kitware

Computer Vision Solutions

Using Open Source Technology and Advanced AI

Did we miss you at the booth?

Scan the QR code to learn more about our computer vision capabilities.

Stay connected:
kitware.com
computervision@kitware.com

Learn More about Kitware's Trusted Computer Vision Platforms

KWIVER

The Kitware Image and Video Exploitation and Retrieval (KWIVER) toolkit is an open source, production-quality image and video exploitation platform. It comes with unrestricted licensing to engage the full-spectrum video research and development community, including academic, industry, and government organizations. KWIVER provides advanced capabilities in video object detection and tracking, in addition to supporting tools for quantitative evaluation, pipeline processing, software builds, and more.

VIAME

Through funding and guidance from the National Oceanic and Atmospheric Administration's (NOAA) Automated Image Analysis Strategic Initiative (AIASI), Kitware has developed the Video and Image Analytics for Marine Environment (VIAME) toolkit. VIAME is an open source software platform for DIY AI that enables end-users to customize cutting-edge, deep learning methods for their specific problems, without any programming or knowledge of how AI works. Used by dozens of marine science labs around the world, it is an evolving toolkit that contains many workflows for generating specialized object detectors, full-frame classifiers, and image mosaics; for performing image and video search combined with rapid detector generation; and for stereo measurement, 3D extraction, and camera calibration.

SMQTK

The Social Media Query Toolkit (SMQTK) is a scalable framework with bundled algorithms for indexing, searching, and query refinement on images and video clips. It contains a user-driven search workflow, known as interactive query refinement (IQR), to easily and quickly refine search results through positive and negative adjudication to locate information of interest. SMQTK includes a web based IQR graphical user interface plus RESTful services.

TeleSculptor

TeleSculptor is an open source, cross-platform desktop application for photogrammetry. It was designed specifically with a focus on aerial video processing leveraging video metadata standards (MISB 0601) for geolocation, but it can handle both images and video either with or without metadata. TeleSculptor uses structure-from-motion techniques to estimate camera parameters and a sparse set of 3D landmarks. TeleSculptor estimates dense depth maps on key frames using multiview stereo techniques. It then fuses those depth maps into a consistent surface mesh, which can be colored from the source imagery. TeleSculptor's flexible plugin architecture allows developers to reconfigure or swap out any algorithm in the pipeline for custom implementations. It is both an end-user application and a research platform.

XAITK

TeleSculptor is an open source, cross-platform desktop application for photogrammetry. It was designed specifically with a focus on aerial video processing leveraging video metadata standards (MISB 0601) for geolocation, but it can handle both images and video either with or without metadata. TeleSculptor uses structure-from-motion techniques to estimate camera parameters and a sparse set of 3D landmarks. TeleSculptor estimates dense depth maps on key frames using multiview stereo techniques.

Leaders in Artificial Intelligence, Machine Learning, and Computer Vision

Our computer vision team is a leader in creating cutting-edge algorithms and software for automated image and video analysis. Our solutions harness the power of artificial intelligence and machine learning (AI/ML) to address challenging problems in converting pixels and other data streams into actionable information. We operate within many domains - **land, sea, air, and space** - providing significant value to government agencies. We understand the difficulties in extracting, interpreting, and utilizing information across images, video, other sensors (MSI, HSI, SAR, LiDAR), metadata, and text, and we recognize the need for robust, affordable solutions. We also seek to advance the fields of AI/ML and computer vision through R&D and collaborative projects that build on our open source software tools.

Air

Our air domain tools are the culmination of years of continuous government funding to address automatic target recognition, high value target tracking, and intelligence, surveillance, and reconnaissance. Deployed operationally, we continue to evolve these tools to address our customer's challenges and apply them to multiple sensors on manned and unmanned aircraft.

Space

Space situational awareness and space domain awareness present difficult, unique challenges for detection, tracking and recognition of objects and debris. Kitware develops sophisticated methods to find and characterize low-observable objects to support key Defense and IC missions.

