



Spatialyze: A Geospatial Video Analytics System with Spatial-Aware Optimizations

Chanwut Kittivorawong*
UC Berkeley
chanwutk@berkeley.edu

Yongming Ge*
UC Berkeley
yongmg@berkeley.edu

Yousef Helal
UC Berkeley
yousefh@berkeley.edu

Alvin Cheung
UC Berkeley
akcheung@cs.berkeley.edu

ABSTRACT

Videos that are shot using commodity hardware such as phones and surveillance cameras record various metadata such as time and location. We encounter such *geospatial videos* on a daily basis and such videos have been growing in volume significantly. Yet, we do not have data management systems that allow users to interact with such data effectively.

In this paper, we describe Spatialyze, a new framework for end-to-end querying of geospatial videos. Spatialyze comes with a domain-specific language where users can construct geospatial video analytic workflows using a 3-step, declarative, *build-filter-observe* paradigm. Internally, Spatialyze leverages the declarative nature of such workflows, the temporal-spatial metadata stored with videos, and physical behavior of real-world objects to optimize the execution of workflows. Our results using real-world videos and workflows show that Spatialyze can reduce execution time by up to 5.3×, while maintaining up to 97.1% accuracy compared to unoptimized execution.

PVLDB Reference Format:

Chanwut Kittivorawong, Yongming Ge, Yousef Helal, and Alvin Cheung. Spatialyze: A Geospatial Video Analytics System with Spatial-Aware Optimizations. PVLDB, 17(9): 2136 - 2148, 2024.
doi:10.14778/3665844.3665846

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://spatialyze.github.io>.

1 INTRODUCTION

Geospatial videos record the locations and periods of the recordings in addition to the visual information. Such videos are prevalent in our daily lives, from surveillance footage to autonomous vehicle (AV) cameras to police body cameras. While the volume of such data has grown tremendously [36, 39], we still lack end-to-end systems that can process and query such data effectively. The rise of machine learning (ML) in recent years has aggravated this issue, with the latest deep learning models capable of carrying out various computer vision tasks, such as object detection [14, 15, 20, 31, 32], multi-objects tracking [5, 7, 9, 38], image depth estimation [17], etc.

All these trends have made geospatial video analytics computationally intensive. For instance, a \$2,299 NVIDIA T4 GPU takes 34 seconds on average to execute a simple workflow of object detection, tracking, and image depth estimation on a 20-second 12-fps video. Modern autonomous driving datasets, such as nuScenes [8], contain 6000 such videos. Running the workflow mentioned above will take 3 full days to run on the entire dataset.

To make matters worse, the lack of programming frameworks and data management systems for geospatial videos has made it challenging for end users to specify their workflows, let alone run them efficiently. For instance, a data journalist writing an article on self-driving car failures would like to examine the footage collected on car cameras to look for a specific behavior, e.g., two cars crossing at an intersection. They will either need to watch all the footage themselves or string together various ML models for video analytics [2, 4, 21, 40] and video-processing libraries (e.g., OpenCV [6], FFmpeg [37]) to construct their workflow. Sadly, the former is infeasible given the amount of video data collected, while the latter requires programming expertise that most users do not possess.

We aim to build a programming framework that bridges video processing and geospatial video analytics. Specifically, as users constructing workflows on geospatial videos are typically interested in the time and location where such videos are taken, and such information is often stored as video metadata, we exploit that to optimize such workflows. Objects captured in videos inherit the physical behavior of real-world objects. For example, when users look for a car at a road intersection, the intersection must be visible in the video frames that the car of interest is visible; therefore, video frames without a visible intersection, which can be determined using geospatial metadata, can be skipped during query processing. We leverage such existing geospatial metadata, the *inherited physical behaviors*, and users' queries to speed up video processing.

Leveraging this insight, we present Spatialyze, a system for geospatial video analytics. To make it easy for users to specify their geospatial video analytics workflows, Spatialyze comes with a conceptual data model where users create and ingest videos into a "world," and users interact with the world by specifying objects (e.g., cars) and scenarios (e.g., cars at an intersection) of interest via Spatialyze's S-Flow, a domain-specific language (DSL) embedded in Python. Spatialyze then efficiently executes the workflow by leveraging various spatial-aware optimization techniques that make use of existing geospatial metadata and assumptions based on the inherited physical behavior of objects in the videos. For instance, Spatialyze's *Road Visibility Pruner* uses the road's visibility as a proxy for objects' visibility to prune out video frames. The *Object Type Pruner* then prunes out objects that are not of interest to the users' workflow. Spatialyze's *Geometry-Based 3D*

*Co-first authors.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 17, No. 9 ISSN 2150-8097.
doi:10.14778/3665844.3665846

Location Estimator speeds up object 3D location estimation by replacing a computationally expensive ML-based approach with a geometry-based approach. Finally, the *Exit Frame Sampler* prunes out unnecessary video frames based on the inherited physical behavior of vehicles and traffic rules. Spatialyze’s optimizations are all driven by various geospatial metadata and real-world physical behavior, with the goal of reducing the number of video frames to be processed and ML operations to be invoked, thus speeding up the video processing runtime.

As we are unaware of end-to-end geospatial video analytics systems, we evaluate different parts of Spatialyze against state-of-the-art (SOTA) video analytic tools (OTIF [4], VIVA [33], EVA [40]), a geospatial data analytic tool (nuScenes devkit [8]), and an aerial drone video sensing platform (SkyQuery [3]) on different geospatial video analytics workflows. We are up to 7.3× faster than EVA in geospatial object detection workflow, 1.06-2.28× faster than OTIF in object tracking, 1.68× faster than VIVA, 1.18× faster than SkyQuery in geospatial object tracking workflow, and 117-716× faster than nuScenes in geospatial data analytics. We also evaluate each of our optimization techniques by executing each of the queries with and without each of the techniques, shown in the results, achieving 2.5-5.3× speed up with 83.4-97.1% accuracy on 3D object tracking.

In sum, we make the following contributions. In §4, we present a conceptual data model and a DSL (S-Flow) for geospatial video analytics. Using S-Flow, users construct their analytic workflows simply by following our *build-filter-observe* paradigm. In §5, we describe the design of Spatialyze, a fast and extensible system to process geospatial videos of arbitrary length in S-Flow. In §6, we develop 4 optimization techniques that leverage geospatial metadata embedded in videos and inherited physical behavior of objects to speed up geospatial video processing in Spatialyze’s geospatial video analytic workflow. In §7, we implemented our techniques in Spatialyze and evaluated it against other SOTA video processing techniques, as well as ablation studies, using real-world videos from nuScenes [8], VIVA [33], and SkyQuery [3].

2 RELATED WORK

Geospatial Video Analytics. Recent use cases for querying geospatial video data (autonomous driving [24], surveillance [18], traffic analysis [16], and transshipment activities [29, 30]) have led to the development of geospatial DBMSs for data storage and retrieval. VisualWorldDB [19] allows users to ingest video data from various sources and make them queryable as a single multidimensional object. It optimizes query execution by jointly compressing overlapping videos and reusing results from them. Apperception [13] provides a Python API for users to input geospatial videos, organize them into a four-dimensional data model, and query video data. While Apperception focuses on organizing and retrieving geospatial video data, it falls short of optimizing query execution. In contrast, Spatialyze provides a comprehensive geospatial video analytic workflow interface for declaratively specifying videos of interest. In addition, it introduces new query optimization techniques by leveraging the inherent geospatial properties of videos.

Video Processing. Researchers have proposed techniques to efficiently detect, track, and estimate 3D locations of objects from

videos. Monodepth2 [17] estimates per-pixel image depth data using only monocular videos. EVA [40] identifies, materializes, and reuses expensive user-defined functions (UDF) on videos for exploratory video analytics with multiple refined and reused queries. However, EVA’s optimization techniques aim at scenarios where multiple queries repeatedly invoke overlapping UDFs. OTIF [4] is a general-purpose object tracking algorithm. It uses a segmentation proxy model to determine if frames or regions of frames contain objects that need to be processed with an object detector, along with a recurrent reduced-rate tracking method that speeds up tracking by reducing frame rate. VIVA [22, 33] is a video analytic system that allows users to specify relationship opportunities of replacing or filtering one ML model with another. It uses these relationships to speed up the video processing time. However, it does not optimize video processing with geospatial metadata by default, leaving substantial performance gains unexploited, as our experiments show. Complexer-YOLO [34] and EagerMOT [23] propose one-step ML models for efficiently tracking 3D objects using LiDAR data, which is larger and not always available compared to cameras and GPS sensors information. Spatialyze overcomes this limitation by not requiring such information to process geospatial workflows, but instead using readily available geospatial metadata to optimize user workflows. In addition, one of our optimization techniques estimates objects’ 3D location 192× faster than the approach that uses Monodepth2 on average. SkyQuery [3] is a drone video sensing platform. It also provides a tool for querying and visualizing sensed aerial drone videos but does not optimize video query processing.

3 USAGE SCENARIO

In this section, we give an overview of Spatialyze using an example. Suppose a data journalist is interested in investigating AV footage reports, and the goal is to find video snippets of other vehicles that crash into the driver’s vehicle at a road intersection.

Naively, the data journalist would watch all the videos individually to identify parts of the videos that are relevant. Even with programming skills, they still need to manually invoke ML models to process large volumes of geospatial videos, and possibly optimize model inference by reducing the number of frames to be processed. They need to then merge these results with geospatial metadata (e.g., road information) to extract vehicle geospatial data, write queries to isolate objects of interest, and link them to the geospatial objects for a cohesive output. This process demands high coding expertise and is slow, tedious, and error-prone.

Spatialyze streamlines this process. Instead of manually writing code to run ML models and extract geospatial information, journalists can use Spatialyze to construct such workflows easily. With Spatialyze, they can construct an analytic workflow with S-Flow’s 3-step paradigm as shown in Listing 1 to: **Build a world**, by integrating the video data with the geospatial metadata (e.g., road network and camera configurations); **Filter the video for parts of interest**, by describing relationships between objects (humans, cars, trucks) and their surrounding geospatial environment (lane, intersection); **Observe the filtered video parts**, in this case, by saving all the filtered video parts into video files for further examination.

Build a world. In line 1, the user creates a world which represents a geospatial virtual environment to be discussed in §4. Line 2 then

initializes the `world` by loading `roadNetwork` data, in this case, a set of the Boston Seaport’s road segments. Lines 3-4 load videos with their camera configurations (locations, rotations, and intrinsics [41]) and associate them to create geospatial videos added into the `world`. Finally, in line 4, they have constructed a data-rich `world` with video cameras, and road segments, all related spatially and temporally.

Filter for the video parts of interest. From the data journalist’s perspective, they can interact with objects and their trackings from each video added to the `world` through `object()` with the `world` initialized. Here, our journalist wants to find video snippets where other vehicles crash into the driver’s vehicle at a road intersection. They do so by adding a filter to the `world` to only include objects that match the provided predicates: that they are cars or trucks (line 7), are within 50 meters of the camera (line 8), are at an intersection (line 8), and are moving towards the camera (in line 9). After `filter` is executed at line 9, the `world` will contain only cars or trucks on intersections that are moving towards the camera, as desired.

