VIAME

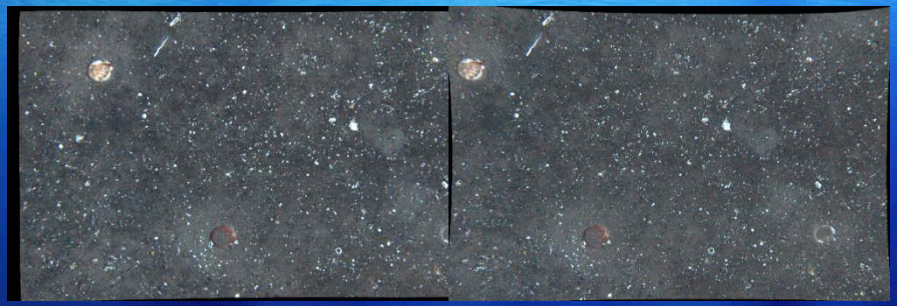


VIAME: An Open-Source Framework for Underwater Image and Video Analytics

Matthew Dawkins[1], Linus Sherrill[1], Keith Fieldhouse[1], Anthony Hoogs[1], Benjamin Richards[2], David Zhang[3], Jon Crall[1], Lakshman Prasad[4], Nathan Lauffenburger[2], Gaoang Wang[5]

[1]Kitware Inc., [2]National Oceanic and Atmospheric Administration, [3]SRI International, [4]Los Alamos National Lab, [5]University of Washington

Kitware Open Source Platforms

- **KWIVER** Kitware Imagery and Video Exploitation and Retrieval
- **VTK** the visualization toolkit
- **ParaView** large data analysis & visualization application
- **ITK** insight image analysis toolkit
- **CMake** cross-platform build system
  - CDash, CTest, CPack, software process tools
- **Resonant/Girder** informatics and information visualization
- **Kiwi & VES** mobile visualization
- IGSTK, CTK, vxl, Open Chemistry Project, VolView, tubeTk, and more…
- **MIDAS** for computational scientific research, testing, and visualization

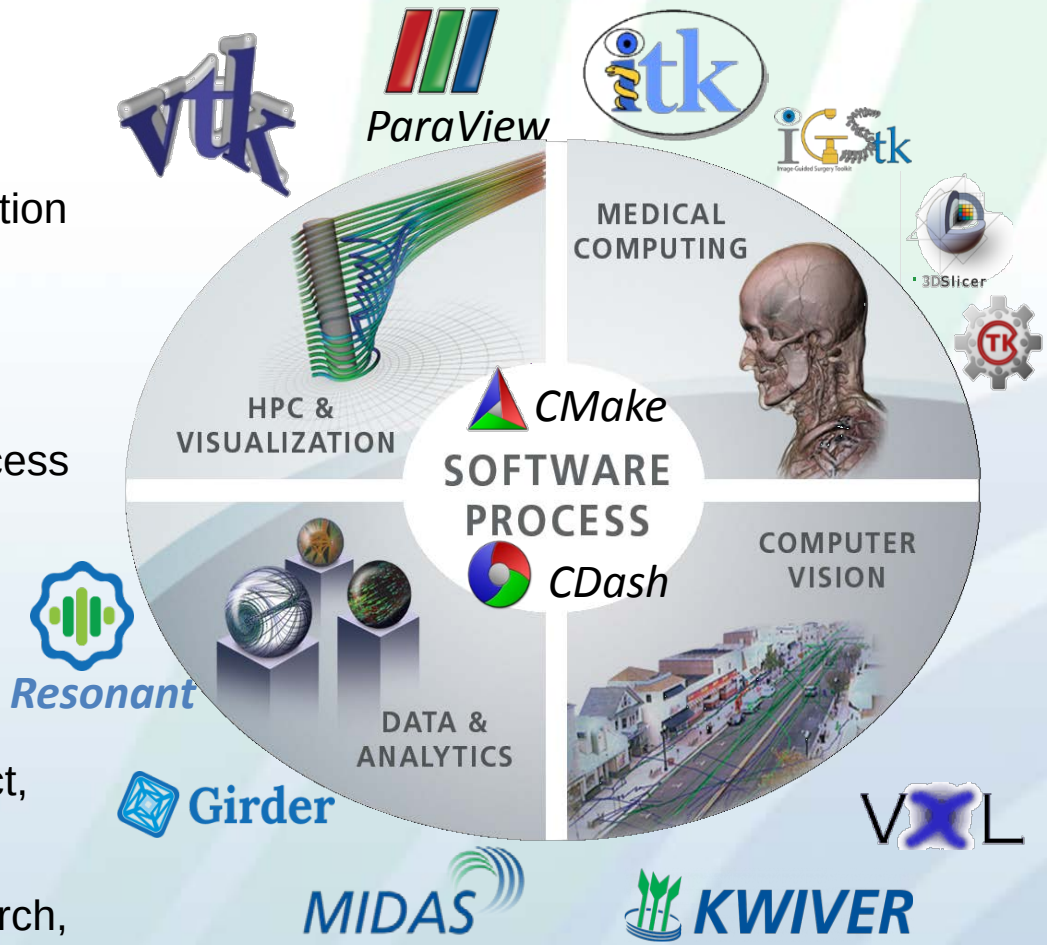

Kitware

Kitware

- 25+ team members
- 12 PhDs
- Founded in 2007
- 35+ contracts

Dr. Anthony Hoogs
anthony.hoogs@kitware.com
518-881-4910

Object and Building Recognition by Function (DARPA)

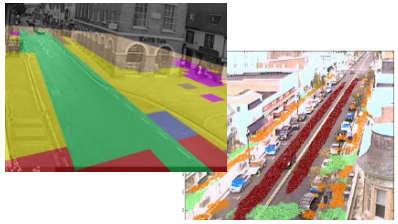

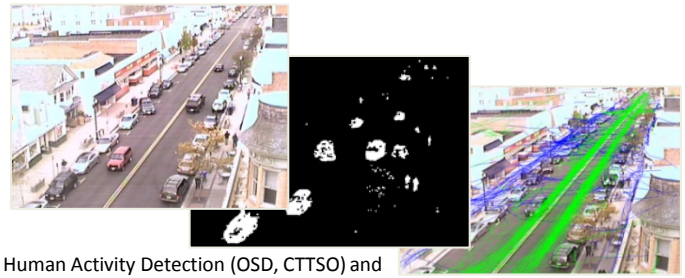

Human Activity Detection (OSD, CTTSO) and Tracking in Wide-Area Video (AFRL)

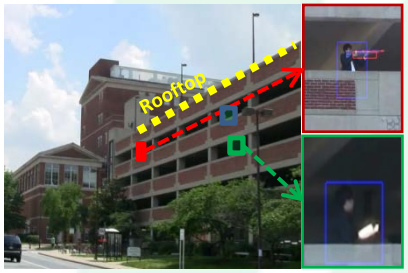

Threat Detection in Video (DARPA)

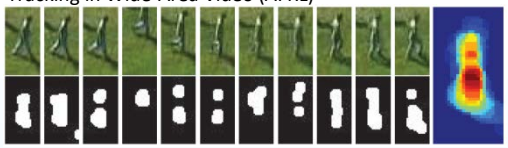

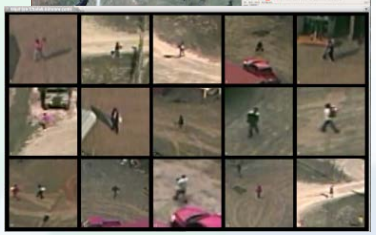

Content-based Video Retrieval by Actions (DARPA VIRAT)

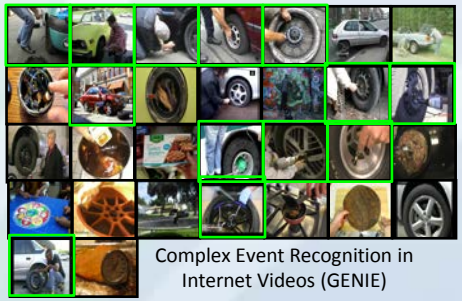

Complex Event Recognition in Internet Videos (GENIE)

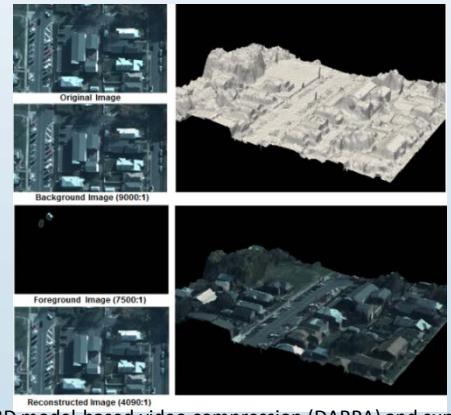

3D model-based video compression (DARPA) and super-resolved 3D reconstruction (DARPA)

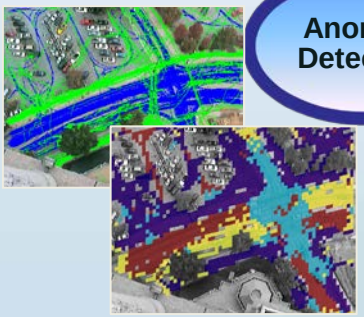

Normalcy Modeling and Anomaly Detection (DARPA PANDA and PerSEAS)

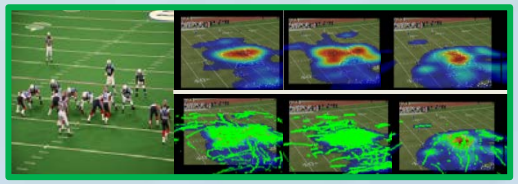

Football Play Recognition (DARPA CARVE)

Wide-Area Motion Imagery Event, Anomaly and Activity Detection (OSD Data to Decisions, DARPA PerSEAS)

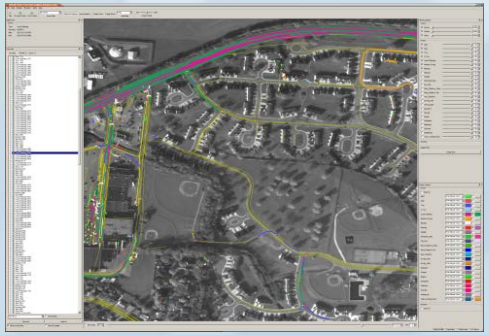

Kitware

VIAME

https://github.com/Kitware/VIAME

- **V**ideo and **I**mage **A**nalytics for the **M**arine **E**nvironment
- Goal: Develop an open-source software platform for NMFS image and video analysis
  - In close coordination with NOAA community

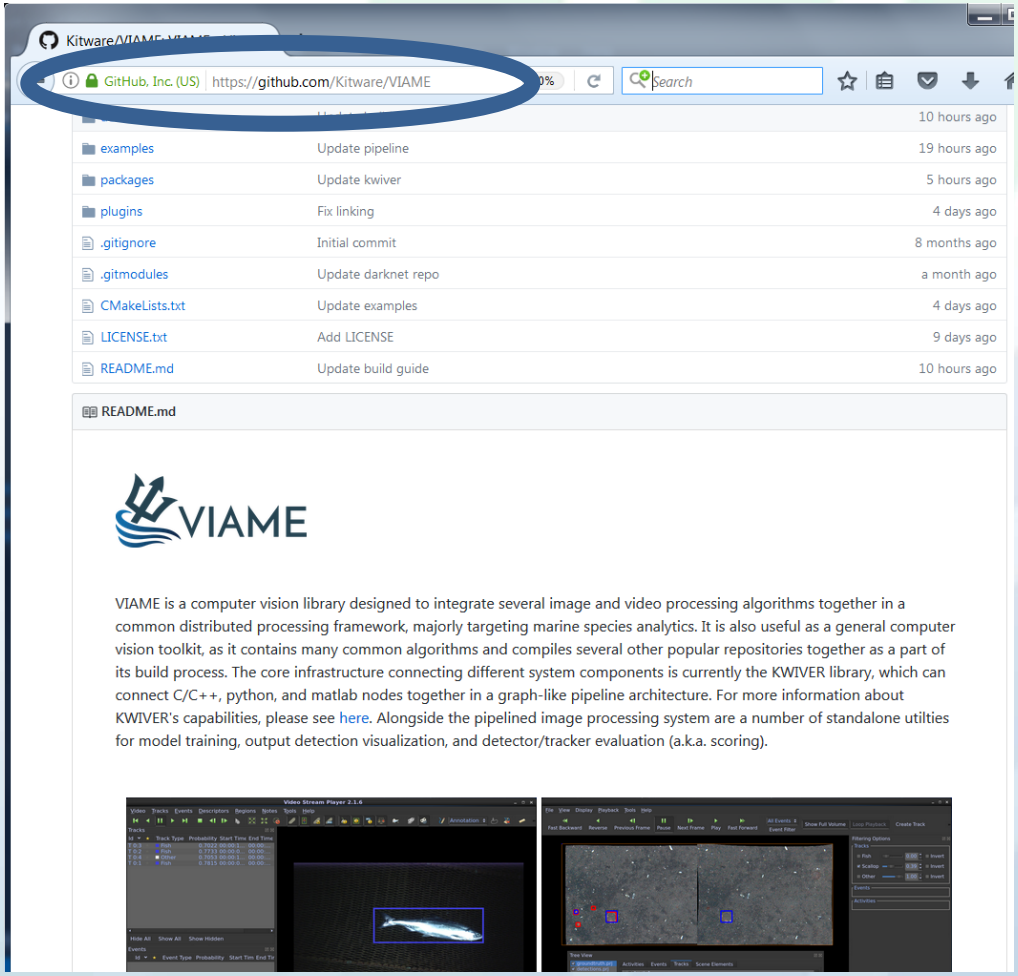

Kitware

Lots of Data, Analytics and Users

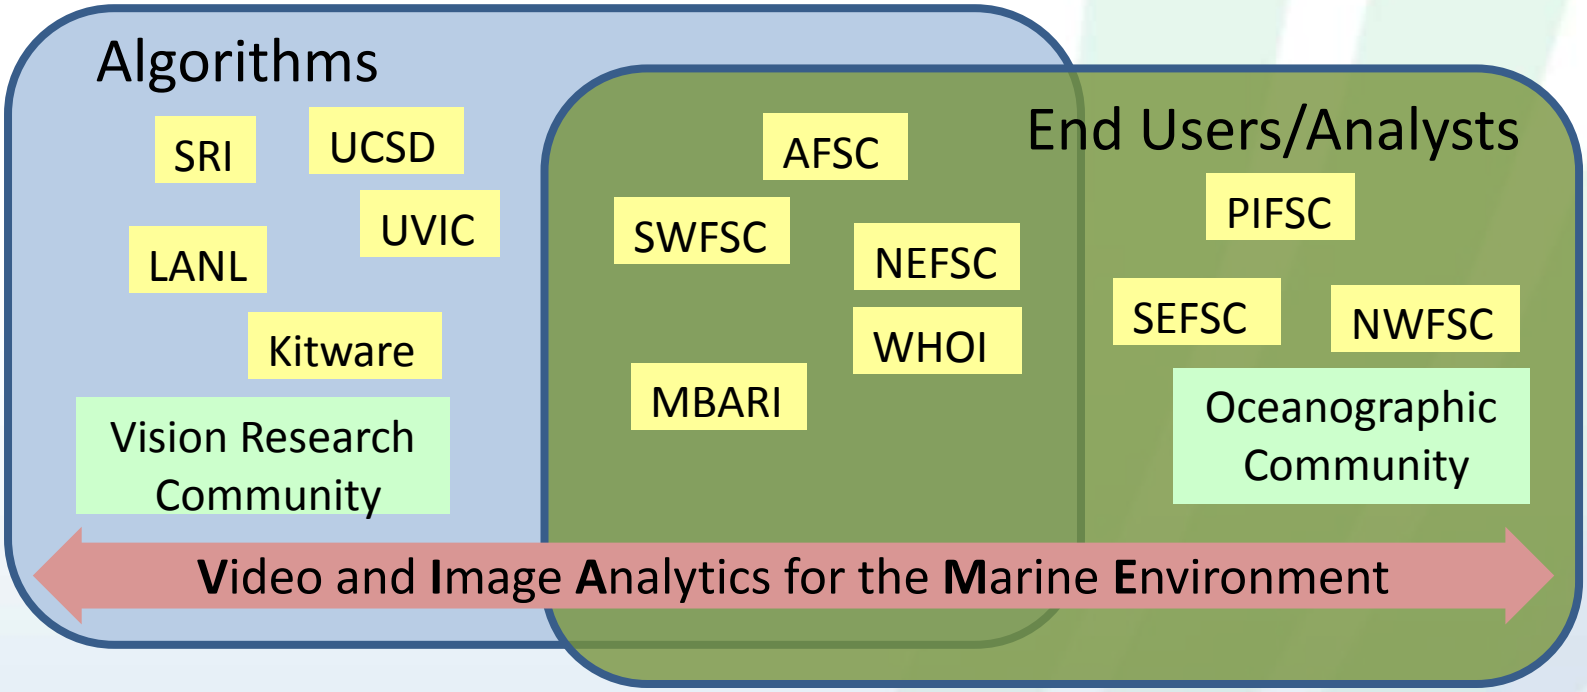

Kitware

Current Image/Video Analytics

Capability	Primary data source	POC	Stereo calibration	Stereo processing	Video	Color, contrast correction	scallop detection	fish detection	fish length, sizing	fish tracking	fish classification	anomaly det.	habitat classification	image segmentation
NW SC CamTrawl	Cam Trawl	Williams	yes	yes	4 Hz	no, grayscale		yes	automatic	yes	yes			
ROV video fish detection and tracking	SWFSC ROV video	Cutter	no	no	30 Hz	yes		yes, DPM (UW)	no	yes (UW student)	desired	desired		
ROV stereo fish measurement	SWFSC ROV GigE stereo	Cutter	yes	yes	2-4 Hz	yes		no	manual	no				
WHOI/NEFSC scallop detector	HABCAM towed rig	Dvora	yes	yes	no	yes	yes							
RPI/Kitware scallop detector	HABCAM towed rig	Hoogs	no	no	no	yes	yes							
SRI fish detection, classification, size	PI FSC MOUSS/BotCam	Ben/Mike	yes, accept cal files	yes	30 Hz	no, grayscale		yes		yes	yes			
SEFSC stereo proc	Drop cams from SEFSC	Thompson	yes	yes	yes			yes, basic background	manual	no	no			
Toyon SBIR I	Drop cams from SEFSC	Thompson	yes	yes	yes			yes, basic HOG	manual	yes	yes			
LANL segmentation and shape analysis	HABCAM towed rig	Lakshman	no	yes	no	no	yes	yes	no	no	yes	yes (image)	Yes	yes (polygonal)
Toyon SBIR II	Still Images AUV, drop, towed	Clarke	yes	yes	no	yes (Hanu)		yes	yes	no	yes			
WHOI/NEFSC habitat classifier	HABCAM towed rig	Dvora	yes	yes	no	yes							yes	
NWFSC clustering	AUV and MOUSS	Clarke	no		no	no						yes	partially	

Green	well-implemented; quantified, comparative performance assessment; ready for integration	Yellow	Existing implementation as mature research code; some performance quantification	Red	preliminary research code with ongoing work against major problems	Gray	idea or concept; no implementation