Land

Ground-based cameras - fixed, on mobile platforms, or phones - provide in-situ intelligence but large volumes of data. Kitware develops and fields advanced capabilities in threat detection, activity recognition, 3D vision, and scene understanding to support force protection, ISR and multimedia analytics.

Sea Surface

Kitware researches, develops, tests, integrates, and deploys advanced maritime capabilities addressing detection, classification, and tracking of objects on the sea surface. We fuse imagery and information to build patterns of life models and understand complex maritime activities.

Undersea

In the undersea domain, cameras can provide close-up, high-resolution information from unmanned platforms. Our open-source AI capabilities can identify, classify, and track organic and non organic objects of interest supporting scientific research and to generate actionable intelligence.

Deep Learning

Through our extensive experience in AI and our early adoption of deep learning, we have made dramatic improvements in object detection, recognition, tracking, activity detection, semantic segmentation, and content-based retrieval. Our expertise focuses on hard visual problems, such as low resolution, very small training sets, rare objects, long-tailed class distributions, large data volumes, real-time processing, and onboard processing.

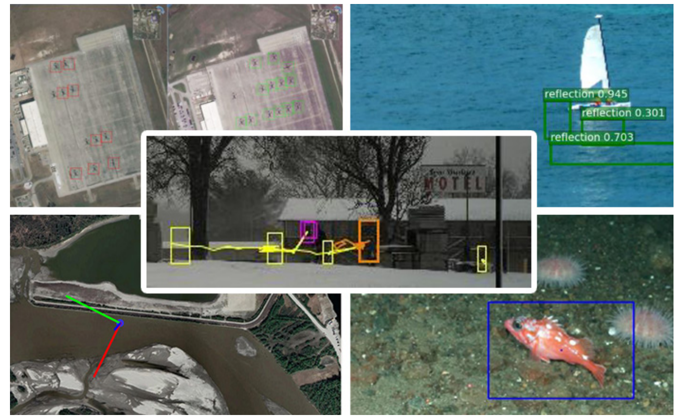

Dataset Collection and Annotation

The growth in deep learning has increased the demand for quality, labeled datasets needed to train models and algorithms. Kitware has developed and cultivated dataset collection, annotation, and curation processes to build powerful capabilities that are unbiased and accurate. Kitware can collect and source datasets, design custom annotation pipelines, and we are able to annotate image, video, text and other data types. Kitware also performs quality assurance that is driven by rigorous metrics to highlight when returns are diminishing.

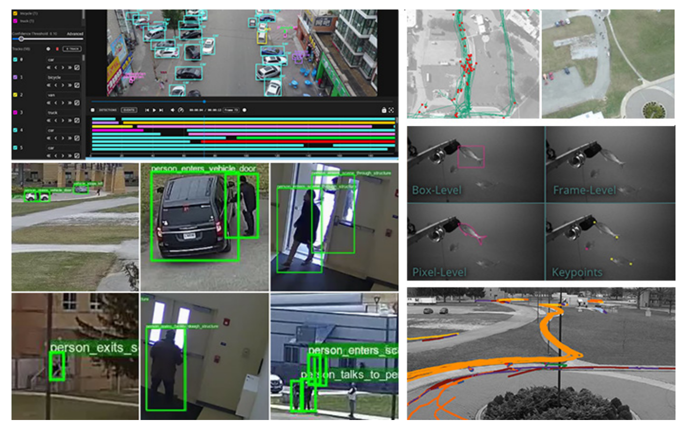

Interactive Do-It-Yourself AI

Using Kitware's interactive DIY AI toolkits, users can rapidly build, test, and deploy novel AI solutions without having expertise in deep learning or computer programming. You can easily train object classifiers using interactive query refinement, without drawing any bounding boxes. Our toolkits also allow you to perform customized, highly specific searches of large image and video archives. Our DIY AI toolkits have both scientific and defense applications and are provided as open source software or with unlimited rights to the Government.