Observe the filtered video parts. After applying the filter, the `world` has fewer objects than when it was initialized. When the journalist calls `saveVideos()` in line 10, `Spatialyze` saves only video snippets with objects in the current `world`, significantly shortening the total video length. This allows the journalist to focus only on relevant segments with highlighted objects of interest. As shown, `S-Flow` provides a declarative interface where users do not specify “how” to find the video snippets of interest; instead, they only need to describe “what” their video snippets of interest look like, and `Spatialyze` will optimize accordingly.

From the users’ perspective, `Spatialyze` automatically tracks objects from videos, right after each of them is added to the `world`. Hence, they can filter the `world` using these tracked objects. Internally, `Spatialyze` does not execute any part of the workflow until users call `saveVideos()`. Deferring executions this way allows `Spatialyze` to analyze the entire workflow for efficient execution.

We next discuss `Spatialyze` in detail (in §5) and how it optimizes workflow execution (in §6) by leveraging videos’ geospatial metadata and the physical behavior of real-world objects.

Listing 1 Geospatial Video Analytic Workflow with `Spatialyze`

```

1 world = World() # OR World(detect=CstmDetector, track=CstmTracker)
2 world.addGeogConstructs(RoadNetwork('road-network/boston-seaport/'))
3 world.addVideo(GeospatialVideo(Video('v0.mp4'), Camera('c0.json')))
4 world.addVideo(GeospatialVideo(Video('v1.mp4'), Camera('c1.json')))
5 obj, cam, intersection = ( world.object(), world.camera(),
6                           world.geogConstruct(type='intersection') )
7 world.filter( ((obj.type == 'car') | (obj.type == 'truck'))
8               & (distance(obj, cam) < 50) & contains(intersection, obj)
9               & headingDiff(obj, cam, between=[135, 225]) )
10 world.saveVideos('output_videos/', addBoundingBoxes=True)

```

4 CONSTRUCTING GEOSPATIAL VIDEO ANALYTIC WORKFLOWS

`Spatialyze` comes with `S-Flow`, a DSL embedded in Python for users to construct their geospatial video data workflow. In this section, we describe `S-Flow` with `Spatialyze`’s conceptual data model.

4.1 Conceptual Data Model

Analyzing videos in a general-purpose programming language requires users to manually extract and manipulate object bounding boxes from video frames and join these with geospatial metadata, which is tedious. Hence, we developed `Spatialyze`’s data model, a

higher-level abstraction, to simplify user interaction with geospatial video. This model includes 3 key concepts: 1) *World*, 2) *Geographic Constructs*, and 3) *Movable Objects*, each discussed further below.

4.1.1 World. As discussed in §3 and in line 1 of Listing 1, users construct a *World* as a first step of `S-Flow`. Inspired by `VisualWorldDB`, a *World* represents a geospatial virtual environment which encompasses *Geographic Constructs* and *Movable Objects*, all coexisting with spatial relationships. These relationships, combined with the `S-Flow`, enable users to describe the videos of interest effectively.

The *World* also reflects the physical world, where *Movable Objects* exhibit realistic behaviors, like cars following lanes and adhering to speed limits. We refer to these phenomena as inherited physical behaviors of *Movable Objects*. We will discuss how `Spatialyze` leverages these behaviors to optimize video processing in §6.

4.1.2 Geographic Constructs. Each *World* has a set of *Geographic Constructs*; each *Geographic Construct* has a spatial property that defines its area, represented as a polygon. It also has non-spatial properties including construct ID and construct type. We model each *Geographic Construct* as $(cid, type, ((x_0, y_0), (x_1, y_1), \dots))$ where cid is its identifier; $type$ is the construct type (e.g., lane, intersection); and x_i, y_i is a vertex of the polygon that represents the construct.

4.1.3 Movable Objects. Besides static properties like *Geographic Constructs*, a *World* may contain multiple *Movable Objects*. Being movable, each had spatiotemporal properties, e.g., its location at a given time. Each also has non-spatiotemporal properties, including object ID and object type, that do not change throughout the object’s life span. For example, a *Movable Object* with object type of “car” will always be a “car.” `Spatialyze` models each *Movable Object* as $(oid, type, ((l_0, r_0, p_0, t_0), (l_1, r_1, p_1, t_1), \dots))$, where oid is its identifier; $type$ is the object type, such as car, truck, or human. At time t_i , $l_i = (x_i, y_i, z_i)$ is its location; r_i is its 3D rotation, defined as a quaternion [26]; p_i is its type-specific properties.

A *Movable Object* with $type = camera$ represents a camera that captures videos. It has a type-specific property $p_i = (it_i, \dots)$, where it_i is the intrinsic [1, 41] of the camera at time t_i . Despite being *Movable Objects*, this paper distinguishes the camera from other types of *Movable Objects* for the following reason. *Movable Objects* with $type = camera$ are parts of users’ inputs into the *World*, while `Spatialyze` infers *Movable Objects* with other types from the visual contents captured by the cameras.

We do not include “video” as a part of our conceptual data model. A video itself does not have any geospatial property, so it does not exist in a *World*. Similarly, none of the objects tracked in the video exists initially in the *World*. Internally only after we bind the video with a camera that captures the video, the video gains geospatial properties from the camera and becomes a part of it. Objects tracked in the video also gain their geospatial properties, as a result of the binding, becoming a *Movable Object* in the *World*. `Spatialyze` derives the objects’ geospatial locations by combining their 2D bounding boxes and the camera’s geospatial information.

4.2 S-Flow

We design `S-Flow` based on the data model presented in §4.1, where users interact and manipulate geospatial objects via *Geographic Constructs* and *Movable Objects*, rather than video frames. `S-Flow`’s

language constructs consist of functions that operate on *World*, *RoadNetwork*, and *GeospatialVideo*. Users compose such functions to construct geospatial video analytics workflows.

4.2.1 Camera. Lines 3-4 of Listing 1 defines a *Movable Object* with *type = camera* as discussed in §4.1.3. Users initialize a *Camera* from a list, where the *i*-th element of the list is a dictionary with 4 fields: translation, rotation, intrinsic, and timestamp. They correspond to l_i , r_i , it_i , and t_i , as defined in §4.1.3.

4.2.2 GeospatialVideo. A *GeospatialVideo* is a video enhanced with geospatial metadata. A *Video* (lines 3-4 of Listing 1) by itself contains a sequence of images where each represents a frame of the video. Such video exists spatially in the *World* after its *Video* object is bound to a *Camera* to obtain its spatial properties, where each frame *i* is an image taken from the camera at location l_i , rotation r_i , intrinsic it_i , and at time t_i . An object in each frame obtains its geospatial properties after its track is inferred from *GeospatialVideo*. These objects then become *Movable Objects* as defined in §4.1.3.

4.2.3 RoadNetwork. A *RoadNetwork* is a collection of road segments. Each road segment is a *Geographic Construct* as defined in §4.1.2. Besides being a *Geographic Construct*, a road segment stores its segment heading, which indicates traffic headings in that road segment. A road segment can have multiple segment headings (curve lane) or no segment heading (intersection). Road segments in the *RoadNetwork* dataset that we use have the following road types: roadsection, intersection, lane, and lanegroup. In S-Flow, we initialize a *RoadNetwork* by loading from a directory containing files of the road network from each type. This road network information is accessible and often provided with AV video datasets [8, 12].

4.2.4 World. Each language construct we introduced so far represents users' data. Bringing the concept of *World* described in §4.1.1 into our DSL, users interact with all of their video data through *Geographic Construct* and *Movable Objects* that exists in a *World*. In a Spatialyze workflow, users interact with the *World* with this 3-step paradigm: *build*, *filter*, and *observe*. First, users *build* a *World* by adding videos and geospatial metadata. Once built, users *filter* the *World* to only contain their objects of interest. Finally, users *observe* the remaining objects through annotated videos.

Build. A *World* is initially empty. Users add *RoadNetwork* using `addGeogConstructs()`, followed by adding *GeospatialVideo* via `addVideo()`. At this point of the workflow, the *World* contains *Geographic Constructs*, and *Cameras*. Spatialyze uses YOLOv5 [20] and StrongSORT [9] for object detection and object tracking by default. Spatialyze also supports using UDFs for object detection or tracking, shown as the comment in line 1 of Listing 1.

Filter. Users then construct predicates to filter objects of their interest. As discussed in §3, Spatialyze has not extracted detections and trackings of objects from the videos yet at this point. Nonetheless, Spatialyze provides an interface for users to interact with these objects as if they already exist. For instance, `object()` refers to an arbitrary *Movable Object* with *type ≠ camera* in the *World* to use in the predicate. Calling multiple `object()` gives users multiple arbitrary *Movable Objects* that appear together in the same video. For existing information, `camera()` refers to an arbitrary *Camera*. Similarly, `geoConstruct(type=...)` refers to a *Geographic Construct* of a certain type. Users can construct a predicate involving multiple

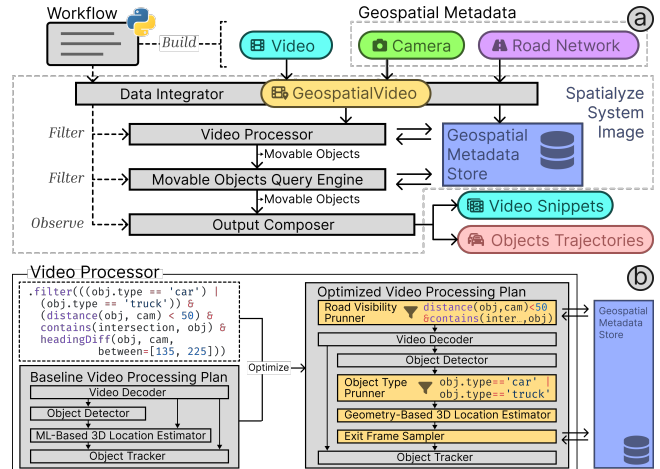


Figure 1: (a) Geospatial Video Analytics Workflow Execution. (b) Processing Plan for the workflow shown in Listing 1.

Movable Objects, a *Camera*, or *Geographic Constructs*, using our provided helper functions. For instance, `contains(intersection, obj)` determines if `intersection` as a polygon contains `obj` as a point. `headingDiff(obj, cam, between=[0, 9])` determines if the difference of the heading of `obj` and the heading of `cam` is between 0 and 9 degrees. In addition, we provide more helper functions and an interface for users to define their own helper functions. Finally, users can chain the `filter` or use boolean operators to combine predicates.

Observe. After filtering their desired *Movable Objects*, users observe the filtered *World*, using one of the following 2 observer functions. First, they can observe the *World* by watching through snippets of the users' own input videos. Using `saveVideos(file, addBoundingBoxes)` function, Spatialyze saves all the video snippets that contain the *Movable Objects* of interest to a file, where `addBoundingBoxes` is a Boolean flag that indicates whether the bounding box of each *Movable Objects* should be shown. Second, they can observe the *World* by getting the *Movable Objects* directly. Using `getObjects()`, Spatialyze returns a list of *Movable Objects*.