VIAME High-Level Components

Multi-Processing Pipeline Framework

Calibration and Stabilization

Detection and Tracking

Classification

Other Analytics

Data Abstraction and Conversion (not a native DB)

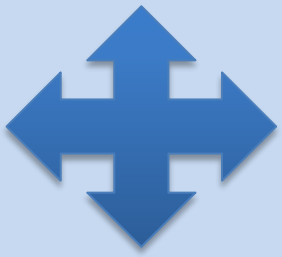

Web based services and tools

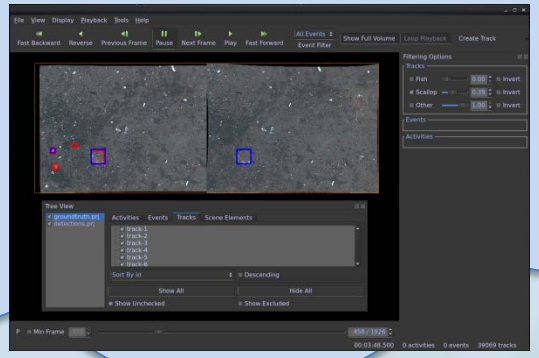

Visualization Tools

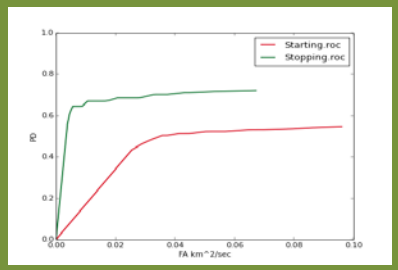

Evaluation and Quantification Tools

Kitware

Goal: Object Detection

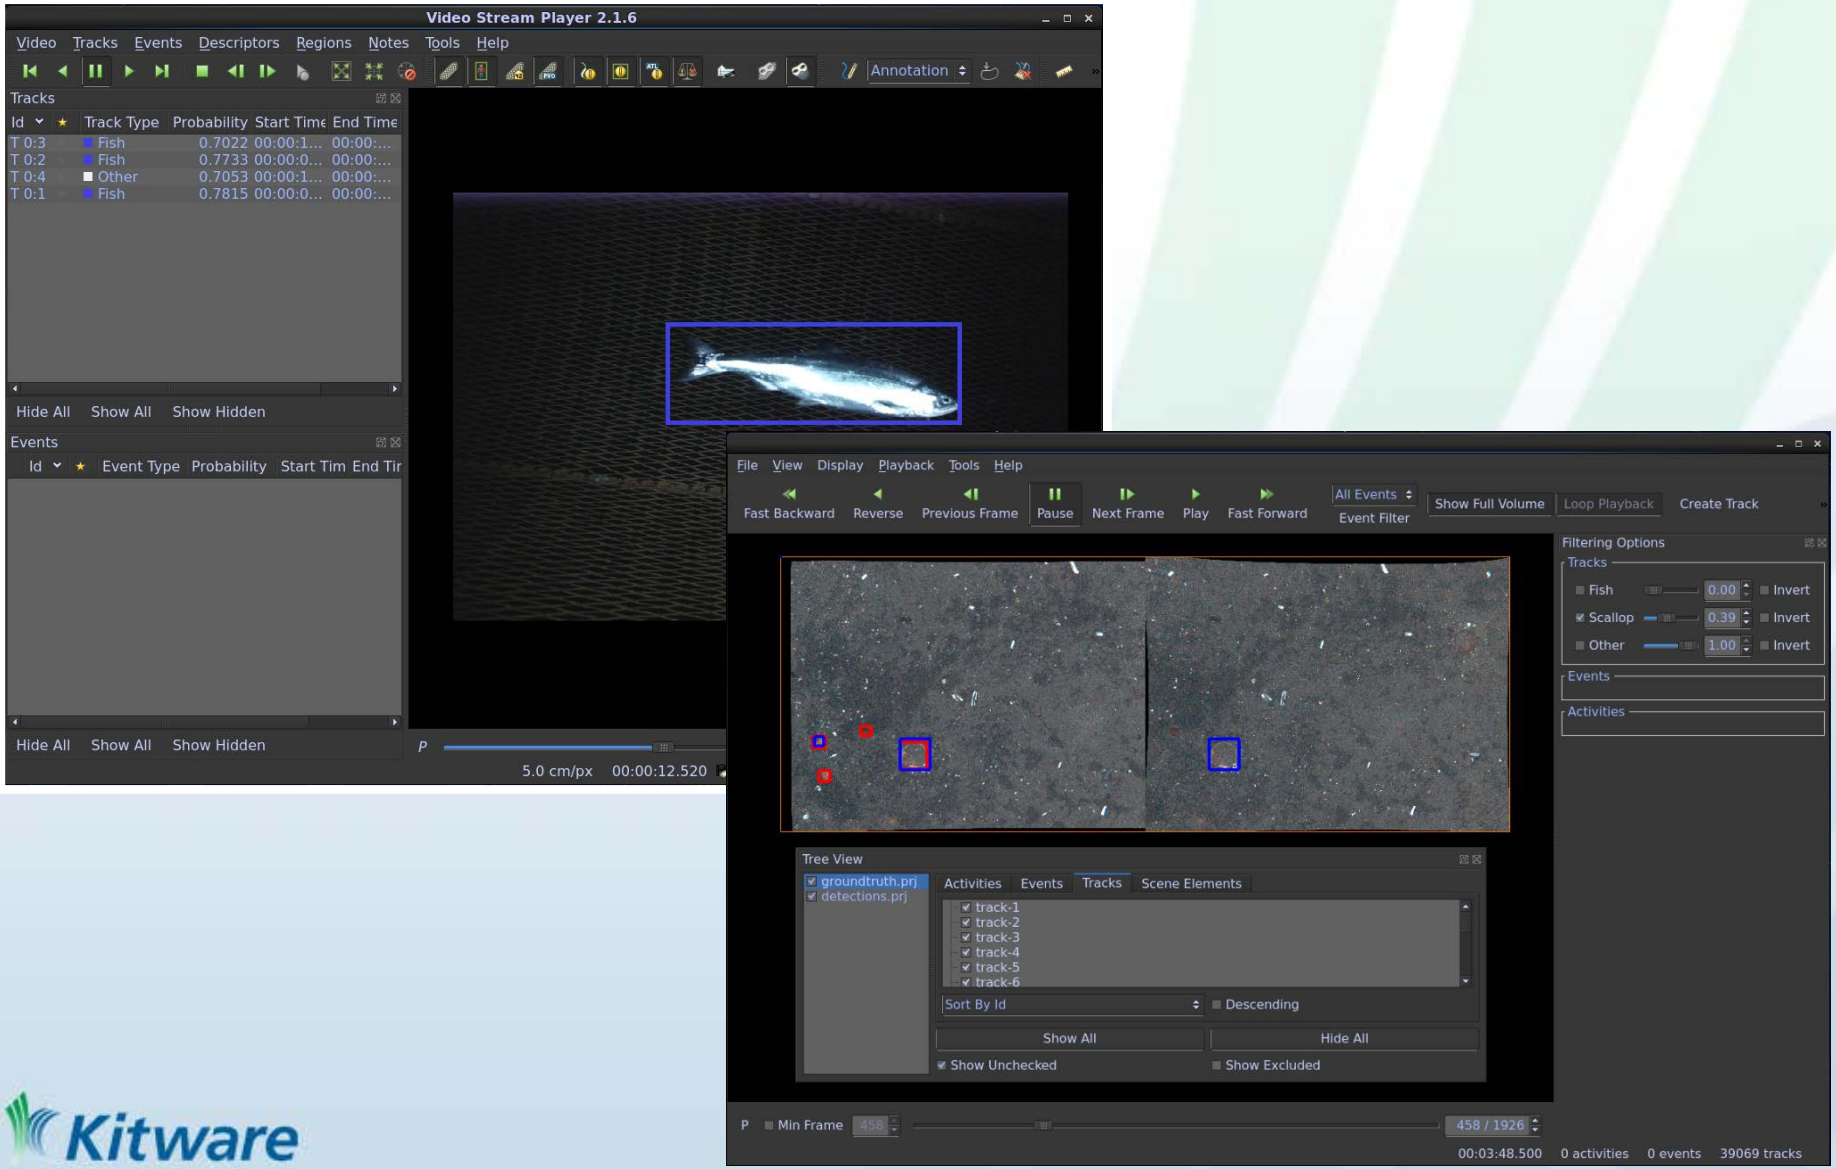

Kitware

Goal: Object Detection Viewing and Adjudication

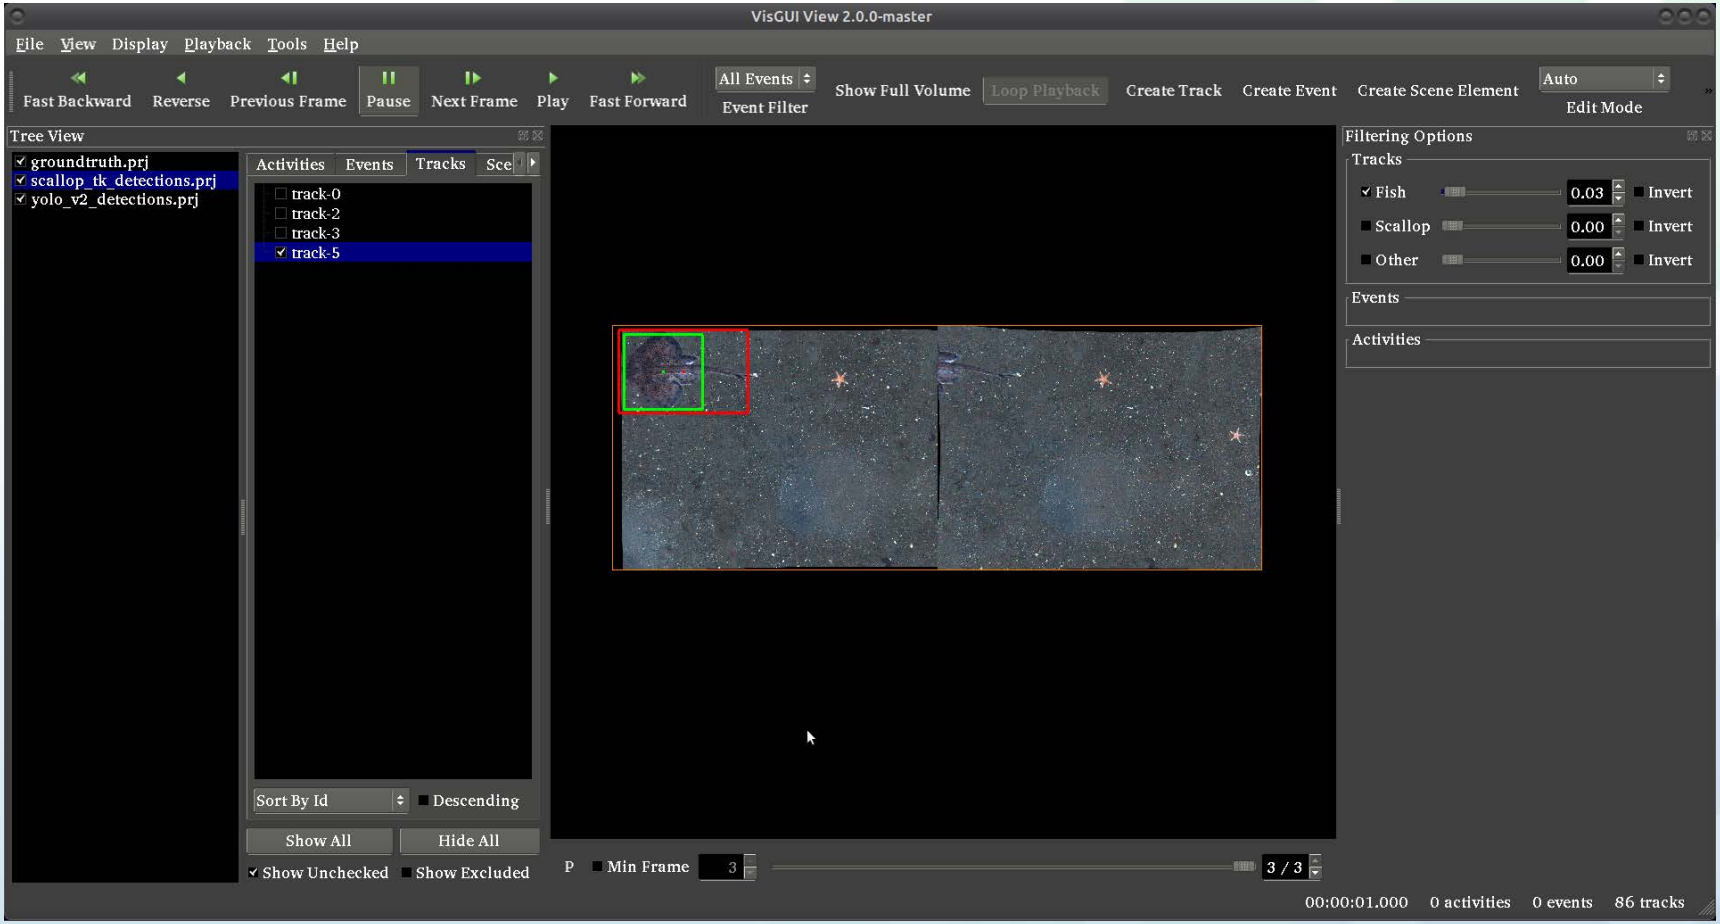

Goal: Object Tracking

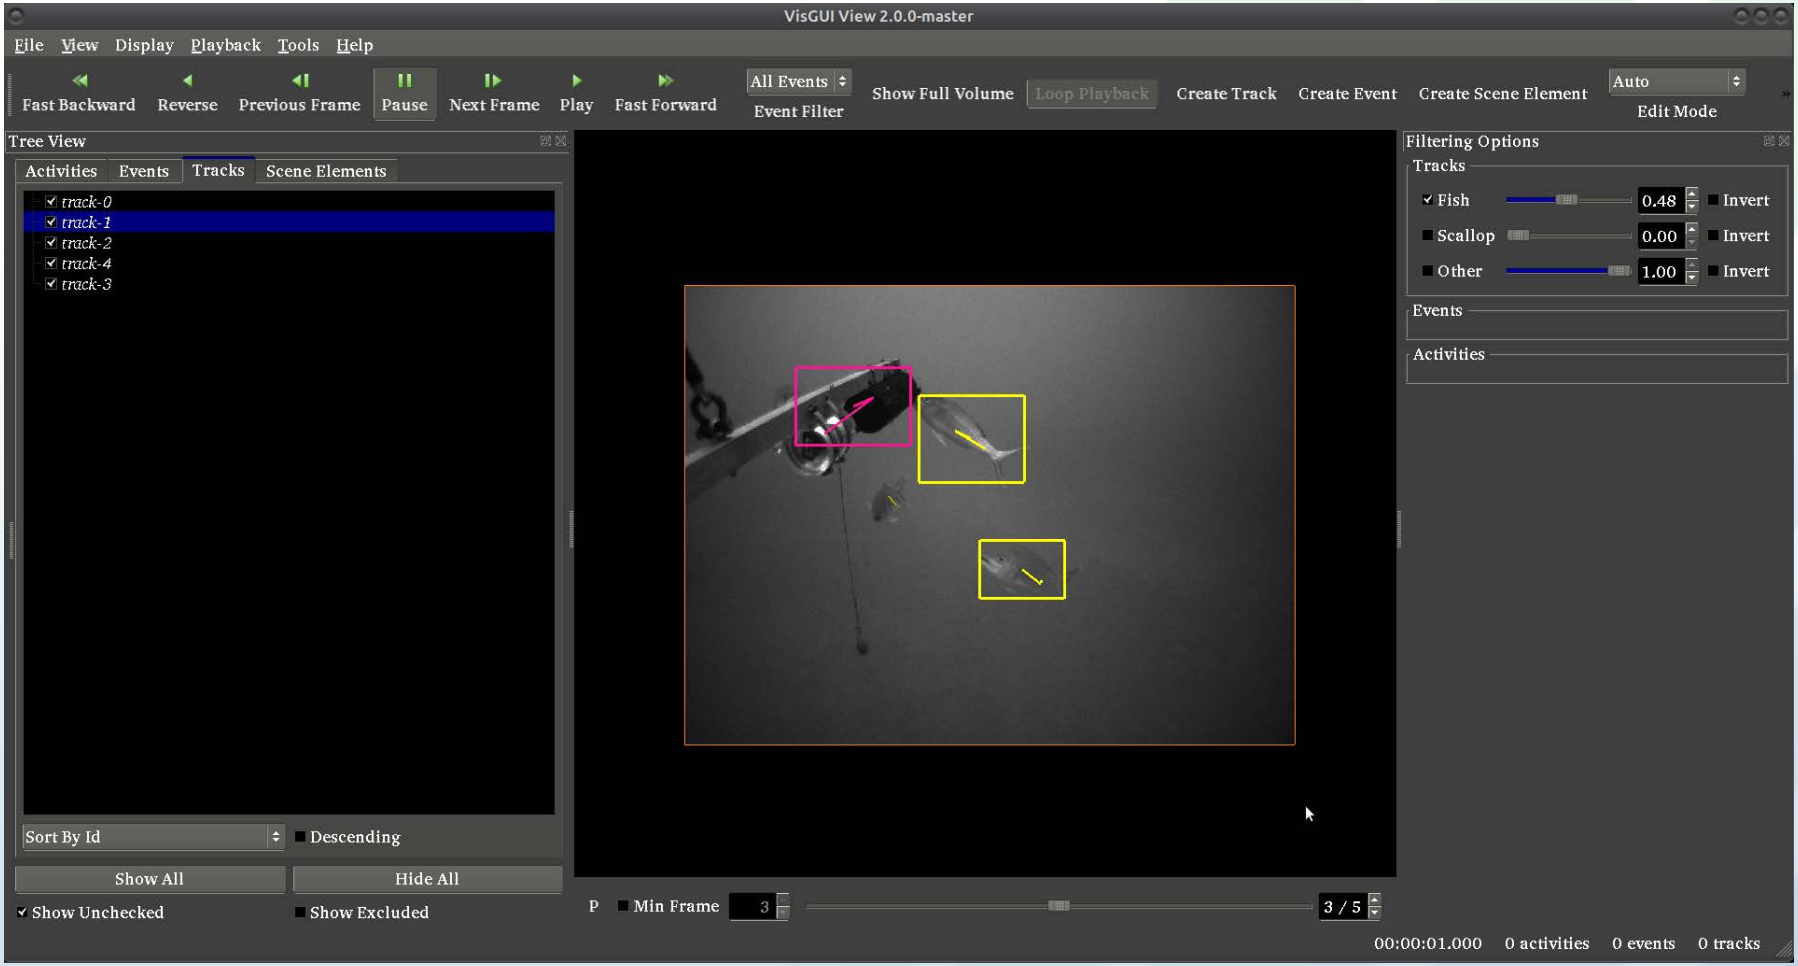

Goal: Fish Measurement

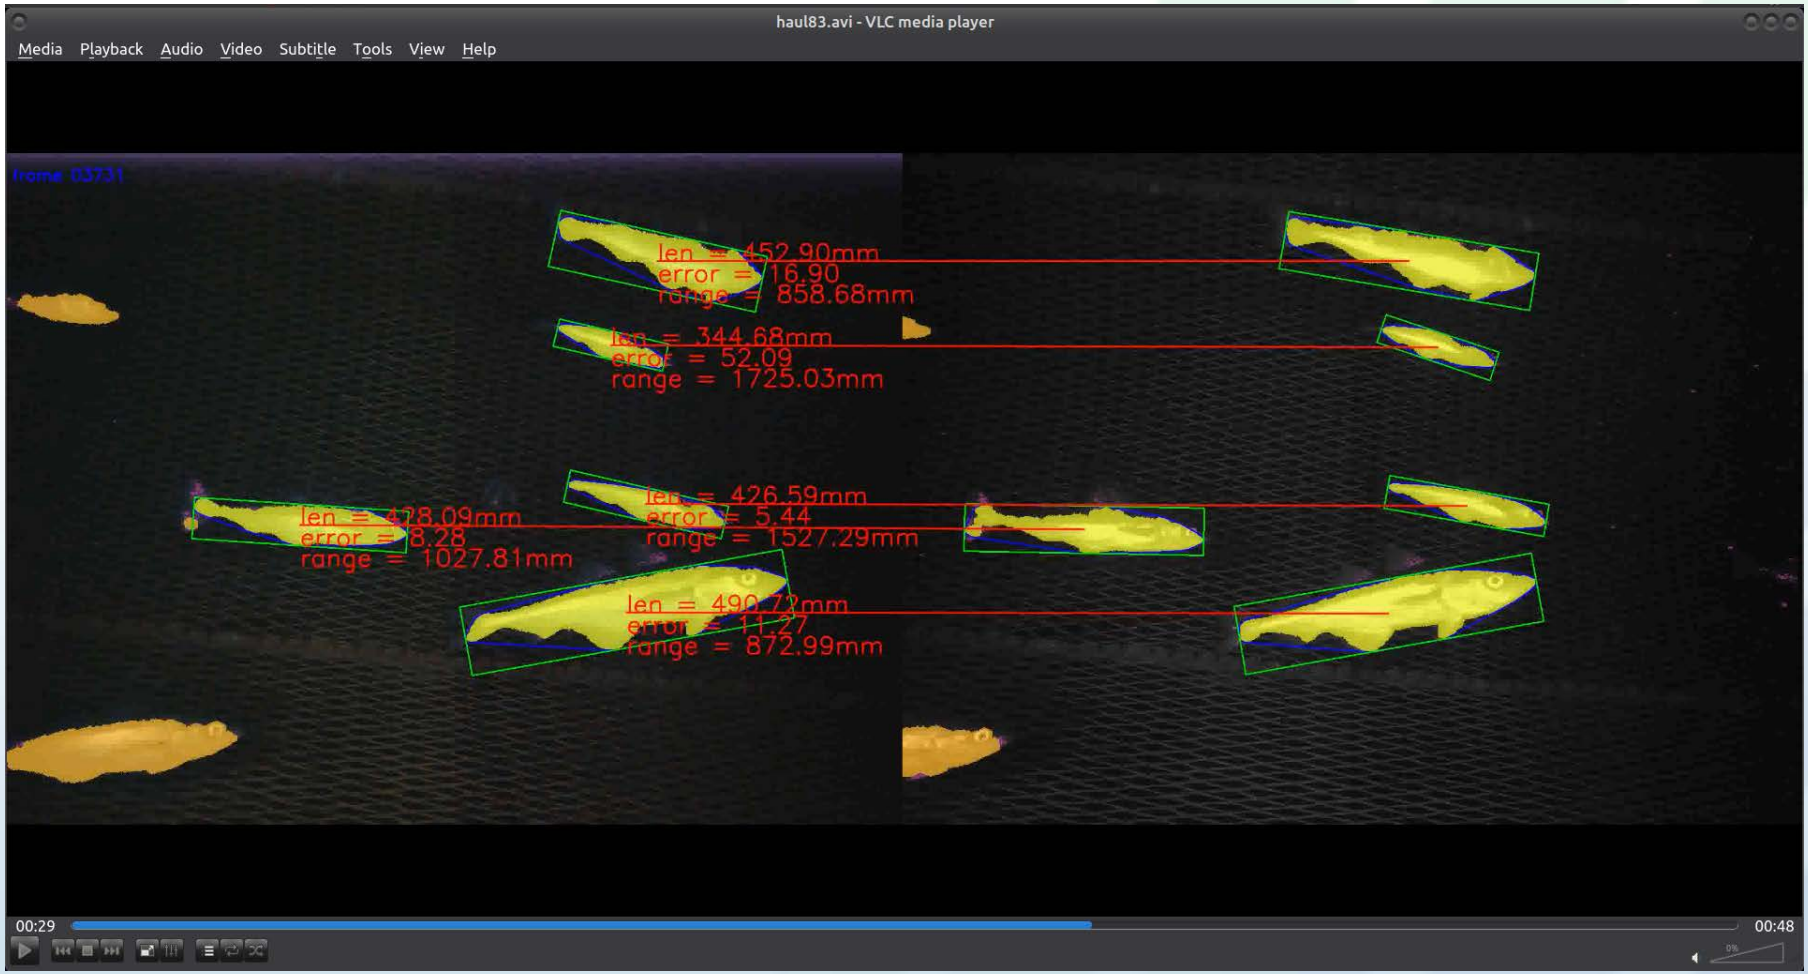

Goal: Image and Video Search

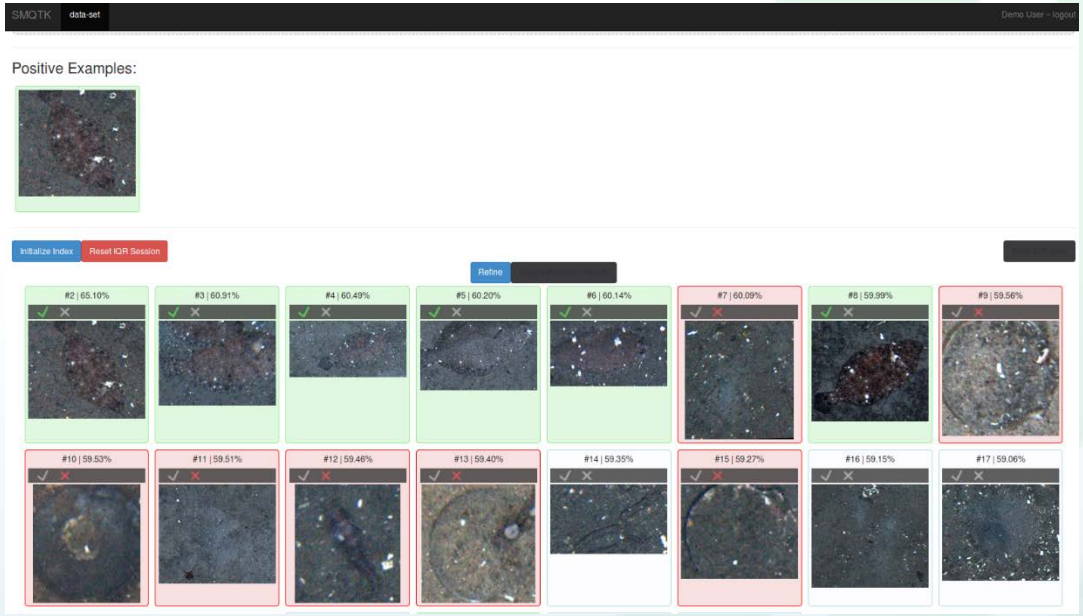

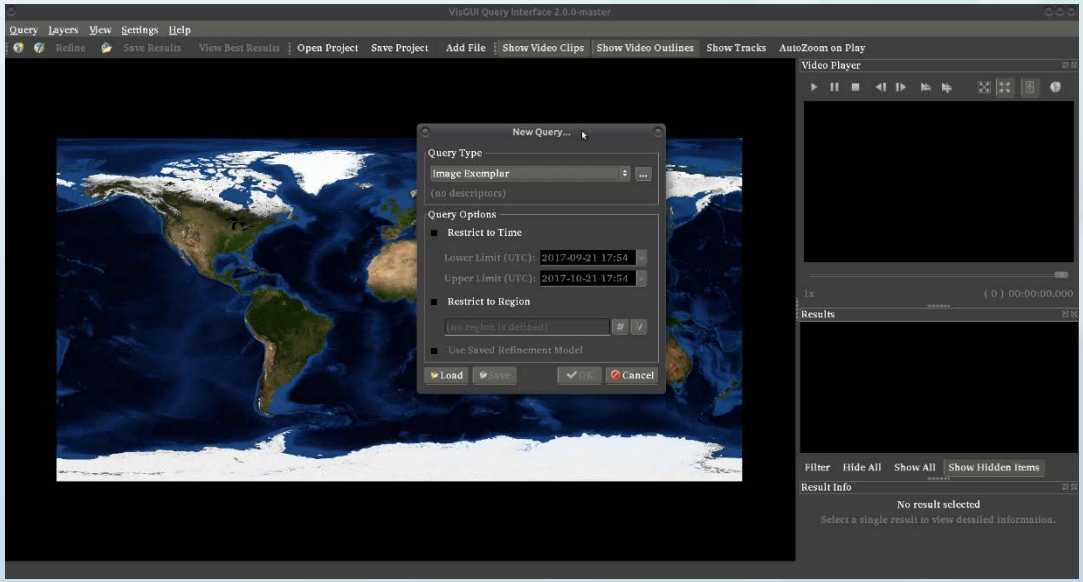

VIAME System Components

Fletch

Builds common computer vision dependencies

Sprokit/KWIVER

Connects up different algorithm nodes in runnable pipelines. Nodes can be implemented in C++, Python, or Matlab.

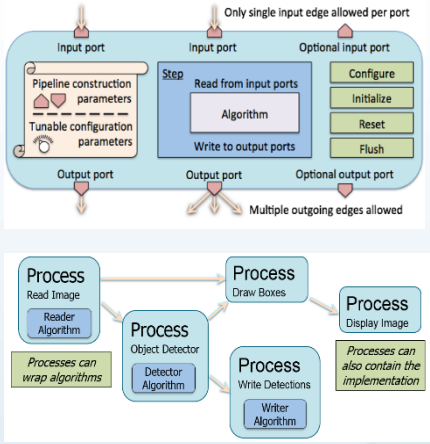

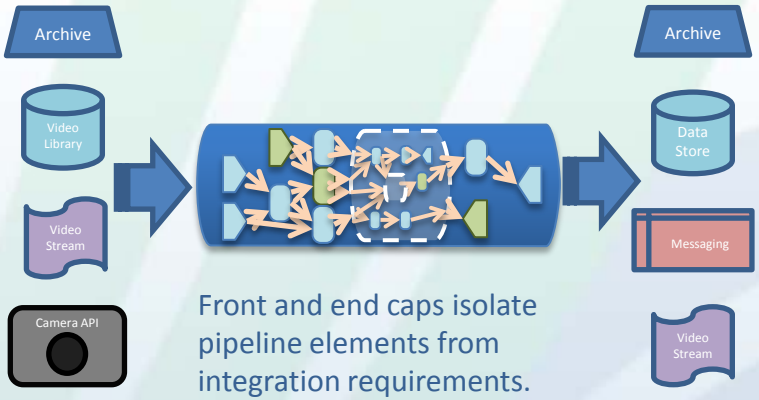

Front and end caps isolate pipeline elements from integration requirements.

```
#ifndef VIAME_SCALLOP_TK_DETECTOR_H
#define VIAME_SCALLOP_TK_DETECTOR_H

#include <plugins/scallop_tk/viame_scallop_tk_export.h>

#include <vital/algo/image_object_detector.h>

namespace viame {

class VIAME_SCALLOP_TK_EXPORT scallop_tk_detector :
  public kwiver::vital::algorithm_impl<
    scallop_tk_detector, kwiver::vital::algo::image_object_detector >
{
public:
  scallop_tk_detector();
  virtual ~scallop_tk_detector();

  // Get the current configuration (parameters) for this detector
  virtual kwiver::vital::config_block_sptr get_configuration() const;

  // Set configurations automatically parsed from input pipeline and config files
  virtual void set_configuration( kwiver::vital::config_block_sptr config );
  virtual bool check_configuration( kwiver::vital::config_block_sptr config ) const;

  // Main detection method
  virtual kwiver::vital::detected_object_set_sptr detect(
    kwiver::vital::image_container_sptr image_data ) const;