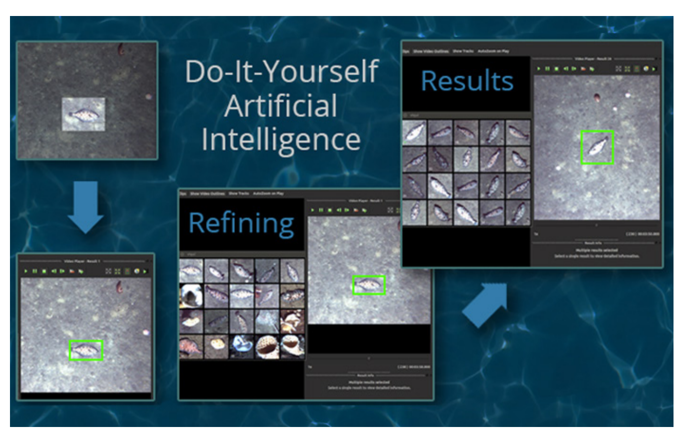

Object Detection, Recognition and Tracking

Our video object detection and tracking tools are the culmination of years of continuous government investment. Our suite of tools can identify and track moving objects in many types of intelligence, surveillance, and reconnaissance data (ISR), including video from ground cameras, aerial platforms, underwater vehicles, robots, and satellites. These trackers are able to perform in challenging settings and address difficult factors, such as low contrast and resolution, moving cameras, and high traffic density, through specialized techniques.

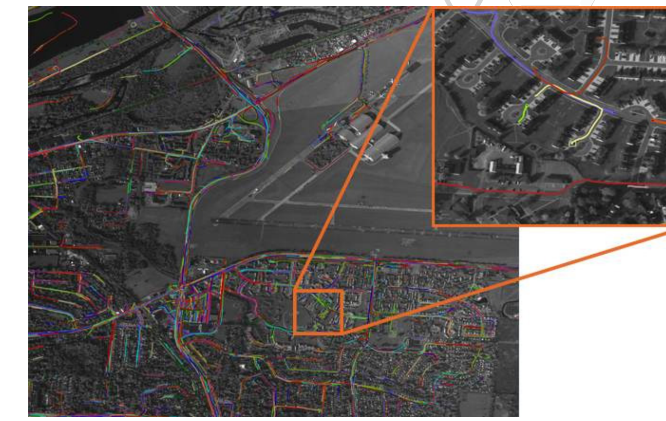

Explainable and Ethical AI

Kitware has developed powerful tools to explore, quantify, and monitor the behavior of deep learning systems. Our team is making deep neural networks explainable and robust when faced with previously-unknown conditions. We are stepping outside of classic AI systems to address domain independent novelty identification, characterization, and adaptation to be able to acknowledge the introduction of unknowns. We also value the need to understand the ethical concerns, impacts, and risks of using AI. Therefore, Kitware is developing methods to understand, formulate and test ethical reasoning algorithms for semi-autonomous applications.

Saliency maps show whether the algorithm is focused on the right object

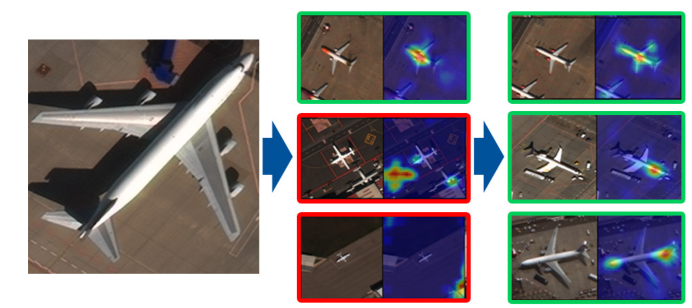

Cyber-Physical Systems (CPS)

Kitware has created state-of-the-art cyber-physical systems that perform onboard, autonomous processing to gather and extract critical data. Using computer vision and deep learning technology, our sensing and analytics systems can overcome the challenges of an ever-changing physical environment. They are customized to solve real-world problems in aerial, ground, and underwater scenarios. These capabilities have been field-tested and proven successful in programs funded by DARPA, AFRL, NOAA, and more.