5 WORKFLOW EXECUTION IN SPATIALYZE

From users' view, Spatialyze filters objects in the *World* after they call `filter()`. Then, they can *observe* the filtered objects afterward, using one of the observer functions described in §4.2.4. Internally, Spatialyze only performs video processing to detect objects and compute tracks for them *as needed* given the workflow. For example, in line 9 of Listing 1, the user filters on objects' moving directions. Executing this requires computing the trajectory of objects in the videos to compute their directions. Given an S-Flow workflow, the goal of Spatialyze is to understand the users' workflow, recognize and execute only the video processing operators needed for users' filter predicates, and return *Movable Objects* that satisfy the filters.

In §4, we discuss a high-level interface for Spatialyze that allows users to construct their geospatial video analytics workflows declaratively. Users follow the *build-filter-observe* pattern without knowing when the actual execution happens. Internally, all execution happens *when the users observe the world*. In this section,

we discuss the challenges in designing Spatialyze system and how Spatialyze generates execution plans for a workflow.

5.1 Challenges in Workflow Execution

We identify 3 challenges in building a geospatial video analytics system, specifically designing a system that is fast, capable of processing arbitrary-length videos, and extensible. First, processing videos with ML functions is slow; Spatialyze must leverage its knowledge of the entire workflow to decide which video frames and ML functions must be executed. Second, one minute of a 1080p 12fps decoded video is 4.4 GB. Two days of such video can already fill up one GCP VM with the largest available memory size of 11,776 GB as of 2023. Spatialyze must scale with video sizes with limited memory. Finally, Spatialyze must be extensible with future ML functions and optimization techniques. We discuss how we address these challenges in §5.2 and §6.

5.2 System Design & Workflow Execution

To minimize execution time, Spatialyze defers ML inferences, which are expensive, until users *observe* the *World*. This way, Spatialyze can optimize the entire workflow execution. Once users *observe* a *World*, Spatialyze starts executing users’ workflow in 4 stages sequentially: ① Data Integrator, ② Video Processor, ③ Movable Objects Query Engine, and ④ Output Composer; as shown in Fig. 1.

5.2.1 Data Integrator. Spatialyze integrates the input data when users *Build* their *World*. It processes the *Geographic Constructs* by creating tables in a geospatial metadata store to store the *RoadNetwork* and create spatial indexing for every *Geographic Construct*. Then, it joins each *Video* and its corresponding *Camera* frame-by-frame, by their frame number, resulting in a *GeospatialVideo*.

5.2.2 Video Processor. This stage takes in users’ *GeospatialVideo* and users’ filter predicates from their workflow’s *Filter*. It then extracts *Movable Objects* from the *GeospatialVideo*. To do so, it first decodes the video to get image frames. It then executes an object detector to extract bounding boxes of objects from each frame. Based on the bounding box, image frame, and *Camera*, it estimates the 3D location of each object. Finally, it executes the object tracker on the bounding boxes and image frame to get *Movable Objects*. By default, Spatialyze uses an off-the-shelf algorithm to implement each step. The video decoder is implemented using OpenCV [6] as it is the fastest option available. The object detector is implemented using YOLO [20], a SOTA object detector used by our prior work [33, 40] with high accuracy, low runtime, and capability to detect multiple types of objects. The object 3D location estimator is implemented using Monodepth2 [17] as it is the SOTA image depth estimation that only requires an image as an input. The object tracker is implemented using StrongSORT [9] as it is a fast and online method, suited for processing large videos.

We designed the video processor to stream video frames through all the above algorithms. Each video frame is pipelined through decoding, object detection, 3D location estimation, and object tracking; hence our video processor only needs to keep $O(1)$ video frames and can handle videos of arbitrary size, addressing the second challenge. While stream processing of videos has been used

previously [7], applying it to a streaming video processor is not trivial. For example, the original YOLO-StrongSORT object tracker [7] streams video frames. However, naively extending the tracker with our optimization techniques requires hard-coding them into the tracker’s implementation. This does not scale to future trackers. To solve this problem (and address the third challenge), we implemented our video processor as a plan of streaming operators shown in Listing 2. Our baseline operators¹ include: ① Video Decoder; ② Object Detector; ③ 3D Location Estimator; and ④ Object Tracker. Specifically, ① decodes each video to get a sequence of RGB images representing video frames. ② detects objects in each frame from ①, returning each object and its object type and 2D bounding box. A *World* is 3 dimensional where all objects have 3D locations; therefore, ③ estimates the 3D location of each object from ② based on its *Camera*. Lastly, ④ tracks objects by associating detections from ③ between frames. We implemented each operator as an iterator function, where each operator takes in streams of per-frame inputs and returns a stream of per-frame outputs. As a result, our video processor can process videos of arbitrary size, including those larger than the available memory. In addition, we implemented a caching system for the operators. In Listing 2, *frames* is an input to both *ObjectDetector* and *ObjectTracker*. Instead of computing the stream of frames twice, Spatialyze iterates through frames once and caches the results. For instance, a frame i is decoded when the *ObjectDetector* processes it. Spatialyze caches the frame i so that the *ObjectTracker* can process it without having to decode the frame again. Spatialyze keeps track of the number of the operators that will be using a cached result and evicts a frame as soon as all of the operators finished processing it; therefore, Spatialyze only caches a few frames at a given time.

Stateless operators like *ObjectDetector* or *Loc3DEstm* process each video frame independently. In contrast, *ObjectTracker* is a stateful operator that keeps track of what objects have been tracked so far at a given frame. To retain states, an ML function can be implemented as a stateful streaming operator using our provided wrapper with minimal changes to the original code of the function.

Listing 2 Baseline Video Processing Plan

```
1 frames = VideoDecoder('videofile.mp4')
2 object2Ds = ObjectDetector(frames)
3 object2DAnd3Ds = Loc3DEstm(frames, object2Ds, cameras)
4 movableObjects = ObjectTracker(object2DAnd3Ds, frames)
```

Executing the plan computes 3D tracks of objects in a video. However, streaming each video through all operators is computationally expensive and not always necessary. Based on the users’ filter predicates from when they *Filter* the *World*, Spatialyze only includes the necessary operators in the execution plan. For instance, a predicate based on object types like `obj.type=="car"` requires objects’ types and hence the execution plan will include lines 1 and 2. Likewise, a predicate with `distance` or `contain` functions requires objects’ 3D location and therefore include lines 1, 2, and 3 in the plan. Finally, a predicate with a `headingDiff` function requires objects’ moving directions and include all lines in the plan.

We implemented each optimization technique in §6 as a streaming operator. Two of the optimization operators execute users’ predicate early on to reduce the number of inputs (video frames and

¹The **streaming operators** internally process videos and geospatial metadata to extract (2D or 3D) object detections and/or tracks. In contrast, users use **predicate operators** (e.g., `contains`, `distance`) to define filter predicates.

objects) into later operators. We designed our video processor that constructs plans using streaming operators; therefore, future optimization techniques and video processing functions can be added, addressing the runtime efficiency and extensibility challenges.

5.2.3 Movable Objects Query Engine. This stage takes in users’ filter predicates from their workflow’s *Filter* and resulting *Movable Objects* from the video processor. It then filters the *Movable Objects* with the users’ predicates in 2 steps. First, it streams the *Movable Objects* into a newly created table in the geospatial metadata store. The table has a spatial index on objects’ tracks and a temporal index on the period, in which the objects exist. Then, it translates the user’s filter predicates into an SQL query and executes it in the geospatial metadata store against the *Movable Objects*, *Cameras*, and *RoadNetwork*. When the user’s predicates involve multiple objects (e.g., `distance(car1,car2)<5`), Spatialyze self-joins the *Movable Objects* table to find all objects that coexist in a given period using the temporal index to speed up the join. When the predicates involve spatial relationships between a *Movable Object* and a *RoadNetwork* (e.g., `contains(road,car)`), the spatial index speeds up the joins between the *Movable Objects* and *RoadNetwork* tables. We implemented the metadata store using MobilityDB [43] as it provides necessary spatial and temporal indexes for movement objects. The index along with its spatiotemporal data type are suited for processing queries on *Movable Objects* and *Geographic Constructs*.

5.2.4 Output Composer. This stage composes the *Movable Objects* query results to users’ preferred formats for them to *observe*. Calling `getObjects`, Spatialyze returns a list of *Movable Objects*. Calling `saveVideos`, Spatialyze encodes the frames containing the *Movable Objects* into video files with colored bounding boxes annotation.

Fig. 1 shows the execution stages for the workflow in Listing 1. It contains all the stages mentioned above. A user first *builds* the world by adding data. Spatialyze’s Data Integrator creates an indexed table consisting of *RoadNetwork* and joins the *Video* and its corresponding *Camera*. Then, Spatialyze’s Video Processor creates an optimized video processing plan with 4 processing operators and 4 optimization operators (to be discussed in §6). Parts of the users’ predicates are executed in *Road Visibility Pruner* and *Object Type Pruner*. The *Movable Objects Query Engine* streams the *Movable Objects* from the Video Processor into the geospatial metadata store and filters them with the users’ predicates. Finally, as the workflow asks for videos, Spatialyze formats and saves the frames that have the filtered objects into video snippet files.

6 VIDEO PROCESSING OPTIMIZATION

SOTA video inference uses ML models to extract the desired information, such as objects’ positions and trajectories within each video. Running such models is computationally expensive. Spatialyze’s video processor operates with 4 optimization techniques to reduce such runtime bottlenecks. To be discussed below, the idea behind our optimization is to leverage the geospatial properties of the videos, and the *Movable Objects*’ inherited physical behaviors. We implement each optimization as a streaming operator in the Video Processing stage in §5.2.2. In addition to the existing 4 processing operators, Spatialyze may place optimization operators in the middle of the video processing plan to reduce the amount of

input into later processing operators based on the filter predicates. **In §6.1**, Spatialyze places the *Road Visibility Pruner* after the Video Decoder, when users specify the *Geographic Constructs* the objects of interest must be in. Doing so reduces the number of input frames for the rest of the operators in the plan. **In §6.2**, If users filter on object types, then Spatialyze places the *Object Type Pruner* right after the Object Detector. It reduces the number of objects that Spatialyze needs to calculate for their 3D locations and tracking. **In §6.3**, Spatialyze replaces the ML-based 3D Location Estimation operator with the much faster *Geometry-Based 3D Location Estimator*, if Spatialyze can assume that the objects of interest are on the ground, such as cars or people. **In §6.4**, Spatialyze places the *Exit Frame Sampler* in between the 3D Location Estimator and the Object Tracker, if the users’ predicates focus on cars. It utilizes the 3D locations of the objects and the user-provided metadata to reduce the number of frames that the Object Tracker processes.