private:
  class priv;
  const std::unique_ptr< priv > d;
};
```

VIAME-Core

Contains domain-specific algorithms in “super-build” comprised of multiple projects and the above.

```
# Process definitions and configs
#
process input
  :: frame_list_input
  :image_list_file       input_files.txt
  :frame_time            0.3333
  :image_reader:type     ocv

process detector
  :: image_object_detector
  :detector:type         ex_fish_detector
  :detector:model1       model_file.xml
  :detector:threshold    0.20

process draw
  :: draw_detected_object_boxes
  :default_line_thickness 3

process disp
  :: view_image
  :annotate_image        true
  :pause_time            2.0
  :title                 VIAME images

# Global pipeline configs
#
config _pipeline:_edge
       :capacity         5

# Connections between processes
#
connect from input.image
        to   detector.image

connect from detector.detected_object_set
        to   draw.detected_object_set
connect from input.image
        to   draw.image

connect from input.timestamp
        to   disp.timestamp
connect from draw.image
        to   disp.image
```

CMake Build System

Why?

- Cross-platform as long as code supports it
- Underlying pipeline code written primarily in C/C++
- Already used by several popular vision projects (OpenCV, VXL, Caffe)
- Free to use
- Kitware

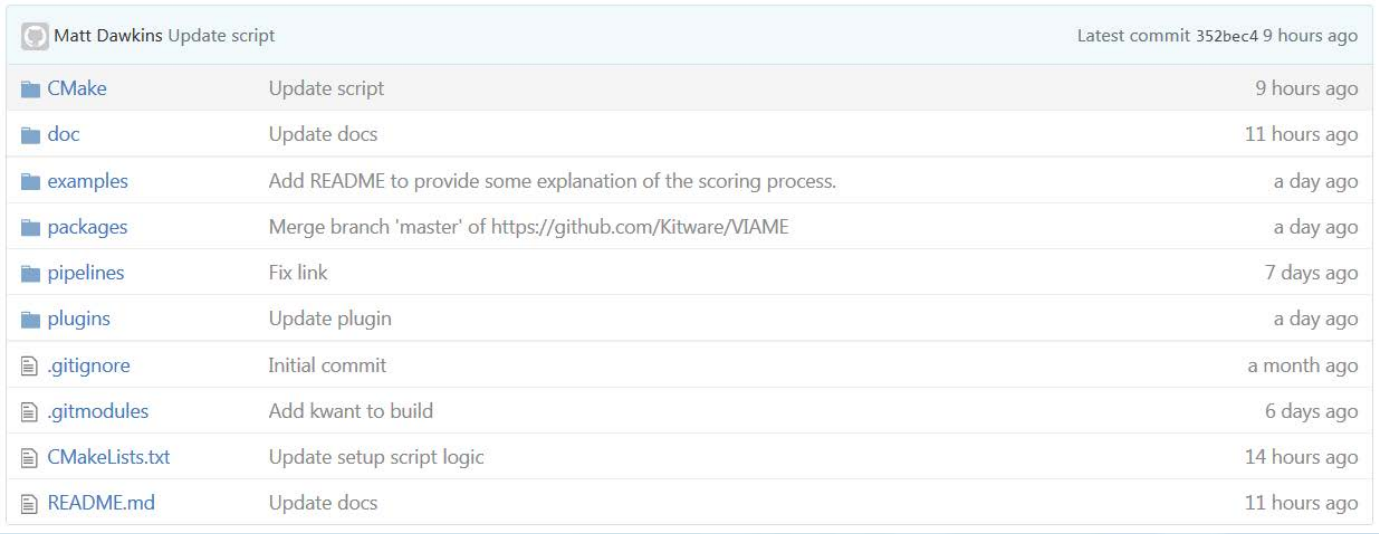

Matt Dawkins Update script		Latest commit 352bec4 9 hours ago
CMake	Update script	9 hours ago
doc	Update docs	11 hours ago
examples	Add README to provide some explanation of the scoring process.	a day ago
packages	Merge branch 'master' of https://github.com/Kitware/VIAME	a day ago
pipelines	Fix link	7 days ago
plugins	Update plugin	a day ago
.gitignore	Initial commit	a month ago
.gitmodules	Add kwant to build	6 days ago
CMakeLists.txt	Update setup script logic	14 hours ago
README.md	Update docs	11 hours ago

Fundamentals: https://cmake.org/runningcmake/

Git Version Control

Alternatives: Use Nothing, Google Drive, CVS, SVN, Mercurial

Why Git? Like CMake, also used by many other vision toolkits. Slightly harder to learn than SVN, but more extensible, distributed.