Complex Activity, Event, and Threat Detection

Kitware's tools recognize high-value events, salient behaviors and anomalies, complex activities, and threats through the interaction and fusion of low-level actions and events in dense cluttered environments. Operating on tracks from WAMI, FMV, MTI or other sources, these algorithms characterize, model, and detect actions, such as people picking up objects and vehicles starting/stopping, along with complex threat indicators such as people transferring between vehicles and multiple vehicles meeting. Many of our tools feature alerts for behavior, activities and events of interest, including highly efficient search through huge data volumes, such as full frame WAMI missions using approximate matching. This allows you to identify actions in massive video streams and archives to detect threats, despite missing data, detection errors and deception.

3D Reconstruction, Point Clouds, and Odometry

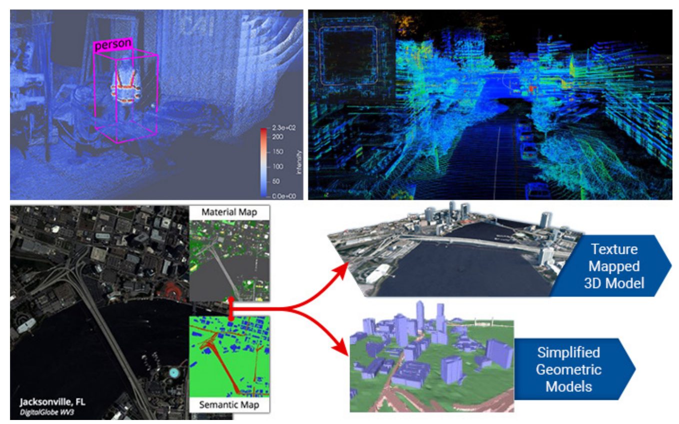

Kitware's algorithms can extract 3D point clouds and surface meshes from video or images, without metadata or calibration information, or exploiting it when available. Operating on these 3D datasets or others from LiDAR and other depth sensors, our methods estimate scene semantics and 3D reconstruction jointly to maximize the accuracy of object classification, visual odometry, and 3D shape. Our open source 3D reconstruction toolkit, Telesculptor, is continuously evolving to incorporate advancements to automatically analyze, visualize, and make measurements from images and video. LiDARView, another open source toolkit developed specifically for LiDAR data, performs 3D point cloud visualization and analysis in order to fuse data, techniques, and algorithms to produce SLAM and other capabilities.

Image and Video Forensics

In this new age of disinformation, it has become critical to validate the integrity and veracity of images, video, and other sources. As photo-manipulation and photo generation techniques are evolving rapidly, we are continuously developing algorithms to automatically detect image and video manipulation that can operate at scale on large data archives. These advanced deep learning algorithms give us the ability to detect inserted, removed, or altered objects, distinguish deep fakes from real images, and identify deleted or inserted frames in videos in a way that exceeds human performance. We continue to extend this work through multiple government programs to detect manipulations in falsified media exploiting text, audio, images, and video.

