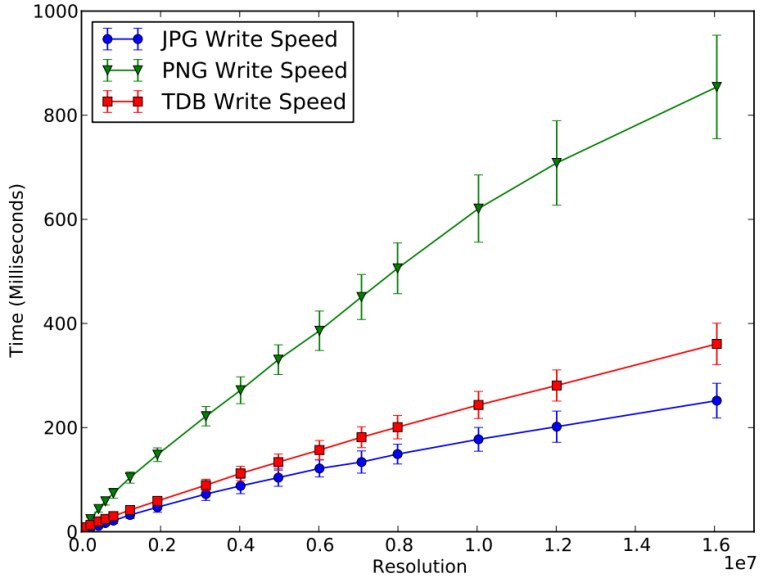

Read Performance

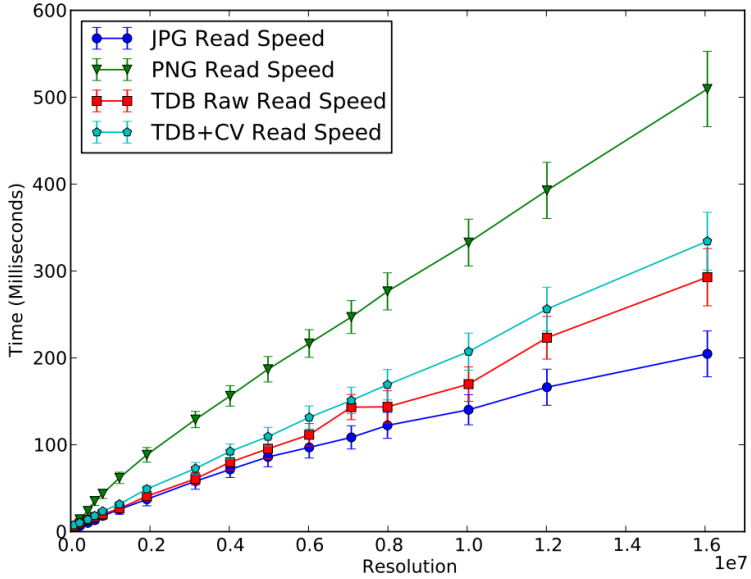

Request Server

Unified and simple client API

Route query to the right component for a coherent user response

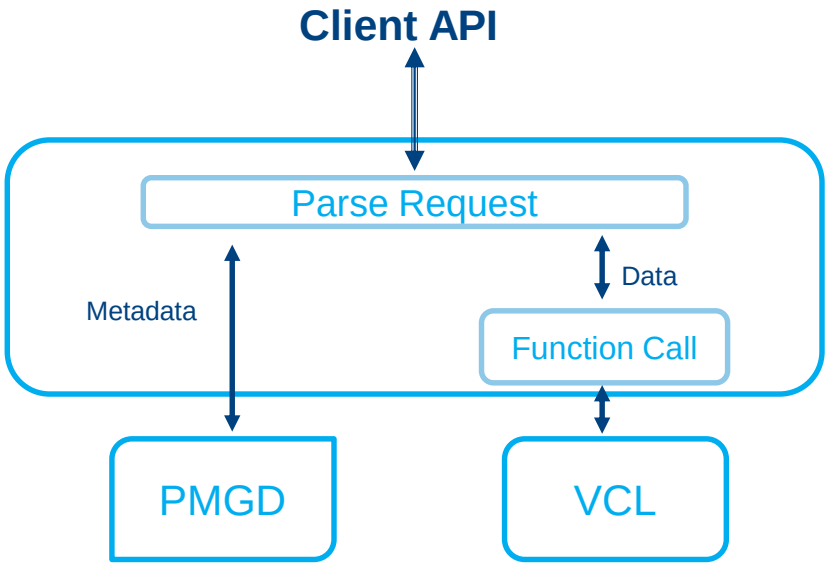

BraTS Challenge - Driving Application

> Personnel > Spyridon (Spyros) Bakas, Ph.D. > MICCAI BraTS 2017: Scope

Multimodal Brain Tumor Segmentation Challenge 2017

• Scope • Relevance • Tasks • Data • Evaluation • Participation Summary • Registration • Previous BraTS • People •

Scope

BraTS has always been focusing on the evaluation of state-of-the-art methods for the segmentation of brain tumors in magnetic resonance imaging (MRI) scans. **BraTS 2017** utilizes multi-institutional pre-operative MRI scans and **focuses on the segmentation of** intrinsically heterogeneous (in appearance, shape, and histology) **brain tumors**, namely gliomas. Furthemore, this year, in order to pinpoint the clinical relevance of this segmentation task, BraTS'17 also focuses **on the prediction of patient overall survival**, via integrative analyses of radiomic features and machine learning algorithms.

IMPORTANT DATES:

5 May	Release of training datasets. — ***(Registration is now closed)***
30 Jun	Release of validation datasets. — ***(Released)*** *View the Leaderboard*

VDMS Alternatives

No one solution to do it all

Intel automotive path

- HDFS for storing data
- Hbase for organizing metadata
- Another layer to make querying using relationships easier

Initial CMU solution

- PostgreSQL database for metadata
- Write their own frame server and use OpenCV
- Still looking for an API

Facebook's Tao + Haystack, Amazon's Neptune + S3

- Large scale but still not optimized for visual data management