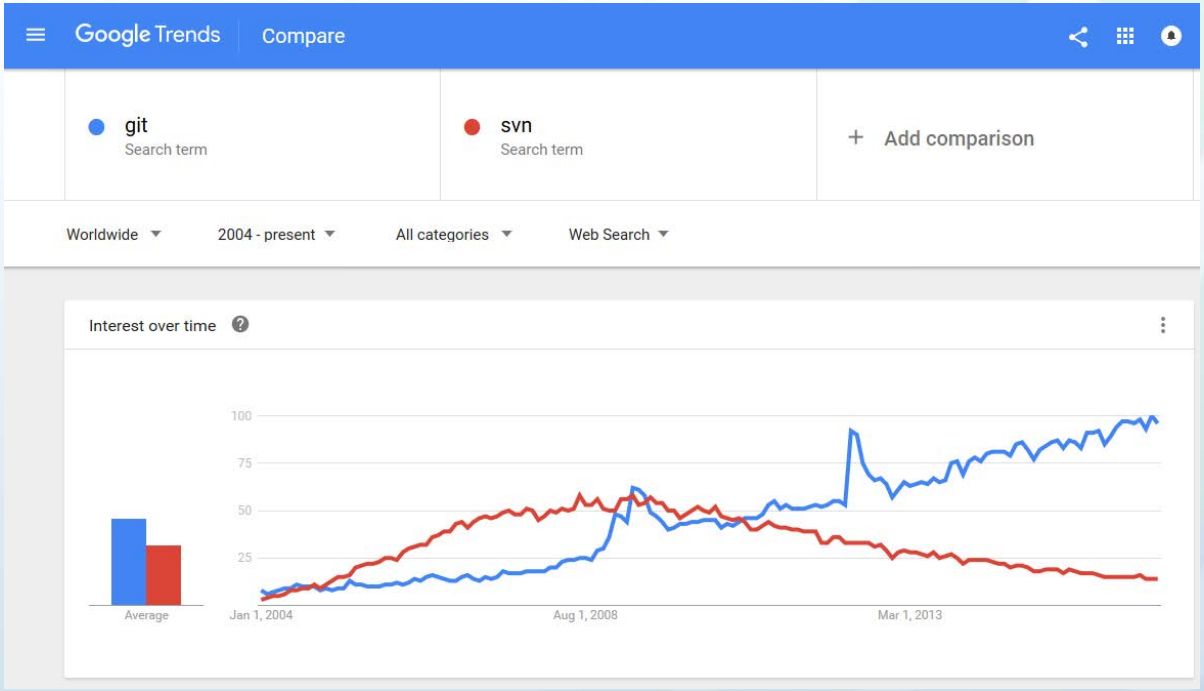

Fundamentals: https://git-scm.com/book/en/v2/Getting-Started-Git-Basics

Super-Build Comprised of Multiple Projects

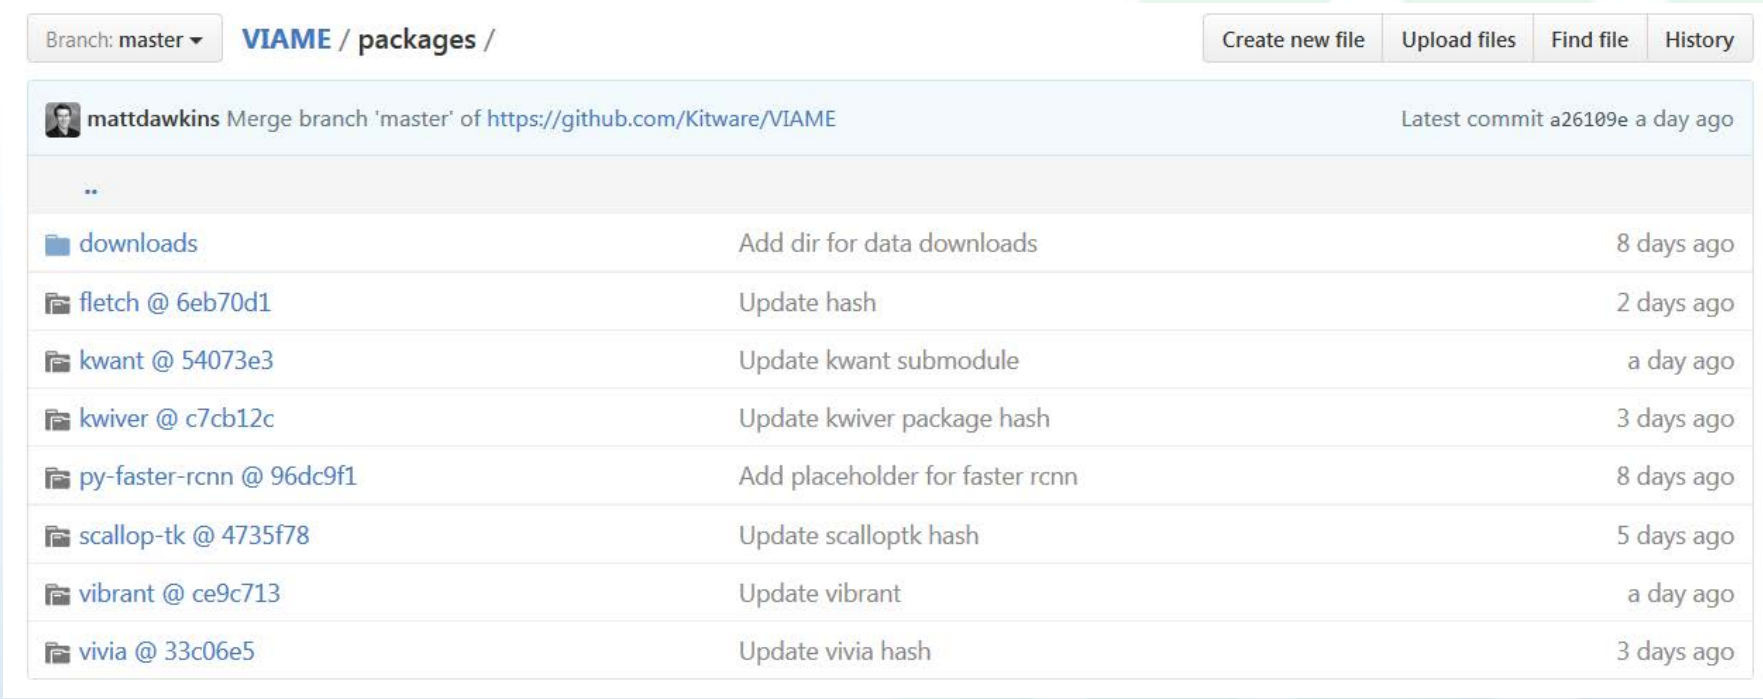

KWIVER.org

Kitware Image and Video Exploitation and Retrieval Toolkit

An Open Source, production-quality video analytics toolkit

Social Multimedia Query ToolKit

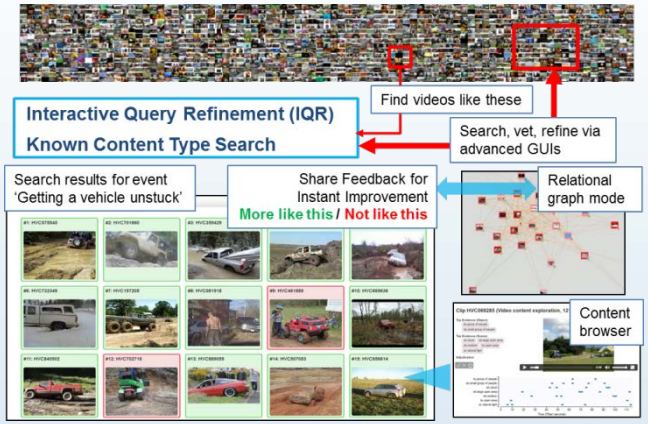

VIBRANT:
Video and Image-Based Retrieval and Analysis Toolkit

Archive Query

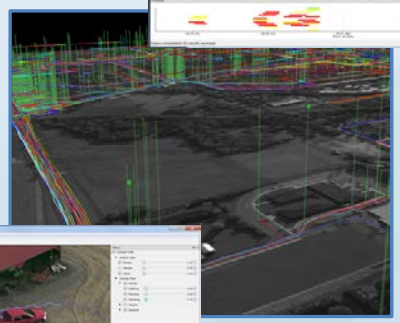

Streaming FMV

Motion-imagery Aerial Photogrammetry Toolkit

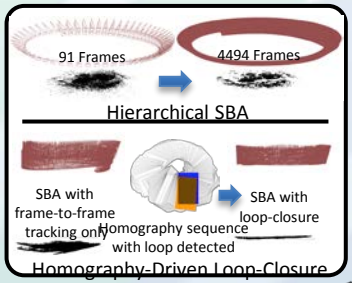

We hope to establish an open-source community for video analytics research and development

Kitware

KWIVER.org

- Source code repositories maintained at GitHub
- Current toolkits available:
  - Motion-imagery Aerial Photogrammetry Toolkit for video stabilization and online bundle adjustment **MAP-TK**
  - Social Multimedia Query Toolkit for visual context extraction and querying for social multimedia **SMQTK**
  - Stream Processing Toolkit to facilitate multi-state, pipelined processing of data streams **Sprokit**
  - Common data structures and abstractions for computer vision and machine learning systems **VITAL**
  - GPU acceleration of core vision algorithms using OpenCL **VisCL**
  - GUIs and sophisticated visualization tools for content automatically extracted from video, based on VTK **ViVIA**
  - Detection and tracking of movers in video **VIBRANT**
  - CMake tools to set up complex build environments and dependencies **Fletch**

Sprokit – A Framework for Streaming Data Processing

- What is Sprokit?
  - Plugin-based **streaming data processing engine** on which to build modular streaming data processing applications (especially video processing applications).
- What does Sprokit do?
  - Chains processing elements into a directed acyclic graph (DAG)
  - Executes a constructed pipeline on streaming data (e.g. video)
  - Provides dynamic construction/configuration via configuration files
  - Supports algorithms written in C++, Matlab, and Python
- Why was Sprokit developed?
  - To build complex streaming algorithms from simple components
  - To replace an older, much more restrictive, pipeline framework
  - Because existing open source frameworks (e.g. Gstreamer, Ecto, etc.) did not meet all requirements in the list above.

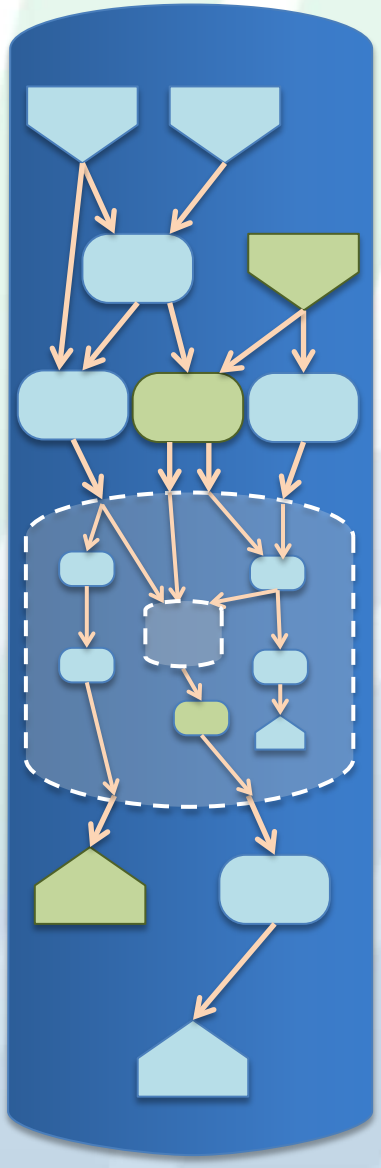

Kitware

Swappable Front and End Caps

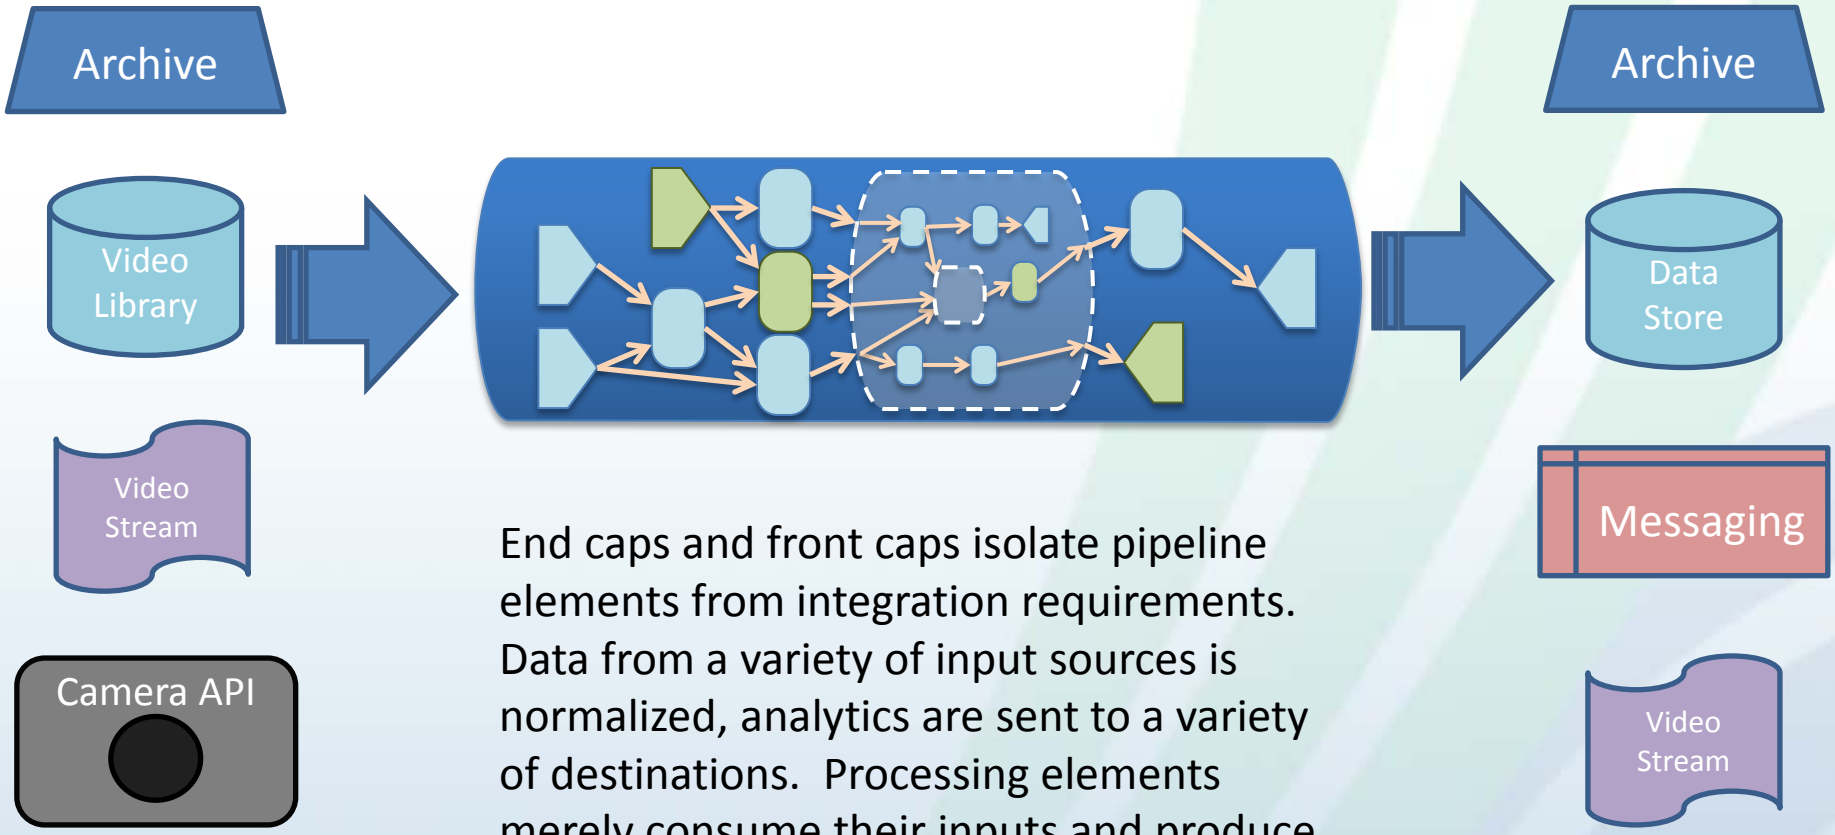

End caps and front caps isolate pipeline elements from integration requirements. Data from a variety of input sources is normalized, analytics are sent to a variety of destinations. Processing elements merely consume their inputs and produce their outputs.

Individual Processing Nodes

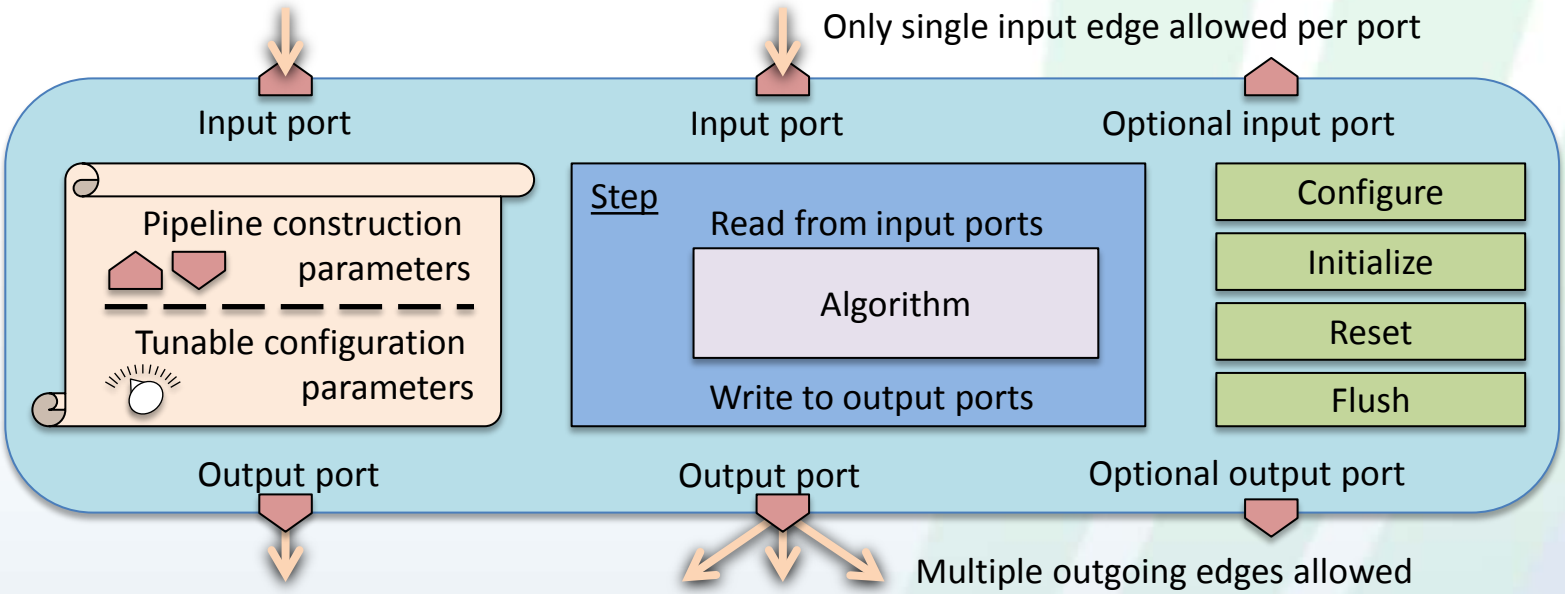

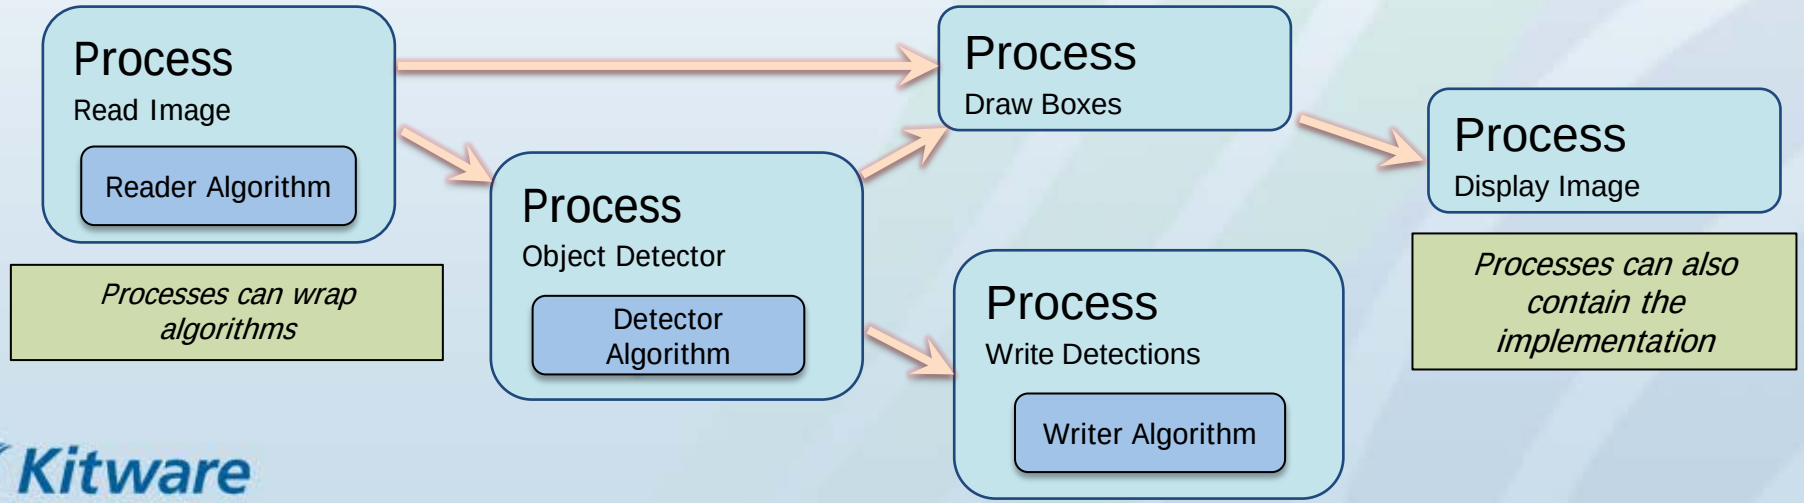

Kitware

VITAL: Common Data Structures for Edges

- Normalize and standardize data structures to facilitate integration across different algorithms
- Isolate system integration issues by providing test fixtures for remote development
- Encourage collaboration by providing a framework for data sharing and replay
- Encourage modular development with pipeline based architecture with common data structured passed on edges