Fig. 1 shows where the video processor executes all the optimization steps. Line 8 in Listing 1 corresponds to *Road Visibility Pruner*. Line 7 filters by object type, so Spatialyze adds *Object Type Pruner*. The objects of interest are cars or trucks, so Spatialyze adds *Exit Frame Sampler* and *Geometry-Based 3D Location Estimator*. In the following sections, we discuss each optimization in detail.

6.1 Road Visibility Pruner

In our experiments, the runtime of ML models [9, 20, 42] in the video processing plan takes 90% of the entire workflow execution runtime, on average. Hence, the goal of the Road Visibility Pruner is to remove frames that do not contain objects of the user’s interest to reduce the overall runtime of the video processing plan.

6.1.1 High-Level Concepts. The *Road Visibility Pruner* uses *Geographic Construct*’s visibility as a proxy for a *Movable Object*’s visibility. Specifically, a predicate includes `contains(road,obj)` and `distance(cam,obj)<d` means that `obj` is not visible in the frames where `road` is not visible within `d` meters; therefore, Spatialyze prunes out those frames. In Listing 1, the *Road Visibility Pruner* partially executes lines 8 where the user searches for “cars within 50 meters at an intersection.” Spatialyze removes all frames that do not contain any visible intersection within 50 meters as these frames will not contain any car of users’ interest.

The *Road Visibility Pruner* only concerns the visibility of *Geographic Constructs*. Therefore, it uses *Geographic Constructs* and *Camera* to decide whether to prune out a video frame. It uses the intrinsic and extrinsic properties of *Camera* [41] to determine the viewable area of each of the cameras according to the world coordinate system.² It then determines if the viewable area of each frame includes any *Geographic Construct* of interest; frames that do not contain such constructs are dropped and will not be processed in the later processing operators. Since the *Road Visibility Pruner* will be the first processing operator if applied, it can reduce the number of frames that need to be processed by subsequent ML models in the plan, and hence can substantially reduce overall execution time.

6.1.2 Algorithm. The *Road Visibility Pruner* works on a per-frame basis. For each video frame i shot by a *Camera* at location l_i with

² World Coordinate System is a 3-dimensional Euclidean space, representing a *World*.

rotation r_i and camera intrinsic it_i , the *Road Visibility Pruner* performs three steps to determine if the frame should be pruned out. First, it computes the 3D viewable space of the *Camera* at frame i . Then, it uses the viewable space to find *Geographic Constructs* that are visible to the *Camera* at frame i . Finally, it uses the users' filter predicates to determine if any *Geographic Construct* of interest is visible in the frame. We explain each step below.

First, the *Road Visibility Pruner* calculates the 3D viewable space of the *Camera* at frame i . The viewable space is a pyramid (Fig. 2 side view). (A) is the location of the camera. (B), (C), (D), and (E) represent the top-left, top-right, bottom-right, and bottom-left of the frame at d meters in front of the camera. We compute (B), (C), (D), and (E) by converting the mentioned 4 corner points of the frame from pixel (2D) to world (3D) coordinates. This is done in 3 steps by converting a point from: 1) *Pixel to Camera Coordinate System*; 2) *Camera to World Coordinate System*; and 3) *Pixel to World Coordinate System*.

Pixel to Camera Coordinate System. The equation for converting a 3D point (x_c, y_c, z_c) from camera coordinate system³ to a 2D point (x_p, y_p) in pixel coordinate system is:

$$\begin{bmatrix} x_p \\ y_p \\ 1 \end{bmatrix} = \begin{bmatrix} f_x & s & x_0 \\ 0 & f_y & y_0 \\ 0 & 0 & 1 \end{bmatrix} \times \begin{bmatrix} x_c \\ y_c \\ z_c \end{bmatrix} \cdot \frac{1}{z_c}, \text{ where } it_i = \begin{bmatrix} f_x & s & x_0 \\ 0 & f_y & y_0 \\ 0 & 0 & 1 \end{bmatrix} \quad (1)$$

where it_i represents the camera intrinsic, a set of parameters representing *Camera's* focal lengths (f_x, f_y), skew coefficient (s), and optical center in pixels (x_0, y_0). The intrinsic is specific to each camera, independent of its location and rotation. It converts 3D points from a camera perspective to 2D points in pixels.

To convert a 2D point in pixel (x_p and y_p), we can find the 3D point in the camera coordinate system (x_c, y_c , and z_c) if we know the distance of the point from the camera (z_c). From Eq. 1, we have:

$$\begin{bmatrix} x_c & y_c & z_c \end{bmatrix}^T = it_i^{-1} \times \begin{bmatrix} x_p & y_p & 1 \end{bmatrix}^T \cdot z_c \quad (2)$$

Modifying Eq. 2 so that the left side is $\begin{bmatrix} x_c & y_c & z_c & 1 \end{bmatrix}^T$,

$$\begin{bmatrix} x_c \\ y_c \\ z_c \\ 1 \end{bmatrix} = \begin{bmatrix} it_i^{-1} & 0 \\ 0 & 1 \end{bmatrix} \times \begin{bmatrix} x_p z_c \\ y_p z_c \\ z_c \\ 1 \end{bmatrix} = C \times \begin{bmatrix} x_p z_c & y_p z_c & z_c & 1 \end{bmatrix}^T, \text{ where } C = \begin{bmatrix} it_i^{-1} & 0 \\ 0 & 1 \end{bmatrix} \quad (3)$$

Camera to World Coordinate System. For this conversion, we set up R , a 3×3 matrix, representing the 3D rotation of the camera that is derived from the camera's rotation quaternion. And, t is a 3D vector, representing the camera's translation. Then, a 3×4 extrinsic matrix $[R|t]$ converts 3D points from camera coordinate system (x_c, y_c, z_c) to 3D points (x, y, z) in world coordinate system:

$$\begin{bmatrix} x & y & z \end{bmatrix}^T = [R|t] \times \begin{bmatrix} x_c & y_c & z_c & 1 \end{bmatrix}^T \quad (4)$$

Pixel to World Coordinate System. Combining the previous 2 conversions in Eq. 3 and Eq. 4 converts 2D points from a pixel coordinate system to 3D points in the world coordinate system:

$$\begin{bmatrix} x & y & z \end{bmatrix}^T = [R|t] \times C \times \begin{bmatrix} x_p z_c & y_p z_c & z_c & 1 \end{bmatrix}^T \quad (5)$$

4 corners of a video frame in the World Coordinate System. To convert the 4 corners of a video frame from pixel to world coordinates, we modify Eq. 5 where we substitute z_c with d , and x_p and y_p with actual pixel values representing the 4 corners of the video frame; w and h are the width and height in pixels of the video frame, and

³ *Camera Coordinate System* is a 3-dimensional Euclidean space from the perspective of a camera. The origin point is at the camera position. Z-axis and X-axis point toward the front and right of the camera. The Y-axis points downward from the camera.

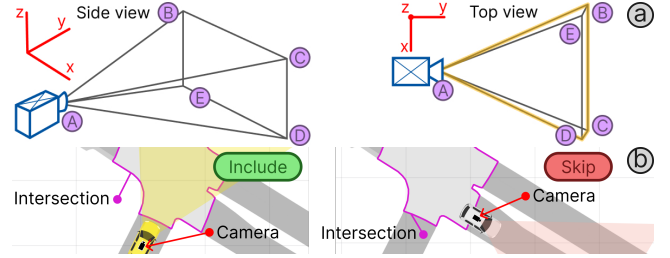


Figure 2: (a) A 3D viewable space of a camera in a pyramid shape and its projected area of the pyramid onto the $z = 0$ plane. (b) In Listing 1, the video processor executes expensive ML models only on video frames with a visible intersection.

each of B, C, D, E is a 3×1 vector representing its 3D location.

$$\begin{bmatrix} B & C & D & E \end{bmatrix} = [R|t] \times C \times \begin{bmatrix} 0 & wd & wd & 0 \\ 0 & 0 & hd & hd \\ d & d & d & d \\ 1 & 1 & 1 & 1 \end{bmatrix} \quad (6)$$

After the corners of each frame are translated to world coordinates, the *Road Visibility Pruner* uses the viewable space to determine the *Geographic Constructs* that are visible by *Camera* at frame i . Each *Geographic Construct's* area is defined using a 2D polygon. Specifically, all polygons representing *Geographic Constructs* are 2D in the $z = 0$ plane. Therefore, the *Road Visibility Pruner* projects the computed 3D viewable space onto the $z = 0$ plane. We define the top-down 2D viewable area as the convex hull [35] of the projected vertices (A), (B), (C), (D), and (E) (yellow highlight in Fig. 2's top view). We then spatially join the 2D viewable area with *Geographic Constructs* in the geospatial metadata store, using their geospatial indexing created in §5.2.1. Any *Geographic Construct* that overlaps with the 2D viewable area is considered visible to the *Camera* at frame i . Once computed, we get `visibleGeogs`, a set of *Geographic Construct* types that are visible to the *Camera* at frame i .

Finally, the *Road Visibility Pruner* uses the users' filter predicates to determine if there is any *Geographic Construct* of interest visible in the frame. Examining the contains predicates in the filter, it assigns each contains to `True` if its first argument is in `visibleGeogs` and `False` otherwise. If the value of the transformed filter predicates is `False`, then the *Road Visibility Pruner* prunes out that frame.

6.1.3 Benefits and Limitations. The *Road Visibility Pruner* has an insignificant overhead runtime compared to the runtime it can save. Our experiments show that the *Road Visibility Pruner* takes 0.1% of the video processing runtime to execute while saving up to 19.6% of the video processing runtime. So in the worst case that *Road Visibility Pruner* cannot prune out any video frames, the runtime increase is still negligible. However, having the *Road Visibility Pruner* in the video processing plan may cause tracking accuracy to drop in some scenarios. Using §3 as an example, if a car starts at an intersection, exits the intersection, and enters another intersection. When the car is not in any intersection, the camera may not capture any frames containing an intersection. The *Road Visibility Pruner* will prune out the period of video frames where no intersection is visible to the camera. The object tracker might consider the car before leaving the intersection and the car after entering another intersection to be two different cars.

6.2 Object Type Pruner

SORT-Family algorithms [5, 9, 38] use the Hungarian method [25] to associate objects between each 2 consecutive video frames. The runtime of the Hungarian method scales with the number of objects to associate in each frame. Reducing the number of objects input into the algorithms reduces the runtime of the object tracker.

6.2.1 High-Level Concepts. Spatialyze only needs to track objects of interest as stated in the workflow. The *Object Type Pruner* prunes out irrelevant objects from being tracked. We analyze users’ filter predicates to look for types of objects that are necessary for computing the workflow outputs. Using the predicates in Listing 1 as an example, the *Object Type Pruner* partially executes the line 7, where users only look for cars and trucks in the workflow outputs. The trajectories of other objects (e.g., humans) will not appear in the final workflow outputs anyway, so tracking them is unnecessary. Therefore, we only need to input objects detected as cars or trucks into the object tracker to reduce its workload.