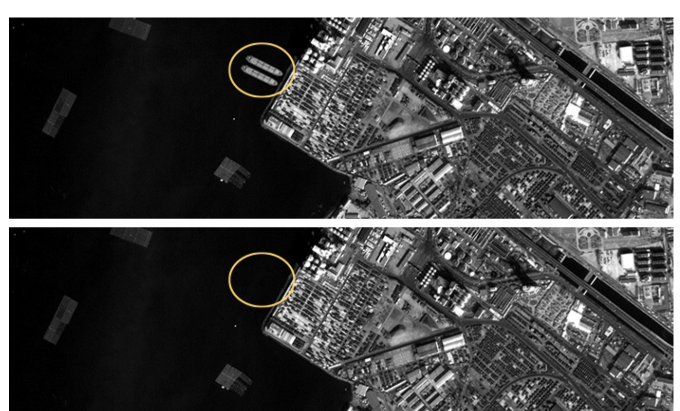

Super Resolution and Enhancement

Kitware's super-resolution techniques enhance single or multiple images to produce higher-resolution, improved images. We use novel methods to compensate for widely spaced views and illumination changes in overhead imagery, particulates and light attenuation in underwater imagery, and other challenges in a variety of domains. The resulting higher-quality images enhance detail, enable advanced exploitation, and improve downstream automated analytics, such as object detection and classification.

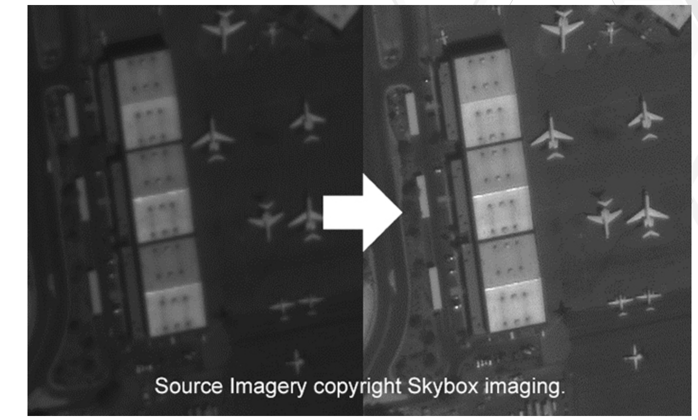
Source Imagery copyright Skybox imaging.

Computational Imaging

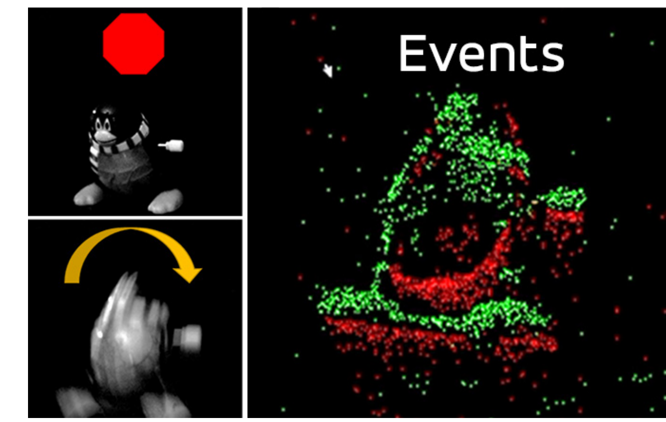

The success or failure of computer vision algorithms is often determined upstream, when images are captured with poor exposure or insufficient resolution that can negatively impact downstream detection, tracking, or recognition. Recognizing that an increasing share of imagery is consumed exclusively by software and may never be presented visually to a human viewer, computational imaging approaches co-optimize sensing and exploitation algorithms to achieve more efficient, effective outcomes than are possible with traditional cameras. Conceptually, this means thinking of the sensor as capturing a data structure from which downstream algorithms can extract meaningful information. Kitware's customers in this growing area of emphasis include IARPA, AFRL, and MDA, for whom we're mitigating atmospheric turbulence and performing recognition on unresolved targets for applications such as biometrics and missile detection.

Scene Understanding

Kitware's knowledge-driven scene understanding capabilities use deep learning techniques to accurately segment scenes into object types. In video, our unique approach defines objects by behavior, rather than appearance, so we can identify areas with similar behaviors. Through observing mover activity, our capabilities can segment a scene into functional object categories that may not be distinguishable by appearance alone. These capabilities are unsupervised so they automatically learn new functional categories without any manual annotations. Semantic scene understanding improves downstream capabilities such as threat detection, anomaly detection, change detection, 3D reconstruction, and more.