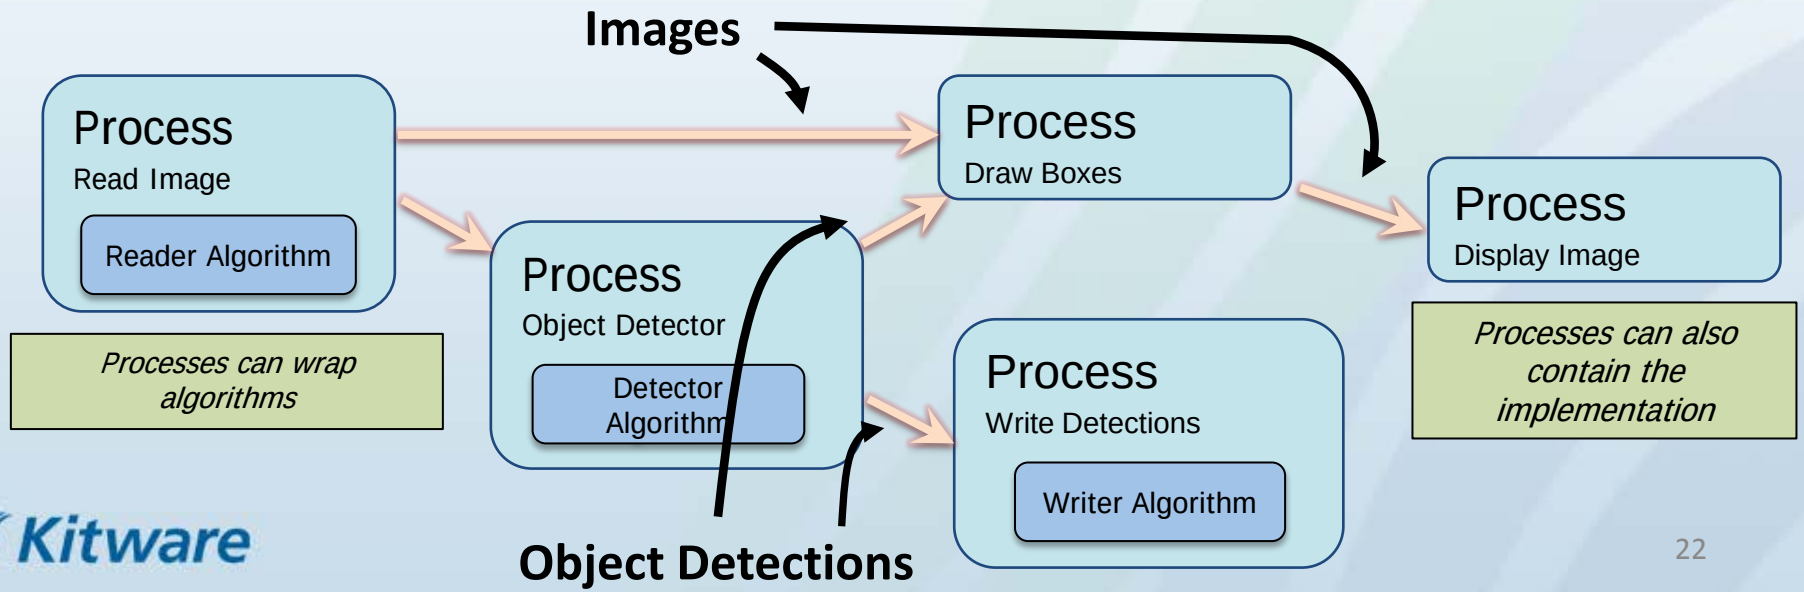

VITAL: Common Example Types

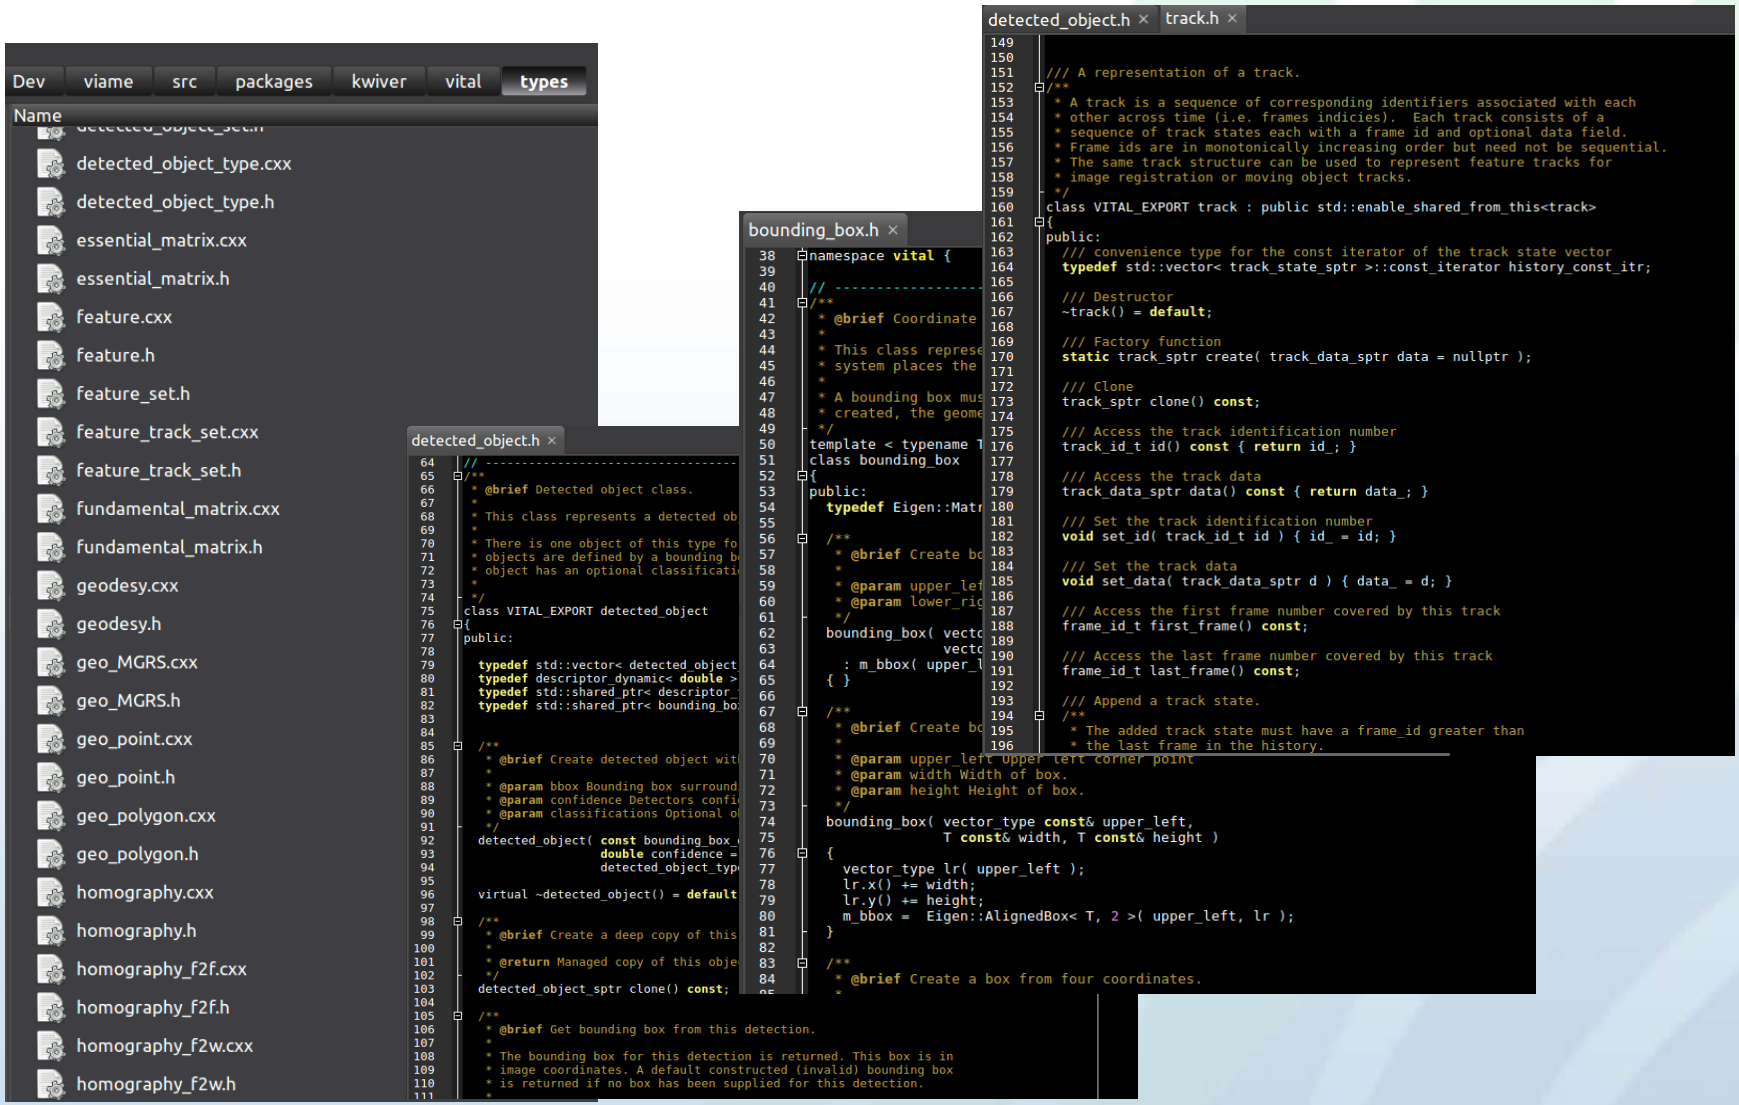

Simple Pipeline Example

```
# ======================================
process input
  :: frame_list_input
  :image_list_file     input_files.txt
  :frame_time          .3333
  :image_reader:type   ocv

# ======================================
process detector
  :: image_object_detector
  :type                                scallop_tk_detector
  :scallop_tk_detector:config_file   config_location

# ======================================
process draw
  :: draw_detected_object_boxes
  :default_line_thickness 3

# ======================================
process disp
  :: view_image
  :annotate_image          true
  :pause_time              0  # 1.0
  :title                   NOAA images
```

```
# ======================================
# global pipeline config
#
config _pipeline:_edge
      :capacity 10

# ======================================
# connections
connect from input.image
      to   detector.image

connect from detector.detected_object_set
      to   draw.detected_object_set
connect from input.image
      to draw.image

connect from detector.detected_object_set
      to archive.detected_object_set

connect from input.timestamp
      to   disp.timestamp
connect from draw.image
      to   disp.image
```

Kitware

Process Definition

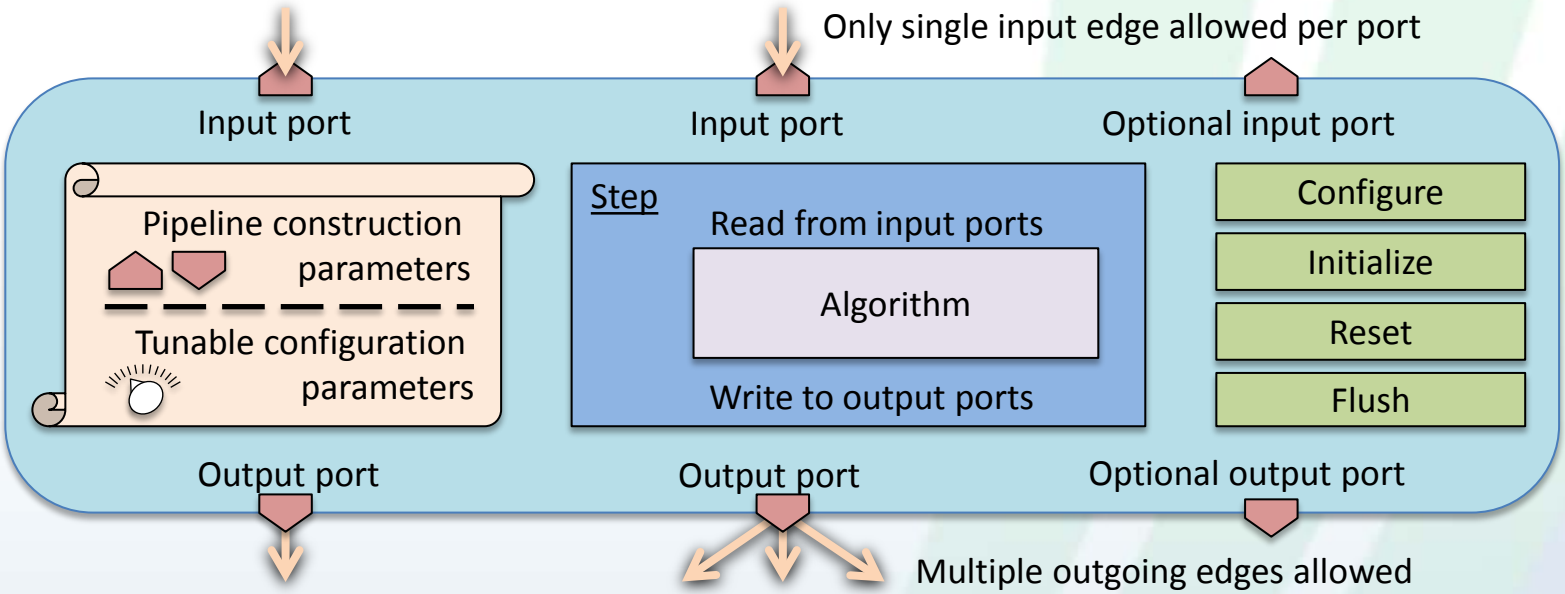

Two separate ways to define:

- Manual: User specifies all input/output ports and each rectangular block in the above (step, configure, initialize, etc…)
- Automatic: Only need to define single function for existing base class APIs with process wrappings
  - For example, object_detector: image in, detections out

Algorithm Concepts

- Application uses abstract algorithm type using a polymorphic model
- Instantiates an implementation based on config
- Implementations are dynamically loadable
- New implementations can be easily added

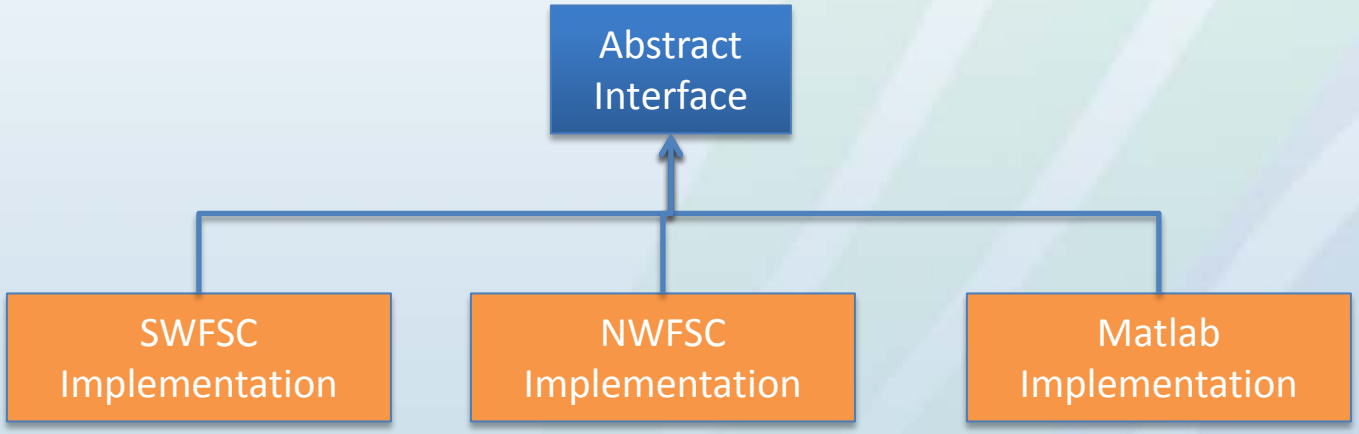

Kitware

Automatic Implementation Method:
Example Base Classes

kwiver/vital/algos/image_object_detection.h

image_object_detector.h

```
// ------------------------------------------------------------------
/**
 * @brief Image object detector base class/
 *
 */
class VITAL_ALGO_EXPORT image_object_detector
: public algorithm_def<image_object_detector>
{
public:
  /// Return the name of this algorithm
  static std::string static_type_name() { return "image_object_detector"; }

  /// Find all objects on the provided image
  /**
   * This method analyzes the supplied image and along with any saved
   * context, returns a vector of detected image objects.
   *
   * \param image_data the image pixels
   * \returns vector of image objects found
   */
  virtual detected_object_set_sptr
      detect( image_container_sptr image_data) const = 0;

protected:
  image_object_detector();
};

/// Shared pointer for generic image_object_detector definition type.
typedef std::shared_ptr<image_object_detector> image_object_detector_sptr;

} } } // end namespace

#endif //VITAL_ALGO_IMAGE_OBJECT_DETECTOR_H_
```

kwiver/sprokit/core/image_object_detector_process.h

image_object_detector_process.h

```
// ------------------------------------------------------------------
/**
 * @brief Image object detector process.
 *
 * \iports
 * \iport{image}
 *
 * \oports
 *
 * \oport{detected_object_set}
 */
class KWIVER_PROCESSES_NO_EXPORT image_object_detector_process
  : public sprokit::process
{
public:
  image_object_detector_process( kwiver::vital::config_block_sptr const& config );
  virtual ~image_object_detector_process();


protected:
  virtual void _configure();
  virtual void _step();

private:
  void make_ports();
  void make_config();

  class priv;
  const std::unique_ptr<priv> d;
}; // end class object_detector_process



} // end namespace

#endif /* ARROWS_PROCESSES_IMAGE_OBJECT_DETECTOR_PROCESS_H */
```