6.2.2 Benefits and Limitations. The *Object Type Pruner* reduces the object tracker’s runtime. It has low overhead runtime of 0.06% and does not require geospatial metadata but can reduce up to 69% of the tracking runtime, which is 26% of the video processing time.

6.3 Geometry-Based 3D Location Estimator

In our video processing plan, Spatialyze estimates an object’s 3D location using its distance from the camera and calculated bounding box. Specifically, we use YOLOv5 [20] to estimate the bounding box and Monodepth2 [17] to estimate the object’s distance from the camera in our baseline plan. Monodepth2 only requires monocular images to estimate each object’s distance from the camera; therefore, it is flexible in general use cases. Spatialyze can always use Monodepth2 to estimate any object’s distance from the camera in a video frame. However, this algorithm is slow and can take up to 48% portion of the total baseline video processing runtime. Therefore, we explored an alternative algorithm that leverages existing geospatial metadata to speed up 3D object location estimation.

6.3.1 High-Level Concepts. *Geometry-Based 3D Location Estimator* leverages *Camera* to estimate an object’s 3D locations from its 2D bounding box. It is a substitution for Monodepth2 to estimate objects’ 3D locations. *Geometry-Based 3D Location Estimator* uses basic geometry calculations and requires 2 assumptions to hold. First, the object of interest must be on the ground. We can then assume that the lower horizontal border of the object’s bounding box is where the object touches the ground. Second, we can identify the equation of the plane that the object of interest is on. In Spatialyze, the plane $z = 0$ is the ground, although our algorithm can be modified to be used with any plane. The middle point of the lower border of each object’s 2D bounding box represents the point that it touches the ground (red ★ in Fig. 3).

6.3.2 Algorithm. *Geometry-Based 3D Location Estimator* performs 2 steps to estimate the 3D location of an object from its 2D bounding box. First, based on the second assumption, *Geometry-Based 3D Location Estimator* finds an equation that represents a *Vector of Possible 3D Locations* of an object. Second, based on the first and




Figure 3: From a car’s 2D location (★) to its 3D location (★).

second assumptions, *Geometry-Based 3D Location Estimator* finds an *Exact 3D Location* from the vector that intersects with the ground.

Vector of Possible 3D Locations. Using Eq. 5, we convert any 2D point in the pixel coordinate system to a 3D point in the world coordinate system, if we know z_c . Since we do not know z_c yet, we can derive a parametric equation representing all possible 3D points (red line in Fig. 3) converted from the 2D point. Let d be the unknown distance of the object from the camera.

$$[x \ y \ z]^T = [R|t] \times C \times [x_p \ d \ y_p \ d \ 1]^T \quad (7)$$

Exact 3D Location. We assume that the bottom line of the object’s bounding box touches the ground; therefore, (x, y, z) also touches the ground. Because the ground is $z = 0$ in our case, we can solve the Eq. 7 for d . Once we know d , we can solve the parametric Eq. 7 for x and y (blue ★ in Fig. 3).

6.3.3 Benefits and Limitations. *Geometry-Based 3D Location Estimator* relies on the assumptions mentioned, which do not always hold, especially, the first one. Nevertheless, Spatialyze can analyze users’ filter predicates and tell whether the types of objects of interest can be assumed to touch the ground, applying the optimization only then. For example, “car,” “bicycle,” or “human” touch the ground, while “traffic light” does not. In addition, *Geometry-Based 3D Location Estimator* is 192× faster compared than Monodepth2 on average. If one of the estimated 3D locations of an object in a video frame ends up being behind the camera, which means that the object does not touch the ground, Spatialyze can fall back to using Monodepth2 for that frame. Otherwise, *Geometry-Based 3D Location Estimator* provides 99.48% runtime reduction on average for object 3D location estimation and 52% for the video processing.

6.4 Exit Frame Sampler

Common object tracking algorithms [5, 9, 38], including that used in Spatialyze in the object tracker (§5.2.2), performs tracking-by-detection [27]. An object detector identifies objects in video frames but does not link detections across frames to represent the same object. The object tracker associates detections in consecutive frames, creating a chain of associated detections that forms an object trajectory, such as a *Movable Object*. However, not all detections are necessary in an object trajectory. For instance, a car travels in a straight line from frame 1 to 3, where it is at location $(0, 1)$, $(0, 2)$, $(0, 3)$ at frame 1, 2, 3. The in-between detection $(0, 2)$ is unnecessary as the detections at frame 1 and frame 3 produce the same trajectory.

Spatialyze’s *Exit Frame Sampler* utilizes users’ geospatial metadata (*Camera* and *RoadNetwork*) to heuristically prune out video frames that are likely to contain such unnecessary detections, by sampling only the frames that may contain at least one necessary detection. As a result, the object tracker only needs to perform data association on fewer frames, significantly reducing its runtime.

6.4.1 *High-Level Concepts.* Suppose a car is visible to a *Camera* and driving in a lane. As discussed in §4.1.1, a *Movable Object* of type “car” has 2 inherited physical behaviors: it follows the lane’s direction and travels at the speed limit regulated by traffic rules. At a particular frame, the *Exit Frame Sampler* leverages these behaviors to estimate the trajectory of the car so that the object tracker need not track any frame until: i) a car exits its current lane; ii) a car exits the *Camera*’s view; iii) a new car enters the *Camera*’s view. These events are *sampleEvents*. The *Exit Frame Sampler* samples the frame when a *sampleEvent* happens for the object tracker to track.

For (i), when the car is in a lane, without object tracking, the *Exit Frame Sampler* assumes that it follows the lane’s direction until it reaches the end of the lane. After it exits the lane, it may enter an intersection and stop traveling in a straight line. To maintain the accuracy of the trajectory, the object tracker needs to perform data association at the frame right before the car exits the lane. In contrast, if the car is already at an intersection, it may not travel in a straight line; therefore, the object tracker cannot skip any frame.

Event (ii) is the first frame that the car becomes invisible to the *Camera*. If this event happens before the car exits its lane, although its trajectory continues along the lane, the object tracker needs to perform data association at the frame before the car becomes invisible to the *Camera* as the end of the car’s trajectory.

For (iii), if a new car enters the camera’s view, the object tracker needs to start tracking the new car at that frame.

If a frame contains multiple cars, the *Exit Frame Sampler* skips to the earliest frame that (i), (ii), or (iii) is estimated to happen for any car, which prevents incorrect tracking for any of them. Our sampling algorithm implements this idea. Given a current frame and its detected cars, it samples the next frame for the object tracker to track. This algorithm starts from the first frame of a video to find the next frame, skips to it, and repeats until the end of the video.

Because the algorithm depends on the 2 mentioned inherited physical behaviors, it only works when users filter for *Movable Objects* with such behaviors, such as vehicles. Therefore, the video processor only adds the *Exit Frame Sampler* when users’ workflow *Filters* on the object type being vehicles.

6.4.2 *Sampling Algorithm.* Listing 3 shows the pseudocode of each *sampleEvent*. `exitsLane` implements (i) when the car exits the lane. Given the car’s location at the current frame (`carLoc`) and the lane that contains the car, we assume that the car drives in the lane straight along the lane’s direction. In line 2, we model the car’s movement as a motion tuple (`carLoc`, `lane.direction`). We compute the intersection between the car’s motion tuple and the lane as the location where the car exits the lane. Given the car’s speed `v`, we know the time it reaches the exit location starting from the current frame in line 3. We sample the frame before the car reaches the exit location. `exitsCamera` implements (ii) when the car exits the camera’s view. Knowing the car’s current location, moving direction, and speed, `carMoves` computes its location at any given frame in line 8. In line 7, we use the viewable area of the *Camera* at each frame derived from §6.1 to find the first frame when the car’s location is not in the camera’s view. This is the frame when the car already exits the camera’s view, so we sample its preceding frame. `newCar` implements (iii) when a new car enters the *Camera*’s view. In line 12, we sample the next frame that has more cars detected

than the current one. The number of cars is a good heuristic to predict that a new car is entering the frame.

At the current frame, the *Exit Frame Sampler* first finds the frame in which we have more detections for *sampleEvent* (iii). For each car, the *Exit Frame Sampler* samples frames for *sampleEvents* (i) and (ii). Among (i), (ii), (iii), the *Exit Frame Sampler* choose the earliest *sampleEvent* as the next frame for the object tracker to track.

Listing 3 Computing the events that might trigger a frame sample

```

1 def exitsLane(currentFrame, carLoc, lane, v, cameras):
2   exitLaneLoc = intersection((carLoc, lane.direction), lane)
3   exitLaneTime = cameras[currentFrame].time + distance(exitLaneLoc, carLoc)/v
4   return lastFrameBefore(exitLaneTime)
5 def exitsCamera(currentFrame, carLoc, lane, v, cameras):
6   nextFrame = currentFrame + 1
7   while cameras[nextFrame].view.contains(carMoves(carLoc, lane.direction, v,
8     cameras[nextFrame].time - cameras[currentFrame].time)): nextFrame += 1
9   return lastFrameBefore(nextFrame)
10 def newCar(currentFrame):
11   nextFrame = currentFrame + 1
12   while numCars(nextFrame) <= numCars(currentFrame): nextFrame += 1
13   return nextFrame

```

Sampling Algorithm Walk-through. Frame number 1 labeled as *currentFrame* in Fig. 4 was taken at 3:51:49.5 PM. Now, we illustrate how the *Exit Frame Sampler* gets the next frame to sample. In this frame, our object detector only detects one car, *Car_A*. The 3D Location Estimator step calculates its 3D location to be (70.92, 74.7, 0). We join the car’s location with the polygons of traffic lanes in our geospatial metadata store to get the geospatial information of traffic lane *Lane_A* that contains the car, specifically, the polygon of *Lane_A* and its direction to be 181° counterclockwise from the east. We speed up this process by using the geospatial index for the polygons. We assume the car’s speed is 25 mph based on common traffic rules. We also have the viewable area of all frames.

First, we compute `newCar(1)` when new cars enter *Camera_B*’s view. In the rest of the videos, only one car is in the frame. We then compute `exitsLane(1, (70.92, 74.7, 0), LaneA, 25 mph, cameras)` when *Car_A* exits *Lane_A*. In line 2, *Car_A*’s motion tuple intersects the polygon of *Lane_A* at location (66.3, 72.7, 0) at 3:51:49.96 PM in line 3. The last frame before this time is frame 22, depicted as `exitsLane` in Fig. 4. Finally, we compute `exitsCamera(1, (70.92, 74.7, 0), LaneA, 25 mph, cameras)` when the *Car_A* exits the *Camera_B*’s view. The *Car_A* exits the camera view at frame 35 computed in line 8. Thus, we sample the frame 34, depicted as `exitsCamera` in Fig. 4.

Among the 3 *sampleEvents*, `exitsLane` in Fig. 4 is the earliest at frame 22. Thus, the *Exit Frame Sampler* samples the frame before *Car_A* exits its lane at frame 22 labeled as *nextFrame* in Fig. 4, skipping all frames in between. The *Exit Frame Sampler* repeats the algorithm to sample frames until it reaches the end of the video.

6.4.3 *Benefits and Limitations.* To evaluate the efficacy of the *Exit Frame Sampler*, we define “skip distance” as the number of frames skipped between two non-skipped frames. We quantify tracking performance using the *Exit Frame Sampler* by summing the runtime of the sampling algorithm and StrongSORT from the current frame to the next non-skipped frame and dividing that by the runtime without running the *Exit Frame Sampler*. After processing all the videos from the nuScenes dataset, we computed the average for all data points with the same skip distance. As Fig. 4 shows, the runtime ratio decreases as the skip distance increases. With an average skip of 3.6 frames, per-frame runtime is 39% of the original.

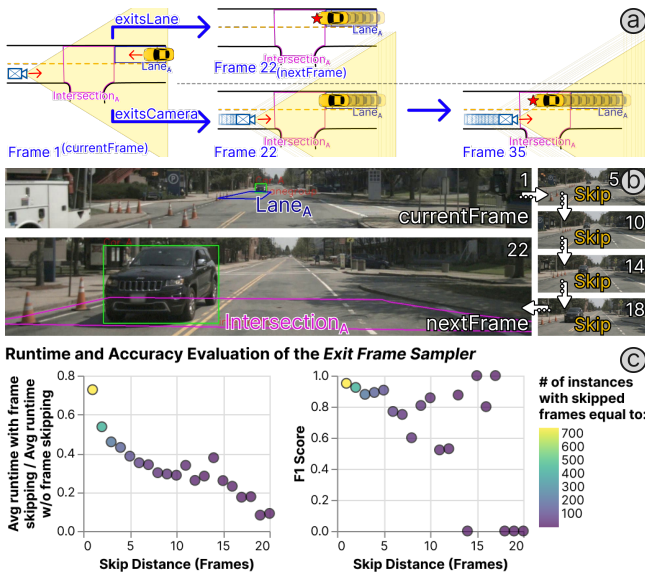


Figure 4: (a) Sample Events exitsLane (i) and exitsCamera (ii). (b) Example Results of our Sampling Algorithm. (c) F1 score and runtime reduction for each skip distance when the *Exit Frame Sampler* can skip at least 1 frame.

Meanwhile, skipping too many frames results in the tracker losing sight of the target vehicles. To assess this risk, we computed the F1 score at each skip distance. The ground truth is the object trackings using StrongSORT without the *Exit Frame Sampler*. Here, “prediction” refers to whether an object at frame f is correctly tracked at frame $(f + \text{skip distance})$ as predicted by the *Exit Frame Sampler*. Fig. 4 shows that the accuracy stays high for skip distances at 13 with the runtime of 28.27% of the original per frame. Hence, we set the maximum skip distance to 13 to maximize this accuracy vs runtime tradeoff. We omit the skip distances over 20 from the graph due to their poor accuracy to focus on effective metrics.

7 EVALUATION

We have implemented a prototype of Spatialyze and evaluated its components against other SOTA systems. We also evaluate our optimization techniques with an ablation study by comparing Spatialyze’s video processing plans with and without each of them. We explain our experiment setup below.

Queries. We implemented Spatialyze’s workflows using four realistic road scenarios, inspired by queries from Scenic [24], VIVA [33], and SkyQuery [3]. These workflows test Spatialyze against real-life situations. Each workflow starts by loading *RoadNetworks* and *Videos* with their *Cameras* from each query’s corresponding datasets, similar to lines 1-4 in Listing 1. Then, we filter the videos with one of the queries described in Table 1, similar to lines 5-9. All queries look for objects closer than 50 meters. Finally, we save the filtered videos into video snippet files, similar to line 10. Table 1 shows the queries in natural language and S-Flow’s filter predicates.

Dataset. We use the multi-modal AV datasets from Scenic [10, 11, 24] to evaluate Spatialyze. They comprise 3 datasets that are needed for executing the 4 queries mentioned: 1) videos from on-vehicle

Table 1: Evaluation query descriptions and their S-Flow predicates (VisualRoad^[vr]; Scenic^[s]; VIVA^[vv]; and SkyQuery^[vv]).

Q1	A pedestrian is at an intersection facing perpendicularly to that of a camera (vr, s). contains(intersection, person) & perpendicular(person, camera)
Q2	2 cars at an intersection moving in opposite directions (vr, s). contains(intersection, [car1, car2]) & opposite(car1, car2)
Q3	Camera moving in the direction opposite to the lane direction, with another car (which is moving in the direction of the lane) within 10 meters of it (vr, s). contains(lane, [camera, car]) & opposite(lane, camera) & sameDirection(lane, car) & distance(camera, car) < 10
Q4	A car and a camera moving in the same direction on lanes; 2 other cars moving together on opposite lanes (vr, s). contains(lane1, [car1, camera]) & sameDirection(car1, camera) & contains(lane2, [car2, car3]) & sameDirection(car2, car3) & opposite(lane1, lane2)
Q5	A pedestrian is at an intersection (vr, s). contains(intersection, person)
Q6	2 cars are at an intersection (vr, s). contains(intersection, [car1, car2])
Q7	A car is on a lane within 10 meters of the camera (vr, s). contains(lane, camera) & distance(camera, car) < 10
Q8	3 cars, each on a lane (vr, s). contains(lane, car1) & contains(lane, car2) & contains(lane, car2)
Q9	A car turning left with a pedestrian at an intersection (vr, vv). contains(intersection, [car, person]) & turnLeft(car)
Q10	A car stopped in a cycling lane (vr, sq). contains(bikeLane, car) & stopped(car)

cameras, 2) the cameras’ movement and specification, and 3) the road network of Boston Seaport where the videos were shot. The former 2 datasets are directly from nuScenes’s training set [8], and the latter is from Scenic. From nuScenes, we randomly sampled 80 out of 467 scenes that are captured at the Boston Seaport from on-vehicle cameras; we chose videos from 3 front cameras for each scene; each video is 20 seconds long at 12 FPS. The runtime of Spatialyze’s execution per video does not depend on the number of input videos; therefore, we have determined that 240 sampled videos sufficiently demonstrate our contributions. To compare Spatialyze with VIVA, we use the VIVA-provided Jackson Square Traffic Camera dataset, which contains 19,469 5-second-long 1080p videos at 30 FPS of a traffic intersection, in addition to the nuScenes dataset. Finally, to compare with SkyQuery, we use their aerial drone geospatial video datasets, which contain 17,853 frames of a 1080p top-down video, along with the per-frame GPS information.

Hardware. In our experiments, we use Google Cloud Compute Engine with Nvidia T4; n1-highmem-4 Machine with 2 vCPUs (2 cores), 26 GB of Memory; and Balanced persistent disk.

7.1 Comparison with Other Systems

Systems selection. Since we are unaware of any end-to-end system that focuses on geospatial video analytics, we chose 5 systems to evaluate Spatialyze, where we selected each system to evaluate the novelty of different components of Spatialyze.

EVA [40]. We evaluate Spatialyze against EVA with its SOTA optimization techniques to speed up video processing.

VIVA [33]. We evaluate Spatialyze against VIVA to emphasize the novelty of our optimization techniques to speed up our end-to-end geospatial video query processing. In addition, the videos, with which VIVA uses to evaluate, are long; we use this fact to evaluate Spatialyze’s capability of processing long videos. Finally, VIVA uses a different ML function [38] to track objects; we evaluate the video processor’s capability of supporting the function.

NuScenes Devkit [8]. As the official tool to process NuScenes’ geospatial movement data, we evaluate the runtime efficiency and the capability of processing large amounts of movement objects of the Movable Objects Query Engine against the NuScenes Devkit.

OTIF [4]. We evaluate Spatialyze against OTIF to emphasize the novelty of our optimization techniques to speed up object tracking.

SkyQuery [3]. We evaluate Spatialyze against SkyQuery’s geospatial video query processor. As SkyQuery uses different object detector, tracker, and 3d-estimator, we evaluate the video processor’s ability to support different ML functions and how the existing optimization techniques perform with the new ML functions.

7.1.1 EVA. We compare the runtime of the Spatialyze pipeline with that of EVA’s implementation using the same queries. EVA is a VDBMS that allows custom user-defined functions (UDFs) along with their runtime optimizations. Since EVA executes the queries on a frame-by-frame basis, it cannot compute object tracks (and hence object directions), so we use EVA to evaluate only the object detection capability of Spatialyze with Q5-8. To use EVA’s caching mechanism, we execute the queries on EVA when they are run in series without resetting EVA and its database (running Q6 directly after Q5, Q7 directly after Q6, etc). The results are shown in Fig. 5. Our goal in comparing against EVA is to evaluate the efficiency of the two system’s query optimization techniques in §6.1 and §6.3.

As shown, Spatialyze performs 2-7.3× faster than EVA on Q5-7, since Spatialyze does not need to run Monodepth2 on each frame, due to the *Object Type Pruner* and *Geometry-Based 3D Location Estimator*. In addition, Spatialyze shaves even more time by employing the *Road Visibility Pruner* to avoid costly YOLO on frames where they are unnecessary. For Q8, our processing time is comparable to EVA. This is because Spatialyze’s Movable Objects Query Engine finds every possible triplet of cars; thus it performs 2 self-joins internally, while EVA only returns video frames with 3 cars or more. We also compare the lines of code (LoC) counts between Spatialyze (14-30) and EVA (58-60), showing that Spatialyze makes query implementation easier and requires less LoC for the same results.

7.1.2 VIVA. We compare the run-time of the Spatialyze workflow with that of VIVA using the traffic application query from VIVA’s evaluation (Q9 in Table 1) as it is VIVA’s only query that involves geospatial content and road network. We run this query against both the VIVA-provided Jackson Square dataset as well as the nuScenes dataset. To match VIVA’s algorithm, we also resize Spatialyze’s input video to a resolution of 360×240 , resample the video at 1 FPS, and employ DeepSORT [38] for tracking rather than StrongSORT to demonstrate Spatialyze’s extensibility. We notice that Spatialyze can execute the same workflow at the videos’ native resolution and FPS while VIVA crashes, which demonstrates Spatialyze’s ability to process large videos and the runtime efficiency of the video processor as a result of Spatialyze’s optimization.

Shown in Fig. 5, Spatialyze performs 1.68× faster compared to VIVA when run on VIVA’s dataset, and 6× faster when run on the nuScenes dataset. We attribute this increase in performance to Spatialyze’s *Object Type Pruner* which eliminates the amount of time spent tracking unnecessary objects. VIVA also spends significantly more time creating an optimization plan.

In addition, we compare the LoC of the implementations of VIVA (56) and Spatialyze (51), with Spatialyze providing a shorter implementation. While we do not have the ground truth to compare with, Spatialyze has higher accuracy in the query results: 45% of tracks returned by Spatialyze are indeed cars turning left. Meanwhile, only 35.7% of the tracks returned by VIVA are correct.