Kitware

Example Algorithms

Kitware: ScallopTK

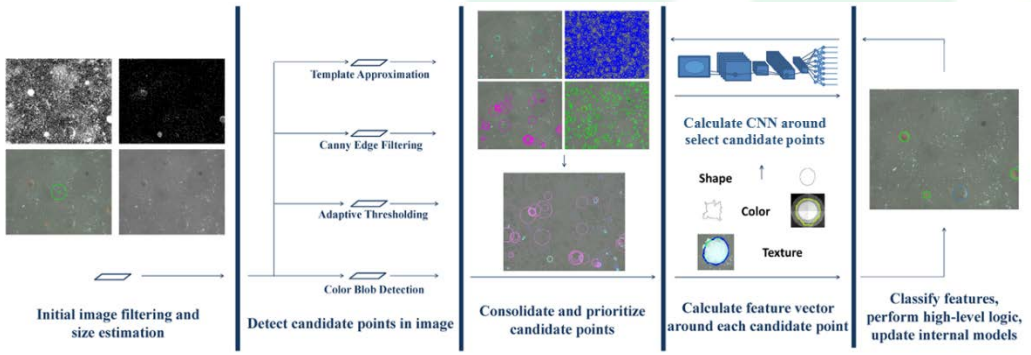

SRI: BenthosDetect

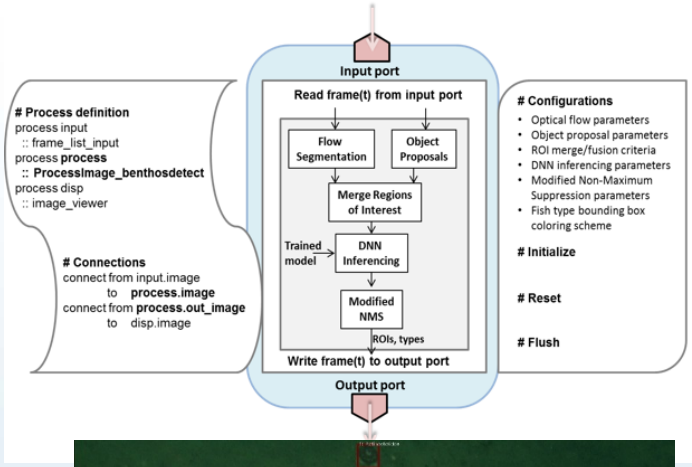

NOAA/UW: FishRuler

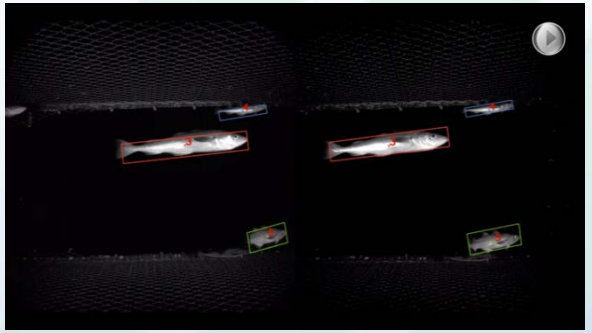

LANL: ScallopFinder

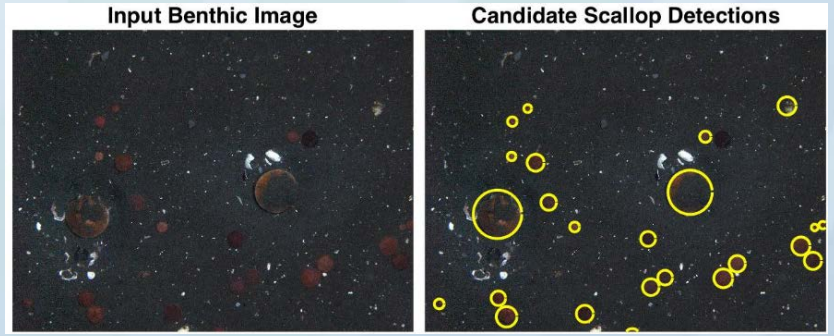

Kitware

Example Algorithms (cont.)

Burn-Out

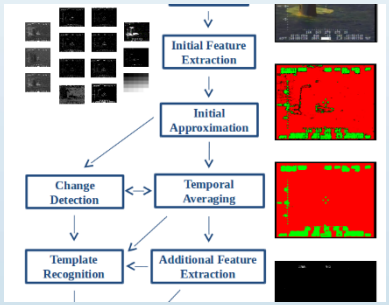

Faster R-CNN

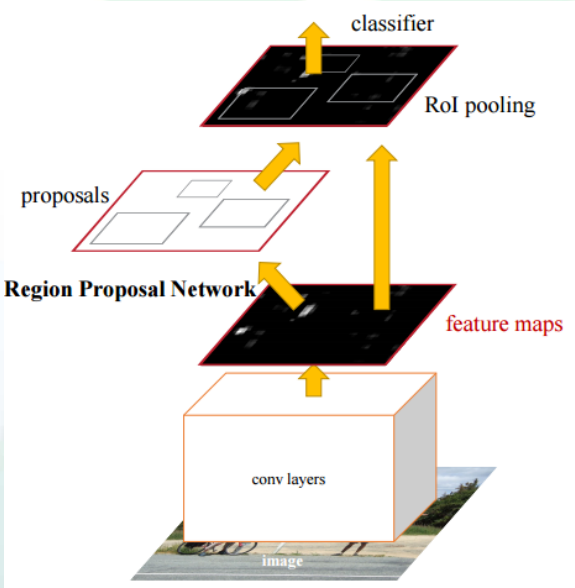

YOLOv2

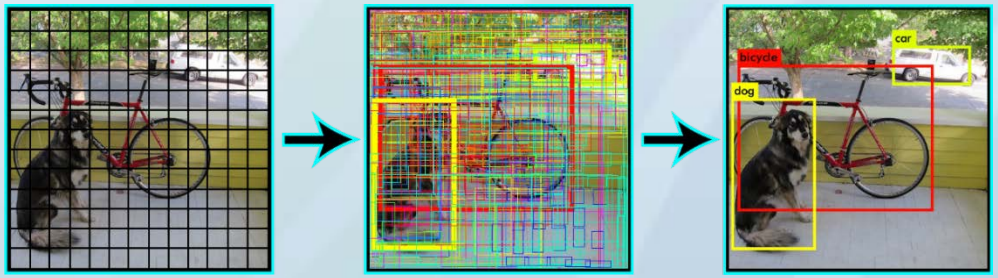

Kitware

Baseline Tracker – Simple Tracker

ingest_video.pipe

```
66          to   detector_writer.image_file_name
67
68  # ================================ CORE TRACKER ====================================
69
70  process detection_descriptor
71   :: compute_track_descriptors
72    :inject_to_detections                        true
73    :computer:type                               burnout
74    relativepath computer:burnout:config_file =  detection_descriptors.conf
75
76  process tracker
77   :: compute_association_matrix
78    :matrix_generator:type                       from_features
79
80    block matrix_generator:from_features:filter
81      :type                                      class_probablity_filter
82      :class_probablity_filter:threshold         0.001
83      :class_probablity_filter:keep_all_classes  false
84      :class_probablity_filter:keep_classes      car;person
85    endblock
86
87  process track_associator
88    :: associate_detections_to_tracks
89    :track_associator:type                       threshold
90    :track_associator:threshold:threshold        100.0
91    :track_associator:threshold:higher_is_better false
92
93  process track_initializer
94    :: initialize_object_tracks
95    :track_initializer:type                      threshold
96
97    block track_initializer:threshold:filter
98      :type                                      class_probablity_filter
99      :class_probablity_filter:threshold         0.001
100     :class_probablity_filter:keep_all_classes  false
101     :class_probablity_filter:keep_classes      car;person
102   endblock
```

Pipeline performs differencing on CNN intermediate features derived from bounding boxes around detections

Baseline Tracker – Tracking the Untrackable

Implemented in PyTorch, Integrated into VIAME

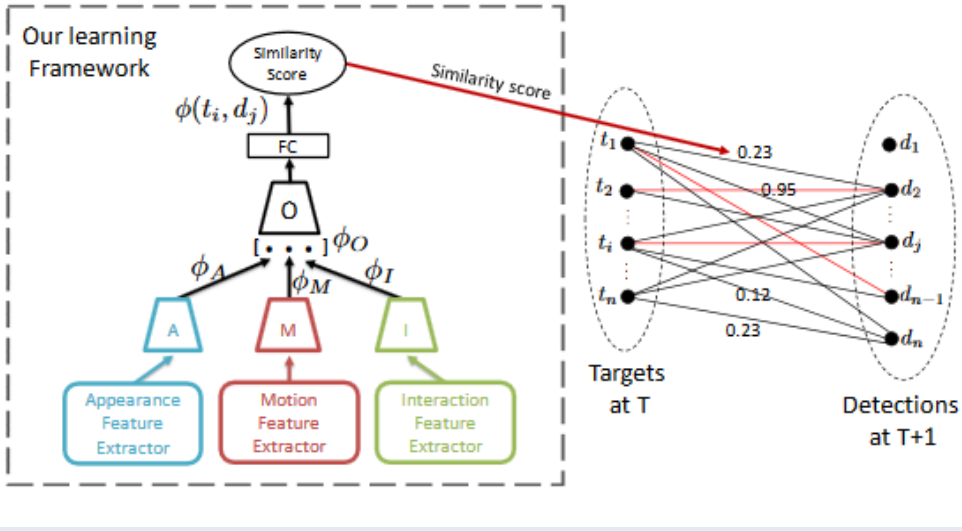

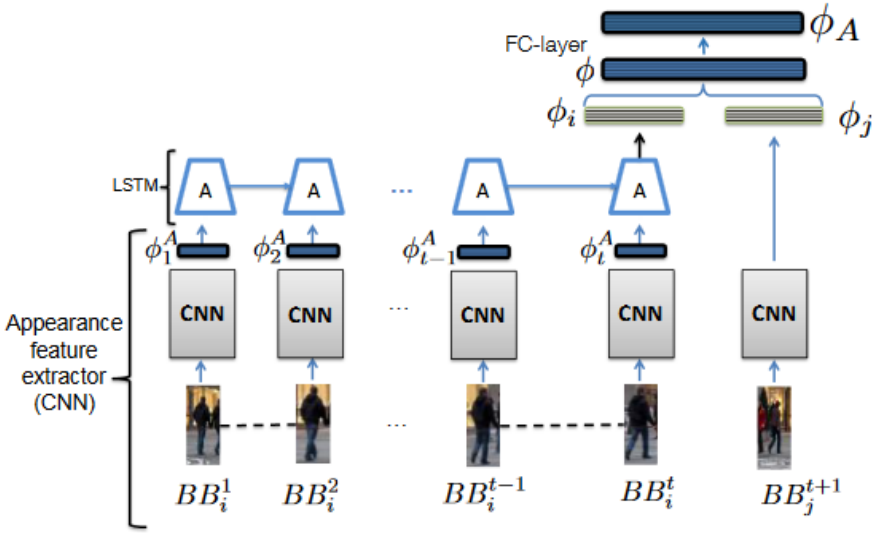

"Tracking the Untrackable: Learning to track multiple cues with long-term dependencies" Sadeghian et al. ICCV 2017

Algorithm Usage

- Application uses abstract algorithm(s) via pointer to base class, can either be used in pipeline files or C++ code. Pipelines can also be embedded in C++ code.
- Configuration info specifies algorithm to use
- High level code is unchanged as different algorithm implementations are used

```
namespace algo = kwiver::vital::algo;
algo::image_io_sptr    m_image_reader;
kwiver::vital::config_block_sptr algo_config = get_config(); // config

// validate config parameters
if ( ! algo::image_io::check_nested_algo_configuration( "image_reader", algo_config ) )
{  // Handle error
}

// instantiate image reader and converter based on config type
algo::image_io::set_nested_algo_configuration( "image_reader", algo_config, m_image_reader);
if ( ! m_image_reader )
{  // Handle error
}

// Read an image
kwiver::vital::image_container_sptr img_c;
img_c = m_image_reader->load( resolved_file_name );
```

Standalone Utilities: GUIs for Imagery

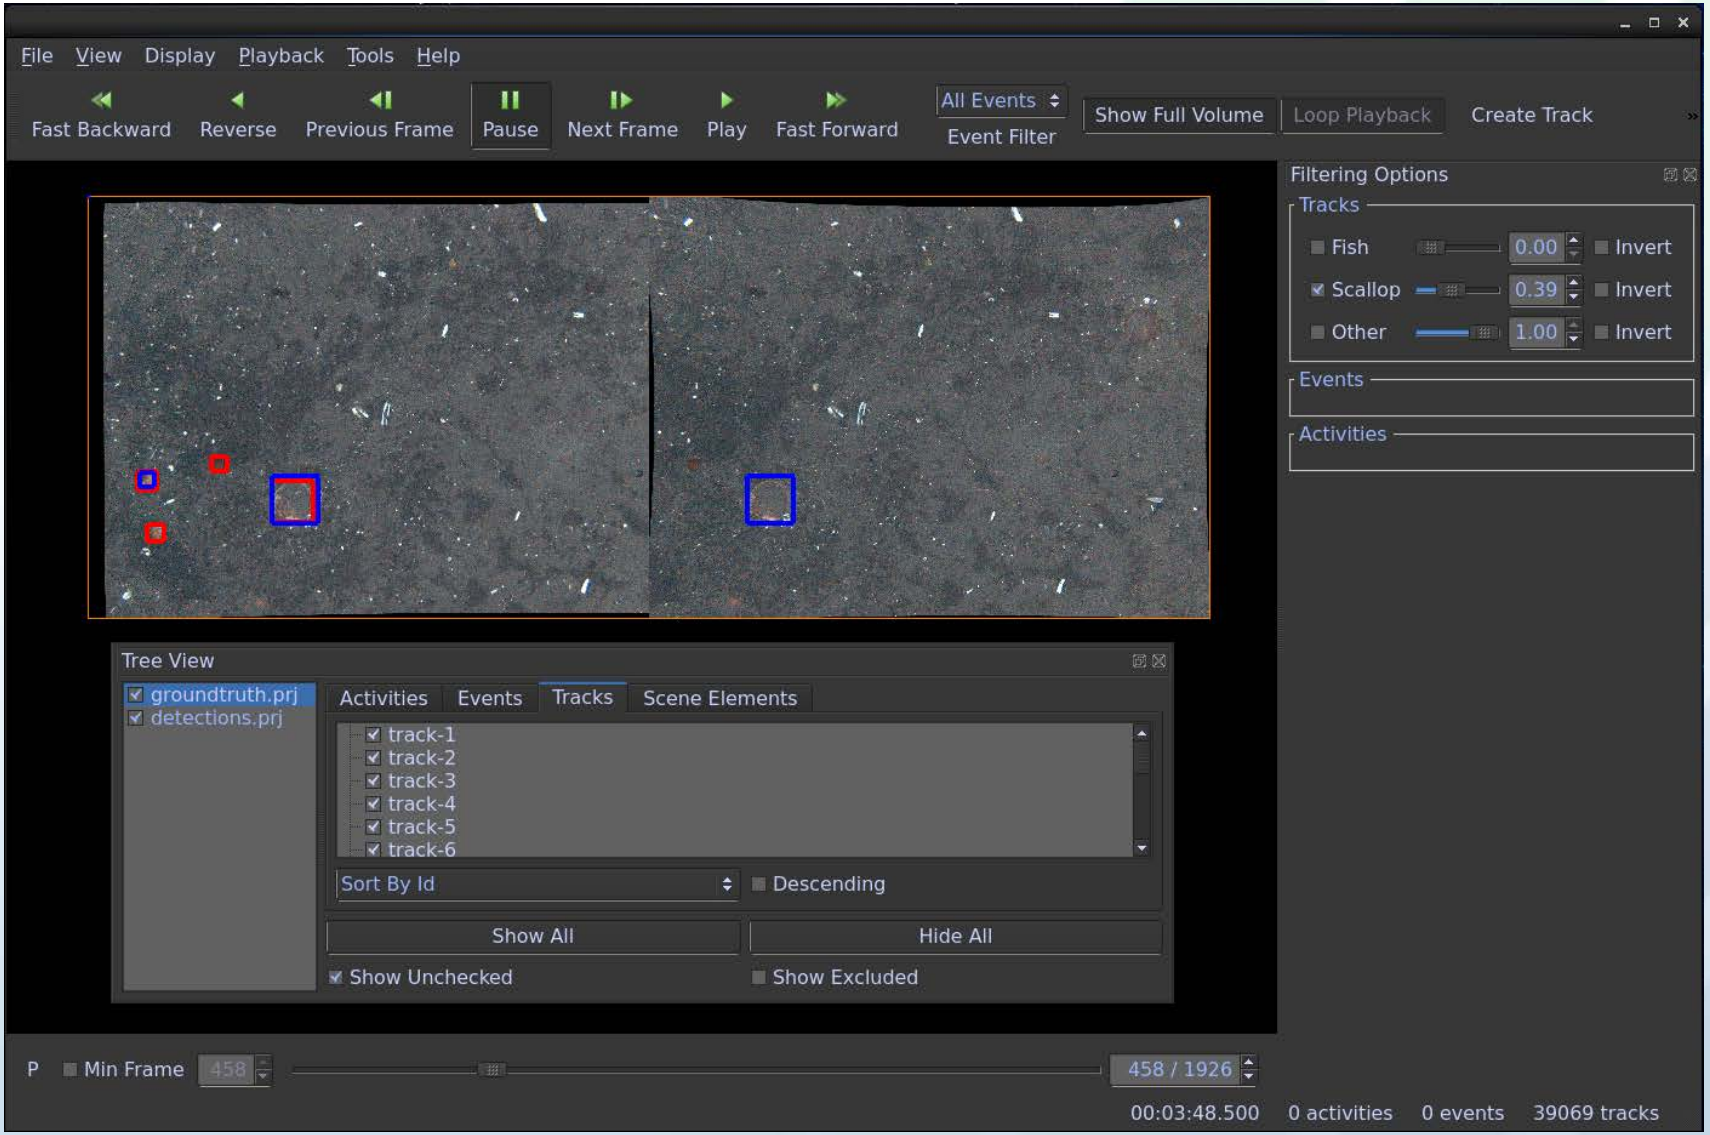

GUIs for Displaying Videos and Annotations

Evaluation and Scoring

- VIAME includes an extensive scoring capability for measuring detection, tracking, and classification on images or video
- Existing annotations can be translated to VIAME-compatible formats

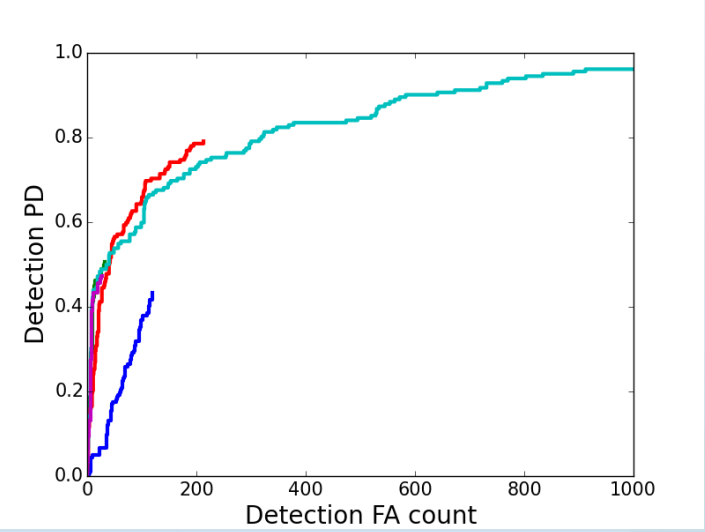


```
Detection-Pd: 0.791209
Detection-FA: 213
Detection-PFA: 0.515738
Frame-NFAR: not computed
Track-Pd: 0.791209
Track-FA: 213
Computed-track-PFA: 0.515738
Track-NFAR: not computed
Avg track (cont., purity ): 1.34, 1
Avg target (cont., purity ): 1.47, 0.79
Track-frame-precision: 0.5

DEPTH-Hash : "a2123cde"
```