7.1.3 nuScenes. We compare the run-time of the *Movable Objects Query Engine* stage with the same queries implemented using the nuScenes Devkit [8]. NuScenes Devkit is a Python API used to query results from the nuScenes dataset [8]. We only compare the Movable Objects Query Engine stage as Devkit operates on already processed videos and ingested object annotations.

The Devkit implementation of Q4 ran out of memory during execution due to the creation of too many Python objects, as a result of materializing all possible combinations of 3 object tuples in memory before filtering them. Shown in Fig. 5, it also takes significantly longer to execute the same query as Spatialyze, with Spatialyze achieving a 117-716× faster runtime. Two of the main factors that allow Spatialyze to achieve this better execution time are the following. First, Spatialyze pre-processes the road network data, generating additional columns and indexes that allow it to avoid costly joins, which contribute greatly to the large execution time of Devkit. Second, certain Devkit functions perform costly linear algebra that is avoided by Spatialyze using the geospatial metadata store. Comparing the line counts between the different systems, it took 55-123 lines in Devkit to implement the queries while it only took 12-29 lines to do so in S-Flow.

7.1.4 OTIF. We compare the run-time of the video processor of Spatialyze with that of OTIF on object tracking. We run each of the two systems on the sampled nuScenes datasets and compare the number of frames processed per second for each system. OTIF’s training is 61m37s. After it is trained, OTIF tracks objects at 17.3 FPS. For Spatialyze, our video processor applies all optimizations elaborated in §6, and it processes videos at 18.3-39.5 FPS to get the tracks for each query. Without counting OTIF’s training time, Spatialyze still tracks objects faster.

7.1.5 SkyQuery. SkyQuery is an aerial drone sensing platform for detecting and tracking vehicles from top-down videos. A user inputs an aerial drone’s video and its geospatial data, then constructs a query to detect and track objects. Spatialyze and SkyQuery both process geospatial videos by 1) detecting objects of interest, 2) finding the objects’ 3D locations, and then 3) tracking the objects.

To compare with SkyQuery, we implemented the 3 steps mentioned as Spatialyze’s video processing steps, and used SkyQuery’s customized YOLOv3 for car detection, their 3D location estimator, and SORT [5] for object tracking. We retained the rest of Spatialyze’s components as well as the optimization steps. Then, we executed both of the systems on Q10 on SkyQuery’s video and geospatial metadata, and compared execution times. As shown in Fig. 5, Spatialyze processes 6.08 FPS, while SkyQuery processes 5.15 FPS, making us 18% faster. We write 15 lines of S-Flow vs 4 lines of SkyQuery to express Q10. The majority of the lines in S-Flow are spent on setting up the world and inputting videos and geospatial metadata. While we input videos and geospatial metadata into Spatialyze using S-Flow, we need to specify such input directories in the command line when starting SkyQuery.

As our evaluation uses the same ML models as SkyQuery, our speedup is entirely due to Spatialyze’s leveraging of the query’s semantics to prune video frames and avoid executing ML functions on them. Specifically, Spatialyze uses the *Road Visibility Pruner* to exclude video frames that do not contain any cycling lanes. We do not evaluate our accuracy in this evaluation as we use the same ML

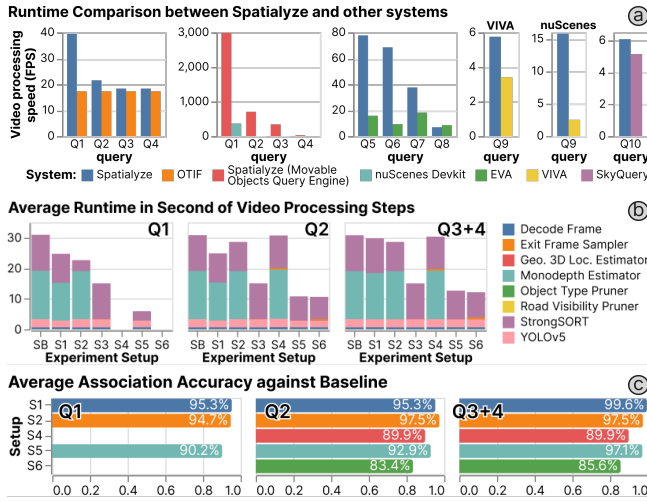


Figure 5: (a) Video frames processed per second for each system. (b) Average video processing time for each video of 20 seconds, comparing each optimization technique. (c) AssA against (SB) vs. other experiment setups.

inference functions as those in SkyQuery. As the only optimization technique applied in this evaluation is the *Road Visibility Pruner*, it only prunes out video frames that do not contain a cycling lane, and thus, does not affect the query result.

7.2 Ablation Study

To evaluate each of our optimization techniques, we manually construct video processing plans for the following 7 experiment setups. (SB) **Baseline** is a Baseline plan without any optimization. It includes the 4 processing operators discussed in §5.2.2. (S1) is the Baseline plan with the *Road Visibility Pruner*. (S2) is the Baseline plan with the *Object Type Pruner*. (S3) is the Baseline plan with the Monocular Depth Estimator replaced with the *Geometry-Based 3D Location Estimator*. (S4) is the Baseline plan with the *Exit Frame Sampler*. (S5) enables all optimizations except the *Exit Frame Sampler*. (S6) enables all optimizations. We evaluated Spatialyze using Q1-4 because Spatialyze can apply all or most of the optimization techniques for the execution of such queries; therefore, we can compare the effects of each optimization technique.

7.2.1 Execution time. The average runtime of our workflow execution without any optimization is 34 seconds for each video. When broken down, 0.01% of the time is spent in the *Data Integrator*; 89.9% in *Video Processor*; 9.5% in *Movable Objects Query Engine*; 0.6% in *Output Composer*. In this section, we evaluate the effect of each of our optimization techniques, by executing Spatialyze on the Q1-Q4. Because the video processing plans for Q3 and Q4 are the same, we do not discuss Q4 in this section.

With all optimization techniques, we achieve 2.5-5.3× speed up over the baseline plan, as shown in Fig. 5. The *Road Visibility Pruner* affects the entire plan. It prunes out 3.8% of the frames that do not contain “lane” in Q3 and 21.5% of the frames that do not contain “intersection” in Q1-2, with 3.2-20.0% runtime reduction. The *Object Type Pruner* only affects StrongSORT. It prunes out

36.5% of objects that are not cars in Q2-4 or trucks and 86.3% that are not pedestrians in Q1, with 7.0-26.3% runtime reduction. The *Geometry-Based 3D Location Estimator* makes the runtime portion of 3D location estimation insignificant (from 48% to 0.55%). Q1 does not execute the *Exit Frame Sampler* as it only works with cars or trucks. The *Exit Frame Sampler* has an overhead runtime of executing §6.4.2; however, it reduces the number of frames into StrongSORT (13.1%), with 0.8% runtime reduction. We also evaluate the execution time with respect to video length by varying the video length input into the workflow. We found that the execution time of the workflow is linearly proportional to video length and Spatialyze can process videos larger than the available memory.

7.2.2 Accuracy. Our optimization techniques potentially reduce object tracking accuracy as a result of dropping frames. To quantify this, we evaluate our optimization techniques using the execution results from (SB) setup as the ground truth. Then, we compare the execution results from other setups against that of the (SB).

We evaluate our tracking accuracy of each experiment setup by comparing the tracking output from the experiment setup with the tracking output from the (SB) setup, using HOTA [28], an evaluation metric for measuring multi-object tracking accuracy. Using the ground truth tracks from a video and prediction tracks from the same video, HOTA returns an Association Accuracy (AssA) score. In this evaluation, the ground truth tracks are the tracking results from the (SB) setup, and the prediction tracks are the tracking results from each of the other setups. We also exclude object detections from the frame pruned by the *Road Visibility Pruner* because this pruning is a part of users’ predicates and will not be output.

In Fig. 5, the *Road Visibility Pruner* (S1) causes a small AssA drop of 0.4% for Q3-4 and 4.7% for Q1-2. The *Object Type Pruner* (S2) also causes a small AssA drop of 2.5% from Q2-4 and 5.3% for Q1. We do not include (S3) into the results since the object tracker only concerns the 2D bounding boxes of objects; *Geometry-Based 3D Location Estimator* does not affect the AssA of the object tracker. With all 3 optimization techniques (S5), the video processor still produces accurate object tracks of 93.4% on average compared to (SB), while speeding up 2.5-5.3× of its runtime. With all optimization techniques (S6), the video processor’s accuracy for object tracking is 84.5% on average, while speeding up 1.65-4.32%, additionally.

8 CONCLUSION

We presented Spatialyze, a system for geospatial video analytics that leverages geospatial metadata and physical properties of objects in videos to speed up video processing. Spatialyze can process videos of arbitrary length by streaming each frame through the operators in workflows expressed using S-Flow. Spatialyze’s DSL allows users to specify their geospatial video workflows declaratively, and our evaluation shows Spatialyze’s efficiency in processing such workflows when compared to SOTA systems.

ACKNOWLEDGMENTS

We thank Daniel Kang, Francisco Romero, Favyen Bastani, Edward Kim, and our anonymous reviewers for their insightful feedback. This work is supported in part by NSF grants IIS-1955488, IIS-2027575, ARO W911NF2110339, ONR N00014-21-1-2724, and DOE awards DE-SC0016260 and DE-SC0021982.

REFERENCES

- [1] Aqeel Anwar. 2022. What are Intrinsic and Extrinsic Camera Parameters in Computer Vision? <https://towardsdatascience.com/what-are-intrinsic-and-extrinsic-camera-parameters-in-computer-vision-7071b72fb8ec>. Accessed: 2023-07-25.
- [2] Favien Bastani, Songtao He, Arjun Balasingam, Karthik Gopalakrishnan, Mohammad Alizadeh, Hari Balakrishnan, Michael Cafarella, Tim Kraska, and Sam Madden. 2020. MIRIS: Fast Object Track Queries in Video. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data* (Portland, OR, USA) (SIGMOD '20). Association for Computing Machinery, New York, NY, USA, 1907–1921. <https://doi.org/10.1145/3318464.3389692>
- [3] Favien Bastani, Songtao He, Ziwen Jiang, Osbert Bastani, and Sam Madden. 2021. SkyQuery: An Aerial Drone Video Sensing Platform. In *Proceedings of the 2021 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software* (Chicago, IL, USA) (Onward! 2021). Association for Computing Machinery, New York, NY, USA, 56–67. <https://doi.org/10.1145/3486607.3486750>
- [4] Favien Bastani and Samuel Madden. 2022. OTIF: Efficient Tracker Pre-Processing over Large Video Datasets. In *Proceedings of the 2022 International Conference on Management of Data* (Philadelphia, PA, USA) (SIGMOD '22). Association for Computing Machinery, New York, NY, USA, 2091–2104. <https://doi.org/10.1145/3514221.3517835>
- [5] Alex Bewley, Zongyuan Ge, Lionel Ott, Fabio Ramos, and Ben Upcroft. 2016. Simple online and realtime tracking. In *2016 IEEE International Conference on Image Processing (ICIP)*. 3464–3468. <https://doi.org/10.1109/ICIP.2016.7533003>
- [6] Gary Bradski. 2000. The OpenCV Library. *Dr. Dobbs' Journal of Software Tools* 25, 11 (Nov. 2000), 120, 122–125. http://www.ddj.com/ftp/2000/2000_11/opencv.txt
- [7] Mikel Broström. 2022. Real-time multi-camera multi-object tracker using YOLOv5 and StrongSORT with OSNet. https://github.com/mikel-brostrom/Yolov5_StrongSORT_OSNet.
- [8] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. 2019. nuScenes: A multimodal dataset for autonomous driving. <https://doi.org/10.48550/ARXIV.1903.11027>
- [9] Yunhao Du, Zhicheng Zhao, Yang Song, Yanyun Zhao, Fei Su, Tao Gong, and Hongying Meng. 2023. StrongSORT: Make DeepSORT Great Again. arXiv:2202.13514 [cs.CV]
- [10] Daniel J. Fremont, Tommaso Dreossi, Shromona Ghosh, Xiangyu Yue, Alberto L. Sangiovanni-Vincentelli, and Sanjit A. Seshia. 2019. Scenic: A Language for Scenario Specification and Scene Generation. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Phoenix, AZ, USA) (PLDI 2019). Association for Computing Machinery, New York, NY, USA, 63–78. <https://doi.org/10.1145/3314221.3314633>
- [11] Daniel J. Fremont, Edward Kim, Tommaso Dreossi, Shromona Ghosh, Xiangyu Yue, Alberto L. Sangiovanni-Vincentelli, and Sanjit A. Seshia. 2023. Scenic: a language for scenario specification and data generation. *Machine Learning* 112, 10 (01 Oct 2023), 3805–3849. <https://doi.org/10.1007/s10994-021-06120-5>
- [12] Jannik Fritsch, Tobias Kühnl, and Andreas Geiger. 2013. A new performance measure and evaluation benchmark for road detection algorithms. In *16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)*. 1693–1700. <https://doi.org/10.1109/ITSC.2013.6728473>
- [13] Yongming Ge, Vanessa Lin, Maureen Daum, Brandon Haynes, Alvin Cheung, and Magdalena Balazinska. 2021. Demonstration of Apperception: A Database Management System for Geospatial Video Data. *Proc. VLDB Endow.* 14, 12 (oct 2021), 2767–2770. <https://doi.org/10.14778/3476311.3476340>
- [14] Ross Girshick. 2015. Fast R-CNN. arXiv:1504.08083 [cs.CV]
- [15] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. 2014. Rich feature hierarchies for accurate object detection and semantic segmentation. arXiv:1311.2524 [cs.CV]
- [16] Derek Gloudeans and Daniel B. Work. 2021. Fast Vehicle Turning-Movement Counting using Localization-based Tracking. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. 4150–4159. <https://doi.org/10.1109/CVPRW53098.2021.00469>
- [17] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel Brostow. 2019. Digging Into Self-Supervised Monocular Depth Estimation. arXiv:1806.01260 [cs.CV]
- [18] Timothy Griffin and Monica K. Miller. 2008. Child Abduction, AMBER Alert, and Crime Control Theater. *Criminal Justice Review* 33, 2 (2008), 159–176. <https://doi.org/10.1177/0734016808316778> arXiv:https://doi.org/10.1177/0734016808316778
- [19] Brandon Haynes, Maureen Daum, Amrita Mazumdar, Magdalena Balazinska, Alvin Cheung, and Luis Ceze. 2020. VisualWorldDB: A DBMS for the Visual World. In *Conference on Innovative Data Systems Research*. 6.
- [20] Glenn Jocher, Ayush Chaurasia, Alex Stoken, Jirka Borovec, NanoCode012, Yonghye Kwon, Kalen Michael, TaoXie, Jiacong Fang, imyhxy, Lorna, Zeng Yifu, 曾逸夫 (Zeng Yifu), Colin Wong, Abhiram V, Diego Montes, Zhiqiang Wang, Cristi Fati, Jebastin Nadar, Laughing, UnglvKitDe, Victor Sonck, tkianai, yxNONG, Piotr Skalski, Adam Hogan, Dhruv Nair, Max Strobel, and Mrinal Jain. 2022. *ultralytics/yolov5: v7.0 - YOLOv5 SOTA Realtime Instance Segmentation*. <https://doi.org/10.5281/zenodo.7347926>
- [21] Daniel Kang, John Emmons, Firas Abuzaid, Peter Bailis, and Matei Zaharia. 2017. NoScope: Optimizing Neural Network Queries over Video at Scale. *Proc. VLDB Endow.* 10, 11 (aug 2017), 1586–1597. <https://doi.org/10.14778/3137628.3137664>
- [22] Daniel Kang, Francisco Romero, Peter D. Bailis, Christos Kozyrakis, and Matei Zaharia. 2022. VIVA: An End-to-End System for Interactive Video Analytics. In *12th Conference on Innovative Data Systems Research, CIDR 2022, Chaminate, CA, USA, January 9-12, 2022*. www.cidrdb.org, 9. <https://www.cidrdb.org/cidr2022/papers/p75-kang.pdf>
- [23] Aleksandr Kim, Aljoša Ošep, and Laura Leal-Taixé. 2021. EagerMOT: 3D Multi-Object Tracking via Sensor Fusion. arXiv:2104.14682 [cs.CV]
- [24] Edward Kim, Jay Shenoy, Sebastian Junges, Daniel Fremont, Alberto Sangiovanni-Vincentelli, and Sanjit Seshia. 2021. Querying Labelled Data with Scenario Programs for Sim-to-Real Validation. arXiv:2112.00206 [cs.CV]
- [25] H. W. Kuhn. 1955. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly* 2, 1-2 (1955), 83–97. <https://doi.org/10.1002/nav.3800020109> arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/nav.3800020109
- [26] Jack B. Kuipers. 1999. *Quaternions and rotation sequences: a primer with applications to orbits, aerospace, and virtual reality*. Princeton Univ. Press, Princeton, NJ. <http://www.worldcat.org/title/quaternions-and-rotation-sequences-a-primer-with-applications-to-orbits-aerospace-and-virtual-reality/oclc/246446345>
- [27] Laura Leal-Taixé. 2014. Multiple object tracking with context awareness. (11 2014), 15–25.
- [28] Jonathon Luiten, Aljoša Ošep, Patrick Dendorfer, Philip Torr, Andreas Geiger, Laura Leal-Taixé, and Bastian Leibe. 2020. HOTA: A Higher Order Metric for Evaluating Multi-object Tracking. *International Journal of Computer Vision* 129, 2 (oct 2020), 548–578. <https://doi.org/10.1007/s11263-020-01375-2>
- [29] Gavin G. McDonald, Christopher Costello, Jennifer Bone, Reniel B. Cabral, Valerie Farabee, Timothy Hochberg, David Kroodmsa, Tracey Mangin, Kyle C. Meng, and Oliver Zahn. 2021. Satellites can reveal global extent of forced labor in the world's fishing fleet. *Proceedings of the National Academy of Sciences* 118, 3 (2021), e2016238117. <https://doi.org/10.1073/pnas.2016238117> arXiv:https://www.pnas.org/doi/pdf/10.1073/pnas.2016238117
- [30] Jaeyoon Park, Jungsam Lee, Katherine Seto, Timothy Hochberg, Brian A. Wong, Nathan A. Miller, Kenji Takasaki, Hiroshi Kubota, Yoshioki Oozeki, Sejal Doshi, Maya Midzik, Quentin Hanich, Brian Sullivan, Paul Woods, and David A. Kroodmsa. 2020. Illuminating dark fishing fleets in North Korea. *Science Advances* 6, 30 (2020), eabb1197. <https://doi.org/10.1126/sciadv.abb1197> arXiv:https://www.science.org/doi/pdf/10.1126/sciadv.abb1197
- [31] Joseph Redmon and Ali Farhadi. 2018. YOLOv3: An Incremental Improvement. arXiv:1804.02767 [cs.CV]
- [32] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2016. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. arXiv:1506.01497 [cs.CV]
- [33] Francisco Romero, Johann Hauswald, Aditi Partap, Daniel Kang, Matei Zaharia, and Christos Kozyrakis. 2022. Optimizing Video Analytics with Declarative Model Relationships. *Proc. VLDB Endow.* 16, 3 (nov 2022), 447–460. <https://doi.org/10.14778/3570690.3570695>
- [34] Martin Simon, Karl Amende, Andrea Kraus, Jens Honer, Timo Sämann, Hauke Kaulbersch, Stefan Milz, and Horst Michael Gross. 2019. Complexer-YOLO: Real-Time 3D Object Detection and Tracking on Semantic Point Clouds. arXiv:1904.07537 [cs.CV]
- [35] Jack Sklansky. 1982. Finding the Convex Hull of a Simple Polygon. *Pattern Recogn. Lett.* 1, 2 (dec 1982), 79–83. [https://doi.org/10.1016/0167-8655\(82\)90016-2](https://doi.org/10.1016/0167-8655(82)90016-2)
- [36] Robert Sparks. 2023. 11 Security Camera Statistics & Data - 2023 update. <https://opticsmag.com/security-camera-statistics/>
- [37] Suramya Tomar. 2006. Converting Video Formats with Ffmpeg. *Linux J.* 2006, 146 (jun 2006), 10.
- [38] Nicolai Wójke, Alex Bewley, and Dietrich Paulus. 2017. Simple Online and Realtime Tracking with a Deep Association Metric. arXiv:1703.07402 [cs.CV]
- [39] Simon Wright. 2023. Autonomous cars generate more than 300 TB of data per year. <https://www.tuxera.com/blog/autonomous-cars-300-tb-of-data-per-year/>
- [40] Zhuangdi Xu, Gaurav Tarlok Kakkar, Joy Arulraj, and Umakishore Ramachandran. 2022. EVA: A Symbolic Approach to Accelerating Exploratory Video Analytics with Materialized Views. In *Proceedings of the 2022 International Conference on Management of Data* (Philadelphia, PA, USA) (SIGMOD '22). Association for Computing Machinery, New York, NY, USA, 602–616. <https://doi.org/10.1145/3514221.3526142>
- [41] Zhengyou Zhang. 2014. *Camera Parameters (Intrinsic, Extrinsic)*. Springer US, Boston, MA, 81–85. https://doi.org/10.1007/978-0-387-31439-6_152
- [42] Kaiyang Zhou, Yongxin Yang, Andrea Cavallaro, and Tao Xiang. 2019. Omni-Scale Feature Learning for Person Re-Identification. arXiv:1905.00953 [cs.CV]
- [43] Esteban Zimányi, Mahmoud Sakr, and Arthur Lesuisse. 2020. MobilityDB: A Mobility Database Based on PostgresSQL and PostGIS. *ACM Trans. Database Syst.* 45, 4, Article 19 (dec 2020), 42 pages. <https://doi.org/10.1145/3406534>