Model Training

Old:

```
python format_data_for_training.py \
  -i ${input_folder}/FishCLEF15/video2/cadb2ec9\#201105051700_0.xml \
  -f clef -o ${output_folder}/formatted_samples \
  -v ${output_folder}/validation  \
  -s clef2_ -e ${input_folder}/clef_exclude.txt \
  ${common_args}

python format_data_for_training.py \
  -i ${input_folder}/FalseEx/filelist.txt  \
  -f habcam -o ${output_folder}/formatted_samples \
  -v ${output_folder}/validation \
  ${common_args}

python format_data_for_training.py \
  -i ${input_folder}/HabCamEx/Groundtruth.txt  \
  -f habcam -o ${output_folder}/formatted_samples \
  -v ${output_folder}/validation \
  --clip-right \
  ${common_args}

# Generate input training list and run training
python generate_headers.py -t YOLOv2 \
  -i ${input_folder} \
  -o ${output_folder} \
  -e ${data_type}

darknet -i ${gpu_id} detector train \
  ${output_folder}/YOLOv2.data \
  config_files/YOLOv2.cfg \
  ../detector_pipelines/models/model2.weights
```

New:

training_pipeline.pipe

```
# Pipeline for training object detector

# ==================== GLOBAL PROPERTIES ========================
# global pipeline config
#
config _pipeline:_edge
       :capacity 5

# ======================== INPUT NODE ==========================
process tdata
  :: training_data_source
  :image_list_file          image_list.txt
  :image_reader:type        ocv

  :groundtruth_file         detections.kw18
  :groundtruth_reader:type  kw18

# ======================= DATA FORMATTER =======================
process split
  :: split_image
  :split_image:type         ocv

connect from tdata.image
        to   split.image

process dm
  :: compute_stereo_depth_map
  :compute_map:type         ocv_bm

connect from split.left_image
        to   dm.left_image
connect from split.left_image
        to   dm.right_image

process merger
  :: merge_image
  :split_image:type         ocv

connect from tdata.image
        to   merger.image1
connect from dm.depth_map
        to   merger.image2

# ====================== TRAINING UTILITY ======================
process stk_trainer
  :: detector_trainer
  :trainer:type         scallop_tk_trainer

connect from merger.net_image
        to   stk_trainer.image
connect from tdata.groundtruth
             stk_trainer.groundtruth
```

pipeline_runner –p training_pipeline.pipe

Interactive Query Refinement:
A User-Driven Search Work Flow

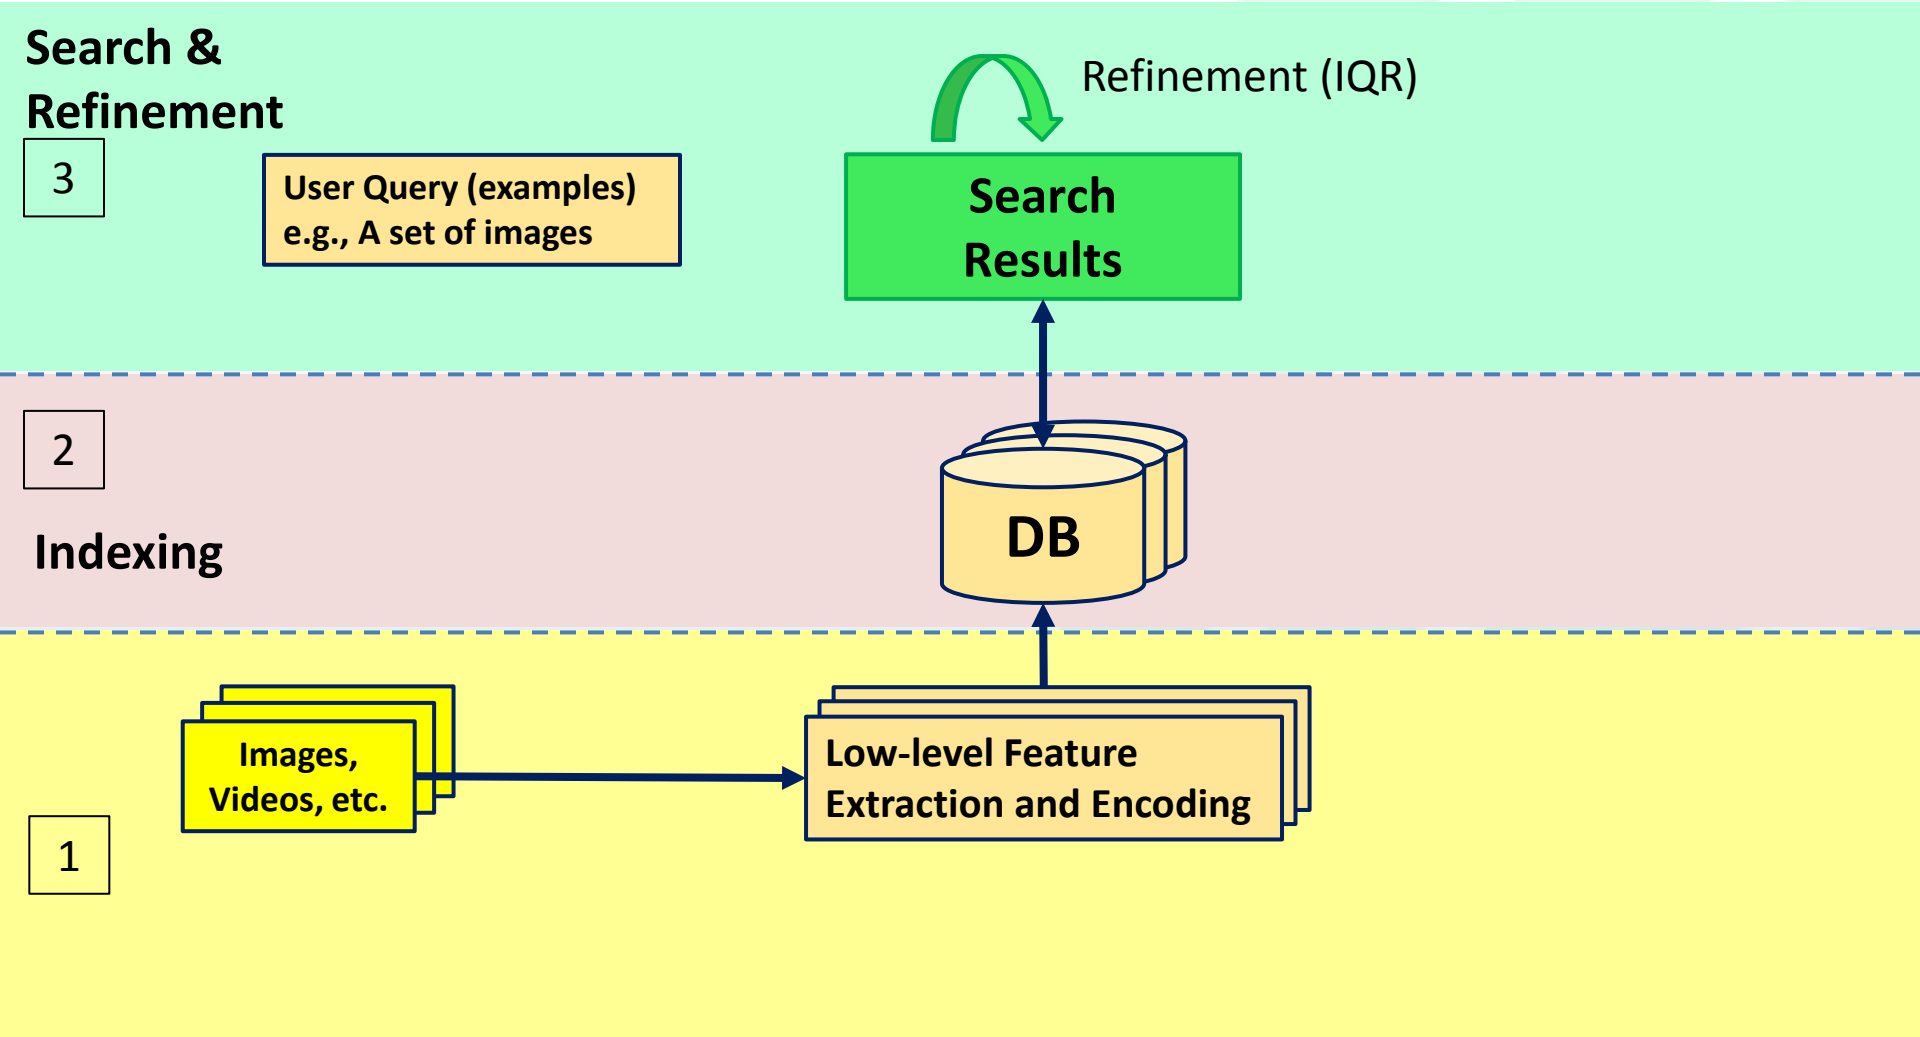

Kitware

+ : good for similarity search
- : no semantic queries

Demo IQR Interface Example

Upload Positive Examples From Disk

(re)Search Data set

Mark Items Relevant or Not Relevant

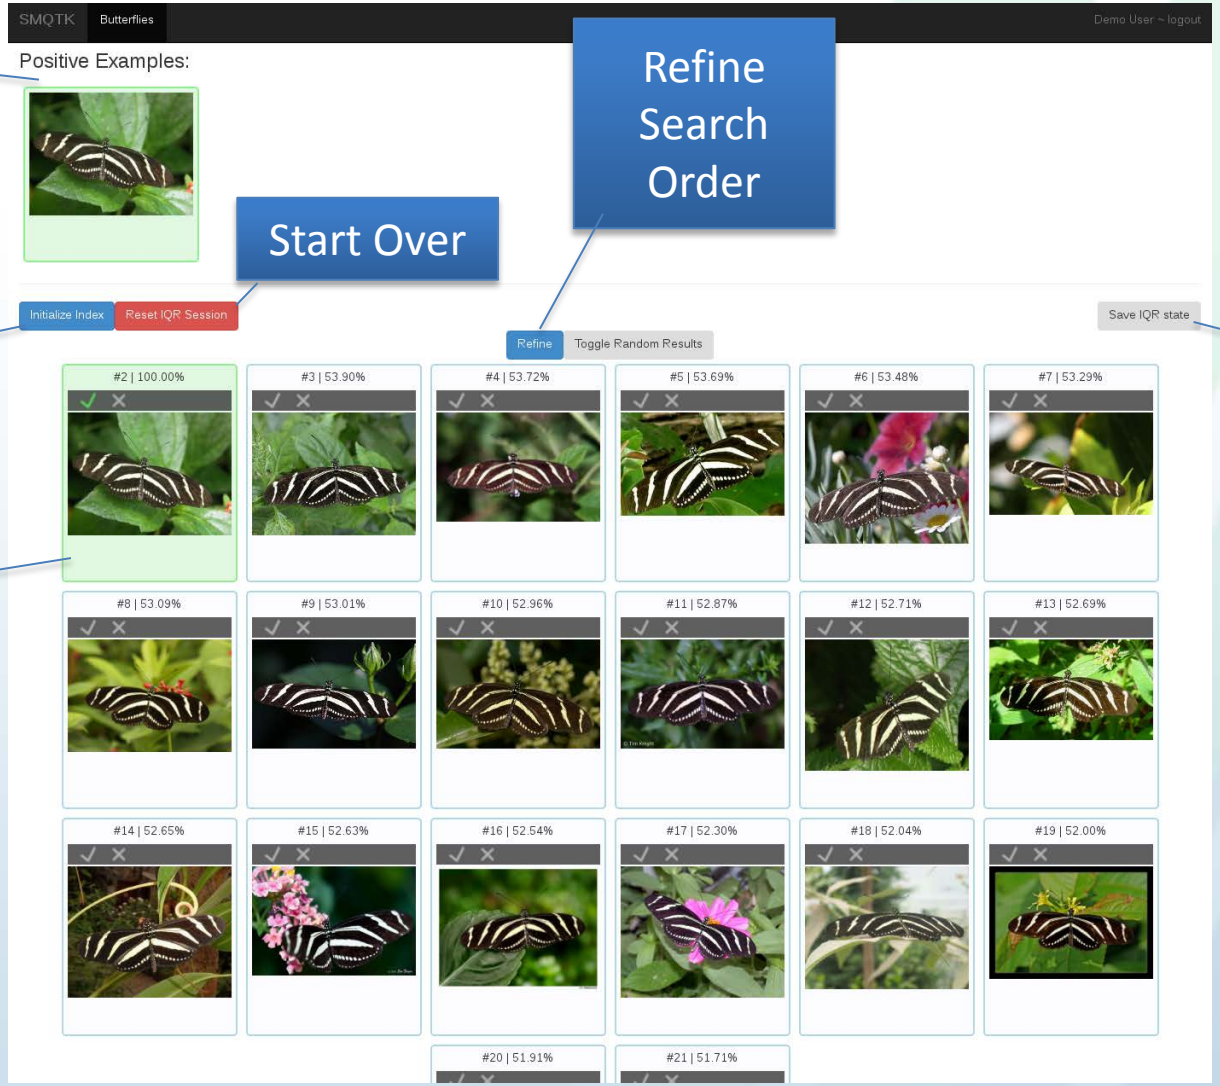

Start Over

Refine Search Order

Export Annotations for Classifier Construction

IQR Example

Start IQR with a single positive exemplar

Dataset contains 832 images with 55-100 images per type.

Use CAFFE AlexNet Layer 7 as an image descriptor

Random Selections from Leeds Butterfly Dataset

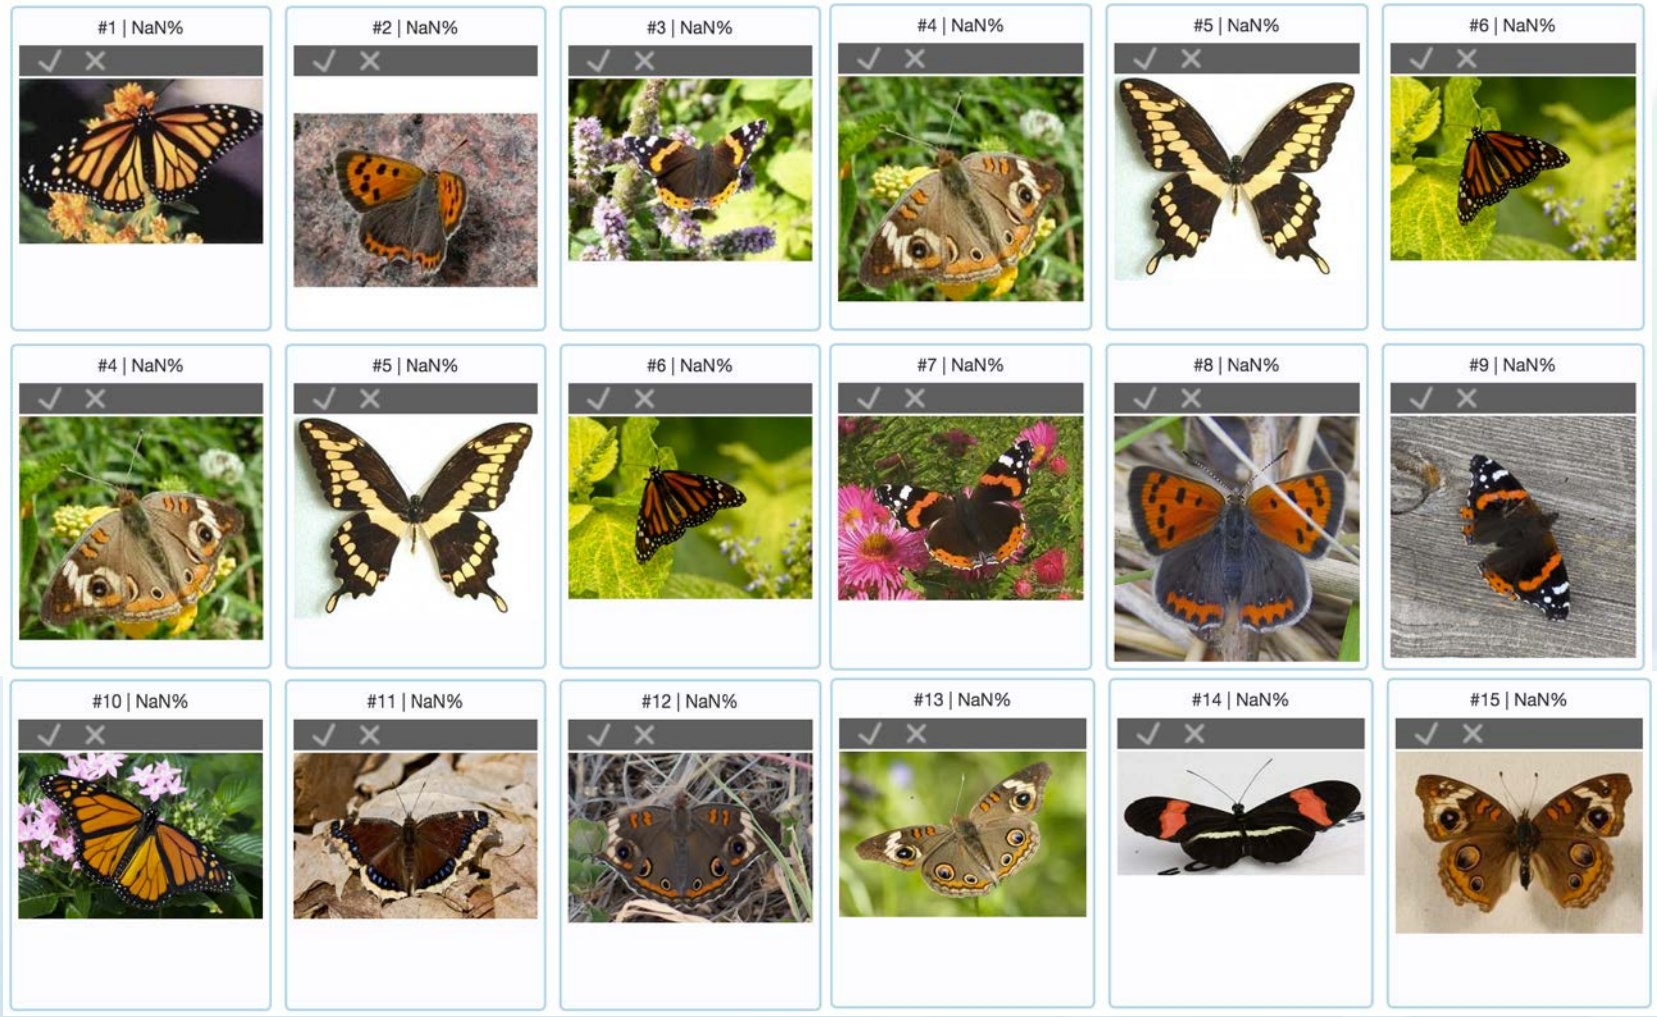

Josiah Wang, Katja Markert, and Mark Everingham
Learning Models for Object Recognition from Natural Language Descriptions
In *Proceedings of the 20th British Machine Vision Conference (BMVC2009)*

Results from Single Exemplar

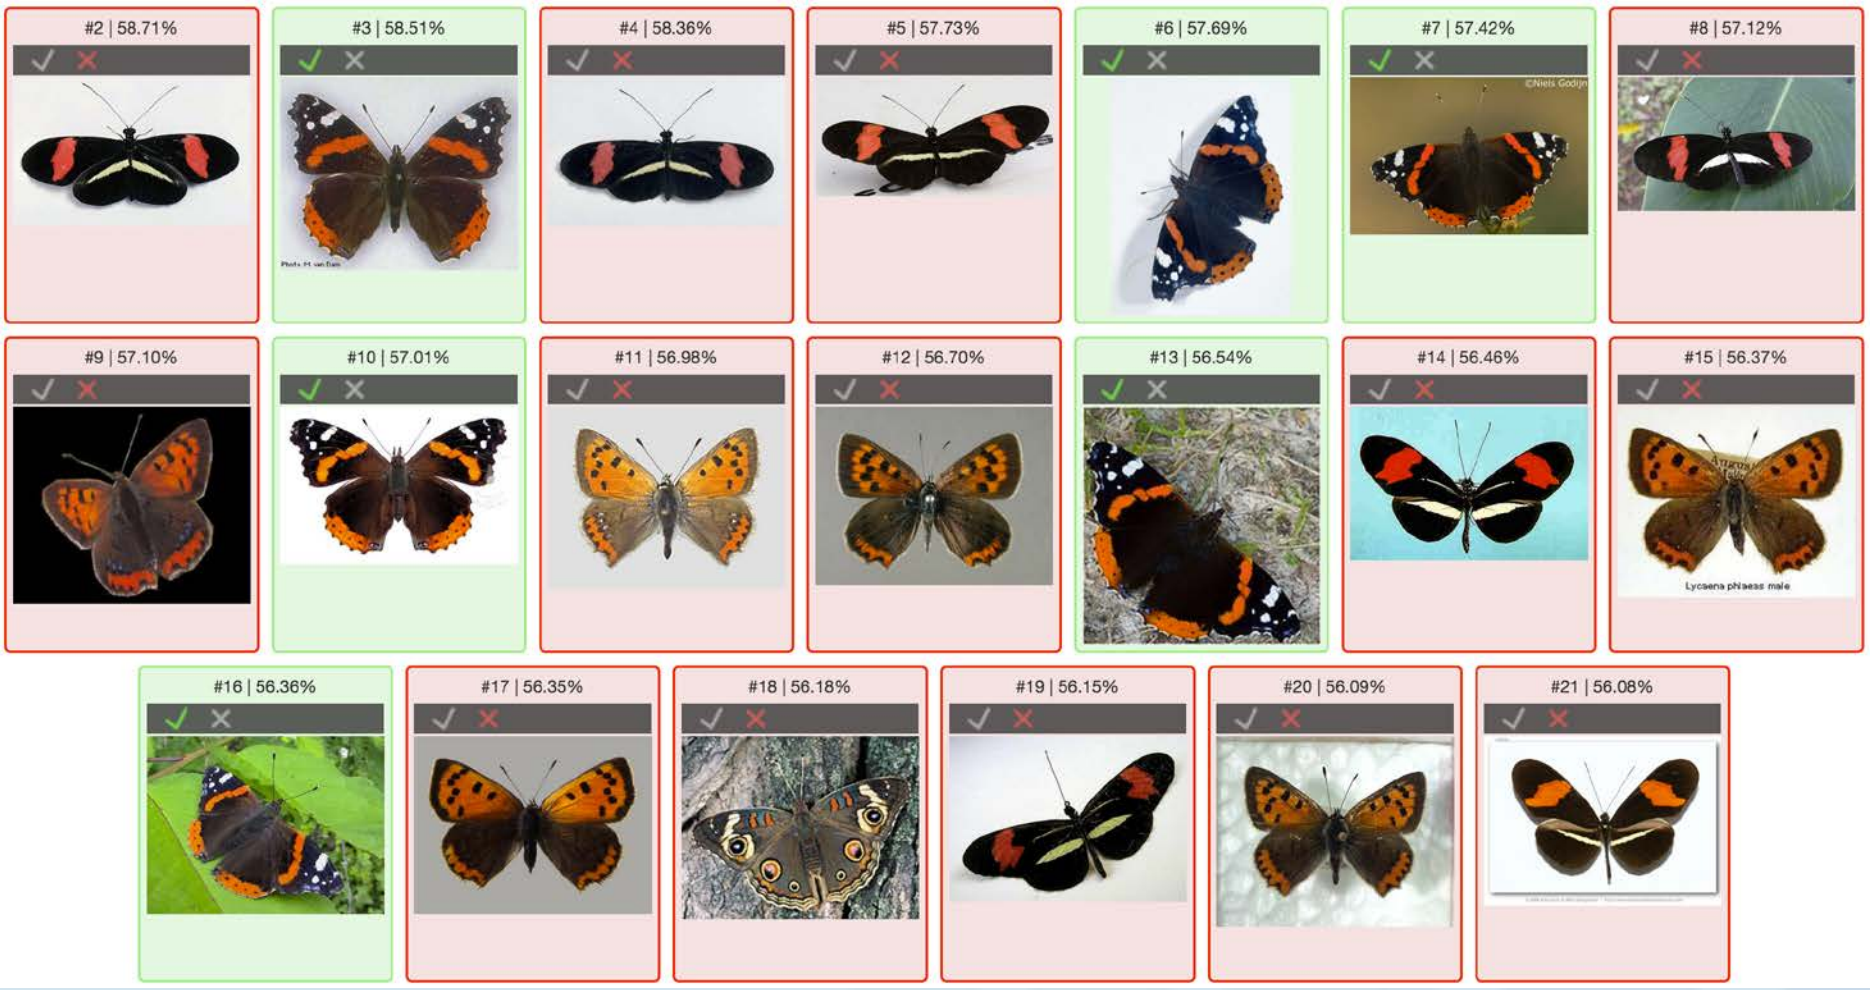

One Refinement Based on Adjudications from Previous Slide

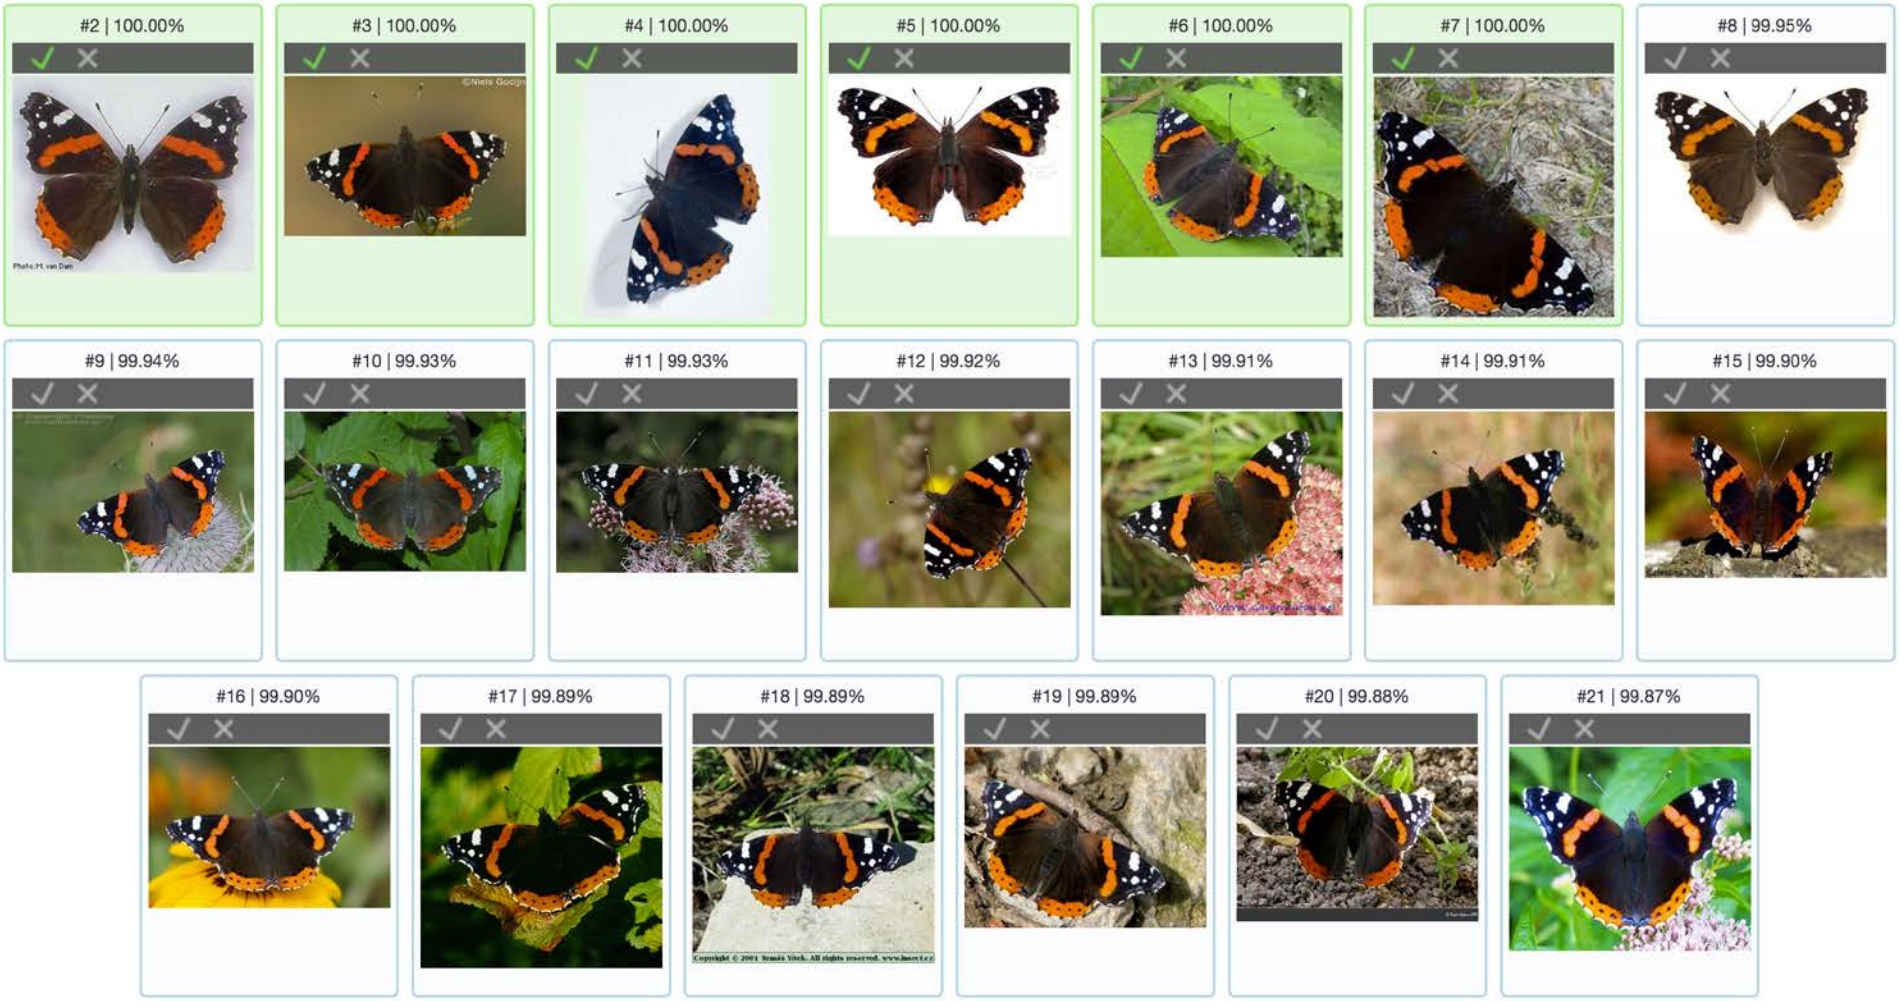

Ongoing and Future Developments

- Improved video handling
- Improved detection and tracking
- Additional stereo processing
  - Calibration
  - Dense 3D reconstruction
- More integrated analytics
  - Anomaly detection and clustering
  - Habitat classification
- Large-scale visualization
- Extend deep learning integration
- Database extensions
- Make system easier to use
- Documentation

Demo

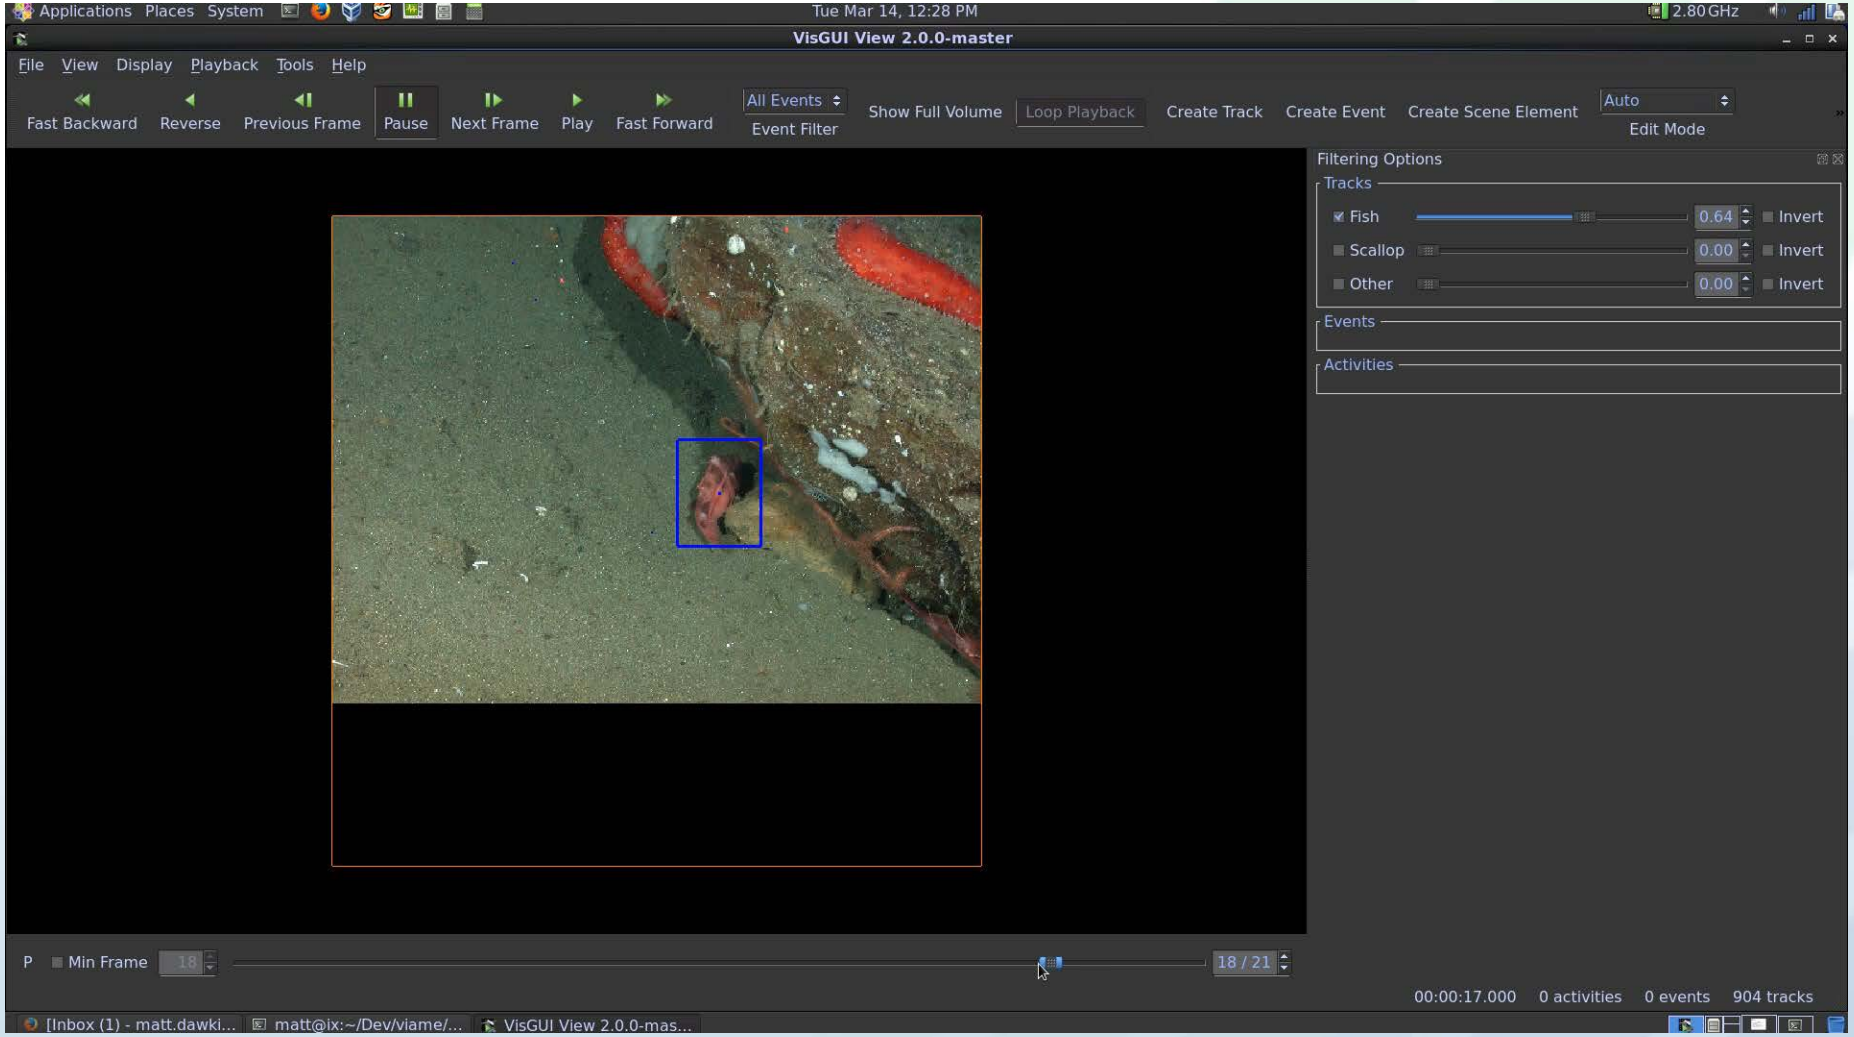

Kitware

VIAME Resources

- VIAME is publicly-available, open-source software
  - viametoolkit.org/
  - Community contributions are highly encouraged, both framework additions and analytics
- Multiple benthic datasets previews are available at marineresearchpartners.com/nmfs_aiasi/Home.html

Thank you NOAA!

Code Repository: https://github.com/Kitware/VIAME

Initial development and testing of VIAME was funded by the NOAA Fisheries Strategic Initiative on Automated Image Analysis. The findings and conclusions in the paper are those of the authors and do not necessarily represent the views of the National Marine Fishery Service, NOAA, or the government of the United States. The use of trade, firm, or corporation names in this publication is for the convenience of the reader and does not constitute an official endorsement or approval of any product or service to the exclusion of others that may be suitable.

Kitware